Performance analysis of nonuniform quantization-based data-hiding

J. E. Vila-Forcén, S. Voloshynovskiy*, O. Koval and T. Pun
University of Geneva, Department of Computer Science,
24 rue Général-Dufour, CH 1211, Geneva, Switzerland

ABSTRACT

In this paper, we tackle the problem of performance improvement of quantization-based data-hiding in the middle-watermark-to-noise ratio (WNR) regime. The objective is to define the quantization-based framework that maximizes the performance of the known-host-state data-hiding in the middle-WNR taking into account the host probability density function (pdf). The experimental results show that the usage of uniform deadzone quantization (UDQ) permits to achieve higher performance than using uniform quantization (UQ) or spread spectrum (SS)-based data-hiding. The performance enhancement is demonstrated for both achievable rate and error probability criteria.

Keywords: data-hiding, quantization-based, uniform quantization, uniform deadzone quantization, quantization index modulation, spread spectrum

1. INTRODUCTION

Digital data-hiding has appeared as an emerging solution for many different applications such as copyright protection, fingerprinting, authentication and tamper proofing. The design of practical data-hiding methods is a complex task, which involves a trade-off among security, visibility and performance in terms of achievable rate of communications or corresponding probability of error. Leaving the optimal solution to this trade-off outside of the scope of this paper, we will mostly concentrate our attention on the maximization of the achievable rate as well as on the analysis of the error probability under the additive white Gaussian noise (AWGN) channel.

In the scope of this study, two classes of practical embedding methods are of particular interest: known-host-statistics and known-host-state[1] techniques. The first group is mainly represented by the SS methods, where the data is embedded into the host image exploiting the host statistics at the encoder. The latter group of methods includes the quantization-based techniques, designed in order to approach a host-interference free communications.[2–4]

The analysis of quantization-based methods demonstrates their high efficiency in terms of both selected performance criteria at high-WNR regime while the in the low-WNRs significant performance loss of known-host-state methods is observed. The reason for that is closely related to the underlying assumption used in the development of these techniques. In particular, the host pdf is not taken into account in the design of quantizers. One can consider the scalar Costa scheme (SCS),[4] where it is explicitly assumed that the host pdf is flat, that corresponds to the assumption of an infinite host variance $\sigma_X^2 \to \infty$. This assumption justifies the selection of the UQ that was proved to be optimal in the rate-distortion sense.[5] However, such a model is not realistic and it does not correspond to the statistical properties of the real data.[6,7]

This incorrect assumption about the statistical properties of the host was already exploited[8,9] in order to reveal the real potential of quantization-based data-hiding techniques. The main obtained results state that, in order to be optimal, one still should provide two different methods of hidden information communications at high- and low-WNRs respectively. In particular, it was shown that the achievable rates in such a protocol can be made arbitrarily close to those provided by the SS techniques at low-WNR and by the SCS at high-WNR, respectively.

Another possibility to adopt the embedding strategy to real data statistics is to use a UDQ that is known to provide better performance than the UQ in terms of empirical achievable rate-distortion pairs.[10] The

*The contact author is S. Voloshynovskiy (email: svolos@cui.unige.ch). http://sip.unige.ch

first attempt to exploit such an embedding strategy in data-hiding was accomplished,[11–13] where performance improvement over the classical UQ-based techniques was reported in terms of the error probability for the AWGN and uniform noise channels assuming that the host pdf is flat within the quantization bin and therefore high-rate quantization assumption conditions are satisfied.

Extending the existing results obtained in,[11–13] we propose an optimally designed UDQ data-hiding scheme that outperforms the SS and UQ-based data-hiding methods in the middle-WNR and matches the performance of UQ-based data-hiding methods in the high-WNR. We select the error probability and the achievable rate of communications as performance measures relaxing the previous assumptions about host pdf flatness within the quantization bin.

The paper is organized as follows: a review of existing data-hiding methods as well as problem formulation are presented in Section 2. The experimental results that demonstrate the performance improvement of the UDQ-based over UQ-based data-hiding are presented in Section 3. Finally, Section 4 concludes the paper.

Notations We use capital letters X to denote scalar random variables, bold capital letters $\mathbf{X}$ to denote N-length vector random variables, corresponding small letters x and $\mathbf{x}$ to denote the realizations of respectively scalar and vector random variables. m represents the message and $\mathcal{M}$ the set of messages. $\mathbf{X} \sim f_{\mathbf{X}}(\mathbf{x})$ denotes the host signal distributed according to $f_{\mathbf{X}}(\mathbf{x})$, $\mathbf{Z} \sim f_{\mathbf{Z}}(\mathbf{z})$ represents the noise, $\mathbf{W} \sim f_{\mathbf{W}}(\mathbf{w})$ the watermark and $\mathbf{V} \sim f_{\mathbf{V}}(\mathbf{v})$ the received signal. The WNR is defined as $\text{WNR} = 10\log_{10}\frac{\sigma_W^2}{\sigma_Z^2}$, where σ_W^2 and σ_Z^2 stand for the variance of the watermark and the noise, respectively. The watermark-to-image ratio (WIR) is defined as $\text{WIR} = 10\log_{10}\frac{\sigma_W^2}{\sigma_X^2}$, where σ_X^2 denotes the variance of the host. The distortion-compensation parameter is denoted as α. The mathematical expectation of a random variable $X \sim p_X(x)$ is designated by $E_X[X]$ or simply by $E[X]$.

2. DATA-HIDING METHODS

Since it was realized that the problem of host interference cancellation plays a crucial role in the maximization of the achievable rate of communications in the data-hiding channels,[14] usually such protocols are modeled using a setup extensively studied by Gel'fand and Pinsker in digital communications. The Gel'fand-Pinsker formulation of the data-hiding problem is presented in Fig. 1, where the encoder produces a watermark $\mathbf{W}$ through $\phi : \mathcal{M}\times\mathcal{X}^N\times\mathcal{K} \rightarrow \mathcal{W}^N$, the stego data $\mathbf{Y}$ is obtained using the embedding mapping $\varphi : \mathcal{W}^N\times\mathcal{X}^N \rightarrow \mathcal{Y}^N$, and the decoder is trying to reveal the communicated message applying $\psi : \mathcal{V}^N \times \mathcal{K} \rightarrow \mathcal{M}$, where K denotes the key and $\mathcal{K}$ designates the set of keys.

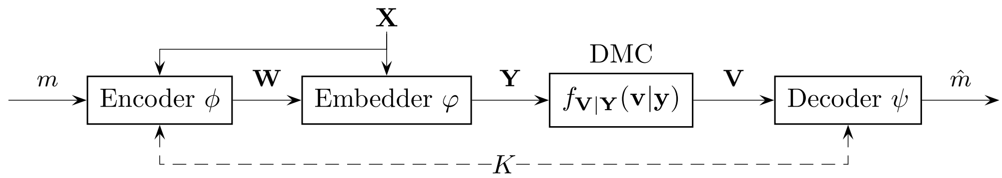

Figure 1. Gel'fand-Pinsker formulation of the data-hiding problem.

In data-hiding, embedding and attacking are constrained to guarantee the acceptable quality of the content. Therefore, the variances of the stego data and the attack are bounded. In this protocol, the key is necessary in order to define a particular codebook used for communications. Nevertheless, key management is outside of the scope of this paper and we will not consider it further.

Costa[2] considered the Gel'fand-Pinsker problem for the independent and identically distributed (i.i.d.) Gaussian formulation and the mean square error (MSE) distance. In the Costa setup, the embedder function and the channel are additive and they are defined by: $Y = W + X$ and $V = Y + Z$. As it was shown,[2] using a random binning-based codebook design and exploiting the knowledge of the channel statistics available at the encoder a priori to transmission, one can approach the capacity of the interference free AWGN channel.

It is interesting to note that existing practical data-hiding methods can be considered as approximations of the Costa coding for various WNR regimes.

2.1. SS-based data-hiding

SS-based data-hiding is a special case of the Costa coding optimized to the low-WNR conditions. The conclusions of this optimization leading to the watermark generation disregarding the host data is a poor performance at high-WNRs. The reason for this is an insufficient number of the codewords generated located in each codebook bin.

The stego data for SS-based data-hiding is obtained using the encoding function ϕ_{SS} as:

$$y = x + w. \tag{1}$$

The embedding distortion is given by $\sigma_W^2 = E[w^2]$.

2.1.1. Error probability

Assuming maximum likelihood (ML) decoding and binary signaling, the error probability is calculated as the corresponding integration over the error region $\overline{\mathcal{R}}_m$ of the equivalent noise pdf $Z_{eq1} = X + Z$.

We define $f_{Z_{eq_0}}(z_{eq0}, \beta, \sigma)$ as the convolution of a Laplacian pdf $\mathcal{L}(0, \beta)$ and a Gaussian pdf $\mathcal{N}(0, \sigma^2)$:

$$f_{Z_{eq_0}}(z_{eq0}, \beta, \sigma^2) = \frac{\beta}{2\sigma\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{-\beta|x|} e^{\frac{-(x - z_{eq0})^2}{2\sigma^2}} dx, \tag{2}$$

where β is the parameter of the Laplacian distribution and σ^2 the variance of the Gaussian pdf.

Thus, the equivalent noise pdf can be expressed using (2) as the convolution between the Laplacian host $X \sim \mathcal{L}(0, \frac{1}{\sigma_X^2})$ and the Gaussian noise $Z \sim \mathcal{N}(0, \sigma_Z^2)$ due to the independence of X and Z:

$$f_{Z_{eq_1}}(z_{eq1}) = f_{Z_{eq_0}}\left(z_{eq1}, \frac{1}{\sigma_X^2}, \sigma_Z^2\right). \tag{3}$$

The error probability is calculated as the integral of the equivalent noise pdf over the error region $\overline{\mathcal{R}}_m$:

$$P_e = E_{p_M}\left[\int_{\overline{\mathcal{R}}_m} f_{Z_{eq_1}}(z_{eq1}) dz_{eq1}\right], \tag{4}$$

where $p_M(m)$ is the marginal probability mass function (pmf) of the input messages and $\overline{\mathcal{R}}_m$ is defined by the likelihood ratio test decoder. Unfortunately no close form solution exists and numerical computations are needed to evaluate (4).

2.1.2. Achievable rates

The achievable rate of the SS-based data-hiding assuming the AWGN channel and Laplacian host X is given by the mutual information:

$$\begin{aligned} I(M; V) &= h(V) - h(V|M) \\ &= h(V) - h(X + Z). \end{aligned} \tag{5}$$

Denoting the equivalent noise $Z_{eq1} = X + Z$ we can write $f_{Z_{eq_1}}(z_{eq1}) = f_{Z_{eq_0}}(z_{eq1}, \frac{1}{\sigma_X^2}, \sigma_Z^2)$. Thus, the differential entropy $h(X + Z)$ is calculated as:

$$h(X + Z) = -E[\log_2 f_{Z_{eq_1}}(z_{eq1})]. \tag{6}$$

The differential entropy of the channel output V is:

$$h(V) = h(W + X + Z) = h(X + Z_{eq2}), \tag{7}$$

where $Z_{eq2} = W + Z$. The watermark pdf is Gaussian in order to maximize the achievable rate, thus $Z_{eq2} \sim \mathcal{N}(0, \sigma_W^2 + \sigma_Z^2)$ in this case. The pdf of V is calculated using (2) as $f_V(v) = f_{Z_{eq_0}}(v, \frac{1}{\sigma_X^2}, \sigma_W^2 + \sigma_Z^2)$. Finally, the achievable rate of communications of the SS-based data-hiding is calculated numerically.

2.2. Quantization-based data-hiding

In practice, it is not convenient to deal with the Costa data-hiding setup because of the Costa codebook exponential complexity. A number of practical algorithms exploit structured codebooks instead of random ones, and the most known discrete approximations to the Costa problem are known as quantization index modulation (QIM), distortion compensated quantization index modulation (DC-QIM)[3] and SCS.[4] The structured codebooks are designed using quantizers (or lattices[15]) in order to achieve host interference cancellation.

In this paper, we restrict our analysis to one dimension quantizers and binary signaling. The stego data for the DC-QIM case is obtained using the encoding function ϕ_{DC-QIM} as:

$$y = x + \alpha'(Q_m(x) - x), \tag{8}$$

where $0 < \alpha' \leq 1$ is the analogue of the Costa optimization parameter α, $Q_m(x)$ is a fixed scalar quantizer and m denotes the embedded message. If $\alpha' = 1$, the DC-QIM (8) simplifies to the QIM data-hiding method.

Under the assumption of a flat host pdf, the quantizers are usually designed to be uniform. Existing results[11–13] already demonstrated performance enhancement for a non-flat host using non-uniform quantization. Here, we assume binary signaling and a Laplacian host pdf. The designed quantizer preserves the uniformity everywhere besides in the vicinity of zero where the bin width (the width of the deadzone) is controlled through the parameter ρ. In Fig. 2 we represent the UDQ scheme†.

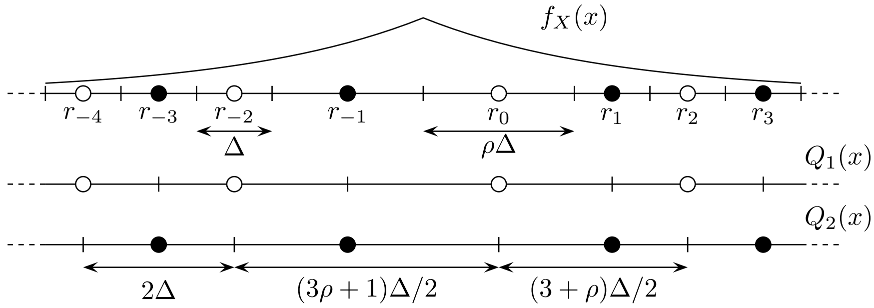

Figure 2. UDQ in one dimension and binary signaling case. The represented host pdf is Laplacian and $\mathcal{M} = \{1, 2\}$.

The reconstruction levels r_i are given by (Fig. 2):

$$r_i = \begin{cases} \rho\frac{\Delta}{2}, & i = 0; \\ \Delta(\rho + i - \frac{1}{2}), & i > 0; \\ -\rho\frac{\Delta}{2}, & i = -1; \\ -\Delta(\rho + |i| - \frac{3}{2}), & i < -1, \end{cases} \qquad i \in \mathcal{Z}, \tag{9}$$

where $\mathcal{Z}$ is the set of integer numbers. The even values of i correspond to the reconstruction levels for one embedded message and the odd reconstruction levels to the another one. For a given message, the quantization thresholds are given by the reconstruction levels of the other possible message.

The stego data pdf for a given embedded message is given by:

$$f_{Y|M}(y|m) = \begin{cases} \sum_i f_X\left(\frac{y - r_i}{1-\alpha'} + r_i\right)\frac{1}{1-\alpha'}, & |y - r_i| < (1-\alpha')\Delta; \\ 0, & \text{otherwise}, \end{cases} \tag{10}$$

where the summation is performed over the even or odd values of i depending on the embedded message. Assuming a Laplacian host pdf, the stego data pdf is represented in Fig. 3 for different values of the distortion compensation parameter. Contrarily to our previous work,[13] the pdf of the stego-data is not assumed to be flat within the quantization bin.

†We use the term QIM to reflect the fact that the UDQ does not possesses a regular structure.

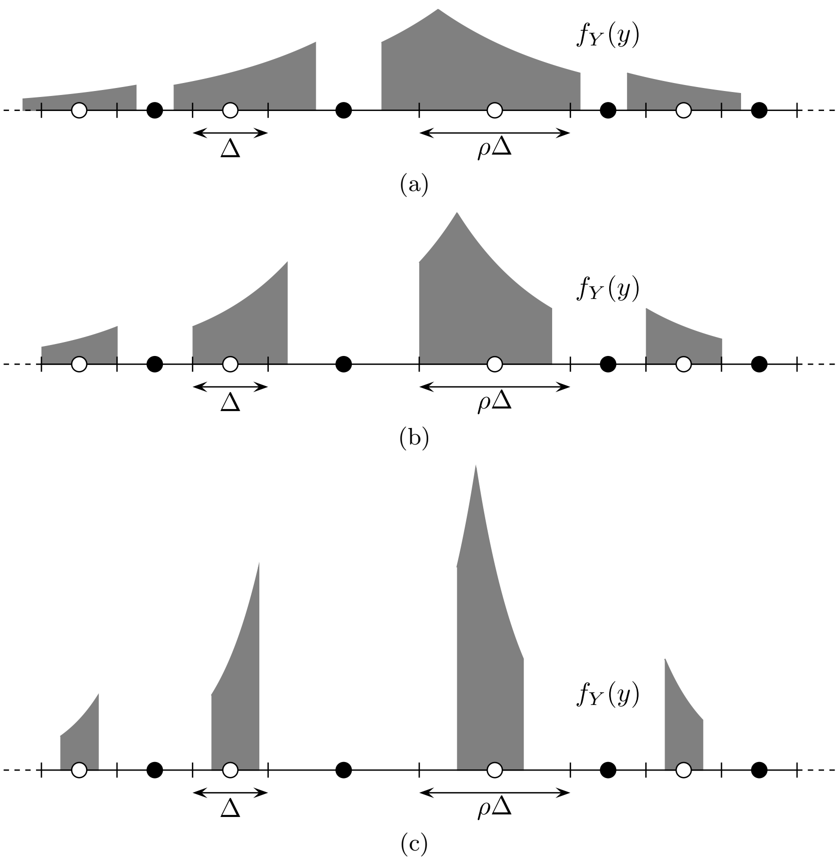

Figure 3. Stego-data pdf for one message assuming a Laplacian host pdf and $\rho = 2$ with (a) $\alpha' = 0.25$, (b) $\alpha' = 0.5$ and (c) $\alpha' = 0.75$.

The embedding distortion is given by:

$$\sigma_W^2 = \int_{-\infty}^{\infty} (x-y)^2 f_X(x)dx = {\alpha'}^2 E_{p_M} \left[\int_{-\infty}^{\infty} (Q_m(x) - x)^2 f_X(x) dx \right]. \tag{11}$$

For a given embedding variance and quantizer deadzone controlling parameter ρ, the value of Δ can be calculated using (11) in order to fix the quantizer design.

2.2.1. Error probability

The error probability measures the performance of the selected data-hiding technique assuming a fixed decoder structure. In general, it is defined as:

$$P_e = E_{p_M} \left[\Pr[\psi(V) \neq m | M = m] \right]. \tag{12}$$

In this paper, we restrict the analysis to a modified minimum Euclidean distance (MD) decoder. The decision boundaries are presented in Fig. 4. They can be calculated according to the following function:

$$T_j = \begin{cases} 0, & j = 0, \\ \operatorname{sgn}(j)\Delta(\rho + |j| - 1), & \text{otherwise}, \end{cases} \qquad j \in \mathcal{Z}, \tag{13}$$

where $\operatorname{sgn}(i)$ is the sign function.

It is possible to observe that the error probability might be nonzero even in a noiseless scenario (Fig. 3(a)), since some part of the stego-data pdf is in the error region if $\alpha' < 0.5$.

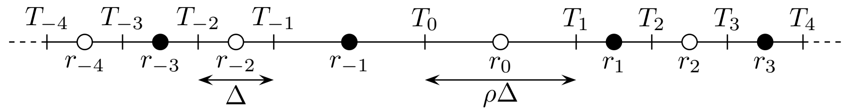

Figure 4. Modified MD decoder decision boundaries T_j and reconstruction points r_i of the UDQ.

Defining $\overline{\mathcal{R}}_m$ as the error region when the message m is transmitted and observing the symmetry of the construction, we can express the error probability in (12) as the expected value of the integral of the received signal pdf $f_V(\cdot)$ in the error region:

$$P_e = E_{p_M}\left[\int_{\overline{\mathcal{R}}_m} f_V(v)dv\right]. \tag{14}$$

Unfortunately, a closed analytical solution to (14) does not exist for the Laplacian host pdf and numerical computations are needed to evaluate it.

2.2.2. Achievable rates

Similarly to the SS-based data-hiding communications, the maximum achievable rate in the case of DC-QIM is given by the maximum of the mutual information:

$$R = \max_{\alpha',\rho} I(M;V). \tag{15}$$

It is known[16] that the mutual information can be expressed as a Kullback-Leibler divergence (KLD):

$$\begin{aligned} I(M;V) &= D(f_{M,V}(m,v)||f_V(v)p_M(m)) \\ &= E_{p_M}\left[\int f_{M,V}(m,v)\log_2\frac{f_{V|M}(v|M=m)}{f_V(v)}dv\right], \end{aligned} \tag{16}$$

where $f_{M,V}(m,v)$ represents the joint pdf of the input message and the channel output and $f_V(v)$ is the marginal pdf of the channel output.

Assuming equiprobable messages in the binary case ($p_M(m) = \frac{1}{2}$), (16) can be rewritten as the KLD between the received pdf when one of the messages has been transmitted and the average output pdf:

$$I(M;V) = D\left(f_{V|M}(v|M=1)||f_V(v)\right) \tag{17}$$

The output pdf is the convolution between the stego-data pdf given by (10) and the AWGN. The output pdf $f_V(v)$ can be calculated as:

$$f_V(v) = E_{p_M}\left[f_{V|M}(v|M=m)\right]. \tag{18}$$

Unfortunately, no closed analytical solution exists and the achievable rates are evaluated numerically.

3. EXPERIMENTAL RESULTS

The UDQ is optimized for the Laplacian host pdf and the AWGN channel. The embedding distortion was fixed to $\sigma_W^2 = 10$ and WIR $= -16$dB. The value of the quantization bin width Δ was calculated numerically according to (11). Under assumption of WIR $\to -\infty$, our results coincides with SCS.[4]

In case of the UDQ, the optimization has been performed on both the compensation parameter α' and the deadzone-to-bin width ratio ρ for each WNR. In Fig. 5(a), the performance comparison for SS, uniform deadzone quantization distortion compensated quantization index modulation (UDQ-DCQIM) and uniform quantization distortion compensated quantization index modulation (UQ-DCQIM) using error probability as cost function is presented.

The proposed UDQ-DCQIM design outperforms the classical UQ-DCQIM. The performance UDQ-DCQIM converges to the one of UQ-DCQIM in the high-WNRs while in the middle and low-WNR UDQ-DCQIM achieves smaller error probability. The obtained error probability performance coincides with the SCS for WIR $\to -\infty$, which evidently is worser because of neglecting actual host pdf. Nevertheless, their performance coincides at the high-WNR.

The achievable rates presented in Fig. 5(b) demonstrate the superior performance of UDQ-based data-hiding over UQ-based and SS. The presented results demonstrate that, if the host pdf is taken into account, the performance enhancement of UQ-DCQIM versus SCS is evident and the difference with SS becomes smaller. While in the high-WNR the UDQ design performance converges to the UQ one, in the middle and low-WNR the UDQ-DCQIM method demonstrates a superior performance than the SS. As in the error probability analysis case, the proposed method obtain the highest possible performance at all WNRs considered.

At low-WNR, it was shown[9] for a Gaussian host that it is possible to achieve the performance of SS-based data-hiding using SCS[4] data-hiding and a proper selection of the compensation parameter. Here, we demonstrate that the maximization of the achievable rate on the compensation parameter and on the design of the quantizers allow to provide a technique that outperforms SS-based data-hiding in all WNRs.

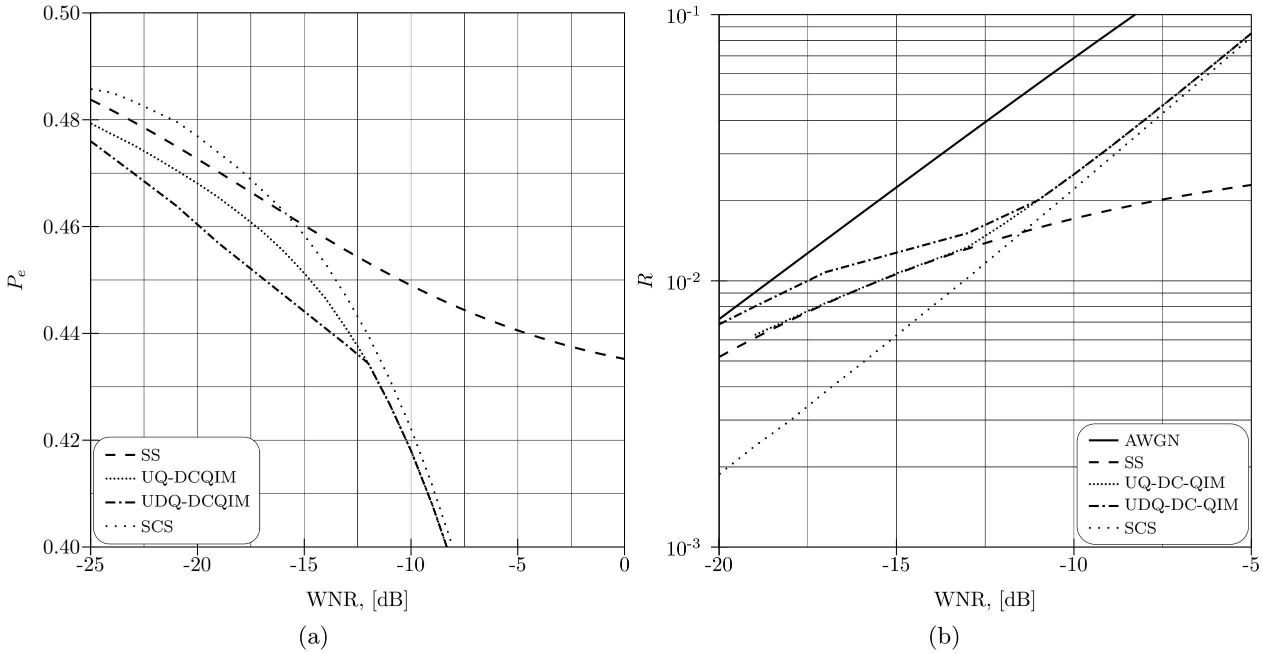

Figure 5. Performance comparison for SS, UQ-QIM and UDQ-QIM assuming WIR $= -16$dB, and SCS (WIR $= -\infty$dB) using a Laplacian host pdf under the AWGN channel using (a) error probability or (b) rate of reliable communications.

4. CONCLUSIONS

In this paper, we have considered the problem of performance improvement for quantization-based data-hiding techniques optimizing the quantizer design. The performance was evaluated assuming a real host pdf following a Laplacian distribution and the AWGN channel on the error probability assuming a MD decoder and on the achievable rates of communications. We propose a UDQ-DCQIM scheme with varying deadzone size and compensation factor. The obtained results demonstrate that the performance assuming a realistic host (i.e., finite WIR) differs from the one presented in the SCS setup in the middle and low-WNR. In that WNR region, typically SS-based data-hiding was used. Nevertheless, the proposed setup demonstrates a better performance than existing data-hiding techniques, allowing to use a unique technique for all WNRs. The demonstrated performance critically relies on the availability of the noise variance at the encoder prior to transmission.

5. ACKNOWLEDGEMENTS

This paper was partially supported by SNF Professorship grant No PP002-68653/1, Interactive Multimodal Information Management (IM2) project and by the European Commission through the IST Programme under Contract IST-2002-507932 ECRYPT and FP6 - 507609 - SIMILAR. The authors are thankful to the members of the Stochastic Image Processing group at University of Geneva.The information in this document reflects only the author's views, is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

REFERENCES

1. 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 Transactions on Signal Processing, Special Issue on Signal Processing for Data Hiding in Digital Media and Secure Content Delivery* **51**, pp. 960–980, Apr. 2003.
2. M. Costa, "Writing on dirty paper," *IEEE Transactions on Information Theory* **29**, pp. 439–441, May 1983.
3. B. Chen and G. W. Wornell, "Quantization index modulation: A class of provably good methods for digital watermarking and information embedding," *IEEE Transactions on Information Theory* **47**, pp. 1423–1443, May 2001.
4. J. J. Eggers, R. Bäuml, R. Tzschoppe, and B. Girod, "Scalar costa scheme for information embedding," *IEEE Transactions on Signal Processing* **51**, pp. 1003–1019, Apr. 2003.
5. H. Gish and J. Pierce, "Asymptotically efficient quantizing," *IEEE Transactions on Information Theory* **14**, pp. 676–683, Sept. 1968.
6. J. Vila-Forcen, O. Koval, and S. Voloshynovskiy, "Distributed single source coding with side information," in *Proceedings of IS&T/SPIE 16th Annual Symposium: Electronic Imaging 2004, Image Processing*, (San Jose, USA), 18-22 Jan. 2004.
7. S. Voloshynovskiy, O. Koval, and T. Pun, "Wavelet-based image denoising using non-stationary stochastic geometrical image priors," in *Proceedings of IS&T/SPIE 15th Annual Symposium: Electronic Imaging, Image and Video Communications and Processing V*, (Santa Clara, California USA), 20-24 Jan. 2003.
8. L. Pérez-Freire and F. Pérez-González, "Spread-spectrum vs. quantization-based data hiding: Misconceptions and implications," in *Proceedings of IS&T/SPIE 17th Annual Symposium: Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VII*, (San Jose, California, USA), Jan. 2005.
9. L. Pérez-Freire, F. Pérez-González, and S. Voloshynovskiy, "Revealing the true achievable rates of scalar costa scheme," in *IEEE International Workshop on Multimedia Signal Processing*, (Siena, Italy), 29 Sept.-1 Oct. 2004.
10. J. K. Su, *Adaptive Rate-Constrained Transform Video Coding*. PhD thesis, School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA, Dec. 1997. Thesis advisor: R. M. Mersereau.
11. O. Koval, S. Voloshynovskiy, F. Deguillaume, F. Pérez-González, and T. Pun, "Robustness improvement of known-host-state watermarking using host statistics," in *Proceedings of IS&T/SPIE 17th Annual Symposium: Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VII*, (San Jose, USA), 16-20 Jan. 2005.
12. O. Koval, S. Voloshynovskiy, F. Pérez-González, F. Deguillame, and T. Pun, "Quantization-based watermarking performance improvement using host statistics: AWGN attack case," in *ACM Multimedia and Security Workshop. Proceedings*, (Magdeburg, Germany), 20-21 Sept. 2004.
13. O. Koval, S. Voloshynovskiy, J. E. Vila-Forcén, F. Pérez-González, F. Deguillaume, and T. Pun, "Robustness improvement of known-host-state watermarking using host statistics," *Signal Processing* , 2005. (accepted).
14. I. J. Cox, M. L. Miller, and J. A. Bloom, *Digital Watermarking*, Morgan Kaufmann Publishers, 2002.
15. U. Erez and R. Zamir, "Lattice decoding can achieve 0.5 log(1+snr) over the additive white gaussian noise channel using nested codes," in *IEEE International Symposium on Information Theory*, p. 125, (Washington DC, USA), June 2001.
16. T. Cover and J. Thomas, *Elements of Information Theory*, Wiley and Sons, New York, 1991.