H. Vicky Zhao, W. Sabrina Lin, and K.J. Ray Liu

Behavior Modeling and Forensics for Multimedia Social Networks

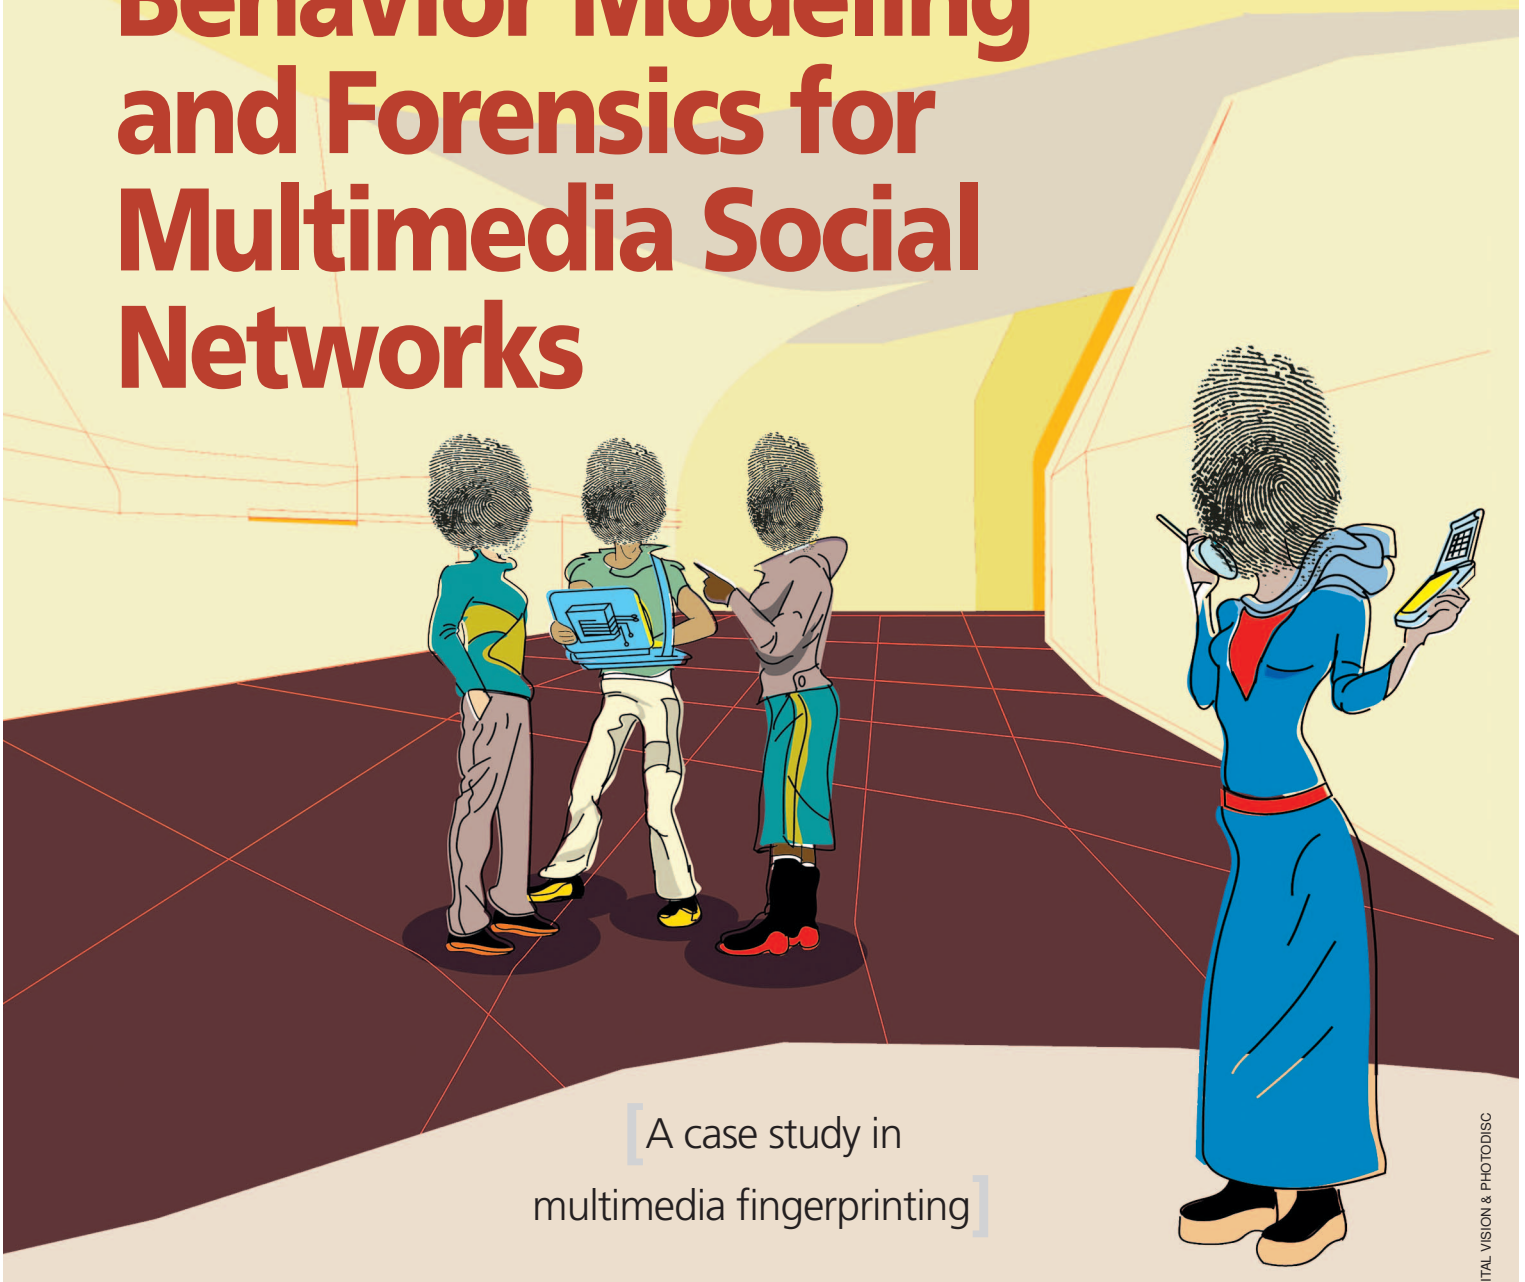

A case study in multimedia fingerprinting

Within the past decade, Internet traffic has shifted dramatically from HTML text pages to multimedia file sharing [1] as illustrated by the emergence of large-scale multimedia social network communities such as Napster, flickr, and YouTube. For example, a study showed that in a campus network, peer-to-peer file sharing can consume 43% of the overall bandwidth, which is about three times of all WWW traffic [2]. This consumption poses new challenges to the efficient, scalable, and robust sharing of multimedia over large and heterogeneous networks. It also significantly affects the copyright industries and raises critical issues of protecting intellectual property rights of multimedia.

This recent increase in Internet traffic adversely affects the user experience for people all across the world. To improve the efficiency of data transmission within multimedia social networks, we must analyze the impact of human factors on multimedia networks, that is, how users interact with and respond to

Digital Object Identifier 10.1109/MSP.2008.930648

1053-5888/09/$25.00©2009IEEE

one another. Such an understanding provides fundamental guidelines to better design of multimedia systems and networking, and to offer more secure and personalized services. The area of human and social dynamics has recently been identified by the U.S. National Science Foundation as one of its five priority areas, which also shows the importance of this emerging interdisciplinary research area.

Factors influencing human behavior have seldom appeared in signal processing disciplines. Therefore, the goals of this tutorial are to illustrate why human factors are important, identify emerging issues strongly related to signal processing, and to demonstrate that signal processing can be effectively used to model, analyze, and perform behavior forensics for multimedia social networks. Since media security and content protection is a major issue, this article illustrates various aspects of issues and problems in multimedia social networks via a case study of human behavior in traitor-tracing multimedia fingerprinting. We focus on the understanding of behavior forensics from signal processing perspective and present a framework to model and analyze user dynamics. The objective is to provide a broad overview of recent advances in behavior modeling and forensics for multimedia social networks.

MULTIMEDIA SOCIAL NETWORKS

A social network is a structure of nodes (including individuals and organizations) that are connected with each other via certain types of relations, for example, values, friendship, conflict, financial exchange, and trade. Figure 1 gives examples of some typical social networks, and Figure 2 demonstrates the complex user dynamics there. People have been studying methodologies to formulate the relationships between members at all scales, from interpersonal to international, and across many disciplines such as sociology, economics, and information science.

In a multimedia social network community, a group of users form a dynamically changing network infrastructure to share and exchange data, often multimedia content, as well as other resources. For example, in a peer-to-peer file-sharing system, users pool together the resources and cooperate with each other to provide an inexpensive, highly scalable, and robust platform for distributed data sharing [3], [4]. However, since participation nature in many multimedia social networks is often voluntary and unregulated, users' full cooperation cannot be guaranteed unless there exist powerful central authorities who mandate and enforce user cooperation. A recent study of Napster and Gnutella showed that many users are free riders and 25% of the users in Gnutella share no files at all [5].

Before multimedia social network communities become successful, they must provide a predictable and satisfactory level of service, and a critical issue to be resolved first is to stimulate cooperation among users [6]. For example, in peer-to-peer file-sharing systems, one possible solution is to use payment-based methods where users pay to consume resources and are paid if they contribute resources [6]. These schemes can effectively stimulate cooperation, but they require tamper-proof hardware or central billing services to track various transactions and implement micropayment. Another form of incentives is to use reputation-based methods to differentiate among users and

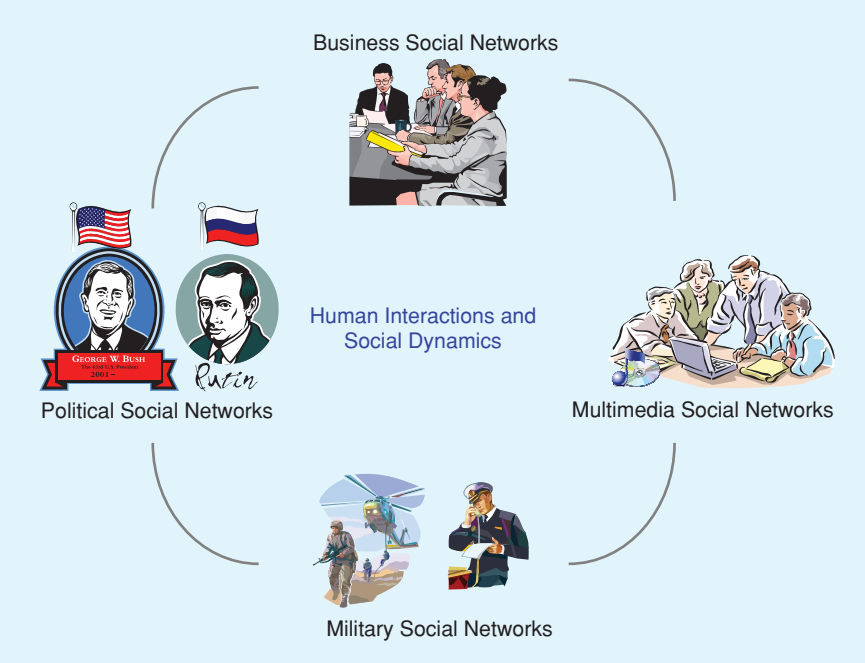

[FIG1] **Examples of social networks.**

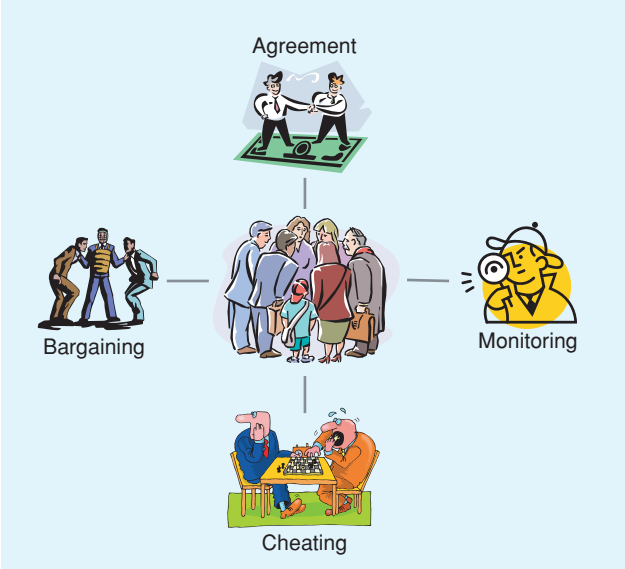

[FIG2] **User dynamics in social networks.**

adopt the differential service model to offer better services to peers who contribute more [6].

By participating in multimedia social networks, users receive rewards by being able to access extra resources from their peers, and they also contribute their own resources. Users aim to maximize their own payoff by participating in multimedia social networks, and different users have different (and often conflicting) objectives. Thus, as demonstrated in Figure 2, an important issue in multimedia social networks is to understand the strategies that users will play when negotiating with each other and study how they achieve fairness. Game theory [7], [8] provides a fundamental tool to study the fairness dynamics among users. The Nash equilibrium provides the optimum strategies from which no user has incentives to deviate.

There are different types of users in multimedia social networks. Rational users are willing to contribute their own resources if cooperation with others can help improve their payoff. They are honest when exchanging information with other users. Unlike rational users, there are also selfish users who wish to consume others' resources with little or no contribution of their own. If necessary, these selfish users might even cheat during the negotiation process in order to maximize their own payoff, as shown in Figure 2. Furthermore, there might exist malicious users whose goal is to attack and sabotage the system. For example, in peer-to-peer file-sharing systems, they tamper with the media files with the intention of making the content useless (the so-called pollution attack) [2]. They can also launch the denial of service (DoS) attack to exhaust other users' resources and make the system unavailable [9]. It is possible that a few malicious users collude with each other to effectively attack the system, for example, the flooding distributed DoS (DDoS) attack in peer-to-peer file-sharing systems. Therefore, cheat prevention and attack resistance are fundamental requirements in order to achieve user cooperation and provide reliable services.

To model and analyze human dynamics in multimedia social networks containing selfish users and malicious users, the first step is to study the strategies that these users use to cheat or attack the system. The next issue is to implement monitoring mechanisms to detect and identify misbehaving users, as illustrated in Figure 2. A challenging issue here is that the monitoring mechanisms should be able to distinguish intentional misbehavior (for example, intentional manipulation of multimedia content) from the innocent (for example, transmission errors and packet loss in erroneous and congested networks). The above investigation will facilitate the design of cheat-proofing and attack-resistant strategies, which make noncooperation nonprofitable, thus unattractive to selfish users, and minimize the damage to the system caused by malicious users.

Because different multimedia social networks have different structures, there are different ways to implement cheat-proofing and attack-resistant cooperation strategies. Some multimedia social networks have a centralized structure where there are one or more entities whom all users trust and who can facilitate interaction among users. For example, the first generation peer-to-peer file-sharing networks (for example, the Napster music file-sharing system) used a set of central servers to provide content indexing and search services [4]. Though these central servers do not have the authorities to enforce user cooperation, they can help monitor users' behavior. For example, they can serve as the central billing entity in the payment-based methods to help track the transactions and identify misbehaving users [4]. Other multimedia social networks have a distributed structure and a flat topology where users take the same role, for example, Gnutella and Chord [4]. In these multimedia social networks, users have to monitor other users and identify misbehavior themselves.

Essentially, multimedia social networks involve a large number of users of different types with different objectives, and modeling and analysis of user dynamics is a fundamental issue to address in multimedia social networks. Such an analysis helps stimulate user cooperation, facilitates the implementation of misbehavior monitoring mechanisms, and provides important guidelines on the design of cheat-proofing and attack-resistant strategies. All these are essential factors to maximize the overall system performance and minimize the damage caused by malicious users. In addition, for different multimedia social networks, different structures will result in different mechanisms to monitor user behavior and to achieve cheat prevention and attack resistance.

BEHAVIOR MODELING IN MULTIMEDIA FINGERPRINTING FORENSICS

Without loss of generality, in this article, we use multimedia fingerprinting as an example to illustrate the modeling and analysis of user behavior in multimedia social networks. In this section, we first introduce the digital fingerprinting technology used to identify the source of illicit copies. Then, we formulate the dynamics among users in multimedia fingerprinting.

As we move to the digital era and experience the convergence of networks, communications and multimedia, scalability in multimedia coding becomes a critical issue to support universal media access and provide rich media access from anywhere using any devices [10]. Scalable video coding encodes video into several bit streams (layers) of different priorities: the base layer contains the most important information and the enhancement layers gradually refine the resolution of the receiver's reconstructed copy. Such a coding structure provides flexible solutions for multimedia transmission and offers adaptivity to heterogeneous networks, varying channel conditions and diverse computing capability at the receiving terminals [10].

In this article, we use temporal scalability as an example, inherent in most current video coding standards such as Moving Pictures Expert Group (MPEG) and H.26x, to demonstrate issues that arise from scalability. Without loss of generality, we consider three-layer temporal scalability and use frame skipping and frame copying to implement temporal decimation and interpolation, respectively. For example, with MPEG encoding, the base layer may include all the I frames, the enhancement

layer one includes all the P frames, and the enhancement layer two contains all the B frames.

MULTIMEDIA FINGERPRINTING FOR TRAITOR TRACING

Digital fingerprinting is an emerging technology that offers proactive post-delivery protection of multimedia [11], [12]. As illustrated in Figure 3, it labels each distributed copy with the corresponding user's identification information, known as a fingerprint, which can be used to trace culprits who use their copies illegally. Traitor-tracing multimedia forensics has compelling commercial applications, for example, in the antipiracy campaign by Hollywood film industries. A preliminary technology based on robust watermarking was adopted in the 2004 Oscar season and successfully captured a few people who helped illegally post Oscar screener copies of movies on the Internet [13]. It is also important for government operations and intelligence agencies to be able to securely and reliably exchange multimedia data and prevent the leakage of confidential information.

In multimedia fingerprinting, fingerprints are embedded into the host signal using traditional data hiding techniques [14]–[16]. Spread spectrum embedding techniques [17], [18] are widely used in the literature due to the robustness against many attacks. With the three-layer temporally scalable coding structure, let S_b, S_{e1}, and S_{e2} be the base layer, the enhancement layer one and the enhancement layer two of the host signal, respectively. For Alice, who subscribes to the low-resolution copy, let $W_b^{(\text{alice})}$ be Alice's unique fingerprint. The content owner distributes to Alice the fingerprinted base layer

$$X_b^{(\text{alice})}(j) = S_b(j) + c(j) \times W_b^{(\text{alice})}(j). \tag{1}$$

Here, $X_b^{(\text{alice})}(j)$, $S_b(j)$, and $W_b^{(\text{alice})}(j)$ are the jth components of the fingerprinted copy, the host signal, and Alice's fingerprint, respectively. c is the just noticeable difference (JND) from human visual models [17], [18] to control the energy and achieve the imperceptibility of the embedded fingerprints. For Bob, who subscribes to the medium resolution copy, he receives the fingerprinted base layer

$$X_b^{(\text{bob})} = S_b + c \times W_b^{(\text{bob})} \tag{2}$$

and the fingerprinted enhancement layer one

$$X_{e1}^{(\text{bob})} = S_{e1} + c \times W_{e1}^{(\text{bob})}, \tag{3}$$

from the content owner, where $W_b^{(\text{bob})}$ and $W_{e1}^{(\text{bob})}$ are Bob's fingerprints embedded in the base layer and the enhancement layer one, respectively. (We drop the component index j to simplify the notations.) Similarly, for Carl, who subscribes to all three layers, the fingerprinted base layer, enhancement layer one and enhancement layer two that he receives are $X_b^{(\text{carl})} = S_b + c \times W_b^{(\text{carl})}$, $X_{e1}^{(\text{carl})} = S_{e1} + c \times W_{e1}^{(\text{carl})}$, and $X_{e2}^{(\text{carl})} = S_{e2} + c \times W_{e2}^{(\text{carl})}$, respectively. $W_b^{(\text{carl})}$, $W_{e1}^{(\text{carl})}$ and

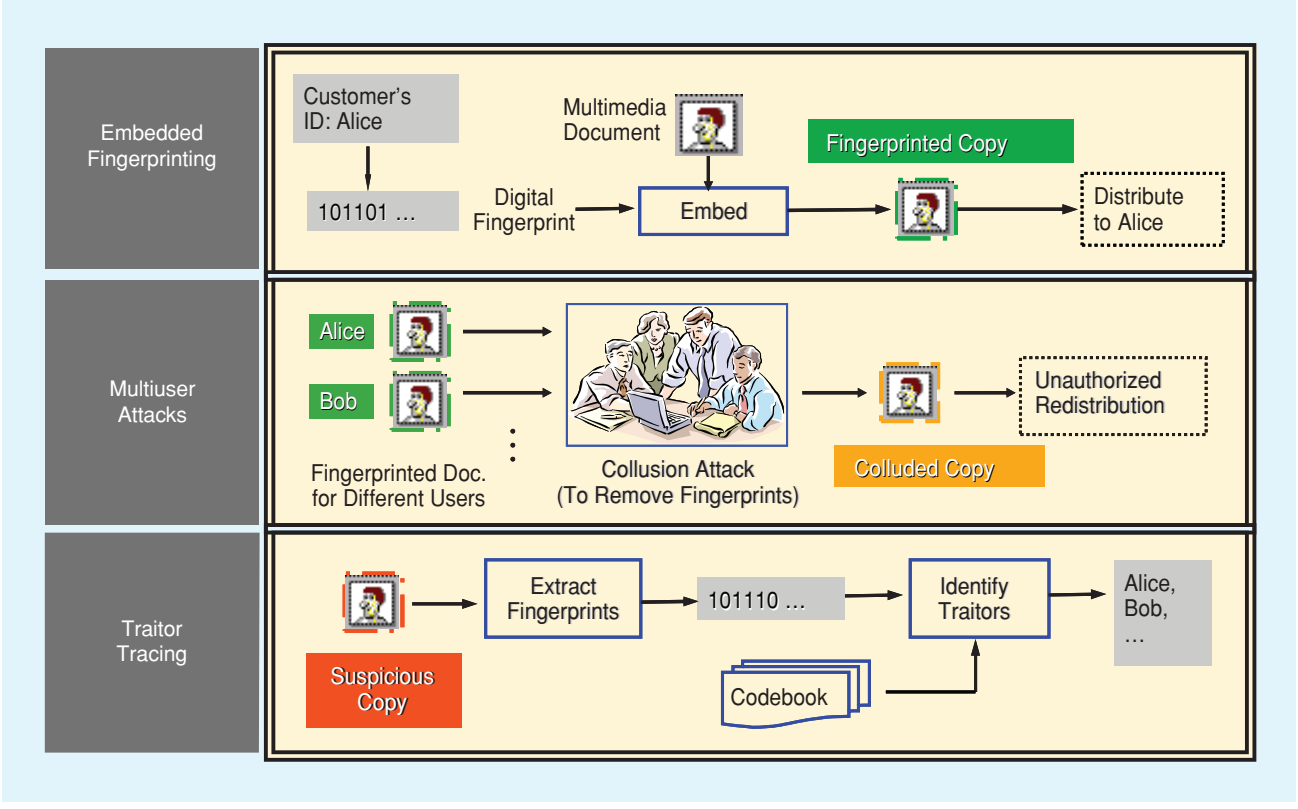

[FIG3] Using embedding fingerprinting for traitor tracing.

$\mathbf{W}_{e2}^{(\text{carl})}$ are Carl's fingerprints that are embedded in the base layer, enhancement layer one and enhancement layer two, respectively. Here, the superscript is the user index and the subscript is the layer index.

Without loss of generality, we consider orthogonal fingerprint modulation [12], where in the same layer, fingerprints assigned to different users are orthogonal to each other and have the same energy. For example, in the above example, with orthogonal fingerprint modulation, we have

$$\begin{aligned}
&\left\langle \mathbf{W}_b^{(\text{alice})}, \mathbf{W}_b^{(\text{bob})} \right\rangle = \left\langle \mathbf{W}_b^{(\text{alice})}, \mathbf{W}_b^{(\text{carl})} \right\rangle = \left\langle \mathbf{W}_b^{(\text{bob})}, \mathbf{W}_b^{(\text{carl})} \right\rangle = 0, \\
&\left\langle \mathbf{W}_{e1}^{(\text{bob})}, \mathbf{W}_{e1}^{(\text{carl})} \right\rangle = 0, \quad \left\| \mathbf{W}_b^{(\text{alice})} \right\|^2 \\
&\qquad = \left\| \mathbf{W}_b^{(\text{bob})} \right\|^2 = \left\| \mathbf{W}_b^{(\text{carl})} \right\|^2, \quad \text{and} \\
&\left\| \mathbf{W}_{e1}^{(\text{bob})} \right\|^2 = \left\| \mathbf{W}_{e1}^{(\text{carl})} \right\|^2. \qquad (4)
\end{aligned}$$

In (4), $\langle x, y \rangle$ is the correlation between x and y, and $\|x\|^2$ returns the Euclidean norm of x.

Once an illegal copy is discovered, the digital rights enforcer first extracts the fingerprint $\mathbf{Y}$ from the colluded copy. Then, the digital rights enforcer uses the correlation-based detection statistic

$$TN^{(i)} = \langle \mathbf{Y}, \mathbf{W}^{(i)} \rangle / \|\mathbf{W}^{(i)}\| \qquad (5)$$

to measure the similarity between the extracted fingerprint $\mathbf{Y}$ and user $\mathbf{u}^{(i)}$'s fingerprint $\mathbf{W}^{(i)}$. The fingerprint detector compares all the detection statistics $\{TN^{(i)}\}$ with a predetermined threshold h and identifies those whose detection statistics are larger than h as colluders.

COLLUSION ATTACKS AND ANTICOLLUSION FINGERPRINT DESIGN

However, protecting digital fingerprints is no longer a traditional security issue with a single adversary. The global nature of Internet has enabled a group of attackers (colluders) to work together and collectively mount attacks to remove the fingerprints. These attacks, known as multiuser collusion, pose serious threats to intellectual property rights. Analysis of the strategies, capabilities, and limitations of attackers is an indispensable and crucial part of research in multimedia security.

Linear collusion is one of the most feasible collusion attacks that may be employed against multimedia fingerprinting [19]–[21]. Given K different fingerprinted signals $\{\mathbf{X}^{(i)}\}$ of the same content, attackers generate $\mathbf{Y} = \sum_k a_k \mathbf{X}^{(k)}$, where the weights satisfy $\sum_k a_k = 1$ to maintain the average intensity of the original multimedia signal (thus the perceptual quality of the attacked copy). With orthogonal fingerprinting, such an averaging attenuates the energy of the kth contributing fingerprint by a factor of a_k^2 and thus reduces colluder k's probability of being detected. In [19], collusion attacks were modeled as averaging differently fingerprinted copies with equal weights (that is, $a_k = 1/K$) followed by the addition of noise. Their work showed that $O(\sqrt{N/\log N})$ colluders are sufficient to defeat the underlying fingerprinting system, where N is the fingerprint length.

In addition to linear averaging, another important class of collusion attacks is based upon operations as taking the minimum, maximum, and median of corresponding components of the fingerprinted signals [22]. For example, given K fingerprinted signals $\{\mathbf{X}^{(i)}\}$, to generate the jth component of the colluded copy $\mathbf{Y}(j)$, colluders use the minimum value of $\mathbf{X}^{(1)}(j), \mathbf{X}^{(2)}(j), \cdots, \mathbf{X}^{(K)}(j)$ and let $\mathbf{Y}(j) = \min(\{\mathbf{X}^{(k)}(j)\})$. Since each fingerprinted copy is expected to have high perceptual quality, colluders have high confidence that $\mathbf{Y}(j)$ is within the JSD range. Similarly, colluders can also let $\mathbf{Y}(j) = \max(\{\mathbf{X}^{(k)}(j)\})$ and take the maximum value of $\{\mathbf{X}^{(i)}(j)\}$. They can also use the median value and select $\mathbf{Y}(j) = \text{median}(\{\mathbf{X}^{(k)}(j)\})$. Detailed analysis of linear and nonlinear collusion attacks on orthogonal fingerprints was provided in [23]. The gradient attack was proposed in [24], which uses the combination of several basic nonlinear collusion attacks in [23]. The work in [25] evaluated the collusion resistance of multimedia fingerprints as a function of system parameters, including fingerprint length, total number of users, and system requirements.

Collusion attacks pose serious threats to multimedia intellectual property rights. To provide reliable and trustworthy traitor-tracing performance, it is of ample importance to design anticollusion fingerprints. In the literature, techniques from a wide range of disciplines were used to improve the fingerprinting system's collusion resistance. A two-layer fingerprint design scheme was proposed in [26] where the inner code from spread spectrum embedding [17], [18] is combined with an outer error-correcting code (ECC) [27]. A permuted subsegment embedding technique and a group-based joint coding and embedding technique were proposed in [28] to improve the collusion resistance of ECC-based multimedia fingerprinting while maintaining the detection efficiency. In [29], finite projective geometry was used to generate codes whose overlap with each other can identify colluding users. The anticollusion code based on combinatorial theories was proposed in [30]. In [31], prior knowledge of the possible collusion patterns was used to improve the collusion resistance of the fingerprinting systems. The anticollusion dithering technique was proposed in [32] to resist multiuser collusion attacks for compressed multimedia. Readers who are interested in anticollusion fingerprint design are referred to [12] for detailed discussion of current research in this area.

BEHAVIOR MODELING AND FORENSICS IN MULTIMEDIA FINGERPRINTING

During collusion, attackers form a unique social network: they share the reward from the illegal usage of multimedia as well as the risk of being captured by the digital rights enforcer. An agreement must be reached regarding how to distribute the risk and the reward before collusion relationship can be established.

However, each colluder prefers the agreement that favors his or her payoff the most, and different colluders have different preferences. To address such a conflict, a critical issue is to decide how to fairly distribute the risk and the reward. In addition, even though all colluders agree so, some colluders might be selfish and wish to break away from their fair-collusion agreement. They might cheat their fellow attackers during the negotiation process in order to minimize their own risk and maximize their own payoff.

On the other hand, to protect their own interests, other colluders may want to identify selfish colluders and exclude them from collaboration. It is of great importance to understand how colluders negotiate with each other to achieve fairness of the attack and study the cheating and the cheat-proofing strategies that colluders may adopt to maximize their own payoff and protect their own interests.

In addition, users in multimedia fingerprinting influence each other's decisions and performance. To maximize their own payoff, users should observe and learn how others play the game and adjust their own strategies accordingly. For example, to maximize the traitor-tracing capability, the digital rights enforcer should explore and utilize as much knowledge about collusion as possible when designing the fingerprints and identifying the colluders. Here, analysis of the colluder dynamics, especially the investigation on how attackers achieve fairness of collusion, provides the digital rights enforcer with important insights on how to probe and use such side information about collusion. Therefore, another important issue in behavior modeling is to understand the techniques that users can use to probe information about how others play the game, study how they adjust their strategies accordingly to maximize their own payoff, and analyze the impact of side information on multimedia social networks.

In the sequel, using multimedia fingerprinting system as an example, we present a general framework by employing a few signal processing techniques to formulate and analyze human behavior in multimedia social networks. We first investigate the fairness dynamics in colluder social networks and analyze how colluders negotiate with each other to achieve fair collusion. We then study how selfish colluders cheat to maximize their own payoff and how other colluders detect such cheating behavior. We will also study side information in multimedia fingerprinting, how the digital rights enforcer can probe and utilize side information to improve the traitor-tracing performance, and how such side information affect the overall fingerprinting systems.

FAIRNESS DYNAMICS IN MULTIMEDIA SOCIAL NETWORKS

In multimedia social networks, by contributing their own resources and cooperating with each other, users are able to access extra resources from their peers and thus receive rewards. Each user aims to maximize his or her own payoff and different users have different objectives. To address this conflict, an important issue is to investigate users' strategies to achieve a notion of fairness. In this section, we use colluder social networks as an example to illustrate the methodologies that can be used to analyze the fairness dynamics among users.

EQUAL-RISK ABSOLUTE FAIRNESS

Colluders receive rewards from the illegal usage of multimedia content, for example, the profit from the unauthorized redistribution of copyrighted materials. They also take the risk of being captured by the digital rights enforcer, which can be considered as the colluders' cost by participating in collusion. In the literature, a commonly used definition of a colluder's utility (payoff) function is his or her chance of not being captured by the digital rights enforcer, and the notion of equal-risk absolute fairness is widely adopted where all colluders agree to share the same risk and have equal probabilities of being detected.

If all colluders receive fingerprinted copies of the same resolution, a simple average of all copies with equal weights reduces the energy of each contributing fingerprint by the same ratio, thus ensuring equal risk of all attackers. When colluders receive fingerprinted copies of different resolutions, it is much more complicated to guarantee equal risk of all colluders, especially when colluders wish to generate a colluded copy of higher resolution.

A SIMPLE EXAMPLE WITH THREE COLLUDERS

For the example with three colluders, Alice, Bob, and Carl, who receive fingerprinted copies of different resolutions, a possible solution of collusion is shown in Figure 4(a), where the colluded copy includes all three layers. Here, the colluders average the three base-layer copies that they have with equal weights 1/3; for the enhancement layer one, they average the two copies from Bob and Carl with equal weights 1/2; and the colluded copy's enhancement layer two equals to that in Carl's copy. Therefore, in the colluded copy, the three fingerprints corresponding to the three attackers have the same energy in the base layer. The enhancement layers contain only Bob and Carl's fingerprints, not the fingerprint identifying Alice. It is obvious that among the three, Carl has the largest probability of being caught and Alice takes the smallest risk. Consequently, the collusion in Figure 4(a) does not achieve equal-risk fairness.

Figure 4(b) shows another possible solution, where the colluded copy contains the base layer only. Here, the colluders average the three copies of the base layer with equal weights 1/3. In this example, the fingerprints corresponding to the three attackers have the same energy in the colluded copy and, therefore, the three attackers have the same probability of being detected. Although the collusion in Figure 4(b) ensures equal-risk fairness, the attacked copy has low resolution.

When there is difference in the resolution of fingerprinted copies due to network and device heterogeneity, how can colluders establish fair multiuser collusion that guarantees the collective equal risk among all attackers while still generating an attacked copy of high resolution? A possible solution is shown in Figure 5. In the base layer of the colluded copy, the three copies are assigned different weights β_1, β_2, and β_3, respectively.

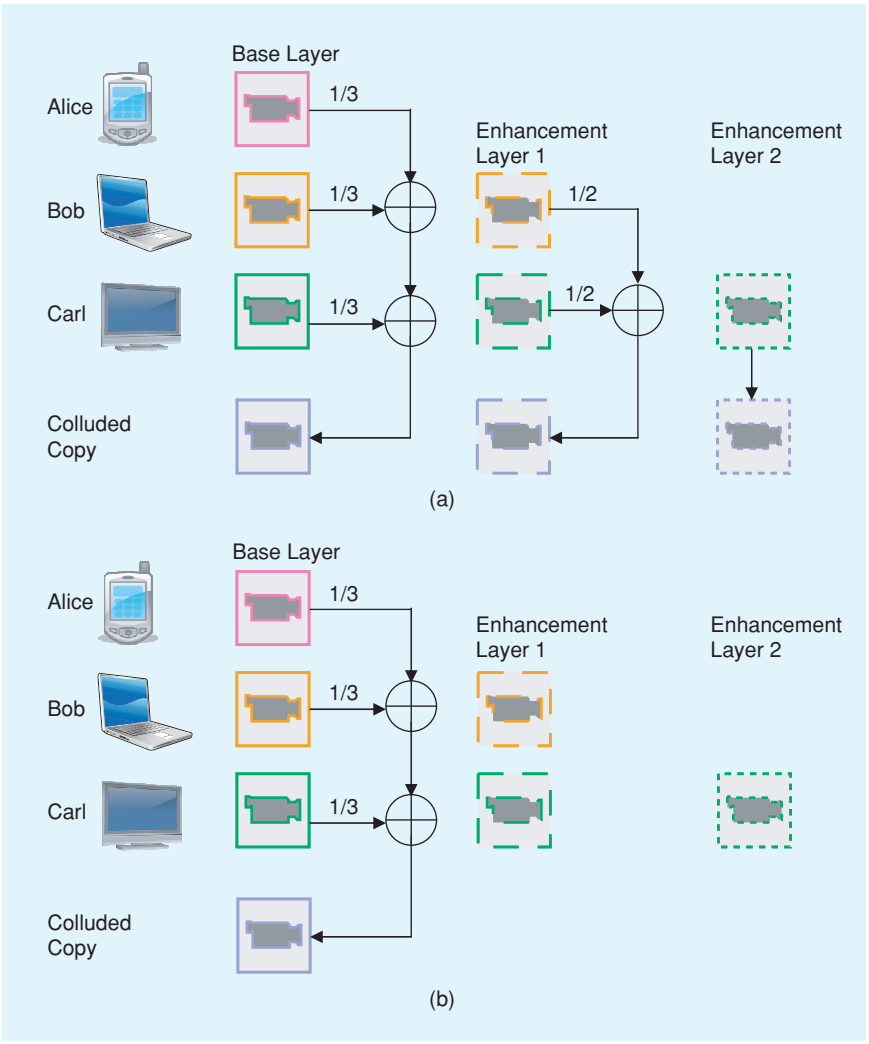

[FIG4] Two solutions of collusion in scalable multimedia fingerprinting.

Similarly, the enhancement layer one in the colluded copy is the average of Bob and Carl's copies with weights α_1 and α_2, respectively. The colluders copy the enhancement layer two in Carl's copy to the colluded copy. To achieve fairness of collusion, Alice, Bob, and Carl select the collusion parameters $\{\alpha_k, \beta_l\}$ such that they have the same probability of being detected.

TWO-STAGE COLLUSION

In general, when colluders receive fingerprinted copies of different resolutions, they apply the two-stage collusion as in Figure 5 to achieve equal-risk absolute fairness. They first apply the intra-group collusion to guarantee that colluders who receive copies of the same resolution have the same probability of being detected. Then, they apply the the inter-group collusion to ensure that colluders who receive copies of different resolutions share the same risk.

To demonstrate how attackers colluder in scalable multimedia fingerprinting, we first introduce the symbols that we use. For user $\mathbf{u}^{(i)}$, let $\mathbf{X}_b^{(i)}$, $\mathbf{X}_{e1}^{(i)}$, and $\mathbf{X}_{e2}^{(i)}$ denote the fingerprinted base layer, enhancement layer one and enhancement layer two, respectively, that $\mathbf{u}^{(i)}$ receives from the content owner. In our notations, the superscript i is the user index and the subscript b, $e1$, or $e2$ is the layer index.

During collusion, the colluders first divide themselves into three subgroups: SC^b includes the indices of those colluders who receive the fingerprinted base layer only; the second subgroup, $SC^{b,e1}$, contains colluders who receive the base layer and the enhancement layer one; and the last group, SC^{all}, includes colluders who receive all three layers. Let $K^b = |SC^b|$, $K^{b,e1} = |SC^{b,e1}|$ and $K^{\text{all}} = |SC^{\text{all}}|$ be the numbers of colluders in SC^b, $SC^{b,e1}$ and SC^{all}, respectively. Here, we use "b," "b, $e1$," and "all" in the superscript to differentiate different subgroups of colluders.

Then, they apply the intra-group collusion, where colluders collude with their fellow attackers in the same subgroup and average different copies of the same resolution with equal weights. In this stage, different subgroups collude independently. This intra-group collusion ensures that colluders who receive fingerprinted copies of the same resolution have the same probability of being detected. In our example, colluders in SC^b generates $\mathbf{X}_b^b = \sum_{k\in SC^b} \mathbf{X}_b^{(k)}/K^b$. Colluders in $SC^{b,e1}$ generates a copy of the base layer $\mathbf{X}_b^{b,e1} = \sum_{k\in SC^{b,e1}} \mathbf{X}_b^{(k)}/K^{b,e1}$ and a copy of the enhancement layer one $\mathbf{X}_{e1}^{b,e1} = \sum_{k\in SC^{b,e1}} \mathbf{X}_{e1}^{(k)}/K^{b,e1}$. Similarly, $\mathbf{X}_b^{\text{all}} = \sum_{k\in SC^{\text{all}}} \mathbf{X}_b^{(k)}/K^{\text{all}}$, $\mathbf{X}_{e1}^{\text{all}} = \sum_{k\in SC^{\text{all}}} \mathbf{X}_{e1}^{(k)}/K^{\text{all}}$, and $\mathbf{X}_{e2}^{\text{all}} = \sum_{k\in SC^{\text{all}}} \mathbf{X}_{e2}^{(k)}/K^{\text{all}}$.

Finally, as illustrated in Figure 5, colluders apply the inter-group collusion and average copies from different subgroups with different weights. This step guarantees that colluders who receive fingerprinted copies of different resolutions have equal risk of being captured. In our example, in the final colluded copy V, the base layer is

$$\mathbf{V}_b = \beta_1 \mathbf{X}_b^b + \beta_2 \mathbf{X}_b^{b,e1} + \beta_3 \mathbf{X}_b^{\text{all}} + \mathbf{n}, \tag{6}$$

where $0 \le \beta_1, \beta_2, \beta_3 \le \beta_1 + \beta_2 + \beta_3 = 1$; the enhancement layer one is

$$\mathbf{V}_{e1} = \alpha_1 \mathbf{X}_{e1}^{b,e1} + \alpha_2 \mathbf{X}_{e1}^{\text{all}} + \mathbf{n}, \tag{7}$$

where $0 \le \alpha_1, \alpha_2 \le \alpha_1 + \alpha_2 = 1$; and the enhancement layer two is

$$\mathbf{V}_{e2} = \mathbf{X}_{e2}^{\text{all}} + \mathbf{n}. \tag{8}$$

In (6)–(8), $\mathbf{n}$ is additive noise to further hinder the detection process.

ACHIEVING EQUAL-RISK FAIRNESS

Given the above two-stage collusion model, to ensure equal risk for all colluders, attackers need to first estimate each colluder's probability of being detected, and then select the collusion parameters $\{\alpha_k, \beta_l\}$ in (6)–(8) accordingly. Therefore, an important step in multiuser collusion is to follow the same fingerprint detection process as the digital rights enforcer and estimate each attacker's chance of being caught. This analysis provides colluders with important guidelines on the selection of collusion parameters to achieve fairness.

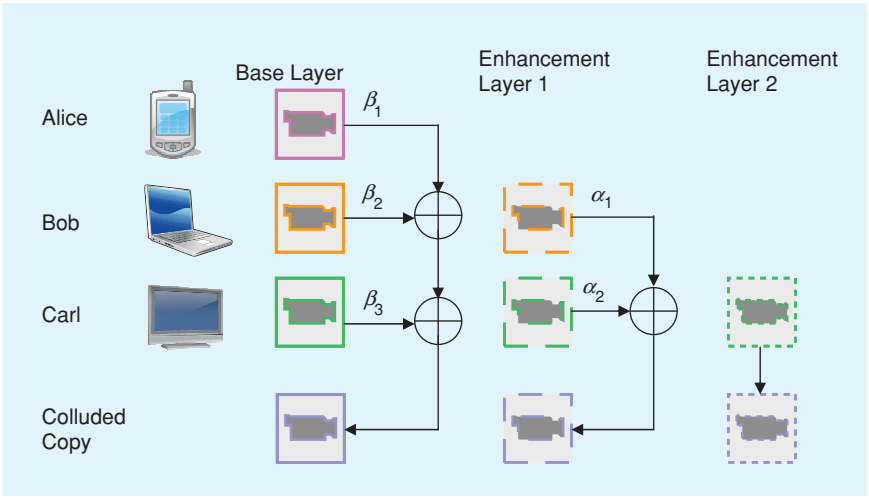

[FIG5] The intra-group and the inter-group collusion attacks.

In the example in Figure 5, let $\mathbf{Y}_b$, $\mathbf{Y}_{e1}$ and $\mathbf{Y}_{e2}$ be the fingerprints extracted from the base layer, enhancement layer one and enhancement layer two, respectively. Since Alice only receives the base layer from the content owner, only $\mathbf{Y}_b$ is used to determine if she participates in collusion. Her detection statistic is

$$TN_c^{(\text{alice})} = \left\langle \mathbf{Y}_b, \mathbf{W}_b^{(\text{alice})} \right\rangle \Big/ \left\| \mathbf{W}_b^{(\text{alice})} \right\|. \tag{9}$$

For Bob, who receives a medium-resolution copy, $\mathbf{Y}_b$ and $\mathbf{Y}_{e1}$ are used collectively to decide if Bob is a colluder, and his detection statistic is

$$\begin{aligned} TN_c^{(\text{bob})} &= \langle \mathbf{Y}, \mathbf{W}^{(\text{bob})} \rangle / \|\mathbf{W}^{(\text{bob})}\| \\ \text{where} \quad \langle \mathbf{Y}, \mathbf{W}^{(\text{bob})} \rangle &= \left\langle \mathbf{Y}_b, \mathbf{W}_b^{(\text{bob})} \right\rangle \\ &\quad + \left\langle \mathbf{Y}_{e1}, \mathbf{W}_{e1}^{(\text{bob})} \right\rangle, \\ \text{and} \quad \|\mathbf{W}^{(\text{bob})}\|^2 &= \left\| \mathbf{W}_b^{(\text{bob})} \right\|^2 + \left\| \mathbf{W}_{e1}^{(\text{bob})} \right\|^2. \end{aligned} \tag{10}$$

Because Carl receives all three layers from the content owner, $\mathbf{Y}_b$, $\mathbf{Y}_{e1}$ and $\mathbf{Y}_{e2}$ will be used collectively to determine if Carl colludes with others, and Carl's detection statistic is

$$\begin{aligned} TN_c^{(\text{carl})} &= \langle \mathbf{Y}, \mathbf{W}^{(\text{carl})} \rangle \Big/ \sqrt{\|\mathbf{W}^{(\text{carl})}\|^2} \\ \text{where} \quad \langle \mathbf{Y}, \mathbf{W}^{(\text{carl})} \rangle &= \left\langle \mathbf{Y}_b, \mathbf{W}_b^{(\text{carl})} \right\rangle + \left\langle \mathbf{Y}_{e1}, \mathbf{W}_{e1}^{(\text{carl})} \right\rangle \\ &\quad + \left\langle \mathbf{Y}_{e2}, \mathbf{W}_{e2}^{(\text{carl})} \right\rangle \\ \text{and} \quad \|\mathbf{W}^{(\text{carl})}\|^2 &= \left\| \mathbf{W}_b^{(c)} \right\|^2 + \left\| \mathbf{W}_{e1}^{(\text{carl})} \right\|^2 + \left\| \mathbf{W}_{e2}^{(\text{carl})} \right\|^2. \end{aligned} \tag{11}$$

In (9)–(11), we use the subscript c to denote the collective detection statistics that use fingerprints extracted from all layers collectively to identify colluders.

With orthogonal fingerprint modulation, if the additive noise $\mathbf{n}$ is i.i.d. Gaussian $\mathcal{N}(0, \sigma_n^2)$, the collective detection statistics follow the normal distributions [33]

$$\begin{aligned} &TN_c^{(\text{alice})} \sim \mathcal{N}\left(\mu^{(a)}, \sigma_n^2\right) \quad \text{with} \quad \mu^{(a)} = \frac{\beta_1 \sqrt{N_b}}{K^b} \sigma_W, \\ &TN_c^{(\text{bob})} \sim \mathcal{N}\left(\mu^{(b)}, \sigma_n^2\right) \quad \text{with} \quad \mu^{(b)} = \frac{\beta_2 N_b + \alpha_1 N_{e1}}{K^{b,e1} \sqrt{N_b + N_{e1}}} \sigma_W, \\ &\text{and} \quad TN_c^{(\text{carl})} \sim \mathcal{N}\left(\mu^{(c)}, \sigma_n^2\right) \quad \text{with} \\ &\qquad \mu^{(c)} = \frac{\beta_3 N_b + \alpha_2 N_{e1} + N_{e2}}{K^{\text{all}} \sqrt{N_b + N_{e1} + N_{e2}}} \sigma_W. \end{aligned} \tag{12}$$

In (12), N_b, N_{e1} and N_{e2} are the lengths of the fingerprints embedded in the base layer, enhancement layer one and enhancement layer two, respectively, and σ_W^2 is the variance of the fingerprints $\mathbf{W}$. Detailed derivations are available in [33]. Therefore, Alice's probability of being detected is

$$P_s^{(\text{alice})} = Q\left(\frac{h - \mu^{(a)}}{\sigma_n}\right), \tag{13}$$

and $P_s^{(\text{bob})}$ and $P_s^{(\text{carl})}$, which are Bob's and Carl's probabilities of being detected, share the similar form. To guarantee that $P_s^{(\text{alice})} = P_s^{(\text{bob})} = P_s^{(\text{carl})}$ and ensure the equal risk of all colluders, it is equivalent to select $\{\alpha_k, \beta_l\}$ such that $\mu^{(a)} = \mu^{(b)} = \mu^{(c)}$. Table 1 [33] lists the constraints on collusion and the selection of collusion parameters to achieve equal-risk absolute fairness when generating a colluded copy of high, medium, and low resolutions, respectively.

UNDERSTANDING THE CONSTRAINTS ON COLLUSION TO ACHIEVE FAIRNESS

From Table 1, if the colluders wish to generate a high-resolution colluded copy while still achieving equal-risk absolute fairness, then $(K^b, K^{b,e1}, K^{\text{all}})$ and (N_b, N_{e1}, N_{e2}) have to satisfy

the constraints (*) in the row of highest resolution, and the colluders should select parameters as in (**). Similarly, if the colluders wish to generate a colluded copy of medium resolution, to achieve equal-risk absolute fairness, $(K^b, K^{b,e1}, K^{\text{all}})$ and (N_b, N_{e1}, N_{e2}) must satisfy the constraint (†) in the row of medium resolution, and the colluders should select the parameters according to (‡) therein. If the colluders only want to generate a low-resolution copy, there are no constraints on $(K^b, K^{b,e1}, K^{\text{all}})$ and (N_b, N_{e1}, N_{e2}), and they should follow (§§) to achieve equal-risk absolute fairness. If we compare the constraints (*), (†) and (§) in Table 1, it is easy to see that generating a colluded copy of higher resolution puts more severe constraints on collusion to guarantee that all colluders have the same risk of being detected.

To have a better visualization of fair collusion, Figure 6 shows an example of the constraints on collusion to ensure the equal risk of all colluders. Define $K = K^b + K^{b,e1} + K^{\text{all}}$ as the total number of colluders. In addition, we let $R^b = K^b/K$, $R^{b,e1} = K^{b,e1}/K$ and $R^{\text{all}} = K^{\text{all}}/K$ denote the percentages of colluders who receive the low-, medium- and high-resolution copies, respectively, and $R^b + R^{b,e1} + R^{\text{all}} = 1$. In Figure 6, the horizontal and the vertical axes are R^b and R^{all}, respectively, and each point in the figure corresponds to a unique triplet $(R^b, R^{b,e1}, R^{\text{all}})$ where $R^{b,e1} = 1 - R^b - R^{\text{all}}$.

In Figure 6(a), the line $\overline{AB}$ is defined as

$$\overline{AB} \triangleq \left\{ (R^b, R^{b,e1}, R^{\text{all}}) : \frac{R^{\text{all}}\sqrt{N_b + N_{e1} + N_{e2}}}{R^b\sqrt{N_b} + R^{b,e1}\sqrt{N_b + N_{e1}} + R^{\text{all}}\sqrt{N_b + N_{e1} + N_{e2}}} = \frac{N_{e2}}{N_b + N_{e1} + N_{e2}} \right\}, \tag{14}$$

which corresponds to the boundary of the second constraint in (*) in Table 1. The line $\overline{CD}$ is

$$\overline{CD} \triangleq \left\{ (R^b, R^{b,e1}, R^{\text{all}}) : \frac{R^b\sqrt{N_b}}{R^b\sqrt{N_b} + R^{b,e1}\sqrt{N_b + N_{e1}} + R^{\text{all}}\sqrt{N_b + N_{e1} + N_{e2}}} = \frac{N_b}{N_b + N_{e1} + N_{e2}} \right\}, \tag{15}$$

which is the boundary of the first constraint in (*) in Table 1. In Figure 6(b), the line $\overline{EF}$ is

$$\overline{EF} \triangleq \left\{ (R^b, R^{b,e1}, R^{\text{all}}) : \frac{R^b\sqrt{N_b}}{R^b\sqrt{N_b} + (R^{b,e1} + R^{\text{all}})\sqrt{N_b + N_{e1}}} = \frac{N_b}{N_b + N_{e1}} \right\}, \tag{16}$$

which is the boundary of the constraint (†) in Table 1 when colluders wish to generate a colluded copy of medium resolution.

From Table 1, if colluders wish to generate a high-resolution colluded copy, $(R^b, R^{b,e1}, R^{\text{all}})$ have to be in the shaded area shown in Figure 6(a) to guarantee that all attackers have the same probability of being detected. To generate a colluded copy of medium resolution, $(R^b, R^{b,e1}, R^{\text{all}})$ have to be in the shaded area shown in Figure 6(b) to ensure that colluders share the same risk. As we can see from Figure 6, generating a colluded copy of higher resolution puts more severe constraints on collusion to achieve

[TABLE 1] CONSTRAINTS ON COLLUSION AND SELECTION OF COLLUSION PARAMETERS TO ACHIEVE EQUAL RISK.

HIGHEST RESOLUTION WITH ALL THREE LAYERS	FAIRNESS CONSTRAINTS	$\begin{cases} \frac{K^b\sqrt{N_b}}{K^b\sqrt{N_b}+K^{b,e1}\sqrt{N_b+N_{e1}}+K^{\text{all}}\sqrt{N_b+N_{e1}+N_{e2}}} \le \frac{N_b}{N_b+N_{e1}+N_{e2}}, \\ \frac{K^{\text{all}}\sqrt{N_b+N_{e1}+N_{e2}}}{K^b\sqrt{N_b}+K^{b,e1}\sqrt{N_b+N_{e1}}+K^{\text{all}}\sqrt{N_b+N_{e1}+N_{e2}}} \ge \frac{N_{e2}}{N_b+N_{e1}+N_{e2}}. \end{cases}$ (*)
	PARAMETER SELECTION	$\begin{cases} \beta_1 = \frac{N_b+N_{e1}+N_{e2}}{N_b} \frac{K^b\sqrt{N_b}}{K^b\sqrt{N_b}+K^{b,e1}\sqrt{N_b+N_{e1}}+K^{\text{all}}\sqrt{N_b+N_{e1}+N_{e2}}}, \\ \beta_2 N_b + \alpha_1 N_{e1} = \frac{(N_b+N_{e1}+N_{e2})K^{b,e1}\sqrt{N_b+N_{e1}}}{K^b\sqrt{N_b}+K^{b,e1}\sqrt{N_b+N_{e1}}+K^{\text{all}}\sqrt{N_b+N_{e1}+N_{e2}}}, \\ \beta_3 = 1-\beta_1-\beta_2,\ \alpha_2 = 1-\alpha_1. \end{cases}$ (**)
MEDIUM RESOLUTION WITH THE BASE LAYER AND THE ENHANCEMENT LAYER ONE	FAIRNESS CONSTRAINTS	$\frac{K^b\sqrt{N_b}}{K^b\sqrt{N_b}+(K^{b,e1}+K^{\text{all}})\sqrt{N_b+N_{e1}}} \le \frac{N_b}{N_b+N_{e1}}.$ (†)
	PARAMETER SELECTION	$\begin{cases} \beta_1 = \frac{N_b+N_{e1}}{N_b} \frac{K_b\sqrt{N_b}}{K^b\sqrt{N_b}+(K^{b,e1}+K^{\text{all}})\sqrt{N_b+N_{e1}}}, \\ \beta_2 = \frac{K^{b,e1}}{K^{b,e1}+K^{\text{all}}}(1-\beta_1),\ \beta_3 = 1-\beta_1-\beta_2, \\ \alpha_1 = \frac{K^{b,e1}}{K^{b,e1}+K^{\text{all}}},\ \alpha_2 = 1-\alpha_1. \end{cases}$ (‡)
LOWEST RESOLUTION WITH THE BASE LAYER ONLY	FAIRNESS CONSTRAINTS	NO CONSTRAINTS ON $(K^b, K^{b,e1}, K^{\text{all}})$ AND (N_b, N_{e1}, N_{e2}). (§)
	PARAMETER SELECTION	$\beta_1 = \frac{K^b}{K^b+K^{b,e1}+K^{\text{all}}},\quad \beta_2 = \frac{K^{b,e1}}{K^b+K^{b,e1}+K^{\text{all}}},\quad \beta_3 = \frac{K^{\text{all}}}{K^b+K^{b,e1}+K^{\text{all}}}.$ (§§)

equal-risk fairness, and it requires that more colluders receive the high-resolution copies from the content owner.

GAME-THEORETIC MODELING OF COLLUDER DYNAMICS

Equal-risk absolute fairness only considers each colluder's risk and ensures that all colluders have the same probability of being detected. During collusion, colluders not only negotiate how to distribute the risk but also bargain how to share the rewards from the illegal usage of multimedia. In addition, rather than absolute fairness, colluders may prefer other ways to distribute the risk and the reward. For example, some colluders may want to benefit more from collusion by taking a higher risk of being detected. In [34], this complex dynamics was modeled as a bargaining problem where colluders negotiate with each other to resolve the conflict, and game theory [7] was used to analyze this negotiation process.

In this game-theoretic framework, colluders first define the utility (payoff) function π, which is a function of a colluder's risk as well as the reward that he or she receives from collusion. A natural definition of the utility function is the expected payoff that a colluder receives by participating in collusion. For colluder $\mathbf{u}^{(i)}$, his or her utility can be given by

$$\pi^{(i)} = -P_s^{(i)} L^{(i)} + (1 - P_s^{(i)}) Rw^{(i)}, \tag{17}$$

where $P_s^{(i)}$ is his or her probability of being detected, $L^{(i)}$ is colluder $\mathbf{u}^{(i)}$'s loss if he or she is captured by the fingerprint detected, and $Rw^{(i)}$ is the reward that $\mathbf{u}^{(i)}$ receives if he or she successfully escapes being detected. Each colluder tries to maximize his or her own utility function during the negotiation process.

Without loss of generality, we use a two-layer multimedia fingerprinting example to demonstrate how colluders bargain during collusion. We assume that there are a total of 250 colluders of which 80 attackers receive the low-resolution copies and the rest have the high-resolution version. For simplicity, we consider a scenario where colluders who receive fingerprinted copies of the same resolution agree to share the same risk and have equal utilities. Therefore, colluders who receive the low-resolution copies act as a single player in the game and they have the same utility π^b, while colluders who have the high-resolution copies act as a single player during the bargaining process and they have the same utility π^{be}.

The second step in the bargaining process is to find the feasible set $\mathbf{S} = \{(\pi^b, \pi^{be}) \in \mathbb{R}^2\}$ of the game, where for every $(\pi^b, \pi^{\text{be}}) \in \mathbf{S}$, it is possible for colluders to act together and obtain the utilities π^b and π^{be}, respectively. For the above mentioned colluder game, Figure 7 shows the feasible set, which is the curve **AB** plus the line **BC**. Note that if colluders select a solution that corresponds to a point on the line **BC**, then they can always find another solution that gives the same π^{be} but a larger π^b. Therefore, in a bargaining situation like this, colluders would always like to settle at a Pareto-optimal point, where no one can further increase his or her utility without decreasing others'. In Figure 7, the Pareto-optimal set includes solutions that correspond to the points on the curve **AB**.

Depending on their definition of fairness and their objectives of collusion, colluders select different collusion strategies. For example, with equal-payoff absolute fairness, colluders select the point where $\pi^b = \pi^{\text{be}}$ and let all attackers have the same utility. Colluders can also select the collusion parameters to maximize the minimum utility that a colluder can receive by participating in collusion, that is,

$$\pi^* = \max_{\beta} \min \left\{ \pi^b, \pi^{\text{be}} \right\}, \tag{18}$$

where β is the collusion parameter in Figure 5. This solution guarantees that by participating in collusion, a colluder can receive at least π^* utilities. The maxsum solution maximizes the sum of all attackers' utilities if they cooperate with each other during collusion. Another popular solution in game theory is the famous Nash bargaining solution

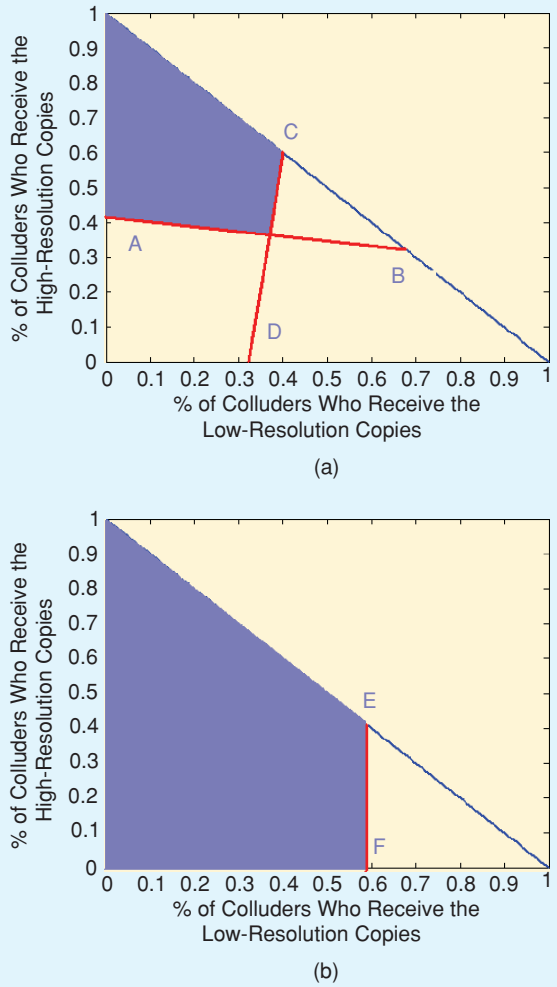

[FIG6] An example of the constraints on collusion to achieve equal-risk absolute fairness. (a) Generating a colluded copy of high resolution. (b) Generating a colluded copy of medium resolution $N_b : N_{e1} : N_{e2} = 1 : 1 : 2$.

(NBS), which aims to achieve proportional fairness. It divides the additional utility between the two players in a ratio that is equal to the rate at which this utility can be transferred [7]. Mathematically, the NBS maximizes

$$g(\pi^b, \pi^{be}) = (\pi^b - \pi^{b*})(\pi^{be} - \pi^{be*}),$$
$$\text{where} \quad \pi^{b*} = \min_{\beta}\{\pi^b\} \quad \text{and} \quad \pi^{be*} = \min_{\beta}\{\pi^{be}\}. \quad (19)$$

Different collusion strategies correspond to different points in the Pareto-optimal set. In the example shown in Figure 7, the equal-payoff absolute fairness and the maxmin strategies give the same result, while the maxsum and the NBSs favor colluders who receive the high-resolution fingerprinted copies more.

CHEATING BEHAVIOR IN MULTIMEDIA SOCIAL NETWORKS

In multimedia social networks, users need to exchange private information with each other when negotiating, and achieving fairness requires that they give each other correct information about their own resources. However, the assumption of fair play may not always hold. Although they might agree so, some users might be selfish and wish to maximize their own payoff. To achieve this goal, they might break away from their agreement and cheat other users during the bargaining process. To improve the overall system performance, it is important to study the cheating and cheat-proofing dynamics among users, investigate the selfish colluders' cheating strategies, and design cheat-proofing mechanisms. In this article, we use multiuser collusion as an example to understand the colluders' cheating and cheat-proofing strategies and study the traitor-within-traitor problem.

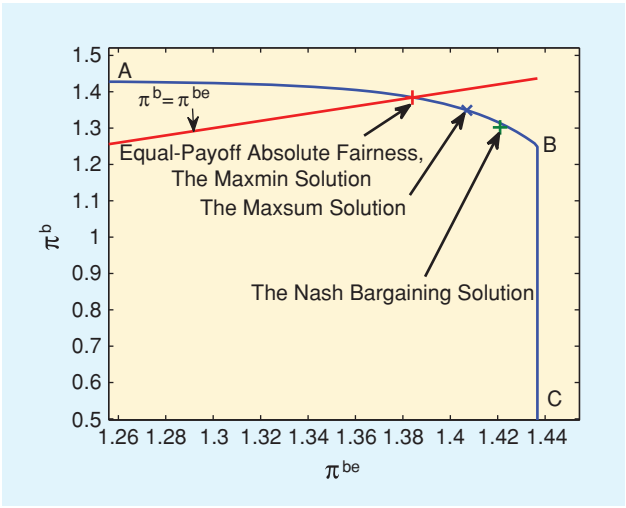

[FIG7] An example of the feasible set and different solutions of the colluder game. The horizontal axis is the utility of colluders who receive the high-resolution copies, and the vertical axis is the utility of colluders who receive the low-resolution copies.

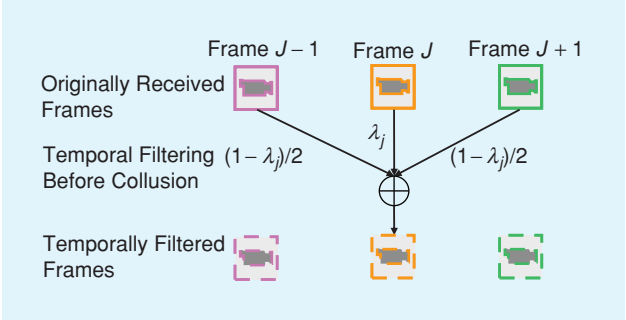

[FIG8] Temporal filtering before multiuser collusion.

In multiuser collusion, colluders need to exchange private information (that is, the resolution of the fingerprinted copies and the fingerprinted coefficients in each frame) with each other to ensure fairness of the attack. Without loss of generality, we use equal-risk absolute fairness as an example. In this scenario, colluders agree to distribute the risk evenly among themselves, while selfish colluders wish to minimize their own probability of being detected. To achieve this goal, selfish colluders process their fingerprinted copies before collusion and contributes the processed copy instead of the originally received ones during collusion. In this section, we focus on the analysis of selfish colluders' cheating strategies and demonstrate a few techniques that selfish colluders can use to minimize their own risk.

RISK MINIMIZATION AND TRADEOFF

For selfish colluders, in order to further lower their risk, one possible solution is to attenuate the energy of the embedded fingerprints even before multiuser collusion. Examples include averaging or swapping neighboring frames to replace each segment of the fingerprinted signal with another, seemingly similar segment from different regions of the content [35]–[37].

For example, consider frame averaging where the selfish colluder uses linear interpolation to generate a temporally filtered and smoothed video. Assume that for colluder $\mathbf{u}^{(i)}$, $\mathbf{X}_j^{(i)}$ is his or her originally received fingerprinted frame j. As shown in Figure 8, for each frame (say j) in the video sequence, the selfish colluder replaces it with a linear combination of the current frame (j), the previous frame ($j-1$) and the next frame ($j+1$) with weights $\{\lambda_j, 1-\lambda_j/2, 1-\lambda_j/2\}$, respectively, and generates a new frame

$$\widetilde{\mathbf{X}}_j^{(i)} = \frac{1-\lambda_j}{2}\mathbf{X}_{j-1}^{(i)} + \lambda_j \mathbf{X}_j^{(i)} + \frac{1-\lambda_j}{2}\mathbf{X}_{j+1}^{(i)}. \quad (20)$$

The selfish colluder repeats this process for all frames in the video sequence and different frames are processed independently during precollusion processing. During collusion, the selfish colluder contributes the temporally filtered copy $\widetilde{\mathbf{X}}_j^{(i)}$ instead of the originally received one $\mathbf{X}_j^{(i)}$. If other colluders do not discover this temporal filtering, same as in the previous section, they average all the fingerprinted copies that they have and add additional noise $\mathbf{n}$ to further hinder the detection process.

During precollusion processing, the selfish colluder wishes to minimize his or her chance of being detected by the fingerprint detector. Meanwhile, temporal filtering should introduce

as little perceptually noticeable distortion as possible to his or her fingerprinted copy. To select the appropriate parameter λ_j in (20), the selfish colluder needs to analyze how temporal filtering changes his or her risk of being detected by the digital rights enforcer and study how it affects the perceptual quality of his or her fingerprinted copy.

From the analysis in [38], with orthogonal fingerprint modulation, if the additive noise $\mathbf{n}$ is i.i.d. Gaussian $\mathcal{N}(0, \sigma_n^2)$, the selfish colluder $\mathbf{u}^{(i)}$'s detection statistic $TN^{(i)}$ follows the normal distribution

$$TN^{(i)} \sim \mathcal{N}\left(\mu^{(i)}, \sigma_n^2\right), \quad \text{where}$$

$$\mu^{(i)} = \sum_j \frac{\left\langle \mathbf{W}_{j-1}^{(i)}, \mathbf{W}_j^{(i)} \right\rangle + \left\langle \mathbf{W}_j^{(i)}, \mathbf{W}_{j+1}^{(i)} \right\rangle}{2K\sqrt{\sum_l \left\| \mathbf{W}_l^{(i)} \right\|^2}} + \sum_j \lambda_j \times \frac{2\left\| \mathbf{W}_j^{(i)} \right\|^2 - \left\langle \mathbf{W}_{j-1}^{(i)}, \mathbf{W}_j^{(i)} \right\rangle - \left\langle \mathbf{W}_j^{(i)}, \mathbf{W}_{j+1}^{(i)} \right\rangle}{2K\sqrt{\sum_l \left\| \mathbf{W}_l^{(i)} \right\|^2}}, \tag{21}$$

and K is the total number of colluders. Therefore, $\mathbf{u}^{(i)}$'s probability of being detected is

$$P_s^{(i)} = Q\left((h - \mu^{(i)})/\sigma_n\right), \tag{22}$$

where $Q(\cdot)$ is the Gaussian tail function and h is a predetermined threshold. Since

$$\left\langle \mathbf{W}_{j-1}^{(i)}, \mathbf{W}_j^{(i)} \right\rangle \le \left\langle \mathbf{W}_j^{(i)}, \mathbf{W}_j^{(i)} \right\rangle = \left\| \mathbf{W}_j^{(i)} \right\|^2 \quad \text{and} \quad \left\langle \mathbf{W}_{j+1}^{(i)}, \mathbf{W}_j^{(i)} \right\rangle \le \left\| \mathbf{W}_j^{(i)} \right\|^2, \tag{23}$$

$\mu^{(i)}$ is a nondecreasing function of λ_j and is minimized when $\lambda_j = 0$ for all j. This is because, by using a smaller λ_j, the interpolated frame $\widetilde{\mathbf{X}}_j^{(i)}$ contains less information of the originally received frame $\mathbf{X}_j^{(i)}$ and, therefore, attenuates the fingerprints embedded in frame j by a larger amount. Thus, from risk minimization's point of view, smaller values of $\{\lambda_j\}$ are preferred.

To analyze how temporal filtering affects the perceptual quality, we calculate the mean square error (MSE) between the filtered frame $\widetilde{\mathbf{X}}_j^{(i)}$ and the originally received one $\mathbf{X}_j^{(i)}$, which is

$$\begin{aligned} \text{MSE}_j &= \left\| \widetilde{\mathbf{X}}_j^{(i)} - \mathbf{X}_j^{(i)} \right\|^2 = \left(\frac{1-\lambda_j}{2}\right)^2 \cdot \phi_j, \\ \text{where} \quad \phi_j &= 4\left\| \mathbf{X}_j^{(i)} \right\|^2 + \left\| \mathbf{X}_{j-1}^{(i)} \right\|^2 + \left\| \mathbf{X}_{j+1}^{(i)} \right\|^2 \\ &\quad - 4\left\langle \mathbf{X}_{j-1}^{(i)}, \mathbf{X}_j^{(i)} \right\rangle - 4\left\langle \mathbf{X}_j^{(i)}, \mathbf{X}_{j+1}^{(i)} \right\rangle \\ &\quad + 2\left\langle \mathbf{X}_{j-1}^{(i)}, \mathbf{X}_{j+1}^{(i)} \right\rangle. \end{aligned} \tag{24}$$

From (24), a larger λ_j implies a smaller MSE and better quality. When $\lambda_j = 1$, $\widetilde{\mathbf{X}}_j^{(i)} = \mathbf{X}_j^{(i)}$ and it corresponds to the scenario where $\mathbf{u}^{(i)}$ does not process his or her copy before collusion. Therefore, from the perceptual quality's point of view, $\mathbf{u}^{(i)}$ should choose a larger λ_j.

To address such tradeoff between the risk and the perceptual quality, the selfish colluder selects the parameters $\{\lambda_j\}$ to minimize his or her chance of being detected by the digital rights enforcer under the constraint that the MSE between the temporally filtered copy $\widetilde{\mathbf{X}}_j^{(i)}$ and the originally received one $\mathbf{X}_j^{(i)}$ is below a predetermined threshold ε. Therefore, for a selfish colluder $\mathbf{u}^{(i)}$, the selection of the parameter $\{\lambda_j\}$ can be modeled as

$$\min_{\{\lambda_j\}} \left\{ \mu^{(i)} = \sum_j \mu_j^{(i)} \right\}$$
$$s.t. \qquad \text{MSE}_j \le \varepsilon, \ 0 \le \lambda_j \le 1, \ j = 1, 2, \ldots, \tag{25}$$

From [38], the solution to the above optimization problem is: for every frame j,

$$\lambda_j^* = \max\left\{0, \ 1 - 2\sqrt{\varepsilon/\phi_j}\right\}, \tag{26}$$

where ϕ_j is in (24). By using $\{\lambda_j^*\}$ during temporal filtering, a selfish colluder minimizes his or her own probability of being detected and ensures that the newly generated frames have small perceptual distortion when compared with the originally received ones (the MSE between these two is no larger than ε).

TRAITOR-WITHIN-TRAITOR DYNAMICS

In addition to temporal filtering, the selfish colluder can use a wide range of techniques to process his or her fingerprinted copy before multiuser collusion. For example, when colluders receive fingerprinted copies of different resolutions, colluders first need to estimate every attacker's probability of being detected. Then they follow the analysis in the section "Fairness Dynamics in Multimedia Social Networks" to select the parameters such that all colluders share the same risk. The assumption here is that all colluders are honest about the resolutions of their fingerprinted copies, and they can correctly estimate each other's risk. Thus, to further reduce their risk, one possible option for selfish colluders is to lie about (for example, change) the resolutions of their copies before collusion [38].

Without loss of generality, we use three-layer temporal scalability as an example and consider a selfish colluder $\mathbf{u}^{(i)}$ who receives a low-resolution copy with the base layer only. During precollusion processing, $\mathbf{u}^{(i)}$ can interpolate the base-layer frames and generate the missing frames in the two enhancement layers. Assume that $\mathbf{X}_{j_1}^{(i)}$ and $\mathbf{X}_{j_3}^{(i)}$ are two adjacent frames in the base layer that $\mathbf{u}^{(i)}$ receives. To forge a frame j_2 in the enhancement layers where $j_1 < j_2 < j_3$, $\mathbf{u}^{(i)}$ can use a simple linear-interpolation-based method and let

$$\widetilde{\mathbf{X}}_{j_2}^{(i)} = \lambda_1 \cdot \mathbf{X}_{j_1}^{(i)} + \lambda_2 \cdot \mathbf{X}_{j_3}^{(i)},$$

$$\text{where} \quad \lambda_1 = \frac{j_3 - j_2}{j_3 - j_1}, \quad \text{and} \quad \lambda_2 = \frac{j_2 - j_1}{j_3 - j_1}. \tag{27}$$

Other complicated algorithms, e.g., motion-based interpolation [39], can also be used.

To analyze the effectiveness of this precollusion processing in reducing $\mathbf{u}^{(i)}$'s risk, we consider two scenarios: when the selfish colluder does not apply precollusion processing and when $\mathbf{u}^{(i)}$ increases the temporal resolution of his or her copy before collusion, and we compare the selfish colluder's probability of being detected in these two scenarios.

IN A MULTIMEDIA SOCIAL NETWORK COMMUNITY, A GROUP OF USERS FORM A DYNAMICALLY CHANGING NETWORK INFRASTRUCTURE TO SHARE AND EXCHANGE DATA, OFTEN MULTIMEDIA CONTENT, AS WELL AS OTHER RESOURCES.

Scenario 1: Without Precollusion Processing

We first consider the scenario when $\mathbf{u}^{(i)}$ does not process his or her copy and contributes the originally received frames during collusion. In this scenario, the analysis is the same as that in the section "Equal-Risk Absolute Fairness." Since $\mathbf{u}^{(i)}$ receives the base layer only, other colluders believe that the fingerprint detector will use the fingerprint extracted from the base layer only to determine if $\mathbf{u}^{(i)}$ participates in collusion. Thus, following the same analysis as in the section "Equal-Risk Absolute Fairness," colluders calculate that $\mathbf{u}^{(i)}$'s risk of being detected is

$$P_s^{(i)} = Q\left(\frac{h - \beta_1\sqrt{N_b}\sigma_W/K^b}{\sigma_n}\right), \tag{28}$$

where h is a predetermined threshold. Then, they follow Table 1 and select $\{\alpha_k, \beta_l\}$ such that $P_s^{(i)}$ is the same as other colluders' probability of being detected. Since $\mathbf{u}^{(i)}$ does not process his or her copy before collusion, other colluders correctly estimate $\mathbf{u}^{(i)}$'s risk, and $P_s^{(i)}$ in (28) is $\mathbf{u}^{(i)}$'s probability of being detected by the fingerprint detector.

Scenario 2: With Precollusion Processing

We then consider the scenario where $\mathbf{u}^{(i)}$ increases the frame rate before multiuser collusion. If other colluders do not discover this cheating behavior, they still believe that the fingerprint detector will use fingerprints extracted from all layers collectively to determine if $\mathbf{u}^{(i)}$ participates in collusion. Based on this assumption, they follow the same analysis as in the section "Equal-Risk Absolute Fairness" and calculate that $\mathbf{u}^{(i)}$'s risk of being detected is

$$\bar{P}_s^{(i)} = Q\left(\frac{h - \bar{\mu}^{(i)}}{\sigma_n}\right), \quad \text{where}$$

$$\bar{\mu}^{(i)} = \frac{\bar{\beta}_3 N_b + \bar{\alpha}_2 N_{e1} + N_{e2}}{\bar{K}^{\text{all}}\sqrt{N_b + N_{e1} + N_{e2}}}\sigma_W. \tag{29}$$

In (29), $\bar{K}^{\text{all}}$ is the total number of colluders who contribute high-resolution fingerprinted copies, including $\mathbf{u}^{(i)}$. Then, they follow Table 1 and select $\{\bar{\alpha}_k, \bar{\beta}_l\}$ such that $\bar{P}_s^{(i)}$ is the same as other colluders' probability of being detected.

However, the fingerprint detector knows that $\mathbf{u}^{(i)}$ receives the fingerprinted base layer only from the content owner. The fingerprint detector believes that if $\mathbf{u}^{(i)}$ is a colluder, the enhancement layers of the colluded copy should not contain $\mathbf{u}^{(i)}$'s identification information. Therefore, the fingerprint detector only uses the fingerprint extracted from the base layer to decide if $\mathbf{u}^{(i)}$ is a colluder. In this case, following the same analysis as in the section "Equal-Risk Absolute Fairness," $\mathbf{u}^{(i)}$'s true probability of being detected is

$$\tilde{P}_s^{(i)} = Q\left(\frac{h - \bar{\beta}_3\sqrt{N_b}\sigma_W/\bar{K}^{\text{all}}}{\sigma_n}\right). \tag{30}$$

Comparing (29) and (30), $\tilde{P}_s^{(i)}$ in (30) does not equal to and is smaller than $\bar{P}_s^{(i)}$ in (29). Other colluders make an error when estimating $P_s^{(i)}$ due to $\mathbf{u}^{(i)}$'s precollusion processing. This estimation error helps $\mathbf{u}^{(i)}$ further lower his or her probability of being detected.

To analyze how precollusion processing affects other colluders, with $\mathbf{u}^{(i)}$'s precollusion processing, for a colluder $\mathbf{u}^{(k)}$ who contributes the originally received copy, following the same analysis, $\mathbf{u}^{(k)}$'s chance of being detected is

$$\tilde{P}_s^{(k)} = Q\left(\frac{h - \bar{\beta}_1\sqrt{N_b}\sigma_W/\bar{K}^b}{\sigma_n}\right). \tag{31}$$

Using the above example, Figure 9 shows the effectiveness of precollusion processing in reducing the selfish colluder's risk. We assume that there are a total of $K = 150$ colluders. Each point on the horizontal axis corresponds to a unique triplet $(K^b, K^{b,e1}, K^{\text{all}})$ on the line $\overline{AB}$ in (14), In Figure 9, we assume that there is only one selfish colluder $\mathbf{u}^{(i)}$ and other colluders do not discover his or her selfish behavior. Figure 9(a) compares $P_s^{(i)}$ in (28) with $\tilde{P}_s^{(i)}$ in (30), that is, $\mathbf{u}^{(i)}$'s probability of being detected with and without precollusion processing. It is obvious that increasing the resolution of the fingerprinted copy can help $\mathbf{u}^{(i)}$ further decrease his or her risk. In Figure 9(b), we plot $\tilde{P}_s^{(i)}$ in (30) and $\tilde{P}_s^{(k)}$ in (31), and we compare the selfish colluder $\mathbf{u}^{(i)}$'s risk with that of another colluder $\mathbf{u}^{(k)}$ who does not apply precollusion processing. It shows that $\mathbf{u}^{(i)}$'s precollusion processing makes others take a much higher probability of being detected and thus increases others' relative risk when compared with $\mathbf{u}^{(i)}$.

Similarly, if the selfish colluder receives not only the base layer but also the enhancement layers, he or she can also drop the enhancement layers and contribute only the low-resolution

copy during collusion. Interested readers can refer to detailed analysis in [38], where it showed that changing the resolution of the fingerprinted copies can help selfish colluders further reduce their probability of being detected, especially when the colluded copy has high resolution. In fact, in some scenarios, precollusion not only increases other colluders' relative risk when compared with that of the selfish colluders, but it may also increase others' absolute risk, that is, their probability of being detected. Therefore, it is not only selfish but also malicious.

CHEAT-PROOFING STRATEGIES IN MULTIMEDIA SOCIAL NETWORKS

In multimedia social networks, due to the selfish nature of human behavior, honestly reporting private information cannot be taken for granted and some users might intentionally cheat others to maximize their own payoff. Therefore, it is important to have cheat-proofing strategies to protect one's own interests. A social network may have different social structures and, therefore, can result in different cheat-proofing strategies. In a centralized multimedia social network where there exists at least one trusted entity (or leader), he or she can help monitor (maintain the order) and detect cheating behavior. However, in a distributed structure where there exists no such trusted entity, users have to detect cheating behavior and identify selfish users themselves. In this section, we consider the development of cheat-proofing strategies, and without loss of generality, we use traitors within traitors in multimedia fingerprinting as an example to illustrate the dynamics.

As we can see from the previous section, in multiuser collusion, precollusion processing is not only a selfish behavior, but can also be a malicious one. To protect their own interests during collusion, it is important for colluders to have cheat-proofing strategies. They must examine all the fingerprinted copies before collusion, detect and identify selfish colluders, and exclude them from collusion. It forces all colluders to keep their fair-play agreement and build trust among attackers. Let us use temporal filtering as an example of the selfish colluders' cheating strategies to illustrate the techniques to detect such temporal filtering and identify selfish colluders who deviate from their agreement. In this section, we first consider a centralized colluder social network with a ringleader whom all colluders trust and investigate how the trusted ringleader can help identify selfish colluders. We then study autonomous selfish colluder identification in the distributed colluder social networks that do not have trusted ringleaders.

The selfish colluder identification scheme should accurately identify all selfish colluders without falsely accusing any others. In addition, note that before deciding with whom to collude, colluders are unwilling to give others copies that contain their identification information. Therefore, selfish colluder identification should also protect the secrecy of all the fingerprinted copies and prevent colluders from accessing the fingerprinted coefficients in others' copies. To meet such an antiframing requirement, all copies must be encrypted appropriately during the selfish identification process. Thus, a challenging issue here is how colluders can detect precollusion processing and identify selfish colluders without knowing the fingerprinted coefficients in others' copies.

CENTRALIZED SOCIAL NETWORKS WITH TRUSTED RINGLEADERS

In the centralized colluder social networks, there exists a ringleader trusted by all colluders. They trust that the ringleader will not give their fingerprinted copies to others, will not frame any colluders, and will not modify the selfish colluder detection and identification results. In this scenario, all colluders give their fingerprinted copies to the ringleader, and the ringleader enforces the collusion by helping them detect selfish behavior.

Accurate identification of selfish colluders requires thorough study of how precollusion processing modifies the fingerprinted signals. Assume that S_j is the original frame j in the video sequence, and $\mathbf{W}^{(i)}$ is user $\mathbf{u}^{(i)}$'s fingerprint that is embedded in frame j. With spread spectrum embedding [17],

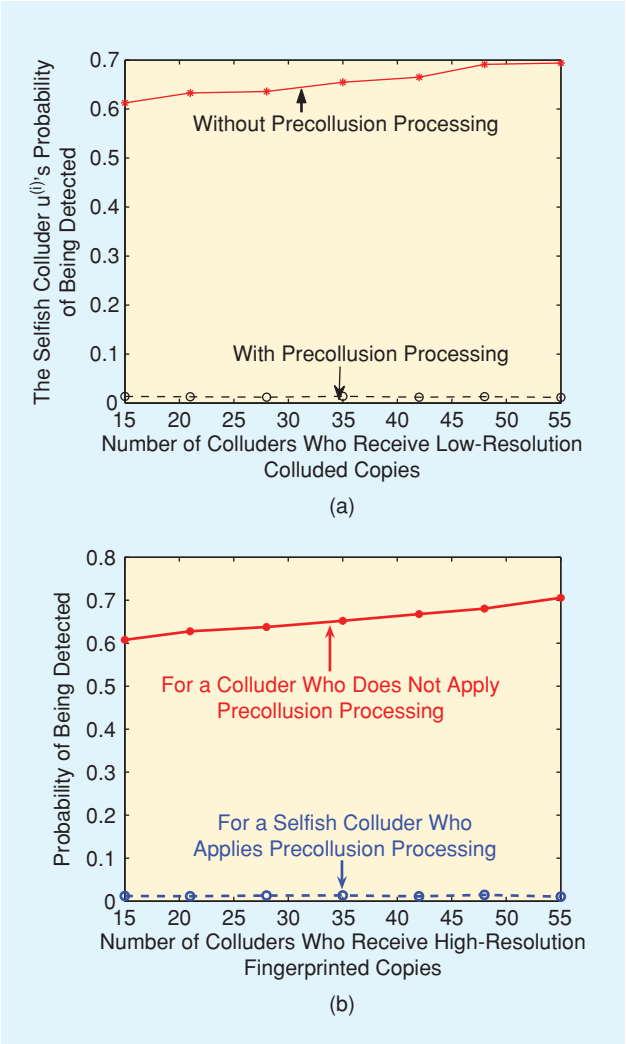

[FIG9] (a) Comparison of the selfish colluder $\mathbf{u}^{(i)}$'s probability of being detected with and without precollusion processing. (b) Comparison of different colluders' probabilities of being detected when $\mathbf{u}^{(i)}$ applies precollusion processing.

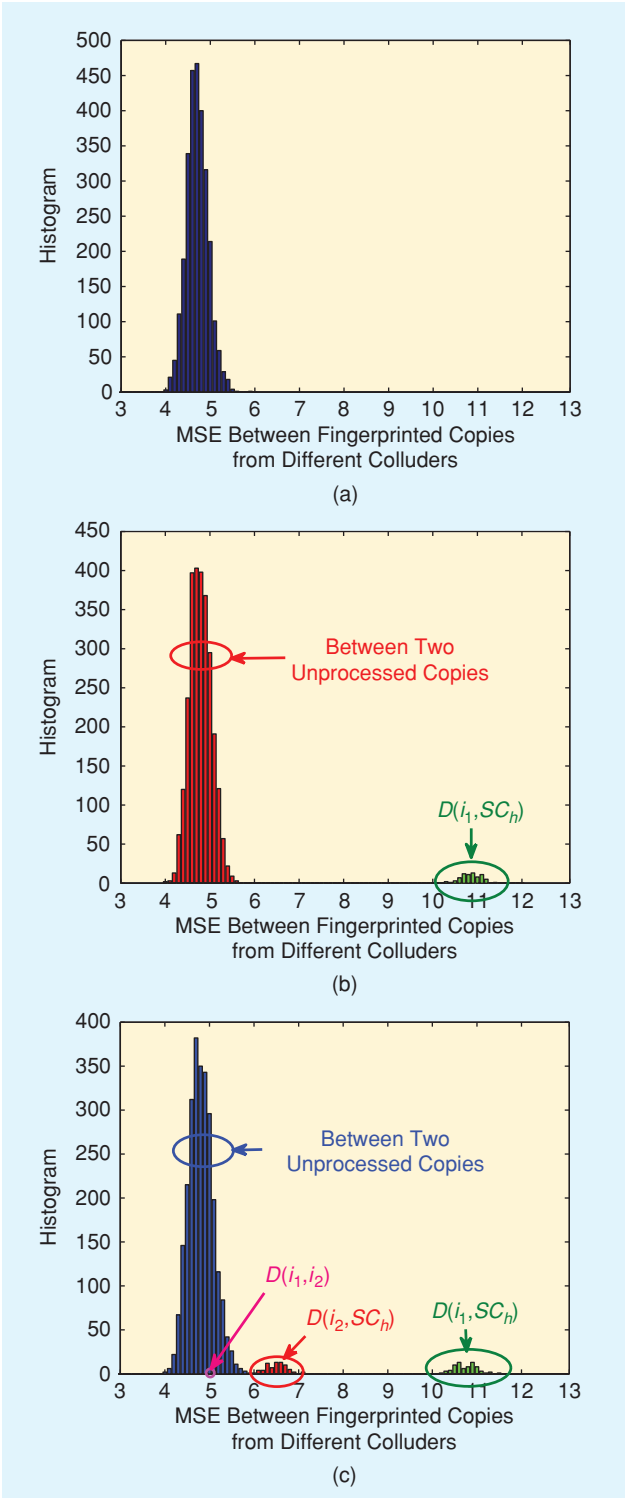

[FIG10] Histograms of $\{D(k,l)\}$. (a) All colluders give each other correct information about their fingerprinted copies. (b) There is one selfish colluder, $\mathbf{u}^{(i_1)}$, who temporally filters his or her copy before multiuser collusion. (c) There are two selfish colluders, $\mathbf{u}^{(i_1)}$ and $\mathbf{u}^{(i_2)}$, who process their copies before multiuser collusion. $\mathbf{u}^{(i_1)}$ chooses $\lambda_j = 0.6031$ in (20), and $\mathbf{u}^{(i_2)}$ selects $\lambda_j = 0.7759$ in (20). SC_h contains the indices of all colluders who do not process their copies before collusion. In (b) and (c), $\mathfrak{D}(i_1, SC_h) = \{D(i_1, l) : l \in SC_h\}$ and $\mathfrak{D}(i_2, SC_h) = \{D(i_2, l) : l \in SC_h\}$.

[18], for three colluders Alice, Bob, and Carl, their received fingerprinted frames are

$$\mathbf{X}_j^{(\text{alice})} = \mathbf{S}_j + \mathbf{W}_j^{(\text{alice})}, \quad \mathbf{X}_j^{(\text{bob})} = \mathbf{S}_j + \mathbf{W}_j^{(\text{bob})}, \quad \text{and}$$
$$\mathbf{X}_j^{(\text{carl})} = \mathbf{S}_j + \mathbf{W}_j^{(\text{carl})}, \tag{32}$$

respectively. (We drop the JND term here to simplify the notations.) Alice and Bob do not process their copies and contribute $\mathbf{X}_j^{(\text{alice})}$ and $\mathbf{X}_j^{(\text{bob})}$ during collusion. Carl uses (20) to temporally filter his copy and contributes

$$\begin{aligned}\widetilde{\mathbf{X}}_j^{(\text{carl})} &= \frac{1-\lambda_j}{2}\mathbf{X}_{j-1}^{(\text{carl})} + \lambda_j \mathbf{X}_j^{(\text{carl})} + \frac{1-\lambda_j}{2}\mathbf{X}_{j+1}^{(\text{carl})} \\ &= \mathbf{S}_j + \Delta\mathbf{S}_j(\lambda_j) + \widetilde{\mathbf{W}}_j^{(\text{carl})}, \\ \text{where} \quad \Delta\mathbf{S}_j(\lambda_j) &= (1-\lambda_j)\left(\frac{\mathbf{S}_{j-1}}{2} + \frac{\mathbf{S}_{j+1}}{2} - \mathbf{S}_j\right), \\ \text{and} \quad \widetilde{\mathbf{W}}_j^{(\text{carl})} &= \frac{1-\lambda_j}{2}\mathbf{W}_{j-1}^{(\text{carl})} + \lambda_j \mathbf{W}_j^{(\text{carl})} \\ &\quad + \frac{1-\lambda_j}{2}\mathbf{W}_{j+1}^{(\text{carl})}.\end{aligned} \tag{33}$$

From (33), temporal filtering not only averages fingerprints embedded in adjacent frames and attenuates their energies, it also filters neighboring host frames and introduces extra distortion $\Delta\mathbf{S}_j(\lambda_j)$ into the host signal.

Define $D(a,b) = \|\mathbf{X}^{(\text{alice})} - \mathbf{X}^{(\text{bob})}\|^2$ and $D(a,c) = \|\mathbf{X}^{(\text{alice})} - \mathbf{X}^{(\text{carl})}\|^2$, where $\|x\|^2$ returns the Euclidean norm of x. From (33), we have

$$\begin{aligned} D_j(a,b) &\approx \left\|\mathbf{W}_j^{(\text{alice})} - \mathbf{W}_j^{(\text{bob})}\right\|^2, \\ \text{and} \quad D_j(a,c) &\approx \left\|\mathbf{W}_j^{(\text{alice})} - \widetilde{\mathbf{W}}_j^{(\text{carl})}\right\|^2 + \|\Delta\mathbf{S}_j(\lambda_j)\|^2, \\ \text{where} \quad \|\Delta\mathbf{S}_j(\lambda_j)\|^2 &= (1-\lambda_j)^2 \times \left\|\frac{\mathbf{S}_{j-1}}{2} + \frac{\mathbf{S}_{j+1}}{2} - \mathbf{S}_j\right\|^2. \end{aligned} \tag{34}$$

As can be seen from (33), $D_j(a,c)$ has a much larger value than $D_j(a,b)$ since $D_j(a,c)$ also includes the extra distortion $\|\Delta\mathbf{S}_j(\lambda_j)\|^2$ due to temporal filtering of adjacent host frames. The difference between $D_j(a,b)$ and $D_j(a,c)$ is more obvious when λ_j takes a smaller value and when the difference between adjacent frames is larger (for example, when the scene of the host video sequence changes fast).

Figure 10 shows examples of the histograms of $\{D(k,l)\}$. As shown in Figure 10(a), when all colluders keep their fair-collusion agreement and give each other correct information of their received copies, all $\{D(k,l)\}$ are from the same distribution with a single mean. On the contrary, if there are selfish colluders who temporally filter their fingerprinted copies before collusion, then $\{D(k,l)\}$ are from two or more distributions with distinct means, as shown in Figure 10(b) and (c). Therefore, study of $\{D(k,l)\}$'s histogram plot can help detect

the existence of selfish colluders. Identification of the selfish colluders requires detailed examination of $D(k,l)$ for each pair of colluders $(\mathbf{u}^{(k)}, \mathbf{u}^{(l)})$. For example, in Figure 10(b), analysis of each individual $D(k,l)$, in particular, those in $\mathfrak{D}(i_1, SC_h)$, will help separate colluders into two subgroups: one includes the selfish colluder $\mathbf{u}^{(i_1)}$, and the other contains those who keep their fair-collusion agreement and contribute the originally received copies. In Figure 10(c), analysis of the two distributions on the right side, $\mathfrak{D}(i_1, SC_h)$ and $\mathfrak{D}(i_2, SC_h)$, can help identify $\mathbf{u}^{(i_1)}$ and $\mathbf{u}^{(i_2)}$ as the selfish colluders. The selfish colluder identification algorithm proposed in [40] can accurately identify all selfish colluders without falsely accusing any others. In addition, for each copy, only the corresponding colluder and the trusted ringleader can access the fingerprinted coefficients. Therefore, the selfish colluder detection and identification algorithm in [40] also protects the secrecy of all the fingerprinted copies and prevents framing attacks.

COLLUSION ATTACKS POSE SERIOUS THREATS TO MULTIMEDIA INTELLECTUAL PROPERTY RIGHTS.

DISTRIBUTED SOCIAL NETWORKS

Now, without a trusted ringleader, colluders form a distributed and peer-structured social network. They have to help each other identify selfish colluders and implement autonomous selfish colluder identification. Assume that $\mathbf{X}_j^{(k)}$ and $\mathbf{X}_j^{(l)}$ are the fingerprinted copies from $\mathbf{u}^{(k)}$ and $\mathbf{u}^{(l)}$, respectively. Without a trusted ringleader, they have to find another colluder $\mathbf{u}^{(i)}$ to help them calculate $D_j(k,l)$. In order to prevent $\mathbf{u}^{(i)}$ from accessing the fingerprinted coefficients in their copies, $\mathbf{u}^{(k)}$ and $\mathbf{u}^{(l)}$ process their copies beforehand and let $\mathbf{u}^{(i)}$ calculate $D(k,l)$ from the processed copies instead of the original ones. This processing should hide information about the fingerprinted coefficients in $\mathbf{X}_j^{(k)}$ and $\mathbf{X}_j^{(l)}$. Meanwhile, it should not change the MSE between these two copies so that $\mathbf{u}^{(i)}$ can calculate the correct $D_j(k,l)$.

Define $f(\cdot)$ as the function that $\mathbf{u}^{(k)}$ and $\mathbf{u}^{(l)}$ use to process $\mathbf{X}_j^{(k)}$ and $\mathbf{X}_j^{(l)}$, and let $\mathbf{Y}^{(k)} = f(\mathbf{X}_j^{(k)})$ and $\mathbf{Y}^{(l)} = f(\mathbf{X}_j^{(l)})$ be the processed copies, respectively. A component-wise, addition-based method [41] can be used to process the fingerprinted copies: $\mathbf{u}^{(k)}$ and $\mathbf{u}^{(l)}$ first generate a noise-like signal $\mathbf{v}_j^{(k,l)}$ and then add $\mathbf{v}_j^{(k,l)}$ to their copies component by component. That is,

$$\mathbf{Y}_j^{(k)} = f\left(\mathbf{X}_j^{(k)}, \mathbf{v}_j^{(k,l)}\right) = \mathbf{X}_j^{(k)} + \mathbf{v}_j^{(k,l)} \quad \text{and}$$
$$\mathbf{Y}_j^{(l)} = f\left(\mathbf{X}_j^{(l)}, \mathbf{v}_j^{(k,l)}\right) = \mathbf{X}_j^{(l)} + \mathbf{v}_j^{(k,l)}, \tag{35}$$

and therefore, $\|\mathbf{Y}_j^{(k)} - \mathbf{Y}_j^{(l)}\|^2 = \|\mathbf{X}_j^{(k)} - \mathbf{X}_j^{(l)}\|^2$. $\mathbf{v}_j^{(k,l)}$ can help protect the secrecy of the fingerprinted coefficients in $\mathbf{X}_j^{(k)}$ and $\mathbf{X}_j^{(l)}$ if it has large amplitude.

Based on the above, as shown in the example in Figure 11, the key steps in the autonomous selfish colluder identification scheme [41] are summarized as follows:

- **Step 1: Grouping:** Colluders randomly divide themselves into two nonoverlapping subgroups SC_1 and SC_2. In the example in Figure 11, colluders 1–5 are in SC_1 and colluders 6–10 are in SC_2. Then, colluders in SC_1 randomly select an assistant (colluder five in Figure 11) to help colluders in SC_2 calculate $\{D_j(k,l)\}_{k,l\in SC_2}$. Similarly, colluder eight is randomly selected to help colluders in SC_1 calculate $\{D_j(k,l)\}_{k,l\in SC_1}$.
- **Step 2: Encryption:** Colluders in SC_1 generate a noise-like signal $\mathbf{v}_j^{(SC_1)}$. Each colluder $\mathbf{u}^{(i)}$ in SC_1 generates a secret key $K^{(i,8)}$ shared with colluder eight. Then, $\mathbf{u}^{(i)}$ uses (35) to process his or her fingerprinted copy $\mathbf{X}_j^{(i)}$ and generates $f(\mathbf{X}_j^{(i)}, \mathbf{v}_j^{(SC_1)})$. Then, $\mathbf{u}^{(i)}$ encrypts his or her copy with key $K^{(i,8)}$ and transmits the encrypted version $\mathrm{Enc}(f(\mathbf{X}_j^{(i)}, \mathbf{v}_j^{(SC_1)}), K^{(i,8)})$ to colluder eight. Colluders in SC_2 follow the same procedure, process and encrypt their fingerprinted copies, and transmit them to colluder five.
- **Step 3: Calculation of $\{D\}$:** After decrypting the bit streams received from all colluders in SC_1, for each pair of colluders $(\mathbf{u}^{(k)}, \mathbf{u}^{(l)})$ in subgroup SC_1, colluder eight calculates $D_j(k,l) = \|f(\mathbf{X}_j^{(k)}, \mathbf{v}_j^{SC_1}) - f(\mathbf{X}_j^{(l)}, \mathbf{v}_j^{(SC_1)})\|^2$. Colluder eight then broadcasts $\{D_j(k,l)\}_{k,l\in SC_1}$ to colluders in SC_1, together with his or her digital signature. Colluder five repeats the same process to help colluders in SC_2 calculate $\{D_j(k,l)\}$ for all $k, l \in SC_2$.
- **Step 4: Selfish colluder identification:** Given $\{D_j(k,l)\}_{k,l\in SC_1}$, colluders in SC_1 apply the same method as in the section "Centralized Social Networks with Trusted Ringleaders" to detect and identify selfish colluders in SC_1. Similarly, attackers in SC_2 examine $\{D_j(k,l)\}_{k,l\in SC_2}$ and identify selfish colluders in SC_2.

Finally, for colluders who do not apply precollusion processing, they combine the detection results from all frames in the sequence and exclude those identified selfish colluders from collusion.

The above autonomous selfish colluder identification scheme can accurately identify all selfish colluders without falsely accusing any others if colluders five and eight in Figure 11 give others correct values of $\{D(k,l)\}$. However, it is possible that a small group of selfish colluders actively attack the scheme by collaborating with each other and manipulating the detection results. For example, in Figure 11, if both colluders one and eight are selfish colluders, then colluder eight can change the values of $\{D_j(1,k)\}_{k=2,3,4,5}$ such that they follow the same distribution as others. In this case, the above selfish colluder identification algorithm cannot identify colluder one as a selfish colluder, and it makes a miss-detection error. Colluder eight can also change $\{D\}$ and let $\{D_j(2,k)\}_{k=1,3,4,5}$ take much larger values than others. Thus, in addition to missing the true selfish colluder one, the above scheme will also falsely accuse colluder two as selfish and make a false-alarm error.

To resist the above attack, colluders 1–5 select not only colluder eight but also colluders seven and nine to help calculate $\{D_j(k,l)\}$, and use majority vote to make the final decision on the identities of selfish colluders. In this scenario, colluders seven and nine will help correct the detection errors due to colluder eight's manipulation of $\{D_j(k,l)\}$, and the proposed selfish colluder identification scheme can still accurately identify all selfish colluders without falsely accusing others [41]. The work in [41] showed that, if less than 15% of the colluders are selfish, the autonomous selfish colluder identification algorithm can correctly identify all selfish colluders without falsely accusing any others.

CHEAT PREVENTION AND ATTACK RESISTANCE ARE FUNDAMENTAL REQUIREMENTS IN ORDER TO ACHIEVE USER COOPERATION AND PROVIDE RELIABLE SERVICES.

LEVERAGING SIDE INFORMATION IN SOCIAL GAMES

In multimedia social networks, to maximize his or her own payoff, each user observes how others play the game and adjusts his or her own strategy accordingly. Thus, side information plays an important role in multimedia social networks, and it is important to understand which side information about others can help a user improve his or her own payoff, study how users probe and utilize such side information, and analyze its impact on the overall system performance. Again, we use traitor tracing in scalable fingerprinting as an example and investigate how side information affects the colluder-detector dynamics and the traitor-tracing performance of multimedia fingerprinting systems.

In the cat-and-mouse game between the colluders and the digital rights enforcer, there are many collusion strategies from which the colluders can select, and the fingerprint detector has numerous choices when detecting fingerprints. To minimize their risk of being detected, based on the available information about the detection procedure, the attackers try by all means to remove the embedded fingerprints under the fairness constraints. Meanwhile, given a colluded copy, the fingerprint detector selects the detection strategy adaptively to maximize the traitor-tracing capability. In this section, we investigate how each player in the game adjusts his or her own strategy based on available information about others' actions to maximize his or her

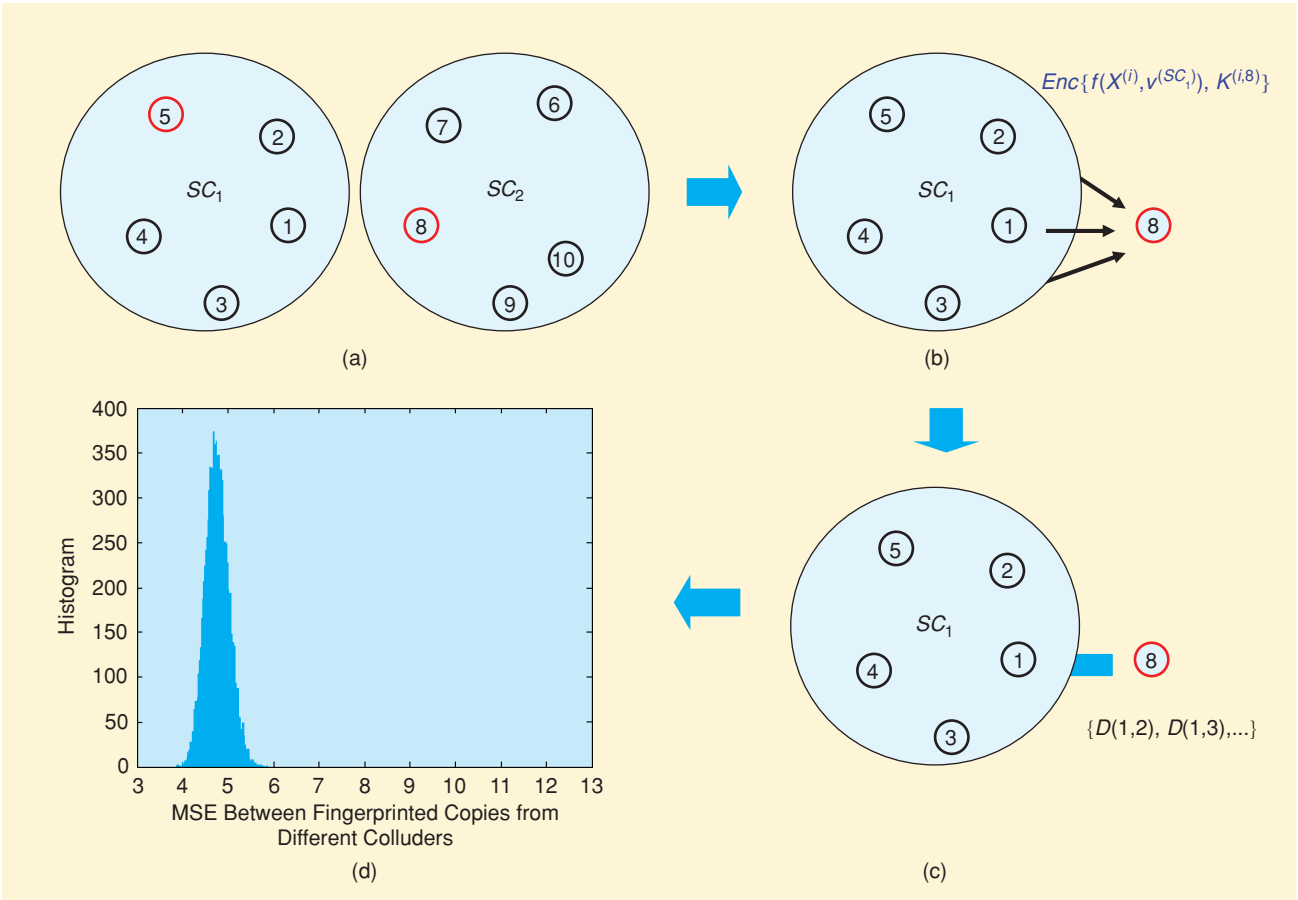

[FIG11] An example of autonomous selfish colluder identification.

own payoff. Without loss of generality, we use equal-risk absolute fairness as an example, and the analysis for other collusion strategies is similar. We further assume that the selfish detection and identification algorithm has accurately identified all selfish colluders, and all attackers that participate in collusion contribute their originally received fingerprinted copies during collusion.

PROBING AND EXPLOITING SIDE INFORMATION

When detecting fingerprints, most prior work simply extended the watermark detection method in digital watermarking and did not consider the unique issues in multiuser collusion. Intuitively, exploring the special characteristics of the colluded copy can help improve the detection performance. Thus, to maximize the success rate of traitor tracing, the fingerprint detector should first examine the colluded copy, probe information about collusion, and utilize this side information to help select the optimum detection strategy.

In a scalable multimedia fingerprinting system, there are various methods to determine if a user participates in collusion. For example, for user $\mathbf{u}^{(i)}$ who receives all three layers, the fingerprint detector can use the collective detection statistics in (11) to measure the similarity between $\mathbf{Y}$ and $\mathbf{W}^{(i)}$. Let SC be the set including the indices of all colluders. Following the same analysis as in (12), with orthogonal fingerprint modulation, if the additive noise $\mathbf{n}$ is i.i.d. Gaussian with zero mean and variance σ_n^2, then the collective detection statistic $TN_c^{(i)}$ in (11) follows the Gaussian distribution [42]

$$TN_c^{(i)} \sim \begin{cases} \mathcal{N}\left(\mu_c^{(i)}, \sigma_n^2\right), & \text{if } i \in SC, \\ \mathcal{N}\left(0, \sigma_n^2\right), & \text{if } i \notin SC, \end{cases}$$
$$\text{where} \quad \mu_c^{(i)} = \frac{(1-\beta_1-\beta_2)N_b + (1-\alpha_1)N_{e1} + N_{e2}}{K^{\text{all}}\sqrt{N_b + N_{e1} + N_{e2}}}\sigma_W. \tag{36}$$

For user $\mathbf{u}^{(i)}$, define $P_s^{(i)}$ as the probability of successfully capturing $\mathbf{u}^{(i)}$ if he or she is guilty, and $P_{fa}^{(i)}$ is the probability of falsely accusing $\mathbf{u}^{(i)}$ if he or she is innocent. With the collective detector in (11), we have

$$P_s^{(i)} = Q\left(\frac{h-\mu_c^{(i)}}{\sigma_n}\right) \quad \text{if} \quad i \in SC, \quad \text{and}$$
$$P_{fa}^{(i)} = Q\left(\frac{h}{\sigma_n}\right) \quad \text{if} \quad i \notin SC. \tag{37}$$

The fingerprint detector can also use the fingerprint extracted from the enhancement layer two, $\mathbf{Y}_{e2}$, to determine if $\mathbf{u}^{(i)}$ is a colluder. In this case, the detection statistic used by the fingerprint detector is

$$TN_{e2}^{(i)} = \left\langle \mathbf{Y}_{e2}, \mathbf{W}_{e2}^{(i)} \right\rangle \Big/ \left\| \mathbf{W}_{e2}^{(i)} \right\|. \tag{38}$$

Following the same analysis as that for the collective detector, $TN_{e2}^{(i)}$ follows the Gaussian distribution

$$TN_{e2}^{(i)} \sim \begin{cases} \mathcal{N}\left(\frac{\sqrt{N_{e2}}}{K^{all}}\sigma_W, \sigma_n^2\right) & \text{if} \quad i \in SC, \\ \mathcal{N}\left(0, \sigma_n^2\right) & \text{if} \quad i \notin SC. \end{cases} \tag{39}$$

Similarly, the fingerprint detector can also use

$$TN_{e1}^{(i)} = \left\langle \mathbf{Y}_{e1}, \mathbf{W}_{e1}^{(i)} \right\rangle \Big/ \left\| \mathbf{W}_{e1}^{(i)} \right\| \quad \text{or}$$
$$TN_b^{(i)} = \left\langle \mathbf{Y}_b, \mathbf{W}_b^{(i)} \right\rangle \Big/ \left\| \mathbf{W}_b^{(i)} \right\| \tag{40}$$

to determine if $\mathbf{u}^{(i)}$ is involved in the attack. The work in [42] showed that

$$TN_{e1}^{(i)} \sim \begin{cases} \mathcal{N}\left((1-\alpha_1)\frac{\sqrt{N_{e1}}}{K^{\text{all}}}\sigma_W, \sigma_n^2\right) & \text{if} \quad i \in SC, \\ \mathcal{N}\left(0, \sigma_n^2\right) & \text{if} \quad i \notin SC, \end{cases}$$
$$\text{and} \quad TN_{e1}^{(i)} \sim \begin{cases} \mathcal{N}\left((1-\beta_1-\beta_2)\frac{\sqrt{N_b}}{K^{\text{all}}}\sigma_W, \sigma_n^2\right) & \text{if} \quad i \in SC, \\ \mathcal{N}\left(0, \sigma_n^2\right) & \text{if} \quad i \notin SC. \end{cases} \tag{41}$$

With $TN_{e2}^{(i)}$, $TN_{e1}^{(i)}$, and $TN_b^{(i)}$, the analysis of $P_s^{(i)}$ and $P_{fa}^{(i)}$ is similar to (37) and thus omitted.

As we can see from the above analysis, the four detection statistics, $TN_c^{(i)}$, $TN_{e2}^{(i)}$, $TN_{e1}^{(i)}$, and $TN_b^{(i)}$, have the same variance but different statistical means, and the one with the largest mean gives the best traitor-tracing performance. Depending on how attackers select the collusion parameters, the detection strategy that gives the best performance may vary from copy to copy, and there is no single detector that outperforms the others in all scenarios.

Figure 12 shows an example of the performance of different detection strategies when detecting colluder $\mathbf{u}^{(i)}$ who receives a high-resolution copy. In this example, when more than 60% of the colluders have high-resolution fingerprinted copies, the collective detector in (11) gives the best performance. This is because in this scenario, $\mathbf{u}^{(i)}$'s fingerprint is spread all over the entire colluded copy, and $\mathbf{W}^{(i)}$'s energy is evenly distributed in the three layers. Therefore, from detection theory [43], fingerprints extracted from all layers should be used collectively during detection to improve the performance. When less than 60% of the colluders receive all three layers, due to the selection of the collusion parameters, a significant portion of $\mathbf{W}^{(i)}$'s energy is in the enhancement layer two, while the other two layers of the colluded copy contain little information of $\mathbf{u}^{(i)}$'s identity. In this scenario, $TN_{e2}^{(i)}$ in (38) gives the best detection performance.

The four detection strategies discussed above use fixed detection statistics to estimate the identities of colluders, and none of them take into consideration how attackers collude and select the collusion parameters. To achieve the optimal performance, the fingerprint detector should first examine the colluded copy and probe such side information about collusion, then uses the best detection statistic with the largest statistical mean to identify colluders. A self-probing detector was proposed in [42] to explore such side information about collusion. As an example, to identify colluders who receive all three layers from the content owner, the key steps in probing side information and selecting the optimum detection statistic are:

- The fingerprint detector first uses the traditional nonprobing detection methods to identify a few suspicious users whose possibilities of participating in collusion are very high. Let $\widehat{SC}$ be the set including the indices of all suspicious users who receive high-resolution copies and are identified in this stage.
- Given $\widehat{SC}$, the detector calculates the sample means of the four detection statistics

$$\hat{\mu}_c = \sum_{i\in\widehat{SC}} TN^{(i)}/|\widehat{SC}|, \quad \hat{\mu}_{e2} = \sum_{i\in\widehat{SC}} TN_{e2}^{(i)}/|\widehat{SC}|,$$
$$\hat{\mu}_{e1} = \sum_{i\in\widehat{SC}} TN_{e1}^{(i)}/|\widehat{SC}|, \quad \text{and} \quad \hat{\mu}_b = \sum_{i\in\widehat{SC}} TN_b^{(i)}/|\widehat{SC}|, \tag{42}$$

where $|A|$ returns the size of the set A.

- The detector compares the four estimated statistical means, $\hat{\mu}_c$, $\hat{\mu}_{e2}$, $\hat{\mu}_{e1}$, and $\hat{\mu}_b$, and selects the detection statistic with the largest sample mean. For example, the collective detector in (11) is chosen if $\hat{\mu}_c$ has the largest value. Then, the fingerprint detector uses the selected detection statistic to make the final decision on the identities of colluders.

MULTIMEDIA SYSTEM DESIGNERS IMPLEMENT ATTACK-RESISTANT AND CHEAT-PROOFING STRATEGIES TO MINIMIZE THE DAMAGE TO AND TO GUARANTEE SATISFACTORY PERFORMANCE OF THE SYSTEM.

Figure 13 compares the performance of three detectors: the collective detector in (11), which always uses the extracted fingerprints from all layers collectively to identify colluders; the optimum detector that has perfect knowledge about the statistical means of the detection statistics and always selects the optimum detection strategy; and the self-probing detector that probes such side information from the colluded copy himself or herself. As shown in Figure 13, information about the statistical mean of the detection statistics can help significantly improve the collusion resistance, and the self-probing detector has approximately the same performance as the optimum detector.

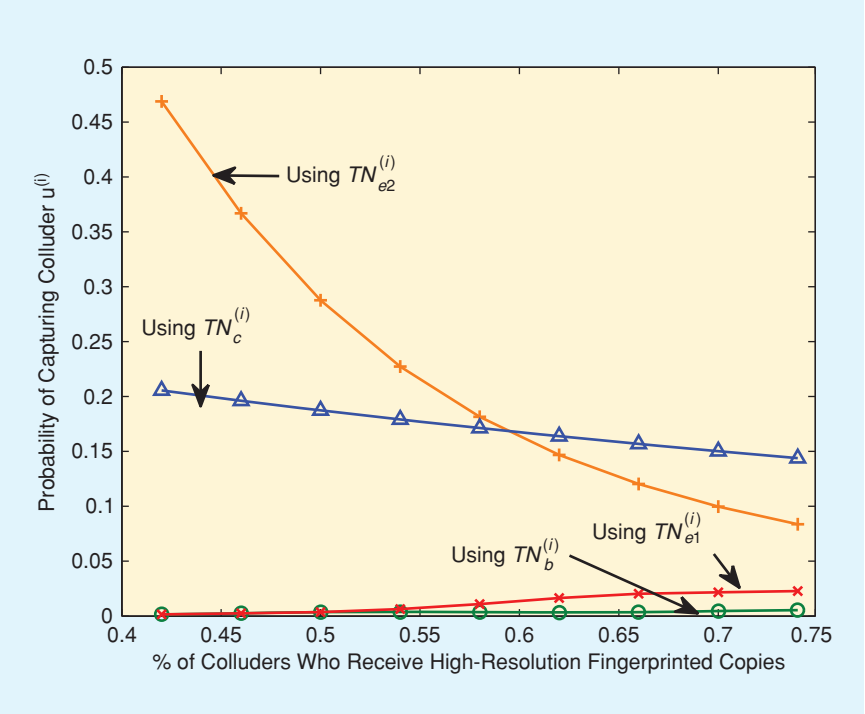

[FIG12] Comparison of different detection statistics (11), (38)–(40). The total number of colluders is fixed as $K = 250$ of which $K^b = 50$ of them receive the low-resolution copies. The horizontal axis (R^{all}) is the percentage of colluders who receive high-resolution fingerprinted copies. Each point on the horizontal axis corresponds to a unique triplet $(K^b, K^{b,e1}, K^{\text{all}})$, where the number of colluders who receive the low-, medium-, and high-resolution fingerprinted copies are $K^b = 50$, $K^{\text{all}} = R^{\text{all}} \times K$ and $K^{b,e1} = K - K^b - K^{\text{all}}$, respectively.

Side information about collusion not only improves the fingerprint detector's traitor-tracing performance, it also affects each colluder's probability of being detected and influences how they collude. Figure 14 shows each colluder's probability of being detected with the self-probing fingerprint detector. From Figure 14, when less than 60% of the colluders receive the high-resolution copies, those colluders who receive all three layers have a much larger probability of being detected than the others. This is because, during collusion, attackers assume that fingerprints extracted from all layers will be used collectively to detect fingerprints, and they select the parameters $\{\alpha_k\}$ and $\{\beta_l\}$ to achieve collective fairness. However, during the colluder identification process, the fingerprint detector probes side information about detection statistics and uses the one that gives the best

collusion resistance. This mismatch causes the difference in different colluders' risk.

IN A SCALABLE MULTIMEDIA FINGERPRINTING SYSTEM, THERE ARE VARIOUS METHODS TO DETERMINE IF A USER PARTICIPATES IN COLLUSION.

GAME-THEORETIC FORMULATION OF ATTACKER-DETECTOR DYNAMICS

Without probing side information, the detector will always use all the frames collectively to identify the colluders, hoping that more frames will give more information about colluders' identities. On the other side of the game, colluders adjust the collusion parameters $\{\alpha_k\}$ and $\{\beta_l\}$ to seek collective fairness. Under such circumstances, the colluders and the fingerprint detector reaches the collective fairness equilibrium. However, side information breaks this equilibrium between the colluders and the fingerprint detector. Both sides need to search for a new equilibrium point, which requires a new framework to model and analyze the complex colluder-detector dynamics. To further analyze the interplay between the colluders and the fingerprint detector, game theory provides fundamental tools to formulate this complex dynamics and facilitate the search of the new equilibrium.

The colluder-detector dynamics can be formulated as a game with two players: the colluders acting as one single player and the fingerprint detector as the other. In this framework, a natural definition of the payoff function is the colluders' probability of being detected, or equivalently, the fingerprint detector's chance of successfully capturing colluders. The two players in this game have conflicting objectives and one player's gain is another's loss.

In such a game, the colluders act first followed by the fingerprint detector. Note that from [42], the self-probing fingerprint detector has approximately the same performance as the optimal detector. Therefore, it is a game with perfect information, where the fingerprint detector is perfectly informed of the colluders' decisions [8]. Consequently, the colluder-detector game can be modeled as a Stackelberg game, where the colluders represents the leader and the fingerprint detector is the follower [8]. The subgame-perfect equilibrium of this game can be found by working backward: first solve for the optimal choice of the fingerprint detector for each possible situation that he or she might face, and then work backward to compute the optimal choice for the colluders. It can be easily shown that this solution is a Nash equilibrium, and each player's actions are optimal at every possible history [44].

For colluder $\mathbf{u}^{(i)}$, define $\mathfrak{D}^{(i)}$ as the set including all the possible detection statistics that can be used to measure the similarity between the extracted fingerprint and $\mathbf{u}^{(i)}$'s fingerprint. For example, when $\mathbf{u}^{(i)}$ receives all three layers of the fingerprinted copy from the content owner, $\mathfrak{D}^{(i)} = \{TN_c^{(i)}, TN_b^{(i)}, TN_{e1}^{(i)}, \text{and } TN_{e2}^{(i)}\}$, where $TN_c^{(i)}$, $TN_b^{(i)}$, $TN_{e1}^{(i)}$, and $TN_{e2}^{(i)}$ are defined in (11), (38), and (40), respectively. Let $P_s^{(i)}$

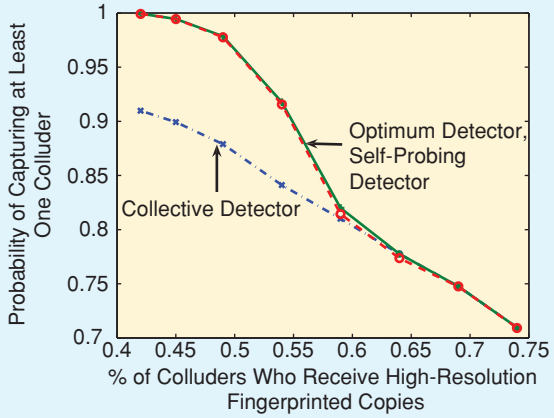

[FIG13] Performance comparison of the collective detector, the optimum detector and the self-probing fingerprint detector. The simulation setup is the same as in Figure 12.

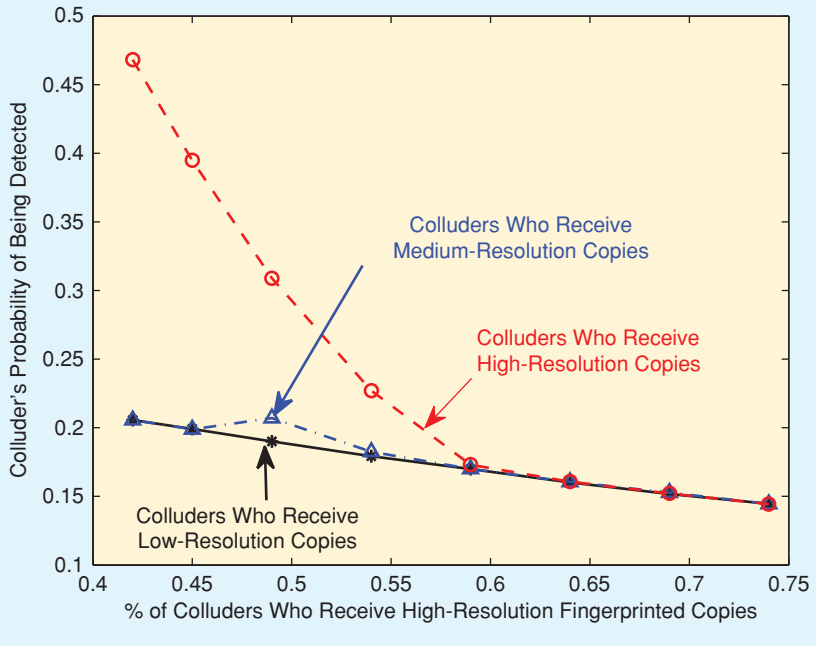

[FIG14] Each colluder's probability of being detected with the self-probing fingerprint detector. The simulation setup is the same as in Figure 12.

be $\mathbf{u}^{(i)}$'s probability of being detected. Mathematically, with equal-risk absolute fairness, to find the subgame-perfect equilibrium of the colluder-detector game, it is equivalent to solve the following min-max problem:

$$\begin{aligned} &\min_{\{\alpha_k,\beta_l\}} \max_{\mathfrak{D}^{(i)}} P_s^{(i)} \\ &\text{s.t.} \quad \max_{\mathfrak{D}^{(i_1)}} P_s^{(i_1)} = \max_{\mathfrak{D}^{(i_2)}} P_s^{(i_2)}, \quad \forall\, i_1, i_2 \in SC, \end{aligned} \tag{43}$$

where SC is the set containing the indices of all colluders. In (43), for every possible set of collusion parameters $\{\alpha_k, \beta_l\}$, $\max_{\mathfrak{D}^{(i)}} P_s^{(i)}$ gives the fingerprint detector's decision when selecting the optimal detection strategy to maximize the traitor-tracing performance; and the minimization operator reflects the colluders' choice of the collusion parameters to minimize their chance of being caught. The constraint $\max_{\mathfrak{D}^{(i_1)}} P_s^{(i_1)} = \max_{\mathfrak{D}^{(i_2)}} P_s^{(i_2)}$ in (43) is the fairness constraint on collusion to ensure the even distribution of risk among colluders.

To find solutions to this min-max problem, for every possible situation that the fingerprint detector might face (that is, for every possible set of collusion parameters $\{\alpha_k, \beta_l\}$), the first step is to analyze $\max_{\mathfrak{D}^{(i)}} P_s^{(i)}$ for every attacker $\mathbf{u}^{(i)}$ and investigate which detection statistic in $\mathfrak{D}^{(i)}$ has the largest statistical mean. This is the same as that of the optimal fingerprint detection in the section "Probing and Exploiting Side Information." The next step is to search for the feasible set, which includes all the possible collusion parameters $\{\alpha_k, \beta_l\}$ that satisfy the fairness constraint $\max_{\mathfrak{D}^{(i_1)}} P_s^{(i_1)} = \max_{\mathfrak{D}^{(i_2)}} P_s^{(i_2)}$ for any two colluders $\mathbf{u}^{(i_1)}$ and $\mathbf{u}^{(i_2)}$. This feasible-set analysis will provide colluders with the constraints on collusion and the selection of collusion parameters to ensure the fair play of the attack. Finally, to minimize their risk, colluders select from the feasible set the collusion parameters that give them the smallest probability of being detected. This min-max solution is a Nash equilibrium of the colluder-detector game [45]: by following this solution, the digital rights enforcer achieves the optimal traitor-tracing performance, and the colluders minimize their risk under the equal-risk absolute fairness constraint.

CONCLUSIONS

In summary, we have discussed recent advances in the study of human dynamics for multimedia social networks, reviewed a few methodologies to investigate the impact of human factors on multimedia security from signal processing perspective, and presented a framework to model and analyze user behavior. Human dynamics plays a vital role in multimedia social networks and must be taken into consideration during the design of multimedia systems. It is important to understand under what conditions users would like to cooperate with each other and how selfish users behave to maximize their own payoff. Possible malicious behavior should also be incorporated in the model to account for malicious users whose goal is to damage and sabotage the system. Equipped with the understanding of human dynamics in social networks, multimedia system designers implement attack-resistant and cheat-proofing strategies to minimize the damage to and to guarantee satisfactory performance of the system.

We hope that the general framework presented in this article will encourage and stimulate researchers from different areas to further explore behavior modeling and forensics for multimedia social networks and beyond. It is an emerging research field with much uncharted territory remains unexplored. We envision that insights from a wide range of disciplines, such as signal processing, game theory, sociology, networking, communications, and economics will help improve our understanding of human dynamics and its impact on multimedia social networks, and ultimately lead to systems with more secure, efficient, and personalized services.

HUMAN DYNAMICS PLAYS A VITAL ROLE IN MULTIMEDIA SOCIAL NETWORKS AND MUST BE TAKEN INTO CONSIDERATION DURING THE DESIGN OF MULTIMEDIA SYSTEMS.

AUTHORS

H. Vicky Zhao (vzhao@ece.ualberta.ca) received the B.S. and M.S. degrees from Tsinghua University, China, in 1997 and 1999, respectively, and the Ph.D. degree from the University of Maryland, College Park, in 2004, all in electrical engineering. She was a research associate with the Department of Electrical and Computer Engineering and the Institute for Systems Research, University of Maryland, College Park from 2005 to 2006. Since 2006, she has been an assistant professor at the University of Alberta, Edmonton, Canada. Since 2007, she has been an adjunct scientist with Telecommunications Research Laboratories. She coauthored the book *Multimedia Fingerprinting Forensics for Traitor Tracing* (Hindawi, 2005). Her research interests include information security and forensics, multimedia, digital communications, and signal processing. She is a Member of the IEEE.

W. Sabrina Lin (wylin@umd.edu) received the B.S. and M.S. degrees in electrical engineering from National Taiwan University in 2002 and 2004, respectively. She is currently pursuing the Ph.D. degree from the University of Maryland, College Park. Her research interests are information security and forensics, multimedia signal processing, and bioinformatics. She received the National Taiwan University Presidential Award in 1999 and 2001 and the University of Maryland Future Faculty Fellowship in 2007.

K.J. Ray Liu (kjrliu@umd.edu) is Distinguished Scholar-Teacher at the University of Maryland, College Park. He is the recipient of numerous honors and awards including best paper awards from the IEEE Signal Processing Society (twice), IEEE Vehicular Technology Society, and EURASIP,

IEEE Signal Processing Distinguishing Lecturer, EURASIP Meritorious Service Award, and National Science Foundation Young Investigator Award. He also received the University of Maryland's Invention of the Year Award and the Poole and Kent Company Senior Faculty Teaching Award. His recent books include *Cooperative Communications and Networking* (Cambridge University Press, 2008); *Resource Allocation for Wireless Networks: Basics, Techniques, and Applications* (Cambridge University Press, 2008); *Ultra-Wideband Communication Systems: The Multiband OFDM Approach* (IEEE-Wiley, 2007); *Network-Aware Security for Group Communications* (Springer-Verlag, 2007); and *Multimedia Fingerprinting Forensics for Traitor Tracing* (Hindawi, 2005). He is a Fellow of the IEEE.

REFERENCES

[1] G.P. Gummadi, R.J. Dunn, S. Saroiu, S.D. Gribble, H.M. Levy, and J. Zahorjan, "Measurement, modeling and analysis of a peer-to-peer file-sharing workload," in *Proc. 19th ACM Symp. Operating Systems Principles (SOSP-19)*, Oct. 2003, pp. 314–329.

[2] J. Liang, R. Kumar, Y. Xi, and K.W. Ross, "Pollution in P2P file sharing systems," *IEEE InfoCom*, vol. 2, pp. 1174–1185, Mar. 2005.

[3] Z. Liu, H. Yu, D. Kundur, and M. Merabti, "On peer-to-peer multimedia content access and distribution," in *Proc. IEEE Int. Conf. Multimedia and Expo*, July 2006, pp. 557–560.

[4] E. Lua, J. Crowcroft, M. Pias, R. Sharma, and S. Lim, "A survey and comparison of Peer-to-Peer overlay network schemes," *IEEE Commun. Surveys Tutorial*, vol. 7, no. 2, pp. 72–93, Mar. 2004.

[5] S. Saroiu, G.P. Gummadi, and S. Gribble, "A measurement study of peer-to-peer file sharing systems," in *Proc. Multimedia Computing and Networking (MMCN)*, Jan. 2002.

[6] C. Buragohain, D. Agrawal, and S. Sur, "A game theoretic framework for incentives in P2P systems," in *Proc. 3rd Int. Conf. Peer-to-Peer Computing*, Sept. 2003, pp. 48–56.

[7] G. Owen, *Game Theory,* 3rd ed. New York: Academic, 1995.

[8] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.

[9] N. Naoumov and K. Ross, "Exploiting P2P systems for DDoS attacks," in *Proc. 1st Int. Conf. Scalable Information Systems*, 2006.

[10] Y. Wang, J. Ostermann, and Y. Zhang, *Video Processing and Communications,* 1st ed. Englewood Cliffs, NJ: Prentice-Hall, 2001.

[11] M. Wu, W. Trappe, Z.J. Wang, and K.J.R. Liu, "Collusion-resistant fingerprinting for multimedia," *IEEE Signal Processing Mag.*, vol. 21, no. 2, pp. 15–27, Mar. 2004.

[12] K.J.R. Liu, W. Trappe, Z.J. Wang, M. Wu, and H. Zhao, *Multimedia Fingerprinting Forensics for Traitor Tracing*, (EURASIP Book Series on Signal Processing and Communications). New York: Hindawi, 2005.

[13] "Movie 'screener' suspect charged, Chicago man violated Hollywood studios copyright," *MSNBC News*, [Online]. Available: http://www.msnbc.msn.com/id/4037016

[14] F. Hartung and M. Kutter, "Multimedia watermarking techniques," *Proc. IEEE*, vol. 87, no. 7, pp. 1079–1107, July 1999.

[15] I. Cox, M. Miller, and J. Bloom, *Digital Watermarking*. San Mateo, CA: Morgan Kaufmann, 2001.

[16] M. Wu and B. Liu, *Multimedia Data Hiding*. New York: Springer-Verlag, Oct. 2002.

[17] C. Podilchuk and W. Zeng, "Image adaptive watermarking using visual models," *IEEE J. Select Areas Commun.*, vol. 16, no. 4, pp. 525–540, May 1998.

[18] I. Cox, J. Killian, F. Leighton, and T. Shamoon, "Secure spread spectrum watermarking for multimedia," *IEEE Trans. Image Processing*, vol. 6, no. 12, pp. 1673–1687, Dec. 1997.

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

[20] J. Killian, T. Leighton, L.R. Matheson, T.G. Shamoon, R. Tajan, and F. Zane, "Resistance of digital watermarks to collusive attacks," Dept. Comput. Sci., Princeton Univ., Princeton, NJ, Tech. Rep. TR-585-98, 1998.

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

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

[23] H. Zhao, M. Wu, Z.J. Wang, and K.J.R. Liu, "Forensic analysis of nonlinear collusion attacks for multimedia fingerprinting," *IEEE Trans. Image Processing*, vol. 14, no. 5, pp. 646–661, May 2005.

[24] D. Kirovski and M.K. Mihcak, "Bounded gaussian fingerprints and the gradient collusion attack," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, vol. 2, pp. 1037–1040, Mar. 2005.

[25] 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. Image Processing*, vol. 14, no. 6, pp. 804–821, June 2005.

[26] F. Zane, "Efficient watermark detection and collusion security," in *Proc. of Financial Cryptography,* (Lecture of Notes in Computer Science, vol. 1962). Feb. 2000, pp. 21–32.

[27] S.B. Wicker, *Error Control Systems for Digital Communication and Storage*. Englewood Cliffs, NJ: Prentice-Hall, 1995.

[28] S. He and M. Wu, "Joint coding and embedding techniques for multimedia fingerprinting," *IEEE Trans. Inform. Forensics and Security*, vol. 1, no. 2, pp. 231–247, June 2006.

[29] J. Dittmann, P. Schmitt, E. Saar, J. Schwenk, and J. Ueberberg, "Combining digital watermarks and collusion secure fingerprints for digital images," *SPIE J. Electron. Imaging*, vol. 9, no. 4, pp. 456–467, Oct. 2000.

[30] W. Trappe, M. Wu, Z. Wang, and K.J.R. Liu, "Anti-collusion fingerprinting for multimedia," *IEEE Trans. Signal Processing*, vol. 51, no. 4, pp. 1069–1087, Apr. 2003.

[31] Z.J. Wang, M. Wu, W. Trappe, and K.J.R. Liu, "Group-oriented fingerprinting for multimedia forensics," *EURASIP J. Appl. Signal Processing, Special Issue on Multimedia Security and Rights Management*, vol. 2004, no. 14, pp. 2142–2162, Nov. 2004.

[32] A. Varna, S. He, A. Swaminathan, M. Wu, H. Lu, and Z. Lu, "Collusion-resistant fingerprinting for compressed multimedia signals," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP)*, vol. 2, Apr. 2007, pp. II 165–168.

[33] H. Zhao and K.J.R. Liu, "Behavior forensics for scalable multiuser collusion: fairness versus effectiveness," *IEEE Trans. Inform. Forensics and Security*, vol. 1, no. 3, pp. 311–329, Sept. 2006.

[34] W.S. Lin, H.V. Zhao, and K.J.R. Liu, "Multi-user collusion behavior forensics: game-theoretic formulation of fairness dynamics," in *Proc. IEEE Int. Conf. Image Processing*, vol. 6, Sept. 2007, pp. 109–112.

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

[36] D. Kiroski and F.A.P. Petitcolas, "Blind pattern matching attack on watermarking systems," *IEEE Trans. Signal Processing*, vol. 51, no. 4, pp. 1045–1053, 2003.

[37] G. Doerr, J.L. Dugelay, and L. Grange, "Exploiting self-similarities to defeat digital watermarking systems: A case study on still images," in *Proc. 2004 ACM Multimedia and Security Workshop*, 2004.

[38] H.V. Zhao and K.J.R. Liu, "Traitor-within-traitor behavior forensics: Strategy and risk minimization," *IEEE Trans. Inform. Forensics and Security*, vol. 1, no. 4, pp. 440–456, Dec. 2006.

[39] S. Baker, R. Gross, I. Matthews, and T. Ishikawa, "Lucas-Kanade 20 years on: A unifying framework," *Int. J. Comput. Vis.*, vol. 56, no. 3, pp. 221–255, Mar. 2004.

[40] H.V. Zhao and K.J.R. Liu, "Selfish colluder detection and identification in traitors within traitors," in *Proc. IEEE Int. Conf. Image Processing*, Oct. 2006, pp. 2289–2292.

[41] H.V. Zhao and K.J.R. Liu, "Autonomous identification of selfish colluders in traitor-within-traitor behavior forensics," *IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, vol. 2, pp. 149–152, Apr. 2007.

[42] W.S. Lin, H.V. Zhao, and K.J.R. Liu, "Scalable multimedia fingerprinting forensics with side information," in *Proc. IEEE Int. Conf. Image Processing*, Oct. 2006, pp. 2293–2296.

[43] H.V. Poor, *An Introduction to Signal Detection and Estimation*, 2nd ed. New York: Springer-Verlag, 1999.

[45] M.J. Osborne and A. Rubinstein, *A Course in Game Theory*. Cambridge, MA: MIT Press, 1994.

[45] W.S. Lin, H.V. Zhao, and K.J.R. Liu, "A game theoretic framework for colluder-detector behavior forensics," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Apr. 2007, vol. 2, pp. 721–724. SP