PARTICLE FILTERING FOR QUANTIZED SENSOR INFORMATION

Rickard Karlsson, Fredrik Gustafsson

Department of Electrical Engineering
Linöping University, Linköping, Sweden,
E-mail: {rickard,fredrik}@isy.liu.se

ABSTRACT

The implication of quantized sensor information on filtering problems is studied. The Cramér-Rao lower bound (CRLB) is derived for estimation and filtering on quantized data. A particle filter (PF) algorithm that approximates the optimal nonlinear filter is provided, and numerical experiments show that the PF attains the CRLB, while second-order optimal Kalman filter (KF) approaches can perform quite bad.

1. INTRODUCTION

Quantization was a well studied topic in *digital signal processing* (DSP) decades ago [1], when the underlying reason was the finite computation precision in micro-processors. Today, new reasons have appeared that motivate a revisit of the area:

- Cheap low-quality sensors have appeared on the market and in many consumer products, this opens up for many new application areas for embedded DSP algorithms where the sensor resolution is much less than the micro-processor resolution.
- The increased use of distributed sensors in *sensor networks* with limited bandwidth.
- Some sensors are naturally quantized such as radar range, vision devices, cogged wheels to measure angular speeds *etc.* With increased performance requirements, quantization effects become important to analyze.
- The renewed interest in nonlinear filtering with the advent of the particle filter [2] enables a tool to take quantization effects into account in the filter design.

In these cases, one can regard the sensor readings as quantized. All sub-sequent 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. Thus, the quantization effects studied in this paper differ from the ones studied decades ago [1].

2. PARAMETER ESTIMATION AND INFORMATION BOUNDS

2.1 Information Bounds

In the sequel, the analysis is heavily based on expressions involving gradients of scalar functions or vector valued functions:

$$\nabla_x g^T(x) = \begin{pmatrix} \frac{\partial g_1}{\partial x_1} & \cdots & \frac{\partial g_m}{\partial x_1} \\ \vdots & & \vdots \\ \frac{\partial g_1}{\partial x_n} & \cdots & \frac{\partial g_m}{\partial x_n} \end{pmatrix},\ g: \mathbb{R}^n \mapsto \mathbb{R}^m. \tag{1}$$

The Laplacian for $g(x,y)$ with $x \in \mathbb{R}^n, y \in \mathbb{R}^m$ is defined as

$$\mathbf{\Delta}_y^x g(x,y) = \nabla_y (\nabla_x g(x,y))^T,\ g: \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}. \tag{2}$$

For an unbiased estimator, $\mathbb{E}(\hat{x}) = x$, the *Cramér-Rao lower bound* (CRLB), [3–5], is given by

$$\operatorname{Cov}(x-\hat{x}) = \mathbb{E}\left((x-\hat{x})(x-\hat{x})^T\right) \succeq J^{-1}(x), \tag{3a}$$

$$J(x) = \mathbb{E}\left(-\mathbf{\Delta}_x^x \log p(y|x)\right), \tag{3b}$$

where $J(x)$ denotes the *Fisher information matrix* (FIM) for the measurement y regarding the stochastic variable x. Also note that an equivalent representation of the information, [4], is

$$\mathbb{E}\left(-\mathbf{\Delta}_x^x \log p(y|x)\right) = \mathbb{E}\left(\nabla_x \log p(y|x)(\nabla_x \log p(y|x))^T\right), \tag{4}$$

Particularly, a Gaussian likelihood $p(y|x)$, with measurement covariance R, gives

$$J(x) = H^T(x)R^{-1}H(x), \text{ where } H^T(x) = \nabla_x h^T(x). \tag{5}$$

2.2 Quantization

Consider now the problem of estimating x from the quantized measurements $y = \mathcal{Q}_m(x+e)$. Explicit expressions for the information for Gaussian noise are derived in the sequel. In this paper the quantization function is restricted to the case of uniform amplitude quantization. In principle, it is implemented as the *midriser* quantizer, as described in [6]. If not saturated it is given as

$$\mathcal{Q}_m(z) = \Delta \left\lfloor \frac{z}{\Delta} \right\rfloor + \frac{\Delta}{2}. \tag{6}$$

Here, $\mathcal{Q}_m(\cdot)$ denotes the nonlinear quantization mapping with m-levels. The $\lfloor\cdot\rfloor$ operator rounds downwards to the nearest integer. To keep a unified notation with the sign quantization $\mathcal{Q}_1(z) = \operatorname{sign}(z)$, the midriser convention will be used, so $y \in \{-m\Delta+\frac{\Delta}{2}, \ldots, (m-1)\Delta+\frac{\Delta}{2}\}$, with $\Delta = 2^{-b}$, using b bits, $2m = 2^b$ levels and $2^b - 1$ thresholds. The sign quantization corresponds to $b = 1$, $m = 1$ and $\Delta = 2$ in this notation.

2.3 The Uniform Additive Approximation

One simple but approximative way to analyze attainable performance of estimators using quantized measurements is based on the assumption of approximating quantization with additive independent uniform noise d_i, [7],

$$y_i = \mathcal{Q}_m(z_i) = \mathcal{Q}_m(h(x_i)+e_i) \approx h(x_i) + e_i + d_i. \tag{7}$$

The independence assumption is not true, but if the variance of the noise e_i is much larger than the quantization resolution ($\operatorname{Var}(e_i) = \sigma^2 \gg \Delta^2$), then this is still a reasonable assumption. Another drawback is that this approach does not include the saturation effects, so in principle $m = \infty$ is assumed. The information in one measurement is thus with this approximation

$$J^{\text{approx}}(x) = \frac{1}{\sigma^2 + \frac{\Delta^2}{12}}. \tag{8}$$

The true information depends on x and includes saturation effects.

2.4 Exact Information After Quantization

2.4.1 Sign Quantizer

In this section the Fisher information for the sign quantizer is derived.

Theorem 1 *Consider the sign quantizer*

$$y = \mathcal{Q}_1(x+e) = \text{sign}(x+e), \quad e \in \mathcal{N}(0, \sigma^2). \tag{9}$$

The Fisher information is

$$J_1(x) = \frac{e^{-\frac{x^2}{\sigma^2}}}{2\pi\sigma^2} \frac{1}{(1-\varrho(-x/\sigma))\varrho(-x/\sigma)}, \tag{10}$$

where $\varrho(x) \triangleq \text{Prob}(X < x)$ *denotes the Gaussian distribution function.*

Proof: see [8].

2.4.2 Multi-Level Quantization

The sign quantizer can be generalized to the multi-level quantization case.

Theorem 2 *Consider the multi-level quantizer.*

$$y = \mathcal{Q}_m(x+e), \quad e \in \mathcal{N}(0, \sigma^2). \tag{11}$$

The Fisher information is

$$\begin{aligned} J_m(x) =& \frac{\left(-\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}(\frac{-m\Delta - x}{\sigma})^2}\right)^2}{\varrho\left(\frac{-m\Delta - x}{\sigma}\right)} \\ &+ \sum_{j=-m+1}^{m-1} \frac{\left(-\frac{1}{\sqrt{2\pi}\sigma}\left(e^{-\frac{1}{2}(\frac{(j+1)\Delta - x}{\sigma})^2} - e^{-\frac{1}{2}(\frac{j\Delta}{\sigma})^2}\right)\right)^2}{\varrho\left(\frac{(j+1)\Delta - x}{\sigma}\right) - \varrho\left(\frac{j\Delta - x}{\sigma}\right)} \\ &+ \frac{\left(\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}(\frac{m\Delta - x}{\sigma})^2}\right)^2}{1 - \varrho\left(\frac{m\Delta - x}{\sigma}\right)}. \end{aligned} \tag{12}$$

Proof: see [8].

The following example illustrates how the information and thus the CRLB depends on the quantization level.

Example 1 (CRLB – Multi-level quantizer) *In Fig. 1, the Fisher information* $J_m(x)$ *is illustrated by plotting the lower bound* $J_m^{-1/2}(x)$ *on the standard deviation for different quantization levels* $\Delta = 2/m$. *Here, the midriser quantizer with additive noise,* $y = \mathcal{Q}_m(x+e)$, $e \in \mathcal{N}(0, \sigma^2)$ *is used with* $\sigma = 0.1$.

2.5 ML-based Estimation

The set of quantized measurements will be denoted $\mathbb{Y}_t = \{y_t^{(i)}\}_{i=1}^N$ and the non-quantized set $\mathbb{Z}_t = \{z_t^{(i)}\}_{i=1}^N$.

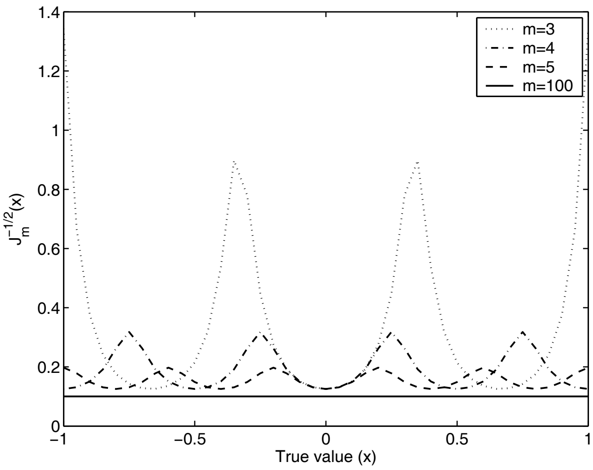

Fig 1: Fisher information used to compute the standard deviation lower bound $J_m^{-1/2}(x)$ as a function of x for different quantization levels $\Delta = 2/m$.

2.5.1 ML for Sign Quantization

Form the log-likelihood as

$$\begin{aligned} \log p(\mathbb{Y}|x) &= \log \prod_{i=1}^N p(y^{(i)}|x) = \sum_{i=1}^N \log p(y^{(i)}|x) = \\ &= N_- \log \varrho(-x/\sigma) + N_+ \log(1 - \varrho(-x/\sigma)), \end{aligned} \tag{13}$$

where N_- and N_+ denote the number of terms with $y^{(i)} = -1$ and $y^{(i)} = +1$ respectively, so that $N_- + N_+ = N$. Maximizing the expression by differentiation yields

$$\varrho\left(-x^{\text{ML}}/\sigma\right) = \frac{N_-}{N_- + N_+} = \frac{N_-}{N}. \tag{14}$$

Since the left hand side is a monotone and increasing function, the estimate, $\hat{x}^{\text{ML}}$, can be found with a line search. For more information on sign quantizers, see for instance [9], where the ML and CRLB for the frequency are calculated for a sinusoidal in noise.

2.5.2 ML for Multi-Level Quantization

The log-likelihood for multi-level quantization is

$$\log p(\mathbb{Y}_N|x) = \sum_{i=1}^N \log p(y^{(i)}|x) = \sum_{j=-m}^m N_j \log p_j(x), \tag{15}$$

where N_j is the number of occurrences of each $y^{(j)}$, so that $\sum_j N_j = N$. The ML estimate is here found numerically by searching for maximum of (15). The probability $p_j(x)$ for $j = -m+1, \ldots, m-1$ is given in [8] as

$$\begin{aligned} p_j(x) &\triangleq \text{Prob}\left(y = j\Delta + \frac{\Delta}{2}\right) \\ &= \text{Prob}(j\Delta < x + e \leq (j+1)\Delta) \\ &= \varrho\left(\frac{(j+1)\Delta - x}{\sigma}\right) - \varrho\left(\frac{j\Delta - x}{\sigma}\right). \end{aligned} \tag{16a}$$

The probability at the end points are calculated as

$$p_{-m}(x) = \varrho\left(\frac{-m\Delta - x}{\sigma}\right), \tag{16b}$$

$$p_{m-1}(x) = 1 - \varrho\left(\frac{m\Delta - x}{\sigma}\right). \tag{16c}$$

3. STATE ESTIMATION AND INFORMATION BOUNDS

For dynamic systems the following model is considered

$$x_{t+1} = f(x_t, w_t), \tag{17a}$$
$$z_t = h(x_t) + e_t, \tag{17b}$$
$$y_t = \mathcal{Q}_m(z_t). \tag{17c}$$

The Bayesian solution to the estimation problem is given by, [10],

$$p(x_{t+1}|\mathbb{Y}_t) = \int_{\mathbb{R}^n} p(x_{t+1}|x_t)p(x_t|\mathbb{Y}_t)\,dx_t, \tag{18a}$$

$$p(x_t|\mathbb{Y}_t) = \frac{p(y_t|x_t)p(x_t|\mathbb{Y}_{t-1})}{p(y_t|\mathbb{Y}_{t-1})}, \tag{18b}$$

where $p(x_{t+1}|\mathbb{Y}_t)$ is the prediction density and $p(x_t|\mathbb{Y}_t)$ the filtering density. The problem is in general not analytically solvable. There are two fundamentally different ways to approach filtering of nonlinear non-Gaussian dynamic systems:

- The *extended Kalman filter* (EKF), [11], that is the sub-optimal filter for an approximate linear Gaussian model, or the optimal *linear* filter for linear non-Gaussian systems.
- Numerical approaches, such as the *particle filter* (PF) [2, 12], that give an arbitrarily good approximation of the optimal solution to the Bayesian filtering problem.

3.1 Posterior CRLB

The theoretical posterior CRLB for a dynamic system is analyzed in [12–15]. Here, a quantized sensor using the system in (17) is considered. From [15], the posterior CRLB is

$$\text{Cov}\left(x_t - \hat{x}_{t|t}\right) = \mathbb{E}\left((x_t - \hat{x}_{t|t})(x_t - \hat{x}_{t|t})^T\right) \succeq P_{t|t}, \tag{19}$$

where $P_{t|t}$ can be retrieved from the recursion

$$P_{t+1|t+1}^{-1} = Q_t^{-1} + J_{t+1} - S_t^T\left(P_{t|t}^{-1} + V_t\right)^{-1} S_t, \tag{20}$$

where

$$V_t = \mathbb{E}\left(-\boldsymbol{\Delta}_{x_t}^{x_t} \log p(x_{t+1}|x_t)\right), \tag{21a}$$
$$S_t = \mathbb{E}\left(-\boldsymbol{\Delta}_{x_t}^{x_{t+1}} \log p(x_{t+1}|x_t)\right), \tag{21b}$$
$$Q_t^{-1} = \mathbb{E}\left(-\boldsymbol{\Delta}_{x_{t+1}}^{x_{t+1}} \log p(x_{t+1}|x_t)\right), \tag{21c}$$
$$J_t = \mathbb{E}\left(-\boldsymbol{\Delta}_{x_t}^{x_t} \log p(y_t|x_t)\right). \tag{21d}$$

Hence, the measurement quantization effects will only affect J_t, which is given by Theorem 2. For linear dynamics with additive Gaussian noise

$$x_{t+1} = F_t x_t + w_t, \tag{22}$$

the following holds

$$V_t = F_t Q_t^{-1} F_t^T, \quad S_t = -F_t^T Q_t^{-1}, \tag{23}$$

where $\text{Cov}(w_t) = Q_t$.

3.2 Kalman Filter for Measurement Quantization

Consider the following linear Gaussian model with quantized observations:

$$x_{t+1} = F_t x_t + G_t w_t, \qquad \text{Cov}(w_t) = Q_t,$$
$$z_t = H_t x_t + e_t, \qquad \text{Var}(e_t) = \sigma^2,$$
$$y_t = \mathcal{Q}_m(z_t).$$

In the sequel the quantized measurement, y_t, is treated as a scalar, but the multi-variable case is covered as long as the measurement noises $e_{t,i}$ are independent, using measurement update iterations in the *Kalman filter* (KF). A sub-optimal EKF approach is based on that quantization can be approximated as additive noise as done in (7). Hence, the likelihood can be evaluated as

$$p(z_t|x_t) = p_{\tilde{e}_t}(z_t - h(x_t)), \tag{24}$$

where $\tilde{e}_t = e_t + n_t$. Hence, for linear systems, the optimal linear filter is given by the Kalman filter with

$$R_t = \sigma_t^2 + \frac{\Delta^2}{12}\mathcal{I}, \tag{25}$$

where $\mathcal{I}$ is the identity matrix. In [16] the finite word-length for Kalman filter implementation is discussed in more detail.

3.3 Particle Filter for Measurement Quantization

The particle filter, [2, 12], here adopted to quantized measurements is given in Alg. 1. Quantization is treated formally correct using its theoretical likelihood in (26).

Alg. 1 The particle filter.

1: Set $t = 0$. For $i = 1, \ldots, N_{\text{PF}}$, initialize the particles, $x_{0|-1}^{(i)} \sim p_{x_0}(x_0)$.

2: For $i = 1, \ldots, N_{\text{PF}}$, evaluate the importance weights $\gamma_t^{(i)} = p(y_t|x_t^{(i)})$ according to the likelihood

$$p(y_t|x_t) = p_j(x_t), \tag{26}$$

where $p_j(x)$ is given in (16).

3: Resample N_{PF} particles with replacement according to,

$$\text{Prob}(x_{t|t}^{(i)} = x_{t|t-1}^{(j)}) = \tilde{\gamma}_t^{(j)},$$

where the normalized weights are given by

$$\tilde{\gamma}_t^{(i)} = \frac{\gamma_t^{(i)}}{\sum_{j=1}^{N_{\text{PF}}} \gamma_t^{(j)}}.$$

4: For $i = 1, \ldots, N_{\text{PF}}$, predict new particles according to

$$x_{t+1|t}^{(i)} \sim p(x_{t+1|t}|x_t^{(i)}).$$

5: Set $t := t + 1$ and iterate from step 2.

For hardware implementations, for instance on efficient resampling algorithms and on the complexity and performance issue for quantized particle filters, see [17, 18]. In [19, 20] the particle filter method is proposed for a sensor fusion method involving quantization, and in [21] smoothing and quantization for audio signals are considered.

Example 2 (Filtering – sign quantizer) *Consider the following scalar system with a sign quantizer*

$$x_{t+1} = F_t x_t + w_t, \qquad x_0 = 0,$$
$$y_t = \mathcal{Q}_1 (x_t + e_t),$$

where

$$F_t = 0.95, Q_t = \text{Var}(w_t) = 0.10^2, R_t = \text{Var}(e_t) = 0.58^2.$$

In Fig. 2 the root mean square error *(RMSE) for the KF and the PF are presented using* 200 *Monte Carlo simulations. The measurement noise in the KF was adjusted in the filter as described in* (25). *The PF used the correct sign quantized likelihood using* 1000 *particles. The theoretical CRLB is also given in Fig. 2, as the solution to* (20)*, which for a general case can be solved using a discrete algebraic Riccati solver. For the scalar case in this example, the covariance (P) can be derived analytically as the solution to*

$$P^2 + (QJ + 1 - F^2)/(JF^2)P - Q/(JF^2) = 0,$$

where $J_t = \frac{2}{\pi\sigma^2}$ *is given in* (10)*, assuming that the evaluation is around* $x = 0$.

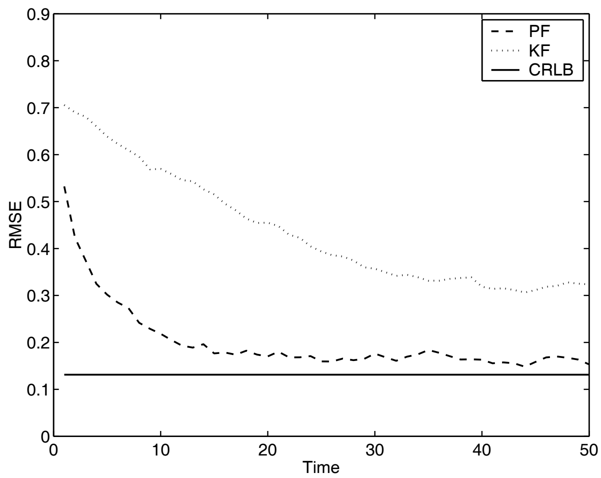

Fig 2: RMSE for the PF and KF for a linear Gaussian system with a sign quantizer in the measurement relation, compared with the CRLB limit.

4. CONCLUSIONS

The implication of quantization on Bayesian and likelihood based approaches to filtering has been studied. A detailed study on the Cramér-Rao lower bound and maximum likelihood estimation was given. Several theoretical results and examples are presented. Finally, a dedicated particle filter was given that applies to arbitrary filtering problems, where independent quantized measurements are given. The particle filter with correct likelihood for quantization attains the Cramér-Rao lower bound, and it is superior to the approximative solution from the Kalman filter.

Acknowledgment

This work was supported by VINNOVA's Center of Excellence ISIS (Information Systems for Industrial Control and Supervision), in particular the partner NIRA Dynamics.

REFERENCES

[1] A. Oppenheim and R. Schafer, *Digital Signal Processing.* Prentice-Hall, 1975.

[2] N. J. Gordon, D. J. Salmond, and A. Smith, "A novel approach to nonlinear/non-Gaussian Bayesian state estimation," in *IEE Proceedings on Radar and Signal Processing*, vol. 140, 1993, pp. 107–113.

[3] H. Cramér, *Mathematical Methods of Statistics.* Princeton, NJ: Princeton University Press, 1946.

[4] S. Kay, *Fundamentals of Statistical Signal Processing.* Prentice Hall, 1993.

[5] E. L. Lehmann, *Theory of Point Estimation.* John Wiley and Sons, 1983.

[6] S. P. Lipshitz, R. A. Wannamker, and J. Vanderkooy, "Quantization and dither: A theoretical survey," *Journal of Audio Eng. Soc*, vol. 40, no. 5, pp. 355–375, May 1992.

[7] B. Widrow, I. Kollar, and M.-C. Liu, "Statistical theory of quantization," *IEEE Trans. Instrum. Meas.*, pp. 353–361, Apr. 1996.

[8] R. Karlsson and F. Gustafsson, "Filtering and estimation for quantized sensor information," Dep. of Electrical Engineering, Tech. Rep. LiTH-ISY-R-2674, Jan. 2005, Linköping University, Sweden, Submitted to IEEE Trans. Signal Processing.

[9] A. Host-Madsen and P. Händel, "Effects of sampling and quantization on single-tone frequency estimation," *IEEE Trans. Signal Processing*, vol. 48, no. 3, pp. 650–662, 2000.

[10] A. H. Jazwinski, *Stochastic Processes and Filtering Theory*, ser. Mathematics in Science and Engineering. Academic Press, 1970, vol. 64.

[11] B. Anderson and J. B. Moore, *Optimal Filtering.* Englewood Cliffs, NJ: Prentice Hall, 1979.

[12] A. Doucet, N. de Freitas, and N. Gordon, Eds., *Sequential Monte Carlo Methods in Practice.* Springer Verlag, 2001.

[13] H. L. Van Trees, *Detection, Estimation and Modulation Theory.* New York: Wiley, 1968.

[14] P. Tichavsky, P. Muravchik, and A. Nehorai, "Posterior Cramér-Rao bounds for discrete-time nonlinear filtering," *IEEE Trans. Signal Processing*, vol. 46, no. 5, pp. 1386–1396, 1998.

[15] N. Bergman, "Recursive Bayesian estimation: Navigation and tracking applications," Dissertations No. 579, Linköping University, Linköping, Sweden, 1999.

[16] D. Williamson, "Finite wordlength design of digital Kalman filters for state estimation," *IEEE Trans. Automat. Contr.*, pp. 930–939, Oct. 1985.

[17] S. Hong, M. Bolic, and P. Djurić, "An efficient fixed-point implementation of residual resampling scheme for high-speed particle filters," *IEEE Signal Processing Lett.*, vol. 11, no. 5, pp. 482–485, 2004.

[18] M. Bolic, S. Hong, and P. Djurić, "Performance and complexity analysis of adaptive particle filtering for tracking applications," in *Conference Record of the Thirty-Sixth Asilomar Conference on Signals, Systems and Computers*, Nov. 2002.

[19] Y. Ruan, P. Willett, and A. Marrs, "Fusion of quantized measurements via particle filtering," in *IEEE Proceedings of Aerospace Conference*, vol. 4, Mar. 2003, pp. 1967–1978.

[20] ——, "Practical fusion of quantized measurements via particle filtering," in *IEE Target Tracking: Algorithms and Applications,*, Mar. 2004, pp. 13–18.

[21] W. Fong and S. Godsill, "Monte Carlo smoothing for nonlinearly distorted signals," in *Proceedings IEEE Conference on Acoustics, Speech and Signal Processing*, vol. 6, 2001, pp. 3997–4000.