COMPLEXITY REGULARIZED VIDEO WATERMARKING VIA QUANTIZATION OF PSEUDO-RANDOM SEMI-GLOBAL LINEAR STATISTICS

Oztan Harmanci[1] *and M. Kıvanç Mıhçak*[2]

[1] Department of Electrical and Computer Engineering, University of Rochester, Rochester, NY, 14627.
E-mail: *harmanci@ece.rochester.edu*
[2] Microsoft Research, Cryptography and Anti-Piracy Group, Redmond, WA, 98052.
E-mail: *kivancm@microsoft.com*

ABSTRACT

In this paper, we propose a semi-blind video watermarking scheme for the verification problem, where mark embedding is carried out by designing a complexity-adaptive watermark signal via solving a constrained optimization problem. By adding the resulting watermark to the unmarked host, we effectively quantize pseudo-random linear statistics of the host in the wavelet domain using a secret codebook. We introduce a stochastic video model and exploit it during the design. In particular, we use the proposed model to generate a "complexity map", which is then employed in solving the underlying optimization problem to "regularize" the watermark. Consequently, the resulting watermark is locally adapted to the statistical complexity of the signal at a coefficient level. Regularization is achieved by solving the underlying optimization problem using an iterative algorithm. We experimentally validate the complexity-adaptive structure of the resulting scheme, while maintaining robustness against numerous attacks, such as low bit rate video compression, mild geometric modifications, etc. This paper can be viewed as a continuation of our previous work [1].

1. INTRODUCTION

We consider the watermark (WM) verification problem for video, where the receiver acts as a detector and makes a binary decision regarding the existence of the WM in the received signal. Note that, the verification problem is different from the decoding problem, where the receiver assumes that WM has been embedded in the received signal and the aim is to reliably extract the embedded bits. Most of the prior art on video watermarking can be considered as applications of existing image watermarking schemes, see for example [2, 3]; in such approaches, image watermarking is applied to some appropriately-selected video frames and the emphasis is on the compatibility with existing video compression standards. On the other hand, these methods may lead to serious security problems due to the strong temporal redundancy in video (for instance, in case of estimation attacks). We advocate a different approach, where video is treated as a 3D signal (with heavy temporal correlations) and the watermarking algorithm is designed accordingly.

This paper is a continuation and an improved version of our earlier work [1]. Both are built on the general framework introduced in [4], such that several significant aspects related to video are taken into consideration. We propose to embed the WM after applying 2D-DWT (discrete wavelet transform) to each video frame. In particular, we use pseudo-random (PR) linear statistics of PR (possibly overlapping) connected 3D regions in the DC subband for WM embedding and detection. For mark embedding, we quantize these statistics and subsequently compute an additive WM via solving a PR optimization problem, such that the statistics of the marked signal are equal to the quantized statistics of the (unmarked) host signal. We assume the presence of unmarked host statistics at the receiver, and accordingly use the correlation detector to verify the presence of the WM (if any). Some similarities may be observed between our approach and "spread-transform" (ST) [5, 6] or "quantization-projection" (QP) [7] methods. However, there are major fundamental differences: Our statistics are produced by non-disjoint sets of host coefficients (i.e., overlapping semi-global connected regions) that have geometric meanings (as opposed to the usage of disjoint sets in ST or QP methods that are not necessarily meaningful geometrically). Consequently, we increase the dimension in which WM is computed, which helps us to avoid estimation attacks and to improve robustness in general.

Although this paper and the method of [1] share the same underlying philosophy of embedding the WM in the PR linear statistics domain, they differ in the WM design. In particular, in our current work, we introduce a "signal-adaptive" design, such that the resulting WM has higher (resp. lower) energy in more (resp. less) complex regions of the host [1]. The rationale is due to the empirical realization of the fact that a typical video host is a highly heterogeneous medium and the designed WM should ideally follow the host behavior (both for perceptual and security reasons). In general, we would like to embed stronger (resp. weaker) WM for more (resp. less) detailed regions; if this is not satisfied, WM could create perceptual artifacts and it may even be possible for an intelligent attacker to produce reasonably accurate estimates of the WM (or the unmarked host) by using signal-estimation or source-separation methods. In practice, there can be both high-texture regions and very smooth regions within the same frame and we need to take this into account in the WM design. This task becomes more complicated when we consider temporal correlations; such correlations can last long depending on the video content.

Based on these observations, we realize that there are two challenges: *(i)* We need a practical method to assess the proper WM strength locally in an input-adaptive fashion (adapted to the signal complexity) *(ii)* We need to incorporate the aforementioned assessment in the WM design accordingly, thereby resulting in stronger (resp. weaker) WM in more (resp. less) desirable regions. To achieve these tasks, we extend the system proposed in [1] as follows: We propose a novel stochastic model for video and utilize it to impose complexity constraints on the WM design; as a result, the WM magnitude is adjusted to the signal complexity at a coefficient level. Hence, our approach yields "complexity-regularized" WM design. Note that, although this paper concentrates on the verification problem, a variant of our design can be used for the decoding problem as well.

The rest of the paper is organized as follows: In Sec. 2, we introduce the notation and necessary background. In Sec. 3, we introduce the proposed video model and show how it can be used to compute per-pixel complexity measures in the DC subband of the wavelet domain. In Sec. 4, we explain how the proposed complexity measures are used in the WM design within the framework of [1]. In sec. 5, we discuss the experimental results and conclusions.

This work was carried out while O. Harmanci was with Microsoft Research.

[1] An analogous result of "*stronger WMs for higher entropy channels*" has previously been obtained in [8] for the game-theoretically optimal power allocation problem for information hiding in parallel Gaussian channels.

2. BACKGROUND AND NOTATION

Boldface lowercase letters and corresponding regular letters with subscripts represent vectors and their individiual elements, respectively; for instance, a_i is the i-th element of the vector $\mathbf{a}$. Let $\mathbf{s} \in \mathbb{R}^n$ denote the length-n unmarked host. Note that, here $\mathbf{s}$ is a vector representation of the concatenation of the DC subbands of all the video frames in the 2D-DWT domain. Similarly, let $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ denote the watermarked and attacked signals, respectively. Here, $\mathbf{x}$ should be perceptually approximately the same as $\mathbf{s}$. The goal of the malicious attacker is to cause errors in detection, via feeding the attacked data $\mathbf{y}$ (which should have perceptually acceptable quality) to the receiver; note that, $\mathbf{y}$ may have been produced from some unmarked data $\mathbf{s}$ or its marked version $\mathbf{x}$. If $\mathbf{y}$ has been generated from $\mathbf{s}$ (resp. $\mathbf{x}$), the attacker aims to increase the probability of receiver's declaring the presence of the WM, denoted by P_F (resp. the probability of receiver's declaring the absence of the WM, denoted by P_M)[2]. Naturally, the receiver's goal is to achieve otherwise. In this work, we consider a symmetric private-key setup; i.e., there exists a secret key that is shared by the embedder and the receiver, which is unknown to the malicious attacker. It should be understood that in all the randomized steps of the proposed scheme, a secure PR number generator has been used with the secret key as the seed.

Most watermarking techniques consider either a "blind" or "non-blind" setup, where the former (resp. the latter) assumes that the receiver has (resp. does not have) access to unmarked host $\mathbf{s}$. Here, we consider an alternative "semi-blind" setup, where the receiver does not have access to $\mathbf{s}$, but rather has access to some side information that is correlated with $\mathbf{s}$ and is a function of the secret key. Thus, by construction, our approach is well-suited for fingerprinting applications. This setup was also considered in [1]; similarly, a hash-aided image watermarking scheme was proposed in [9].

Next, we briefly explain the basics of our "watermarking via statistics quantization" approach, which forms the skeleton of the scheme proposed in this paper. For further details, we refer the interested reader to [1, 4].

2.1 Watermark Embedding and Detection

We embed the WM by effectively changing the statistics of (possibly overlapping) 3D connected regions with PR locations. Currently, we confine ourselves to rectangular prisms; however, our approach allows the usage of regions with arbitrary shapes. We first generate sizes and locations of m rectangular prisms; the i-th prism is represented by the set $R_i \subset \{1, 2, \ldots, n\}$, which defines the indices of the coefficients that belong to it, $1 \le i \le m$. The i-th unmarked statistic μ_i of $\mathbf{s}$ is given by the linear weighted combination of $\{s_j\}$ with smooth PR jointly-Gaussian weights, that are bandlimited to f_t radians by construction (here $\{s_j\}$ are in the region R_i). We represent the weights for R_i with $\mathbf{t}^i \in \mathbb{R}^n$, where $t^i_j = 0$ if $j \notin R_i$. Hence, we write $\mu_i = \sum_{j=1}^n s_j t^i_j = <\mathbf{s}, \mathbf{t}^i>$, which leads to $\mu = \mathbf{T}\mathbf{s}$, where $\mu \in \mathbb{R}^m$ is the unmarked statistics vector, and $\mathbf{T} \in \mathbb{R}^{m \times n}$ is formed such that its i-th row is $\mathbf{t}^i$.

We design the additive WM sequence $\mathbf{w} \in \mathbb{R}^n$, such that the watermarked signal $\mathbf{x} = \mathbf{s} + \mathbf{w}$ has statistics $\hat{\mu} \in \mathbb{R}^m$, that are the quantized version of μ. Currently, we use a scalar uniform quantizer with step size δ; however, in general any high-dimensional quantizer that maps μ to $\hat{\mu}$ can be used. We compute $\mathbf{w}$ such that $\mathbf{T}(\mathbf{s} + \mathbf{w}) = \hat{\mu}$ and $||\mathbf{w}|| = ||\mathbf{x} - \mathbf{s}||$ is minimized; i.e., we solve

$$\min_{\mathbf{w}} ||\mathbf{w}|| \quad \text{s.t.} \quad \mathbf{T}\mathbf{x} = \hat{\mu} \Leftrightarrow \mathbf{T}\mathbf{w} = \hat{\mu} - \mu. \tag{1}$$

Assuming that $\mathbf{T}$ is full-rank (which is almost always satisfied in practice with our parameter selection), the solution to (1) is given by the well-known *minimum-norm* result:

$$\mathbf{w}_{MN} = \mathbf{T}^T \left(\mathbf{T}\mathbf{T}^T\right)^{-1} (\hat{\mu} - \mu). \tag{2}$$

[2] P_F and P_M are conventionally used for probabilities of false alarm and miss in detection theory.

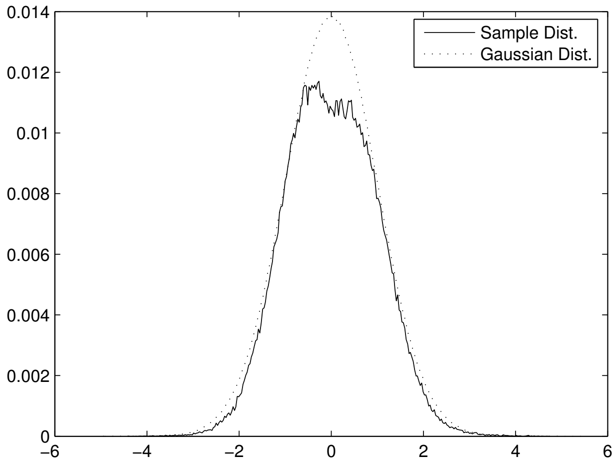

Figure 1: Sample distribution(solid) normalized according to the proposed spatial dependency and theoretical Gaussian distribution(dotted) for 10000 randomly selected samples in DWT DC subband for $M = 5$. Kurtosis of the distribution is 2.9093

The vector μ is sent as side information to the receiver per our semi-blind assumption. Having the input $\mathbf{y}$, the receiver makes a binary decision regarding the existence of the WM using the statistics of $\mathbf{y}$ (represented by $\tilde{\mu} = \mathbf{T}\mathbf{y}$) in a correlation detector. The normalized correlation value is given by $\gamma = \frac{<\tilde{\mu} - \mu, \hat{\mu} - \mu>}{||\hat{\mu} - \mu||^2}$. Having chosen a threshold τ ($0 < \tau < 1$), the detector declares the presence (resp. the absence) of the WM if $\gamma > \tau$ (resp. $\gamma < \tau$).

Note that, the solution (2) yields optimal $\mathbf{w}$ in the sense of l^2 norm in the presence of the constraint of $\mathbf{T}\mathbf{x} = \hat{\mu}$. However, this does not necessarily yield the best solution in the sense of perceptual quality and security. It is possible to solve (1) in the presence of other meaningful constraints. In particular, in Sec. 4, we propose an iterative algorithm, such that $\mathbf{w}$ is designed to achieve $\mathbf{T}\mathbf{x} = \hat{\mu}$, while being sufficiently smooth and locally-adapted to the signal complexity; this forms the crux of our work. In practice, signal complexity is computed at a coefficient level by imposing a statistical model on the video, which is the topic of the next section.

3. PROPOSED STATISTICAL MODEL FOR VIDEO

3.1 Spatial dependencies

In order to capture spatial correlations within each frame, we employ the statistical image model (also known as the EQ model) initially proposed in [10] for compression, subsequently used in [11] and [8] for denoising and information-hiding-capacity computation, respectively. Specifically, we assume a locally approximately i.i.d. (independent identically distributed) Gaussian model for each coefficient; in practice, we find approximate ML (maximum-likelihood) estimates of the mean and variance of each coefficient relying on the i.i.d. approximation within a window of size $M \times M$. In fig. 1, we show the histogram of randomly selected 10000 coefficients (normalized by their estimated means and variances) in the DC subband of a typical video. The figure shows that we have a Gaussian-like operational distribution, thereby empirically justifying our modeling assumptions; the corresponding sample kurtosis is 2.9093 which is close to the ideal value of 3.00 for Gaussian distribution.

Let $C_s(l_x, l_y, l_t)$ denote the *spatial complexity* (i.e., a measure of spatial details) of the DC-subband wavelet coefficient $c(l_x, l_y, l_t)$ in frame l_t with spatial coordinates (l_x, l_y). Motivated by the standard information-theoretic entropy-rate expression for Gaussian sequences[12], we propose to use $C_s(l_x, l_y, l_t) = \frac{1}{2}\log\left[2\pi e \hat{\sigma}_s^2(l_x, l_y, l_t)\right]$, where $\hat{\sigma}_s^2(l_x, l_y, l_t)$ is the approximate ML estimate of the variance of $c(l_x, l_y, l_t)$ using the EQ model within each frame (in practice, we use $M = 5$ for estimation); for fur-

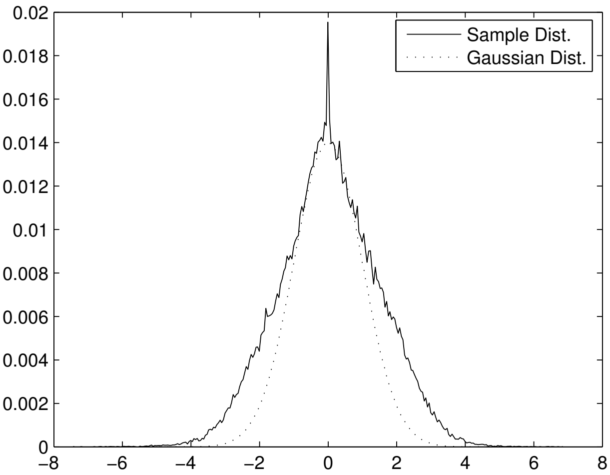

Figure 2: Sample distribution(solid) normalized according to the proposed temporal dependency and theoretical Gaussian distribution(dotted) for 10000 randomly selected samples in DWT DC subband for $N = 13$. Kurtosis of the distribution is 3.1074

ther details on the estimation of the underlying variance field, see [10, 11, 8].

3.2 Temporal dependencies

Suppose, we have a method that gives us accurate per-pixel motion vectors. Generally, optical flow or block matching techniques are used to achieve this task; in practice, we use the method described in [13] to obtain the motion vector field. Let $\mathbf{MV}\left(l_x,l_y,l_t,\Delta\right)$ denote the 2-dimensional motion vector (with horizontal and vertical components $MV_x(\cdot,\cdot,\cdot,\cdot)$ and $MV_y(\cdot,\cdot,\cdot,\cdot)$ respectively), that maps $c\left(l_x,l_y,l_t\right)$ to the corresponding pixel in the $l_t+\Delta$'th frame. Thus, if the motion field is accurate, we should have $c\left(l_x,l_y,l_t\right) \approx c\left(l_x+MV_x\left(l_x,l_y,l_t,\Delta\right),l_y+MV_y\left(l_x,l_y,l_t,\Delta\right),l_t+\Delta\right)$.

Given the DC subband coefficients $\left\{c\left(l_x,l_y,l_t\right)\right\}$ and the corresponding motion vectors $\left\{\mathbf{MV}\left(l_x,l_y,l_t\right)\right\}$, we assume a *locally-stationary first-order Gaussian AR* (auto regressive) process along the direction of the motion field:

$$c\left(l_x,l_y,l_t\right)=\rho\left(l_x,l_y,l_t\right)\cdot c\left(l_x+MV_x,l_y+MV_y,l_t-1\right)+g\left(l_x,l_y,l_t\right),$$

where we dropped the parameters of MV_x and MV_y for convenience. Here, $\left\{g\left(l_x,l_y,l_t\right)\right\}$ are locally i.i.d. and distributed with $\mathcal{N}\left(0,\sigma_t^2\left(l_x,l_y,l_t\right)\right)$. In practice, we apply the locally-stationarity assumption within a temporal window of length-N, where $N=13$. For each coefficient, we first find the least-squares (LS) estimate $\hat{\rho}\left(l_x,l_y,l_t\right)$ of the correlation coefficient $\rho\left(l_x,l_y,l_t\right)$ using the temporal neighbors. Then, we find an approximate ML estimate $\hat{\sigma}_t^2\left(l_x,l_y,l_t\right)$ of $\sigma_t^2\left(l_x,l_y,l_t\right)$, after approximate whitening using $\hat{\rho}\left(l_x,l_y,l_t\right)$. In fig. 2, we show the histogram of randomly selected 10000 $\left\{g\left(l_x,l_y,l_t\right)\right\}$ samples, normalized by their estimated variances. The figure experimentally verifies our Gaussian approximation; the corresponding sample kurtosis is 3.1074 (close to the ideal value of 3.00). Relying on a similar rationale of sec. 3.1, we use $C_t\left(l_x,l_y,l_t\right)=\frac{1}{2}\log\left[2\pi e\hat{\sigma}_t^2\left(l_x,l_y,l_t\right)\right]$ as the temporal complexity measure.

4. COMPLEXITY-REGULARIZED WATERMARK DESIGN

Once the minimum norm solution $\mathbf{w}_{MN}$ is found by (2), we apply the following iterative algorithm to find the final WM:

Set $\mathbf{w}^0=\mathbf{w}_{MN}=\mathbf{T^T}(\mathbf{TT^T})^{-1}d$. For K steps, do

1. $\mathbf{w}^1 = process\left(\mathbf{w}^0\right)$,
2. $\mathbf{w}_d=\mathbf{w}_1-\mathbf{w}_{MN}$
3. $\mathbf{w}_n=\mathbf{w}_d-\mathbf{T^T}(\mathbf{TT^T})^{-1}\mathbf{Tw}_d$
4. $\mathbf{w}_0'=\mathbf{w}_n+\mathbf{w}_{MN}$
5. if $\left\|\mathbf{w}_0'-\mathbf{w}_0\right\|<\varepsilon$ stop, else $\mathbf{w}_0=\mathbf{w}_0'$, go to 1.

Here, the function $process(\cdot)$ indicates pre-processing of the input WM candidate $\mathbf{w}^0$, such that the output is perceptually acceptable and locally-adapted to signal complexity (i.e., complexity-regularized). Ideally, we would like to have $\mathbf{w}^1=\mathbf{w}^0$ at step 1, but this may not be the case in general. To achieve the task of complexity regularization, we first apply element-wise multiplication to the entries of $\mathbf{w}^0$, where the scaling factor for the location $\left(l_x,l_y,l_t\right)$ is the corresponding per-pixel *cumulative complexity measure* $C_c(\cdot,\cdot,\cdot)$, which is shifted weighted average of the spatial and temporal complexity measures: $C_c\left(l_x,l_y,l_t\right)=\alpha C_s\left(l_x,l_y,l_t\right)+\beta C_t\left(l_x,l_y,l_t\right)+\theta$. Here, α and β are chosen experimentally, and θ is chosen to such that the mean of C_c is 1.0;note that this (linear) functional form is a heuristic and is not necessarily optimal in general. Nevertheless, this operation ensures that the WM in the low (resp. high) complexity regions is attenuated (resp. amplified). The next operation that constitutes the function $process(\cdot)$ is ideal low pass filtering, via which smoothness is imposed on the WM. After finding the "correction WM" $\mathbf{w}^d$ on $\mathbf{w}_0$ as a result of step 2, we project it onto the nullspace of $\mathbf{T}$ in step 3 and do the update in step 4. Note that, $\mathbf{w}_{MN}$ lives in the range space of $\mathbf{T}^T$, which is orthogonal to the nullspace of $\mathbf{T}$; thus, we necessarily have $\mathbf{Tw}_n=0$, which leads to $\mathbf{T}\left(\mathbf{w}_n+\mathbf{w}_{MN}\right)=\mathbf{Tw}_{MN}=\hat{\mu}-\mu$, which is crucial in maintaining the quantization condition of $\mathbf{Tx}=\hat{\mu}$.

Remark: We experimentally observed that we can have poor convergence rate (or no convergence at all) if the host is considerably "uncomplex". This is intuitively clear, since in that case the algorithm tries to globally attenuate the WM energy in step 1, which contradicts with the quantization condition. To avoid such problematic situations, we impose a complexity constraint on each candidate prism, whose statistic is to be quantized. In particular, given a candidate prism, if the number of coefficients, whose cumulative complexity measures are larger than a pre-specified threshold, is large enough, we declare that prism to be complex-enough for watermarking and proceed; otherwise we re-select another PR prism. With this adaptive rectangle selection, we observed that the proposed iterative algorithm always converged in practice.

5. EXPERIMENTAL RESULTS AND DISCUSSIONS

We experimentally study the performance of the proposed algorithm and [1]. We used 4 various videos with different spatial and temporal content at 640×480 resolution. The length of each region is selected as 300 frames. Complexity map is created at 2×2 spatial DWT decomposition and further downsampled to match the actual embedding resolution. Detailed parameter set can be found in [1]. The α and β as introduced in 4 are both set to 0.003.

We first compare the WM that is generated by the proposed system and [1]. Figure 3 shows a frame from an artificially edited video. The lower portion of the video is set to a constant value. Lower left pane and and lower right pane show the WM generated by [1] and the proposed system respectively (both magnified by the same amount for visibility). It is easily seen that the proposed system avoids embedding WM in the edited region.

Then we compare the correlation detector performance under the geometric attacks of rotation and cropping. We take a random region from each of the 4 different videos and determine the decoding performance for 100 users. We show the cumulative Receiver Operating Characteristics(ROC) curves in figures 4 and 5 for 5.4 degrees rotation and 16 percent areal cropping respectively. We see that there is no significant performance difference at all.

Discussion: We proposed a novel method for complexity regularized watermark embedding via quantization of pseudo-random semi-global linear statistics. The proposed algorithm uses a stochastic video model to generate a complexity map. We use the complex-

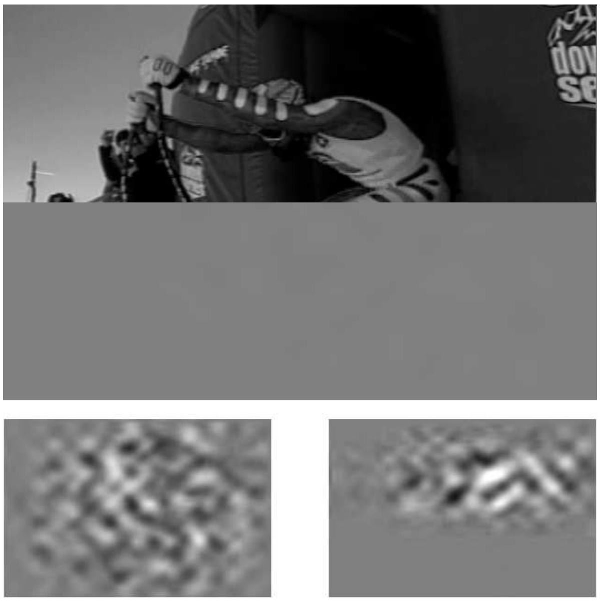

Figure 3: Top: An artificially edited video frame such that the bottom half is constant; Lower left: WM generated by the algorithm [1]($\times$50 for visibility); Lower right: WM generated by the proposed algorithm ($\times$50 for visibility)

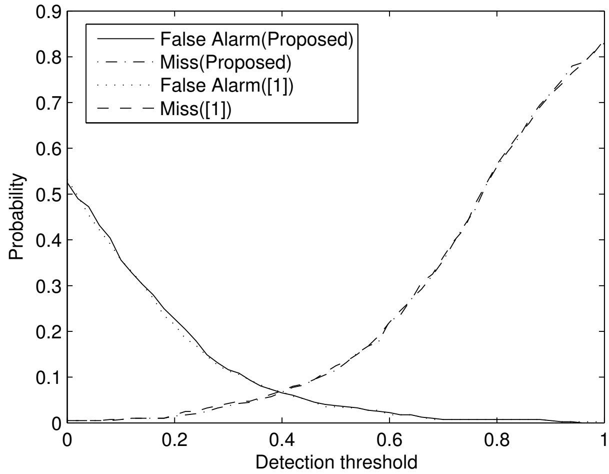

Figure 4: ROC curves comparing the performance of the proposed system and system in [1] under 5.4 degrees central rotational attack performed with 400 users and 4 various videos

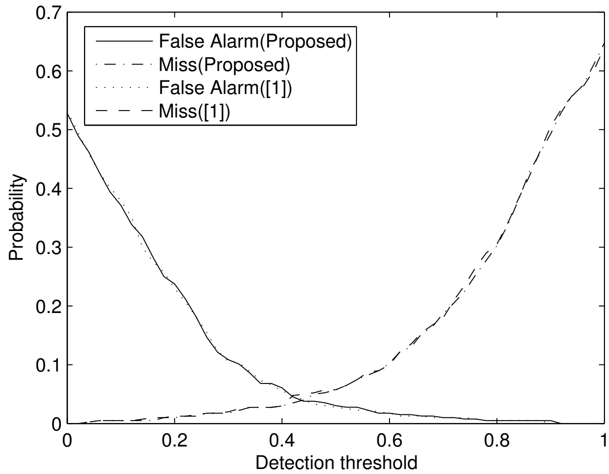

Figure 5: ROC curves comparing the performance of the proposed system and system in [1] under 16 percent areal cropping followed by scaling attack performed with 400 users and 4 various videos

ity map to locally adapt the WM to the signal. Experiments show that the proposed video model is accurate and it is as robust against reasonable geometric attacks as [1] which does not perform complexity regularization.

REFERENCES

[1] M. Kucukgoz, O. Harmanci, M. K. Mıhçak, and R. Venkatesan, "Robust Video Watermarking via Optimization Algorithm for Quantization of Pseudo-Random Semi-Global Statistics," *Proc. SPIE 2005*, San Jose, CA, 2005.

[2] G. Depovere, T. Kalker, J. Haitsma, M. Maes, L. de Strycker, P. Termont, J. Vandewege, A. Langell, C. Alm, P. Norman, G. O'Reilly, B. Howes, H. Vaanholt, R. Hintzen, P. Donnelly, A. Hudson, "The VIVA project: digital watermarking for broadcast monitoring," *Proc. ICIP'99*, vol. 2, pp. 202–205, Oct. 1999.

[3] A. M. Alattar, E. T. Lin, and M. U. Celik, "Digital watermarking of low bit-rate advanced simple profile MPEG-4 compressed video," *IEEE Trans. on Circuits and Systems for Video Technology*, vol. 13, no. 8, pp. 787–800, Aug. 2003.

[4] M. K. Mıhçak, R. Venkatesan and M. Kesal, "Watermarking via Optimization Algorithms for Quantizing Randomized Statistics of Image Regions," Proceedings of the Fortieth Annual Allerton Conference on Communication, Control and Computing, Monticello, IL, Oct. 2002.

[5] J. Eggers, "Information Embedding and Digital Watermarking as Communication with Side Information," *Ph.D. thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg*, 2001.

[6] B. Chen and G. W. Wornell, "Quantization index modulation: a class of provably good methods for digital watermarking and information embedding," *IEEE Trans. Inform. Theory*, vol. 47, no. 4, pp. 1423–1443, May 2001.

[7] F. Pérez-González, F. Balado, and J. R. Hernández, "Performance analysis of existing and new methods for data hiding with known-host information in additive channels," *IEEE Trans. on Signal Processing*, vol. 51, no. 4, pp. 960–980, Apr. 2003.

[8] P. Moulin and M. K. Mıhçak, "A Framework for Evaluating the Data Hiding Capacity of Image Sources," *IEEE Transactions on Image Processing*, vol. 11, no. 9, pp. 1029-1042, Sep. 2002.

[9] J. L. Cannons and P. Moulin, "Design and Statistical Analysis of a Hash-Aided Image Watermarking System," *IEEE Trans. on Image Processing*, Oct. 2004.

[10] S. LoPresto, K. Ramchandran, and M. T. Orchard, "Image coding based on mixture modeling of wavelet coefficients and a fast estimationquantization framework," in *Proc. IEEE Data Compression Conf.*, Snowbird, Utah, pp. 221-230, 1997.

[11] M. K. Mıhçak, I. Kozintsev, K. Ramchandran and P. Moulin, "Low complexity image denoising based on statistical modeling of wavelet coefficients," *IEEE Signal Processing Letters*, vol. 6, no. 12, pp. 300–303, Dec. 1999.

[12] Thomas M. Cover and Joy A. Thomas, "Elements of Information Theory", John Wiley & Sons, Inc., 1991.

[13] S. B. Kang, M. Uyttendaele, S. A. J. Winder and R. Szeliski, "High Dynamic Range Video," *Proceedings of ACM Siggraph 2003.*