

ASYMPTOTIC PERFORMANCE ANALYSIS FOR 1-BIT BAYESIAN SMOOTHING

Lin Zhang, Manuel Stein, Josef A. Nossek

Institute for Circuit Theory and Signal Processing, Technische Universität München, Germany
E-mail: lin.zhang@tum.de, manuel.stein@tum.de, josef.a.nossek@tum.de

ABSTRACT

Energy-efficient signal processing systems require estimation methods operating on data collected with low-complexity devices. Using analog-to-digital converters (ADC) with 1-bit amplitude resolution has been identified as a possible option in order to obtain low power consumption. The 1-bit performance loss, in comparison to an ideal receiver with ∞ -bit ADC, is well-established and moderate for low SNR applications ($2/\pi$ or -1.96 dB). Recently it has been shown that for parameter estimation with state-space models the 1-bit performance loss with Bayesian filtering can be significantly smaller ($\sqrt{2/\pi}$ or -0.98 dB). Here we extend the analysis to Bayesian smoothing where additional measurements are used to reconstruct the current state of the system parameter. Our results show that a 1-bit receiver performing smoothing is able to outperform an ideal ∞ -bit system carrying out filtering by the cost of an additional processing delay Δ .

Index Terms— Bayesian smoothing, state-space estimation, 1-bit signal processing.

1. INTRODUCTION

Parameter estimation finds its application in diverse problems like radar, communications and robotics, where one faces the problem of measuring the current state of the system parameters from measurements. Due to the fact that mobile processing systems are required to operate under strict power and hardware constraints, using measurement data that is obtained with coarse resolution ADCs has gained attention [1, 2, 3]. In particular, a 1-bit ADC provides a sampling device which can be operated energy-efficient at high sampling rates [4]. The non-linear operation of 1-bit hard-limiting results in a performance loss which is moderate in the low SNR regime ($2/\pi$ or -1.96 dB) [5]. As in many applications the parameters change slowly compared to the sampling rate, they can be assumed to be constant within one measurement block and to vary slightly from block to block. If the evolution of the parameters is taken into account by formulating a state-space model and performing filtering [6, 7, 8, 9, 10], it turns out that the performance loss can be reduced to -0.98 dB when using a 1-bit quantizer in the low SNR regime [11].

Pursuing this line of work, we extend the 1-bit performance analysis to Bayesian smoothing. Besides the observations of passed blocks, with smoothing Δ additional measurements are taken into account and the estimate is therefore provided with a delay of at least Δ . In practice technical requirements specify a delay budget which can be used in order to perform all required calculations. The simplicity of 1-bit ADCs and the fact that the resulting binary receive signals allow to speed up the digital hardware, enables to use a portion of the delay budget to acquire Δ additional measurements. In order to reduce the mean square error (MSE), these observations can be incorporated into the estimation algorithm by implementing smoothing.

Our discussion on the performance of 1-bit smoothing is based on Bayesian bounds for the achievable MSE. After the definition of a generic system model, we outline the basic concepts of Bayesian filtering, prediction and smoothing. A recursive Bayesian Cramér-Rao lower bound (BCRLB) is reviewed for these methods. Then we compare the performance of the 1-bit system with an ideal ∞ -bit receiver. Finally, we show the performance gain which is possible with 1-bit smoothing in comparison to ∞ -bit filtering.

2. PROBLEM STATEMENT

2.1. System Model

For the discussion we assume $k = 1, 2, \dots, K$ receive blocks with the receive signal $\mathbf{y}_k \in \mathbb{R}^N$ in the k -th block following the measurement probability density function (PDF) $p(\mathbf{y}_k|\boldsymbol{\theta}_k)$, where $\boldsymbol{\theta}_k \in \mathbb{R}^M$ is the k -th block parameter. The parameter follows a Markov model [12], described by a transition PDF $p(\boldsymbol{\theta}_k|\boldsymbol{\theta}_{k-1})$. We assign a prior PDF $p(\boldsymbol{\theta}_0)$ to the parameter vector before the initial measurement block $k = 1$.

2.2. Bayesian Filtering

The goal of filtering is to calculate the filtering PDF $p(\boldsymbol{\theta}_k|\mathbf{Y}_k)$ of the current state $\boldsymbol{\theta}_k$ given the measurement matrix

$$\mathbf{Y}_k = [\mathbf{y}_1 \quad \mathbf{y}_2 \quad \dots \quad \mathbf{y}_k] \in \mathbb{R}^{N \times k}, \quad (1)$$

containing all of the received samples up to block k . From the filtering PDF the conditional mean estimator (CME)

$$\hat{\theta}_{k|k}(\mathbf{Y}_k) = \mathbb{E}_{\theta_k|\mathbf{Y}_k}[\theta_k] \quad (2)$$

can be deduced. The CME minimizes the filtering MSE

$$\mathbf{B}_{k|k} = \mathbb{E}_{\theta_k, \mathbf{Y}_k} \left[(\theta_k - \hat{\theta}_{k|k}(\mathbf{Y}_k)) (\theta_k - \hat{\theta}_{k|k}(\mathbf{Y}_k))^T \right]. \quad (3)$$

2.3. Bayesian Prediction

Bayesian prediction is referred to as the inference of the parameter $\theta_{l|k}$ in the l -th block with $k < l$ [13]. The dynamic model of the parameter $p(\theta_k|\theta_{k-1})$ is decisive under these circumstances. The CME is calculated from $p(\theta_l|\mathbf{Y}_k)$ by

$$\hat{\theta}_{l|k}(\mathbf{Y}_k) = \mathbb{E}_{\theta_l|\mathbf{Y}_k}[\theta_l] \quad (4)$$

and minimizes the MSE

$$\mathbf{B}_{l|k} = \mathbb{E}_{\theta_l, \mathbf{Y}_k} \left[(\theta_l - \hat{\theta}_{l|k}(\mathbf{Y}_k)) (\theta_l - \hat{\theta}_{l|k}(\mathbf{Y}_k))^T \right]. \quad (5)$$

2.4. Bayesian Smoothing

Bayesian smoothing reconstructs the state $\hat{\theta}_{l|k}$ in block l by taking into account future measurements \mathbf{Y}_k up to the k -th block [13]. The expression for the CME and the MSE is the same as in (4) and (5) while for smoothing $l < k$.

3. BAYESIAN CRAMÉR-RAO BOUNDS

3.1. Filtering Bayesian Cramér-Rao Lower Bound

For Bayesian filtering the MSE can be bounded by the Bayesian Cramér-Rao lower bound (BCRLB) [14, 15]

$$\mathbf{B}_{k|k} \succeq \mathbf{J}_{k|k}^{-1}. \quad (6)$$

The required filtering Bayesian information matrix (BIM) $\mathbf{J}_{k|k}$ can be calculated efficiently in a recursive manner [16]

$$\mathbf{J}_{k|k} = \mathbf{D}_k^{22} + \mathbb{E}_{\theta_k} [\mathbf{F}(\theta_k)] - \mathbf{D}_k^{21} (\mathbf{J}_{k-1|k-1} + \mathbf{D}_k^{11})^{-1} \mathbf{D}_k^{12} \quad (7)$$

with

$$\mathbf{D}_k^{11} = \mathbb{E}_{\theta_{k-1}, \theta_k} \left[\frac{\partial \ln p(\theta_k|\theta_{k-1})}{\partial \theta_{k-1}} \left(\frac{\partial \ln p(\theta_k|\theta_{k-1})}{\partial \theta_{k-1}} \right)^T \right] \quad (8)$$

$$\begin{aligned} \mathbf{D}_k^{12} &= \mathbb{E}_{\theta_{k-1}, \theta_k} \left[\frac{\partial \ln p(\theta_k|\theta_{k-1})}{\partial \theta_{k-1}} \left(\frac{\partial \ln p(\theta_k|\theta_{k-1})}{\partial \theta_k} \right)^T \right] \\ &= \mathbf{D}_k^{21\top} \end{aligned} \quad (9)$$

$$\mathbf{D}_k^{22} = \mathbb{E}_{\theta_{k-1}, \theta_k} \left[\frac{\partial \ln p(\theta_k|\theta_{k-1})}{\partial \theta_k} \left(\frac{\partial \ln p(\theta_k|\theta_{k-1})}{\partial \theta_k} \right)^T \right] \quad (10)$$

and the Fisher information matrix having the entries [17]

$$[\mathbf{F}(\theta_k)]_{ij} = \mathbb{E}_{\mathbf{y}_k|\theta_k} \left[\frac{\partial \ln p(\mathbf{y}_k|\theta_k)}{\partial \theta_i} \frac{\partial \ln p(\mathbf{y}_k|\theta_k)}{\partial \theta_j} \right]. \quad (11)$$

The initial filtering BIM can be computed with the prior PDF

$$\mathbf{J}_{0|0} = \mathbb{E}_{\theta_0} \left[\left(\frac{\partial \ln p(\theta_0)}{\partial \theta_0} \right)^2 \right]. \quad (12)$$

3.2. Prediction Bayesian Cramér-Rao Lower Bound

The performance of prediction can be lower bounded by

$$\mathbf{B}_{l|k} \succeq \mathbf{J}_{l|k}^{-1}, \quad l > k. \quad (13)$$

The recursive calculation of the prediction BIM is obtained with a similar routine like the filtering BIM [18], i.e.,

$$\mathbf{J}_{l|k} = \mathbf{D}_l^{22} - \mathbf{D}_l^{12} (\mathbf{D}_l^{11} + \mathbf{J}_{l-1|k})^{-1} \mathbf{D}_l^{21}. \quad (14)$$

The recursion (14) starts from the filtering BIM $\mathbf{J}_{k|k}$.

3.3. Smoothing Bayesian Cramér-Rao Lower Bound

The smoothing MSE can be lower bounded by [18]

$$\mathbf{B}_{l|k} \succeq \mathbf{J}_{l|k}^{-1}, \quad l < k, \quad (15)$$

where $\mathbf{J}_{l|k}$ is the smoothing BIM with backward recursion

$$\begin{aligned} \mathbf{J}_{l|k} &= \mathbf{J}_{l|l} + \mathbf{D}_{l+1}^{11} \\ &\quad - \mathbf{D}_{l+1}^{21} (\mathbf{D}_{l+1}^{22} + \mathbf{J}_{l+1|k} - \mathbf{J}_{l+1|l})^{-1} \mathbf{D}_{l+1}^{12} \end{aligned} \quad (16)$$

which starts with the filtering BIM $\mathbf{J}_{k|k}$. Substituting the one-step version $\mathbf{J}_{k+1|k}$ of the prediction BIM (14) into (16) a compact form of the smoothing BIM can be obtained

$$\mathbf{J}_{l|k} = \mathbf{J}_{l|l} + \kappa(l|k), \quad (17)$$

where the smoothing gain $\kappa(l|k)$ follows the recursive rule

$$\begin{aligned} \kappa(l|k) &= \mathbf{D}_{l+1}^{11} - \mathbf{D}_{l+1}^{21} (\mathbf{D}_{l+1}^{22} \\ &\quad + \mathbb{E}_{\theta_{l+1}} [\mathbf{F}(\theta_{l+1})] + \kappa(l+1|k))^{-1} \mathbf{D}_{l+1}^{12} \end{aligned} \quad (18)$$

with the initial value

$$\kappa(k-1|k) = \mathbf{D}_k^{11} - \mathbf{D}_k^{21} (\mathbf{D}_k^{22} + \mathbb{E}_{\theta_k} [\mathbf{F}(\theta_k)])^{-1} \mathbf{D}_k^{12}. \quad (19)$$

Note that the smoothing BIM $\mathbf{J}_{l|k}$ is the filtering BIM $\mathbf{J}_{l|l}$ plus the gain factor $\kappa(l|k)$, which represents the additional information obtained by using smoothing instead of filtering.

4. GAUSSIAN RANDOM WALK

For a compact and simple performance analysis we consider a Gaussian random walk model. For the 1-bit receiver the information about the amplitude of the measurement is discarded.

4.1. Gaussian Random Walk

The unquantized measurement model of the random walk is

$$y_k = \theta_k + \eta_k, \quad (20)$$

where we assume that $\eta_k \sim \mathcal{N}(0, \sigma_\eta^2)$. The parameter θ_k evolves over time by following the autoregressive model

$$\theta_k = \alpha\theta_{k-1} + z_k, \quad (21)$$

where $z_k \sim \mathcal{N}(0, \sigma_z^2)$. The prior distribution $p(\theta_0)$ is assumed to be Gaussian, i.e. $\theta_0 \sim \mathcal{N}(\mu_0, \sigma_0^2)$. The required expressions for (7), (14) and (16) are

$$D_k^{11} = \frac{\alpha^2}{\sigma_z^2}, \quad D_k^{22} = \frac{1}{\sigma_z^2}, \quad D_k^{12} = D_k^{21} = -\frac{\alpha}{\sigma_z^2}. \quad (22)$$

and

$$\mathbb{E}_{\theta_k} [F(\theta_k)] = \frac{1}{\sigma_\eta^2}. \quad (23)$$

The filtering BIM can be explicitly written as

$$J_{k|k} = \frac{1}{\sigma_z^2} + \frac{1}{\sigma_\eta^2} - \left(\frac{\alpha}{\sigma_z^2}\right)^2 \left(J_{k-1|k-1} + \frac{\alpha^2}{\sigma_z^2}\right)^{-1}. \quad (24)$$

For prediction ($l > k$) the BIM is

$$J_{l|k} = \frac{1}{\sigma_z^2} - \left(\frac{\alpha}{\sigma_z^2}\right)^2 \left(\frac{\alpha^2}{\sigma_z^2} + J_{l-1|k}\right)^{-1}, \quad (25)$$

while for smoothing ($l < k$) the BIM is given by

$$J_{l|k} = J_{l|l} + \kappa(l|k) \quad (26)$$

with the smoothing gain

$$\kappa(l|k) = \frac{\alpha^2}{\sigma_z^2} - \left(\frac{\alpha}{\sigma_z^2}\right)^2 \left(\frac{1}{\sigma_z^2} + \frac{1}{\sigma_\eta^2} + \kappa(l+1|k)\right)^{-1}. \quad (27)$$

4.2. 1-bit Gaussian Random Walk

The 1-bit receiver operates on a hard-limited measurement

$$r_k = \text{sign}(\theta_k + \eta_k). \quad (28)$$

Under this circumstances the expected FIM

$$\mathbb{E}_{\theta_k} [F_q(\theta_k)] = \int_{-\infty}^{\infty} F_q(\theta_k) p_{\theta_k}(\theta_k) d\theta_k \quad (29)$$

has to be calculated with

$$F_q(\theta_k) = \frac{1}{2\pi\sigma_\eta^2} \frac{e^{-\frac{\theta_k^2}{\sigma_\eta^2}}}{\mathcal{Q}\left(\frac{\theta_k}{\sigma_\eta}\right) \mathcal{Q}\left(-\frac{\theta_k}{\sigma_\eta}\right)}, \quad (30)$$

where $\mathcal{Q}(\cdot)$ is the Q-function. For the evaluation of (29), note that the mean and variance of θ_k evolve according to [11]

$$\mathbb{E}_{\theta_k} [\theta_k] = \alpha^k \mu_0 \quad (31)$$

$$\text{Var} [\theta_k] = \alpha^{2k} \sigma_0^2 + \left(\sum_{i=1}^{k-1} \alpha^{2(k-i)}\right) \sigma_z^2, \quad (32)$$

such that

$$\sigma_{\theta, \infty}^2 = \lim_{k \rightarrow \infty} \text{Var} [\theta_k] = \frac{1}{1 - \alpha^2} \sigma_z^2. \quad (33)$$

The filtering BIM for the quantized random walk is

$$J_{k|k, q} = \frac{1}{\sigma_z^2} + \mathbb{E}_{\theta_k} [F_q(\theta_k)] - \left(\frac{\alpha}{\sigma_z^2}\right)^2 \left(J_{k-1|k-1, q} + \frac{\alpha^2}{\sigma_z^2}\right)^{-1}. \quad (34)$$

For prediction ($l > k$) the quantized BIM is

$$J_{l|k, q} = \frac{1}{\sigma_z^2} - \left(\frac{\alpha}{\sigma_z^2}\right)^2 \left(\frac{\alpha^2}{\sigma_z^2} + J_{l-1|k, q}\right)^{-1}, \quad (35)$$

whereas for smoothing ($l < k$) the quantized BIM is

$$J_{l|k, q} = J_{l|l, q} + \kappa_q(l|k) \quad (36)$$

with

$$\kappa_q(l|k) = \frac{\alpha^2}{\sigma_z^2} - \left(\frac{\alpha}{\sigma_z^2}\right)^2. \quad (37)$$

$$\left(\frac{1}{\sigma_z^2} + \mathbb{E}_{\theta_{l+1}} [F_q(\theta_{l+1})] - J_{l+1|l+1, q} + J_{l+1|k, q}\right)^{-1}. \quad (38)$$

4.3. Unquantized vs. Quantized Random Walk

For the comparison of the performance which can be obtained with the unquantized and the quantized random walk model, we consider the BCRLB of the system. Note that the BCRLB is in general not achievable, such that the performance analysis forms an approximation. After an initial transient phase the filtering algorithm reaches a steady-state, where the estimation error saturates [11] and

$$J_{k|k, \infty} = J_{k+1|k+1, \infty}, \quad (39)$$

$$J_{k|k, q} = J_{k+1|k+1, q}. \quad (40)$$

Therefore, we define the asymptotic filtering BIMs as follows

$$\tilde{J}_{k|k, \infty} = \lim_{k \rightarrow \infty} J_{k|k, \infty}, \quad (41)$$

$$\tilde{J}_{k|k, q} = \lim_{k \rightarrow \infty} J_{k|k, q}. \quad (42)$$

Also for Bayesian smoothing a steady-state is obtained for the unquantized and the quantized system, where

$$J_{k|k+\Delta, \infty} = J_{k+1|k+1+\Delta, \infty}, \quad (43)$$

$$J_{k|k+\Delta, q} = J_{k+1|k+1+\Delta, q} \quad (44)$$

with Δ being a delay being introduced by smoothing. Therefore, we define the asymptotic smoothing BIMs by

$$\tilde{J}_{k|k+\Delta,\infty} = \lim_{k \rightarrow \infty, \Delta \rightarrow \infty} J_{k|k+\Delta,\infty}, \quad (45)$$

$$\tilde{J}_{k|k+\Delta,q} = \lim_{k \rightarrow \infty, \Delta \rightarrow \infty} J_{k|k+\Delta,q}. \quad (46)$$

The 1-bit quantization loss of Bayesian smoothing in steady-state can be defined by the information ratio

$$\rho_{\text{SL}} = \frac{\tilde{J}_{k|k+\Delta,q}}{\tilde{J}_{k|k+\Delta,\infty}}. \quad (47)$$

Analogously the 1-bit quantization loss of Bayesian filtering is denoted by

$$\rho_{\text{F}} = \frac{\tilde{J}_{k|k,q}}{\tilde{J}_{k|k,\infty}}. \quad (48)$$

For the comparison of the steady-state estimation performance of 1-bit smoothing with respect to unquantized filtering we introduce the information ratio

$$\rho_{\text{S}} = \frac{\tilde{J}_{k|k+\Delta,q}}{\tilde{J}_{k|k,\infty}}. \quad (49)$$

5. RESULT

For the visualization of the expression (47) we choose $\alpha = 1 - 10^{-5}$, $\sigma_{\theta_0} = 1$ and $\sigma_{\eta} = 1$. The result depicted in Fig.1 shows that the quantization loss applying Bayesian smoothing approaches an asymptotic value of -0.98 dB in the low SNR domain. The quantization loss caused by using 1-bit smoothing instead of unquantized smoothing is the same as for the application of filtering [11]. The reduction of the performance loss from -1.96 dB without a state-space model [5] to -0.98 dB is due to the additional dynamic model (21), which is not affected by the 1-bit quantizer. Fig. 2 illustrates the behaviour of the performance ratio (48) and (49) when α approaches one. We can see that ρ_{S} is in general larger than ρ_{F} . The performance ratio ρ_{S} becomes positive in the low to medium SNR regime for α close to one. This means that the performance loss caused by quantization of the measurements can be compensated by performing smoothing. For the low SNR regime it is possible to outperform the ideal ∞ -bit filtering algorithm by 2 dB by using a 1-bit smoothing receiver. However note that this result comes with an additional delay Δ as more measurements have to be taken into account.

6. CONCLUSION

We have characterized the effect of coarse quantization onto the state-space estimation performance with smoothing. For a compact analysis the BCRLB has been used in order to approximate the MSEs with efficient processing algorithms. For




Fig. 1. Performance ratio ρ_{SL} vs. SNR




Fig. 2. Performance ratio ρ_{S} and ρ_{F} vs. SNR ($\alpha \rightarrow 1$)

the smoothing BIM a compact formulation which explicitly contains the smoothing gain factor κ has been derived. For the example of a random walk model we carried out a performance analysis for 1-bit Bayesian smoothing by comparing to an ideal system with high measurement resolution. With respect to the performance of ∞ -bit Bayesian filtering we have found that the 1-bit quantization-loss can be completely compensated by considering more measurements while performing smoothing. With the example of a random walk, we have shown that it is possible to significantly outperform an unquantized filtering system by a 1-bit smoothing receiver at the cost of an additional processing delay Δ .

7. REFERENCES

- [1] T. M. Lok and V. K. W. Wei, "Channel Estimation with Quantized Observations," *IEEE Int. Symp. Inf. Theory*, pp. 333, Aug. 1998.
- [2] A. Host-Madsen and P. Handel, "Effects of sampling and quantization on single-tone frequency estimation," *IEEE Trans. Signal Process.*, vol. 48, no. 3, pp. 650–662, Mar. 2000.
- [3] A. Mezghani, F. Antreich and J. A. Nossek, "Multiple parameter estimation with quantized channel output," *Int. ITG Workshop Smart Antennas (WSA)*, pp. 143 – 150, 2010.
- [4] R. H. Walden, "Analog-to-digital converter survey and analysis," *IEEE J. Sel. Areas Commun.*, vol. 17, no. 4, pp. 539–550, Apr. 1999.
- [5] J. H. Van Vleck and D. Middleton, "The spectrum of clipped noise," *Proc. IEEE*, vol. 54, no. 1, pp. 2–19, 1966.
- [6] K. A. Clements and R. Haddad, "Approximate estimation for systems with quantized data," *IEEE Trans. Autom. Control*, vol. 17, no. 2, pp. 235–239, 1972.
- [7] E. Sviestins and T. Wigren, "Optimal recursive state estimation with quantized measurements," *IEEE Trans. Autom. Control*, vol. 45, no. 4, pp. 762–767, 2000.
- [8] R. Karlsson and F. Gustafsson, "Particle filtering for quantized sensor information," *Europ. Sig. Proc. Conf. (EUSIPCO)*, Turkey, 2005.
- [9] Z. Duan, V.P. Jilkov and X. R. Li, "Posterior Cramer-Rao Bounds for State Estimation with Quantized Measurement," *Southeast. Symp. Sys. Theory (SSST)*, pp.376–380, March 2008.
- [10] Z. Duan, V.P. Jilkov and X. R. Li, "State estimation with quantized measurements: Approximate MMSE approach," *Inter. Conf. Inf. Fusion*, pp. 1–6, July 2008.
- [11] M. Stein, A. Kürzl, A. Mezghani and J. A. Nossek, "Asymptotic Parameter Tracking Performance with Measurement Data of 1-bit Resolution," *to appear in the IEEE Trans. Signal Process.*, preprint available: <http://arxiv.org/abs/1412.3964>, 2015.
- [12] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013.
- [13] Z. Chen, "Bayesian filtering: From Kalman filters to particle filters, and beyond," *Stat.*, vol. 182, no. 1, pp. 1–69, 2003.
- [14] B. Ristic, S. Arulampalam and N. Gordon, *Beyond the Kalman Filtery*. Artech House, 2004.
- [15] H. L. Van Trees and K. L. Bell, *Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking*. 1st ed. Piscataway, NJ: Wiley-IEEE Press, 2007.
- [16] P. Tichavsky, C. H. Muravchik and A. Nehorai, "Posterior Cramer-Rao bounds for discrete-time nonlinear filtering," *IEEE Trans. Signal Process.*, vol. 46, no. 5, pp. 1386–1396, May 1998.
- [17] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Prentice Hall, 1993.
- [18] M. Šimandl, J. Královec and P. Tichavský, "Filtering, Predictive, and Smoothing Cramér-Rao Bounds for Discrete-time Nonlinear Dynamic Systems," *Automatica*, vol. 37, no. 11, pp. 1703–1716, 2001.