A Lightweight Rao-Cauchy Detector for Additive Watermarking in the DWT-Domain

Roland Kwitt, Peter Meerwald and Andreas Uhl
Department of Computer Sciences
University of Salzburg
Jakob-Haringer-Str. 2, A-5020 Salzburg, Austria
{rkwitt,pmeerw,uhl}@cosy.sbg.ac.at

ABSTRACT

This paper presents a lightweight, asymptotically optimal blind detector for additive spread-spectrum watermark detection in the DWT domain. In our approach, the marginal distributions of the DWT detail subband coefficients are modeled by one-parameter Cauchy distributions and we assume no knowledge of the watermark embedding power. We derive a Rao hypothesis test to detect watermarks of unknown amplitude in Cauchy noise and show that the proposed detector is competitive with the Generalized Gaussian detector, yet is more efficient in terms of required computations.

Categories and Subject Descriptors

I.4.10 [**Image Processing and Computer Vision**]: Statistical

General Terms

Algorithms, Performance, Security

Keywords

Watermarking, Wavelet, Statistical Signal Detection

1. INTRODUCTION

Watermarking has been proposed as a technology to ensure copyright protection by embedding an imperceptible, yet detectable signal in digital multimedia content such as images or video. For blind watermarking, i.e. when detection is performed without reference to the unwatermarked host signal, the host interferes with the watermark signal. Hence, informed watermark embedding and modeling the host signal is crucial for detection performance [18, 7].

Transform domains – such as the Discrete Cosine Transformation (DCT) or the Discrete Wavelet Transformation (DWT) domain – facilitate modeling human perception and permit selection of significant signal components for watermark embedding. The perceptual characteristics and distributions of transform domain coefficients has been extensively studied for image compression [3].

Many approaches for optimal detection of additive watermarks embedded in transform coefficients have been proposed in literature so far [10, 21, 6, 19, 4]. For blind watermarking, the host transform coefficients are considered as noise from the viewpoint of signal detection. If we assume Gaussian noise, it is known that the optimal detector is the straightforward linear-correlation (LC) detector [13].

Unfortunately, DCT and DWT coefficients do not obey a Gaussian law in general, which renders the LC detector suboptimal in these situations. A first approach, exploiting the fact that DCT or DWT coefficients do not follow a Gaussian law is proposed in [10]. The authors derive an optimal detector for an additive bipolar watermark sequence in DCT transform coefficients following a Generalized Gaussian Distribution (GGD). In [4], it is shown that the low- to mid-frequency DCT coefficients excluding the DC coefficient can also be modeled by the family of symmetric alpha-stable distributions and a detector is derived for Cauchy distributed DCT coefficients by following the same scheme as it is presented in [10]. However, both approaches are based on the strong assumption that the watermark embedding power is known to the detector. In [21], a new watermark detector based on the Rao hypothesis test [22] is proposed for watermark detection in Generalized Gaussian distributed noise. The detector is asymptotically optimal (e.g. for large data records) and does not depend on knowledge about the embedding power any more.

In this work we derive another form of the Rao detector based on the assumption that DWT detail subband coefficients approximately follow a one-parameter Cauchy distribution. Our approach is motivated by the fact that current detectors which rely on the GGD are computationally expensive and require a cumbersome parameter estimation procedure. The Cauchy model however leads to a simple detector, which is competitive with the state-of-the art detectors in this field. Detection runtime requirements are important to certain applications. While [5] aims to reduce the length of the watermark sequence, we try to reduce the computational effort per step. For our discussion on the proposed detector, we go without any perceptual modeling, although our approach can be easily combined with the framework of [16] for example.

The remainder of the paper is structured as follows: In Section 2 we discuss the statistical model of our approach,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
MM&Sec'08, September 22–23, Oxford, United Kingdom.
Copyright 2008 ACM 978-1-60558-058-6/08/09 ...$5.00.

followed by the derivation of the detector in Section 3. In Section 4, we present experimental detection results and evaluate the performance of our detector under JPEG and JPEG2000 attacks. Moreover, we discuss computational issues related to our detector and current state-of-the-art detectors. Finally, Section 5 concludes the paper with a discussion on open problems and an outlook on further research.

2. STATISTICAL MODEL

In this section we briefly review some results on modeling the marginal distributions of DWT detail subband coefficients by univariate probability distributions. It is commonly accepted that the marginal distributions of the subband coefficients of natural images are highly non-Gaussian but can be well modeled by the GGD (see [17, 23, 3]). The probability density function (PDF) of the GGD is given by

$$p(x|b,c) = \frac{c}{2b\Gamma(1/c)} \exp\left(-\left|\frac{x}{a}\right|^c\right), \tag{1}$$

with $-\infty < x < \infty$ and $b, c > 0$, where we have used the general parametrization of [20]. In contrast to the Gaussian distribution (which arises as a special case of the GGD for $c = 2$), the GGD is a leptokurtic distribution which allows heavy-tails. Although the GGD is generally the best known model for the DWT detail subband coefficients, the computational cost for estimating the two distribution parameters, shape c and scale b, is rather high. For example, maximum-likelihood estimation (MLE) requires to find the root of a highly non-linear equation [9] (see Section 4.2). To avoid this computationally intensive procedure in watermarking applications, the shape parameter is often set to a fixed value (e.g. $c = 0.5$ [10]) and the same parameter setting is used for all subbands [2]. However, it is obvious that this reduces the flexibility of the GGD, which can result in a loss of detection performance. In this work, we propose to use the Cauchy distribution, which is a member of the family of symmetric alpha-stable (SαS) distributions, as an alternative distribution to model the DWT detail subband coefficients. In case of low- to mid-frequency DCT coefficients the Cauchy distribution has already been successfully employed for blind DCT-domain spread-spectrum watermarking [4]. The PDF of the Cauchy distribution with location parameter $-\infty < \delta < \infty$ and shape parameter $\gamma > 0$ is given by [15]

$$p(x|\gamma,\delta) = \frac{1}{\pi} \frac{\gamma}{\gamma^2 + (x-\delta)^2}, \tag{2}$$

with $-\infty < x < \infty$. The Cauchy distribution with $\delta = 0$ (which is symmetric around zero) will be abbreviated by $p(x|\gamma) := p(x|\gamma, 0)$. In contrast to the Gaussian distribution, the tails of the Cauchy distribution decay at a rate slower than exponential, hence we observe heavy-tails in the PDF. Figure 1 shows the PDFs of two Cauchy distributions with different shape parameters and $\delta = 0$. Regarding the estimation of the distribution parameter γ, maximum-likelihood estimation is straightforward. Given that $x[1], \ldots, x[N]$ denote realizations of N i.i.d. random variables following a Cauchy distribution with $\delta = 0$, the estimate of γ (denoted by $\hat{\gamma}$), is given as the solution to (see [15])

$$\frac{1}{N} \sum_{t=1}^{N} \frac{2}{1 + (x[t]/\hat{\gamma})^2} - 1 = 0, \tag{3}$$

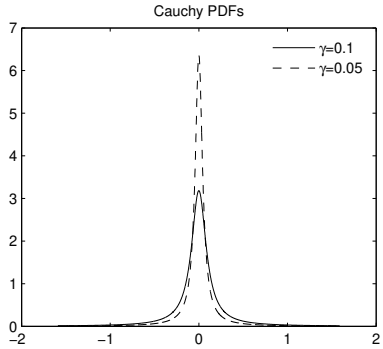

Figure 1: Two exemplary PDFs of the Cauchy distribution with $\delta = 0$.

which has to be solved numerically using some root-finding algorithm (see Section 4.2). The motivation for preferring the Cauchy distribution over the GGD as an underlying model for the DWT coefficients has two particular reasons: First, the simple nature of the Cauchy PDF leads to a simple detection statistic on the one hand (see Section 3) and parameter estimation of γ is less complex than parameter estimation for the GGD on the other hand.

To exemplify that it is reasonable to model the DWT coefficients by a Cauchy distribution, we show Quantile-Quantile plots (Q-Q Plots) for the second-scale horizontal detail subband coefficients in Figure 2 ($\Phi(\cdot)$ denotes the sample quantile function and $F(\cdot)$ denotes the CDF of the Cauchy distribution). In case of a good fit, the data points should follow the straight line, since the empirical and theoretical quantiles would coincide. As we can see from Figure 2, the fit is actually pretty good, especially in the middle regions. However, we notice slight deviations in the tail regions, since the Cauchy distribution is too heavy-tailed.

3. THE DETECTION PROBLEM

In this section, we derive a new, asymptotically optimal detector for additive spread-spectrum watermarking in the DWT domain. We assume that a bipolar watermark sequence is embedded in the transform coefficients and that the watermark embedding power is unknown at the detection stage.

Before we go on, we introduce some notation and define our signal detection problem. For a J-scale pyramidal DWT we obtain three detail subbands per decomposition level $j < J$, denoted by $\mathbf{H}_j$ (horizontal detail subband), $\mathbf{V}_j$ (vertical detail subband) and $\mathbf{D}_j$ (diagonal detail subband). The detail subbands are given in matrix notation. The transform coefficients for the horizontal detail subband $\mathbf{H}_j$ for example will be denoted by $h^j[l,k], 1 \le l,k \le M_j$, where M_j^2 denotes the number of coefficients for a subband on decomposition level j (without loss of generality we assume square images). When it is not necessary to speak of a specific subband, N is the number of subband coefficients and the coefficients are given by $x[1], \ldots, x[N]$ (vector notation). This vector arises by simple concatenation of the row vectors of the appropriate subband coefficient matrix. By adhering to this convention, the elements of the bipolar watermark sequence used for marking an arbitrary subband are denoted by $w[t], 1 \le t \le N$ with $w[t] \in \{+1, -1\}$. For

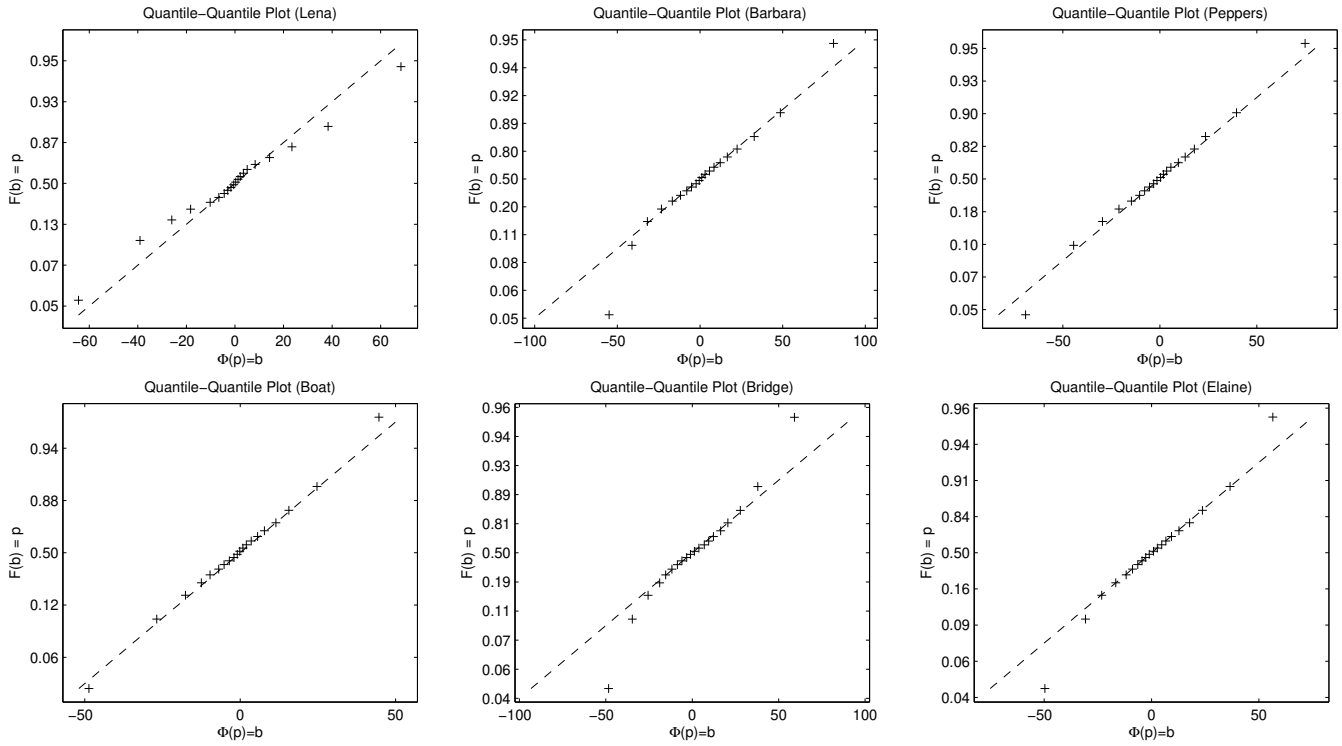

Figure 2: Cauchy Q-Q plots of H_2 subband coefficients

the rest of the paper, small boldface letters denote vectors, big boldface vectors denote matrices. The rule for additive embedding of the watermark sequence in the transform coefficients is given by

$$y[t] = x[t] + \alpha w[t], \quad t \in 1, \ldots, N \tag{4}$$

where $\alpha \in \mathbb{R}$ denotes the watermark embedding power, $y[t]$ denotes a watermarked transform coefficient and $x[t]$ denotes a host image transform coefficient. Now, given that the original DWT subband coefficients are considered as a realization of i.i.d. random variables following a Cauchy distribution with parameter γ and $\delta = 0$, our signal detection problem can be formulated as the detection of a deterministic signal of unknown amplitude (i.e. our watermark) in Cauchy distributed noise (i.e. the subband coefficients) with unknown shape parameter γ. This is actually a composite hypothesis testing problem, which can be formulated as a two-sided parameter test. The null- ($\mathcal{H}_0$) and alternative hypothesis ($\mathcal{H}_1$) of this parameter test are given by

$$\begin{aligned} \mathcal{H}_0 &: \alpha = 0, \gamma \text{ (no/other watermark)} \\ \mathcal{H}_1 &: \alpha \neq 0, \gamma \text{ (watermarked).} \end{aligned} \tag{5}$$

Since this test is two-sided, we know that no uniformly most-powerful (UMP) test exists here [14]. An approach to tackle the problem of unknown amplitude is to construct a Generalized Likelihood-Ratio Test (GLRT), where the unknown parameter is replaced by its MLE. However, deriving a GLRT in our situation is tricky, since we would have to estimate the watermark embedding power under $\mathcal{H}_1$ using ML estimation. In [10], knowledge of the watermark embedding power is assumed and the Bayes decision rule is applied to derive a statistic for the detection of a bipolar watermark sequence of unknown amplitude in Generalized-Gaussian distributed noise. In contrast to that, we do not assume knowledge of α at the detector and follow the same approach as presented in [21] or more generally in [11], which leads to a so called Rao hypothesis test. Since we rely on the assumption that the noise follows a Cauchy distribution, our detector will be termed the Rao-Cauchy (RC) detector.

3.1 The Rao-Cauchy (RC) Detector

In case we have large data records (which we can safely assume in image watermarking), two hypothesis tests exist, namely the Wald and the Rao test, which show the same asymptotic performance as the GLRT [13]. More precisely, these tests are equivalent to the GLRT when $N \to \infty$ (sample size). The Rao test has the advantage that it only requires computation of the MLEs under the null-hypothesis, whereas a GLRT would require computation of the MLEs under both hypothesis. The Rao test thus simplifies our problem considerably, since we do not have to find the MLE of α under $\mathcal{H}_1$ any more. However, we have to deal with one additional parameter γ, which is of no direct concern, but has an effect on the PDFs under $\mathcal{H}_0$ and $\mathcal{H}_1$. Such a parameter is called a *nuisance* parameter. In our case, this parameter is unknown and has to be estimated. With regards to this setup, the Rao test statistic ρ is given by [22, 13]

$$\rho(\mathbf{y}) = \left.\frac{\partial \log p(\mathbf{y}|\theta)}{\partial \alpha}\right|^{T}_{\theta=\hat{\theta}} \mathbf{V}(\hat{\theta}) \left.\frac{\partial \log p(\mathbf{y}|\theta)}{\partial \alpha}\right|_{\theta=\hat{\theta}}, \tag{6}$$

with $\theta = [\alpha \ \gamma]$ and $\hat{\theta}$ denotes the MLE of θ under $\mathcal{H}_0$. Since we know that under the null-hypothesis the watermark embedding power is $\alpha = 0$, we have $\hat{\theta} = [0 \ \hat{\gamma}]$. The term

$\mathbf{V}(\hat{\theta})$ is given by

$$\mathbf{V}(\hat{\theta}) = \mathbf{V}(0,\hat{\gamma}) = \left[\mathbf{I}_{\alpha\alpha}(0,\hat{\gamma}) - \mathbf{I}_{\alpha\gamma}^{T}(0,\hat{\gamma})\mathbf{I}_{\gamma\gamma}(0,\hat{\gamma})^{-1}\mathbf{I}_{\alpha\gamma}(0,\hat{\gamma})\right]^{-1}. \quad (7)$$

where $\mathbf{I}_{\alpha\alpha}(\alpha,\gamma)$ and $\mathbf{I}_{\alpha\gamma}(\alpha,\gamma)$ are partitions of the Fisher information matrix. The general form of the Fisher information matrix for a p-dimensional parameter vector $\theta = [\theta_1,\ldots,\theta_p]$ is given by [12]

$$\mathbf{I}_{ij}(\theta) = -\mathbb{E}\left(\frac{\partial^2 \log p(\mathbf{y}|\theta)}{\partial\theta_i \partial\theta_j}\right), \quad (8)$$

with $i = 1,\ldots,p$ and $j = 1,\ldots,p$. In our setup, we have $\theta = [\alpha\ \gamma]$, hence the parameter vector is two-dimensional. A theorem, which is of fundamental importance for the derivation of our detector is presented in [11]. It states, that in case the noise PDF is symmetric, the element $\mathbf{I}_{\alpha\gamma}$ of the Fisher information matrix becomes zero. Due to the asymptotic equivalence of the GLRT and Rao hypothesis test [22], this further implies that the Rao detector has the same asymptotic performance as a clairvoyant GLRT (i.e. where γ is known). Since the PDF of a Cauchy distribution with $\delta = 0$ is actually symmetric around zero, we can exploit this theorem and the term $\mathbf{V}(0,\hat{\gamma})$ reduces to

$$\mathbf{V}(0,\hat{\gamma}) = \mathbf{I}_{\alpha\alpha}(0,\hat{\gamma}). \quad (9)$$

By assuming that the transform coefficients are realizations of N i.i.d. random variables following a Cauchy distribution, the Rao test of Eq. (6) can now be formulated as

$$\rho(\mathbf{y}) = \left[\sum_{t=1}^{N} \frac{\partial \log p(y[t] - \alpha w[t]|\hat{\gamma})}{\partial\alpha}\bigg|_{\alpha=0}\right]^2 \mathbf{I}_{\alpha\alpha}^{-1}(0,\hat{\gamma}). \quad (10)$$

Depending on whether the null- or alternative-hypothesis is actually true, the detection statistic ρ follows [11]

$$\rho \overset{a}{\sim} \begin{cases} \chi_1^2, & \text{under } \mathcal{H}_0 \\ \chi_{1,\lambda}^2, & \text{under } \mathcal{H}_1 \end{cases} \quad (11)$$

where χ_1^2 denotes the Chi-Square distribution with one degree of freedom and $\chi_{1,\lambda}^2$ denotes a *non-central* Chi-Square distribution with one degree of freedom and non-centrality parameter λ. The $\overset{a}{\sim}$ denotes that this is the asymptotic distribution of ρ. Next, we derive the Rao test statistic for the Cauchy distribution and later return to the computation of the non-centrality parameter λ. In the first step we calculate the numerator of Eq. (10) as

$$\left[\sum_{t=1}^{N} \frac{\partial \log p(y[t] - \alpha w[t]|\hat{\gamma})}{\partial\alpha}\bigg|_{\alpha=0}\right]^2 \overset{(2)}{=} \left[-\sum_{t=1}^{N} \frac{2y[t]w[t]}{\hat{\gamma}^2\left(1+\frac{y[t]^2}{\hat{\gamma}^2}\right)}\right]^2 = 4\left[\sum_{t=1}^{N} \frac{y[t]w[t]}{\hat{\gamma}^2 + y[t]^2}\right]^2. \quad (12)$$

Next, we derive $\mathbf{I}_{\alpha\alpha}(\alpha,\gamma)$ and evaluate it at $\mathbf{I}(0,\hat{\gamma})$. According to [11, 13] we have

$$\mathbf{I}_{\alpha\alpha}(0,\gamma) = i(\alpha)\sum_{i=1}^{N} w[t]^2 \quad (13)$$

with $i(\alpha)$ (termed the *intrinsic accuracy* [8]) given by

$$i(\alpha) = \int_{-\infty}^{\infty} \frac{1}{p(y|\gamma)}\left(\frac{dp(y|\gamma)}{dy}\right)^2 dy. \quad (14)$$

Inserting the PDF of the Cauchy distribution (see Eq. (2)) and solving the integral yields

$$i(\alpha) = \frac{1}{2\gamma^2}. \quad (15)$$

Since we embed a bipolar watermark sequence, i.e. $\forall t : w[t] \in \{+1,-1\}$, the final detection statistic of the Rao detector becomes

$$\rho(\mathbf{y}) \overset{(12),(13),(15)}{=} \left[\sum_{t=1}^{N} \frac{y[t]w[t]}{\hat{\gamma}^2 + y[t]^2}\right]^2 \frac{8\hat{\gamma}^2}{N}. \quad (16)$$

In order to study the theoretical performance of the detector, it is further necessary to derive a closed form solution for the non-centrality parameter λ which is required for the computation of $\chi_{1,\lambda}^2$. From [13] we know that

$$\lambda = \alpha^2\left[\mathbf{I}_{\alpha\alpha}(0,\gamma) - \mathbf{I}_{\alpha\gamma}^{T}(0,\gamma)\mathbf{I}_{\gamma\gamma}^{-1}(0,\gamma)\mathbf{I}_{\alpha\gamma}(0,\gamma)\right]. \quad (17)$$

Again, using the fact that $\mathbf{I}_{\alpha\gamma}(\alpha,\gamma) = 0$, we finally obtain

$$\lambda = \alpha^2\mathbf{I}_{\alpha\alpha}(0,\gamma) \overset{(13)}{=} \frac{N\alpha^2}{2\gamma^2}. \quad (18)$$

This result can now be used to obtain theoretical Receiver-Operating Characteristic (ROC) curves, where we plot the probability of false alarm (P_f) against the probability of missing the watermark (P_m). P_f is defined as the probability that $\rho(\mathbf{y})$ is greater than a given threshold T, although H_0 is actually true and P_m is defined as the probability that $\rho(\mathbf{y})$ is smaller than T although H_1 is true. Obviously, P_f and P_m can be calculated on the basis of Eq. (11). For example, P_f follows from

$$P_f = \mathbb{P}\{\rho > T|\mathcal{H}_0\} = \mathrm{Q}_{\chi_1^2}(T), \quad (19)$$

where $\mathrm{Q}_{\chi_1^2}(\cdot)$ denotes the Q-function to compute right-tail probability (i.e. the probability of exceeding a given value) for a Chi-Square random variable. However, from [1] we know that this right-tail probability can also be expressed in terms of the Q-function to express right-tail probabilities of the Gaussian distribution using the identity $\mathrm{Q}_{\chi_1^2}(x) = 2\,\mathrm{Q}(\sqrt{x})$. Hence, we can express P_f as

$$P_f = 2\,\mathrm{Q}(\sqrt{T}). \quad (20)$$

The Q-function and its relation to the complementary error function $\mathrm{erfc}(\cdot)$ is given by

$$\mathrm{Q}(a) = \frac{1}{2\sqrt{\pi}}\int_a^{\infty} \exp\left(-x^2\right) dx = \frac{1}{2}\,\mathrm{erfc}\left(\frac{a}{\sqrt{2}}\right), \quad (21)$$

which is quite useful for implementation purposes. Accordingly, if a random variable follows a non-central Chi-Square distribution with one-degree of freedom and parameter λ, this is equivalent to the square of a Gaussian random variable with mean $\sqrt{\lambda}$ and variance $\sigma^2 = 1$. The probability of missing the watermark can again be expressed in terms of the Q-function as

$$P_m = 1 - P_d = 1 - \mathbb{P}(\rho > T|\mathcal{H}_1) = 1 - \mathrm{Q}(\sqrt{T} - \sqrt{\lambda}) + \mathrm{Q}(\sqrt{T} + \sqrt{\lambda}), \quad (22)$$

where P_d denotes the probability of detection. We can now substitute $\sqrt{T} = \mathrm{Q}^{-1}(P_f/2)$ into Eq. (22) and establish the connection between P_f and P_m as

$$P_m = 1 - \mathrm{Q}(\mathrm{Q}^{-1}(P_f/2) - \sqrt{\lambda}) - \mathrm{Q}(\mathrm{Q}^{-1}(P_f/2) + \sqrt{\lambda}). \quad (23)$$

In order to compute a threshold for the proposed detector, we simply have to manipulate Eq. (20) to obtain an explicit expression for T as

$$T = \left[\mathrm{Q}^{-1} \left(\frac{P_f}{2} \right) \right]^2 \tag{24}$$

and decide in favor of $\mathcal{H}_1$ if $\rho > T$ for a given probability of false-alarm P_f.

3.2 Some Notes on Multichannel Detection

So far, we have only addressed the detection problem when just one detail-subband is used for embedding the watermark sequence. However, in case we want to mark more than one subband we have to discuss how to combine the detector responses and how to determine a suitable global detection threshold. We will consider the straightforward approach of simply summing up the detector responses of each channel (i.e. subband). In order to derive a model for the global detection statistic, we assume that the detector responses of each channel are independent. This is reasonable in a sense, since the DWT detail coefficients of each subband are actually expansion coefficients of a projection of the original signal onto orthogonal subspaces of $L^2(\mathbb{R})$. Further, the watermark sequences are independent. This assumption allows to exploit the reproductivity property of the Chi-Square distribution, namely that the sum of Chi-Square random variables is again Chi-Square distributed. Formally, if we have N random variables $\rho_1, \ldots, \rho_N$, which all follow Chi-Square distributions χ^2_1 with one degree of freedom, we have

$$\sum_{i=1}^{N} \rho_i \sim \chi^2_v, \tag{25}$$

with $v = N$. We can now determine a threshold for the summed up detector responses using the inverse of the Chi-Square CDF. For the other detectors studied (linear-correlation, Generalized Gaussian, Cauchy), multichannel detection can be accomplished accordingly, since we know that the detection statistics are asymptotically Gaussian and a sum of Gaussian random variables is again Gaussian.

4. EXPERIMENTAL RESULTS

In this section, we evaluate the performance of the proposed Rao-Cauchy (RC) detector and compare it to the traditional LC detector, the Generalized Gaussian (GG) detector of [10] and the Cauchy detector derived in [4]. The implementation[1] of the detectors and all experiments were conducted in MATLAB. In order to determine experimental ROC curves we use Monte-Carlo simulation with 10^4 randomly generated watermarks to obtain the detection statistics under $\mathcal{H}_0$ (no watermark) and $\mathcal{H}_1$ (watermarked). For all detectors we use a Chi-Square Goodness-of-Fit (GoF) test at the 5% significance level to verify that the detection statistic distributions under $\mathcal{H}_0$ are either Gaussian (in case of the LC, GG and Cauchy detector) or Chi-Square with one degree of freedom (in case of the RC detector). Further, we estimate the distribution parameters of the detection statistics and check if they conform to the theoretical values. This is especially important under $\mathcal{H}_0$, since a strong deviation between the theoretical and experimental values won't allow to set the probability of false-alarm. After this check, we determine the threshold T for the RC detector according to Eq. (24) and estimate λ from the detector responses under $\mathcal{H}_1$. We can then determine the experimental probability of miss $\hat{P}_m$ from Eq. (23). Plotting P_f against $\hat{P}_m$ results in the desired experimental ROC curves. The detection thresholds for the other detector are chosen according to [2, 4, 10] and the probability of miss is calculated in the same way.

[1]Software will be available under: `http://www.wavelab.at/sources`

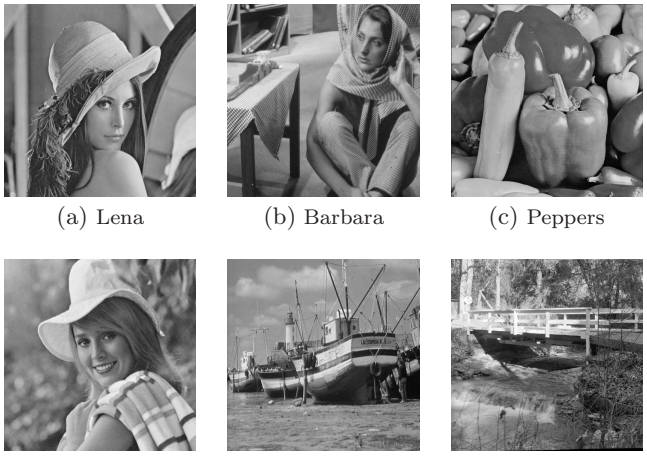

(a) Lena (b) Barbara (c) Peppers

(d) Elaine (e) Boat (f) Bridge

Figure 3: Test images

Regarding the choice of DWT wavelet filter, decomposition depth and embedding subband, we employ the biorthogonal CDF 9/7 filter and embed the watermark sequence in the horizontal detail subband on level two ($\mathbf{H}_2$). The embedding power α of the watermark is chosen so that a Document-to-Watermark Ratio (DWR) of 20dB, 23dB or 25dB is obtained. The resulting image PSNRs vary from ∼47dB to ∼54dB as we can read off Table 4. Due to this high PSNR values no visual artifacts can be noticed in the watermarked images. The bipolar pseudo-random watermark sequence is generated using the Mersenne twister random number generator. We use six 128×128 pixel grayscale test images for our experiments, which are shown in Figure 3.

Regarding the GoF test results for the detection statistics, the null-hypothesis of the GoF test (either Gaussian or Chi-Square with one degree of freedom) could not be rejected for the majority of all cases. Table 2 lists the GoF test results of the RC detector statistics when no watermark is present ($\mathcal{H}_0$ of the detector). Since we determine the experimental detection statistics when no watermark is present by embedding a randomly generated watermark sequence and then try to detect 10^4 different watermark sequences, the test outcomes can vary for different DWRs. A "reject" signifies that the GoF test null-hypothesis was rejected and a "−" signifies no rejection. As we can see, the null-hypothesis cannot be rejected in all cases. We do not list the test outcomes for the other detectors, since these check results have already been reported in the literature.

The experimental ROC curves of all six test images are shown in Figure 4 for a DWR of 25dB. Following the analysis in [21], the watermark power was set to $\alpha = 1$ for the Cauchy and GG detector to simulate the situation that nothing is known about the embedding strength at the detector. With-

DWR [dB]	PSNR [dB]					
	Lena	Barbara	Peppers	Boat	Bridge	Elaine
20	48.41	47.95	47.88	50.40	49.51	49.28
23	50.93	50.65	50.59	51.92	51.44	51.32
25	51.82	51.59	51.55	53.79	52.67	52.40

Table 1: Document-to-Watermark Ratio (DWR) and resulting PSNR for the test images

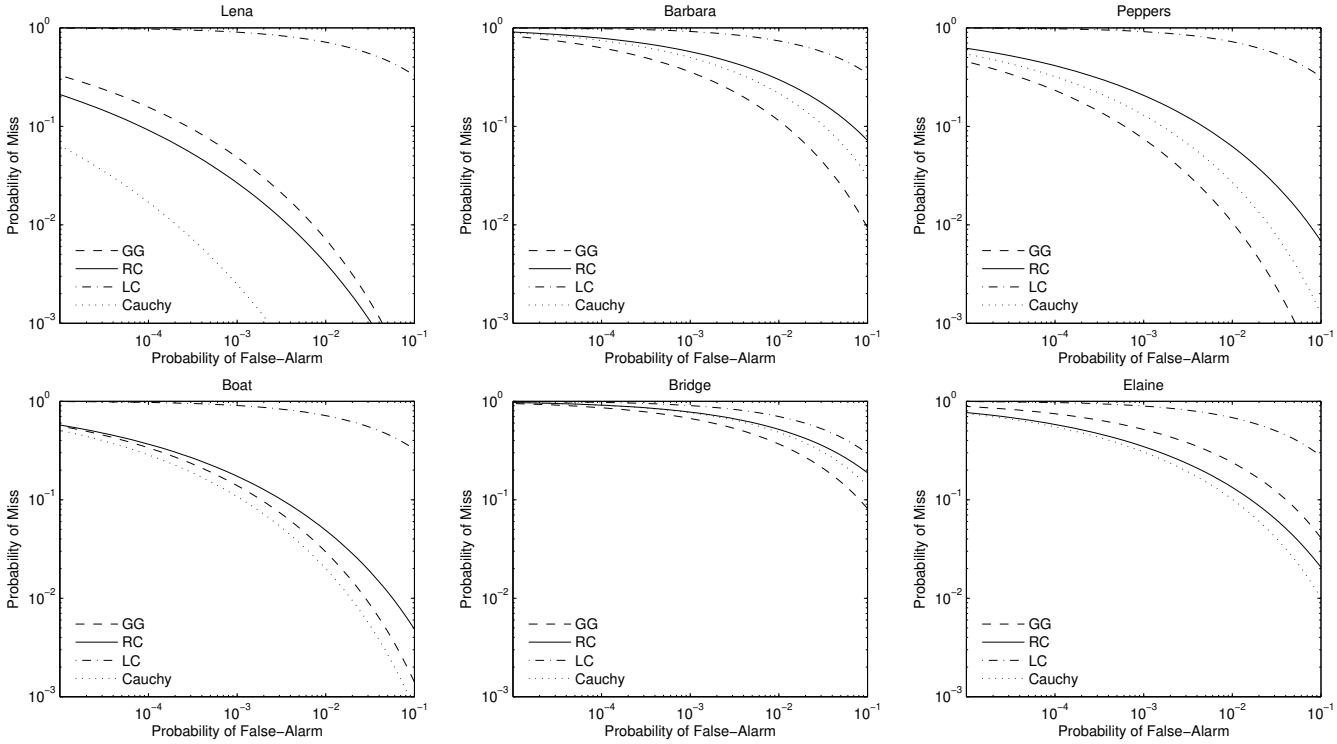

Figure 4: Experimental ROC curves at DWR 25dB

Images	DWR [dB]		
	20	23	25
Lena	–	–	–
Barbara	–	–	–
Peppers	–	–	–
Boat	–	–	–
Bridge	–	–	–
Elaine	–	–	–

Table 2: Chi-Square GoF test results at the 5% significance level

out any attacks, the Cauchy and GG detector show the best performance for our test images. The proposed Rao detector shows competitive performance in all cases and is even better than the GG detector for Lena and Elaine. In Section 4.2 we will further see that the Rao detector requires less computational effort than both the GG and Cauchy detector. We point out, that the quite good results for the Cauchy detector are caused by the fact that a DWR of 25dB leads to a true embedding power close to 1. This additionally confirms that the Cauchy distribution is a good model for the detail-subband coefficients as we have already seen in the Q-Q plots of Figure 2.

4.1 Performance Under Attacks

To study the performance of the RC detector under attacks, we simulate two attack scenarios: first, JPEG compression with quality factors ranging from $Q = 50$ to $Q = 90$ and second, JPEG2000 compression using Jasper (default settings) with compression ratios varying from 0.8bpp (rate 0.1) to 2.4bpp (rate 0.3). Again, we run Monte-Carlo simulations with 10^4 randomly generated watermarks to determine the detection statistics under the null- and alternative-hypothesis. First, we take a look at the JPEG compression results with a quality factor of $Q = 50$, which are shown in Figure 5 for the six test images at a DWR of 20dB. The resulting PSNRs of the compressed images are given on top of each figure. In this setting with a rather low quality value, the proposed RC detector outperforms the GG detector in all cases and is competitive to the Cauchy detector. Generally, we observe the behavior that as we increase the quality factor Q and choose a lower embedding strength (i.e. a higher DWR), the results became more similar to Figure 4, which is obvious since the impact of the attack is reduced. Due to space limitations, we cannot present the corresponding ROC plots here. The reason why the Cauchy distribution related detectors show such a good performance in this attack scenario is probably related to the fact that

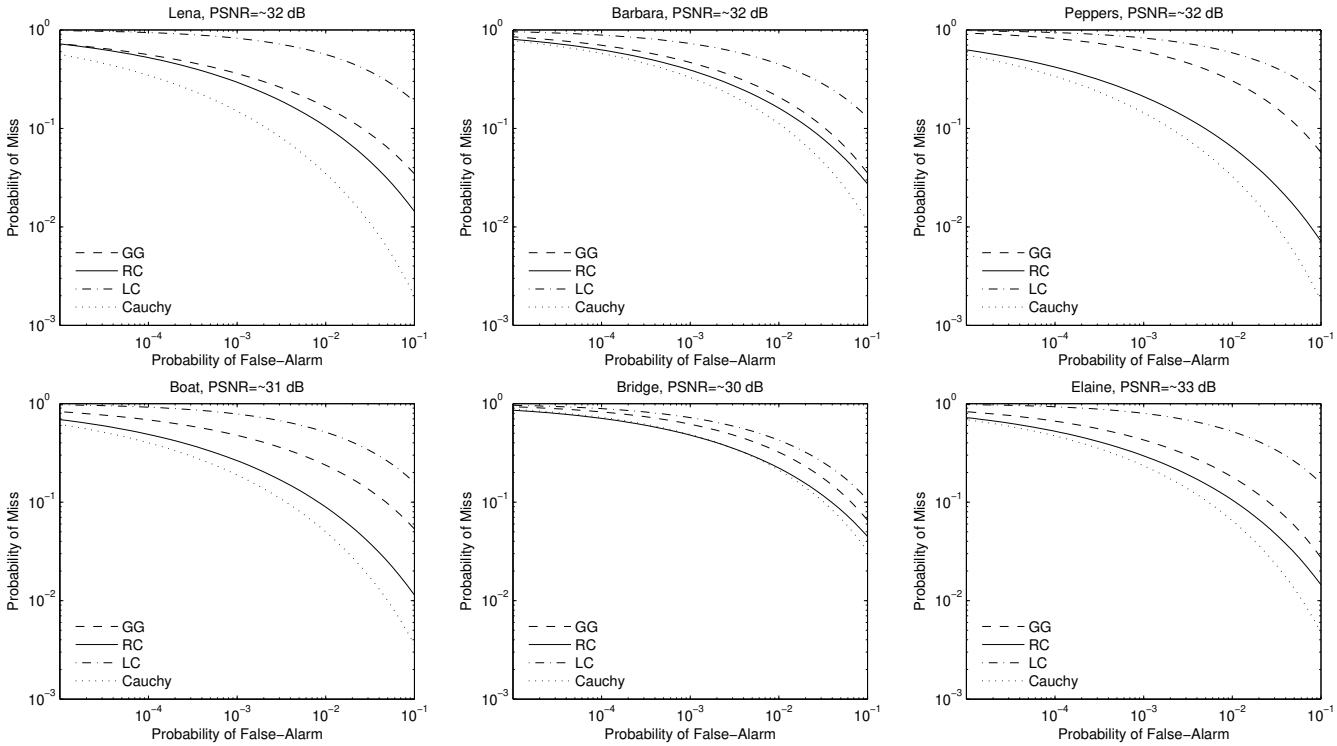

Figure 5: Performance under a JPEG compression attack with $Q = 50$ and DWR 20dB

the JPEG attack alters the embedding strength of the watermark as well as the subband statistics, which results in a performance loss for the GG detector.

Next, we study the performance under JPEG2000 attacks. The ROC curves for a compression rate of 2.4bpp at a DWR of 23dB are shown in Figure 6. In contrast to the JPEG compression attacks, we notice that the average PSNR is higher here, ranging from 36dB to 42dB. When the compression rate is increased so that the average PSNR goes down to the PSNR after JPEG compression, the detection results become almost unacceptable. This behavior can be attributed to the fact that JPEG2000 operates in the wavelet domain and possibly has a stronger influence on the detail subband statistics than traditional JPEG compression. In addition to that, we are marking just one detail subband here. Nevertheless, the RC detector shows good performance even under the JPEG2000 attack and again outperforms the GG and of course the LC detector.

4.2 Computational Analysis

In order to justify the term *lightweight* detector, we take a look at the computational effort which is required to employ the RC detector and compare it to the effort required for the other detectors. This includes a discussion of the number of required arithmetic operations to calculate the detection statistics, parameter estimation issues and the determination of detection thresholds. By arithmetic operations we understand the number of additions & subtractions $(+, -)$, multiplications & divisions $(\times, \div)$ and logarithms & exponentiations $(\log, \text{pow})$. For the sake of completeness, the detection statistics of the LC, GG and Cauchy detector are listed below:

- Linear-Correlation (LC) Detector

$$\rho(\mathbf{y}) = \frac{1}{N}\sum_{t=1}^{N} y[t]w[t] \tag{26}$$

- Generalized Gaussian (GG) Detector

$$\rho(\mathbf{y}) = \frac{1}{b^c}\sum_{t=1}^{N} \left(|y[t]|^c - |y[t] - \alpha w[t]|^c\right) \tag{27}$$

- Cauchy Detector

$$\rho(\mathbf{y}) = \sum_{t=1}^{N} \log\left(\frac{\gamma^2 + y[t]^2}{\gamma^2 + (y[t] - \alpha w[t])^2}\right) \tag{28}$$

In Table 3 we provide the number of operations as a function of the input vector length N. From these numbers it is obvious that the LC detector is by far the simplest in terms of arithmetic operations, since it involves only summations and multiplications of floating point numbers. Only the watermarked coefficients and the watermark sequence itself are involved. However, the RC detector is only slightly more expensive, since the exponentiations in Eq. (16) merely involve integer exponents and additions as well as multiplications which can be very efficiently performed with few CPU cycles. In contrast to that, the Cauchy detector requires $2N + 1$ computations of the logarithm and the GG detector even requires exponentiations with floating point numbers, which is very expensive in terms of CPU cycles.

Regarding parameter estimation issues, the LC detector is again the simplest one, since it requires no parameter estimation at all (see Eq. (26)), followed by the RC and Cauchy

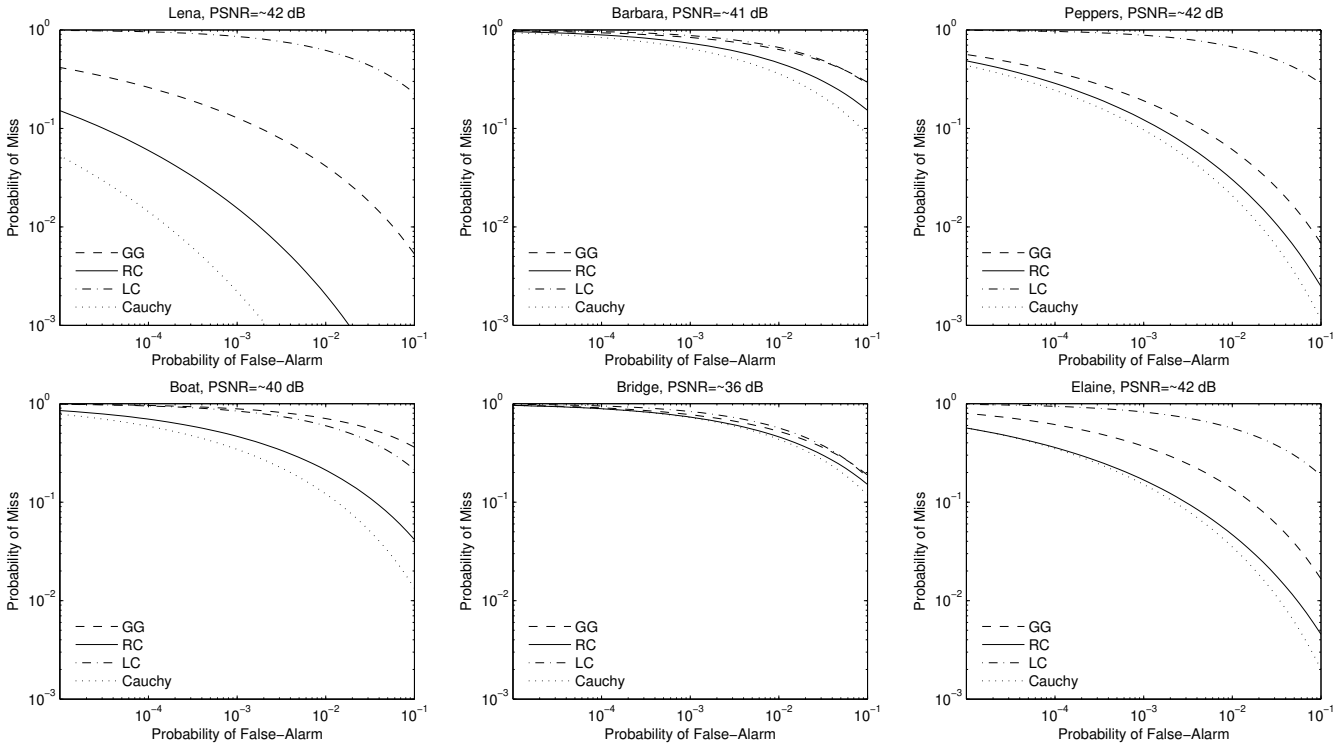

Figure 6: Performance under a JPEG2000 compression attack with 2.4bpp and a DWR of 23dB

Detector	Operations		
	$+,-$	$\times,\div$	pow , log
LC	N	N	
RC	$2N$	$N+1$	2
GG [10]	$2N$	1	$2N+1$
Cauchy [4]	$3N$	N	$2N+1$

Table 3: Number of arithmetic operations

detector, which both require to estimate the shape parameter γ of the Cauchy distribution. The MLE of γ can be numerically computed by solving Eq. (3). For our experiments, we implemented ML estimation using the Newton-Raphson root finding procedure. For that purpose, define the left hand side of (3) as a function of $\hat{\gamma}$, denoted by $g(\hat{\gamma})$. The estimate of γ in the k-th iteration of the Newton-Raphson algorithm is computed by

$$\gamma_{k+1} = \gamma_k - \frac{g(\gamma_k)}{g'(\gamma_k)}. \tag{29}$$

Here, $g'(\cdot)$ denotes the first derivative of g w.r.t. γ, given by

$$g'(\gamma) = \frac{4\gamma}{N} \sum_{t=1}^{N} \frac{x[t]^2}{\left(\gamma^2 + x[t]^2\right)^2}, \tag{30}$$

where the $x[t]$ denote N realizations of i.i.d. Cauchy random variables. As starting value $\hat{\gamma}_1$ for the numerical calculation, we used estimation values based on sample quantiles [15] to obtain a first estimate of γ as

$$\hat{\gamma}_1 = 0.5(x_p - x_{1-p}) \tan(\pi(1-p)), \tag{31}$$

with $0.5 < p < 1$ and x_p, x_{1-p} denoting the sample quantiles.

Moment estimation is not possible in case of the Cauchy distribution, since the moments do not exist. In case of the GGD, ML estimation of the shape c and scale parameter b requires to find the roots of the transcendental equation [9]

$$1 + \frac{\psi(1/\hat{c}) + \log\left(\frac{\hat{c}}{N}\sum_{t=1}^{N}|x[t]|^{\hat{c}}\right)}{\hat{c}} - \frac{\sum_{t=1}^{N}|x[t]|^{\hat{c}}\log(|x[t]|)}{\sum_{t=1}^{N}|x[t]|^{\hat{c}}} = 0, \tag{32}$$

where the $x[t]$ now denote N realizations of i.i.d. GGD random variables and $\psi(\cdot)$ denotes the digamma function [1], i.e. the logarithmic derivative of the gamma function. Inserting the MLE of c into

$$\hat{b} = \left(\frac{\hat{c}}{N}\sum_{t=1}^{N}|x[t]|^{\hat{c}}\right)^{\frac{1}{\hat{c}}} \tag{33}$$

then gives the MLE of b. Good starting values are usually obtained from moment estimates [20]. For our experiments, we used the Newton-Raphson algorithm proposed in [9] to solve the ML equation for c. Taking into account that this includes a numerical computation of the first polygamma function and exponentiations by floating point number, it is obvious that finding the solution to Eq. (32) requires more computational effort than solving Eq. (3). Our estimation experiments confirm that the Cauchy MLE procedure is faster by a factor of four.

Finally, we cover the effort for the determination of detection thresholds. In case of the LC, GG and Cauchy detector, we have to compute the mean and variance of the normally distributed detection statistic under $\mathcal{H}_0$ to determine a suit-

able threshold (see [2] for example). In contrast to that, the RC detector does not require to compute detection statistic parameters, since under $\mathcal{H}_0$ we know that $\rho \sim \chi^2_1$ and the threshold can easily be calculated from Eq. (24).

5. CONCLUSION

In this paper we presented a new, asymptotically optimal blind detector for additive spread-spectrum watermarking in the DWT domain. We showed that DWT detail subband coefficients can be modeled reasonably by a one-parameter Cauchy distribution and derived a watermark detector on the basis of the Rao hypothesis test without depending on knowledge of the watermark embedding strength. Estimation of the shape parameter of the Cauchy distribution is simple and computationally less expensive than estimation of the shape and scale parameter of the GGD. The computation of the detection statistic further requires fewer arithmetic operations than the GG detector and even the determination of detection thresholds is more straightforward. Our experimental results without any attacks indicate superior performance over the LC detector in all test scenarios and at least competitive performance compared to the GG and Cauchy detector. However, under JPEG and JPEG2000 compression attacks, the RC detector shows robust behavior and even better detection results than the GG detector, especially at high compression rates. Nevertheless, we have to point out that the reduced flexibility of the Cauchy distribution in relation to the GGD can result in lower detection performance in cases where the subband coefficient statistics strongly deviate from the Cauchy model. But setting the shape parameter of the GGD to a fixed value would equally lead to reduced flexibility. Further work on this topic includes the incorporation of a perceptual model, a more comprehensive performance evaluation under more attack scenarios and a thorough discussion on multichannel detection issues.

Acknowledgments

Supported by Austrian Science Fund project FWF-P19159-N13.

6. REFERENCES

[1] M. Abramowitz and I. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York, 1964.

[2] M. Barni and F. Bartolini. *Watermarking Systems Engineering*. Marcel Dekker, 2004.

[3] K. A. Birney and T. R. Fischer. On the modeling of DCT and subband image data for compression. *IEEE Transactions on Image Processing*, 4(2):186–193, Feb. 1995.

[4] A. Briassouli, P. Tsakalides, and A. Stouraitis. Hidden Messages in Heavy-Tails: DCT-Domain Watermark Detection Using Alpha-Stable Models. *IEEE Transactions on Multimedia*, 7(4):700–715, Aug. 2005.

[5] R. Chandramouli and N. D. Memon. On sequential watermark detection. *IEEE Transactions on Signal Processing*, 51(4):1034–1044, Apr. 2003.

[6] Q. Cheng and T. S. Huang. An additive approach to transform-domain information hiding and optimum detection structure. *IEEE Transactions on Multimedia*, 3(3):273–284, Sept. 2001.

[7] M. H. M. Costa. Writing on dirty paper. *IEEE Transactions on Information Theory*, 29(3):439–441, May 1983.

[8] D. Cox and D. Hinkley. *Theoretical Statistics*. Chapman & Hall/CRC, 1974.

[9] M. Do and M. Vetterli. Wavelet-based texture retrieval using Generalized Gaussian density and Kullback-Leibler distance. *IEEE Transactions on Image Processing*, 11(2):146–158, 2002.

[10] J. R. Hernández, M. Amado, and F. Pérez-González. DCT-domain watermarking techniques for still images: Detector performance analysis and a new structure. *IEEE Transactions on Image Processing*, 9(1):55–68, Jan. 2000.

[11] S. M. Kay. Asymptotically optimal detection in incompletely characterized non-gaussian noise. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 37(5):627–633, May 1989.

[12] S. M. Kay. *Fundamentals of Statistical Signal Processing: Estimation Theory*, volume 1. Prentice-Hall, 1993.

[13] S. M. Kay. *Fundamentals of Statistical Signal Processing: Detection Theory*, volume 2. Prentice-Hall, 1998.

[14] S. M. Kendall and A. Stuart. *The Advanced Theory of Statistics: Inference and Relationship*, volume 2. Macmillan, 1979.

[15] K. Krishnamoorthy. *Handbook of Statistical Distributions with Applications*. Chapman & Hall, 2006.

[16] W. Liu, L. Dong, and W. Zeng. Optimum detection for spread-spectrum watermarking that employs self-masking. *IEEE Transactions on Information Forensics and Security*, 2(4):645–654, Dec. 2007.

[17] S. Mallat. A theory for multiresolution signal decomposition: the wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, July 1989.

[18] H. S. Malvar and D. A. F. Florencio. Improved spread spectrum: A new modulation technique for robust watermarking. *IEEE Transactions on Signal Processing*, 51(4):898–905, Apr. 2003.

[19] N. Merhav and E. Sabbag. Optimal watermark embedding and detection strategies under limited detection resources. *IEEE Transactions on Information Theory*, 54(1):255–274, Jan. 2008.

[20] S. Nadarajah. A generalized normal distribution. *Journal of Applied Statistics*, 32(7):685–694, Sept. 2005.

[21] A. Nikolaidis and I. Pitas. Asymptotically optimal detection for additive watermarking in the DCT and DWT domains. *IEEE Transactions on Image Processing*, 12(5):563–571, May 2003.

[22] C. R. Rao. *Linear Statistical Inference and Its Applications*. Probability and Mathematical Statistics. Wiley, 1973.

[23] G. Van de Wouwer, P. Scheunders, and D. Van Dyck. Statistical texture characterization from discrete wavelet representations. *IEEE Transactions on Image Processing*, 8(4):592–598, Apr. 1999.