Automatica 49 (2013) 554–560

Contents lists available at SciVerse ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

# Generating dithering noise for maximum likelihood estimation from quantized data[☆]

Fredrik Gustafsson[a,1], Rickard Karlsson[b]

[a] Department of Elec. Eng., Linköping University, Sweden
[b] Nira Dynamics AB, Linköping, Sweden

ARTICLE INFO

*Article history:*
Received 6 July 2010
Received in revised form
20 December 2011
Accepted 17 September 2012
Available online 14 December 2012

*Keywords:*
Maximum likelihood
Estimation
Quantization

ABSTRACT

The Quantization Theorem I (QT I) implies that the likelihood function can be reconstructed from quantized sensor observations, given that appropriate dithering noise is added before quantization. We present constructive algorithms to generate such dithering noise. The application to *maximum likelihood estimation* (MLE) is studied in particular. In short, dithering has the same role for amplitude quantization as an anti-alias filter has for sampling, in that it enables perfect reconstruction of the dithered but unquantized signal's likelihood function. Without dithering, the likelihood function suffers from a kind of aliasing expressed as a counterpart to Poisson's summation formula which makes the exact MLE intractable to compute. With dithering, it is demonstrated that standard MLE algorithms can be re-used on a smoothed likelihood function of the original signal, and statistically efficiency is obtained. The implication of dithering to the Cramér–Rao Lower Bound (CRLB) is studied, and illustrative examples are provided.


## 1. Introduction

Quantization was a well studied topic some decades ago, Oppenheim and Schafer (1975), when the underlying reason was the finite precision in electronics and micro-processors. Today, new reasons have appeared that motivate a revisit of the area. One is that cheap low-quality sensors have appeared on the market which opens up many new application areas for embedded algorithms, where the sensor resolution is much worse than the micro-processor resolution. Some sensors are naturally quantized such as radar range, vision (pixel quantization), cogged wheels to measure angular speeds, etc. Sensor networks is one hot research topic where this work fits in. The conclusion is that quite advanced pre-processing at the sensor node is possible to mitigate the effects of sensor quantization.

This contribution regards the sensor readings as the only instance where quantization effects are important. All subsequent computations are done with floating point precision, or in fixed-point arithmetics with adaptive scaling of all numbers, which means that internal quantization effects can be neglected. As one example in this direction, Aysal, Coates, and Rabbat (2008) shows that dithering helps a network to reach a consensus for estimating a signal mean when quantized samples from different nodes are communicated. The paper (Wornell, Papadopoulos, & Oppenheim, 2001) motivates the parameter estimation problem in sensor networks further, and develops some feedback strategies to the sensors in the same spirit as in Agüero, Goodwin, and Yuz (2007). Statistical treatment of quantization effects was developed in Widrow (1956) and Widrow, Kollar, and Liu (1996), and the newer statistical analysis as surveyed in Wannamaker, Lipshitz, Vanderkooy, and Wright (2000) and Widrow and Kollar (2008). They show that quantization adds two kind of errors to the measurement, the first one is a direct effect that can be modeled as *additive uniform noise* (AUN), and the other one is an intrinsic alias like uncertainty, where fast variations in the *probability density function* (PDF) of the measurement noise are folded to low "frequencies" in the CF domain (see Section 3.2). It is also known, Wannamaker et al. (2000), that adding $r$ uniformly distributed samples as dithering noise enables reconstruction of all the first $r$ moments of the unquantized signal, when uniform quantization is considered, where applications to system identification is studied in Gustafsson and Karlsson (2009a,b). This is known as Quantization Theorem II (QT II). However, reconstruction of moments is easier than reconstruction of the complete amplitude distribution and in particular the likelihood function. The aliasing in the amplitude distribution can be avoided by adding proper dithering noise, and this enables that

[☆] The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Brett Ninness under the direction of Editor Torsten Söderström. This work was supported by the Swedish Research Council through the center of excellence CADICS and project grant *"Fundamental issues in sensor fusion"*.

*E-mail addresses:* fredrik@isy.liu.se (F. Gustafsson), rickard.karlsson@niradynamics.se (R. Karlsson).

[1] Tel.: +46 13282706; fax: +46 13282622.

0005-1098/$ – see front matter © 2012 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2012.11.028

the likelihood function can be reconstructed from quantized sensor observations (QT I). Since this noise plays the same role in quantization and PDF reconstruction as an anti-alias filter does for sampling and reconstruction, such dithering noise will be referred to as *band-limited* (BL) noise. The theoretical requirements for a BL noise were given in the 1950s in Widrow (1956) and later revisited in Sripad and Snyder (1977).

A first contribution is to describe constructive ways to generate such noise. Two methods are provided for generating BL dithering noise, and one concrete algorithm based on accept–reject sampling. For most cases adding dithering noise inevitably destroys information. This is in perfect analogy with low pass filtering used to avoid frequency aliasing. Information is thus lost, but it is at least not misinterpreted leading to an estimation bias. For the case of parameter estimation in 1-bit quantized samples, Dabeer and Karnik (2006) and Dabeer and Masry (2008) design the dithering noise by MSE optimization, trading off a bias decrease to the variance increase due to dithering. Adding a suitably designed dithering noise should simplify the derivation of estimators. A second contribution is to utilize this new class of dithering noise to maximum likelihood estimation problems. The dithering noise implies that the discrete PDF of the quantized signal can be computed by convolving the original likelihood with first the dithering noise PDF and then the PDF of a uniform distribution. That is, quantization can under some conditions be regarded as adding uniformly distributed noise without approximation regarding moment calculations and the likelihood can be reconstructed by low pass filtering the discrete PDF.

The paper is organized as follows: Section 2 formalizes the problem definition and provides some simple motivating examples. Section 3 summarizes the most important concepts from statistical quantization theory needed for the derivations. Section 4 presents the new methods to generate band-limited noise. In Section 5, ML-estimation for different quantization cases are presented, and a performance bound is calculated. Section 6 concludes the paper.

## 2. Motivation and problem formulation

### 2.1. Signal model

The signal model in this contribution includes a dithering noise $d_k$ that is added to a stochastic signal $z_k$ before it is quantized to $q_k$,

$$y_k = z_k + d_k, \tag{1a}$$

$$q_k = \mathcal{Q}_m y_k. \tag{1b}$$

Here, $\mathcal{Q}_m(\cdot)$ denotes the quantization operator

$$\mathcal{Q}_m y_k = \begin{cases} -m\Delta + \dfrac{\Delta}{2}, & y_k < -(m-1)\Delta, \\ \Delta \left\lfloor \dfrac{y_k}{\Delta} \right\rfloor + \dfrac{\Delta}{2}, & -(m-1)\Delta \le y_k < (m-1)\Delta, \\ m\Delta - \dfrac{\Delta}{2}, & y_k \ge (m-1)\Delta, \end{cases}$$

where the floor operator $\lfloor x \rfloor$ denotes the largest integer that is smaller than or equal to $x$, $\mathcal{Q}_1(y)$ defines the binary quantizer and where $\mathcal{Q}_\infty(y)$ is defined as the unsaturated quantization function.

### 2.2. Maximum likelihood estimation

Consider now an estimation problem, where the signal $z_k(\theta)$, $k = 1, 2, \ldots, N$ depends on a vector $\theta$ of unknown parameters and the problem is to make inference of $\theta$ from the quantized but possibly dithered observations

$$q_k = \mathcal{Q}_m(z_k(\theta) + d_k), \quad k = 1, 2, \ldots, N. \tag{2}$$

Without quantization, there are many standard methods for solving the *maximum likelihood* (ML) method, where the *expectation maximization* (EM) algorithm is possibly the most common one. For quantization two approaches will be discussed in the following: applying the ML method directly or using dithering to simplify the problem.

#### 2.2.1. General MLE for quantization

Parameter estimation using the ML method for quantization has been discussed in for instance Balogh, Kollar, and Sarhegyi (2010), Gustafsson and Karlsson (2009b) and Karlsson (2005). A simple but instructive example is the unknown signal mean model $z_k(\theta) = x + e_k$, where $e_k$ denotes white measurement noise and the parameter vector $\theta = (x, \sigma_e^2)$ contains the mean $x$ and possible also the variance $\sigma_e^2 = \mathbb{E}(e_k^2)$. This example can be extended to linear models $z_k = Hx + e_k$, a nonlinear model $z_k = h(x) + e_k$, nonlinear filtering $z_k = h(x_k) + e_k$ (where the state $x_k$ varies over time according to a dynamic model) and system identification, where $\theta$ contains parameters in a dynamic model. The ML method is applicable to all these cases, where the ML *estimate* (MLE) is defined as

$$\hat{\theta}^{\text{ML}} = \arg\max_\theta \begin{cases} \prod_{k=1}^{N} p_{z|\theta}(z_k), & \text{if unquantized}, \\ \prod_{k=1}^{N} p_{q|\theta}(q_k), & \text{if quantized}. \end{cases} \tag{3}$$

The MLE is statistically efficient, defined so that the bias as well as the variance tend to zero as the number of samples $N$ increases. Computing the MLE using the likelihood $p_{q|\theta}(q_k)$ directly is mathematically intractable in most cases and also suffers from the curse of dimensionality. Example 1 illustrates one principal problem.

**Example 1.** Consider the problem of estimating $\theta$ in the Gaussian distribution $z_k(\theta) \sim \mathcal{N}(\theta, \theta^2)$. The MLE can be shown to be

$$\hat{\theta}^{\text{ML}} = -\frac{\bar{z}}{2} + \sqrt{\frac{\bar{z}^2}{4} + \overline{z^2}}, \tag{4}$$

where $\bar{z} = \frac{1}{N}\sum_{k=1}^{N} z_k$ and $\overline{z^2} = \frac{1}{N}\sum_{k=1}^{N} z_k^2$. For quantized observations $q_k(\theta) = \mathcal{Q}_\infty(z_k(\theta))$ and $\Delta = 1$, each observation is mapped to an integer $i$, and the likelihood for each quantized observation $q_k = i\Delta + \Delta/2$ is given by

$$p_{i|\theta}(q_k) = \int_{i\Delta}^{(i+1)\Delta} \frac{1}{\sqrt{2\pi}\theta} e^{-\frac{(q_k - \theta - z)^2}{2\theta^2}} dz. \tag{5}$$

The likelihood function can be differentiated in $\theta$, however in general it can be quite a lot of work doing so and an exhaustive search over the integer space $i = 0, \pm 1, \pm 2, \ldots$ is required to compute the MLE.

#### 2.2.2. MLE using dithering

As indicated in (2) dithering can simplify estimation when dealing with quantized signals. In Section 4 a special class of dithering noise that allows perfect reconstruction of the dithered but unquantized signal is introduced, so the MLE can be computed using the following likelihood

$$p_{q|\theta}(q_k) = p_{z|\theta} \star p_d \star p_u(q_k), \tag{6}$$

where $\star$ denotes convolution, $p_d$ is the PDF of the dithering noise, and $p_u$ denotes the PDF of a uniform distribution. That is, the function to be minimized is the original likelihood function in (3), smoothed with the PDF of dithering noise and a uniform distribution. Example 2 illustrates this.