Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

A comparative investigation of information loss due to variable quantization on parameter estimation of compound distribution ⋆

Ruslan Seifullaev *, **Steffi Knorn** *,**, **Anders Ahlén** *

* *Division of Signals and Systems, Department of Engineering Sciences, Uppsala University, Sweden. (e-mail: ruslan.seifullaev@angstrom.uu.se, steffi.knorn@angstrom.uu.se, anders.ahlen@angstrom.uu.se)*
** *Institute for Automation Engineering, Otto von Guericke University Magdeburg, Germany.*

Abstract: In this paper we study the problem of how quantization may affect the maximum likelihood estimation of the parameters of a probability density function representing a compound distribution. We consider and compare three different approaches to design a variable quantizer allowing to guarantee a predefined loss of Fisher information which is used as a measure of the information loss due to quantization. We also propose the approximations which characterize the asymptotic behavior of the loss allowing a significant reduction of the computational complexity.

Keywords: Estimation, quantized systems, communication networks

1. INTRODUCTION

Due to radical changes in the industry, often called the fourth industrial revolution, more and more wireless devices are used in industrial applications, including wireless sensor networks, which provide information to control the plant. Hence, it must be ensured that sensor information is transferred successfully, which highly depends on the available radio channel. Thus, a better understanding of the channel propagation characteristics and associated statistical models are important in such applications.

Fading models based on compound distributions are most suitable to describe the radio channel power gain over long time horizons, Agrawal et al. (2014). A compound distribution may arise as the convolution of a lognormal (LN) distribution that is used to model a shadowing component of the channel power gain and a Gamma (G) distribution that models a fast fading component. The choice of a correct compound distribution can have a considerable impact on the latency, the energy consumption, and the average bit error rate (BER) of the wireless sensor network, see Olofsson et al. (2016); Agrawal et al. (2014); Shao and Beaulieu (2010); Gungor and Hancke (2009); Olofsson and Ahlén (2018); Croonenbroeck et al. (2017).

Obtain accurate parameter estimates is essential for the design of wireless control systems for industrial use. The parameters of the compound distribution can be estimated by the maximum likelihood (ML) method providing consistent estimates, Olofsson et al. (2016); Dogandzic and Jin (2004). However, the received signal strength (RSS) is measured using sensors with a certain resolution, i.e., data are received in quantized bins. When dealing with quantized observations, some amount of information that the observations carry about the unknown parameters may be lost. Then the appropriate quantization interval should be chosen to bound the maximal information loss.

The authors' previous paper, Seifullaev et al. (2019), studied the case of uniform quantization. But the results can be significantly improved when a variable quantizer is used. This paper considers three different approaches for variable quantizer designs: random perturbations, contraction mapping, and two-sided uniform quantization. For the latter two, an approximation of the loss of the information is proposed that significantly reduces the computational complexity. As a measure of the information lost due to quantization, we use the Fisher information (a classical measure of information that observations carry about unknown parameters of a distribution).

The rest of the paper is structured as follows. In Section 2 the problem of how to characterize the loss of Fisher information due to quantization is considered. Section 3 describes the results for uniform quantization. In Section 4 the methods of how to design the variable quantizers are considered and compared. Section 5 draws conclusions.

2. PROBLEM FORMULATION

Consider the compound distribution

$$p(y \mid \sigma, m) = \int_{-\infty}^{\infty} p_1(y - v \mid m)\, p_0(v \mid \sigma) dv, \qquad (1)$$

⋆ This work was supported by the Swedish Research Council (VR) under contract Dnr: 2017-04186

where y is continuous received signal power in dBm, and p_1 and p_0 are the dB representations of the G- and LN-distributions with the parameters m and σ, respectively, see Olofsson et al. (2016); Agrawal et al. (2014); Olofsson and Ahlén (2018). In the dB-domain, the G-distribution is

$$p_1(y\,|\,m) = \frac{m^m}{\mu\Gamma(m)}\, \mathrm{e}^{m\frac{y-\bar{y}}{\mu}}\, \mathrm{e}^{-m\,\mathrm{e}^{\frac{y-\bar{y}}{\mu}}}, \tag{2}$$

where $\Gamma(\cdot)$ denotes the gamma function, $\mu = 10/\ln 10$, $\bar{y}$ is the corresponding mean power in dBm and $m \geq 1$ is the Nakagami-m fading parameter. The LN-distribution in the dB-domain transforms to the normal (N) distribution

$$p_0(y\,|\,\sigma) = \frac{1}{\sqrt{2\pi}\sigma}\, \mathrm{e}^{\frac{-y^2}{2\sigma^2}} \tag{3}$$

with zero mean and standard deviation $\sigma > 0$.

Without loss of generality we assume $\bar{y} = 0$ in (2). For simplicity we will consider the case when the standard deviation σ is unknown, and the parameter m is fixed. We consider $\sigma \in (0, 5]$ (Region-of-Interest), see Olofsson and Ahlén (2018). The results for unknown m or unknown m and σ can be obtained similarly.

Consider a set of continuous measurements $\mathcal{Y} = [y_1^{\mathrm{c}}, \ldots, y_K^{\mathrm{c}}]$ and a likelihood function $\mathcal{P}(\sigma\,|\,\mathcal{Y}) = \prod_{i=1}^{K} p(y_i^{\mathrm{c}}\,|\,\sigma)$. We assume the observations y_i^c to be independent and identically distributed. We use ML estimation to estimate the unknown parameter σ, i.e., $\hat{\sigma} = \arg\max_\sigma \mathcal{P}(\sigma\,|\,\mathcal{Y}) = \arg\max_\sigma \ln \mathcal{P}(\sigma\,|\,\mathcal{Y})$. The amount of information that the observations carry about the unknown parameter can be measured by the Fisher information $\mathcal{I}_\sigma = \mathrm{E}\left\{[\Psi(\sigma\,|\,\mathcal{Y})]^2 \;|\,\sigma\right\}$, where $\Psi(\sigma\,|\,\mathcal{Y}) = \dfrac{\partial}{\partial\sigma} \ln \mathcal{P}(\sigma\,|\,\mathcal{Y})$ is the score function.

Under the assumption that the observations $\mathcal{Y}$ are independent it follows that

$$\mathcal{I}_\sigma = K i_\sigma, \tag{4}$$

where K is the number of observations and

$$i_\sigma = \mathrm{E}\left\{[\psi(y\,|\,\sigma)]^2 \;|\,\sigma\right\} = \int_{-\infty}^{\infty} [\psi(y\,|\,\sigma)]^2\, p(y\,|\,\sigma) dy, \tag{5}$$

is the Fisher information number, $\psi(y\,|\,\sigma) = \frac{\partial p(y\,|\,\sigma)}{\partial\sigma} \frac{1}{p(y\,|\,\sigma)}$.

Assume that the set of received measurements is obtained from a coarse quantizer, i.e., instead of y_i^c we will consider the points y_i, $i = 1, \ldots, K$, that are obtained in bins. The k-th bin interval is given by $I_k = (z_k, z_{k+1})$ $k = 1, \ldots, N$, where $z_k = \bar{y}_k - \frac{\Delta}{2}$, $\Delta > 0$ is the width of the k-th bin, and N is the total number of bins. Here $\bar{y}_k$ is the middle point of the k-th bin, and we assume that $y_i = \bar{y}_k$ for $y_i \in I_k$, $i = 1, \ldots, K$. Then the distribution corresponding to $p(y\,|\,\sigma)$ is defined by

$$q(y\,|\,\sigma) = \frac{1}{\Delta} \int_{I_k} p(\zeta|\,\sigma) d\zeta, \quad y \in I_k. \tag{6}$$

The problem is to characterize the quality of estimation based on quantized observations by investigating the loss of Fisher information due to quantization, i.e.,

$$d_\sigma^\Delta = i_\sigma - i_\sigma^\Delta, \tag{7}$$

where

$$i_\sigma^\Delta = \mathrm{E}\left\{\left[\psi^\Delta(y\,|\,\sigma)\right]^2 \,\middle|\, \sigma\right\}, \tag{8}$$

and $\psi^\Delta(y\,|\,\sigma) = \frac{\partial q(y\,|\,\sigma)}{\partial\sigma} \frac{1}{q(y\,|\,\sigma)}$.

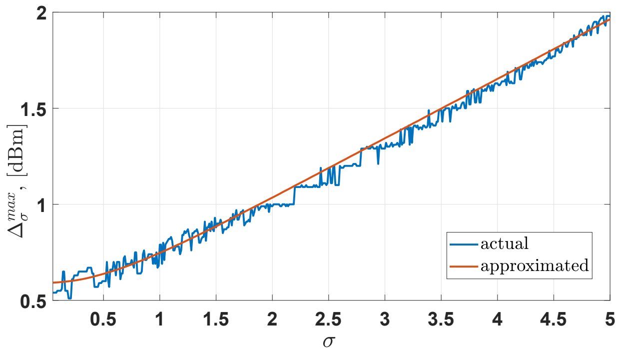

Fig. 1. The maximum value of Δ for which the relative loss of i_σ does not exceed 2%.

3. UNIFORM QUANTIZATION

Let us start with the case when all bins have the same width, i.e., uniform quantization. When the bin intervals are sufficiently small, the asymptotic behavior of the loss d_σ^Δ can be characterized by the following theorem (see (Seifullaev et al., 2019, Them 1).)

Theorem 1. The loss of Fisher information due to quantization $d_\sigma^\Delta \geq 0$ can be assessed as

$$\lim_{\Delta\to 0} \frac{d_\sigma^\Delta}{\Delta^2} = \frac{1}{12} \mathrm{E}\left\{\left[\frac{\partial\psi(y\,|\,\theta)}{\partial y}\right]^2 \,\middle|\, \sigma\right\}. \tag{9}$$

Denote $\tilde{d}_\sigma^\Delta = \frac{\Delta^2}{12} \mathrm{E}\left\{\left[\frac{\partial\psi(y\,|\,\sigma)}{\partial y}\right]^2 \,\middle|\, \sigma\right\}$. From Theorem 1 we have that $d_\sigma^\Delta \sim \tilde{d}_\sigma^\Delta$ for small Δ.

The detailed investigation on how uniform quantization influences on the loss of Fisher information can be found in Seifullaev et al. (2019), concluding that more accurate measurements are needed for small values of σ, i.e., in case of peaky distributions. E.g. Fig. 1 illustrates the maximum values of Δ guaranteeing that the relative loss $r_\sigma^\Delta = \frac{d_\sigma^\Delta}{i_\sigma}$ does not exceed 2%, i.e., $\Delta_\sigma^{\max} = \max\left\{\Delta \,\middle|\, r_\sigma^\Delta \leq 0.02\right\}$. Fig. 1 also shows that in case of 2% relative loss the approximation of the loss obtained with Theorem 1 is sufficiently good, which greatly reduces the computational load when characterizing a quantization interval which guarantees an appropriate quality of the ML estimates.

Consider the compound distribution $p(y\,|\,\sigma)$ depicted in Fig. 2. We see that for small σ the compound distribution $p(y\,|\,\sigma)$ has small support and a steep bell shaped curvature, when for large σ the support is wide and the distribution is more symmetric and flat.

Denote by $S(\sigma)$ the width of 99.9%-support of the distribution $p(y\,|\,\sigma)$, i.e., $S(\sigma) = y_{\mathrm{right}} - y_{\mathrm{left}}$, where $\int_{-\infty}^{y_{\mathrm{left}}} p(y\,|\,\sigma) dy = \int_{y_{\mathrm{right}}}^{\infty} p(y\,|\,\sigma) dy = 0.0005$. Then for every fixed σ^* the number of quantized bins covering the 99.9%-support of $p(y\,|\,\sigma^*)$ and provide at most 2% relative loss of Fisher information can be computed as $\lceil S(\sigma^*)/\Delta_{\sigma^*}^{\max}\rceil$, where $\lceil z\rceil$ denotes the least integer greater than z. We see that with the growth of σ the length of the bin intervals increases approximately by factor four (see Fig. 1). At the same time, the width of support also grows by factor two, see Fig. 3. Hence, the number of bins decreases raughly by half when σ increases from 0 to 5.

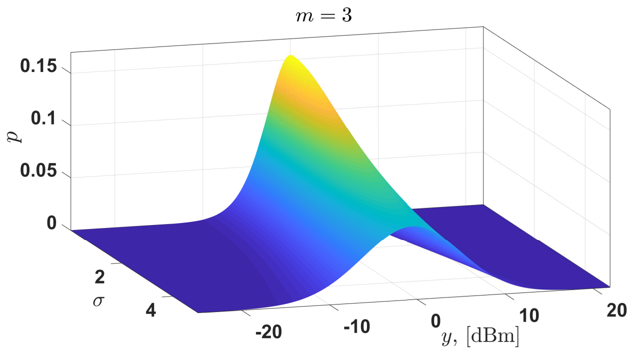

Fig. 2. The compound distribution $p(y \mid \sigma)$.

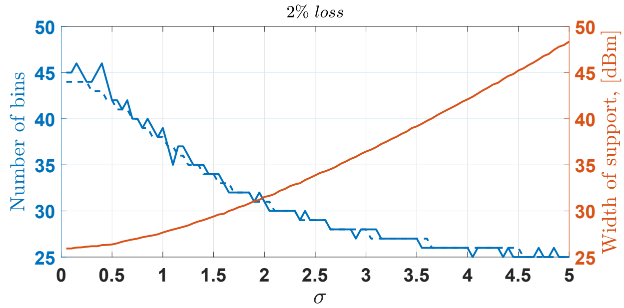

Fig. 3. Blue lines: the number of quantized bins providing at most 2% relative loss (actual values (solid line) and approximated ones (dashed line)). Red line: $S(\sigma)$.

4. VARIABLE QUANTIZATION

In a case of uniform quantization we have the same width Δ for all bins, i.e., the same measurement resolution for all intervals of y. At the same time, not all intervals affect the Fisher information equally. For instance, we can intuitively expect that for the boundary bins, where the distribution has a very low slope, we do not need the same accuracy while the intervals with a steeper slope of p require a more fine-grained quantizer. Thus, we expect that using a variable quantizer may result in a lower loss of Fisher information keeping the same number of bins. For variable quantization the k-th bin interval $\tilde{I}_k$ is denoted by $\tilde{I}_k = (\tilde{z}_k, \tilde{z}_{k+1})$, $k = 1, \ldots, N$, where the end points, $\tilde{z}_1$ and $\tilde{z}_{N+1}$, are fixed, i.e., $\tilde{z}_1 = z_1 = y_{\text{left}}$ and $\tilde{z}_{N+1} = z_{N+1} = y_{\text{right}}$. The goal of the variable quantizer is to place the points $\tilde{y}_k = \frac{\tilde{z}_k + \tilde{z}_{k+1}}{2}$ by selecting $\tilde{z}_k$, $k = 2, \ldots, N$ in an optimal way. Below, we consider and compare three different variable quantizers.

4.1 Random perturbations

The first method to obtain the sequence $\{\tilde{z}_k\}_{k=1}^{N+1}$ from $\{z_k\}_{k=1}^{N+1}$ providing a lower loss of Fisher information compared to uniform quantization is random perturbations, where the points $\tilde{z}_k$ are varied iteratively starting with $\tilde{z}_k = z_k$ for $k = 2, \ldots, N$, followed by randomly perturbing one or several points yielding z'_k and calculating the loss of information for z'_k. At each iteration, if the loss using z'_k is lower than the previous one, then the new locations z'_k are accepted, i.e, set $\tilde{z}_k = z'_k$. Otherwise, the new locations are rejected, i.e., using $\tilde{z}_k$. The advantage of this method is the simple design of the quantizer and its easy implementation.

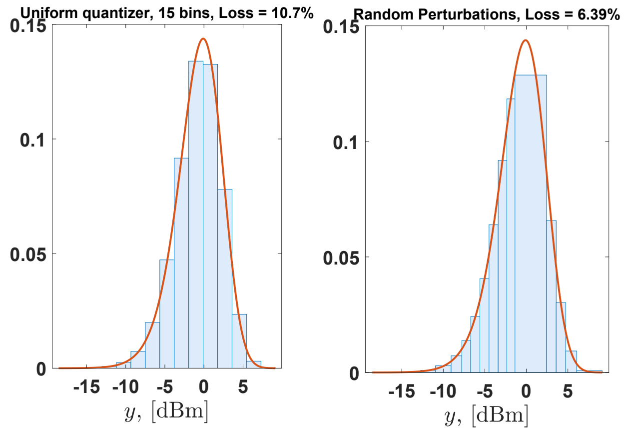

Fig. 4. Variable quantization obtained from the uniform quantizer using random perturbations for $\sigma = 1$.

The result of such a random approach is shown in Fig. 4, where 700 iterations for $\sigma = 1$ and $N = 15$ bins where used. The loss of Fisher information decreased from 10.7% to 6.39% and is expected to decrease further for more iterations. But since the main disadvantage of this method is its very high computational complexity as we need to compute the actual loss at each iteration, this method is acceptable to get some intuition on possible suitable quantization intervals but is inappropriate in practice.

4.2 Contraction mapping

To reduce the computational complexity we want to find a non-iterative method that allows deriving an approximation of the loss as well. We assume now that the variable quantizer is implemented by applying a differentiable mapping $F : \mathbb{R} \to \mathbb{R}$, i.e., $\tilde{z}_k = F(z_k)$, for $k = 2, \ldots, N$. The loss of Fisher information due to variable quantization can then be approximated, see Theorem 3 in Poor (1988):

Proposition 1. Assume that the function $G(y) = \frac{dF(y)}{dy}$ is invertible. Then the loss of Fisher information due to variable quantization can be assessed as

$$\lim_{\Delta \to 0} \frac{d_\sigma^\Delta}{\Delta^2} = \frac{1}{12} \mathrm{E} \left\{ \left[G^{-1}(y) \frac{\partial \psi(y \mid \theta)}{\partial y} \right]^2 \,\middle|\, \sigma \right\}, \tag{10}$$

where G^{-1} is the inverse of G.

The main challenge is to design the mapping F that gives an advantage in the loss compared to uniform quantization. Since, intuitively, one needs more accurate measurements for the intervals where the distribution has a high slope, we define F as a two-sided contraction mapping with the fixed points corresponding to the largest slope of p:

$$F(y) = \begin{cases} y_*^{\mathrm{l}} \, \mathrm{e}^{-a^{\mathrm{l}}(y - y_*^{\mathrm{l}})}, & y \leqslant y_{\mathrm{m}}, \\ y_*^{\mathrm{r}} \, \mathrm{e}^{\, a^{\mathrm{r}}(y - y_*^{\mathrm{r}})}, & y > y_{\mathrm{m}}, \end{cases} \tag{11}$$

where y_{m} is a mode (peak value) of the distribution p, $y_*^{\mathrm{l}} = \max_{y_{\text{left}} < y \leqslant \bar{y}_{\mathrm{m}}} \left| \frac{dp(y \mid \sigma)}{dy} \right|$, $y_*^{\mathrm{r}} = \max_{\bar{y}_{\mathrm{m}} < y < y_{\text{right}}} \left| \frac{dp(y \mid \sigma)}{dy} \right|$, $a^{\mathrm{l}} = \frac{\log y_{\text{left}} - \log y_*^{\mathrm{l}}}{y_*^{\mathrm{l}} - y_{\text{left}}}$ and $a^{\mathrm{r}} = \frac{\log y_{\text{right}} - \log y_*^{\mathrm{r}}}{y_{\text{right}} - y_*^{\mathrm{r}}}$.

Proposition 2. The mapping (11) is a contraction with:

- the points y_*^{l} and y_*^{r} are the fixed points, i.e., $F(y_*^{\mathrm{l}}) = y_*^{\mathrm{l}}$ and $F(y_*^{\mathrm{r}}) = y_*^{\mathrm{r}}$,

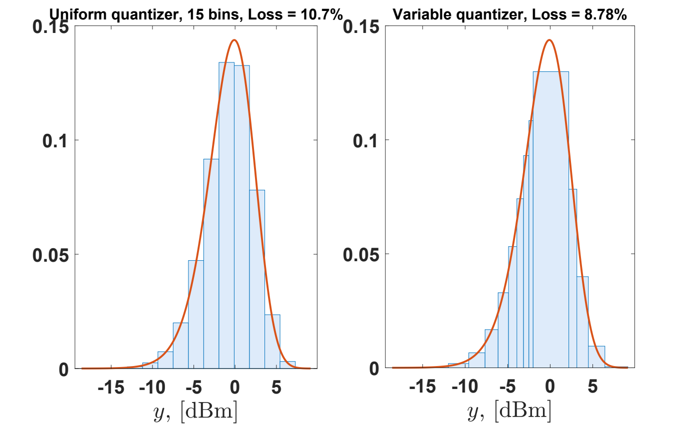

Fig. 5. Variable quantization obtained from the uniform one using the contraction mapping for $\sigma = 1$.

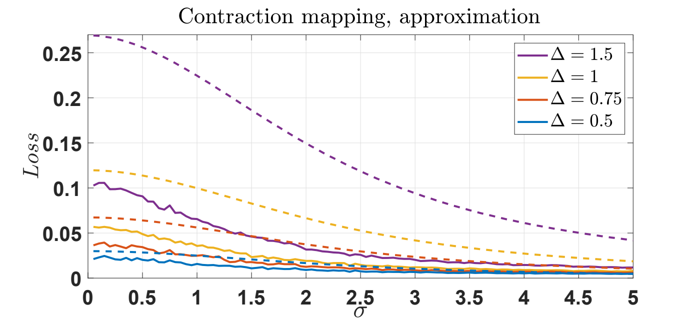

Fig. 6. The loss of Fisher information due to variable quantization based on contraction mapping: the actual loss (solid lines) and approximated as in (10) (dashed).

- for all $y \in (y_{\text{left}}, y_{\text{right}})$ one has

$$\begin{cases} y \ < F(y) < y_*^{\text{l}}, & \text{for } y_{\text{left}} < y < y_*^{\text{l}} \\ y_*^{\text{l}} < F(y) < y, & \text{for } \ \ y_*^{\text{l}} < y \leqslant y_{\text{m}} \\ y \ < F(y) < y_*^{\text{r}}, & \text{for } \ \ y_{\text{m}} < y < y_*^{\text{r}} \\ y_*^{\text{r}} < F(y) < y, & \text{for } \ \ y_*^{\text{r}} < y < y_{\text{right}} \end{cases}$$

Proof. By direct calculations.

Note that (11) has a gap at $\bar{y}_{\text{m}}$. But Proposition 1 is still satisfied, where the expectation integral in (10) is a sum of two integrals on the intervals $(y_{\text{left}}, \bar{y}_{\text{m}})$ and $(\bar{y}_{\text{m}}, y_{\text{right}})$.

Fig. 5 shows that for $\sigma = 1$ and $N = 15$ the loss of Fisher information is decreased from 10.7% to 8.78%. The obtained loss is greater than for the case of random perturbations (see Fig. 4). However, the implementation of the quantizer based on the contraction mapping is much more simple. Moreover, using approximation (10) allows to significantly reduce the computational complexity. See Fig. 6 for the actual and approximated relative losses for different values of Δ, where Δ is the step of the uniform quantizer equivalent to an average step of the variable quantizer with the same number of bins and width of support. Fig. 7 illustrates the loss due to variable quantization based on mapping (11) compared with the loss due to uniform quantization showing that variable quantization gives better results only for the large values of Δ while for $\Delta < 1$ (approximately) the uniform quantizer shows better performance since for small Δ the number of bins is large. Hence, the loss within the large middle bin becomes significant compared to the other bins. At the same time from Fig. 6 we can conclude that the approximation (10) is relatively good only for $\Delta < 0.5$. Thus, the contraction mapping cannot be considered as a generally suitable approach to design variable quantizers.

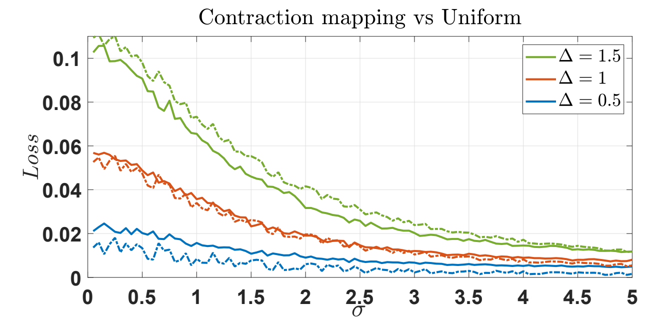

Fig. 7. The loss of Fisher information: variable quantization based on (11) (solid lines) and uniform quantization (dashed lines).

4.3 Two-sided uniform quantization

Figs. 4 and 5 reveal that the resulting variable quantizers have one wide middle bin and two large bins corresponding to the tails of the distribution p. This indicates that only a little amount of Fisher information is lost within these intervals when using quantized measurements. Let us consider the problem more formally next.

From (5) we know that the Fisher information number is the integral of the function $f(y) = [\psi(y \,|\, \sigma)]^2 p(y \,|\, \sigma)$ in a case of continuous measurements, and $f^{\Delta}(y) = [\psi^{\Delta}(y \,|\, \theta)]^2 p(y \,|\, \theta)$ in a case of quantized ones. The plot of $f(y)$ is shown in Fig. 8 (blue line), showing that the function $f(y)$ has three peaks, where the left and right peaks are different due to asymmetry of the distribution p. At the same time we know that the function $\psi^{\Delta}(y \,|\, \theta) = \psi_k = \text{const}$ for $y \in \tilde{I}_k$. Hence, the function $f^{\Delta}(y)$ is piecewise-continuous and constitutes scaled copies of the distribution $p(y \,|\, \sigma)$ within the intervals $\tilde{I}_k$ (see the red curve in Fig. 9). Since both functions, $f(y)$ and $p(y \,|\, \sigma)$, have peaks close to the peak value $\bar{y}_{\text{m}}$ of the function $f(y)$, one can find a relatively large interval $(\bar{y}_{\text{l}}, \bar{y}_{\text{r}})$ such that

$$\int_{\bar{y}_{\text{l}}}^{\bar{y}_{\text{r}}} f(y)dy \approx \int_{\bar{y}_{\text{l}}}^{\bar{y}_{\text{r}}} f^{\Delta}(y)dy. \tag{12}$$

Also we know that at the points y_{left} and y_{right} both functions, $f(y)$ and $p(y \,|\, \sigma)$, have small values. Then definitely it will be possible to find the points y_{l} and y_{r} such that the integrals $\int_{y_{\text{left}}}^{y_{\text{l}}} [f(y) - f^{\Delta}(y)]dy$ and $\int_{y_r}^{y_{\text{right}}} [f(y) - f^{\Delta}(y)]dy$ are sufficiently small.

As stated in Section 4.1, the random perturbations approach gives some intuition on suitable quantization intervals. Fig. 4 shows that the bins within the interval $(y_{\text{l}}, \bar{y}_{\text{l}})$ have approximately the same width. The same is true for the interval $(\bar{y}_{\text{r}}, y_{\text{r}})$. Hence, for simplicity, we can consider the uniform quantizers within these intervals.

Therefore, the design of the two-sided uniform quantizer can be performed in two steps (see Fig. 8). The first step

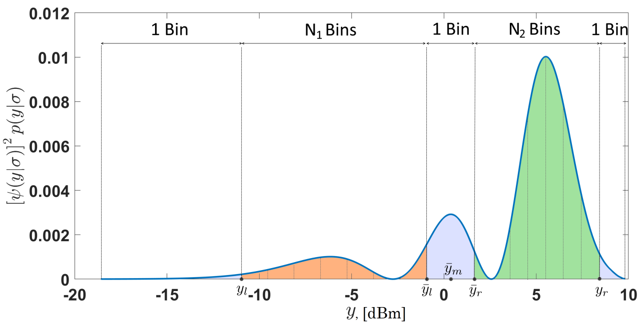

Fig. 8. Two-sided uniform quantizer design. The blue curve is the function $f(y)$ that should be integrated to obtain the Fisher information number i_σ.

is to choose three relatively large bins (blue areas): one middle bin containing the peak value $\bar{y}_{\rm m}$ of $f(y)$, and two bins corresponding to the tails of p. The second step is to specify the number of bins for two intervals of uniform quantization (orange and green areas). Since the peaks are asymmetric a smaller step for the larger peak (green area) should be taken, whereas a larger step can be selected for the wider peak (orange area). As a result, we obtain the following variable quantizer:

- $\tilde{I}_1 = (y_{\rm left}, y_{\rm l})$ and $\tilde{I}_N = (y_{\rm r}, y_{\rm right})$,
- N_1 uniform bins $\tilde{I}_k$ for $k = 2, \ldots, k_* - 1$,
- N_2 uniform bins $\tilde{I}_k$ for $k = k_* + 1, \ldots, N - 1$,
- $\tilde{I}_{k_*} = (\bar{y}_{\rm l}, \bar{y}_{\rm r})$, where $k_* = N - N_1 - N_2 - 2$.

The main challenge of such quantizer is to find the parameters $(y_{\rm l}, \bar{y}_{\rm l}, \bar{y}_{\rm r}, y_{\rm r}, N_1, N_2)$ giving the smallest value of the loss. One way is to simply find the parameters iteratively. However, such an approach requires a lot of computational resources and might hence offer little benefit compared to the iterative method proposed in Section 4.1. In this paper, we therefore propose the following parameter settings:

- The parameters $y_{\rm l}$ and $y_{\rm r}$ can be chosen such that
$$\int_{y_{\rm left}}^{y_{\rm l}} f(y)dy = \int_{y_{\rm r}}^{y_{\rm right}} f(y)dy = \varepsilon,$$
where we choose $\varepsilon = 0.01$.
- We can see from Fig. 8 that the left and right peaks of the function $f(y)$ are asymmetric. Then the quantizer step for the interval with a higher peak, i.e., the interval $(\bar{y}_{\rm r}, y_{\rm r})$ for $\sigma = 1$, should definitely be taken smaller. However, for higher peaks, the width of support is smaller as well. For the region of interest $\sigma \in (0, 5]$, a good choice of the numbers of bins, N_1 and N_2, is $N_1 = N_2 = \frac{N-3}{2}$ for odd N, and $N_1 + 1 = N_2 = \frac{N-2}{2}$ for even N. By choosing the same numbers of bins for the left and the right intervals, $(y_{\rm l}, \bar{y}_{\rm l})$ and $(\bar{y}_{\rm r}, y_{\rm r})$, we automatically obtain a smaller quantizer step for the interval with a higher peak of $f(y)$.
- The choice of $\bar{y}_{\rm l}$ and $\bar{y}_{\rm r}$ depends on the number of bins. Indeed, if N is too small then the error between the integrals in (12) is relatively small compared to the loss within the intervals $(y_{\rm l}, \bar{y}_{\rm l})$ and $(\bar{y}_{\rm r}, y_{\rm r})$. Then we can take the bin $(\bar{y}_{\rm l}, \bar{y}_{\rm r})$ relatively wide which decreases the step of both uniform quantizers. If N is too large, then the loss due to the uniform quantizers is small, and hence, the error between the integrals in (12) becomes significant. We should then make the middle bin narrower. In this paper we propose the following choice of $\bar{y}_{\rm l}$ and $\bar{y}_{\rm r}$:
$$\bar{y}_{\rm l} = \bar{y}_{\rm m} - d, \quad \bar{y}_{\rm r} = \bar{y}_{\rm m} + d,$$
where d is inversely proportional to N.

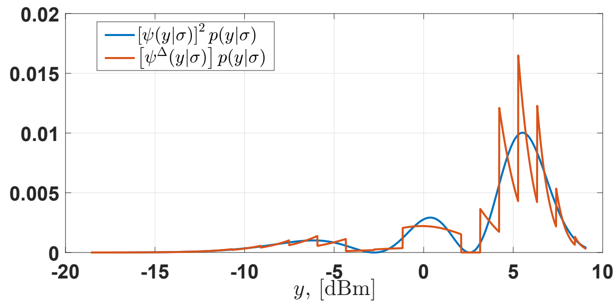

Fig. 9. The plots of $f(y)$ and $f^\Delta(y)$ for $\sigma = 1$ and $N = 15$ bins.

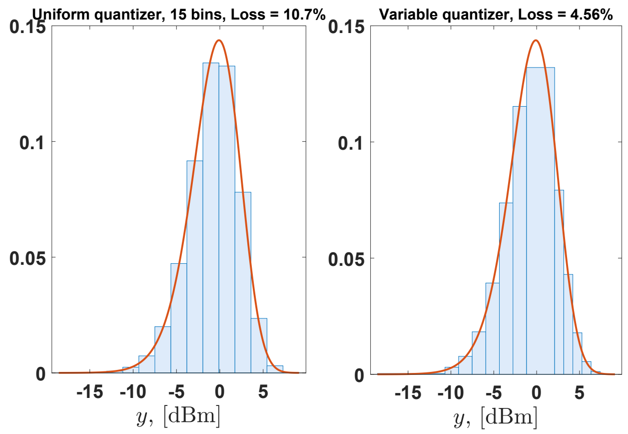

Fig. 10. Two-sided uniform quantization (right figure) compared to uniform quantization (left figure) for $\sigma = 1$, and $N = 15$.

The plots of $f(y)$ and $f^\Delta(y)$ for the parameters designed as above are shown in Fig. 9. In Fig. 10 we can see that for $\sigma = 1$ and $N = 15$ bins the loss of Fisher information is reduced from 10.7% (uniform quantization) to 4.56% (two-sided uniform quantization).

For the two-sided uniform quantizer, the loss can be approximated based on Theorem 1. Introduce the following notations: $L_1 = (y_{\rm left}, y_{\rm l}) \cup (\bar{y}_{\rm l}, \bar{y}_{\rm r}) \cup (y_{\rm r}, y_{\rm right})$ and $L_2 = (y_{\rm l}, \bar{y}_{\rm l}) \cup (\bar{y}_{\rm r}, y_r)$. Since we have three relatively large bins within L_1, for these intervals the actual loss should be computed, while for the intervals L_2 where the uniform quantizers are used the loss can be preferably approximated.

Theorem 2. The loss of Fisher information due to two-sided uniform quantization can be approximated as
$$d_\sigma^\Delta \approx \int_{L_1} \left([\psi(y\,|\,\sigma)]^2 - \left[\psi^\Delta(y\,|\,\sigma)\right]^2 \right) p(y\,|\,\sigma)dy + \tilde{d}_\sigma^\Delta,$$
where $\lim_{\Delta\to 0} \dfrac{\tilde{d}_\sigma^\Delta}{\Delta^2} = \dfrac{1}{12} \displaystyle\int_{L_2} \left[\frac{\partial \psi(y\,|\,\theta)}{\partial y} \right]^2 p(y\,|\,\sigma)dy.$

Proof. Omitted for brevity.

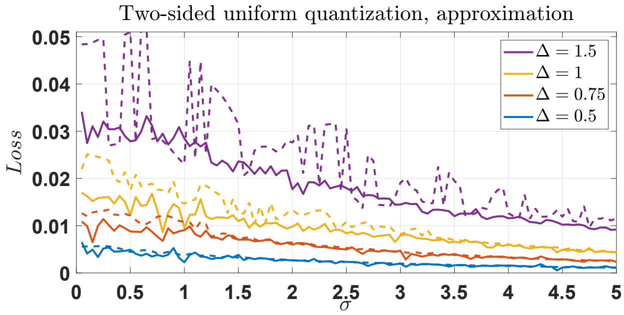

Fig. 11. The loss of Fisher information due to two-sided uniform quantization: the actual loss (solid lines) and the approximated loss (dashed lines).

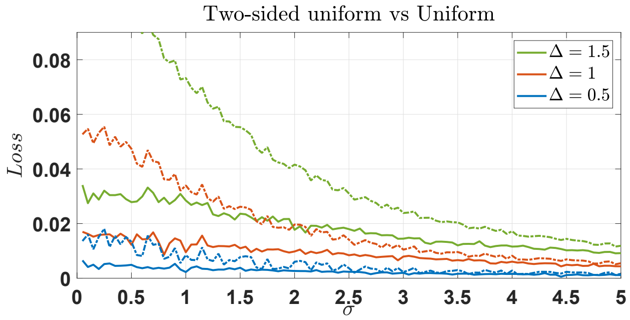

Fig. 12. The loss of Fisher information: two-sided uniform quantization (solid lines) and uniform quantization (dashed lines).

The actual and approximated relative losses are shown in Fig. 11 for different values of the average bin width Δ. We can see that the approximation is sufficiently accurate for $\Delta < 1$. Fig. 12 illustrates the loss due to two-sided uniform quantization compared with the loss due to uniform quantization. We see that the two-sided uniform quantizer shows better performance.

4.4 Comparison of results

The obtained variable quantizers are compared in Fig. 13 which illustrates the minimal number of quantized bins providing at most 2% loss of Fisher information.

For the random quantizer we have used 700 iterations and reduced the number of bins compared to the uniform quantizer. We definitely could obtain even better results by increasing the number of iterations, however, as was noticed above, the computational complexity of the random method is extremely high. In Fig. 13 we can also see that for the quantizer based on contraction mapping the approximation is inaccurate for large σ since with the growth of σ the number of bins decreases, leading to larger values of Δ (see also Fig. 6). For the two-sided uniform quantizer, the approximation is sufficiently good for all σ. Another advantage that provides an easier implementation of the proposed two-sided uniform quantizer is that for all values σ we obtain a similar number of bins providing at most 2% loss. Therefore, we can conclude that using two-sided uniform quantization allows us to significantly reduce the minimal number of quantized bins which guarantees an appropriate quality of estimates.

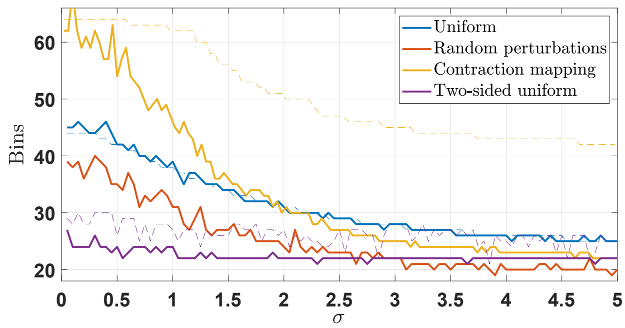

Fig. 13. The number of quantized bins providing at most 2% loss of Fisher information: actual values (solid lines) and approximated values (dashed lines).

5. CONCLUSIONS

The problem of choosing quantization intervals for the compound distribution to provide an appropriate loss of Fisher information is considered. Using variable quantization instead of uniform one can significantly reduce the information loss. Three different approaches to design the variable quantizer are proposed, where the two-sided uniform quantization method showed the best performance.

REFERENCES

Agrawal, P., Ahlén, A., Olofsson, T., and Gidlund, M. (2014). Long term channel characterization for energy efficient transmission in industrial environments. *IEEE Trans. Commun.*, 62(8), 3004–3014.

Croonenbroeck, R., Underberg, L., Wulf, A., and Kays, R. (2017). Measurements for the development of an enhanced model for wireless channels in industrial environments. In *Proc. 13th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications*. Savoy HotelRome; Italy.

Dogandzic, A. and Jin, J. (2004). Maximum likelihood estimation of statistical properties of composite gamma-lognormal fading channels. *IEEE Trans. Signal Process.*, 52(10), 2940–2945.

Gungor, V. and Hancke, G. (2009). Industrial wireless sensor networks: challenges, design principles, and technical approaches. *IEEE Trans. Ind. Electron.*, 56(10), 4258–4265.

Olofsson, T. and Ahlén, A. (2018). Computing probability density functions of compound distributions: A comparative investigation. In *Proc. IEEE International Conference on Communications*. Kansas City, MO, USA.

Olofsson, T., Ahlén, A., and Gidlund, M. (2016). Modeling of the fading statistics of wireless sensor network channels in industrial environments. *IEEE Trans. Signal Process.*, 64(12), 3021–3034.

Poor, V. (1988). Fine quantization in signal detection and estimation. *IEEE Trans. Inf. Theory*, 34, 960–972.

Seifullaev, R., Knorn, S., and Ahlén, A. (2019). The effect of uniform quantization on parameter estimation of compound distributions. *IEEE Control Systems Letters*, 3(4), 1032–1037.

Shao, H. and Beaulieu, N. (2010). An analytical method for calculating the bit error rate performance of rake reception in UWB multipath fading channels. *IEEE Trans. Commun.*, 58(4), 1112–1120.