Estimation with Wireless Sensor Networks:
Censoring and Quantization Perspectives

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Eric James Msechu

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Professor Georgios B. Giannakis, Advisor

June 2011

Acknowledgments

First, I thank God for health, strength, and guidance he provides in everything I do. In my graduate studies I have benefited from interactions with many individuals some of whom I will acknowledge by names.

I am grateful to Professor G. B. Giannakis who served as my academic advisor during my doctoral studies. His energy, enthusiasm, and frankness have helped me grow as a researcher. Collaboration with Prof. Giannakis has given me a lasting impression on the virtues of rigor and attention to detail in academic enquiry. I am grateful for the generous amount of time he invested in my training and for his patience with my shortcomings.

I am thankful to Profs. Stergios Roumeliotis, Tryphon Georgiou, and Mostafa Kaveh for agreeing to serve on my PhD final examination committee.

Research collaborators during my doctoral studies, including Profs. Stergios Roumeliotis and Alejandro Ribeiro, with whom I have had exciting and fruitful exchange of ideas, are gratefully acknowledged. Further gratitude is due to the graduate students and visitors to Professor Giannakis' Spinoff research group including: Prof. Antonio Garcia-Marques, Dr. Alfonso Cano-Pleite, Prof. Geert Leus, Prof. Seung-Jun Kim, Dr. Daniele Angelosante, Dr. Vassilis Kekatos, Ioannis Schizas, Juan Andres Bazerque, Shahrokh Farahmand, Hao Zhu, Nikolaos Gatsis, Dr. Emiliano Dall'Anese, Gonzalo Mateos, Brian Baingana, Morteza Mardani, Pedro-Andres Forero, Ketan Rajawat and Nasim Soltani.

Nashukuru familia yangu kwa upendo na uvumilivu wao. Mke wangu mpenzi Aisia; wanangu Aika na James; dada Helena na June; kaka Robert; mama Elfriede; wajomba, mama na baba wadogo na wakubwa; wakwe zangu mama Aichi na baba Joseph; shemeji Eliainasoe, Malaki, na Mringi. Dada Eunice na shemeji Charles waliotupokea kwa upendo Minnesota nawapa shukrani za pekee. Rafiki wa karibu na familia zao Emma; Thadeus na Pudensiana; Ndaga na Veronica; Elizabeth na John; Fortunata na Heladius. Mwisho nashukuru familia ya kikristo Minnesota kwa upendo miaka yote tuliyoshirikiana kukua kiroho. Wote mjisikie fahari kufanikiwa kwangu kumaliza masomo ya shahada hii ya juu.

Eric J. Msechu
Minneapolis, June 10, 2011

To my beloved parents, Elfriede & James Sr.

Abstract

In the last decade there has been an increase in application areas for wireless sensor networks (WSNs), which can be attributed to the advances in the enabling sensor technology. These advances include integrated circuit miniaturization and mass-production of highly-reliable hardware for sensing, processing, and data storage at a lower cost. In many emerging applications, massive amounts of data are acquired by a large number of low-cost sensing devices. The design of signal processing algorithms for these WSNs, unlike in wireless networks designed for communications, face a different set of challenges due to resource constraints sensor nodes must adhere to. These include: (i) limited on-board memory for storage; (ii) limited energy source, typically based on irreplaceable battery cells; (iii) radios with limited transmission range; and (iv) stringent data rates either due to a need to save energy or due to limited radio-frequency bandwidth allocated to sensor networks.

This work addresses distributed data-reduction at sensor nodes using a combination of measurement-censoring and measurement quantization. The WSN is envisioned for decentralized estimation of either a vector of unknown parameters in a maximum likelihood framework, or, for decentralized estimation of a random signal using Bayesian optimality criteria. Early research effort in data-reduction methods involved using a centralized computation platform directing selection of the most informative data and focusing computational and communication resources toward the selected data only. Robustness against failure of the central computation unit, as well as the need for iterative data-selection and data-gathering in some applications (e.g., real-time navigation systems), motivates a rethinking of the centralized data-selection approach. Recently, research focus has been on collaborative signal processing in sensor neighborhoods for the data-reduction step. It is in this spirit that investigation of methods for *sensor node-based* data reduction is pursued, where each sensor decides whether and what to communicate.

The scope of this dissertation encompasses distributed algorithms for the measurement-reduction step and algorithms for estimation which are amenable to either in-network or fusion center-based implementation using a WSN. Clearly-defined optimality criteria are used as the foundation for development of algorithms for data reduction and estimation. Performance analysis is provided and corroborated using simulated and real-world test cases illustrating the potential of the novel methods.

Contents

List of Figures

List of Algorithms

List of Tables

List of Acronyms

ARE	Algebraic Riccati equation
BCRLB	Bayesian Cramér-Rao lower bound
BQKF	Batch-quantized Kalman filter
CCA	Canonical correlation analysis
ccdf	Complementary cumulative distribution function
cMAP	Censored maximum a posteriori probability
cMLE	Censored maximum likelihood estimator
CRLB	Cramér-Rao lower bound
DOE	Design of optimal experiments
ECM	Error covariance matrix
FC	Fusion center
FCF	Fractional contribution factor
FIM	Fisher's information matrix
IQKF	Iteratively-quantized Kalman filter
KF	Kalman filter
LS	Least-squares
MAP	Maximum a posteriori probability
ML	Maximum likelihood
MLE	Maximum likelihood estimator
MMSE	Minimum mean-square error
MSE	Mean-square error
NEES	Normalized estimation error squared
PCA	Principal component analysis
pdf	Probability density function
qcMAP	Quantized-censored maximum a posteriori probability
qcMLE	Quantized-censored maximum likelihood estimator
QKF	Quantized Kalman filter
SNR	Signal-to-noise ratio
SoI	Sign of innovations
WSN	Wireless sensor network

Chapter 1

Introduction

The application of statistical signal processing techniques to detection and estimation tasks using wireless sensor networks (WSNs) has been a rich field of research for several decades. WSNs are used for monitoring disaster areas, for monitoring integrity of machines or infrastructure, or for a various military surveillance applications – see [6,43,47] for an overview of application areas. It is typical in these WSN applications to have massive data sizes, orders of magnitude larger than the process model dimension; see, e.g., [40]. Furthermore, the sensing nodes are typically battery-operated with limited range, limited processing power, and capable of low-to-moderate data communication rates [47]. The upside of a network built from sensors with limited resources is that, for many applications, it is economically feasible to deploy a large number of sensors as evidenced by the plethora of applications relying on networks of wireless-connected, cheap, (semi)autonomous sensors [2,6]. Broader overview of the challenges facing the design of application-specific wireless sensor networks are given in [13,68]

Earlier work in these application areas treated the ensemble of remote sensors as data acquisition devices while the data processing and decision-making entities resided at a centralized computational platform, referred to as the fusion center (FC). In the past decade, the focus in research efforts has shifted from viewing a WSN as a network for acquiring data for a centralized intelligent entity, to equipping the sensor nodes with processing capabilities. When viewed as a collaborative, distributed, intelligent network with data acquisition, data

processing, and decision-making capabilities, the WSN offers advantages that could surpass those offered by the centralized approach. Some of the application advantages that have resulted from this paradigm shift include: robustness against individual node failure; ability to reconstruct, in real-time, time-varying and spatially-distributed random fields without massive data collection at the FC; and agile target-tracking based on the most informative active sensors.

The WSNs can be classified as FC-based or ad hoc depending on communication constraints and the capabilities of the sensor nodes. To a large extent, communication range limitations of the sensor nodes (due to e.g., the range of the on-board radio or transmission power restrictions) influences how and whether sensor-to-sensor and sensor-to-FC communications are feasible. In applications where an individual sensor node's transmission footprint is small, sensor-to-sensor collaboration in neighborhoods forms the backbone for information exchange in ad hoc WSNs. For the FC-based WSNs, inter-sensor communication could be used to facilitate local signal processing in neighborhoods before data transmission to the FC. There are applications where only sensor-to-FC communication is sufficient.

This thesis focuses on collaborative signal estimation using either ad hoc or FC-based WSNs with reduced data exchanges among sensors or between sensors and the FC. Data reduction is achieved by either *censoring, quantization, or a combination of censoring and quantization* of the measurements at the sensor nodes. The censoring and quantization reduces the data transmitted by the sensors to conform with bandwidth constraints posed by the wireless link. The data reduction also leads to transmission power savings.

1.0.1 Overview of censoring and quantization for estimation

One of the challenges data reduction methods face when employed by WSNs is that the measurements are scattered across the distributed sensor nodes, and thus joint measurement processing is infeasible. That is, methods for data reduction in WSNs have to be performed either in sensor neighborhoods or separately at each sensor node. There is a large number of dimensionality-reduction methods that are applicable to joint data processing such as principal component analysis (PCA), which maps correlated large-dimensional data

to lower-dimensional uncorrelated data capturing the most variability; or canonical correlation analysis (CCA), which maps data into a form with maximal correlation between two sets of variables. Background on PCA and CCA is standard in statistics texts, see e.g., [24, 26]. Distributed versions of these methods for estimation by using collaboration among neighboring sensor nodes lead to distributed dimensionality reduction algorithms such as distributed PCA [16,17,70] and distributed CCA [72,85]. In these methods, several rounds of iterative processing involving message exchanges between neighboring sensors lead to decentralized data reduction algorithms that converge (after a finite number of iterations) approximately to the centralized processing alternatives. When it is desired to have data reduction methods with less data exchanges among sensor nodes by avoiding iterative processing and inter-sensor data exchanges, a re-thinking of the collaborative data-compression framework is needed.

This work proposes data reduction alternatives that can ambitiously be performed *autonomously* on a per sensor basis without collaboration (via measurement censoring) or collaboratively in sensor neighborhoods *without iterations* (via measurement quantization).

Fig. 1.1 illustrates interval censoring of a measurement $y^* \in \mathbb{R}$. It is seen that datum y^* is censored if $y^* \in \big(\tau_{-1}, \tau_1\big]$; otherwise, it is not censored, i.e., if $y^* \notin \big(\tau_{-1}, \tau_1\big]$, the measurement is $y = y^*$. The reduced data support at the output of the censoring $\big\{y \in \mathbb{R} : y \in (-\infty, \tau_{-1}] \cup \{0\} \cup (\tau_1, \infty]\big\}$ constitutes the data reduction in the censoring approach. Censoring as a technique for data reduction has found widespread applications in distributed detection settings [1, 5, 63]. Exploration of analogous application to distributed estimation using WSNs is studied in this thesis.

Further data reduction after the censoring is also possible using quantization of the uncensored measurements $y \neq 0$. The steps of quantization and censoring are depicted in Fig. 1.2. If only the uncensored measurements are quantized, the finite alphabet codewords are allocated only to the intervals $(-\infty, \tau_{-1}]$ and $(\tau_1, \infty]$. The partition and codewords for a generic quantizer are depicted in Fig. 1.3.

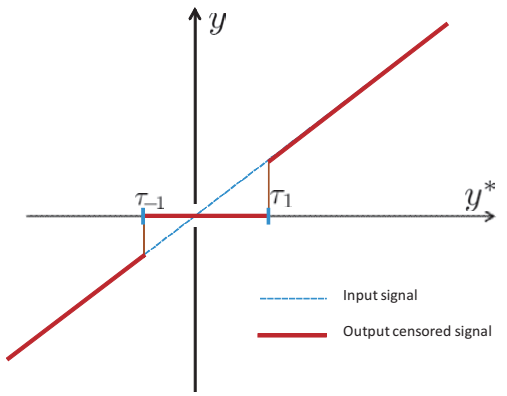

Figure 1.1: Interval censoring.

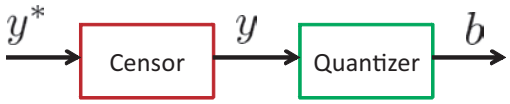

Figure 1.2: Censoring and quantization.

1.0.2 Thesis scope

The first research thrust addresses *parameter* estimation with data reduction using measurement censoring and subsequently with quantization of the uncensored measurements. Collaborative maximum likelihood (ML) estimation for a vector of deterministic parameters using censored data is studied first. Techniques for data reduction via deliberate censoring of measurements from a known noise probability density function (pdf) are explored to find a trade-off between data censoring and estimation mean-square error (MSE) performance. The work also addresses estimation from censored and quantized measurements.

In applied statistics, *parameter* estimation based on censored data is a well-studied problem [4]. Both parametric [4, 60] and nonparametric [12, 33] regression are well-documented in the econometrics and statistics literature. Asymptotic performance analysis [4, 58], as well as algorithmic developments [64] show the viability of ML estimation based on censored data. A recent overview in the context of survival analysis and its applications can be found in [44]. The work on censored regression [4, 25, 28] presumes data to be inherently censored, and focuses on centralized estimation. In order to be applicable to collaborative estimation

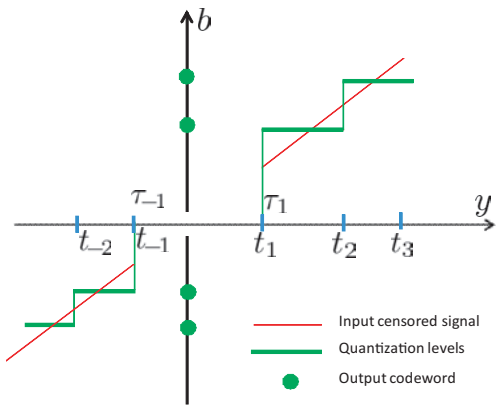

Figure 1.3: Quantization after censoring.

where measurements are distributed in a WSN, a censoring scheme that is amenable to the decentralized nature of the data is devised.

The key points addressed in this research thrust include:

- presenting measurement censoring as a method for distributed data reduction for estimation with WSNs;
- developing low complexity, low-latency algorithms for the estimation of deterministic unknown signals based on censored data;
- combining censoring with quantization to offer data reduction alternatives to purely censoring- or purely quantization-based data reduction;
- characterizing the trade-off between degree of censoring and MSE performance for different settings including variable signal-to-noise ratio (SNR), number of nodes in the WSN, and different signal models; and
- comparison of estimator performance with existing methods for data reduction and for estimation as well as Cramér-Rao lower bound (CRLB) analysis.

The second thrust focuses on modeling the unknown signal parameters as a random vector with a known pdf. Then following Bayesian estimation theory, censor-estimator and quantizer-censor-estimators analogous to those in the first thrust are developed based on

maximum a posteriori probability (MAP) optimality criterion. MAP estimation based on quantized measurements has been studied in [71,76]. However, the combination of censoring and quantization in this work offers even greater flexibility in the data reduction step. Explicit algorithmic formulations are given and the differences, especially in the censoring step, with the ML approach are highlighted.

In addition to derivations similar to those stated in the ML approach, the censored and quantized-censored MAP approach leads to:

- an analytical method for choosing the censoring thresholds which results in quantifiable average number of censored measurements; and
- an information metric based on Fisher's information matrix that quantifies the advantages of censoring when compared with selection-based approaches to data reduction.

The third thrust of the research focuses on *state* estimation from noisy measurements that are interval constrained by a quantization rule. Collaborative state estimation using WSNs for both linear and nonlinear models have been extensively studied [78,83,84]. State estimation using quantized observations is a non-linear estimation problem that can be solved using e.g., particle filters [7,14]. However, the need for low computational-complexity algorithms compel investigation of alternative approaches.

In [67, 80] state estimation based on quantized measurements is documented. In [67] binary quantization of measurements with properly designed quantization thresholds lead to Kalman filter (KF)-like algorithms whose MSE performance, for a low signal-to-noise ratio (SNR), is competitive with the clairvoyant KF based on the analog-amplitude observations [67]. Even though promising, the approach of [67] is limited to a particular 1-bit per observation quantizer. The work in this research thrust builds on and considerably broadens the scope of [67] by addressing the middle ground between estimators based on severely quantized (1-bit) data and those based on un-quantized data. The end result is quantizer-estimator structures that offer desirable trade-offs between bandwidth requirements (dictating the number of quantization bits used for inter-sensor communications), and overall tracking performance (assessed by the MSE). The work on state estimation

based on quantized measurements was published in [55–57].

The main contributions of this thrust include:

- development of batch and iterative quantization of measurements from a dynamical signal model that are amenable to be implemented using an ad hoc WSN. The batch and iterative quantization are depicted in Figs. 1.4(a) and 1.4(b);
- derivation of Kalman filter-like quantizer-estimators that trade-off data rate for MSE performance; and
- MSE performance analysis of the quantized estimators based on the state estimation error covariance matrices.

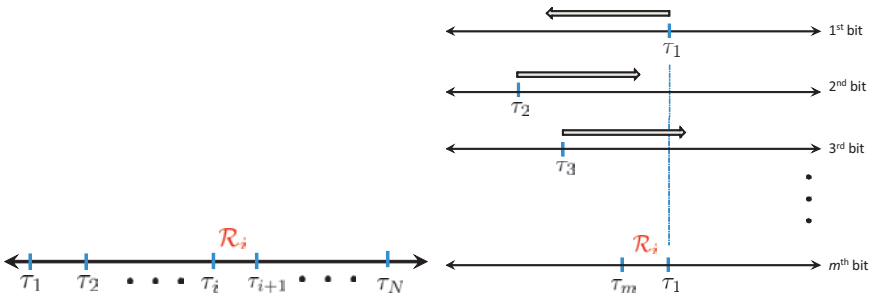

(a) (b)

Figure 1.4: (a) Batch quantization; (b) Iterative quantization

Quantizing observations to estimate a parameter of interest, is *not* the same as quantizing a signal for later reconstruction [20]. Instead of a reconstruction algorithm, the objective in quantization is finding, e.g., MSE-optimal or MAP-optimal estimators using quantized observations [21, 59, 66, 71]. For both deterministic parameter and random (parameter or state) estimation studied in this thesis, the measurements are quantized using the Lloyd-Max algorithm [48].

1.0.3 Thesis outline

In Chapter 2 estimation of an unknown deterministic signal based on noisy measurements is studied using the ML estimation criterion. Overview of ML estimation with all measurement

data, and ML estimation with data selected based on the design of optimal experiments (DOE) [32] are given. Then, distributed algorithms for maximum likelihood estimators (MLEs) using censored data and using quantized-censored data are given. Comparison of MSE performance of the different estimators is also provided, along with the associated CRLB analysis. This work is published in [53].

Chapter 3 details analogous results for random signal models based on MAP estimation - see also published work in [52]. The consequences of the different modeling assumptions (cf. MLE approach) on the censoring and quantization techniques are highlighted. Bayesian CRLB (BCRLB) analysis is also performed for both censor-estimator and quantizer-censored estimator approaches.

State estimation based on multi-bit quantized measurements is the subject of Chapter 4. Analytical results on MSE performance of the proposed quantized state estimators are derived and compared with the optimal state estimator based on unquantized measurements, i.e., Kalman filter (KF).

Simulations are used throughout Chapters 2–4 to corroborate the analytical results, and to offer insights into the differences and merits of the proposed algorithms.

The last part of the thesis in Chapter 5 gives a summary of the thesis' contributions, and discusses areas for future research.

Chapter 2

Censored Maximum Likelihood Estimation

2.1 Model and problem statement

Consider a wireless sensor network with K sensor nodes which are uniformly deployed over the area of interest. The k-th sensor is denoted by $\mathcal{S}_k$ for $k = 1, 2, \ldots, K$. The measurement $y_k^* \in \mathbb{R}$ by $\mathcal{S}_k$ is related to a vector of unknown parameters $\boldsymbol{\theta} \in \mathbb{R}^p$ by the linear regression model

$$y_k^* = \mathbf{h}_k^T \boldsymbol{\theta} + v_k, \tag{2.1}$$

where $\mathbf{h}_k \in \mathbb{R}^p$ denotes the vector of regressors and the corresponding additive noise v_k is assumed to be uncorrelated, zero mean, and Gaussian distributed with variance σ_k^2. Only the uncorrelatedness of the noise is sufficient for the following exposition, assuming that the noise variance is known. However, we will focus on independent, identically distributed (iid), zero mean, additive Gaussian noise $\mathcal{N}\left(v_k; 0, \sigma^2\right)$.

The main objective of this chapter is to estimate the parameter vector $\boldsymbol{\theta}$ at the FC based on a selected subset $\bar{K}$ of the K measurements $\{y_k^*\}_{k=1}^K$. The selection criterion is such that the selected sensors have the largest reduction in estimation mean-square error (MSE). The FC is assumed to know the regressions $\{\mathbf{h}_k\}_{k=1}^K$, the noise variance σ^2, and the

desired number of uncensored measurements $\bar{K}$. This information could be acquired in a training phase or from knowledge of the physics of the problem

If all the measurements $\{y_k^*\}_{i=1}^K$ are available at the FC, then an estimator for $\boldsymbol{\theta}$ is obtained as the solution to the least squares (LS) problem

$$\hat{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta}\in\mathbb{R}^p} \sum_{k=1}^{K} \left(y_k^* - \mathbf{h}_k^T\boldsymbol{\theta}\right)^2. \tag{2.2}$$

The LS estimator from solution of (2.2), for the Gaussian noise assumption, enjoys several other optimality properties: it is the minimum variance unbiased estimator and also the maximum likelihood estimator (MLE). A thorough discourse on linear estimation can be found in the text [34]. In this work, we will focus on maximum-likelihood (ML) estimation and study the asymptotic MSE performance of the proposed estimator for a large number of sensors K. A solution to the LS problem (2.2) is given in closed form as

$$\hat{\boldsymbol{\theta}}_{ls}(K) = \Big(\sum_{k=1}^{K} \mathbf{h}_k\mathbf{h}_k^T\Big)^{-1}\sum_{k=1}^{K} y_k\mathbf{h}_k \tag{2.3}$$

where the estimator covariance matrix is

$$\begin{aligned}\mathbf{C}_{ls}(K) =& \mathbb{E}[(\hat{\boldsymbol{\theta}}_{ls}(K) - \boldsymbol{\theta})(\hat{\boldsymbol{\theta}}_{ls}(K) - \boldsymbol{\theta})^T] \\ =& \sigma^2\Big(\sum_{k=1}^{K} \mathbf{h}_k\mathbf{h}_k^T\Big)^{-1}.\end{aligned} \tag{2.4}$$

The full-data estimator MSE, given by the trace $\mathrm{tr}\{\mathbf{C}_{ls}(K)\}$, will benchmark the performance of the proposed MLE based on censored data.

In a resource-scarce setting where sensor nodes have limited computation and communication resources, a judicious choice of a subset $\bar{K} < K$ of the sensor measurements to be transmitted from the remote sensors to the fusion center would lead to substantial savings on the sensor-to-FC communication costs without significant degradation of the resultant estimator's MSE when compared with the MSE from the full-data MLE. It is this intuition that drives effort toward sensor selection methods for decentralized estimation using WSNs.

Before delving into censoring-based estimation, the selection-based estimator of [32] will be outlined to highlight its strengths and limitations in a WSN setting.

2.1.1 Overview of data selection via optimal designs

The problem of measurement selection has a wide range of applications. Overview of the ideas behind information-based sensor selection is provided by [10, 69, 82]. The measurement selection problem is studied under various names: in applied statistics where it is known as design of optimal experiments (DOE) [36, 62, 81]; in machine learning [46]; in robotics with applications to estimation and control [27, 51]; in robotics for path-planning for the mobile robots [38, 74]; for estimation with static WSNs where it is known as sensor selection [19, 29, 32]; and for the related problem of sensor scheduling or for coverage problems [23]. An overview of the sensor selection techniques with applications to WSNs is provided in [69]. The DOE theory [62], which inspired the sensor selection for collaborative estimation in [32], finds widespread use in fields related to the experimental sciences (e.g., biostatistics) where, motivated by the need to reduce the cost of running experiments, only a few most informative measurements are selected to be run.

Posing the data selection problem for linear estimation under the design of optimal experiment (DOE) framework [62], the authors in [32] select the most informative sensors by optimizing an MSE-based or information theoretic criterion derived solely from the assumed model and without acquiring any measurements first. An overview of the DOE-based data reduction and estimation will be presented in this section.

If a subset $\bar{K}$ of the K measurements is used for the LS estimation, the resultant LS estimator $\hat{\boldsymbol{\theta}}_{ls}(\bar{K})$ will have a covariance matrix with the same structure as that in (2.4), while retaining only the $\bar{K}$ regressors corresponding to the selected sensor nodes. The problem of selecting $\bar{K}$ out of the K measurements which have leads to the lowest MSE of any selection $\bar{K}$ can be expressed as

$$\begin{aligned} \boldsymbol{s}^* = \underset{\boldsymbol{s}\in\{0,\,1\}^K}{\arg\min}\ & \operatorname{tr}\left(\sum_{k=1}^{K} s_k\,\mathbf{h}_k\mathbf{h}_k^T\right)^{-1} \\ \text{s. t. } & \sum_{k=1}^{K} s_k \le \bar{K}, \end{aligned} \tag{2.5}$$

where the integer variables $s_k \in \{0, 1\}$ are the selection indicators [32]. Only measurements with $s_k^* = 1$ are used for estimation.

The optimization problem (2.5) is non-convex due to the binary constraints on s_k. The binary integer constraints $\{0,1\}^K$ in (2.5) are typically replaced by interval constraints $[0\ 1]^K$ which lead to the following convex optimization problem.

$$\begin{aligned} \boldsymbol{s}^* = \underset{\boldsymbol{s}\in[0,\,1]^K}{\arg\min}\ \mathrm{tr}\left(\sum_{k=1}^{K} s_k\,\mathbf{h}_k\mathbf{h}_k^T\right)^{-1} \\ \text{s. t. } \sum_{k=1}^{K} s_k \leq \bar{K}. \end{aligned} \tag{2.6}$$

The solution to the convex problem 2.6 which may be fractional is then mapped into $\{0,1\}^K$ to recover a suboptimal solution to the original selection problem (2.5).

Optimization problems of the form (2.6) are known in applied statistics as A-optimal design problems [62], [32]. Efficient interior point implementation for solving (2.6) has complexity that scales as $\mathcal{O}(K^3)$ [32], where K is the dimension of the design parameter vector $\boldsymbol{s}$. It should be noted that, in the A-optimal design problem, the solution depends solely on the regressors $\{\mathbf{h}_k\}_{k=1}^K$. Thus, complete design of the data collection approach can be done before actual data measurements are acquired. One advantage of such design-before-sensing approach is that only the selected sensor nodes acquire measurements. In the nonlinear regression models the estimator covariance is usually a function of the measurements. In order to use the A-optimal design for measurement selection, approximations such as linearization are typically done. Therefore this model-driven approach to data selection does not accommodate a data-dependent covariance matrix without resorting to model approximation.

In the next section *data-driven* measurement reduction is proposed whereby all sensors acquire measurements and a subset of the measurements are censored. Only the uncensored measurements are transmitted to the FC. Data reduction using censoring is thus attractive in applications where the energy cost involved in sensing is much smaller than the cost of communicating measurements to the FC.

2.2 Data censoring

With $\mathcal{R}_k := [\tau_{1k}, \tau_{2k})$ denoting the censoring interval, consider the following censored counterpart of y_k^*

$$y_k := \begin{cases} y_k^*, & \text{if} \quad y_k^* \in \mathcal{R}_k \\ 0, & \text{otherwise.} \end{cases} \tag{2.7}$$

Only uncensored measurements, i.e., $y_k \neq 0$, are transmitted to the FC.

Censoring as a low-complexity method for data-reduction has been used in FC-based distributed detection [63], [1], where instead of censoring y_k^*, the censored data are local log-likelihood ratios. Finding analogous censoring rules of the form (2.7) for decentralized estimation is, therefore well-motivated. Data censoring can be further motivated as an approximate solution to the LS fitting problem. Suppose temporarily that $\boldsymbol{\theta}$ is known and the goal is to select at most $\bar{K}$ regressors from $\{\mathbf{h}_k\}_{k=1}^K$ for which the selected $\{\mathbf{h}_k^T\boldsymbol{\theta}\}_{k=1}^{\bar{K}}$ fit the measurements $\{y_k^*\}_{k=1}^K$ best in the LS sense. This leads to a constrained binary LS problem

$$\boldsymbol{s}^* = \underset{\boldsymbol{s}\,\in\,\{0,1\}^K}{\arg\min} \sum_{k=1}^{K} \big(y_k^* - s_k\,\mathbf{h}_k^T\boldsymbol{\theta}\big)^2 \tag{2.8a}$$

$$\text{s. to} \quad \sum_{k=1}^{K} s_k \leq \bar{K} \tag{2.8b}$$

where $\boldsymbol{s}^*$ comprises a vector of binary $\{0,1\}$ selection variables. Two difficulties to solving (2.8) are evident: (i) for large K, integer programming solvers (e.g., branch and bound [9,79]) are known to incur prohibitive computational burden; and (ii) problem (2.8) presumes knowledge of $\boldsymbol{\theta}$, which is the unknown to be estimated from the selected data.

By replacing the integer constraints $\{0,1\}$ by interval ones $[0,1]$, an approximate but convex alternative minimization problem to (2.8) can be formed for which efficient solvers are available even for large K values. To this end, consider using uniformly picked measurements from a small set $L \ll \bar{K}$ of sensors to form an initial LS estimate $\bar{\boldsymbol{\theta}} := \hat{\boldsymbol{\theta}}_{ls}(L)$ as in (2.2). Denoting the Lagrange multiplier corresponding to (2.8b) by $\lambda > 0$, a linear-

quadratic optimization problem is obtained

$$\boldsymbol{s}^*(\lambda) = \underset{\boldsymbol{s}\in[0,1]^K}{\arg\min} \frac{1}{2}\sum_{k=1}^{K}\left(y_k^* - s_k\,\mathbf{h}_k^T\bar{\boldsymbol{\theta}}\right)^2 + \lambda\left(\sum_{k=1}^{K} s_k - \bar{K}\right). \tag{2.9}$$

Problem (2.9) is decomposable into K sub-problems for which the k-th one attains its minimum at

$$s_k^*(\lambda^*) := \frac{(y_k^*\,\mathbf{h}_k^T\bar{\boldsymbol{\theta}} - \lambda^*)}{(\mathbf{h}_k^T\bar{\boldsymbol{\theta}})^2}$$

where λ^* denotes the optimum multiplier. Slicing s_k^* leads to the selection $\hat{s}_k := \mathbf{1}_{\{s_k^*>0\}} = \mathbf{1}_{\{y_k^*\,\mathbf{h}_k^T\bar{\boldsymbol{\theta}}>\lambda^*\}}$, where $\mathbf{1}_{\{\cdot\}}$ denotes the indicator function. The condition $y_k^*\mathbf{h}_k^T\bar{\boldsymbol{\theta}} > \lambda^*$ can be recast as $|y_k^*| > \lambda^*/|\mathbf{h}_k^T\bar{\boldsymbol{\theta}}|$ when y_k^* and $\mathbf{h}_k^T\bar{\boldsymbol{\theta}}$ have the same sign. The *per sensor* measurement censoring rule is then summarized as

$$(y_k, \hat{s}_k) = \begin{cases} (y_k^*, 1), & \text{if} \quad |y_k^*| > \tau_k(\lambda^*), \quad \tau_k(\lambda^*) := \frac{\lambda^*}{|\mathbf{h}_k^T\bar{\boldsymbol{\theta}}|} \\ (0,0), & \text{otherwise.} \end{cases} \tag{2.10}$$

Supposing that the optimum λ^* and an initial estimate $\bar{\boldsymbol{\theta}}$ are available at each sensor, censoring can be implemented autonomously at each sensor by using (2.10).

What remains to specify the censoring rule is to find λ^* which satisfies the inequality constraint (2.8b) at least on average. To this end, note that the number of uncensored measurements, $\sum_{k=1}^{K}\hat{s}_k$, is a random variable. Bounding $\mathbb{E}\left[\sum_{k=1}^{K}\hat{s}_k\right] = \sum_{k=1}^{K}\mathbb{E}[\hat{s}_k]$ not to exceed the desired value $\bar{K}$, yields

$$\begin{aligned}\sum_{k=1}^{K}\mathbb{E}[\hat{s}_k] =& \sum_{k=1}^{K}\Pr\left(|y_k^*| > \tau_k(\lambda)\right) \\ =& K - \sum_{k=1}^{K} Q\left(\frac{-\tau_k(\lambda) - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right) - Q\left(\frac{\tau_k(\lambda) - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right) \leq \bar{K}\end{aligned} \tag{2.11}$$

where $Q(\cdot)$ is the complementary cumulative distribution function (ccdf) of the Gaussian pdf. It can be shown that

$$f_k(\lambda) := Q\left(\frac{-\tau_k(\lambda) - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right) - Q\left(\frac{\tau_k(\lambda) - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right) \tag{2.12}$$

is a monotonically increasing function of λ. It thus follows that upon replacing $\boldsymbol{\theta}$ by $\bar{\boldsymbol{\theta}}$, a one-dimensional grid search readily yields the wanted λ^*. For notational brevity, τ_k will henceforth be used in place of $\tau_k(\lambda^*)$, and s_k in place of $\hat{s}_k$.

Remark 2.1 Regardless of how the censoring intervals are chosen, data reduction in the censoring approach does not rely on the linearity of the regression function in (4.1). Consequently, the data selection rule in (2.10) readily extends to measurements of the form $y_k^* = h_k(\boldsymbol{\theta}) + v_k$, with known nonlinear functions $h_k : \mathbb{R}^p \to \mathbb{R}$. In contrast, DOE-based selection critically relies on the linearity of the regression function during the selection step so that the data selection problem is solvable off-line at the FC prior to sensors acquiring measurements.

Remark 2.2 Even though only the uncensored measurements are transmitted to the FC, all sensors need to acquire measurements. This is in contrast with the selection-based data reduction where only the selected sensors acquire measurements. Data reduction using censoring is thus attractive in applications where the energy cost involved in sensing is much smaller than the cost of communicating measurements to the FC. From [47] energy required for different sensor operations are obtained for the Mica sensors. The values given in Table 2.1 show that energy used for write/erase operations, sensor reception, and data storage are an order of magnitude larger than that used in sensing, reading from memory, or idle listening. The Watt-hour (Wh) values in Table 2.1 are deduced from the Ampere-hours (Ah) values in [47, Table 2] by using the nominal Mica operating voltage of 3 volts. The energy expended in data acquisition per sample, i.e., sensing (1 sample), is the lowest of the six operations tabulated thus justifying censoring as an energy-efficient data reduction method.

2.3 ML estimation with censored data

In order to find the MLE with censored data, the joint pdf of $(\mathbf{y}, \boldsymbol{s})$ in (2.10) has to be found. Clearly, each measurement at $\mathcal{S}_k$ has pdf $p\left(y_k^*; \boldsymbol{\theta}\right) = \mathcal{N}\left(y_k^*;\, \mathbf{h}_k^T\boldsymbol{\theta},\, \sigma^2\right)$. Since s_k is a $0-1$ random variable, the joint censored data pdf $p(\mathbf{y}, \boldsymbol{s}; \boldsymbol{\theta})$ is given as described in the following lemma.

Lemma 2.1 *With $\boldsymbol{\theta}$ in (4.1) modeled as deterministic, the censored data are independently*

Table 2.1: Energy required for different Mica sensor operations

Operation	Nominal energy consumption (nWh)
Memory write/erase	250.00
Packet transmission	60.00
Packet reception	24.00
Radio listening (1 ms)	3.75
Memory read	3.33
Sensing (1 sample)	3.24

distributed as

$$p(\mathbf{y}, \boldsymbol{s}; \boldsymbol{\theta}) = \prod_{k=1}^{K} p(y_k, s_k; \boldsymbol{\theta})$$

where

$$p(y_k, s_k; \boldsymbol{\theta}) = \left[\mathcal{N}(y_k; \mathbf{h}_k^T\boldsymbol{\theta}, \sigma^2)\right]^{s_k} \left[\Pr(y_k = 0; \boldsymbol{\theta})\right]^{(1-s_k)} . \tag{2.13}$$

Proof: Since the data $\{y_k^*\}_{k=1}^K$ are independently distributed and the censoring is performed separately per sensor $\mathcal{S}_k$, it follows readily that the censored data $\{y_k, s_k\}_{k=1}^K$ are independently distributed as well. □

The probability of a measurement being censored is obtained by using the Gaussian ccdf $Q(\cdot)$ as

$$\begin{aligned}\Pr\left[y_k = 0; \boldsymbol{\theta}\right] &= \Pr\left[-\tau_k < y_k^* \leq \tau_k; \boldsymbol{\theta}\right] \\ &= Q\left(\frac{-\tau_k - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right) - Q\left(\frac{\tau_k - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right) .\end{aligned} \tag{2.14}$$

Fig. 2.1 illustrates an example of the pdf $p\left(y_k, s_k; \theta\right)$ for $\theta = 1.0$, $\sigma^2 = 0.8$, and $\tau_k = 2$. The reduced support of the data pdf is what leads to transmission power savings. From Lemma 2.1, the joint censored-data pdf is given by

$$\prod_{k=1}^{K} p(y_k, s_k; \boldsymbol{\theta}) = \prod_{k=1}^{K} \left[\mathcal{N}(y_k; \mathbf{h}_k^T\boldsymbol{\theta}, \sigma^2)\right]^{s_k} \left[\Pr(y_k = 0; \boldsymbol{\theta})\right]^{(1-s_k)}. \tag{2.15}$$

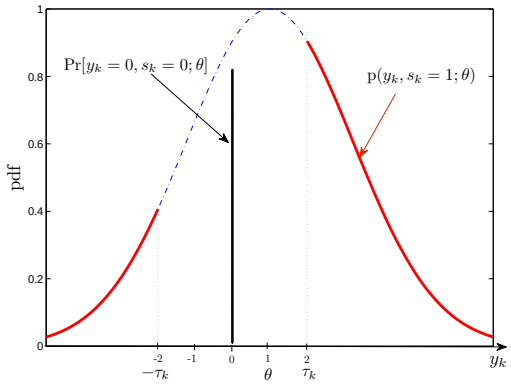

Figure 2.1: Censored data pdf.

Taking the log of the joint pdf, $\ell_K(\boldsymbol{\theta}) := \log p(\mathbf{y}, \boldsymbol{s}; \boldsymbol{\theta})$, and using (2.14) and (2.15), it follows that

$$\ell_K(\boldsymbol{\theta}) = \sum_{k=1}^{K} \Big\{ s_k \ln \mathcal{N}\left(y_k; \mathbf{h}_k^T \boldsymbol{\theta}, \sigma^2\right) + (1 - s_k) \ln \Pr(y_k = 0; \boldsymbol{\theta}) \Big\} \tag{2.16}$$

$$= \sum_{k=1}^{K} \left\{ -\frac{s_k}{2\sigma^2} \left(y_k - \mathbf{h}_k^T \boldsymbol{\theta}\right)^2 - \frac{s_k}{2} \ln(2\pi\sigma^2) + (1 - s_k) \ln \left[Q\left(\frac{-\tau_k - \mathbf{h}_k^T \boldsymbol{\theta}}{\sigma}\right) - Q\left(\frac{\tau_k - \mathbf{h}_k^T \boldsymbol{\theta}}{\sigma}\right) \right] \right\}.$$

Remark 2. If all the measurements $\{y_k^*\}_{k=1}^K$ are uncensored, $-\ell_K(\boldsymbol{\theta})$ reduces to the LS cost in (2.2). On the other hand, if all K measurements are censored, $\ell_K(\boldsymbol{\theta})$ reduces to the log-pdf of binary quantized data [65] with all identical binary symbols. Both of these extreme cases have low probability of occurring for K large.

The censored MLE (cMLE) is obtained as [cf. (2.16)]

$$\hat{\boldsymbol{\theta}}_{cMLE} = \arg \max_{\boldsymbol{\theta} \in \mathbb{R}^p} \ell_K(\boldsymbol{\theta}) \,. \tag{2.17}$$

The ensuing proposition gives conditions that guarantee a unique solution to the maximization (2.17) for a given set of measurements $(\mathbf{y}, \boldsymbol{s})$, based on the log-concavity of the Gaussian pdf and the log-concavity of the Gaussian ccdf. The proof is given in Appendix A.1.

Proposition 2.1 *If* $\sum_{k=1}^{K} \mathbf{h}_k \mathbf{h}_k^T \succ \mathbf{0}$, *the log-likelihood* $\ell_K(\boldsymbol{\theta})$ *in (2.16) is strictly concave.*

2.3.1 Newton's method for ML estimation

Before discussing the Newton algorithm for finding the MLE we summarize the gradient $\mathbf{g}_K(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}} \ell_K(\boldsymbol{\theta})$ and Hessian $\mathbf{H}_K(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T \ell_K(\boldsymbol{\theta})$ each indexed by the total number of sensors K. The gradient $\mathbf{g}_K(\boldsymbol{\theta}) \in \mathbb{R}^p$, and the Hessian $\mathbf{H}_K(\boldsymbol{\theta}) \in \mathbb{S}^{p \times p}$ are obtained as

$$\mathbf{g}_K(\boldsymbol{\theta}) = \sum_{k=1}^{K} \beta_k(\boldsymbol{\theta}) \mathbf{h}_k \tag{2.18a}$$

$$\mathbf{H}_K(\boldsymbol{\theta}) = \sum_{k=1}^{K} \gamma_k(\boldsymbol{\theta}) \mathbf{h}_k \mathbf{h}_k^T \tag{2.18b}$$

where $\beta_k(\boldsymbol{\theta})$ and $\gamma_k(\boldsymbol{\theta})$ are defined in Appendix A.2.

Starting with the initialization $\hat{\boldsymbol{\theta}}^{(0)} = \bar{\boldsymbol{\theta}}$, iterations of the damped Newton method [11] are given for $t = 0, 1, \ldots$ as

$$\hat{\boldsymbol{\theta}}^{(t+1)} = \hat{\boldsymbol{\theta}}^{(t)} + \alpha^{(t)} \left[\mathbf{H}_K(\hat{\boldsymbol{\theta}}^{(t)}) \right]^{-1} \mathbf{g}_K(\hat{\boldsymbol{\theta}}^{(t)}). \tag{2.19}$$

The damping factor $0 < \alpha^{(t)} \leq 1$ is chosen at each step t using a line search rule to ensures that the objective function is non-decreasing, i.e., $\ell_K(\hat{\boldsymbol{\theta}}^{(t+1)}) \geq \ell_K(\hat{\boldsymbol{\theta}}^{(t)})$ and that steps of sufficient size is taken. With the definitions $\Delta\hat{\boldsymbol{\theta}}^{(t)} := \left[\mathbf{H}_K(\hat{\boldsymbol{\theta}}^{(t)}) \right]^{-1} \mathbf{g}_K(\hat{\boldsymbol{\theta}}^{(t)})$ and $d^{(t)} := \mathbf{g}_K^T(\hat{\boldsymbol{\theta}}^{(t)}) \Delta\hat{\boldsymbol{\theta}}^{(t)}$, the Armijo's rule from [11] is tabulated as Algorithm 2.1.

Algorithm 2.1 Backtracking line search (Armojo's rule)

Require: $0 < \mu < 0.5$ and $0 < c < 1$

Line search:

while $\ell_K(\hat{\boldsymbol{\theta}}^{(t)} + \alpha \Delta\hat{\boldsymbol{\theta}}^{(t)}) - \ell_K(\hat{\boldsymbol{\theta}}^{(t)}) < \mu\, \alpha\, d^{(t)}$ **do**

  $\alpha \leftarrow \alpha\, c$

end while

Output: $\alpha^{(t)} = \alpha$

Since the ML cost function $\ell_K(\boldsymbol{\theta})$ is a concave function, the initialization is not critical for the convergence. The cost associated with storage and computation for the Newton's algorithm are dominated by storage and inversion of the Hessian matrix. Cholesky factorization of the $p \times p$ Hessian matrix generally scale as $\mathcal{O}(p^3)$. This attendant computation

cost is outweighed by the advantages offered by Newton's method, namely: rapid convergence, insensitivity to initialization, and guaranteed convergence for the strictly concave cost function $\ell_K(\boldsymbol{\theta})$.

2.3.2 cMLE algorithms with WSNs

Algorithms for the censoring and estimation steps amenable to WSN implementation are detailed in this section. Only transmissions between the FC and the sensor nodes are considered; that is, no inter-sensor transmissions take place in order to lower the communication costs and minimize the overhead of the overall transmission protocol. As a result, the communication cost is quantified by the number of sensor-to-FC transmissions. A round-robin, slotted-time sensor schedule is envisioned such that sensor $\mathcal{S}_k$ transmits at the k-th time slot T_k. If $\mathcal{S}_k$ censors its measurement, $\mathcal{S}_{k+1}$ transmits next. More elaborate sensor scheduling protocols are also possible; see e.g., [3]. Fig. 2.2 illustrates the sensor-FC communication set-up as well as the slotted-time transmission with measurement censoring. It is further posited that sensor-FC transmissions are practically error-free – a condition certainly satisfied using sufficiently powerful channel coding schemes.

The computation and communication steps that constitute sensor censoring are tabulated as Algorithm 2.2. It should be noted that prior to censoring, the FC calculates the threshold parameter λ and broadcasts it to all K sensors. Each $\mathcal{S}_k$ then autonomously decides whether or not to censor its acquired measurement. Unlike measurement selection outlined in Section 2.1.1, where only the selected $\bar{K}$ sensors acquire measurements, in the censoring approach all K sensors need to acquire measurements, but only $\bar{K}$ sensors transmit to the FC. Censoring would, consequently, be attractive for applications where the resource costs for sensing are outweighed by costs associated with transmission to the FC [47, Table 2].

Algorithm 2.3 summarizes a damped-Newton iteration for implementing the maximization in (2.17), details of which can be found in standard optimization texts such as [11]. The cost associated with storage and computations for Newton's iterations is dominated by the storage and inversion of $\mathbf{H}_K(\boldsymbol{\theta})$. Cholesky factorization of the $p \times p$ Hessian matrix scales

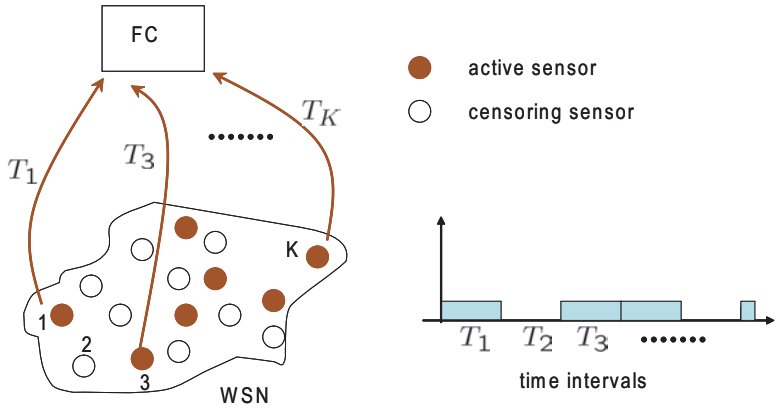

Figure 2.2: Data transmission from WSN to FC with censoring.

as $\mathcal{O}(p^3)$. By comparison, efficient interior point implementation for solving the DOE-based sensor selection scales as $\mathcal{O}(K^3)$ [32]. Since $K \gg p$, the proposed data censoring approach incurs lower computational cost compared to the sensor selection approach.

We next derive the Cramér-Rao lower bound from the pdf (2.15) for the variance of any unbiased estimator from the data given in (2.10).

2.3.3 Cramér-Rao bound and threshold selection

The Cramér-Rao lower bound (CRLB) from the likelihood function will serve two purposes: (i) give an analytical variance lower bound for the cMLE and (ii) provide an objective function for an alternative method for selection of the censoring thresholds $\boldsymbol{\tau} := [\tau_1, \ldots, \tau_K]^T$. First, the Fisher's information matrix (FIM) is obtained.

The FIM defined as $\mathbf{I}_K(\boldsymbol{\theta}) := -\mathbb{E}\left[\nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T \ell_K(\boldsymbol{\theta})\right]$ is obtained as $\mathbf{I}_K(\boldsymbol{\theta}) = \sum_{k=1}^K \bar{\gamma}_k(\boldsymbol{\theta}) \mathbf{h}_k \mathbf{h}_k^T$, where $\bar{\gamma}_k(\boldsymbol{\theta})$ is defined in Appendix A.3. It should be noted that the FIM is a function of the censoring threshold $\boldsymbol{\tau} := [\tau_{11}, \tau_{21}, \ldots, \tau_{1k}, \tau_{2k}, \ldots, \tau_{1K}, \tau_{2K}]^T$ since $\bar{\gamma}_k(\boldsymbol{\theta})$ is a function of $[\tau_{1k}, \tau_{2k}]$ [cf. (A.1c) and (A.2)]. Using the FIM, the CRLB can be found as follows.

Lemma 2.2 *For any unbiased estimator $\hat{\boldsymbol{\theta}}$ of the parameter vector $\boldsymbol{\theta}$ based on $(\mathbf{y}, \boldsymbol{s})$ distributed according to the pdf $p(\mathbf{y}, \boldsymbol{s}; \boldsymbol{\theta})$, it holds that*

$$\mathbb{E}\left[(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^T\right] - [\mathbf{I}_K(\boldsymbol{\theta})]^{-1} \succeq \mathbf{0}.$$

The expectation is taken with respect to the censored data pdf $p(y_k, s_k; \boldsymbol{\theta})$.

Algorithm 2.2 Censoring (cMLE)

Require: FC knows $\{\mathbf{h}_k\}_{k=1}^K$, $\bar{K}$; $\mathcal{S}_k$ knows $\mathbf{h}_k$, y_k^*

Initialization:

FC: Polls $L \ll K$ sensors for $\{y_\ell^*\}_{\ell=1}^L$, calculates $\bar{\boldsymbol{\theta}}$ using (2.3)

FC: Finds λ from (2.11) using a grid search

FC: Broadcasts (λ, $\bar{\boldsymbol{\theta}}$)

Sensor Censoring:

for $k = 1, 2, \ldots, K$ **do**

$\mathcal{S}_k$: Receives (λ, $\bar{\boldsymbol{\theta}}$), form threshold $\tau_k = \lambda/|\mathbf{h}_k^T\bar{\boldsymbol{\theta}}|$

$\mathcal{S}_k$: Censors y_k^* to get y_k using (2.10)

if $y_k \neq 0$ **then**

$\mathcal{S}_k$: Transmit y_k to FC

else

$\mathcal{S}_k$: Stays idle

end if

end for

The CRLB will be used in the simulations of Section 2.5 as a benchmark to assess how far the variance of the (asymptotically) unbiased cMLE is from the theoretical bound.

The FIM can be written in recursive form as follows $\mathbf{I}_K = \mathbf{I}_{K-1} + \bar{\gamma}_K \mathbf{h}_K \mathbf{h}_K^T$, where, to simplify the notation, the dependence on $\boldsymbol{\tau}$ and $\boldsymbol{\theta}$ has been dropped. Then using the matrix inversion lemma, and defining $\mathbf{C}_K := [\mathbf{I}_K]^{-1}$, a recursion for sequential computation of CRLB is obtained as follows:

$$\begin{aligned}\mathbf{C}_K &= \left(\mathbf{I}_{K-1} + \bar{\gamma}_K \mathbf{h}_K \mathbf{h}_K^T\right)^{-1} \\ &= \mathbf{C}_{K-1} - \frac{\mathbf{C}_{K-1}\mathbf{h}_K\mathbf{h}_K^T\mathbf{C}_{K-1}}{(\bar{\gamma}_K)^{-1} + \mathbf{h}_K^T\mathbf{C}_{K-1}\mathbf{h}_k},\end{aligned} \tag{2.20}$$

Information-theory based data selection strategies [46] for inference with Gaussian processes lead to sequential data selection criteria with structure resembling the recursion in (2.20). In [39] sequential data selection designs are advocated for a more general problem. In this work, we envision the interval censoring approach for estimation where the ensemble of uncensored data is collected first before estimation. The censoring approach leverages the asymptotic limit of a large sensor network which guarantees that there is enough data for

Algorithm 2.3 Estimation (cMLE)

Require: Gradient tolerance $\epsilon > 0$; maximum iterations $t_{\max}$

Data Reception:

FC: Receive $\{y_k\}_{k=1}^K$, set $y_k = 0$ if no data received at time T_k

if $y_k = 0$ **then**

FC: $s_k = 0$

else

FC: $s_k = 1$

end if

Parameter Estimation:

FC: Initialize $t \leftarrow 0$, $\hat{\boldsymbol{\theta}}^{(0)} = \bar{\boldsymbol{\theta}}$

repeat

FC: Calculate $\ell_K(\hat{\boldsymbol{\theta}}^{(t)})$, $\mathbf{g}_K(\hat{\boldsymbol{\theta}}^{(t)})$, and $\mathbf{H}_K(\hat{\boldsymbol{\theta}}^{(t)})$ from (2.16), (2.18a), and (2.18b)

FC: Find $\alpha^{(t)}$ using Armijo's rule [cf. Algorithm 2.1]

FC: $\hat{\boldsymbol{\theta}}^{(t)} \leftarrow \hat{\boldsymbol{\theta}}^{(t)} + \alpha^{(t)} \left[\mathbf{H}_K(\hat{\boldsymbol{\theta}}^{(t)})\right]^{-1} \mathbf{g}_K(\hat{\boldsymbol{\theta}}^{(t)})$

FC: $t \leftarrow t + 1$

until $\|\mathbf{g}_K(\hat{\boldsymbol{\theta}}^{(t)})\|_2 \leq \epsilon$ **or** $t = t_{\max}$

FC: Set $\hat{\boldsymbol{\theta}}_{cMLE} = \hat{\boldsymbol{\theta}}^{(t)}$

the cMLE optimization problem to be solvable. Nevertheless, the formula in (2.20) enables a system designer to know the minimum network size K that would be needed to achieve a desired estimation variance.

Finding $\boldsymbol{\tau}$ that maximizes the determinant $|\mathbf{I}(\boldsymbol{\tau}, \boldsymbol{\theta})|$ leads to an alternative method for selection of censor thresholds,

$$\hat{\boldsymbol{\tau}} = \arg \max_{\boldsymbol{\tau} \in \mathbb{R}^K,\, \boldsymbol{\tau} \succ \mathbf{0}} |\mathbf{I}_K(\boldsymbol{\tau}, \boldsymbol{\theta})| \tag{2.21}$$
$$\text{s. t. } K - \sum_{k=1}^{K} \left[Q\left(\frac{-\tau_k - \mathbf{h}_k^T \boldsymbol{\theta}}{\sigma}\right) - Q\left(\frac{\tau_k - \mathbf{h}_k^T \boldsymbol{\theta}}{\sigma}\right) \right] \leq \bar{K}.$$

The functions $\bar{\gamma}_k(\cdot, \boldsymbol{\theta})$ in (A.2) are non-concave w.r.t. τ_k which makes the optimization problem in (2.21) non-convex and difficult to solve. However, using a gradient-ascent algorithm (e.g., Newton's method) leads to a stationary point of the Lagrangian obtained from (2.21). In distributed detection with a censored log-likelihood function [63], similar

challenges for optimization of censoring thresholds are encountered. Two points are worth stressing: (i) the optimal thresholds from (2.21), with the substitution $\boldsymbol{\theta} = \bar{\boldsymbol{\theta}}$ only leads to maximization of the approximate FIM and does not guarantee that the MSE of the resultant MLE is minimized; (ii) related work in distributed estimation with quantization [65] also finds that non-identical threshold values of approximately the same size as $|\mathbf{h}_k^T \boldsymbol{\theta}|$ offer sufficiently good MSE performance.

Empirical results from simulations suggest that choice of censor thresholds that are of same order of magnitude as $|\mathbf{h}_k^T \bar{\boldsymbol{\theta}}|$ has comparable MSE as those obtained from stationary points from (2.21). Consequently, in the simulations the thresholds $\tau_k = \lambda / |\mathbf{h}_k^T \bar{\boldsymbol{\theta}}|$ will be used, where λ is obtained by using a grid search on the monotonic function $f_k(\lambda)$ in (2.12).

2.4 ML estimation with quantized-censored data

By quantizing the uncensored measurements, further transmission power and bandwidth savings can be effected since only a limited number of bits describing the uncensored measurements need to be transmitted to the FC. Incorporating the benefits of quantization to the censored data for ML estimation is the subject of this section. If y_k is uncensored, its conditional pdf is expressed using Bayes rule as

$$p\left(y_k | s_k = 1; \boldsymbol{\theta}\right) = \frac{\mathcal{N}\left(y_k; \mathbf{h}_k^T \boldsymbol{\theta}, \sigma^2\right)}{\Pr[|y_k| > \tau_k]} \mathbf{1}_{\{|y_k| \rhd \tau_k\}} \tag{2.22}$$

where $\Pr[|y_k| > \tau_k] = 1 - \Pr[y_k = 0]$. Uncensored measurements, distributed according to (2.22), form the input to the quantizer.

The quantizer design is complicated by the unknown parameter $\boldsymbol{\theta}$ in (2.22). The initial estimate $\bar{\boldsymbol{\theta}}$ will be used instead. For an MSE distortion criterion, globally optimal scalar quantization is provided by the Lloyd 1 algorithm [18, 48]. The Lloyd algorithm converges exponentially fast provided that the log-pdf is concave with a finite second moment [35]. Motivated by this strong result, Lloyd's iterative algorithm will be adapted here for determining encoder partitions. Since the quantized data are not intended to be reconstructed, but rather used for estimating $\boldsymbol{\theta}$, a uniquely-defined codebook corresponding to uniquely identifiable intervals is sufficient. The codebook is defined as $\mathcal{J} := \{\pm 1, \pm 2, \ldots, \pm \frac{L}{2}\}$,

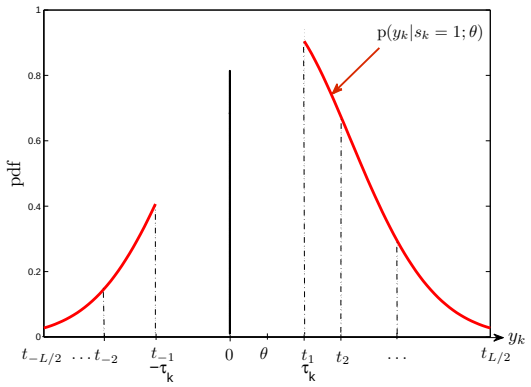

Figure 2.3: L-level quantizer thresholds for uncensored data pdf.

where $|\mathcal{J}| = L$ is the cardinality of the codebook. The set of partition thresholds $\{t_j\}_{j\in\mathcal{J}}$ is found using Lloyd's algorithm; see also [21]. The quantization step following the censoring rule (2.10) yields the finite-alphabet data

$$b_k := \begin{cases} j, & \text{if } y_k \in (t_j, t_{j+1}]\,,\ j \in \mathcal{J} \\ s_k, & \text{otherwise}\,. \end{cases} \tag{2.23}$$

Fig. 2.3 depicts generic L-level quantizer intervals for a truncated Gaussian pdf. In fact, for $j \leq -1$ the corresponding encoder partition is $(t_{j-1},\ t_j]$, whereas for $j \geq 1$ the partition is $(t_j,\ t_{j+1}]$ but the compact notation used in (2.23) will be retained for brevity. Note from Fig. 2.3 that $t_{-1} = -\tau_k$ and $t_1 = \tau_k$.

Following quantization, the digital sensor measurements are $\boldsymbol{b} := [b_1, b_2, \ldots, b_K]^T$, $\boldsymbol{s} := [s_1, s_2, \ldots, s_K]^T$, and $\boldsymbol{d}_j := [d_{j1}, d_{j2}, \ldots d_{jK}]^T$, where $d_{jk} := \mathbf{1}_{\{b_k=j\}}$ and $\sum_{j\in\mathcal{J}} d_{jk} = s_k, \forall k = 1, \ldots, K$. These censored-quantized data are distributed according to the conditional pmf

$$\Pr\left[\boldsymbol{s}, \boldsymbol{b}, \{\boldsymbol{d}_j\}_{j\in\mathcal{J}}; \boldsymbol{\theta}\right] = \prod_{k=1}^{K} \left\{ \Pr\left[s_k = 0; \boldsymbol{\theta}\right]^{(1-s_k)} \prod_{j\in\mathcal{J}} \Pr\left[b_k = j; \boldsymbol{\theta}\right]^{d_{jk}} \right\}. \tag{2.24}$$

By the same logic as in the censored-data case in Lemma 2.1, since quantization is performed separately per sensor node, a product-form pmf results. Note that for each uncensored k, only one of the $L = |\mathcal{J}|$ indicators $\{d_{jk}\}_{j\in\mathcal{J}}$ is nonzero corresponding to the codeword $j \in \mathcal{J}$

from the quantization. The block diagram in Fig. 2.4 illustrates the two stages of censoring and quantization.

Since for each sensor $\mathcal{S}_k$ the conditional pdf in (2.22) has a different approximate mean $\mathbf{h}_k^T\bar{\boldsymbol{\theta}}$ (and consequently a different truncation point $\tau_k = \lambda/|\mathbf{h}_k^T\bar{\boldsymbol{\theta}}|$), Lloyd's algorithm needs to be run separately at each sensor $\mathcal{S}_k$ for which $y_k \neq 0$.

Upon considering the log-pmf $\ell_K^Q(\boldsymbol{\theta}) := \log \Pr[\boldsymbol{s}, \boldsymbol{b}, \{\boldsymbol{d}_j\}_{j\in\mathcal{J}}; \boldsymbol{\theta}]$, it follows that [cf. (2.24)]

$$\ell_K^Q(\boldsymbol{\theta}) = \sum_{k=1}^{K}\Big\{\sum_{j\in\mathcal{J}} d_{jk}\ln\left[Q(\bar{z}_{1jk}) - Q(\bar{z}_{2jk})\right] + (1-s_k)\ln\left[Q(z_{1k}) - Q(z_{2k})\right]\Big\}$$

where $\bar{z}_{1jk}$ and $\bar{z}_{2jk}$ are specified in Appendix A.4. Function $\ell_K^Q(\boldsymbol{\theta})$ is concave following the proof of Proposition 2.1. The corresponding gradient $\mathbf{g}_K^Q(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}}\ell_K^Q(\boldsymbol{\theta})$ and Hessian $\mathbf{H}_K^Q(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}}\nabla_{\boldsymbol{\theta}}^T\ell_K^Q(\boldsymbol{\theta})$ are given by

$$\mathbf{g}_K^Q(\boldsymbol{\theta}) := \sum_{k=1}^{K}\beta_k^Q(\boldsymbol{\theta})\mathbf{h}_k \tag{2.25a}$$

$$\mathbf{H}_K^Q(\boldsymbol{\theta}) := \sum_{k=1}^{K}\gamma_k^Q(\boldsymbol{\theta})\mathbf{h}_k\,\mathbf{h}_k^T \tag{2.25b}$$

where $\beta_k^Q(\boldsymbol{\theta})$ and $\gamma_k^Q(\boldsymbol{\theta})$ are defined in Appendix A.4.

The quantized-censored MLE (qcMLE) is consequently obtained as

$$\hat{\boldsymbol{\theta}}_{qcMLE} = \arg\max_{\boldsymbol{\theta}\in\mathbb{R}^p}\ell_K^Q(\boldsymbol{\theta}).$$

It should be noted that the estimator $\hat{\boldsymbol{\theta}}_{qcMLE}$ will be a function of not only the censoring thresholds $\{\tau_k\}_{k=1}^K$ but also of the quantizer thresholds $\{t_j\}_{j\in\mathcal{J}}$. Intuitively, the more finely-quantized the uncensored measurements are (i.e., the larger L is), the closer the MSE of the qcMLE would be to that of the unquantized cMLE - this will be observed in the numerical studies of Section 2.5.

Analogous to the cMLE approach in Lemma 2.2, the CRLB for the quantized data pdf is given by Lemma 2.3.

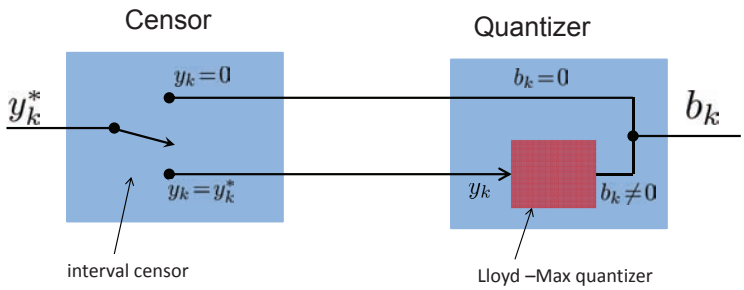

Figure 2.4: Censoring and quantization at sensor $\mathcal{S}_k$.

Lemma 2.3 *If $\hat{\boldsymbol{\theta}}$ denotes any unbiased estimator for $\boldsymbol{\theta}$ based on K data from the pdf (2.22), then*

$$\mathbb{E}\left[(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta})(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta})^T\right]-\left[\mathbf{I}_K^Q(\boldsymbol{\theta})\right]^{-1}\succeq \mathbf{0}\,.$$

The functional form of the FIM $\mathbf{I}_K^Q(\boldsymbol{\theta}):=-\mathbb{E}\left[\nabla_{\boldsymbol{\theta}}\nabla_{\boldsymbol{\theta}}^T\ell_K^Q(\boldsymbol{\theta})\right]$ is summarized in Appendix A.5.

2.4.1 Quantized-censored MLE algorithms

The qcMLE algorithm will have the same censoring step as in Algorithm 2.2 with a quantization step to follow before the MLE, which will now be based entirely on digital data $\{s_k, b_k, \{d_{jk}\}_{j\in\mathcal{J}}\}_{k=1}^K$. The FC needs to store a table of the codebook used by each of the K sensors in order to decode the received codewords $\{b_k\}_{k=1}^K$. The censor-quantizer and estimator schemes are summarized in Algorithm 2.4 and Algorithm 2.5, respectively.

2.5 Numerical studies: MLE

Three metrics are used to compare the performance of different estimators: (i) average fraction of active sensors $\rho=\bar{K}/K$, obtained from averaging N Monte Carlo runs; (ii) the empirical MSE:= $(1/N)\sum_{n=1}^N\|\hat{\boldsymbol{\theta}}(n)-\boldsymbol{\theta}_0\|_2^2$, where $\boldsymbol{\theta}_0$ is the true parameter vector, and $\hat{\boldsymbol{\theta}}(n)$ is the estimate obtained from the n-th Monte Carlo run; and (iii) the signal-to-noise ratio, SNR:=$(1/\sigma^2)\sum_{k=1}^K\|\mathbf{h}_k\|_2^2$. Simulations are performed for the model in (4.1), and compare performance of the cMLE, the A-optimal DOE-based MLE (from [32, Sec. V]),

Algorithm 2.4 Censoring and quantization (qcMLE)

Require: FC knows $\{\mathbf{h}_k\}_{k=1}^K$, $\bar{K}$; $\mathcal{S}_k$ knows $\mathbf{h}_k$, y_k^*

Initialization:

FC: Polls $L \ll K$ sensors for $\{y_\ell^*\}_{\ell=1}^L$, calculates $\bar{\boldsymbol{\theta}}$ using (2.3)

FC: Finds λ from (2.11) using a grid search

FC: Broadcasts (λ, $\bar{\boldsymbol{\theta}}$)

for $k = 1, 2, \ldots, K$ **do**

Censoring:

$\mathcal{S}_k$: Receives (λ, $\bar{\boldsymbol{\theta}}$), form threshold $\tau_k = \lambda / |\mathbf{h}_k^T \bar{\boldsymbol{\theta}}|$

$\mathcal{S}_k$: Censor y_k^* to get y_k using (2.10)

Quantizing:

if $y_k \neq 0$ **then**

$\mathcal{S}_k$: Quantize y_k using (2.23) to get b_k, which is transmitted to FC at time interval T_k

else

$\mathcal{S}_k$: Stays idle

end if

end for

Algorithm 2.5 Estimation (qcMLE)

Require: Gradient tolerance $\epsilon > 0$; maximum iterations $t_{\max}$

Data Reception:

FC: Receives $\{b_k\}_{k=1}^{K}$, where $b_k{=}0$ if no data received at time T_k

if $b_k = 0$ **then**

FC: $s_k = 0$ and $d_{jk} = 0,\ \forall\, j \in \mathcal{J}$

else

FC: $s_k = 1$, $d_{jk} = 1$ for $b_k = j$, and $d_{ik} = 0,\ \forall\, i \in \mathcal{J}/j$

end if

FC: **Parameter Estimation:**

FC: Initialize $t \leftarrow 0$, $\hat{\boldsymbol{\theta}}^{(0)} = \bar{\boldsymbol{\theta}}$

repeat

FC: Calculate $\ell_K^Q(\hat{\boldsymbol{\theta}}^{(t)})$, $\mathbf{g}_K^Q(\hat{\boldsymbol{\theta}}^{(t)})$, and $\mathbf{H}_K^Q(\hat{\boldsymbol{\theta}}^{(t)})$

FC: Use Armijo rule [cf. Algorithm 2.1] to find $\alpha^{(t)}$

FC: $\hat{\boldsymbol{\theta}}^{(t)} \leftarrow \hat{\boldsymbol{\theta}}^{(t)} + \alpha^{(t)} \left[\mathbf{H}_K^Q(\hat{\boldsymbol{\theta}}^{(t)})\right]^{-1} \mathbf{g}_K^Q(\hat{\boldsymbol{\theta}}^{(t)})$

FC: $t \leftarrow t + 1$

until $\|\mathbf{g}_K^Q(\hat{\boldsymbol{\theta}}^{(t)})\| \le \epsilon$ OR $t = t_{\max}$

FC: Set $\hat{\boldsymbol{\theta}}_{cMLE} = \hat{\boldsymbol{\theta}}^{(t)}$

the MLE based on full data, and the CRLB. The A-optimal DOE selects the $\bar{K}$ regressors that minimize the MSE of $\hat{\boldsymbol{\theta}}_{ls}(\bar{K})$ in (2.2). Unless otherwise stated, the regressors $\{\mathbf{h}_k\}_{k=1}^K$ are picked uniformly over $[-1, 1]$, and kept fixed for all the Monte Carlo runs.

2.5.1 Simulation study I

In the first simulation study two scenarios are investigated. The regressors $\{\mathbf{h}_k\}_{k=1}^K$ were picked uniformly over the $[-1, 1]$ range, and kept fixed for all the Monte Carlo runs.

The first scenario in Fig. 2.5 depicts the variation of MSE with ρ for a fixed SNR value (given by $\sigma^2 = 2.2.5$). A WSN with $K = 100$, with a large dimensional parameter space $p = 10$ and with $N = 100$ Monte Carlo simulation runs. Comparison of the MSE values are done for the cMLE, MLE of the A-optimal design, MLE from uniformly selected sensors data, and the full data MLE. With iid noise and uniformly distributed regressors, a sensor selection scheme that uniformly picks $\bar{K} < K$ sensors would, *on average*, yield MSE performance as good as any ingenious selection of $\bar{K}$ measurements. However, the *censoring-based* selection benefits from additional partial information deduced from the untransmitted (censored) measurements, namely $\Pr[s_k = 0; \boldsymbol{\theta}]$, in the pdf (2.13). Thus, it is expected that cMLE will perform better than the uniform selection. This intuition is corroborated by simulation tests where *censoring-based* MLE outperforms the A-optimal MLE. In fact, the MSE values for cMLE almost coincide with the CRLB over the entire range of ρ values. In this scenario A-optimal selection has no MSE performance advantage over random measurement selection.

The second scenario in Fig. 2.6 investigates the variation of SNR with ρ for a fixed MSE target. The SNR is changed by varying σ^2 values. The number of uncensored measurements $\bar{K} := \rho K$ needed to achieve MSE of 1.00 ± 0.25 for the given range of SNR values are depicted in Fig. 2.6. The simulations are done for $p = 2$, using $K = 400$ sensors, whereby for each SNR value $N = 1000$ Monte Carlo runs were used. It is seen that the cMLE requires fewer measurements to achieve MSE value of 1.00 ± 0.25 than A-optimal MLE over the -30 dB to -22 dB SNR range.

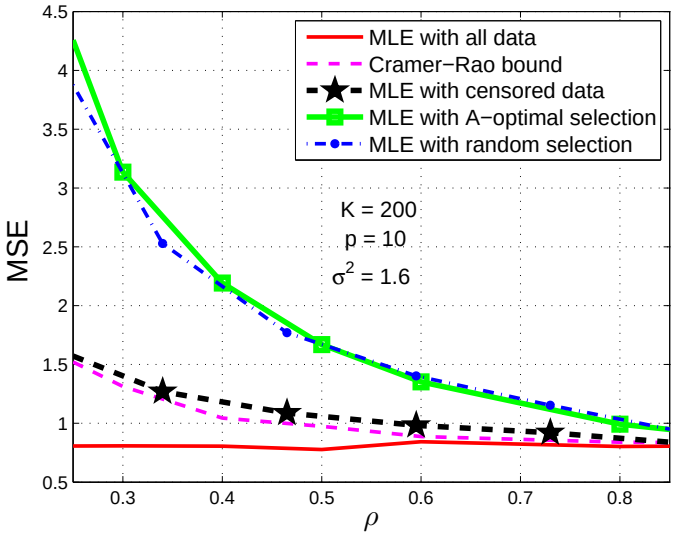

Figure 2.5: ρ vs. MSE: fixed SNR.

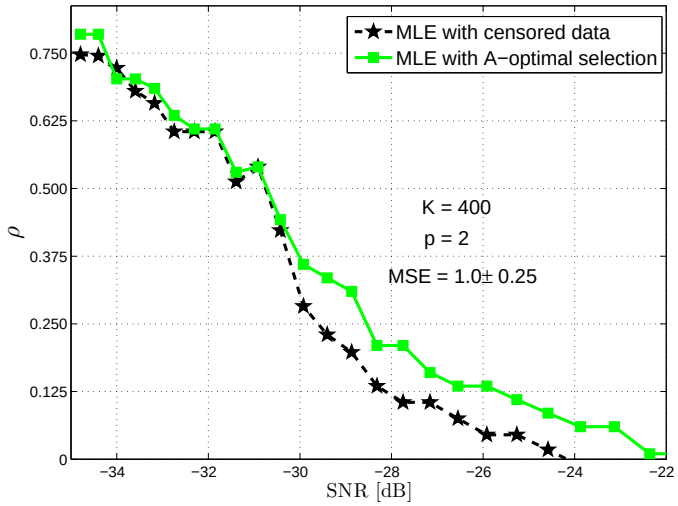

Figure 2.6: ρ vs. SNR: fixed MSE.

2.5.2 Simulation study II

Figs. 2.7 and 2.8 compare ρ vs. MSE for a fixed SNR, using two different distributions of regressor l_2-norms $\|\mathbf{h}_k\|_2$. The regressor norms are depicted in corresponding plots. In the first case of Fig. 2.7, the regressors have comparable l_2-norms which are termed homogeneous regressors. The MSE of cMLE is close to that of the full-data MLE even for low ρ values. A-optimal MLE here exhibits slightly larger MSE than the random selection. The reason is that since homogeneous regressors have comparable effect on the estimator's covariance matrix, there is no advantage to selecting them using an optimization criterion over picking them uniformly at random. Incidentally, the regressors in Fig. 2.5 were homogeneously distributed as well, hence the near-coincidence of the A-optimal MLE and random-selection MLE curves.

A test case with the same settings as in Fig. 2.7 but with some regressors having a larger l_2-norm than others is depicted in Fig. 2.8. This scenario with heterogeneous regressors is encountered, e.g., in source localization where measurements from sensors closer to the source typically achieve larger MSE reduction than those distant from the source. A-optimal selection excels in the heterogeneous regressors case since it picks measurements whose regressors yield larger reduction on the MSE. It can be seen from Fig. 2.8 that the MSE of cMLE and of the A-optimal MLE are indistinguishable for all ρ values, and both are close to the MSE from the full-data MLE.

2.5.3 Intel lab data

In Fig. 2.9 the cMLE algorithm is tested for linear regression using the Intel Berkeley dataset from [40]. Data from 40 of the 54 deployed sensors are used. The temperature measurement is modeled as $t(k) = \theta_1 + k\theta_2 + v_k$, where k indexes the specific time instants (epochs) for which measurements were acquired. Sliding windows of data sizes $K = 1,000$ (corresponding to about 13 minutes of data acquisition at 1 sample every 31 seconds) are used. The combination of small data frames and use of polynomial regressors leads to a model that fits the dataset quite well; this point was also observed in [22] where for larger data frames polynomials of higher degree or Gaussian kernels were used as regressors.

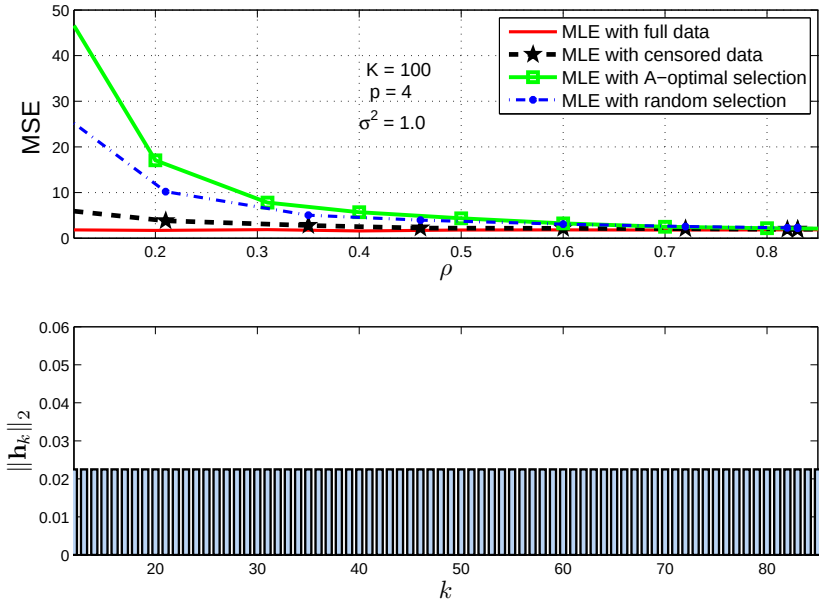

Figure 2.7: ρ vs. MSE: homogeneous regressors.

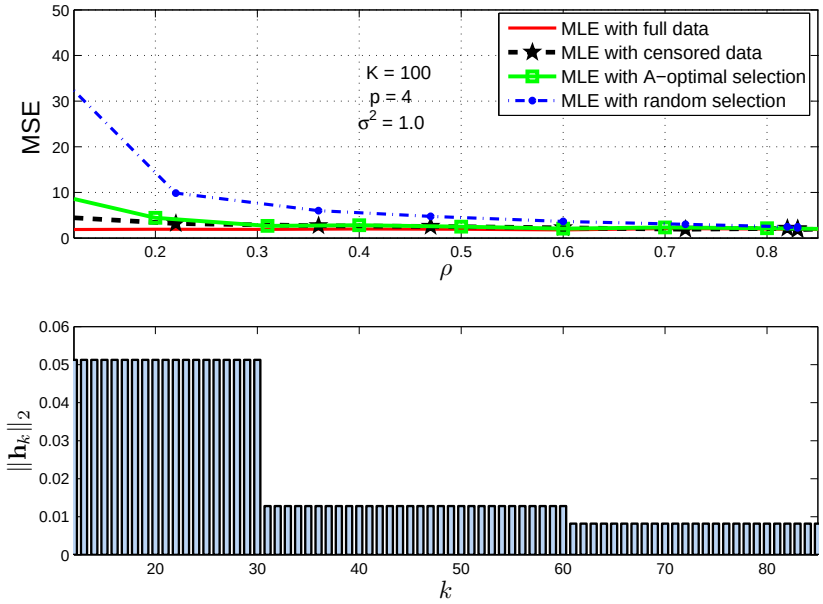

Figure 2.8: ρ vs. MSE: heterogeneous regressors.

However, the linear regression model suffices for the slow temperature variations in the time scales envisioned.

Initially, 500 data samples are used to estimate the trend $\bar{y} = \frac{1}{500}\sum_{i=1}^{500} y_i^*$ and the variance $\bar{\sigma}^2 = \frac{1}{500}\sum_{i=1}^{500}(y_i^* - \bar{y})^2$. Assuming that the modeling errors and measurement noise are Gaussian, with the estimated variance $\bar{\sigma}^2$, the MSE performance of the cMLE for different $\bar{K}$ values is compared to the A-optimal MLE as well as the benchmark MLE based on all K measurements. It should be noted that the regressors in this case are heterogeneous and, as expected, the A-optimal MLE performs well. Fig. 2.9 shows, however, that for moderate values of $\bar{K}$ cMLE clearly outperforms the A-optimal MLE. The expected poor performance of uniform selection in heterogeneous regressor scenario is also verified.

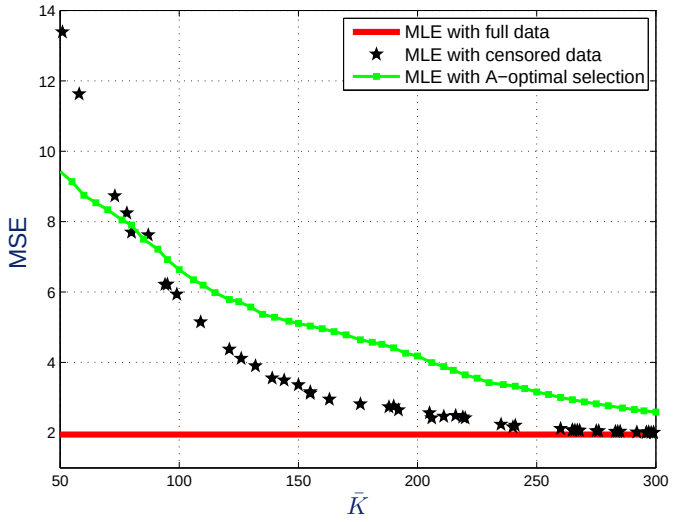

Figure 2.9: MSE vs. $\bar{K}$ for the Intel lab temperature data.

2.6 Concluding remarks

Data selection to save resources for parameter estimation with wireless sensor networks was explored in this chapter. Formulated as a censoring problem, a sensor-centric data selection method was introduced. A maximum likelihood iterative estimator was developed incorporating knowledge of the censoring technique. Furthermore, quantization of the uncensored

measurements was investigated to further reduce the amount of information from the sensor nodes to the FC. Performance analysis based on the Cramér-Rao bound, and numerical studies with synthetic and real-world data demonstrated the potential of measurement censoring as a data-reduction technique for estimation with large sensor networks. In the next chapter, modeling the unknown parameters as a random vector with a known pdf will be pursued. Bayesian estimator counterparts for censored, and censored-quantized schemes will be explored.

Chapter 3

Censored Maximum a Posteriori Estimation

3.1 Censored Maximum a Posteriori estimation

3.1.1 MAP estimation with uncensored data

Consider that the measurement at sensor node $\mathcal{S}_k$ follows the linear regression model $y_k^* = \mathbf{h}_k^T\boldsymbol{\theta} + v_k$ similar to (4.1). However, instead of the unknown parameter vector $\boldsymbol{\theta}$ being deterministic, let it be a Gaussian random vector with pdf $p(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta}; \mu_\theta, \sigma_\theta^2)$. Then the MAP estimator of $\boldsymbol{\theta}$ based on the uncensored data $\{y_k^*\}_{k=1}^K$ is given by

$$\hat{\boldsymbol{\theta}}_{MAP}(K) = \arg\max_{\boldsymbol{\theta}\in\mathbb{R}^p} p(\boldsymbol{\theta}|\mathbf{y}^*) \tag{3.1}$$

where $\mathbf{y}^* := [y_1^*, \ldots, y_K^*]^T$, $p(\boldsymbol{\theta}|\mathbf{y}^*) \propto \prod_{k=1}^K p(y_k^*|\boldsymbol{\theta})p(\boldsymbol{\theta})$, and $p(y_k^*|\boldsymbol{\theta}) = \mathcal{N}(y_k^*; \mathbf{h}_k^T\boldsymbol{\theta}, \sigma^2)$.

Paralleling the MLE approach, the MAP estimator (3.1) can be obtained as a minimizer of $-\ln p(\boldsymbol{\theta}|\mathbf{y}^*)$, which leads to a regularized LS problem

$$\hat{\boldsymbol{\theta}}_{MAP}(K) = \arg\min_{\boldsymbol{\theta}\in\mathbb{R}^p} \left\{ \frac{1}{\sigma^2}\sum_{k=1}^K (y_k^* - \mathbf{h}_k^T\boldsymbol{\theta})^2 + (\boldsymbol{\theta}-\boldsymbol{\mu}_\theta)^T\mathbf{C}_{\boldsymbol{\theta}}^{-1}(\boldsymbol{\theta}-\boldsymbol{\mu}_\theta) \right\}$$

and admits the closed-form solution

$$\hat{\boldsymbol{\theta}}_{MAP}(K) = \left(\sum_{k=1}^K \mathbf{h}_k\mathbf{h}_k^T + \sigma^2\mathbf{C}_\theta^{-1}\right)^{-1}\left(\sum_{k=1}^K y_k^*\mathbf{h}_k + \sigma^2\mathbf{C}_\theta^{-1}\boldsymbol{\mu}_\theta\right) \tag{3.2}$$

with estimator error covariance matrix

$$\begin{aligned}\mathbf{C}_{MAP}(K) &:= \mathbb{E}\left[(\hat{\boldsymbol{\theta}}_{MAP}(K)-\boldsymbol{\theta})(\hat{\boldsymbol{\theta}}_{MAP}(K)-\boldsymbol{\theta})^T\right] \\ &= \sigma^2\left(\sum_{k=1}^{K}\mathbf{h}_k\mathbf{h}_k^T+\sigma^2\mathbf{C}_\theta^{-1}\right)^{-1} \end{aligned} \tag{3.3}$$

and estimator variance given by $\mathrm{tr}[\mathbf{C}_{MAP}(K)]$. It should be noted that the MAP estimator for the linear-Gaussian model, given that the prior $p(\boldsymbol{\theta})$ is a known Gaussian pdf, coincides with the Bayesian minimum MSE estimator [34, Ch. 11]. The MAP estimator (3.2) will benchmark performance of MAP estimators based on censored and quantized data. Furthermore, if $\mathbf{C}_\theta^{-1}=\mathbf{0}$, meaning that the prior pdf of $\boldsymbol{\theta}$ is degenerate, the MAP estimator in (3.2) coincides with the MLE when $\boldsymbol{\theta}$ is modeled as deterministic [34].

3.1.2 MAP estimation with data selection

Motivated by the need to reduce data transmission from the sensor nodes to the FC in a Bayesian estimation setting, the data selection approach can be applied also to MAP estimation [32], with the goal of finding a selection-based MAP (sMAP) estimator

$$\hat{\boldsymbol{\theta}}_{sMAP}(\bar{K}) = \arg\max_{\boldsymbol{\theta}\in\mathbb{R}^p} p(\boldsymbol{\theta}|\mathbf{y},\boldsymbol{s})$$

where $\mathbf{y}\in\mathbb{R}^K$ has zero entries corresponding to measurements not selected, and $\boldsymbol{s}\in\{0,1\}^K$ denotes the measurement selection vector as in the sMLE setting of Section 2.1.1.

Following [32], entries of $\boldsymbol{s}$ are selected to approximately minimize the sMAP estimator variance $\mathrm{tr}\Big\{\Big(\sum_{k=1}^K s_k\mathbf{h}_k\mathbf{h}_k^T+\sigma^2\mathbf{C}_\theta^{-1}\Big)^{-1}\Big\}$, subject to the constraint $\sum_{k=1}^K s_k=\bar{K}$. This sensor selection method corresponds to the so-termed Bayesian A-optimal design, and its centralized implementation in [32] incurs computational complexity $\mathcal{O}(K^3)$. It is solved off-line at the FC prior to sensors acquiring any measurements since the selection criterion does not depend on the measurements.

With a given set of $\bar{K}$ selected measurements $\{y_k^*\colon\ s_k=1\}$ sent to the FC, the sMAP estimator for the linear-Gaussian model (4.1) is obtained from the MAP estimator (3.2)-(3.3). For non-linear regression or non-Gaussian noise models, linearization would be needed to

use the centralized sensor selection criterion of [32]. In the ensuing section a novel *decentralized* sensor data selection is developed by using measurement censoring whose overall complexity scales as $\mathcal{O}(K)$.

3.2 Data censoring and MAP estimation

Measurement reduction via censoring is detailed first for a random $\boldsymbol{\theta}$ as an alternative to data selection that is both computationally efficient and decentralized. Analysis of the Bayesian MSE for the resultant censored MAP (cMAP) estimator is presented afterwards.

3.2.1 Data censoring

Censoring when $\boldsymbol{\theta}$ is random differs from that of cMLE. Using $p(\boldsymbol{\theta}) = \mathcal{N}(\boldsymbol{\theta}; \mu_\theta, \mathbf{C}_\theta)$ as prior, it follows that $p(y_k^*) = \mathcal{N}(y_k^*; \mu_{y_k^*}, \sigma^2_{y_k^*})$, where $\mu_{y_k^*} := \mathbf{h}_k^T \boldsymbol{\mu}_\theta$ and $\sigma^2_{y_k^*} := \mathbf{h}_k^T \mathbf{C}_\theta \mathbf{h}_k + \sigma^2$. Letting $\tau_k > 0$ denote the censoring threshold at $\mathcal{S}_k$, and leveraging knowledge of the data pdf, consider the following per-sensor censoring rule

$$y_k = \begin{cases} y_k^*, & \text{if} \quad \frac{|y_k^* - \mu_{y_k^*}|}{\sigma_{y_k^*}} > \tau_k \\ 0, & \text{otherwise}\,. \end{cases} \tag{3.4}$$

The normalized measurements $z_k := (y_k^* - \mu_{y_k^*})/\sigma_{y_k^*}$ have pdf $p(z_k) = \mathcal{N}(z_k; 0, 1)$.

Letting $s_k := \mathbf{1}_{\{y_k \neq 0\}}$, the probability that a measurement is not censored is clearly

$$\Pr[s_k = 1] = \Pr[|z_k| > \tau_k] = 2Q(\tau_k) \tag{3.5}$$

while the expected number of uncensored measurements is given by

$$\mathbb{E}\Big[\sum_{k=1}^K s_k\Big] = \sum_{k=1}^K \Pr[|z_k| > \tau_k] = 2\sum_{k=1}^K Q(\tau_k)\,.$$

If the expected number of uncensored measurements is bounded to be at most $\bar{K}$, it follows that

$$\frac{1}{K}\sum_{k=1}^K Q(\tau_k) \leq \frac{1}{2}\rho$$

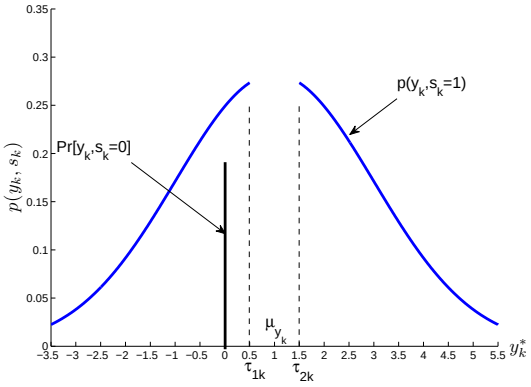

Figure 3.1: Truncated Gaussian pdf from the censored data model

where $\rho := \bar{K}/K$ denotes the fraction of uncensored measurements. Since z_k are identically distributed, identical thresholds $\tau_k = \tau$ for $k = 1, \ldots, K$, will be used, leading to the threshold selection condition $Q(\tau) \leq \frac{1}{2}\rho$, which implies that

$$\tau := Q^{-1}(\rho/2) \tag{3.6}$$

is the censoring threshold for which the average number of uncensored measurements is $\bar{K}$. Censoring in (3.4) can alternatively be expressed as

$$(y_k,\ s_k) := \begin{cases} (0,\ 0), & \text{if} \quad \tau_{1k} < y_k^* \leq \tau_{2k} \\ (y_k^*,\ 1), & \text{otherwise} \end{cases} \tag{3.7}$$

where $\tau_{1k} := -\tau\,\sigma_{y_k^*} + \mu_{y_k^*}$, and $\tau_{2k} := \tau\,\sigma_{y_k^*} + \mu_{y_k^*}$. Fig. 3.1 depicts an example of the censored-data pdf for $t = 0.5$, $\sigma^2_{y_k^*} = 1$, and $\mu_{y_k^*} = 1$ leading to $\tau_{1k} = 0.5$ and $\tau_{2k} = 1.5$. Note the difference in censor-thresholds adopted by cMAP with those used for cMLE in Section 2.2.

3.2.2 MAP estimation with censored data

The measurements $(\mathbf{y}, \boldsymbol{s})$ used in estimation are distributed according to the conditional pdf $p(\mathbf{y}, \boldsymbol{s}|\boldsymbol{\theta})$ given by the following lemma whose proof is quite similar to that of Lemma 2.1, and is thus omitted.

Lemma 3.1 *For the latent data model (4.1) and with the censoring per sensor performed as in (3.7), the censored data* $\{y_k, s_k\}_{k=1}^K$ *are conditionally independent given* $\boldsymbol{\theta}$*; i.e.,* $p(\mathbf{y}, \boldsymbol{s}|\boldsymbol{\theta}) = \prod_{k=1}^K p(y_k, s_k|\boldsymbol{\theta})$.

Using Bayes' rule, it follows that $p(\boldsymbol{\theta}|\mathbf{y}, \boldsymbol{s}) \propto p(\mathbf{y}, \boldsymbol{s}|\boldsymbol{\theta})p(\boldsymbol{\theta})$, and the cMAP estimator is

$$\hat{\boldsymbol{\theta}}_{cMAP} = \arg\max_{\boldsymbol{\theta}\in\mathbb{R}^p} p(\boldsymbol{\theta}|\mathbf{y}, \boldsymbol{s}) = \arg\max_{\boldsymbol{\theta}\in\mathbb{R}^p} \prod_{k=1}^K p(y_k, s_k|\boldsymbol{\theta})p(\boldsymbol{\theta})\,. \tag{3.8}$$

From (4.1) it follows that the conditional pdf in (3.8) is [cf. (3.7)]

$$p(y_k, s_k|\boldsymbol{\theta}) = \left[\mathcal{N}(y_k; \mathbf{h}_k^T\boldsymbol{\theta}, \sigma^2)\right]^{s_k} \left[\Pr[y_k = 0|\boldsymbol{\theta}]\right]^{1-s_k}$$

where

$$\Pr[y_k = 0|\boldsymbol{\theta}] = Q\left(\frac{\tau_{1k} - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right) - Q\left(\frac{\tau_{2k} - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right).$$

Using $z_{1k} := (\tau_{1k} - \mathbf{h}_k^T\boldsymbol{\theta})/\sigma$ and $z_{2k} := (\tau_{2k} - \mathbf{h}_k^T\boldsymbol{\theta})/\sigma$, the product terms in (3.8) can be written compactly as

$$p(\mathbf{y}, \boldsymbol{s}|\boldsymbol{\theta})\, p(\boldsymbol{\theta}) = \prod_{k=1}^K \left[\mathcal{N}(y_k; \mathbf{h}_k^T\boldsymbol{\theta}, \sigma^2)\right]^{s_k} \left[Q(z_{1k}) - Q(z_{2k})\right]^{1-s_k} \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\mu}_\theta, \mathbf{C}_\theta)\,. \tag{3.9}$$

Taking logarithm of both sides of (3.9) yields $\ell_K^B(\boldsymbol{\theta}) := \log\left[p(\mathbf{y}, \boldsymbol{s}|\boldsymbol{\theta})p(\boldsymbol{\theta})\right]$ which is given by

$$\ell_K^B(\boldsymbol{\theta}) = \sum_{k=1}^K \left\{ \frac{s_k}{2\sigma^2}\left(y_k - \mathbf{h}_k^T\boldsymbol{\theta}\right)^2 + \frac{s_k}{2}\ln\left(2\pi\sigma^2\right) - (1-s_k)\ln\left[Q\left(\frac{\tau_{1k} - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right) - Q\left(\frac{\tau_{2k} - \mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}\right)\right]\right\}$$
$$+ \frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)^T\mathbf{C}_\theta^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta) + \frac{1}{2}\ln\left((2\pi)^p|\mathbf{C}_\theta|\right). \tag{3.10}$$

Based on (3.10), the cMAP estimator (3.8) can thus be obtained as

$$\hat{\boldsymbol{\theta}}_{cMAP} = \arg\max_{\boldsymbol{\theta}\in\mathbb{R}^p} \ell_K^B(\boldsymbol{\theta})\,. \tag{3.11}$$

The following proposition, whose proof is an extension of that of Proposition 2.1 given in Appendix A.1 and is thus omitted, summarizes conditions ensuring a unique solution to the optimization problem in (3.11).

Proposition 3.1 *Function* $\ell_K^B(\boldsymbol{\theta})$ *in (3.10) is strictly concave if at least one of the following conditions holds: (i)* $\sum_{k=1}^K \mathbf{h}_k\mathbf{h}_k^T \succ \mathbf{0}$; *and/or (ii)* $\mathbf{C}_\theta \succ \mathbf{0}$.

3.2.3 Censored MAP algorithm

The function $\ell_K^B(\boldsymbol{\theta})$ is twice differentiable with gradient $\mathbf{g}_K^B(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}} \ell_K^B(\boldsymbol{\theta})$ and Hessian $\mathbf{g}_K^B(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T \ell_K^B(\boldsymbol{\theta})$ given by

$$\mathbf{g}_K^B(\boldsymbol{\theta}) = \sum_{k=1}^{K} \beta_k^B(\boldsymbol{\theta}) \mathbf{h}_k + \mathbf{C}_\theta^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta) \tag{3.12a}$$

$$\mathbf{H}_K^B(\boldsymbol{\theta}) = \sum_{k=1}^{K} \gamma_k^B(\boldsymbol{\theta}) \mathbf{h}_k \, \mathbf{h}_k^T + \mathbf{C}_\theta^{-1} \tag{3.12b}$$

where $\beta_k^B(\boldsymbol{\theta})$ and $\gamma_k^B(\boldsymbol{\theta})$ are defined in Appendix B.2.

The details of how the censoring and iterative estimation are implemented with an FC-based WSN are quite similar to those of the cMLE approach, except that now both FC and sensors are assumed to know the prior pdf, and the FC broadcasts the censoring threshold τ instead of λ used in the cMLE. The MAP censoring and estimation schemes are tabulated as Algorithms 3.1 and 3.2, respectively.

Algorithm 3.1 Censoring (cMAP)

Require: FC knows $\{\mathbf{h}_k\}_{k=1}^K$, σ^2, $\boldsymbol{\mu}_\theta$, $\mathbf{C}_\theta$, ρ

$\mathcal{S}_k$ knows $\mathbf{h}_k$, $\boldsymbol{\mu}_\theta$, $\mathbf{C}_\theta$, σ^2, y_k^*

Initialization:

FC: Calculates $\tau := Q^{-1}(\rho/2)$

FC: Broadcasts τ

Sensor Censoring:

for $k = 1, 2, \ldots, K$ **do**

$\mathcal{S}_k$: Receives τ, and calculates thresholds τ_{1k} and τ_{2k}

$\mathcal{S}_k$: Censors y_k^* to find y_k using (3.7)

if $y_k \neq 0$ **then**

$\mathcal{S}_k$: Transmit y_k to FC

else

$\mathcal{S}_k$: Stays idle

end if

end for

Algorithm 3.2 Estimation (cMAP)

Require: Gradient tolerance $\epsilon > 0$; maximum iterations $t_{\max}$

Data Reception:

FC: Receive $\{y_k\}_{k=1}^K$, where $y_k = 0$ if no data received at time T_k

if $y_k = 0$ **then**

FC: $s_k = 0$

else

FC: $s_k = 1$

end if

Estimation:

FC: Initialize $t \leftarrow 0$, draw $\hat{\boldsymbol{\theta}}^{(0)}$ from $\mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\mu}_\theta, \mathbf{C}_\theta)$

repeat

FC: Calculate $\ell_K^B(\hat{\boldsymbol{\theta}}^{(t)})$, $\mathbf{g}_K^B(\hat{\boldsymbol{\theta}}^{(t)})$, and $\mathbf{H}_K^B(\hat{\boldsymbol{\theta}}^{(t)})$ using (3.10), (3.12a), and (3.12b)

FC: Find $\alpha^{(t)}$ using Armijo's line search rule [9]

FC: $\hat{\boldsymbol{\theta}}^{(t)} \leftarrow \hat{\boldsymbol{\theta}}^{(t)} + \alpha^{(t)} \left[\mathbf{H}_K^B(\hat{\boldsymbol{\theta}}^{(t)})\right]^{-1} \mathbf{g}_K^B(\hat{\boldsymbol{\theta}}^{(t)})$

FC: $t \leftarrow t + 1$

until $\|\mathbf{g}_K^B(\hat{\boldsymbol{\theta}}^{(t)})\| \leq \epsilon$ OR $t = t_{\max}$

FC: Set $\hat{\boldsymbol{\theta}}_{cMAP} = \hat{\boldsymbol{\theta}}^{(t)}$

3.2.4 Performance analysis of the cMAP estimator

The BCRLB provides a lower bound on the Bayesian MSE of an estimator [77, p. 72]. The BCRLB is based on the FIM $\mathbf{I}_K^B$, which is defined as

$$\begin{aligned}\mathbf{I}_K^B &:= \mathbb{E}\left[\nabla_\theta \ln p(\mathbf{y}, \boldsymbol{s}, \boldsymbol{\theta})\ \nabla_\theta^T \ln p(\mathbf{y}, \boldsymbol{s}, \boldsymbol{\theta})\right] \\ &= \sum_{k=1}^K \bar{\gamma}_k^B \mathbf{h}_k \mathbf{h}_k^T + \mathbf{C}_\theta^{-1}. \end{aligned} \tag{3.13}$$

Derivation of the FIM in (3.13), including the definition of $\bar{\gamma}_k$, is detailed in Appendix B.3. The BCRLB for the censored-data model provides a lower bound on the MSE of any Bayesian estimator based on the censored data as summarized next.

Lemma 3.2 *Given the prior pdf $p(\boldsymbol{\theta})$, for any estimator $\hat{\boldsymbol{\theta}}$ of $\boldsymbol{\theta}$ based on $\{\mathbf{y}, \boldsymbol{s}\}$ drawn from*

the pdf $p(\mathbf{y}, \boldsymbol{s})$*, it holds that*

$$\mathbb{E}\left[(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^T\right] - \left[\mathbf{I}_K^B\right]^{-1} \succeq \mathbf{0}$$

where $\mathbb{E}[\cdot]$ *is over the joint pdf* $p(\mathbf{y}, \boldsymbol{s}, \boldsymbol{\theta})$*.*

The scalar $\bar{\gamma}_k^B$ in (3.13) is related to the contribution of the k-th measurement to the "increase" of the information in the FIM expression (3.13). Let $\bar{\gamma}_k^B$ for $\sigma^2 = 1, \|\mathbf{h}_k\|_2 = 1, \forall k$, be called the fractional contribution factor (FCF). For the censored data FIM in (3.13), the FCF obtained using (B.1) is approximately the same for $k = 1, 2, \ldots, K$. The variation of $\bar{\gamma}_k^B$ with ρ can be deduced after substituting $\tau = Q^{-1}(\rho/2)$ from (3.6) into (B.1).

If $\bar{K}$ out of the K measurements were uniformly selected at random, it can be readily shown that the FIM would be $\mathbf{I}_K^B = (\rho/\sigma^2)\sum_{k=1}^K \mathbf{h}_k\mathbf{h}_k^T + \mathbf{C}_\theta^{-1}$. Thus, the FCF corresponding to random selection for $\sigma^2 = 1$ equals ρ. For the full-data pdf, each measurement's FCF is $\rho = 1$. Fig. 3.2 illustrates the variation of FCF with ρ for the censored-data FIM, i.e, $\text{FCF}_{\text{censor}} = \bar{\gamma}_k^B$; for the uncensored (or full)-data FIM, i.e. $\text{FCF}_{\text{full}} = 1$; and for the random selection FIM, i.e. $\text{FCF}_{\text{select}} = \rho$. The figure is plotted for a scalar θ, with $p(\theta) = \mathcal{N}(\theta; 0, C_\theta)$. It can be seen that with ρ value of about 0.7, the censored-data FCF equals approximately that of the full-data pdf. The linear variation of the FCF ρ for the random selection serves to highlight the advantage of using censoring over random selection.

For the Bayesian A-optimal design in [32], the FCF is obtainable only numerically due to lack of a closed-form expression for the selection indices $\boldsymbol{s}$. Since the selected $\bar{K}$ sensors effect the largest reduction in MSE, A-optimal design's FCF will generally be larger than the FCF of ρ achieved by random selection. However, the censored data pdf accounts not only for the $\bar{K}$ selected (uncensored) measurements, but also for the probability information from the $K - \bar{K}$ censored measurements, namely, $\Pr[s_k = 0] = 1 - \Pr[s_k = 1]$ from (3.5), and thus it should have an FCF larger than that corresponding to the A-optimal selection. In Section 3.4, comparison of the MSE, which is related to the FIM via the BCRLB, will corroborate this intuition.

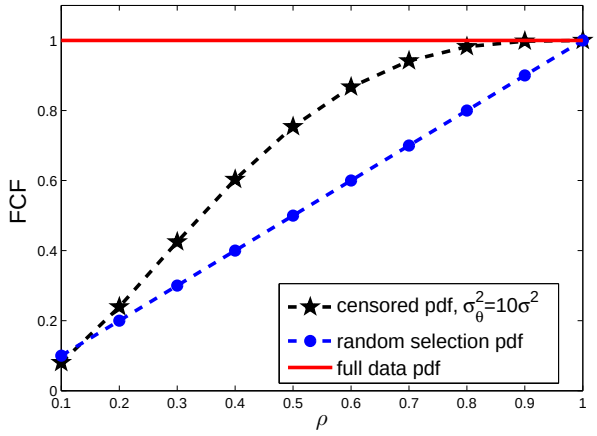

Figure 3.2: Fractional contribution factor (FCF) of a measurement to FIM.

3.3 MAP estimation with quantized-censored data

If the $\mathcal{S}_k$ measurement is not censored, Bayes' rule yields the conditional pdf $p(y_k|s_k=1)$ as a truncated Gaussian with mean $\mu_{y_k} := \mathbf{h}_k^T \boldsymbol{\mu}_\theta$, and variance $\sigma_{y_k}^2 := \mathbf{h}_k^T \mathbf{C}_{\boldsymbol{\theta}} \mathbf{h}_k + \sigma^2$; that is,

$$p(y_k|s_k=1) = \frac{\mathcal{N}(y_k; \mu_{y_k}, \sigma_{y_k}^2)}{\Pr[y_k \neq 0]} \mathbf{1}_{\{y_k \neq 0\}} .$$

Using (3.5) and (3.6), it clearly follows that $\Pr[y_k \neq 0] = \rho$. Due to symmetry of the conditional pdf $p(y_k|s_k=1)$, it follows readily that $\mathbb{E}[y_k|s_k=1] = \mathbf{h}_k^T \boldsymbol{\mu}_\theta$, and $\mathrm{cov}(y_k, y_l|s_k=1, s_l=1) = \mathbf{h}_k^T \mathbf{C}_{\boldsymbol{\theta}} \mathbf{h}_l + \sigma^2 \delta_{lk}$. Unlike the qcMLE approach in (2.22), in the quantized censored MAP (qcMAP) case, sensor $\mathcal{S}_k$ has knowledge of the truncated Gaussian pdf $p(y_k|s_k=1)$ and does not need to invoke the estimate $\bar{\boldsymbol{\theta}}$ for the quantization step. Each $y_k \neq 0$ is quantized using an L-level Lloyd-Max quantizer yielding the finite-alphabet data b_k, s_k, and $\{\boldsymbol{d}_j\}_{j\in\mathcal{J}}$. The conditional pmf for the quantized data is given by

$$\Pr\left[\boldsymbol{s}, \boldsymbol{b}, \{\boldsymbol{d}_j\}_{j\in\mathcal{J}}|\boldsymbol{\theta}\right] = \prod_{k=1}^{K} \left\{ \Pr\left[s_k=0|\boldsymbol{\theta}\right]^{(1-s_k)} \prod_{j\in\mathcal{J}} \Pr\left[b_k=j|\boldsymbol{\theta}\right]^{d_{jk}} \right\}$$

where $\mathbf{b}$, $\boldsymbol{s}$, and $\boldsymbol{d}_j$ are defined as in Section 2.4. The quantizer partitions are illustrated in Fig. 3.3.

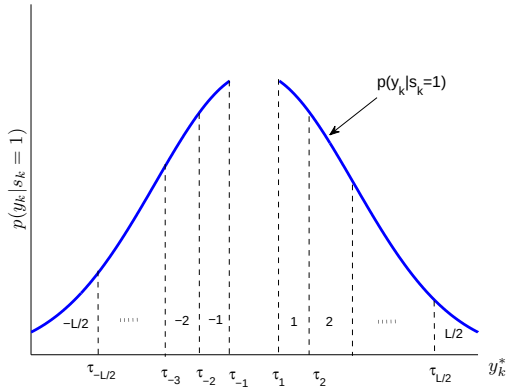

Figure 3.3: L-level quantization for uncensored data

Since the posterior pdf is $p(\boldsymbol{\theta}|\boldsymbol{s},\boldsymbol{b},\boldsymbol{d}) \propto \Pr[\boldsymbol{s},\boldsymbol{b},\boldsymbol{d}|\boldsymbol{\theta}]\, p(\boldsymbol{\theta})$, its logarithm, ignoring additive terms not dependent on $\boldsymbol{\theta}$, is given by

$$\ell_K^{BQ}(\boldsymbol{\theta}) = \sum_{k=1}^{K}\left\{(1-s_k)\ln\left[Q(z_{1k})-Q(z_{2k})\right] + \sum_{j\in\mathcal{J}} d_{jk}\ln\left[Q(\bar{z}_{1jk})-Q(\bar{z}_{2jk})\right]\right\} - \frac{1}{2}(\boldsymbol{\theta}-\boldsymbol{\mu}_\theta)^T\mathbf{C}_\theta^{-1}(\boldsymbol{\theta}-\boldsymbol{\mu}_\theta) \tag{3.14}$$

where $\bar{z}_{1jk} := (\tau_j - \mathbf{h}_k^T\boldsymbol{\theta})/\sigma$, and $\bar{z}_{2jk} := (\tau_{j+1} - \mathbf{h}_k^T\boldsymbol{\theta})/\sigma$.

The gradient and Hessian of the objective function (3.14) are detailed in Appendix B.4. Algorithms for qcMAP censoring and estimation can be obtained by replacing y_k with b_k and $\ell^B(\cdot), \mathbf{g}_K^B(\cdot), \mathbf{H}_K^B(\cdot)$ with $\ell^{BQ}(\cdot), \mathbf{g}_K^{BQ}(\cdot), \mathbf{H}_K^{BQ}(\cdot)$ (cf. Appendix B.4) respectively, in Algorithms 3.1 and 3.2. The resulting qcMAP censor-quantizer and estimation algorithms are tabulated in Algorithms 3.3 and 3.4, respectively. The expression for the FIM of the quantized-censored pmf is also provided in Appendix B.4.

3.4 Numerical studies: MAP

Simulations were carried out for the linear model (4.1). Similar to the ML approach, the three metrics used for performance comparison are: (i) the number of active sensors, ρ; (ii) empirical MSE:= $\frac{1}{R}\sum_{r=1}^{R}\frac{1}{N}\sum_{n=1}^{N}\|\hat{\boldsymbol{\theta}}(n,r)-\boldsymbol{\theta}(r)\|^2$, where $\boldsymbol{\theta}(r)$ is r-th realization of $\boldsymbol{\theta}$ from the prior, and $\hat{\boldsymbol{\theta}}(n,r)$ is the corresponding estimate from the n-th Monte Carlo

Algorithm 3.3 Censoring and quantization (qcMAP)

Require: FC knows $\{\mathbf{h}_k\}_{k=1}^K$, σ^2, $\boldsymbol{\mu}_\theta$, $\mathbf{C}_\theta$, ρ

$\mathcal{S}_k$ knows $\mathbf{h}_k$, $\boldsymbol{\mu}_\theta$, $\mathbf{C}_\theta$, σ^2, y_k^*

Initialization:

FC: Calculates $\tau := Q^{-1}(\rho/2)$

FC: Broadcasts τ

for $k = 1, 2, \ldots, K$ **do**

Censoring:

$\mathcal{S}_k$: Receives τ, and calculates thresholds τ_{1k} and τ_{2k}

$\mathcal{S}_k$: Censors y_k^* to find y_k using (3.7)

Quantizing:

if $y_k \neq 0$ **then**

$\mathcal{S}_k$: Quantize y_k to get b_k, which is transmitted to FC at time interval T_k

else

$\mathcal{S}_k$: Stays idle

end if

end for

Algorithm 3.4 Estimation (qcMAP)

Require: Gradient tolerance $\epsilon > 0$; maximum iterations $t_{\max}$

Data Reception:

FC: Receives $\{b_k\}_{k=1}^K$, where $b_k{=}0$ if no data received at time T_k

if $b_k = 0$ **then**

FC: $s_k = 0$ and $d_{jk} = 0$, $\forall\, j \in \mathcal{J}$

else

FC: $s_k = 1$, $d_{jk} = 1$ for $b_k = j$, and $d_{ik} = 0$, $\forall\, i \in \mathcal{J}/j$

end if

FC: **Parameter Estimation:**

FC: Initialize $t \leftarrow 0$, $\hat{\boldsymbol{\theta}}^{(0)} = \bar{\boldsymbol{\theta}}$

repeat

FC: Calculate $\ell_K^{BQ}(\hat{\boldsymbol{\theta}}^{(t)})$, $\mathbf{g}_K^{BQ}(\hat{\boldsymbol{\theta}}^{(t)})$, and $\mathbf{H}_K^{BQ}(\hat{\boldsymbol{\theta}}^{(t)})$

FC: Use Armijo rule [cf. Algorithm 2.1] to find $\alpha^{(t)}$

FC: $\hat{\boldsymbol{\theta}}^{(t)} \leftarrow \hat{\boldsymbol{\theta}}^{(t)} + \alpha^{(t)} \left[\mathbf{H}_K^{BQ}(\hat{\boldsymbol{\theta}}^{(t)})\right]^{-1} \mathbf{g}_K^{BQ}(\hat{\boldsymbol{\theta}}^{(t)})$

FC: $t \leftarrow t + 1$

until $\|\mathbf{g}_K^{BQ}(\hat{\boldsymbol{\theta}}^{(t)})\| \leq \epsilon$ OR $t = t_{\max}$

FC: Set $\hat{\boldsymbol{\theta}}_{qcMAP} = \hat{\boldsymbol{\theta}}^{(t)}$

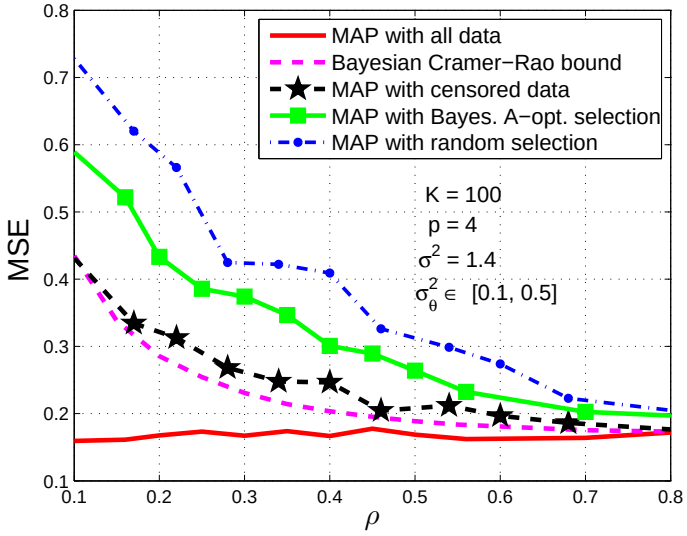

Figure 3.4: ρ vs. MSE: fixed SNR, cMAP estimator

run; and (iii) SNR:$=\frac{1}{\sigma^2}\sum_{k=1}^{K}\|\mathbf{h}_k\|^2$. The full-data MAP estimator and the BCRB were used to benchmark the performance comparison. A random selection-based estimator, and the Bayesian A-optimal design MAP estimator were compared to the cMAP and qcMAP estimators.

3.4.1 Degree of censoring vs. MSE

In the first simulation case ρ vs. MSE performance is studied for different settings. The first scenario in Fig. 3.4 shows that the cMAP estimator has variance that is close to the BCRB over the whole range of ρ values. It is also clear that Bayesian A-optimal design-based MAP estimator has worse variance than the cMAP estimator.

The second scenario in Figs. 3.5 and 3.6 a comparison of the estimators for homogeneous and heterogeneous regressors is done analogously to the MLE approach. In Figs. 3.5 the regressors have comparable ℓ_2-norms (i.e., homogeneous regressors). It is seen that the cMAP estimator outperforms estimators based on random selection and Bayesian A-optimal design. Random selection has comparable performance to Bayesian A-optimal design since for homogeneous sensors uniform selection performs as well as any optimization-based selection. Since cMAP uses the additional knowledge from censored measurements, its performance

is improved without requiring extra data transmission. It can be seen that cMAP is almost coincident with the BCRB even at low ρ values.

The heterogeneous case with different ℓ_2-norms for regressors across sensors is depicted in Fig. 3.6. It is seen that the Bayesian A-optimal design performs quite well and random selection exhibits the worst MSE performance. This is because in this case a sensor's location has an effect on its contribution to the MSE reduction, and Bayesian A-optimal design selection optimally selects measurements which reduce the MSE maximally. The cMAP performs as well as Bayesian A-optimal design at lower selection computation complexity.

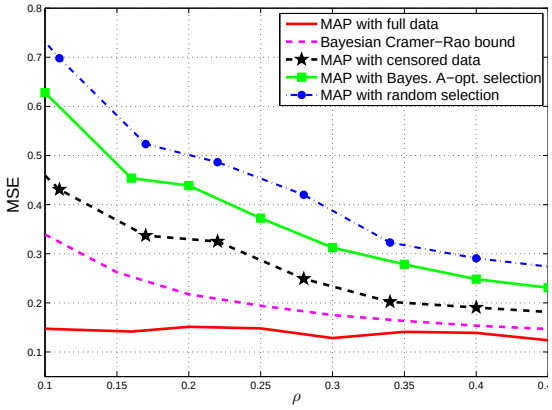

Figure 3.5: ρ vs. MSE: homogeneous regressors.

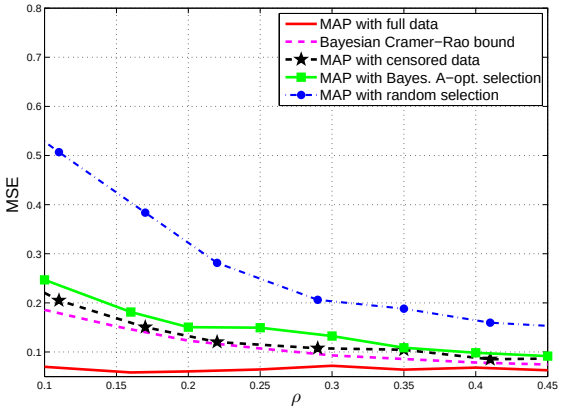

Figure 3.6: ρ vs. MSE: heterogeneous regressors.

3.4.2 High SNR vs. Low SNR performance

In Fig. 3.7 it is seen that with SNR = 6.02 dB, the cMAP estimator achieves the full-data MAP estimator MSE at moderate ρ values. However, as SNR degrades so does the performance of selection-based or censoring-based MAP estimators. For example, at SNR = 0 dB, there is an MSE gap between cMAP and full-data MAP estimator even for large ρ values. The A-optimal design has an even larger MSE gap to the full-data MAP when SNR = 0 dB.

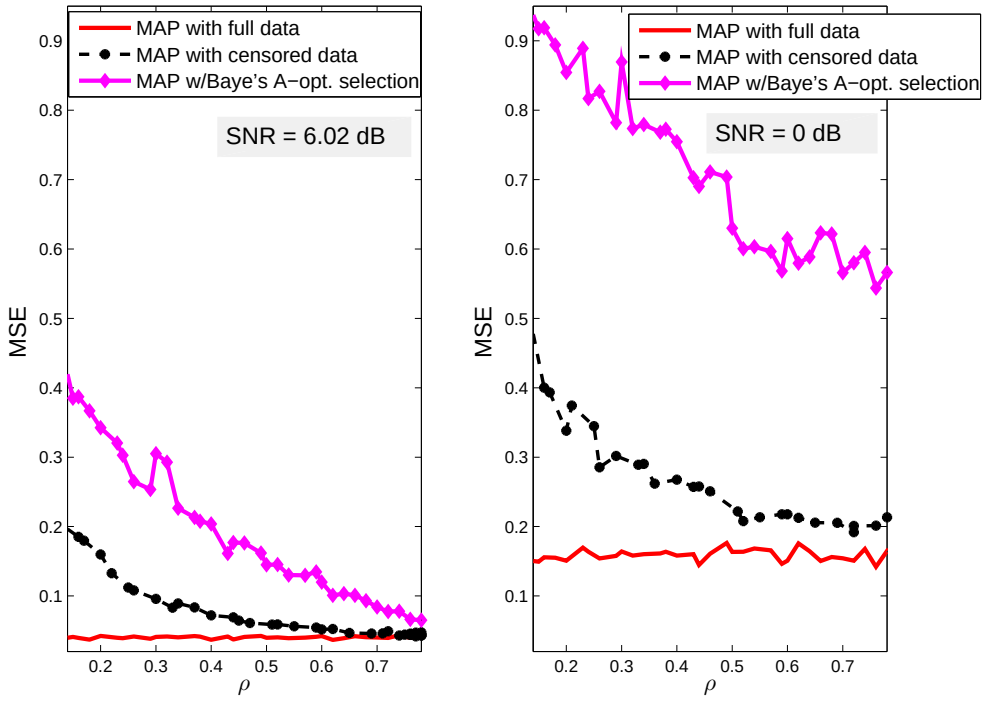

Figure 3.7: ρ vs. MSE: high and low SNR cases

3.4.3 Quantized-censored MSE performance

In the last simulation setting, Figs. 3.8 and 3.9 compares full-data MAP, cMAP, and qcMAP with $L = 2$, 4 and 6. Fig. 3.9 shows a magnified portion of curves in Fig. 3.8. It is seen that with $L = 4$ (corresponding to 2-bit quantization of uncensored data) the MSE of qcMAP and of the unquantized cMAP almost coincide. As is the case in the purely quantizer-estimator approach of [57], there is diminishing return in MSE gains with increase in the number of quantization levels. In fact, for this particular example, the difference in MSE performance of qcMLE for $L = 4$ and for $L = 6$ is not significant enough to justify the increased data-rate of a 6-level quantizer.

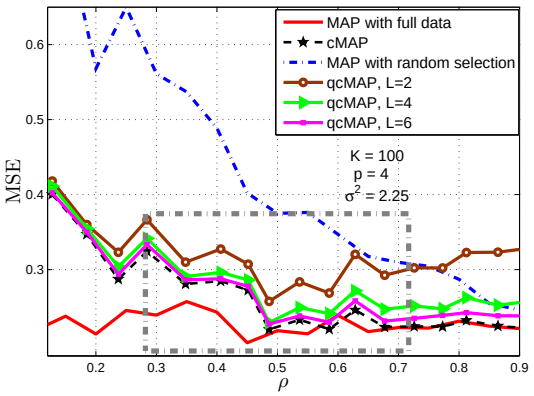

Figure 3.8: ρ vs. MSE: fixed SNR, qcMAP estimator.

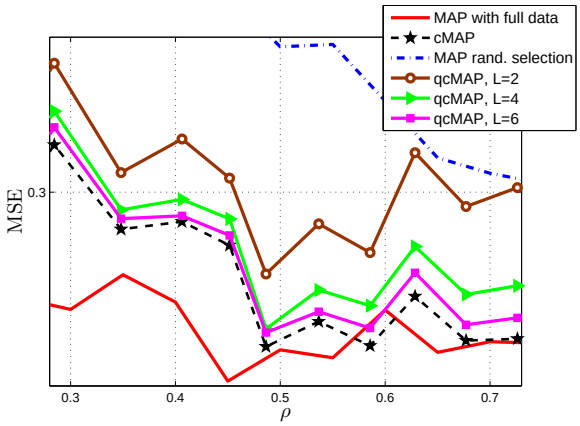

Figure 3.9: ρ vs. MSE: magnified part of Fig. 3.8.

3.5 Concluding Remarks

In this chapter, data selection for random parameter estimation was investigated. Extending the work of Chapter 2, measurement censoring and quantization as sensor-centric data selection techniques were studied for Bayesian maximum a posteriori estimation approach. Analogous to the deterministic parameter estimation case, it was shown that the computations involved in censoring are much lower than alternative centralized data-selection techniques. Starting with the censored-data model, an optimal fusion rule for MAP esti-

mation was derived which not only accounts for the uncensored measurements, but also for the underlying censoring technique used to acquire the measurements. Furthermore, it was shown that this approach leads to lower overall mean-square error than existing alternatives, in particular for identical sensor measurement distributions. Quantization of the uncensored measurements was also explored and showed that only a few quantization levels are necessary for estimator MSE performance to be comparable to that of the unquantized case. Explicit algorithmic formulations amenable to wireless sensor network implementation were also developed. In the next chapter, the challenging problem of state estimation using wireless sensor networks is studied as a natural extension of the Bayesian modeling studied in this chapter.

Chapter 4

Quantized State Estimation

4.1 Models and problem statement

Consider an ad-hoc WSN whose K sensor nodes $\{\mathcal{S}_k\}_{k=1}^K$ are deployed to estimate a multivariate discrete-time random process $\boldsymbol{\theta}(n) \in \mathbb{R}^p$. The state equation is given as

$$\boldsymbol{\theta}(n) = \mathbf{A}(n)\,\boldsymbol{\theta}(n-1) + \mathbf{u}(n) \tag{4.1}$$

where $\mathbf{A}(n) \in \mathbb{R}^{p\times p}$ denotes the state transition matrix and $\mathbf{u}(n)$ the driving noise, assumed to be a zero-mean white Gaussian process with covariance matrix $\mathbb{E}\{\mathbf{u}(n)\,\mathbf{u}^T(m)\} = \mathbf{C}_u(n)\delta_{nm}$.

Each sensor records scalar observations $y_k(n)$ adhering to a linear measurement equation

$$y_k(n) = \mathbf{h}_k^T(n)\,\boldsymbol{\theta}(n) + v_k(n) \tag{4.2}$$

where k is the sensor index, $\mathbf{h}_k(n) \in \mathbb{R}^p$ denotes the regression vector, and $v_k(n)$ is a temporally and spatially white zero-mean Gaussian noise process with covariance $\mathbb{E}\{v_k(n)\,v_l(m)\} = c_v(n)\delta_{nm}\delta_{kl}$. It is further assumed that $\mathbf{u}(n)$ is independent of both $v_k(n)$ and $\boldsymbol{\theta}(0)$.

Supposing that $\mathbf{A}(n)$, $\mathbf{C}_u(n)$, $\mathbf{h}_k(n)$ and $c_v(n)$ are available $\forall\ n, k$ from the physics of the problem, the goal of the WSN is for each sensor to form an estimate of $\boldsymbol{\theta}(n)$ to be used, e.g., in a habitat monitoring [47], or as a first step in a distributed control setup [31].

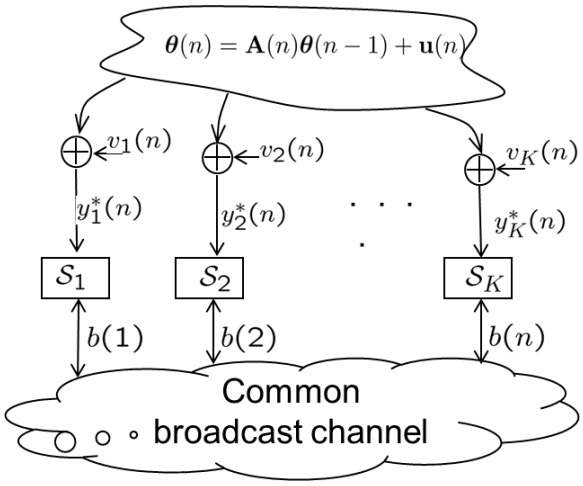

Figure 4.1: Sensor broadcast schedule in an *ad hoc* WSN setup.

Estimating $\boldsymbol{\theta}(n)$ necessitates each sensor $\mathcal{S}_k$ to communicate $y_k(n)$ to all other sensors $\{\mathcal{S}_l\}_{l=1,l\neq k}^{K}$. Communication takes place over the shared wireless channel that we assume can afford transmission of a single packet of m bits per time slot n. This leads to a one-to-one correspondence between time n and sensor index k and allows us to drop the sensor argument k in (4.2).

Each sensor in the network does both data gathering (sensing) and signal processing (estimation from broadcast data and local observations). The collection of cooperating sensor nodes forms an *ad hoc* network without a central fusion unit. Thus, signal processing at each node based on data exchanged among sensors as well as local observations constitutes the distributed filtering at hand.

The decision as to which sensor $\mathcal{S}_k = \mathcal{S}(n)$ broadcasts at time n, and consequently which observation $y_k(n) = y(n)$ is transmitted, depends on the underlying scheduling algorithm. Time-division multiple-access is assumed and for simplicity each sensor's broadcast footprint reaches the entire collaborating set of sensors. More elaborate sensor scheduling techniques - see e.g., [23, 50] and references therein - can be considered but are beyond the scope of this thesis.

In Fig. 4.1 a depiction of the ad hoc WSN highlighting the round-robin broadcast schedule for each sensor node is presented.

4.1.1 MMSE estimation with quantized observations

In order to effect digital inter-sensor communication in the bandwidth-limited WSN, the observations $y(n) \in \mathbb{R}$ are quantized. With $\mathcal{B}$ denoting a finite set of quantization messages[1], we investigate quantization rules $\mathbf{q}_n[\cdot]$ of the form

$$b(n) := \mathbf{q}_n[y(n)], \quad \text{where} \quad \mathbf{q}_n : \mathbb{R} \to \mathcal{B}. \tag{4.3}$$

Given current and past messages $\mathbf{b}_{1:n} := \{b(1), b(2), \ldots, b(n)\}$, we are interested in developing estimates $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ of the state $\boldsymbol{\theta}(n)$. The error covariance matrix (ECM) of the estimator is defined as $\mathbf{M}(n|\mathbf{b}_{1:n}) := \mathbb{E}\{[\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) - \boldsymbol{\theta}(n)][\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) - \boldsymbol{\theta}(n)]^T\}$, and the MSE of $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ is given by the trace of the ECM, i.e., $\mathrm{tr}\{\mathbf{M}(n|\mathbf{b}_{1:n})\}$. As is well known, see e.g., [49, Chapter 5], the wanted minimum MSE (MMSE) estimator is given by the conditional mean

$$\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) := \mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}\} = \int_{\mathbb{R}^p} \boldsymbol{\theta}(n)\, p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}]\, d\boldsymbol{\theta}(n)\,. \tag{4.4}$$

To obtain a closed-form expression for $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$, the posterior distribution $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}]$ has to be known and the integral in (4.4) needs to be computable. In principle, $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}]$ can be obtained from the state-observation model in (4.1)-(4.2) using the prediction-correction steps [P1]-[C1] outlined next.

[P1] Prediction step. With $p[\boldsymbol{\theta}(n-1)|\mathbf{b}_{1:n-1}]$ known, the prior pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}]$ follows from the theorem of total probability

$$p\,[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}] = \int_{\mathbb{R}^p} p[\boldsymbol{\theta}(n)|\boldsymbol{\theta}(n-1), \mathbf{b}_{1:n-1}] p[\boldsymbol{\theta}(n-1)|\mathbf{b}_{1:n-1}] d\boldsymbol{\theta}(n-1)\,. \tag{4.5}$$

Note that conditioning on $\boldsymbol{\theta}(n-1)$ in $p[\boldsymbol{\theta}(n)|\boldsymbol{\theta}(n-1), \mathbf{b}_{1:n-1}]$, renders conditioning on $\mathbf{b}_{1:n-1}$ redundant [cf. (4.1)]; consequently $p[\boldsymbol{\theta}(n)|\boldsymbol{\theta}(n-1), \mathbf{b}_{1:n-1}] = p[\boldsymbol{\theta}(n)|\boldsymbol{\theta}(n-1)] = \mathcal{N}[\boldsymbol{\theta}(n); \mathbf{A}(n)\boldsymbol{\theta}(n-1), \mathbf{C}_u(n)]$. The pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}]$ can be used to predict the state $\boldsymbol{\theta}(n)$ using past messages $\mathbf{b}_{1:n-1}$ and the state evolution model (4.1).

[1] These messages are represented by the set $\{1, 2, \ldots, |\mathcal{B}|\}$, where each quantization message, i.e., binary codeword $b(n) \in \mathcal{B}$, can be represented by $\log_2|\mathcal{B}|$ bits.

[C1] Correction step. When new observations become available, we write $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}] = p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}, b(n)]$ using Bayes' rule as

$$p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}] = p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}] \frac{\Pr\{b(n)|\boldsymbol{\theta}(n), \mathbf{b}_{1:n-1}\}}{\Pr\{b(n)|\mathbf{b}_{1:n-1}\}} \tag{4.6}$$

where the probability $\Pr\{b(n)|\boldsymbol{\theta}(n), \mathbf{b}_{1:n-1}\}$ can be obtained using the quantization rule $\mathbf{q}_n[y(n)]$ in (4.3). Likewise, $\Pr\{b(n)|\mathbf{b}_{1:n-1}\}$ depends on $\mathbf{q}_n[y(n)]$ and the prediction pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}]$; in Sections 4.2 and 4.3 specific quantizer structures and resulting pdfs are explored.

The Kalman filter (KF) is a recursive algorithm of the form [P1]-[C1]. In the framework of this work KF corresponds to the transmission of un-quantized messages $b(n) = y(n)$, i.e., a clairvoyant scenario without bandwidth constraints. In this case, the pdfs in (4.5) and (4.6) are Gaussian and it suffices to propagate their means and covariance matrices leading to the Kalman recursions [P2]-[C2].

[P2] Prediction step. Consider the predicted estimate $\hat{\boldsymbol{\theta}}(n|\mathbf{y}_{1:n-1}) := \mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{y}_{1:n-1}\}$ and let $\mathbf{M}(n|\mathbf{y}_{1:n-1}) := \mathbb{E}\left\{[\hat{\boldsymbol{\theta}}(n|\mathbf{y}_{1:n-1}) - \boldsymbol{\theta}(n)][\hat{\boldsymbol{\theta}}(n|\mathbf{y}_{1:n-1}) - \boldsymbol{\theta}(n)]^T\right\}$ denote the corresponding ECM. Given the previous estimate $\hat{\boldsymbol{\theta}}(n\text{–}1|\mathbf{y}_{1:n-1})$ and its ECM $\mathbf{M}(n\text{–}1|\mathbf{y}_{1:n-1})$, we have

$$\hat{\boldsymbol{\theta}}(n|\mathbf{y}_{1:n-1}) = \mathbf{A}(n)\hat{\boldsymbol{\theta}}(n-1|\mathbf{y}_{1:n-1}) \tag{4.7}$$

$$\mathbf{M}(n|\mathbf{y}_{1:n-1}) = \mathbf{A}(n)\mathbf{M}(n-1|\mathbf{y}_{1:n-1})\mathbf{A}^T(n) + \mathbf{C}_u(n)\,. \tag{4.8}$$

[C2] KF correction step. Consider the predicted data $\hat{y}(n|\mathbf{y}_{1:n-1}) := \mathbb{E}\{y(n)|\mathbf{y}_{1:n-1}\} = \mathbf{h}^T(n)\hat{\boldsymbol{\theta}}(n|\mathbf{y}_{1:n-1})$ and their innovation $\tilde{y}(n|\mathbf{y}_{1:n-1}) := y(n) - \hat{y}(n|\mathbf{y}_{1:n-1})$. Then $\hat{\boldsymbol{\theta}}(n|\mathbf{y}_{1:n})$ in (4.4) and its ECM obey

$$\hat{\boldsymbol{\theta}}(n|\mathbf{y}_{1:n}) = \hat{\boldsymbol{\theta}}(n|\mathbf{y}_{1:n\text{–}1}) + \frac{\mathbf{M}(n|\mathbf{y}_{1:n\text{–}1})\mathbf{h}(n)}{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{y}_{1:n\text{–}1})\mathbf{h}(n)+c_v(n)}\,\tilde{y}(n|\mathbf{y}_{1:n\text{–}1}) \tag{4.9}$$

$$\mathbf{M}(n|\mathbf{y}_{1:n}) = \mathbf{M}(n|\mathbf{y}_{1:n-1}) - \frac{\mathbf{M}(n|\mathbf{y}_{1:n-1})\mathbf{h}(n)\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{y}_{1:n-1})}{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{y}_{1:n-1})\mathbf{h}(n) + c_v(n)}\,. \tag{4.10}$$

Computations for the KF iterations in [P2]-[C2] are simpler than for the general iteration [P1]-[C1] with quantized observations. Indeed, while [P2]-[C2] requires a few algebraic operations per time-step n, [P1]-[C1] requires: (i) numerical integration to obtain the prediction

pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}]$ in (4.5); (ii) numerical update of the posterior pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}]$ using Bayes' rule in (4.6); and, (iii) numerical integration to evaluate the expectation in (4.4) and obtain $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$.

This high computational cost is inherent to non-linear models thus motivating approximate filters e.g., the extended KF, the unscented KF, and the particle filter. (Non-linearity in this work is due to quantization of the observations.) An alternative workaround is by Gaussian approximation of the prior pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}]$, see e.g., [37]. This simplifies coping with the potentially intractable pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}]$ by tracking its mean and covariance matrix which results in the following iteration [P3]-[C3]:

[P3] Prediction step. Define the predicted estimate $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}\}$ and the corresponding ECM $\mathbf{M}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{[\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) - \boldsymbol{\theta}(n)][\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) - \boldsymbol{\theta}(n)]^T\}$. Given the previous estimate $\hat{\boldsymbol{\theta}}(n-1|\mathbf{b}_{1:n-1})$ and its ECM $\mathbf{M}(n-1|\mathbf{b}_{1:n-1})$, linearity of the expectation operator yields

$$\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) = \mathbf{A}(n)\hat{\boldsymbol{\theta}}(n-1|\mathbf{b}_{1:n-1}) \tag{4.11}$$

$$\mathbf{M}(n|\mathbf{b}_{1:n-1}) = \mathbf{A}(n)\mathbf{M}(n-1|\mathbf{b}_{1:n-1})\mathbf{A}^T(n) + \mathbf{C}_u(n)\,. \tag{4.12}$$

[C3] Correction step. Adopt the approximation $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}] = \mathcal{N}[\boldsymbol{\theta}(n); \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}), \mathbf{M}(n|\mathbf{b}_{1:n-1})]$. As in [C1], write $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}] = p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}, b(n)]$ and use Bayes' rule to obtain

$$\begin{aligned} p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}] =& \mathcal{N}[\boldsymbol{\theta}(n); \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}), \mathbf{M}(n|\mathbf{b}_{1:n-1})] \\ &\times \frac{\Pr\{b(n)|\boldsymbol{\theta}(n), \mathbf{b}_{1:n-1})\}}{\Pr\{b(n)|\mathbf{b}_{1:n-1})\}}. \end{aligned} \tag{4.13}$$

Numerator and denominator of the second term in (4.13) can be found using the approximate prior pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}] = \mathcal{N}[\boldsymbol{\theta}(n); \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}), \mathbf{M}(n|\mathbf{b}_{1:n-1})]$ - see [67, Remark 3] for justification of the Gaussian approximation. Estimator $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ is then obtained by evaluating the integral in (4.4).

The computational cost incurred by [P3]-[C3] is between that of [P1]-[C1] and [P2]-[C2]. To obtain $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ using [P3]-[C3] we do not need to evaluate the integral in (4.5) but we still need to apply Bayes' rule in (4.13) and evaluate the integral in (4.4). The KF iteration

[P2]-[C2] on the other hand, evaluates the exact MMSE with small computational cost. For our purposes it represents a clairvoyant benchmark in terms of MSE performance and computational cost.

In this thesis, batch (Section 4.2) and iterative (Section 4.3) joint quantization-estimation approaches are pursued. In each case our objectives are : (i) to show that the approximate MMSE estimation [P3]-[C3] can be further simplified yielding a filter with comparable computational cost to the KF; and (ii) to compare the MSE of [P3]-[C3] using quantized observations, with the MSE of the KF in [P2]-[C2]. We will further contend that the MSE penalty associated with state estimates based on quantized data is small even with coarse quantization down to a few bits per observation.

4.2 Kalman filtering with batch quantized observations

In this section, each observation $y(n)$ is quantized by partitioning the observation space $\mathbb{R}$ into N intervals $\mathcal{R}_i := \big[\tau_i(n), \tau_{i+1}(n)\big)$, where $i \in \mathcal{B} := \{1, \ldots, N\}$. The quantizer $\mathbf{q}[\cdot]$ is thus specified through the thresholds $\{\tau_i(n)\}_{i=1}^{N+1}$, where $\tau_1(n) = -\infty$, $\tau_i(n) < \tau_{i+1}(n)$, and $\tau_{N+1}(n) = +\infty$. Consider the estimate $\hat{y}(n|\mathbf{b}_{1:n-1}) = \mathbb{E}\{y(n)|\mathbf{b}_{1:n-1}\}$, the corresponding innovation sequence $\tilde{y}(n|\mathbf{b}_{1:n-1}) := y(n) - \hat{y}(n|\mathbf{b}_{1:n-1})$, and the quantization rule

$$b(n) = i, \qquad \text{iff} \quad \tilde{y}(n|\mathbf{b}_{1:n-1}) \in \big[\tau_i(n),\ \tau_{i+1}(n)\big)\,. \tag{4.14}$$

In order to compute $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ [cf. (4.4)], the pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}]$ is needed. From (4.13), the distributions $\Pr\{b(n) = i|\mathbf{b}_{1:n-1}\}$ and $\Pr\{b(n) = i|\boldsymbol{\theta}(n), \mathbf{b}_{1:n-1}\}$, characterized in the following equations, need to be known. Using the quantization rule in (4.14) the events $\{b(n) = i\}$ and $\{\tau_i(n) \leq \tilde{y}(n|\mathbf{b}_{1:n-1}) < \tau_{i+1}(n)\}$ are equivalent and consequently

$$\Pr\{b(n) = i|\boldsymbol{\theta}(n), \mathbf{b}_{1:n-1}\} = \Pr\left\{\tau_i(n) \leq \tilde{y}(n|\mathbf{b}_{1:n-1}) < \tau_{i+1}(n)|\boldsymbol{\theta}(n), \mathbf{b}_{1:n-1}\right\}. \tag{4.15}$$

Given $\{\boldsymbol{\theta}(n), \mathbf{b}_{1:n-1}\}$, the innovation $\tilde{y}(n|\mathbf{b}_{1:n-1}) := y(n) - \hat{y}(n|\mathbf{b}_{1:n-1}) = \mathbf{h}^T(n)[\boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})] + v(n)$ has conditional pdf

$$p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\boldsymbol{\theta}(n), \mathbf{b}_{1:n-1}] = \mathcal{N}\big[\tilde{y}(n|\mathbf{b}_{1:n-1}); \mathbf{h}^T(n)\tilde{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}), c_v(n)\big]$$

, where $\tilde{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) := \boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})$. Using this pdf we can rewrite (4.15) in terms of the Gaussian tail function $Q(\cdot)$ as

$$\begin{aligned}\Pr\{b(n)=i|\boldsymbol{\theta}(n),\mathbf{b}_{1:n-1}\} \qquad (4.16)\\ = Q\left(\frac{\tau_i(n)-\mathbf{h}^T(n)\tilde{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})}{\sqrt{c_v(n)}}\right) - Q\left(\frac{\tau_{i+1}(n)-\mathbf{h}^T(n)\tilde{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})}{\sqrt{c_v(n)}}\right).\end{aligned}$$

Likewise, $\Pr\{b(n)=i|\mathbf{b}_{1:n-1}\} = \Pr\{\tau_i(n) \le \tilde{y}(n|\mathbf{b}_{1:n-1}) < \tau_{i+1}(n)|\mathbf{b}_{1:n-1}\}$ which is identical to (4.15) except for the conditioning variables. Note that, unlike (4.15), the conditional pdf $p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}]$ is non-Gaussian. It can though be approximated by using a Gaussian prior for $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}]$ as follows.

If $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}] = \mathcal{N}[\boldsymbol{\theta}(n); \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}), \mathbf{M}(n|\mathbf{b}_{1:n-1})]$, then the prior pdf $p[y(n)|\mathbf{b}_{1:n-1}]$ is also normal with mean $\hat{y}(n|\mathbf{b}_{1:n-1})$ and variance $\sigma_y^2(n) := \mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)$ [cf. (4.2)]. Since the innovation is defined as $\tilde{y}(n|\mathbf{b}_{1:n-1}) := y(n) - \hat{y}(n|\mathbf{b}_{1:n-1})$, we have $p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}] = \mathcal{N}[\tilde{y}(n|\mathbf{b}_{1:n-1});\, 0, \sigma_y^2(n)]$ and we can thus write

$$\begin{aligned}\Pr\{b(n)=i|\mathbf{b}_{1:n-1}\} &= Q[\tau_i(n)/\sigma_y(n)] - Q[\tau_{i+1}(n)/\sigma_y(n)]\\ &= Q[\Delta_i(n)] - Q[\Delta_{i+1}(n)] \qquad (4.17)\end{aligned}$$

where $\Delta_i(n) := \tau_i(n)/\sigma_y(n)$ are thresholds normalized by the standard deviation of the observations.

Substituting (4.16) and (4.17) into (4.13) we obtain an expression for $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n}]$ that can be used in (4.4) to obtain $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$. It is remarkable that with the Gaussian assumption for $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}]$, (4.4) can be found analytically as we show in the next proposition.

Proposition 4.1 *Consider the state space model (4.1)-(4.2) and quantized observations $b(n)$ defined as in (4.14). Suppose that the predicted estimate $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})$ and corresponding ECM $\mathbf{M}(n|\mathbf{b}_{1:n-1})$ are given. If the prediction pdf is $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}] = \mathcal{N}[\boldsymbol{\theta}(n); \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}), \mathbf{M}(n|\mathbf{b}_{1:n-1})]$, the MMSE estimator $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ in (4.4) can be obtained as follows:*

[C4] Batch Quantized (BQ)KF correction step. *Define $\sigma_y^2(n) := \mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)$ and normalized thresholds $\Delta_i(n) := \tau_i(n)/\sigma_y(n)$. Furthermore, define in terms*

of the conditional mean and variance of the innovation $\tilde{y}(n|\mathbf{b}_{1:n-1})$ *given* $b(n) = i$ *the ratios*

$$\begin{aligned}\alpha_i(n) :=& \frac{\mathbb{E}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\}}{\sigma_y(n)} \\ =& \frac{1}{\sqrt{2\pi}} \frac{\exp\left[-\Delta_i^2(n)/2\right] - \exp\left[-\Delta_{i+1}^2(n)/2\right]}{Q[\Delta_i(n)] - Q[\Delta_{i+1}(n)]} \end{aligned} \tag{4.18}$$

$$\begin{aligned}\beta_i(n) :=& 1 - \frac{\mathrm{var}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\}}{\sigma_y^2(n)} \\ =& \alpha_i^2(n) - \frac{1}{\sqrt{2\pi}} \frac{\Delta_i(n)e^{-\Delta_i^2(n)/2} - \Delta_{i+1}(n)e^{-\Delta_{i+1}^2(n)/2}}{Q[\Delta_i(n)] - Q[\Delta_{i+1}(n)]}\,. \end{aligned} \tag{4.19}$$

Then the resultant estimate and its ECM are given by

$$\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) = \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) + \alpha_i(n) \frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)}{\sqrt{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}} \tag{4.20}$$

$$\mathbf{M}(n|\mathbf{b}_{1:n}) = \mathbf{M}(n|\mathbf{b}_{1:n-1}) - \beta_i(n) \frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})}{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}\,. \tag{4.21}$$

Proof: See Appendix C.1.

Proposition 4.1 suggests an algorithm for finding (approximate) MMSE state estimates using quantized observations. The resulting batch-quantized Kalman filter (BQKF) consists of recursive application of the prediction step [P3] and the correction step [C4] in a collaborating setting as follows. The active sensor at time n, i.e., one scheduled to sense $y(n)$ and broadcast a quantized version $b(n)$, is assumed to broadcast within reach of all collaborating sensors. Each sensor, as stated earlier, keeps track of $\mathbf{A}(n)$, $\mathbf{C}_u(n)$, $\mathbf{h}_k(n)$ and $c_v(n)$ and thus can run the predictor step [P3]. Upon receiving $b(n)$ from the active sensor, all sensors execute the update step [C4].

The BQKF retains a notable resemblance to the clairvoyant KF [P2]-[C2] which is based on analog-amplitude observations. In order to highlight the similarities and differences of the respective correction steps [C2] and [C4], we define $\tilde{\alpha}_i(n|\mathbf{b}_{1:n-1}) := \sqrt{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}\,\alpha_i(n)$, $\bar{\beta}(n) := \mathbb{E}\{\beta_i(n)|\mathbf{b}_{1:n-1}\}$, and $\bar{\mathbf{M}}(n|\mathbf{b}_{1:n}) :=$

$\mathbb{E}\{\mathbf{M}(n|\mathbf{b}_{1:n})|\mathbf{b}_{1:n-1}\}$. From (4.20) and (4.21), it follows that

$$\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) = \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) + \frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)}{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}\,\tilde{\alpha}_i(n|\mathbf{b}_{1:n-1})$$

$$\bar{\mathbf{M}}(n|\mathbf{b}_{1:n}) = \mathbf{M}(n|\mathbf{b}_{1:n-1}) - \bar{\beta}(n)\frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})}{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}$$

where $\tilde{\alpha}_i(n|\mathbf{b}_{1:n-1})$ is the BQKF counterpart of the innovation $\tilde{y}(n|\mathbf{y}_{1:n-1})$ in KF [cf. C2] and $\bar{\beta}(n)$ is obtained from (4.19) as follows:

$$\begin{aligned}\bar{\beta}(n) :&= \mathbb{E}_{b(n)}\big\{\beta_i(n)|\mathbf{b}_{1:n-1}\big\} \\ &= \sum_{i=1}^{N}\big[Q[\Delta_i(n)] - Q[\Delta_{i+1}(n)]\big]\beta_i(n) \\ &= \sum_{i=1}^{N}\big[Q[\Delta_i(n)] - Q[\Delta_{i+1}(n)]\big]\alpha_i^2(n). \end{aligned} \tag{4.22}$$

Remark 4.1 Comparison of the ECM corrections for the KF in (4.10) with those for the BQKF in (4.21) reveal that they are identical except for the $\beta_i(n)$ factor in (4.21). The similarity is quantified by defining the ECM reduction per correction step [cf. (4.21)]

$$\Delta\mathbf{M}(n) := \mathbf{M}(n|\mathbf{b}_{1:n-1}) - \mathbf{M}(n|\mathbf{b}_{1:n}) = \beta_i(n)\frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})}{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}. \tag{4.23}$$

If we use $y(n)$ instead of $b(n)$ in the correction step, the ECM reduction will be [cf. (4.10)]

$$\begin{aligned}\Delta\mathbf{M}^{\mathrm{K}}(n) &:= \mathbf{M}(n|\mathbf{b}_{1:n-1}) - \mathbf{M}(n|\mathbf{b}_{1:n-1}, y(n)) \\ &= \frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})}{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}. \end{aligned} \tag{4.24}$$

Comparing (4.23) with (4.24) we see that $\Delta\mathbf{M}(n) = \beta_i(n)\Delta\mathbf{M}^{\mathrm{K}}(n)$; that is the ECM reduction achieved by the BQKF is $\beta_i(n)$ times that of the clairvoyant KF. From (4.19), $0 < \beta_i(n) < 1$ for any possible selection of thresholds, consistent with the fact that $b(n)$ is a coarse representation of $y(n)$.

4.2.1 Quantizer design for batch quantized KF

The ECM $\mathbf{M}(n|\mathbf{b}_{1:n})$ in (4.21), and consequently the variance reduction $\Delta\mathbf{M}(n)$ in (4.23), depend on $b(n)$. We define the optimal quantizer as the one that maximizes the average

variance reduction, i.e.,

$$\{\Delta_i^*(n)\}_{i=2}^N := \max_{\{\Delta_i(n)\}_{i=2}^N} \mathbb{E}_{b(n)}\big\{\Delta\mathbf{M}(n)|\mathbf{b}_{1:n-1}\big\}$$
$$= \arg \max_{\{\Delta_i(n)\}_{i=2}^N} \mathbb{E}_{b(n)}\big\{\beta_i(n)|\mathbf{b}_{1:n-1}\big\} \tag{4.25}$$

where $\mathbb{E}_{b(n)}\{\cdot|\mathbf{b}_{1:n-1}\}$ denotes expectation with respect to $\Pr\{b(n)|\mathbf{b}_{1:n-1}\}$; and in establishing the second equality we used the fact that given $\mathbf{b}_{1:n-1}$, $\mathbf{M}(n|\mathbf{b}_{1:n-1})$ is conditionally independent of $b(n)$ [cf. (4.12)]. The last expectation in (4.25) can be evaluated by substituting $\alpha_i(n)$ from (4.18) in (4.22) leading to

$$\bar{\beta}(n) := \mathbb{E}_{b(n)}\big\{\beta_i(n)|\mathbf{b}_{1:n-1}\big\}$$
$$= \frac{1}{2\pi}\sum_{i=1}^{N} \frac{\big[\exp(-\Delta_i^2(n)/2) - \exp(-\Delta_{i+1}^2(n)/2)\big]^2}{Q[\Delta_i(n)] - Q[\Delta_{i+1}(n)]}\,. \tag{4.26}$$

The optimal thresholds $\{\Delta_i^*(n)\}_{i=2}^N$ in (4.25) can be obtained as the maximizers of $\bar{\beta}(n)$ as detailed next in Proposition 4.2. An alternative approach to the optimal quantization is obtained by using the definition (4.19), $\sigma_y^2(n)[1-\beta_i(n)] = \mathrm{var}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\}$. Proposition 4.2 shows that the thresholds $\{\Delta_i^*(n)\}_{i=2}^N$ in (4.25) define the optimal quantizer of the innovation $\tilde{y}(n|\mathbf{b}_{1:n-1})$ with an MSE distortion.

Proposition 4.2 *Consider the problem of optimal quantization of the innovation* $\tilde{y}(n|\mathbf{b}_{1:n-1})$. *If* $b(n) = i$, *the reconstructed innovation is* $\hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n) = i\} = \sigma_y(n)\,\alpha_i(n)$ *[cf. (4.18)]. We adopt an MSE distortion conditioned on* $\mathbf{b}_{1:n-1}$ *and define the optimal quantizer of* $\tilde{y}(n|\mathbf{b}_{1:n-1})$ *as*

$$\{\Delta_i^\dagger(n)\}_{i=2}^N := \arg \min_{\{\Delta_i(n)\}_{i=2}^N} \mathrm{d}[\tilde{y}(n|\mathbf{b}_{1:n-1}), \hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n-1})]$$
$$:= \arg \min_{\{\Delta_i(n)\}_{i=2}^N} \mathbb{E}\{[\tilde{y}(n|\mathbf{b}_{1:n-1}) - \hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n-1})]^2|\mathbf{b}_{1:n-1}\}. \tag{4.27}$$

The optimal thresholds in (4.25) and (4.27) are equal, i.e., $\{\Delta_i^*(n)\}_{i=2}^N = \{\Delta_i^\dagger(n)\}_{i=2}^N$.

Proof: Since from definition (4.19) $\mathrm{var}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\} = \sigma_y^2(n)\,[1-\beta_i(n)]$,

$$\mathbb{E}_{b(n)}\big\{\mathrm{var}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)\}|\mathbf{b}_{1:n-1}\big\} = \sigma_y^2(n)\big[1 - \mathbb{E}_{b(n)}\big\{\beta(n)|\mathbf{b}_{1:n-1}\big\}\big]$$

it follows that the thresholds that maximize $\bar{\beta}(n)$ can be obtained from the thresholds that minimize the expectation of the conditional variance, i.e.,

$$\begin{aligned} &\arg \max_{\{\Delta_i(n)\}_{i=2}^N} \mathbb{E}_{b(n)}\big\{\beta(n)|\mathbf{b}_{1:n-1}\big\} \\ &= \arg \min_{\{\Delta_i(n)\}_{i=2}^N} \mathbb{E}_{b(n)}\big\{\mathrm{var}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\}|\mathbf{b}_{1:n-1}\big\}. \end{aligned} \tag{4.28}$$

It remains only to show that the minimization in the second expression of (4.28) is the conditional MSE of $\hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n-1})$. Writing the variance explicitly with $\mathcal{R}_i := \big[\tau_i(n), \tau_{i+1}(n)\big)$, we have

$$\begin{aligned} &\mathrm{var}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\} \\ &= \int_{\mathcal{R}_i} [\tilde{y}(n|\mathbf{b}_{1:n-1}) - \mathbb{E}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\}]^2 \\ &\quad \times p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i]\, d\tilde{y}(n|\mathbf{b}_{1:n-1}). \end{aligned} \tag{4.29}$$

Given that $\mathbb{E}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\} = \hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n})$ the conditional expectation of (4.29) is written as

$$\begin{aligned} &\mathbb{E}_{b(n)}\big\{\mathrm{var}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\}|\mathbf{b}_{1:n-1}\big\} \\ &= \sum_{i=1}^{N} \int_{\mathcal{R}_i} \big[\tilde{y}(n|\mathbf{b}_{1:n-1}) - \hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n-1})\big]^2 \\ &\qquad \times p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i] d\tilde{y}(n|\mathbf{b}_{1:n-1}) \Pr\{b(n)=i|\mathbf{b}_{1:n-1}\}. \end{aligned} \tag{4.30}$$

Since $p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i]=0$ for $\tilde{y}(n|\mathbf{b}_{1:n-1}) \notin \mathcal{R}_i$, then $p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i]\Pr\{b(n)=i|\mathbf{b}_{1:n-1}\} = p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}]$ if $\tilde{y}(n|\mathbf{b}_{1:n-1}) \in \mathcal{R}_i$, and zero otherwise. Thus, (4.30) can be written as

$$\begin{aligned} &\mathbb{E}_{b(n)}\big\{\mathrm{var}\{\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}, b(n)=i\}|\mathbf{b}_{1:n-1}\big\} \\ &\qquad = \sum_{i=1}^{N} \int_{\mathcal{R}_i} [\tilde{y}(n|\mathbf{b}_{1:n-1}) - \hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n-1})]^2 p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}] d\tilde{y}(n|\mathbf{b}_{1:n-1}) \\ &\qquad = \mathrm{d}[\tilde{y}(n|\mathbf{b}_{1:n-1}), \hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n-1})] \end{aligned}$$

where the second equality comes from the definition of the distortion metric in (4.27). Since the optimization arguments in (4.25) and (4.27) are equal the proof follows. □

Table 4.1: Quantization thresholds for Gaussian pdf, $\Delta_i = -\Delta_{N+2-i},\ \forall\, i \in \{1, 2, \ldots, N\}$

N	Δ_2	Δ_3
2	0	∞

N	Δ_3	Δ_4	Δ_5
4	0	0.982	∞

N	Δ_5	Δ_6	Δ_7	Δ_8	Δ_9
8	0	0.501	1.050	1.748	∞

N	Δ_9	Δ_{10}	Δ_{11}	Δ_{12}	Δ_{13}	Δ_{14}	Δ_{15}	Δ_{16}	Δ_{17}
16	0	0.258	0.522	0.800	1.099	1.437	1.844	2.401	∞

Table 4.2: $\bar{\beta}(n)$ values for batch quantization

N ($\log_2(N)$)	2 (1 bit)	4 (2 bits)	8 (3 bits)	16 (4 bits)
$\bar{\beta}(n)$	0.637=$2/\pi$	0.883	0.966	0.991

We emphasize the subtle difference between the thresholds $\{\Delta_i^*(n)\}_{i=2}^N$ in (4.25) and $\{\Delta_i^\dagger(n)\}_{i=2}^N$ in (4.27). While the former minimizes the MSE of the state estimates $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$, the latter minimizes the reconstruction error when estimating the innovation $\tilde{y}(n|\mathbf{b}_{1:n-1})$ by $\hat{\tilde{y}}^{(i)}(n|\mathbf{b}_{1:n-1})$. Even though these two criteria are different, the corresponding optimal thresholds coincide. Indeed, Proposition 4.2 asserts that the optimal strategy for estimating $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ is to quantize $\tilde{y}(n|\mathbf{b}_{1:n-1})$ with minimum MSE distortion.

The optimization problem in (4.27) has a well known solution given by the Lloyd-Max quantizer [42, 48]. Since $p[\tilde{y}(n|\mathbf{b}_{1:n-1})|\mathbf{b}_{1:n-1}] \approx \mathcal{N}[\tilde{y}(n|\mathbf{b}_{1:n-1}); 0, \sigma_{y_n}^2]$, the use of normalized thresholds $\Delta_i := \tau_i(n)/\sigma_{y_n}$ leads to quantization of $\tilde{y}(n|\mathbf{b}_{1:n-1})/\sigma_{y_n}$ which corresponds to quantization of zero-mean, unit variance Gaussian random variables. The optimal normalized thresholds values for $N = 2$, $N = 4$, $N = 8$, and $N = 16$, from [48], are given in Table 4.1 and the corresponding $\bar{\beta}(n)$ obtained using (4.26) are summarized in Table 4.2. We can see that even quantizing to a single bit has $\bar{\beta}(n) = 2/\pi \approx 0.64$ and quantizing to more than 4 bits, for which $\bar{\beta}(n) \approx 0.99$, seems rather unjustified.

4.2.2 Binary quantized Kalman filter (1-QKF)

A particularly interesting case results when $N = 2$, which amounts to binary quantization. In this case there is a single threshold $\tau(n) := \tau_2(n)$ to be determined (since $\tau_1(n) = -\infty$ and $\tau_3(n) = +\infty$), and we have $\mathcal{R}_1 = \big(-\infty, \tau(n)\big)$ and $\mathcal{R}_2 = \big[\tau(n), +\infty\big)$. Upon defining $\Delta(n) := \tau(n)/[\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)]^{1/2}$ the variables $\alpha_{1,2}(n)$ and $\beta_{1,2}(n)$ take on simpler expressions given by [cf. (4.18) and (4.19)]

$$\alpha_{1,2}(n) = \mp \frac{1}{\sqrt{2\pi}} \frac{\exp\big[-\Delta^2(n)/2\big]}{Q[\mp\Delta(n)]}$$

$$\beta_{1,2}(n) = \alpha_{1,2}^2(n) \pm \frac{1}{\sqrt{2\pi}} \frac{\Delta(n)\exp\big[-\Delta^2(n)/2\big]}{Q[\mp\Delta(n)]}$$

where $\mp$ signifies the use of minus for $b(n) = 1$ and plus for $b(n) = 2$ (and vice versa for $\pm$). Interestingly, the expected performance factor simplifies to [cf. (4.26)]

$$\begin{aligned}\bar{\beta}(n) &= \frac{1}{2\pi}\left[\frac{\exp^2\big[-\Delta^2(n)/2\big]}{Q[\Delta(n)]} + \frac{\exp^2\big[-\Delta^2(n)/2\big]}{Q[-\Delta(n)]}\right] \\ &= \frac{1}{2\pi}\frac{\exp[-\Delta^2(n)]}{Q[\Delta(n)]Q[-\Delta(n)]}\,. \end{aligned} \tag{4.31}$$

Appendix shows that $\Delta(n) = 0$ maximizes $\bar{\beta}(n)$ in (4.31). With $\Delta(n) = 0$, $\bar{\beta}(n) = 2/\pi$, corroborating the result for $N = 2$ in Table 4.2. For the optimally selected thresholds $\tau(n) = 0$ it is convenient to change the notation and define the message, $b(n)$ in (4.14), as the sign of the innovation

$$b(n) = \text{sign}[\tilde{y}(n|\mathbf{b}_{1:n-1})] \tag{4.32}$$

which is equivalent to $b(n) = -1$ if $\tilde{y}(n|\mathbf{b}_{1:n-1}) \in \mathcal{R}_1$ and $b(n) = +1$ if $\tilde{y}(n|\mathbf{b}_{1:n-1}) \in \mathcal{R}_2$. In this case, the correction step [C4] takes on the rather simple form presented next.

[C5] 1-bit quantized KF (1-QKF) correction step. The state estimate and ECM are given by

$$\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) = \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) + \sqrt{\frac{2}{\pi}}\frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)}{\sqrt{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}}b(n) \tag{4.33}$$

$$\mathbf{M}(n|\mathbf{b}_{1:n}) = \mathbf{M}(n|\mathbf{b}_{1:n-1}) - \frac{2}{\pi}\frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})}{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}. \tag{4.34}$$

The iteration [P3]-[C5] is the Sign of Innovations KF (SoI-KF) introduced in [67]. The simplicity of [C5] suggests an alternative approach for iterative multi-bit quantization which is pursued in the next section.

4.3 Kalman filtering with iteratively quantized observations

In this section's iteratively quantized Kalman filter (IQKF), sensors rely on m-bit binary messages $\mathbf{b}(n) := \mathbf{b}^{(1:m)}(n) := [b^{(1)}(n), \ldots, b^{(m)}(n)]$, with the i-th bit $b^{(i)}(n)$ defined as the sign of innovations [cf. (4.36)] conditioned on the previous messages $\mathbf{b}_{1:n-1} := [\mathbf{b}(1), \ldots, \mathbf{b}(n-1)]$ and previous bits $\mathbf{b}^{(1:i-1)}(n) := [b^{(1)}(n), \ldots, b^{(i-1)}(n)]$ of the current (n-th) message. Specifically, let

$$\begin{aligned} \hat{y}^{(0)}(n|\mathbf{b}_{1:n-1}) &:= \hat{y}(n|\mathbf{b}_{1:n-1}) = \mathbb{E}\left\{y(n)|\mathbf{b}_{1:n-1}\right\} \\ \hat{y}^{(i)}(n|\mathbf{b}_{1:n-1}) &:= \mathbb{E}\left\{y(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\}, \quad \text{for } i \geq 1 \end{aligned} \tag{4.35}$$

stand for MMSE estimates of $y(n)$ using past messages $\mathbf{b}_{1:n-1}$ and the first i bits of the current message denoted as $\mathbf{b}^{(1:i)}(n)$. The i-th bit of the current message, $b^{(i)}(n)$, is obtained as

$$\begin{aligned} b^{(i)}(n) &:= \text{sign}\left[y(n) - \hat{y}^{(i-1)}(n|\mathbf{b}_{1:n-1})\right] \\ &:= \text{sign}\left[\tilde{y}^{(i-1)}(n|\mathbf{b}_{1:n-1})\right]. \end{aligned} \tag{4.36}$$

Our goal here is to derive an iterative algorithm to obtain estimates

$$\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\left\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\} \tag{4.37}$$

based on $\hat{\boldsymbol{\theta}}^{(i-1)}(n|\mathbf{b}_{1:n-1})$ and $b^{(i)}(n)$. When using m bits, $\mathbf{b}^{(1:m)}(n)$, we will refer to the resulting algorithm as m-IQKF.

At this point we should note that the 1-IQKF coincides with the BQKF with $N = 2$ quantization regions and optimally selected threshold $\tau(n) = 0$, as in Section 4.2.2. Thus, when $m = 1$ the estimates in (4.37) can be obtained using the iteration [P2]-[C5]. Since the definition of $b^{(i)}(n)$ is a straightforward extension of the corresponding definition for

the 1-QKF [cf. (4.36) and (4.32)], one option for $i \geq 1$ would be to set $\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1})$ equal to $\mathbf{h}^T(n)\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1})$ and then apply iterative correction steps of the form [C5]. However, $\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1}) \neq \mathbf{h}^T(n)\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1})$ if $i \geq 1$ as explained next. Using the observations (4.2), the definitions of $\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1})$ in (4.35) and $\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1})$ in (4.37), we obtain

$$\begin{aligned}\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1}) &= \mathbb{E}\left\{\mathbf{h}^T(n)\boldsymbol{\theta}(n) + v(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\} \\ &= \mathbf{h}^T(n)\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1}) + \mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\}. \end{aligned} \tag{4.38}$$

The noise estimate $\mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\}$ is not necessarily zero for $i \geq 1$; see also (4.48). (The converse is true for the 1-IQKF where $\mathbb{E}\{v(n)|\mathbf{b}_{1:n-1}\} = \mathbb{E}\{v(n)\} = 0$.) Therefore, in order to obtain $\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1})$ we need $\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1})$ as well as $\mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\}$. In order to keep track of the noise estimate and its covariance, we augment the state with the noise as described in the next section.

4.3.1 State augmentation

Through state augmentation the estimates $\mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\}$ can be obtained so that the predicted observation $\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1})$ in (4.38) can be evaluated even for $i > 1$ (i.e., multi-bit quantization). Specifically, appending observation noise $v(n)$ to the state vector $\boldsymbol{\theta}(n)$ we construct an augmented state vector $\breve{\boldsymbol{\theta}}(n) := [\boldsymbol{\theta}^T(n), v(n)]^T$. Correspondingly, we define augmented driving noise $\breve{\mathbf{u}}(n) := [\mathbf{u}^T(n), v(n)]^T$, state propagation matrix $\breve{\mathbf{A}}(n) := \begin{pmatrix} \mathbf{A}(n) & \mathbf{0} \\ \mathbf{0}^T & 0 \end{pmatrix}$, and observation vector $\breve{\mathbf{h}}(n) := [\mathbf{h}^T(n), 1]^T$. The model in (4.1)-(4.2) can, consequently, be rewritten with the following augmented state and observation equations

$$\breve{\boldsymbol{\theta}}(n) = \breve{\mathbf{A}}(n)\,\breve{\boldsymbol{\theta}}(n-1) + \breve{\mathbf{u}}(n) \tag{4.39}$$

$$\breve{y}(n) = \breve{\mathbf{h}}^T(n)\,\breve{\boldsymbol{\theta}}(n) + \breve{v}(n) \tag{4.40}$$

where the new observation noise is $\breve{v}(n) = 0,\ \forall n$ (by construction) and can be thought of as Gaussian noise with variance $c_{\breve{v}}(n) = 0$. Note that the covariance matrix of the augmented driving noise is a block-diagonal matrix $\mathbf{C}_{\breve{\mathbf{u}}}(n) \in \mathbb{R}^{(p+1)\times(p+1)}$ with $[\mathbf{C}_{\breve{\mathbf{u}}}(n)]_{1:p,1:p} = \mathbf{C}_{\mathbf{u}}(n)$ and $[\mathbf{C}_{\breve{\mathbf{u}}}(n)]_{p+1,p+1} = c_v(n)$.

The augmented state formulation (4.39)-(4.40) increases the dimension of the state vector but is otherwise equivalent to (4.1)-(4.2). However, it has the appealing property that MMSE estimates of the augmented state $\breve{\boldsymbol{\theta}}(n)$ contain MMSE estimates of the original state $\boldsymbol{\theta}(n)$ and the observation noise $v(n)$. In particular, state-augmentation allows simple computation of $\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1})$ in (4.35) as detailed in the following lemma:

Lemma 4.1 *Consider the MMSE estimate* $\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1}) = \mathbb{E}\left\{\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\}$ *and the predicted augmented observation estimate* $\hat{\breve{y}}^{(i)}(n|\mathbf{b}_{1:n-1}) = \mathbb{E}\left\{\breve{y}(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\}$. *The predicted observation MMSE estimate* $\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1})$ *in (4.35) can be obtained as*

$$\hat{y}^{(i)}(n|\mathbf{b}_{1:n-1}) = \hat{\breve{y}}^{(i)}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{h}}^T(n)\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1}). \tag{4.41}$$

Proof: The observation $y(n)$ in (4.2) and the augmented observation $\breve{y}(n)$ in (4.40) are the same by construction. That is,

$$\begin{aligned}\breve{y}(n) &= \breve{\mathbf{h}}^T(n)\breve{\boldsymbol{\theta}}(n) + \breve{v}(n) = \breve{\mathbf{h}}^T(n)\breve{\boldsymbol{\theta}}(n) \\ &= \mathbf{h}^T(n)\boldsymbol{\theta}(n) + v(n) = y(n)\end{aligned} \tag{4.42}$$

where the first equality follows from (4.40), the second one from $\breve{v}(n) = 0$, the third one from the definitions of $\breve{\mathbf{h}}(n)$ and $\breve{\boldsymbol{\theta}}(n)$, and the last one from (4.2). Taking expectation with respect to $p[y(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)]$ in the first, third and fifth terms of the equality in (4.42) results in (4.41). □

Based on Lemma 4.1 we can now use the augmented state estimates to find the observation estimates. Combining this with the 1-QKF recursion [P3]-[C5], we obtain the MMSE estimates in (4.37) using the algorithm detailed in the next proposition.

Proposition 4.3 *Consider the augmented state space model in (4.39) and (4.40). Define the augmented state estimates* $\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}\}$ *and* $\hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\left\{\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i-1)}(n)\right\}$. *Let* $\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{[\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) - \breve{\boldsymbol{\theta}}(n)][\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) - \breve{\boldsymbol{\theta}}(n)]^T\}$ *and* $\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{[\hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) - \breve{\boldsymbol{\theta}}(n)][\hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) - \breve{\boldsymbol{\theta}}(n)]^T\}$ *denote the corresponding ECMs for* $i = 1, \ldots, m$. *And construct the messages* $\mathbf{b}^{(1:m)}(n)$ *in (4.36) as*

$$b^{(i)}(n) := \text{sign}\left[y(n) - \breve{\mathbf{h}}^T(n)\hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\right]. \tag{4.43}$$

Furthermore, find the estimate $\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1})$ *from the following recursion:*

[P4] *Given the previous estimate* $\hat{\breve{\boldsymbol{\theta}}}(n-1|\mathbf{b}_{1:n-1})$ *and its ECM* $\breve{\mathbf{M}}(n-1|\mathbf{b}_{1:n-1})$*, form*

$$\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{A}}(n)\hat{\breve{\boldsymbol{\theta}}}(n-1|\mathbf{b}_{1:n-1}) \tag{4.44}$$

$$\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{A}}(n)\breve{\mathbf{M}}(n-1|\mathbf{b}_{1:n-1})\breve{\mathbf{A}}^T(n) + \mathbf{C}_{\breve{\mathbf{u}}}(n)\,. \tag{4.45}$$

[C6] *Assuming that* $p[\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i-1)}(n)] = \mathcal{N}[\breve{\boldsymbol{\theta}}(n); \hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1}), \breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})]$*, the MMSE estimate* $\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\}$ *and the corresponding ECM* $\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n})$ *are obtained by iterative application of*

$$\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1}) = \hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) + \sqrt{\frac{2}{\pi}}\frac{\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}{\sqrt{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}}\, b^{(i)}(n) \tag{4.46}$$

$$\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) - \frac{2}{\pi}\frac{\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})}{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)} \tag{4.47}$$

where we used the definitions $\hat{\breve{\boldsymbol{\theta}}}^{(0)}(n|\mathbf{b}_{1:n-1}) := \hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1})$ *and* $\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1}) := \breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1})$*. For time index* n*, (4.46) and (4.47) are repeated* m*-times.*

The MMSE estimate of $\breve{\boldsymbol{\theta}}(n)$ *given* $\mathbf{b}_{1:n}$ *is*

$$\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n}) = \mathbb{E}\left\{\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n}\right\} = \mathbb{E}\left\{\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:m)}(n)\right\} = \hat{\breve{\boldsymbol{\theta}}}^{(m)}(n|\mathbf{b}_{1:n-1}).$$

The corresponding ECM is $\breve{\mathbf{M}}(n|\mathbf{b}_{1:n}) := \breve{\mathbf{M}}^{(m)}(n|\mathbf{b}_{1:n-1})$*.*

Proof: See Appendix C.2.

The state estimates $\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1})$ in (4.37) are the first p components of the augmented state estimate $\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1})$, i.e., $\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1}) = [\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1})]_{1:p}$ and the noise estimate is the $(p+1)$-st component $\mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\} = [\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1})]_{p+1}$.

Corollary 4.1 *For* $i = 1$*, it holds that* $\mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\} \neq 0$.

Proof: The $(p+1)$-st entry of $\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1})$ in (4.46) is the noise estimate $\mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right\}$. Thus, for $i=1$

$$\begin{aligned}\mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}, b(n)\right\} &= \mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}\right\} + \sqrt{\frac{2}{\pi}} \frac{[\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)]_{p+1}}{\sqrt{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}}\, b(n) \\ &= \sqrt{\frac{2}{\pi}} \frac{c_v(n)}{\sqrt{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)}}\, b(n) \neq 0\,. \end{aligned} \tag{4.48}$$

The last equality uses $\mathbb{E}\left\{v(n)|\mathbf{b}_{1:n-1}\right\} = \mathbb{E}\{v(n)\} = 0$ and $\left[\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\right]_{p+1} = c_v(n)$. □

The similarity of the m-IQKF in Proposition 4.3 with the clairvoyant KF based on un-quantized observations (cf. [P2]-[C2]) is even more remarkable than the similarity between BQKF and KF. As in the BQKF, the ECM updates in (4.47) are identical to the ECM updates of the KF except for the scale factor $2/\pi$, suggesting that the MSE penalty due to quantization is small. The ECMs of the m-IQKF iterates are independent of the message sequence $\mathbf{b}_{1:n}$, which is not the case for the BQKF. This property shared by the KF strengthens the variance performance claims of the m-IQKF. While the small variance penalty factors summarized in Table 4.2 are valid on average for BQKF, the variance penalty factors of the m-IQKF hold true for all message sequences $\mathbf{b}_{1:n}$. The variance penalties associated with m-IQKF are investigated in the following corollaries.

Corollary 4.2 *Consider the m-IQKF algorithm in Proposition 4.3 and define the ECM reduction after $i=1,\ldots,m$ iterations as $\Delta\breve{\mathbf{M}}_i(n) := \breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1}) - \breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1})$, where $\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1}) := \breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1})$. With $c_i := 1-(1-\frac{2}{\pi})^i$, the error covariance reduction after i iterations is given as*

$$\Delta\breve{\mathbf{M}}_i(n) = c_i \frac{\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})}{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}\,. \tag{4.49}$$

Proof: We first write $\Delta\breve{\mathbf{M}}_i(n)$ as a summation of differences between successive ECM

matrices

$$\begin{aligned}\Delta\breve{\mathbf{M}}_i(n) :=& \breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1}) - \breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1}) \\ =& \sum_{j=1}^{i}\big[\breve{\mathbf{M}}^{(j-1)}(n|\mathbf{b}_{1:n-1}) - \breve{\mathbf{M}}^{(j)}(n|\mathbf{b}_{1:n-1})\big] \\ =& \frac{2}{\pi}\sum_{j=1}^{i}\left[\frac{\breve{\mathbf{M}}^{(j-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(j-1)}(n|\mathbf{b}_{1:n-1})}{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(j-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}\right] \end{aligned} \tag{4.50}$$

where the last equality follows from (4.47). Next, we recursively obtain the product $\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)$ by multiplying the ECM in (4.47) by $\breve{\mathbf{h}}(n)$ to obtain

$$\begin{aligned}\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n) =& \breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n) \\ & - \frac{2}{\pi}\frac{\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)} \qquad (4.51) \\ =& \left(1-\frac{2}{\pi}\right)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n) \qquad (4.52)\end{aligned}$$

where we canceled out factors $\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)$ in the numerator and denominator of the second term in (4.51). Repeating steps (4.51)-(4.52) for $\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)$ with decreasing index i yields

$$\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n) = \left(1-\frac{2}{\pi}\right)^i \breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\,.$$

Thus, substituting $\breve{\mathbf{M}}^{(j-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n) = \big(1-\frac{2}{\pi}\big)^{j-1}\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)$ into (4.50) yields

$$\Delta\breve{\mathbf{M}}_i(n) = \frac{2}{\pi}\Big[\sum_{j=1}^{i}\left(1-\frac{2}{\pi}\right)^{j-1}\Big]\frac{\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})}{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}$$

and upon invoking the geometric sum identity $\frac{2}{\pi}\sum_{j=1}^{i}\big(1-\frac{2}{\pi}\big)^{j-1} = 1-(1-\frac{2}{\pi})^i$, (4.49) follows. □

For the clairvoyant KF based on un-quantized observations the ECM reduction $\Delta\mathbf{M}^{\mathrm{K}}(n)$ is given by (4.24), and can be seen to correspond to the reduction $\Delta\breve{\mathbf{M}}_i(n)$ in (4.49) with $c_m = 1$. Therefore, Corollary 4.2 asserts that the ECM reduction achieved by quantizing to m bits is c_m times smaller than the one achieved with analog-amplitude observations.

Values of c_m are shown in Table 4.3. With only 4 quantization bits, the value of c_m is just 2% less than the value for the clairvoyant KF ($c_m = 1$).

Using Corollary 4.2 we can relate the predicted MSE $\mathrm{tr}\{\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1})\}$ and the corrected MSE $\mathrm{tr}\{\breve{\mathbf{M}}(n|\mathbf{b}_{1:n})\}$ after observing the m bits of the n-th observation, $\mathbf{b}^{(1:m)}(n)$. Corollary 4.3 summarizes this result.

Corollary 4.3 *The predicted and corrected ECMs of the m-IQKF are related as*

$$\breve{\mathbf{M}}(n|\mathbf{b}_{1:n}) = \breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1}) - c_m \frac{\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1})}{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)} \tag{4.53}$$

with $\breve{\mathbf{M}}(n|\mathbf{b}_{1:n}) = \breve{\mathbf{M}}^{(m)}(n|\mathbf{b}_{1:n-1})$ *and* $\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})$.

Proof: Write $\breve{\mathbf{M}}(n|\mathbf{b}_{1:n}) = \breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1}) - \Delta\breve{\mathbf{M}}_m(n)$ and use Corollary 4.2. Recall that by definition $\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{M}}^{(0)}(n|\mathbf{b}_{1:n-1})$ and $\breve{\mathbf{M}}(n|\mathbf{b}_{1:n}) = \breve{\mathbf{M}}^{(m)}(n|\mathbf{b}_{1:n-1})$. □

Equation (4.53) is instructive in terms of understanding the MSE performance of m-IQKF. In particular, note that for $c_m = 1$, (4.53) coincides with the correction ECM of the KF in (4.10). Furthermore, as $m \to \infty$, we find $c_m \to 1$, implying that for infinite number of quantization bits we recover the clairvoyant KF. This important consistency result is summarized in the following corollary.

Corollary 4.4 *As the number of quantization bits $m \to \infty$, the correction step ECM at time n of the m-IQKF converges to the ECM of the clairvoyant KF for which $\mathbf{b}_{1:n-1} = \mathbf{y}_{1:n-1}$, i.e.,*

$$\lim_{m\to\infty} \breve{\mathbf{M}}^{(m)}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{M}}(n|\mathbf{y}_{1:n})\,. \tag{4.54}$$

Proof: As $m \to \infty$, we have $\lim_{m\to\infty} c_m = \lim_{m\to\infty}\left[1 - (1 - \frac{2}{\pi})^m\right] = 1$. For $c_m = 1$, $\breve{\mathbf{M}}^{(m)}(n|\mathbf{b}_{1:n-1})$ [cf. (4.53)] and $\breve{\mathbf{M}}(n|\mathbf{y}_{1:n})$ [cf. (4.10)] become identical, which results in (4.54). □

Corollary 4.4 establishes that m-IQKF asymptotically (in the number of quantization bits) achieves the per correction step MSE performance of the clairvoyant KF. As demonstrated by simulations in Section 4.5, even a small number of quantization bits ($m = 2$

Table 4.3: Per step factor, c_m, for iterative quantization

bits, m	1	2	3	4
c_m	$0.637{=}2/\pi$	0.868	0.952	0.983

or $m = 3$) renders the MSE performance of m-IQKF indistinguishable from that of the clairvoyant KF.

We remark that Corollary 4.2 quantifies the *per time-step* ECM reduction for the proposed filter. Because of error accumulation however, noticeable differences could emerge between the MSEs of the m-IQKF given by $\mathrm{tr}\{\mathbf{M}(n|\mathbf{b}_{1:n})\}$ and that of the clairvoyant KF, $\mathrm{tr}\{\mathbf{M}(n|\mathbf{y}_{1:n})\}$, as time progresses.

4.3.2 IQKF for vector observations

So far, a low-complexity scalar quantizer was employed to digitize the scalar analog-valued observations $y(n)$ at each sensor. A more general case is when sensor observations are vector-valued, $\mathbf{y}(n) \in \mathbb{R}^q,\ q > 1$ (for vector state $\mathbf{x}(n) \in \mathbb{R}^p,\ p > 1$), for which an MSE optimal estimator-quantizer would entail vector quantization of $\mathbf{y}(n)$ using an estimation-based distortion metric [20]. Optimal state estimation based on vector quantized observations is left for future work.

In this section, an optimal bit allocation scheme for iterative scalar quantization of components of $\mathbf{y}(n)$ is introduced whose performance will be compared against the clairvoyant Kalman filter based on un-quantized $\mathbf{y}(n)$. The vector observation equation is given by $\mathbf{y}(n) = \mathbf{H}(n)\,\boldsymbol{\theta}(n) + \mathbf{v}(n)$ where $\mathbf{H}(n) \in \mathbb{R}^{(q\times p)}$, $\mathbf{y}(n) := [y(n,1),\ldots,y(n,q)]$, and $\mathbf{v}(n) := [v(n,1),\ldots,v(n,q)]$ with $\mathbf{C}_v(n) := E\{\mathbf{v}(n)\mathbf{v}^T(n)\}$. If $\mathbf{C}_v(n) \neq \mathbf{I}$ a noise whitening transformation of $\mathbf{y}(n)$ is performed by setting $\mathbf{y}^{'}(n) := \mathbf{C}_v^{-1/2}(n)\mathbf{y}(n)$ resulting in

$$\begin{aligned}\mathbf{y}^{'}(n) &= \mathbf{C}_v^{-1/2}(n)\mathbf{H}(n)\,\boldsymbol{\theta}(n) + \mathbf{C}_v^{-1/2}(n)\mathbf{v}(n)\\ &:= \mathbf{H}^{'}(n)\,\boldsymbol{\theta}(n) + \mathbf{v}^{'}(n)\end{aligned} \tag{4.55}$$

where the noise term $\mathbf{v}^{'}(n) := \mathbf{C}_v^{-1/2}(n)\mathbf{v}(n)$ is white.

If the matrix $\mathbf{H}^{'}(n)$ is tall (i.e., $q > p$), optimal dimensionality reduction can be performed in (4.55) using the QR-factorization $\mathbf{H}^{'}(n) = \mathbf{Q}_1(n)\mathbf{R}_1(n)$, where $\mathbf{Q}_1(n) \in \mathbb{R}^{(q\times p)}$ has p orthonormal columns and $\mathbf{R}_1(n) \in \mathbb{R}^{(p\times p)}$ is upper triangular [11, pg. 682]. Projection of the q-dimensional $\mathbf{y}(n)$ onto the p-dimensional space spanned by the rows of $\mathbf{Q}_1(n)$ is given by

$$\begin{aligned}\mathbf{Q}_1^T(n)\mathbf{y}^{'}(n) &= \mathbf{Q}_1^T(n)\mathbf{H}^{'}(n)\,\boldsymbol{\theta}(n) + \mathbf{Q}_1^T(n)\mathbf{v}^{'}(n)\\ &= \mathbf{R}_1(n)\,\boldsymbol{\theta}(n) + \mathbf{Q}_1^T(n)\mathbf{v}^{'}(n)\end{aligned} \tag{4.56}$$

where $\mathbf{Q}_1^T(n)\mathbf{H}^{'}(n) = \mathbf{Q}_1^T(n)\mathbf{Q}_1(n)\mathbf{R}_1(n) = \mathbf{R}_1(n)$ and $\mathbf{Q}_1^T(n)\,E\{\mathbf{v}(n)\mathbf{v}^T(n)\}\,\mathbf{Q}_1(n) = \mathbf{I}$. Thus, using (4.56) we can henceforth assume that $\mathbf{y}(n)$ is p-dimensional with white noise, i.e., after using the correspondences $\mathbf{y}(n) \leftrightarrow \mathbf{Q}_1^T(n)\mathbf{y}^{'}(n)$, $\mathbf{H}(n) \leftrightarrow \mathbf{Q}_1^T(n)\mathbf{H}^{'}(n)$, and $\mathbf{v}(n) \leftrightarrow \mathbf{Q}_1^T(n)\mathbf{v}^{'}(n)$.

Let m bits be used for quantizing the entries $\{y(n,l)\}_{l=1}^{p}$ of $\mathbf{y}(n) \in \mathbb{R}^p$. Clearly, there are p^m possible bit allocation choices. To reduce the complexity of a bit allocation scheme, the iterative quantizer of (4.43), performing sequential binary quantization, will be employed. Iterative scalar quantization and MMSE estimation entails selecting the entries of $\mathbf{y}(n)$ to be quantized and estimation by the m-IQKF algorithm of Proposition 4.3. The resultant MMSE is given by the trace of the first p entries of (4.47) i.e., $\mathrm{tr}\{\mathbf{M}^{(i)}(n|\mathbf{b}_{1:n-1})\}$. Iterative bit allocation of the m bits leads to p choices of scalars to quantize for each of the m bits leading to $p \times m$ choices - significantly less than the exponential number of searches (p^m) needed for an optimal bit allocation. Proposition 4.4 summarizes the i-th bit allocation for the m-IQKF of [P4]-[C6].

Proposition 4.4 *Consider vector observations* $\mathbf{y}(n) = \mathbf{H}(n)\boldsymbol{\theta}(n) + \mathbf{v}(n)$. *Let* $\mathbf{h}^T(n,l)$ *denote the* l*-th row of* $\mathbf{H}(n) \in \mathbb{R}^{p\times p}$ *and* $E\{\mathbf{v}(n)\mathbf{v}^T(n)\} = \mathbf{I}$. *Given* m*-bits for iterative quantization of the components of* $\mathbf{y}(n)$ *(i.e.,* $\{y(n,l)\}_{l=1}^{p}$*) for the m-IQKF of [P4]-[C6][cf. Section 4.3.1], the MMSE of estimator of* $\boldsymbol{\theta}(n)$ *is achieved by allocating the* i*-th bit to the component* $y(n,l^{(i)})$*, where* $l^{(i)}$ *is given as*

$$l^{(i)} = \arg\max_{l=1,2,\ldots,p} \frac{\mathbf{h}^T(n,l)\left[\mathbf{M}^{(i-1)}(n|\mathbf{b}_{1:n-1})\right]^2\mathbf{h}(n,l)}{\mathbf{h}^T(n,l)\mathbf{M}^{(i-1)}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n,l)+1}. \tag{4.57}$$

Proof: From (4.47), the estimator MSE reduction due to i-th quantization bit $b^{(i)}(n)$ is given as

$$\text{trace}\left\{\frac{2}{\pi}\frac{\mathbf{M}^{(i-1)}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)\mathbf{h}^T(n)\mathbf{M}^{(i-1)}(n|\mathbf{b}_{1:n-1})}{\mathbf{h}^T(n)\mathbf{M}^{(i-1)}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)+c_v(n)}\right\}$$
$$=\frac{2}{\pi}\frac{\mathbf{h}^T(n)\left[\mathbf{M}^{(i-1)}(n|\mathbf{b}_{1:n-1})\right]^2\mathbf{h}(n)}{\mathbf{h}^T(n)\mathbf{M}^{(i-1)}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)+c_v(n)}. \tag{4.58}$$

Using $c_v(n)=1$ and minimizing (4.58) over the choices of $\mathbf{h}(n):=\mathbf{h}(n,l)$ corresponds to selecting scalar components $\{y(n,l)\}_{l=1}^p$ to be quantized as $b^{(i)}(n)=\text{sign}\{y(n,l)-\mathbf{h}^T(n,l)\hat{\boldsymbol{\theta}}^{(i)}(n|\mathbf{b}_{1:n-1})\}$. The sequence of minimizations for each bit $i=1,2,\ldots,m$ is therefore obtained by (4.57). □

4.3.3 Performance analysis of the m-IQKF

We have quantified the *per correction step* ECM reduction $\Delta\breve{\mathbf{M}}_m(n)$ in the m-IQKF as a function of the number of bits m used for iteratively quantizing the scalar observations $y(n)$ in Section 4.3.1. We next compare the MSE performance of m-IQKF with that of the clairvoyant KF, when both prediction and correction steps are considered, by deriving the continuous-time Riccati equations for both filters.

Consider first the discrete-time algebraic Riccati equation (ARE) for the m-IQKF [P4]-[C6]. Note that regardless of the structure of $\breve{\mathbf{M}}(n-1|\mathbf{b}_{1:n-1})$, substituting for $\breve{\mathbf{A}}(n)$ and $\mathbf{C}_{\breve{\mathbf{u}}}(n)$ in (4.45) will result in

$$\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1})=\begin{pmatrix}\mathbf{M}(n|\mathbf{b}_{1:n-1}) & \mathbf{0}\\ \mathbf{0}^T & c_v(n)\end{pmatrix} \tag{4.59}$$

with $\mathbf{M}(n|\mathbf{b}_{1:n-1})=\mathbf{A}(n)\mathbf{M}(n-1|\mathbf{b}_{1:n-1})\mathbf{A}^T(n)+\mathbf{C}_{\mathbf{u}}(n)$. Equation (4.59) shows that the predicted ECM for the augmented-state has a block-diagonal structure.

To simplify notation let $\breve{\mathbf{M}}(n+1):=\breve{\mathbf{M}}(n+1|\mathbf{b}_{1:n})$, $\breve{\mathbf{M}}(n):=\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1})$, $\mathbf{M}(n+1):=\mathbf{M}(n+1|\mathbf{b}_{1:n})$, $\mathbf{M}(n):=\mathbf{M}(n|\mathbf{b}_{1:n-1})$, and substitute (4.53) in (4.45) to obtain the m-IQKF

ARE for the ECM of $\breve{\mathbf{x}}(n)$ as

$$\breve{\mathbf{M}}(n+1) = \breve{\mathbf{A}}(n+1)\breve{\mathbf{M}}(n)\breve{\mathbf{A}}^T(n+1) + \mathbf{C}_{\breve{\mathbf{u}}}(n+1) - c_m \frac{\breve{\mathbf{A}}(n+1)\breve{\mathbf{M}}(n)\breve{\mathbf{h}}(n)\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}(n)\breve{\mathbf{A}}^T(n+1)}{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}(n)\breve{\mathbf{h}}(n)}. \tag{4.60}$$

Substituting for $\breve{\mathbf{A}}(n+1)$, $\mathbf{C}_{\breve{\mathbf{u}}}(n+1)$, $\breve{\mathbf{h}}(n)$ and $\breve{\mathbf{M}}(n)$ in (4.60), and equating both sides of the upper $p \times p$ block-diagonal sub-matrices of the resulting expressions, yields the ARE for the ECM of $\mathbf{x}(n)$ as

$$\mathbf{M}(n+1) = \mathbf{A}(n+1)\mathbf{M}(n)\mathbf{A}^T(n+1) + \mathbf{C}_{\mathbf{u}}(n+1) - c_m \frac{\mathbf{A}(n+1)\mathbf{M}(n)\mathbf{h}(n)\mathbf{h}^T(n)\mathbf{M}(n)\mathbf{A}^T(n+1)}{\mathbf{h}^T(n)\mathbf{M}(n)\mathbf{h}(n) + c_v(n)}. \tag{4.61}$$

Interestingly, the resulting ARE for the ECM of the m-IQKF estimate $\hat{\mathbf{x}}(n|\mathbf{b}_{1:n})$ [cf. (4.61)], becomes identical to the ARE of the clairvoyant KF [49, Ch. 5.11], as $\lim_{m\to\infty} c_m = 1$. The implication of these relations will become clear from the continuous-time ARE for the m-IQKF given in Proposition 4.5, where (4.1) and (4.2) are viewed as discrete-time equivalent of a continuous-time state space model.

Proposition 4.5 *The continuous-time ECM* $\mathbf{M}_c(t) = \mathbf{M}_c(nT_s) := \lim_{T_s\to 0} \mathbf{M}(n;T_s)$ *for the m-IQKF of (4.44)-(4.47) is the solution of the following differential (Riccati) equation:*

$$\dot{\mathbf{M}}_c(t) = \mathbf{A}_c(t)\mathbf{M}_c(t) + \mathbf{M}_c(t)\mathbf{A}_c^T(t) + \mathbf{C}_u(t) - \mathbf{M}_c(t)\mathbf{h}_c(t)\left[\frac{c_{v_c}(t)}{c_m}\right]^{-1}\mathbf{h}_c^T(t)\mathbf{M}_c(t)\,. \tag{4.62}$$

Proof: The derivation steps from (4.61) to (4.62) are the same as those followed for deriving the KF Riccati equation from the corresponding ECMs [49, pp. 259], [67]. The only difference is that the scale factor, c_m, in front of the third term on the RHS of (4.61) for the KF equals to 1. □

When comparing (4.62) to the clairvoyant KF ARE [49, pp. 259]

$$\dot{\mathbf{M}}_c^K(t) = \mathbf{A}_c(t)\mathbf{M}_c^K(t) + \mathbf{M}_c^K(t)\mathbf{A}_c^T(t) + \mathbf{C}_u(t) - \mathbf{M}_c^K(t)\mathbf{h}_c(t)\left[c_{v_c}(t)\right]^{-1}\mathbf{h}_c^T(t)\mathbf{M}_c^K(t) \tag{4.63}$$

Table 4.4: Noise variance penalty for iterative quantization.

bits, m	1	2	3	4
$(1/c_m - 1)100\%$	57.08%	15.21%	5.04%	1.77%

where $\mathbf{M}_c^K(t)$ denotes the continuous-time ECM of KF, it is evident that the variance of the continuous-time observation noise of the m-IQKF is amplified by $1/c_m$ compared to that of the clairvoyant KF. This is the price paid for using quantized observations, $b(n)$, instead of the analog-amplitude ones, $y(n)$. Note that $2/\pi \leq c_m = 1 - (1 - 2/\pi)^m \leq 1$. Table 4.4 shows the percentage observation noise variance increase versus the number of quantization bits, m. Surprisingly, even with only 4-bit quantization, the observation noise variance increase is less than 2% of the analog-amplitude observations variance.

In the next section, simulations for the two quantizer-estimator approaches (batch and iterative) are provided in order to corroborate the analytical statements and validate model consistency.

4.4 Quantized Kalman filter algorithms

The predictor and corrector steps of the batch quantized Kalman filter algorithms are summarized in this section. Analogous algorithms can be obtained for the iteratively quantized Kalman filter by replacing the batch quantizer with the iterative quantizer.

The active sensor at discrete time n is determined by finding $k = \text{mod}_K(n)$, where $\text{mod}_K(\cdot)$ represents modulo K operation. The steps for sensing, quantization, transmission in an active sensor is given in Fig. 4.2(a). Furthermore, the operations performed at all sensors to obtain the estimates are given in Fig. 4.2(b).

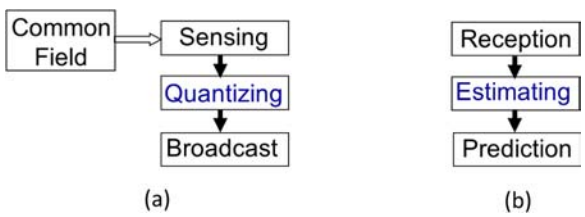

Figure 4.2: (a) Active sensor operations (b) All sensors operations.

Algorithm 4.1 tabulates the steps performed at time n from acquiring a measurement $y^*(n)$ to broadcasting the message $b(n)$ from quantized innovations $\tilde{y}(n|\mathbf{b}_{1:n-1})$. The corrector step is tabulated as Algorithm 4.2.

Algorithm 4.1 Prediction and Quantization (BQKF)

Require: Active sensor $\mathcal{S}_k$ knows $\mathbf{h}(n)$, $y^*(n)$

All sensors $\{\mathcal{S}_k\}_{k=1}^K$ know $\hat{\boldsymbol{\theta}}(n-1|\mathbf{b}_{1:n-1})$, $\mathbf{M}(n-1|\mathbf{b}_{1:n-1})$

Predictor:

$\{\mathcal{S}_k\}_{k=1}^K$: $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) = \mathbf{A}(n)\hat{\boldsymbol{\theta}}(n-1|\mathbf{b}_{1:n-1})$

$\mathbf{M}(n|\mathbf{b}_{1:n-1}) = \mathbf{A}(n)\mathbf{M}(n-1|\mathbf{b}_{1:n-1})\mathbf{A}^T(n) + \mathbf{C}_u(n)$

$\{\mathcal{S}_k\}_{k=1}^K$: Calculate $\mathbf{k}_b(n) := \frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)}{\sqrt{\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)+c_v(n)}}$

Sensing:

$\mathcal{S}_k$: Acquire $y^*(n)$

Quantization:

$\mathcal{S}_k$: Calculate innovations $\tilde{y}(n|\mathbf{b}_{1:n-1}) := y^*(n) - \mathbf{h}^T(n)\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})$

$\mathcal{S}_k$: Quantize $\tilde{y}(n|\mathbf{b}_{1:n-1})$ using (4.14) to get $b(n) = i$

$\mathcal{S}_k$: Broadcast i

Algorithm 4.2 Estimation (BQKF)

Require: All sensors $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})$, $\mathbf{M}(n|\mathbf{b}_{1:n-1})$ and $\mathbf{k}_b(n)$

Data Reception:

FC: Receive $b(n) = i$

Corrector:

$\{\mathcal{S}_k\}_{k=1}^K$: Calculate $\alpha_i(n)$ and $\beta_i(n)$ using (4.18) and (4.19), respectively

$\{\mathcal{S}_k\}_{k=1}^K$: $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) = \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) + \alpha_i(n)\mathbf{k}_b(n)$

$\mathbf{M}(n|\mathbf{b}_{1:n}) = \mathbf{M}(n|\mathbf{b}_{1:n-1}) - \beta_i(n)\mathbf{k}_b(n)\mathbf{k}_b^T(n)$

4.5 Numerical studies: QKF

Numerical studies will be used to verify the performance of the two quantizer-estimator approaches proposed. In both approaches the Bayesian MMSE performance of the QKF will be compared to that of unquantized KF case. A linear state space model will be studied.

In the first study, linear measurements will be used whereas the second study will adopt non-linear measurements model and employ the extended (quantized) KF approach.

4.5.1 Linear observations

The two-dimensional state space model simulated is

$$\dot{\boldsymbol{\theta}}_c(t) := \begin{pmatrix} \dot{x}_{c_1}(t) \\ \dot{x}_{c_2}(t) \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} x_{c_1}(t) \\ x_{c_2}(t) \end{pmatrix} + \begin{pmatrix} 0 \\ 1 \end{pmatrix} u_c(t)\,.$$

Noisy measurements at sensor $\mathcal{S}_k$ are given as $y_k(t) = x_{c_1}(t) + \theta_k x_{c_2}(t) + v_k(t)$ where θ_k is a parameter at sensor k. A discrete-time equivalent model is used with sampling time $T_s = 1$.

WSN data corresponding to $K{=}2$ sensors with $[\theta_1 \;\; \theta_2] = [0.1 \;\; 0.2]$ are simulated. The measurement and state driving noise variances are $c_v(t) = 1$ and $c_u(t) = 1$. The MSE plots in Fig. 4.9(b) and Fig. 4.11(a) illustrate the evolution of the MSEs, obtained from the trace of the respective ECMs, against the time index n.

Figs. 4.9(a) and 4.10(a) correspond to BQKF whereas Figs. 4.9(b) and 4.11(a) correspond to IQKF. First, a discussion of the BQKF simulation results is highlighted.

In Fig. 4.9(a) BQKF MSE, i.e., $\mathrm{tr}\{\mathbf{M}(n|\mathbf{b}_{1:n})\}$, for $1, 2$ and 3 bits is compared with the MSE of the clairvoyant KF. The plots demonstrate the MSE improvement offered by 2 bits of quantization compared to the 1-bit BQKF case. Quantizing with more than 2 bits offers little MSE improvement.

Model consistency tests comparing the empirical MSEs with analytical MSEs are depicted in Fig. 4.10(a). The analytical MSEs are compared with empirically computed MSEs (sample average of the squared estimation errors) for both predictor MSE and estimator MSE. Note that the analytical MSE is obtained from the trace of the ECM $\mathbf{M}(n|\mathbf{b}_{1:n}) := \mathbb{E}\{[\boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})][\boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})]^T|\mathbf{b}_{1:n}\}$. The empirical MSE is the sample estimator of $\mathrm{tr}\{\mathbf{M}(n|\mathbf{b}_{1:n})\}$ obtained as a sample average of the squared estimation errors. The plots depict both predictor MSE $\mathrm{tr}\{\mathbf{M}(n|\mathbf{b}_{1:n-1})\}$ and estimator MSE $\mathrm{tr}\{\mathbf{M}(n|\mathbf{b}_{1:n})\}$. This consistency test reveals that the empirical and analytical MSEs are nearly identical.

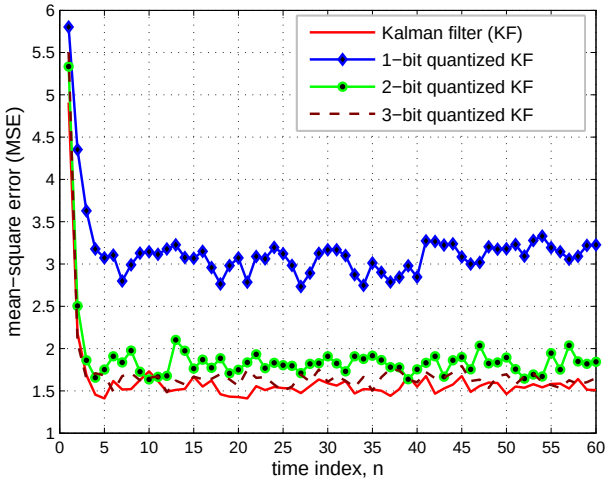

Figure 4.3: MSEs for 1-3 bits of QKF and MSE for KF.

In the IQKF case, the first simulation setup, Fig. 4.9(b) the MSE of the IQKF is compared to the MSE of the clairvoyant KF. The plots demonstrate that there is a substantial MSE reduction when going from 1 bit of quantization to 2 bits and little MSE reduction for higher number of quantization bits. With 2 bits of quantization the MSE of IQKF is virtually coincident with that of the clairvoyant KF as was postulated by the analytical values in Table 4.3 whereby $c_2 = 86.8\%$.

Model consistency checks depicted in Fig. 4.11(a) reveals that the empirical and analytical MSEs are nearly identical for the IQKF case as well.

4.5.2 Non-linear observations

Target tracking in $\mathbb{R}^2$ is simulated here with the target position $\mathbf{s}(n) := [s_1(n), s_2(n)]^T$ and velocity $\mathbf{w}(n) := [w_1(n), w_2(n)]^T$ forming a state vector $\boldsymbol{\theta}(n) := [\mathbf{s}(n)^T, \mathbf{w}(n)^T]^T$ with state equation

$$\begin{pmatrix} \mathbf{s}(n) \\ \mathbf{w}(n) \end{pmatrix} = \begin{pmatrix} 1 & 0 & T_s & 0 \\ 0 & 1 & 0 & T_s \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \mathbf{s}(n-1) \\ \mathbf{w}(n-1) \end{pmatrix} + \begin{pmatrix} T_s^2/2 & 0 \\ 0 & T_s^2/2 \\ T_s & 0 \\ 0 & T_s \end{pmatrix} \mathbf{u}(n) \tag{4.64}$$

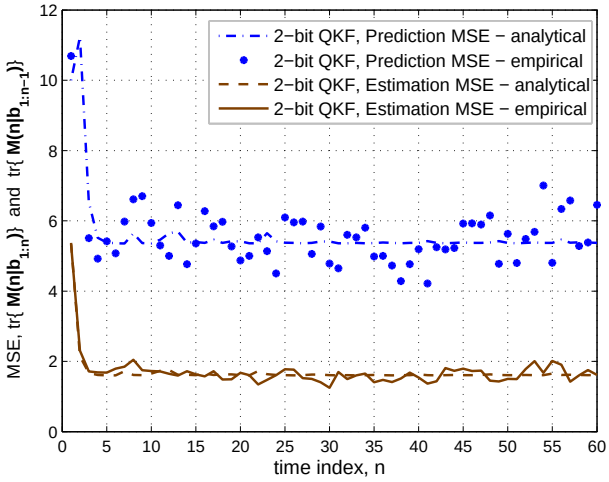

Figure 4.4: Analytical vs. Empirical MSEs.

where $T_s = 1$ denotes sampling period. The model describes a constant-velocity tracking set-up [8, pp. 82], where the acceleration term is captured by the noise term $\mathbf{u}(n)$. Uniformly deployed sensors take noisy Euclidean distance measurements of the target. Sensor $\mathcal{S}_k$, located at position $\boldsymbol{\theta}^k$ measures

$$y_k(n) = \|\boldsymbol{\theta}(n) - \boldsymbol{\theta}^k\| + v(n). \tag{4.65}$$

Linearizing (4.65) about a generic state prediction $\hat{\boldsymbol{\theta}}(n|n-1)$ in similar fashion to the extended Kalman filter (E)KF, we obtain

$$y_k(n) \approx \mathbf{h}_k^T(n)\boldsymbol{\theta}(n) + \mathbf{y}_k^0(n) + v(n) \tag{4.66}$$

where $\mathbf{h}_k(n) = \frac{\hat{\boldsymbol{\theta}}(n|n-1)-\boldsymbol{\theta}^k}{\|\hat{\boldsymbol{\theta}}(n|n-1)-\boldsymbol{\theta}^k\|}$ and $\mathbf{y}_k^0(n)$ is a function of $\hat{\boldsymbol{\theta}}(n|n-1)$ and $\boldsymbol{\theta}^k$; The linearized observations (4.66) together with (4.64) are amenable to the IQKF and BQKF algorithms of Sections 4.2 and 4.3.

WSN data corresponding to K=100 sensors are simulated. All plots generated illustrate the evolution of the MSEs, obtained from the trace of the respective ECMs, against the time index n. Figs. 4.9(a), 4.10(a), and 4.10(b) are for the batch quantized KF whereas Figs. 4.9(b), 4.11(a), and 4.11(b) are for the iteratively quantized KF. The MSE for a linear state space model with vector measurements using IQKF based on the bit allocation of Section 4.3.2 is shown in Fig. 4.8.

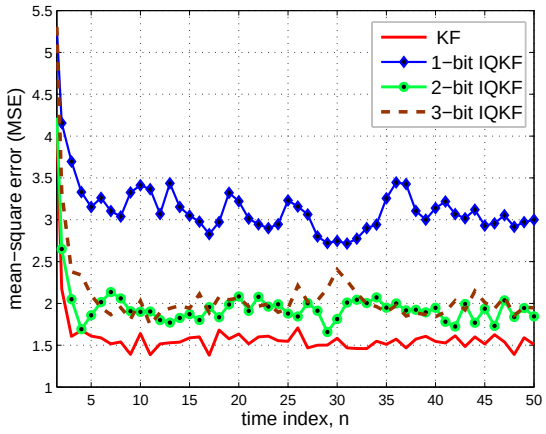

Figure 4.5: IQKF estimator MSE $\mathrm{tr}\{\mathbf{M}(n|\mathbf{b}_{1:n})\}$ from Monte Carlo data for 1-3 bits of quantization; MSE for clairvoyant KF is included for comparison.

In the first simulation setup, Fig. 4.9(a) depicts the BQKF MSE. The simulations are performed for $1, 2$ and 3 bits and the respective MSEs are compared with the MSE of the clairvoyant KF. The plots demonstrate the MSE improvement offered by 2 quantization bits compared to the 1-bit BQKF case, even for the extended KF approach. It is also evident that quantization to more than 2 bits offers little MSE improvement. Alternative model consistency checks comparing the empirically obtained MSEs with analytical MSEs are depicted in Fig. 4.10(a). The consistency check reveals that the empirical and analytical MSEs are nearly identical.

Fig. 4.10(b) shows alternative model consistency tests for the BQKF using the normalized estimation error squared (NEES) tests of [8, Ch. 5.4]. NEES $r(n) := [\boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})]^T[\mathbf{M}(n|\mathbf{b}_{1:n})]^{-1}[\boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})]$ is postulated to have a χ^2 pdf with p degrees of freedom (since the p-dimensional $\tilde{\boldsymbol{\theta}}(n) := \boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ is assumed zero-mean Gaussian with covariance $\mathbf{M}(n|\mathbf{b}_{1:n})$ if the estimator is consistent with the model). Under the hypothesis that the estimator is consistent, L realizations of the NEES statistics $\{r_i(n)\}_{i=1}^L$ each χ^2 distributed with p degrees of freedom, lead to a χ^2 distribution with Lp degrees of freedom. This is checked by running Monte Carlo simulations and computing the sample average NEES $\hat{r}(n) := \frac{1}{L}\sum_{i=1}^L r_i(n)$ and then defining an acceptance (confidence) region (for the consistent hypothesis). If $l \leq \hat{r}(n) \leq u$, then the estimator is consistent; lower and

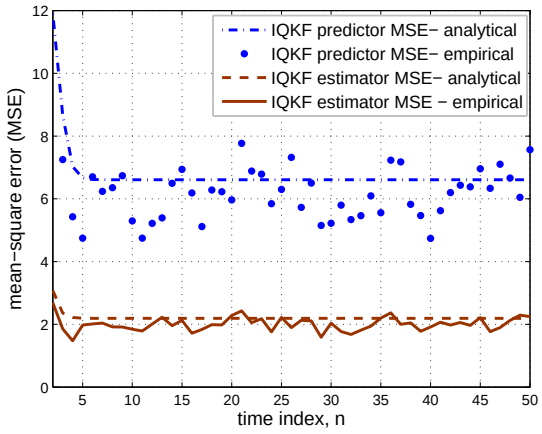

Figure 4.6: Comparing empirical MSE from Monte Carlo data and analytical MSE, for predictor ECM $\mathbf{M}(n|\mathbf{b}_{1:n-1})$ and estimator ECM $\mathbf{M}(n|\mathbf{b}_{1:n})$.

upper bounds l and u are obtained from $\Pr\{\hat{r}_n \in [l,\ u]\} = 1-\alpha$, where $1-\alpha$ is the probability of acceptance region. Using $L = 100$ realizations, state space of $p = 2$ dimensions, and $\alpha = 0.05$ (i.e., 95% region), we observed that about 80% of 500 time samples simulated are within the 95% acceptance region.

In Fig. 4.9(b), the MSE of the IQKF is compared with the MSE of the clairvoyant KF. Once again we observe a substantial MSE reduction when going from 1 to 2 quantization bits and very little performance gain for higher number of bits as was the case for BQKF. With 2 quantization bits the MSE performance is virtually coincident with that of the clairvoyant KF as was postulated by the analytical values in Table 4.3 whereby $c_2 = 86.8\%$.

In Figs. 4.11(a) and 4.11(b), model consistency of the IQKF estimator is checked. The IQKF algorithm's analytical MSE is compared with the empirical MSE evaluated through Monte Carlo runs depicted in Fig. 4.11(a). The analytical and the empirical MSEs match well. The alternative NEES consistency test for the iteratively quantized KF is presented in Fig. 4.11(b), where $p = 4$, 500 time samples, and $\alpha = 0.05$ are used. About 75% of 500 time samples lie inside the 95% acceptance region.

MSE comparison between 2-BQKF and 2-IQKF is depicted in Fig. 4.7. The respective MSEs are very close. However, the BQKF exhibits slightly smaller MSE which happens because the iterative approach uses the Gaussian approximation on $p\left[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}, \mathbf{b}^{(1:i)}(n)\right]$

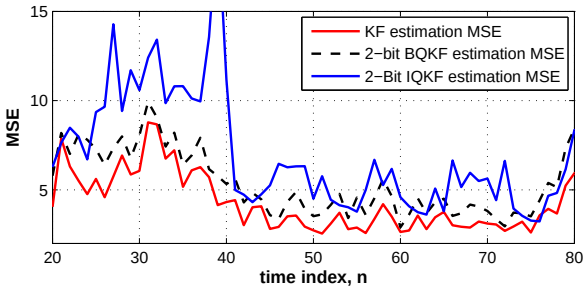

Figure 4.7: Batch vs. iterative quantization

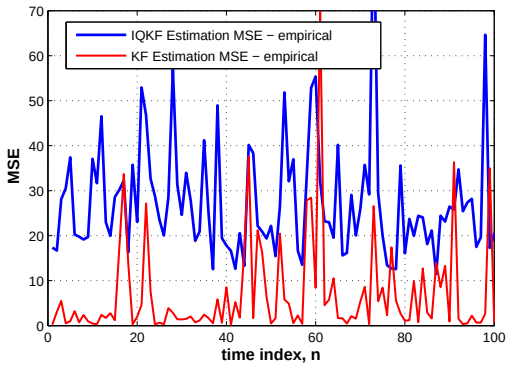

Figure 4.8: IQKF: MSE for iteratively quantized vector observations vs. clairvoyant KF

a number of times (m) per time-step (Proposition 4.3), whereas the batch approach invokes this approximation only once per time-step (Proposition 4.1). Figure 4.8 demonstrates that the low complexity iterative scalar quantization has potential for application in vector observation cases as well.

4.6 Concluding remarks

Recursive state estimators based on quantized observations were considered. Multi-bit quantization effected by either an iterative binary quantizer or a single-shot batch quantization of the measurement innovations to a block of bits was studied. Motivated by the need to find quantifiable trade-offs between estimation performance and number of quantization bits for distributed estimation, it was shown that quantization using 2 to 3 bits improves the performance of Kalman-like algorithms to virtually coincide with the Kalman

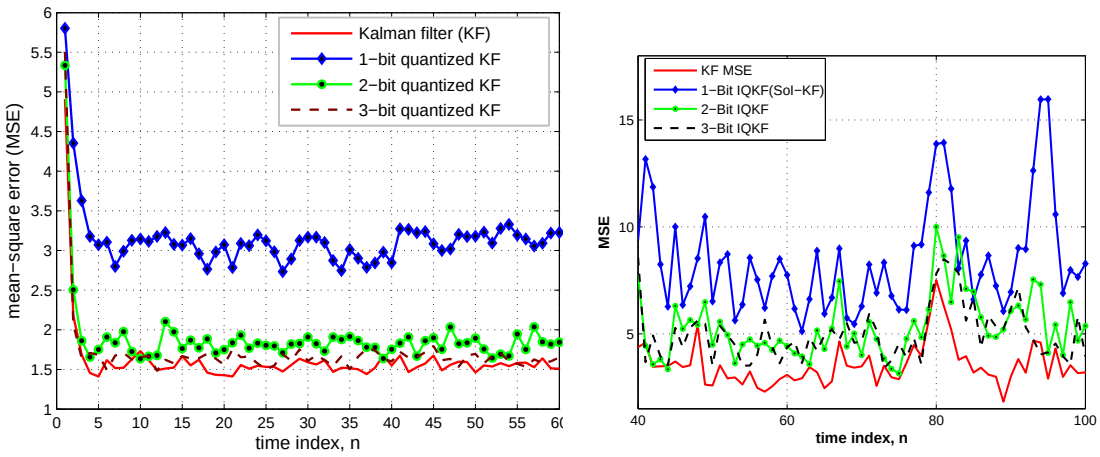

(a) (b)

Figure 4.9: (a) Batch quantized KF, (b) IQKF vs. clairvoyant KF: Empirical MSE

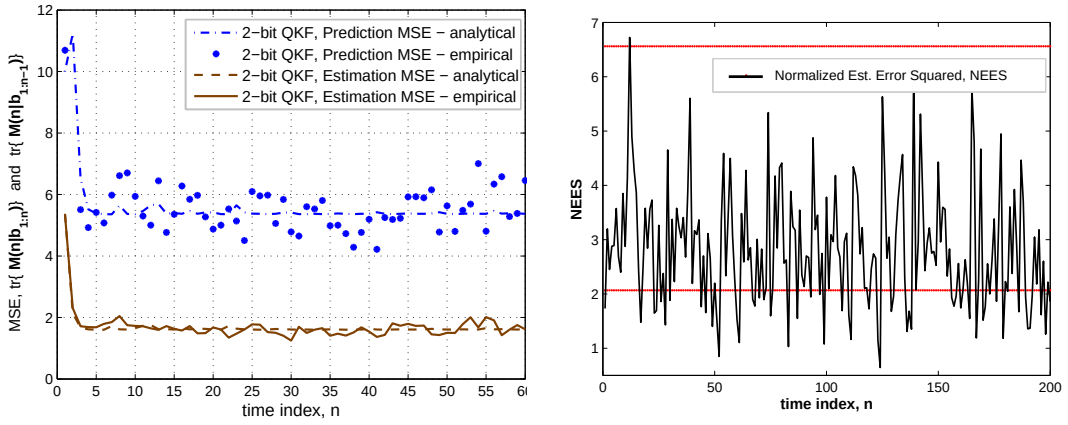

(a) (b)

Figure 4.10: BQKF consistency test: (a) Empirical MSE vs. analytical MSE (b) Normalized estimation error squared

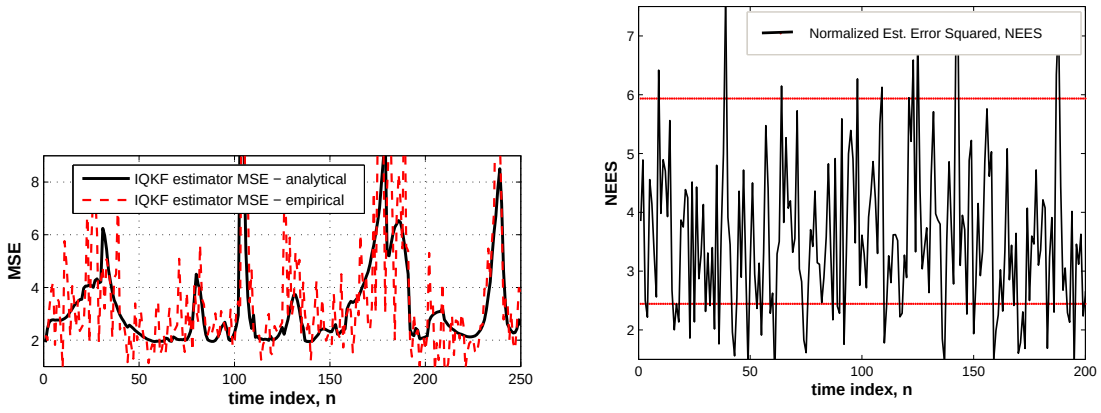

(a) (b)

Figure 4.11: IQKF consistency test: (a) Empirical MSE vs analytical MSE (b) Normalized estimation error squared

filter based on unquantized measurements.

Chapter 5

Conclusions and future research

This thesis explored the problem of distributed estimation based on low-complexity data-reduction techniques. In this chapter, the contributions contained in this work are summarized highlighting the trade-offs needed for different assumptions on signal models ranging from static linear regression to dynamical state space models.

In the first part of the thesis, Chapter 2, optimal estimation of a deterministic signal parameter based on censored and censored-quantized measurements is presented. It is shown that censoring can be motivated from a least-squares data-fitting criterion. The resultant censoring is then shown to be amenable to distributed estimation. Performance analysis based on the Cramér-Rao lower bound and numerical studies with simulated and real-world data demonstrated the potential of data censoring for measurement reduction in estimation with a large wireless sensor network. Conditions were established to guarantee uniqueness of numerical solution to the maximum likelihood estimation problems. The maximum likelihood estimator based on censoring was then shown to outperform estimators based on sensor selection, especially if the regressors in the signal model have comparable l_2 norms.

In Chapter 3, the signal was modeled as random with a known Gaussian pdf and censored-estimators and quantized-censored estimators were developed. It turns out that the known pdf of the signal parameters simplifies the design of both the censoring scheme and the quantizer compared to the deterministic signal model case. Furthermore, using Fisher's information matrix, the merits of censoring as a data reduction alternative to data

selection were analytically quantified. Analogous to the maximum likelihood approach, it was shown that if the regressors at different sensors have comparable l_2-norms, the censoring approach leads to estimators outperforming their DOE-based counterparts. In fact, simulations corroborated the intuition that uniform measurement selection would perform competitively with DOE-based selection if the regressors have identical l_2 norms; see also [54]. An extension to quantization of the uncensored data was further studied, which showed that $1-2$ bits of quantization are sufficient, in most cases, to achieve MSE performance competitive with that offered by estimators based on unquantized measurements.

The last part of the thesis in Chapter 4 focused on state estimation based on two novel multi-bit quantization techniques: batch and sequential quantization. Recursive state estimators based on quantized observations were developed. Motivated by the need to find quantifiable trade-offs between estimation performance and number of quantization bits for decentralized estimation, it was shown that quantization using 2 to 3 bits leads the performance of Kalman-like algorithms to virtually coincide with the optimal state estimator (Kalman filter), with only minimal increase in computation. Numerical simulations were used to compare the novel filter's true covariances with analytical ones, to check model consistency issues, as well as for performance comparison of the two quantization-estimation approaches. The mean-square error performance for both batch and iteratively quantized estimators were found to be close to the clairvoyant KF based on unquantized measurements for tracking examples involving the Gauss-Markov model.

Fruitful research directions in censored-estimators and censor-quantizer-estimators could extend the censoring idea in two directions: estimation with nonlinear regression models and real-time tracking of time-varying signals. These directions will broaden the viability of measurement censoring as a data selection technique to a wider range of inference problems. The time-varying signal could be estimated by sequentially updating the censoring thresholds as was done for the quantizer thresholds in the quantized Kalman filtering approach of Chapter 4.

Further investigation regarding the decentralized state estimation based on quantized measurements include studying the use of particle filters in place of the Gaussian approxi-

mants of the predictor pdf. Robustness issues of the estimation algorithms to noisy inter-sensor channels is also worth studying.

Censoring of noisy measurements from the state space model resemble an approach from Kalman filtering based on intermittently received measurements in [41, 75]. In [41, 75], it was established that there is a lower threshold on the fraction of measurements received below which the stability of the Kalman filter is not guaranteed. An interesting problem to study is stability of the error covariance matrix from the Riccati equation resulting from tracking a time-varying signal based on a limited set of measurements due to the censored measurements. It is conjectured that there would be a corresponding censoring rate ρ below which, with high probability, the Kalman filter will diverge. Preliminary numerical studies using simulations indicate that to be the case. Rigorous stability analysis paralleling that of [75] is thus well motivated.

Bibliography

[1] P. Addesso, S. Marano, and V. Matta, "Sequential Sampling in Sensor Networks for Detection with Censoring Nodes," *IEEE Transactions on Signal Processing*, vol. 55, pp. 5497–5505, November 2007.

[2] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless Sensor Networks: A Survey," *Computer Networks*, vol. 38, no. 4, pp. 393–422, 2002.

[3] M. Ali, U. Saif, A. Dunkels, T. Voigt, K. Roemer, K. Langendoen, J. Polastre, and Z. A. Uzmi, "Medium Access Control Issues in Sensor Networks," *ACM SIGCOMM Computer Communication Review*, vol. 36, pp. 33–36, April 2006.

[4] T. Amemiya, "Regression Analysis When the Dependent Variable is Truncated Normal," *Econometrica*, vol. 41, pp. 997–1016, November 1973.

[5] S. Appadwedula, V. V. Veeravalli, and D. L. Jones, "Decentralized Detection with Censoring Sensors," *IEEE Transactions on Signal Processing*, vol. 56, pp. 1362–1373, April 2008.

[6] T. Arampatzis, J. Lygeros, and S. Manesis, "A Survey of Applications of Wireless Sensors and Wireless Sensor Networks," in *Proc. of the IEEE Intl. Symposium on Intelligent Control*, Limassol, Cyprus, pp. 719–724, June 2005.

[7] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A Tutorial on Particle Filters for Online Nonlinear/Non-Gaussian Bayesian Tracking," *IEEE Transactions on Signal Processing*, vol. 50, pp. 174–188, October 2002.

[8] Y. Bar-Shalom and X. Li, *Estimation and Tracking: Principles, Techniques, and Software.* Artech House, 1993.

[9] D. P. Bertsekas, *Nonlinear Programming.* Belmont, MA: Athena Scientific, 1999.

[10] F. Bian, D. Kempe, and R. Govindan, "Utility-based Sensor Selection," in *Proc. of Intl. Conf. Information Processing in Sensor Networks*, Nashville, TN, pp. 11–18, April 2006.

[11] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY: Cambridge University Press, 2004.

[12] J. Buckley and I. James, "Linear Regression with Censored Data," *Biometrika*, vol. 66, pp. 429–436, December 1979.

[13] C.-Y. Chong and S. P. Kumar, "Sensor Networks: Evolution, Opportunities, and Challenges," *Proceedings of the IEEE*, vol. 91, pp. 1247–1256, August 2003.

[14] P. M. Djuric, M. Vemula, and M. F. Bugallo, "Target Tracking by Particle Filtering in Binary Sensor Networks," *IEEE Transactions on Signal Processing*, vol. 56, pp. 2229–2238, June 2008.

[15] J. Doob, *Stochastic Processes*. John Wiley & Sons Inc., 1953.

[16] J. Fang and H. Li, "Optimal/Near–optimal Dimensionality Reduction for Distributed Estimation in Homogeneous and Certain Inhomogeneous Scenarios," *IEEE Transactions on Signal Processing*, vol. 58, pp. 4339–4353, August 2010.

[17] M. Gastpar, P. L. Dragotti, and M. Vetterli, "The distributed Karhunen–Loeve transform," *IEEE Transactions on Information Theory*, vol. 52, pp. 5177–5196, December 2006.

[18] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*, Kluwer Academic Publishers: Norwell, MA, 1992.

[19] D. Golovin, M. Faulkner, and A. Krause, "Online Distributed Sensor Selection," in *Proc. of Intl. Conf. on Information Processing in Sensor Networks*, Stockholm, Sweden, pp. 220–231, April 2010.

[20] R. M. Gray, "Quantization in Task–Driven Sensing and Distributed Processing," in *Proc. of Intl. Conference on Acoustics, Speech, and Signal Processing*, Toulouse, France, vol. 5, pp. 215–219, May 2006.

[21] J. Gubner, "Distributed Estimation and Quantization," *IEEE Transactions on Information Theory*, vol. 39, pp. 1456–1459, July 1993.

[22] C. Guestrin, P. Bodik, R. Thibaux, M. Paskin, and S. Madden, "Distributed Regression: An Efficient Framework for Modeling Sensor Network Data," in *Proc. of Intl. Symposium on Information Proc. in Sensor Networks*, Berkeley, CA, pp. 1–10, April 2004.

[23] V. Gupta, T. Chung, B. Hassibi, and R. M. Murray, "On a Stochastic Sensor Selection Algorithm with Applications in Sensor Scheduling and Dynamic Sensor Coverage," *Automatica*, vol. 42, pp. 251–260, February 2006.

[24] W. Haerdle and L. Simar, *Applied Multivariate Statistical Analysis*. Berlin, Germany: Springer-Verlag, 2nd ed., 2007.

[25] J. B. Hampshire and J. W. Strohbehn, "Tobit Maximum–Likelihood Estimation for Stochastic Time Series Affected by Receiver Saturation," *IEEE Transactions on Information Theory*, vol. 38, pp. 457–469, March 1992.

[26] R. J. Harris, *A Primer of Multivariate Statistics*. Lawrence Erlbaum Assoc., N.J., 3rd ed., 2001.

[27] G. E. Hovland and B. J. McCarragher, "Dynamic Sensor Selection for Robotic Systems," in *Proc. of Intl. Conf. on Robotics and Automation*, Albuquerque, NM, pp. 272–277, April 1997.

[28] B. Ibarz-Gabardos and P. J. Zufiria, "A Kalman Filter with Censored Data," in *Proc. of Intl. Workshop on Intelligent Signal Processing*, Faro, Portugal, pp. 74–79, September 2005.

[29] V. Isler and R. Bajcsy, "The Sensor Selection Problem for Bounded Uncertainty Sensing Models," *IEEE Transactions on Automation Science and Engineering*, vol. 3, pp. 372–381, October 2006.

[30] K. Ito and K. Xiong, "Gaussian Filters for Nnlinear Filtering Problems," *IEEE Transactions on Automatic Control*, vol. 45, pp. 910–927, May 2000.

[31] A. Jadbabaie, J. Lin, and A. Morse, "Coordination of Groups of Mobile Autonomous Agents using Nearest Neighbor Rules," *IEEE Transactions on Automatic Control*, vol. 48, pp. 988–1001, June 2003.

[32] S. Joshi and S. Boyd, "Sensor Selection via Convex Optimization," *IEEE Transactions on Signal Processing*, vol. 57, pp. 451–462, February 2009.

[33] E. L. Kaplan and P. Meier, "Nonparametric Estimation from Incomplete Observations," *Journal of the American Statistical Association*, vol. 53, pp. 457–481, June 1958.

[34] S. M. Kay, *Fundamentals of Statistical Signal Processing - Estimation Theory.* New York, NY: Prentice Hall, 1993.

[35] J. Kieffer, "Exponential Rate of Convergence for Lloyd's Method I," *IEEE Transactions on Information Theory*, vol. 28, pp. 205–210, March 1982.

[36] J. Kiefer, "Optimum Experimental Designs," in *Journal of the Royal Statistical Society, Series B,* vol. 21, pp. 272–319, 1959.

[37] J. Kotecha and P. Djuric, "Gaussian Particle Filtering," *IEEE Transactions on Signal Processing*, vol. 51, pp. 2602–2612, October 2003.

[38] A. Krause and C. Guestrin, "Near-optimal Observation Selection using Submodular Functions," in *Proc. of National Conf. on Artificial Intelligence*, vol. 22, Vancouver, Canada, pp. 1650–1654, July 2007.

[39] A. Krause, A. Singh, and C. Guestrin, "Near-optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies," *The Journal of Machine Learning Research*, vol. 9, pp. 235–284, February 2008.

[40] Intel-Berkeley Lab, "Intel Lab Data" [online] available: `http://db.csail.mit.edu/labdata/labdata.html`, 2004.

[41] X. Liu and A. Goldsmith, "Kalman Filtering with Partial Observation Losses," in *Proc. of Conference on Decision and Control*, vol. 4, pp. 4180–4186, December 2004.

[42] S. Lloyd, "Least Squares Quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, March 1982.

[43] J. P. Lynch and K. J. Loh, "A Summary Review of Wireless Sensors and Sensor Networks for Structural Health Monitoring," *Shock and Vibration Digest*, vol. 38, pp. 91–130, March 2006.

[44] Z. Ma and A. W. Krings, "Survival Analysis Approach to Reliability, Survivability and Prognostics and Health Management," in *Proc. of Aerospace Conference*, Montana, MT, pp. 1–20, March 2008.

[45] D. J. C. MacKay, "Bayesian Interpolation," *Neural Computations*, vol. 4, pp. 415–447, May 1992.

[46] D. J. C. MacKay, "Information–based Objective Functions for Active Data Selection," *Neural Computation*, vol. 4, pp. 590–604, July 1992.

[47] A. Mainwaring, D. Culler, J. Polastre, R. Szewczyk, and J. Anderson, "Wireless Sensor Networks for Habitat Monitoring," in *Proc. of the 1st ACM Intern. Workshop on Wireless Sensor Networks and Applications*, New York, NY, pp. 88–97, September 2002.

[48] J. Max, "Quantizing for Minimum Distortion," *IRE Transactions on Information Theory*, vol. 6, pp. 7–12, March 1960.

[49] P. S. Maybeck, *Stochastic Models, Estimation and Control – Vol. 1.* New York, NY: Academic Press, 1979.

[50] A. I. Mourikis and S. I. Roumeliotis, "Optimal Sensing Strategies for Mobile Robot Formations: Resource Constrained Localization," in *Proc. of Robotics: Sci. and Sys. Conf.*, Cambridge, MA, pp. 281–288, 2005.

[51] A. I. Mourikis and S. I. Roumeliotis, "Optimal Sensor Scheduling for Resource-Constrained Localization of Mobile Robot Formations," *IEEE Transactions on Robotics*, vol. 22, pp. 917–931, October 2006.

[52] E. J. Msechu and G. B. Giannakis, "Decentralized Data Selection for MAP Estimation: a Censoring and Quantization Approach," in *Proc. of Intl. Conference on Information Fusion*, Chicago, IL, July 2011.

[53] E. J. Msechu and G. B. Giannakis, "Distributed Measurement Censoring for Estimation with Wireless Sensor Networks," in *Proc. of Intl. Workshop on Signal Processing Advances in Wireless Communications*, San Francisco, CA, pp. 171–175, June 2011.

[54] E. J. Msechu and G. B. Giannakis, "Sensor-centric Data Reduction for Estimation with WSNs via Censoring and Quantization," *IEEE Transactions on Signal Processing*, May 2011 (submitted).

[55] E. J. Msechu, A. Ribeiro, S. I. Roumeliotis, and G. B. Giannakis, "Distributed Kalman Filtering based on Quantized Innovations," in *Proc. of Intl. Conference on Acoustics, Speech and Signal Processing*, Las Vegas, NV, pp. 3293–3296, April 2008.

[56] E. J. Msechu, S. I. Roumeliotis, A. Ribeiro, and G. B. Giannakis, "Distributed Iteratively Quantized Kalman Filtering for Wireless Sensor Networks," in *Proc. of Asilomar Conference on Signals, Systems and Computers*, Pacific Grove, CA, pp. 646–650, November 2007.

[57] E. J. Msechu, S. I. Roumeliotis, A. Ribeiro, and G. B. Giannakis, "Decentralized Quantized Kalman Filtering with Scalable Communication Cost," *IEEE Transactions on Signal Processing*, vol. 56, pp. 3727–3741, August 2008.

[58] W. K. Newey and D. L. McFadden, "Large Sample Estimation and Hypothesis Testing," vol. 4 of *Handbook of Econometrics*, pp. 2111–2245, Amsterdam, The Netherlands: Elsevier, 1994.

[59] H. Papadopoulos, G. Wornell, and A. Oppenheim, "Sequential Signal Encoding from Noisy Measurements using Quantizers with Dynamic Bias Control," *IEEE Transactions on Information Theory*, vol. 47, pp. 978–1002, 2001.

[60] J. L. Powell, "Least Absolute Deviations Estimation for the Censored Regression Model," *Journal of Econometrics*, vol. 25, pp. 303–325, July 1984.

[61] A. Prékopa, "Logarithmic Concave Measures and Related Topics," *Stochastic Programming*, pp. 63–82, 1980.

[62] F. Pukelsheim, *Optimal Design of Experiments*. Philadelphia, PA: Society for Industrial Mathematics, 2006.

[63] C. Rago, P. Willett, and Y. Bar-Shalom, "Censoring Sensors: A Low-Communication-Rate Scheme for Distributed Detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 32, pp. 554–568, April 1996.

[64] C. R. Rao and L. C. Zhao, "Recent Contributions to Censored Regression Models," *Metrika*, vol. 42, pp. 203–213, December 1995.

[65] A. Ribeiro and G. B. Giannakis, "Bandwidth–constrained Distributed Estimation for Wireless Sensor Networks, Part I: Gaussian Case," *IEEE Transactions on Signal Processing*, vol. 54, pp. 2784–2796, July 2006.

[66] A. Ribeiro and G. B. Giannakis, "Bandwidth–constrained Distributed Estimation for Wireless Sensor Networks, Part II: Unknown pdf," *IEEE Transactions on Signal Processing*, vol. 54, pp. 1131–1143, March 2006.

[67] A. Ribeiro, G. B. Giannakis, and S. I. Roumeliotis, "SOI-KF: Distributed Kalman Filtering with Low–Cost Communications Using the Sign of Innovations," *IEEE Transactions on Signal Processing*, vol. 54, pp. 4782–4795, December 2006.

[68] K. Romer and F. Mattern, "The Design Space of Wireless Sensor Networks," *IEEE Transactions on Wireless Communications*, vol. 11, pp. 54–61, December 2004.

[69] H. Rowaihy, S. R, M. Johnson, D. Verma, A. Bar-Noy, T. Brown, and T. La Porta, "A Survey of Sensor Selection Schemes in Wireless Sensor Networks," in *Proc. of SPIE Defense and Security Symposium*, Orlando, FL, p. 6562., April 2007.

[70] O. Roy and M. Vetterli, "Dimensionality Reduction for Distributed Estimation in the Infinite Dimensional Regime," *IEEE Transactions on Information Theory*, vol. 54, pp. 1655–1669, April 2008.

[71] S. F. A. Shah, A. Ribeiro and G. B. Giannakis, "Bandwidth-constrained MAP Estimation for Wireless Sensor Networks," in *Proc. of Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, pp. 215–219, November 2005.

[72] I. D. Schizas, G. B. Giannakis, and Z.-Q. Luo, "Distributed Estimation using Reduced-Dimensionality Sensor Observations," *IEEE Transactions on Signal Processing*, vol. 55, pp. 4284–4299, August 2007.

[73] M. K. Simon and M. S. Alouini, *Digital Communications over Fading Channels*. Hoboken, NJ: Wiley-Interscience, 2000.

[74] A. Singh, A. Krause, C. Guestrin, and W. Kaiser "Efficient Informative Sensing using Multiple Robots," *Journal of Artificial Intelligence Research*, vol. 34, pp. 707–755, April 2009.

[75] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I. Jordan, and S. S. Sastry, "Kalman Filtering with Intermittent Observations," *IEEE Transactions on Automatic Control*, vol. 49, pp. 1453–1464, September 2004.

[76] N. Trawny, *Cooperative Localization: On Motion-induced Initialization and Joint State Estimation under Communication Constraints*. PhD thesis, University of Minnesota, Minnesota, USA, 2010.

[77] H. L. V. Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: Wiley, 1968.

[78] B. Wang, M. Li, H. B. Lim, D. Ma, and C. Fu, "Energy Efficient Information Processing in Wireless Sensor Networks," in *Guide to Wireless Sensor Networks*, Computer Communications and Networks, pp. 1–26, Springer, London, 2009.

[79] W. J. Welch, "Branch-and-bound Search for Experimental Designs based on D-Optimality and Other Criteria," *Technometrics*, vol. 24, pp. 41–48, February 1982.

[80] J.-J. Xiao, A. Ribeiro, Z.-Q. Luo, and G. B. Giannakis, "Distributed Compression-Estimation Using Wireless Sensor networks," *IEEE Signal Processing Magazine*, vol. 23, pp. 27–41, July 2006.

[81] K. Yu, J. Bi and V. Tresp, "Active Learning via Transductive Experimental Design," in *Proc. of Intl. Conf. on Machine learning*, Pittsburgh, PA, pp. 1081–1088, June 2006.

[82] F. Zhao, J. Liu, J. Liu, L. Guibas, and J. Reich, "Collaborative Signal and Information Processing: An Information-directed Approach," *Proc. of the IEEE*, vol. 91, pp. 1199–1209, August 2003.

[83] T. Zhao and A. Nehorai, "Distributed Sequential Bayesian Estimation of a Diffusive Source in Wireless Sensor Networks," *IEEE Transactions on Signal Processing*, vol. 55, pp. 1511–1524, April 2007.

[84] T. Zhao and A. Nehorai, "Information–driven Distributed Maximum Likelihood Estimation based on Gauss-Newton Method in Wireless Sensor Networks," *IEEE Transactions on Signal Processing*, vol. 55, pp. 4669–4682, September 2007.

[85] H. Zhu, I. D. Schizas, and G. B. Giannakis, "Power–efficient Dimensionality Reduction for Distributed Channel–aware Kalman Tracking using WSNs," *IEEE Transactions on Signal Processing*, vol. 57, pp. 3193–3207, August 2009.

Appendix A

Appendices: cMLE

A.1 Proof of proposition 2.1

It can be readily proved that a doubly-differentiable function of the form $f(\boldsymbol{\theta}) = \mathrm{e}^{-Q(\boldsymbol{\theta})}$ is strictly log-concave if and only if $\nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T Q(\boldsymbol{\theta}) \succeq \mathbf{0}$. This simple result will be used to show that $\sum_{k=1}^K \mathbf{h}_k \mathbf{h}_k^T \succeq \mathbf{0}$ is sufficient for log-concavity of both the Gaussian pdf and the corresponding ccdf.

Proof: Since $Q(\boldsymbol{\theta}) = (1/2\sigma^2) \sum_{k=1}^K \left(y_k^* - \mathbf{h}_k^T \boldsymbol{\theta}\right)^2$ corresponds to the negative of the exponent of the Gaussian pdf of the data in (4.1), then $\nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T Q(\boldsymbol{\theta}) = (1/\sigma^2) \sum_{k=1}^K \mathbf{h}_k \mathbf{h}_k^T \succeq \mathbf{0}$ is sufficient for log-concavity of the uncensored components of the joint pdf (2.15). What remains is to show that this condition is sufficient for log-concavity of the ccdf component in (2.15) as well. The k-th term of the ccdf component can be written as

$$\Pr[-\tau_k < y_k^* \leq \tau_k] = \int_{-\tau_k}^{\tau_k} p(y; \boldsymbol{\theta})\, dy\,.$$

It follows readily that the pdf $p(y_k^*; \boldsymbol{\theta}) := \mathcal{N}(y_k^*; \mathbf{h}_k^T \boldsymbol{\theta}, \sigma^2)$ is a log-concave function of y_k as well. From [61, Theorem 2], integration of a log-concave function preserves log-concavity. Consequently, $\ell_K(\boldsymbol{\theta})$ is strictly concave if $\sum_{k=1}^K \mathbf{h}_k \mathbf{h}_k^T \succ \mathbf{0}$. □

A.2 Gradient and Hessian in the cMLE algorithm

The first and second derivatives of $\ell_K(\boldsymbol{\theta})$, yield the gradient $\mathbf{g}_K(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}}\ell_K(\boldsymbol{\theta})$ and the Hessian $\mathbf{H}_K(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\theta}}\nabla_{\boldsymbol{\theta}}^T\ell_K(\boldsymbol{\theta})$ in (2.18a) and (2.18b), respectively, where

$$\beta_k(\boldsymbol{\theta}) := s_k\frac{\left(y_k-\mathbf{h}_k^T\boldsymbol{\theta}\right)}{\sigma^2}+\frac{(1-s_k)}{\sigma}\frac{\phi(z_{1k})-\phi(z_{2k})}{Q(z_{1k})-Q(z_{2k})} \tag{A.1a}$$

$$\gamma_k(\boldsymbol{\theta}) := -\frac{1}{\sigma^2}s_k-\frac{(1-s_k)}{\sigma^2}\left\{\left(\frac{\phi(z_{1k})-\phi(z_{2k})}{Q(z_{1k})-Q(z_{2k})}\right)^2 - \frac{z_{1k}\phi(z_{1k})-z_{2k}\phi(z_{2k})}{Q(-z_{1k})-Q(z_{2k})}\right\} \tag{A.1b}$$

$$z_{1k} := \frac{-\tau_k-\mathbf{h}_k^T\boldsymbol{\theta}}{\sigma}, \qquad z_{2k} := \frac{\tau_k-\mathbf{h}_k^T\boldsymbol{\theta}}{\sigma} \tag{A.1c}$$

$$\phi(z) := \mathcal{N}(z;0,1) = \frac{e^{-z^2/2}}{\sqrt{2\pi}}, \qquad Q(z) := \int_z^\infty \phi(u)\,du. \tag{A.1d}$$

A.3 The FIM in lemma 2.2

Using the definitions in (A.1c) and (A.1d), the pdf $p(y_k, s_k;\boldsymbol{\theta})$ in Lemma 2.1 yields

$$\nabla_{\boldsymbol{\theta}}\ln p(y_k, s_k;\boldsymbol{\theta}) = \frac{\mathbf{h}_k}{\sigma}\left\{s_k\frac{(y_k-\mathbf{h}_k^T\boldsymbol{\theta})}{\sigma}+(1-s_k)\frac{\phi(z_{1k})-\phi(z_{2k})}{Q(z_{1k})-Q(z_{2k})}\right\}$$

from which the FIM can be expressed as

$$\mathbf{I}_K(\boldsymbol{\theta}) := \mathbb{E}\left[\nabla_{\boldsymbol{\theta}}\ln p(\mathbf{y},\mathbf{s};\boldsymbol{\theta})\nabla_{\boldsymbol{\theta}}^T\ln p(\mathbf{y},\mathbf{s};\boldsymbol{\theta})\right] = \sum_{k=1}^K \bar{\gamma}_k(\boldsymbol{\theta})\mathbf{h}_k\mathbf{h}_k^T$$

which simplifies to $\mathbf{I}_K(\boldsymbol{\theta})$ in Lemma 2.1, with the parameter vector $\bar{\gamma}_k(\boldsymbol{\theta})$ given by

$$\begin{aligned}\bar{\gamma}_k(\boldsymbol{\theta}) &:= \frac{1}{\sigma^2}\left\{\mathbb{E}\left[s_k(y_k-\mathbf{h}_k^T\boldsymbol{\theta})^2\right]+\mathbb{E}[1-s_k]\frac{\left(\phi(z_{1k})-\phi(z_{2k})\right)^2}{\left(Q(z_{1k})-Q(z_{2k})\right)^2}\right\} \\ &= \frac{1}{\sigma^2}\left\{1-[Q(z_{1k})-Q(z_{2k})]+\frac{\left(\phi(z_{1k})-\phi(z_{2k})\right)^2}{Q(z_{1k})-Q(z_{2k})}-[z_{1k}\phi(z_{1k})-z_{2k}\phi(z_{2k})]\right\}. \end{aligned} \tag{A.2}$$

A.4 Gradient and Hessian in qcMLE algorithm

The first and second derivatives of $\ell_K^Q(\boldsymbol{\theta})$ lead to the gradient $\mathbf{g}_K^Q(\boldsymbol{\theta})$ and the Hessian $\mathbf{H}_K^Q(\boldsymbol{\theta})$ in (2.25a) and (2.25b), respectively, with $\beta_k^Q(\boldsymbol{\theta})$ and $\gamma_k^Q(\boldsymbol{\theta})$ defined as

$$\begin{aligned}\beta_k^Q(\boldsymbol{\theta}) &= \sum_{j\in\mathcal{J}} \frac{d_{jk}}{\sigma}\frac{\phi(\bar{z}_{1jk})-\phi(\bar{z}_{2jk})}{Q(\bar{z}_{1jk})-Q(\bar{z}_{2jk})} + \frac{(1-s_k)}{\sigma}\frac{\phi(z_{1k})-\phi(z_{2k})}{Q(z_{1k})-Q(z_{2k})}\\ \gamma_k^Q(\boldsymbol{\theta}) &= -\frac{(1-s_k)}{\sigma^2}\left[\left(\frac{\phi(z_{1k})-\phi(z_{2k})}{Q(z_{1k})-Q(z_{2k})}\right)^2 - \frac{z_{1k}\phi(z_{1k})-z_{2k}\phi(z_{2k})}{Q(z_{1k})-Q(z_{2k})}\right]\\ &\quad - \sum_{j\in\mathcal{J}}\frac{d_{jk}}{\sigma^2}\left[\left(\frac{\phi(\bar{z}_{1jk})-\phi(\bar{z}_{2jk})}{Q(\bar{z}_{1jk})-Q(\bar{z}_{2jk})}\right)^2 - \frac{\bar{z}_{1jk}\phi(\bar{z}_{1jk})-\bar{z}_{2jk}\phi(\bar{z}_{2jk})}{Q(\bar{z}_{1jk})-Q(\bar{z}_{2jk})}\right]\end{aligned}$$

where $\bar{z}_{1jk} := (t_j - \mathbf{h}_k^T\boldsymbol{\theta})/\sigma$, and $\bar{z}_{2jk} := (t_{j+1} - \mathbf{h}_k^T\boldsymbol{\theta})/\sigma$.

A.5 The FIM in quantized-censored MLE case

Using the pmf (2.24), and following similar derivation steps as in Appendix A.3, the FIM is obtained as

$$\begin{aligned}\mathbf{I}_K^Q(\boldsymbol{\theta}) &:= \mathbb{E}\left[\nabla_{\boldsymbol{\theta}}\ln\Pr\left[\boldsymbol{s},\boldsymbol{b},\{\boldsymbol{d}_j\}_{j\in\mathcal{J}};\boldsymbol{\theta}\right]\nabla_{\boldsymbol{\theta}}^T\ln\Pr\left[\boldsymbol{s},\boldsymbol{b},\{\boldsymbol{d}_j\}_{j\in\mathcal{J}};\boldsymbol{\theta}\right]\right]\\ &= \sum_{k=1}^{K}\bar{\gamma}_k^Q(\boldsymbol{\theta})\mathbf{h}_k\mathbf{h}_k^T\end{aligned}$$

where

$$\begin{aligned}\bar{\gamma}_k^Q(\boldsymbol{\theta}) := \frac{1}{\sigma^2}\Bigg\{&\frac{\left(\phi(z_{1k})-\phi(z_{2k})\right)^2}{Q(z_{1k})-Q(z_{2k})} - \left(z_{1k}\phi(z_{1k})-z_{2k}\phi(z_{2k})\right)\\ &+\sum_{j\in\mathcal{J}}\left[\frac{\left(\phi(z_{1jk})-\phi(z_{2jk})\right)^2}{Q(z_{1jk})-Q(z_{2jk})} - \left(z_{1jk}\phi(z_{1jk})-z_{2jk}\phi(z_{2jk})\right)\right]\Bigg\}.\end{aligned} \tag{A.3}$$

Appendix B

Appendices: cMAP

B.1 Proof of Proposition 3.1

The proof relies on establishing the following: (i) the prior pdf $p(\boldsymbol{\theta})$ is a log-concave function of $\boldsymbol{\theta}$, and (ii) the Gaussian pdf $p(\mathbf{y}^*|\boldsymbol{\theta})$ is log-concave, given $\sum_{k=1}^{K} \mathbf{h}_k \mathbf{h}_k^T \succeq \mathbf{0}$.

Proof: (i) Since $\mathbf{C}_{\boldsymbol{\theta}}$ is a covariance function, i.e., $\mathbf{C}_{\boldsymbol{\theta}} \succeq \mathbf{0}$, the quadratic term in (3.10) is nonnegative; (ii) A doubly-differentiable function of the form $f(\boldsymbol{\theta}) = \mathrm{e}^{-Q(\boldsymbol{\theta})}$ is strictly log-concave if and only if $\nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T Q(\boldsymbol{\theta}) \succeq \mathbf{0}$. This simple result will be used to show that the condition (i) of Proposition 3.1 is sufficient for log-concavity of both the Gaussian pdf and the corresponding ccdf.

If $Q(\boldsymbol{\theta}) = (1/2\sigma^2) \sum_{k=1}^{K} \left(y_k^* - \mathbf{h}_k^T \boldsymbol{\theta}\right)^2$ corresponding to the negative of the exponent of the Gaussian pdf of the model (4.1), then $\nabla_{\boldsymbol{\theta}} \nabla_{\boldsymbol{\theta}}^T Q(\boldsymbol{\theta}) = (1/\sigma^2) \sum_{k=1}^{K} \mathbf{h}_k \mathbf{h}_k^T \succeq \mathbf{0}$ is sufficient for log-concavity of the uncensored components of the joint pdf (3.9). What remains is to show that this condition is sufficient for log-concavity of the ccdf component in (3.9) as well.

The k-th term of the ccdf component can be written as

$$\Pr[|y_k^*| < \tau_k] = \int_{-\tau_k}^{\tau_k} p(y; \boldsymbol{\theta})\, dy\,.$$

It follows readily that the pdf $p(y_k^*; \boldsymbol{\theta}) := \mathcal{N}(y_k^*; \mathbf{h}_k^T \boldsymbol{\theta}, \sigma^2)$ is a log-concave function of y_k as well. From [61, Thrm. 2], integration of a log-concave function preserves the log-concavity. The ccdf is therefore log-concave. Consequently, $\ell^B(\boldsymbol{\theta})$ is strictly convex if either $\mathbf{C}_{\boldsymbol{\theta}} \succ \mathbf{0}$, or $\sum_{k=1}^{K} \mathbf{h}_k \mathbf{h}_k^T \succ \mathbf{0}$. □

B.2 Gradient and Hessian in the cMAP algorithm

The first and second derivatives of the objective $\ell_K^B(\boldsymbol{\theta})$ give the gradient $\mathbf{g}_K^B(\boldsymbol{\theta})$ and the Hessian $\mathbf{H}_K^B(\boldsymbol{\theta})$ in (3.12a) and (3.12b), respectively, where

$$\beta_k^B(\boldsymbol{\theta}) := -\frac{s_k}{\sigma^2}\left(y_k - \mathbf{h}_k^T\boldsymbol{\theta}\right) - \frac{(1-s_k)}{\sigma}\frac{\phi(z_{1k}) - \phi(z_{2k})}{Q(z_{1k}) - Q(z_{2k})}$$

$$\gamma_k^B(\boldsymbol{\theta}) := \frac{s_k}{\sigma^2} + \frac{(1-s_k)}{\sigma^2}\left\{\left(\frac{\phi(z_{1k}) - \phi(z_{2k})}{Q(z_{1k}) - Q(z_{2k})}\right)^2 - \frac{z_{1k}\phi(z_{1k}) - z_{2k}\phi(z_{2k})}{Q(z_{1k}) - Q(z_{2k})}\right\}.$$

B.3 The FIM in Lemma 3.2

The FIM $\mathbf{I}_K^B$ defined in (3.13) is found using the pdf $p(\mathbf{y}, \boldsymbol{s}, \boldsymbol{\theta}) = \prod_{k=1}^K p(y_k, s_k|\boldsymbol{\theta})\, p(\boldsymbol{\theta})$, whereby

$$\nabla_{\boldsymbol{\theta}} \ln p(y_k, s_k|\boldsymbol{\theta}) = \left\{\frac{s_k}{\sigma^2}\left(y_k - \mathbf{h}_k^T\boldsymbol{\theta}\right) + \frac{(1-s_k)}{\sigma}\frac{\phi(z_{1k}) - \phi(z_{2k})}{Q(z_{1k}) - Q(z_{2k})}\right\}\mathbf{h}_k$$

from which

$$\mathbf{I}_K^B(\boldsymbol{\theta}) := \mathbb{E}\left[\nabla_{\boldsymbol{\theta}} \ln p(\mathbf{y}, \boldsymbol{s}|\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}}^T \ln p(\mathbf{y}, \boldsymbol{s}|\boldsymbol{\theta})\right] = \sum_{k=1}^K \bar{\gamma}_k^B(\boldsymbol{\theta}) \mathbf{h}_k \mathbf{h}_k^T$$

$$\begin{aligned}\bar{\gamma}_k^B(\boldsymbol{\theta}) &:= \frac{\mathbb{E}\left[s_k(y_k - \mathbf{h}_k^T\boldsymbol{\theta})^2|\boldsymbol{\theta}\right]}{\sigma^2} + \frac{\mathbb{E}[(1-s_k)|\boldsymbol{\theta}]}{\sigma^2}\frac{\left(\phi(z_{1k}) - \phi(z_{2k})\right)^2}{\left(Q(z_{1k}) - Q(z_{2k})\right)^2} \\ &= \frac{1}{\sigma^2}\left\{1 - [Q(z_{1k}) - Q(z_{2k})] + \frac{\left(\phi(z_{1k}) - \phi(z_{2k})\right)^2}{Q(z_{1k}) - Q(z_{2k})} - [z_{1k}\phi(z_{1k}) - z_{2k}\phi(z_{2k})]\right\}.\end{aligned}$$

The FIM is thus given by $\mathbf{I}_K = \mathbb{E}[\mathbf{I}_K(\boldsymbol{\theta})] + \mathbf{C}_\theta^{-1} = \sum_{k=1}^K \mathbb{E}[\bar{\gamma}_k^B(\boldsymbol{\theta})]\mathbf{h}_k\mathbf{h}_k^T + \mathbf{C}_\theta^{-1}$. Numerical integration is necessary to compute the expectation with respect to $\boldsymbol{\theta}$. To this end, an approximation amounting to $\bar{\gamma}_k^B := \mathbb{E}\left[\bar{\gamma}_k^B(\boldsymbol{\theta})\right] \approx \bar{\gamma}_k^B(\mathbb{E}[\boldsymbol{\theta}]) = \bar{\gamma}_k^B(\boldsymbol{\mu}_\theta)$ is used. Such an approximation is justifiable in at least two cases: (i) if the prior $p(\boldsymbol{\theta})$ is approximately constant over the support of $\boldsymbol{\theta}$; see e.g., [45] for the justification of using a point estimate in lieu of a difficult integration in Bayesian models; and (ii) if the Taylor expansion of $\gamma_k(\boldsymbol{\theta})$ about $\boldsymbol{\mu}_\theta$ up to the quadratic term provides a sufficiently accurate approximation; that is,

$$\begin{aligned}\bar{\gamma}_k^B(\boldsymbol{\theta}) \approx \bar{\gamma}_k^B(\boldsymbol{\mu}_\theta) &+ (\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)^T \nabla_\theta \bar{\gamma}_k^B(\boldsymbol{\mu}_\theta) \\ &+ \frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)^T \nabla_\theta \nabla_\theta^T \bar{\gamma}_k^B(\mu_\theta)(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta).\end{aligned}$$

Taking expectation of $\bar{\gamma}_k^B(\boldsymbol{\theta})$ and using $u_k := (\sigma_{y_k^*}/\sigma)t$ leads to

$$\bar{\gamma}_k^B \approx \frac{2}{\sigma^2}\left[Q(u_k) - u_k\,\phi(u_k)\right]. \tag{B.1}$$

B.4 Gradient, Hessian, and FIM for the qcMAP estimator

From the first and second derivatives of the function in (3.14), the gradient and Hessian are

$$\mathbf{g}_K^{BQ}(\boldsymbol{\theta}) = \sum_{k=1}^{K} \beta_k^Q(\boldsymbol{\theta})\mathbf{h}_k + \mathbf{C}_\theta^{-1}(\boldsymbol{\theta} - \boldsymbol{\mu}_\theta)$$
$$\mathbf{H}_K^{BQ}(\boldsymbol{\theta}) = \sum_{k=1}^{K} \gamma_k^Q(\boldsymbol{\theta})\mathbf{h}_k\,\mathbf{h}_k^T + \mathbf{C}_\theta^{-1}$$

where $\beta_k^Q(\boldsymbol{\theta})$ and $\gamma_k^Q(\boldsymbol{\theta})$ are defined in Appendix A.4, whereby now $\boldsymbol{\theta}$ is considered random.

The FIM is obtained from the function $\mathbf{I}_K^{BQ}(\boldsymbol{\theta})$, which is defined as

$$\begin{aligned}\mathbf{I}_K^{BQ}(\boldsymbol{\theta}) &:= \mathbb{E}\Big[\nabla_{\boldsymbol{\theta}} \ln \Pr[\boldsymbol{s}, \boldsymbol{b}, \{\boldsymbol{d}_j\}_{j\in\mathcal{J}}|\boldsymbol{\theta}]\, \nabla_{\boldsymbol{\theta}}^T \ln \Pr[\boldsymbol{s}, \boldsymbol{b}, \{\boldsymbol{d}_j\}_{j\in\mathcal{J}}|\boldsymbol{\theta}]\Big] \\ &= \sum_{k=1}^{K} \bar{\gamma}_k^Q(\boldsymbol{\theta})\, \mathbf{h}_k\mathbf{h}_k^T\end{aligned}$$

where $\bar{\gamma}_k^Q(\boldsymbol{\theta})$ is obtained from (A.2) assuming $\boldsymbol{\theta}$ is random. Then the FIM for the quantized-censored data is obtained as $\mathbf{I}_K^{BQ} := \mathbb{E}\left[\mathbf{I}_K^{BQ}(\boldsymbol{\theta})\right] + \mathbf{C}_\theta^{-1}$, which simplifies to

$$\begin{aligned}\mathbf{I}_K^{BQ} &= \sum_{k=1}^{K} \mathbb{E}[\bar{\gamma}_k^Q(\boldsymbol{\theta})]\mathbf{h}_k\mathbf{h}_k^T + \mathbf{C}_\theta^{-1} \\ &= \sum_{k=1}^{K} \bar{\gamma}_k^{BQ}\mathbf{h}_k\mathbf{h}_k^T + \mathbf{C}_\theta^{-1}\,.\end{aligned}$$

Numerical integration is necessary to compute the expectation $\bar{\gamma}_k^{BQ} := \mathbb{E}[\bar{\gamma}_k^Q(\boldsymbol{\theta})]$ [cf. Appendix B.3].

Appendix C

Appendices: QKF

C.1 Proof of proposition 4.1: Derivation of BQKF

The proof will rely on a result known as iterated conditional expectation [15, pp. 37] which asserts that: Given random variables $\boldsymbol{\theta} \in \mathbb{R}^p, y \in \mathcal{R}_i \subset \mathbb{R}$, where $\mathcal{R}_i$ and $\mathbb{R}$ are Borel fields and a function $g(\cdot)$ defined on $\mathbb{R}^p$, it holds that

$$\mathbb{E}\{g(\boldsymbol{\theta}) \mid y \in \mathcal{R}_i\} = \mathbb{E}\{\mathbb{E}\{g(\boldsymbol{\theta}) \mid Y\} \mid y \in \mathcal{R}_i\} \tag{C.1}$$

where Y denotes a random variable in $\mathbb{R}$ and y is its realization. The proof of (C.1) is given by

$$\begin{aligned}
\mathbb{E}\{\mathbb{E}\{g(\boldsymbol{\theta})|Y\}|y \in \mathcal{R}_i\} &= \int_{\mathbb{R}} \left[\int_{\mathbb{R}^p} g(\boldsymbol{\theta}) p(\boldsymbol{\theta}|Y=\zeta) d\boldsymbol{\theta} \right] p(\zeta|y \in \mathcal{R}_i) d\zeta \\
&= \int_{\mathcal{R}_i} \left[\int_{\mathbb{R}^p} g(\boldsymbol{\theta}) p(\boldsymbol{\theta}|Y=\zeta) d\boldsymbol{\theta} \right] \frac{p(\zeta)}{\Pr\{y \in \mathcal{R}_i\}} d\zeta \qquad \text{(C.2)} \\
&= \int_{\mathbb{R}^p} g(\boldsymbol{\theta}) \left[\int_{\mathcal{R}_i} \frac{p(\boldsymbol{\theta}, \zeta)}{\Pr\{y \in \mathcal{R}_i\}} d\zeta \right] d\boldsymbol{\theta} \\
&= \int_{\mathbb{R}^p} g(\boldsymbol{\theta}) p(\boldsymbol{\theta}|y \in \mathcal{R}_i) d\boldsymbol{\theta} = \mathbb{E}\{g(\boldsymbol{\theta})|y \in \mathcal{R}_i\} \qquad \text{(C.3)}
\end{aligned}$$

where in (C.2) and (C.3) we have used the conditional pdfs

$$\begin{aligned}
p(\zeta|y \in \mathcal{R}_i) &= \begin{cases} \frac{p(\zeta)}{\Pr\{y \in \mathcal{R}_i\}}, & \text{if} \quad \zeta \in \mathcal{R}_i \\ 0, & \text{otherwise} \end{cases} \\
p(\boldsymbol{\theta}|y \in \mathcal{R}_i) &= \int_{\mathcal{R}_i} \frac{p(\boldsymbol{\theta}, \zeta)}{\Pr\{y \in \mathcal{R}_i\}} d\zeta. \qquad \blacksquare
\end{aligned}$$

Proof: Subsequent derivations will use substitutions $\boldsymbol{\theta} \leftrightarrow \boldsymbol{\theta}(n)$; $Y \leftrightarrow [\mathbf{b}_{1:n-1}, \tilde{y}(n)]$; and

$y \leftrightarrow [\mathbf{b}_{1:n-1}, \tilde{y}(n) \in \mathcal{R}_i]$. To derive $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) = \mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}, b(n)\}$ in (4.20) for $b(n) = i$ where $\mathcal{R}_i = \big[\tau_i(n), \tau_{i+1}(n)\big)$ (i.e., $\tilde{y}(n) \in \mathcal{R}_i$), let $g(\boldsymbol{\theta}(n)) = \boldsymbol{\theta}(n)$ which, from (C.1), leads to

$$\begin{aligned}&\mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1},\, b(n)=i\}\\ &= \mathbb{E}\{\,\mathbb{E}\{\boldsymbol{\theta}(n)\,|\,\mathbf{b}_{1:n-1},\, \tilde{y}(n)\}\,|\mathbf{b}_{1:n-1}, b(n)=i\}\,.\end{aligned} \tag{C.4}$$

We will first evaluate the inner expectation in (C.4) for which the pdf $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}, \tilde{y}(n)]$ is needed. If $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}] = \mathcal{N}[\boldsymbol{\theta}(n); \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}), \mathbf{M}(n|\mathbf{b}_{1:n-1})]$, where $\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}\}$ and $\mathbf{M}(n|\mathbf{b}_{1:n-1}) = \mathbb{E}\{[\boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})][\boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})]^T\}$ with $\tilde{\boldsymbol{\theta}}(n) := \boldsymbol{\theta}(n) - \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})$ and $\tilde{y}(n) := \mathbf{h}^T(n)\tilde{\boldsymbol{\theta}}(n) + v(n)$, then $p\,[\tilde{y}(n)|\mathbf{b}_{1:n-1}] = \mathcal{N}[\,\tilde{y}(n);\, 0,\, \sigma^2_{y_n}]$, where $\sigma^2_{y_n} = \mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n) + c_v(n)$. Since $\tilde{y}(n) = \mathbf{h}^T(n)\tilde{\boldsymbol{\theta}}(n) + v(n)$, the joint conditional pdf of $\boldsymbol{\theta}(n), \tilde{y}(n)$ is Gaussian, i.e.,

$$\begin{aligned}p\,[\boldsymbol{\theta}(n),\, \tilde{y}(n)|\mathbf{b}_{1:n-1}] = \mathcal{N}\Bigg[&\boldsymbol{\theta}(n), \tilde{y}(n); \begin{pmatrix}\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})\\ 0\end{pmatrix},\\ &\begin{pmatrix}\mathbf{M}(n|\mathbf{b}_{1:n-1}) & \mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)\\ \mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1}) & \sigma^2_{y_n}\end{pmatrix}\Bigg]\end{aligned} \tag{C.5}$$

and thus the inner expectation becomes

$$\begin{aligned}\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n}) &:= \mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1},\, \tilde{y}(n)\}\\ &= \mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1}\} + \mathbb{E}\{\tilde{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1},\, \tilde{y}(n)\}\\ &= \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) \quad + \mathbf{k}^c(n)\,\tilde{y}(n)\end{aligned} \tag{C.6}$$

where

$$\mathbf{k}^c(n) := \frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)}{\sigma^2_{y_n}} \tag{C.7}$$

since it is the conditional mean of jointly Gaussian random variables [cf. (C.5)] (similar to the KF in a Gauss-Markov model [34, pp. 472]). The outer expectation in (C.4) follows from (C.6) as

$$\begin{aligned}\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n}) &= \mathbb{E}\{\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})|\mathbf{b}_{1:n-1},\, b(n)=i\}\\ &= \hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1}) + \mathbf{k}^c(n)\mathbb{E}\{\,\tilde{y}(n)|\mathbf{b}_{1:n-1}, b(n)=i\}.\end{aligned} \tag{C.8}$$

The pdf $p[\tilde{y}(n)|\mathbf{b}_{1:n-1}, b(n)=i] = p[\tilde{y}(n)|\mathbf{b}_{1:n-1}, \tilde{y}(n) \in \mathcal{R}_i]$ is obtained from Bayes' theorem as

$$p[\tilde{y}(n)|\mathbf{b}_{1:n-1}, b(n)=i] = \begin{cases}\frac{p[\tilde{y}(n)|\mathbf{b}_{1:n-1}]}{\Pr\{\tilde{y}(n)\in\mathcal{R}_i|\mathbf{b}_{1:n-1}\}}, & \text{if } \tilde{y}(n)\in\mathcal{R}_i\\ 0, & \text{otherwise}\end{cases} \tag{C.9}$$

and thus the expectation in the second term of (C.8) becomes

$$\mathbb{E}\{\,\tilde{y}(n)\,|\,\mathbf{b}_{1:n-1},\,b(n)=i\}=\frac{\int_{\mathcal{R}_i}\tilde{y}(n)\;p\,[\,\tilde{y}(n)|\mathbf{b}_{1:n-1}]\;d\tilde{y}(n)}{\Pr\{\,\tilde{y}(n)\in\mathcal{R}_i\,|\,\mathbf{b}_{1:n-1}\}}\,. \tag{C.10}$$

Since $p[\tilde{y}(n)|\mathbf{b}_{1:n-1}]=\mathcal{N}[\,\tilde{y}(n);0,\sigma^2_{y_n}]$, and $\{\tilde{y}(n)\in\mathcal{R}_i\}\equiv\{b(n)=i\}$, (C.10) becomes

$$\begin{aligned}\mathbb{E}\{\tilde{y}(n)|\mathbf{b}_{1:n-1},b(n)=i\}=&\frac{\int_{\tau_i(n)}^{\tau_{i+1}(n)}\tilde{y}(n)\mathcal{N}[\tilde{y}(n);0,\sigma^2_{y_n}]d\tilde{y}(n)}{\Pr\{\tau_i(n)\leq\tilde{y}(n)<\tau_{i+1}(n)|\mathbf{b}_{1:n-1})\}}\\ =&\frac{\sigma_{y_n}}{\sqrt{2\pi}}\frac{\exp\left[-\frac{\tau_i^2(n)}{2\,\sigma^2_{y_n}}\right]-\exp\left[-\frac{\tau_{i+1}^2(n)}{2\,\sigma^2_{y_n}}\right]}{Q\left[\frac{\tau_i(n)}{\sigma_{y_n}}\right]-Q\left[\frac{\tau_{i+1}(n)}{\sigma_{y_n}}\right]}\\ =&\frac{\sigma_{y_n}}{\sqrt{2\pi}}\frac{\exp\left[-\Delta_i^2(n)/2\right]-\exp\left[-\Delta_{i+1}^2(n)/2\right]}{Q\left[\Delta_i(n)\right]-Q\left[\Delta_{i+1}(n)\right]}:=\sigma_{y_n}\alpha_i(n).\end{aligned} \tag{C.11}$$

where $\Delta_i(n):=\tau_i(n)/\sigma_{y_n}$ and the last equality follows from the definition of $\alpha_i(n)$ in (4.18). Substituting (C.11) in (C.8), $\mathbf{k}^c(n)$ from (C.7) and $\sigma^2_{y_n}:=\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)+c_v(n)$, we obtain [cf. (4.20)]

$$\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})=\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n-1})+\alpha_i(n)\,\frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\mathbf{h}(n)}{\sigma_{y_n}}\,. \tag{C.12}$$

We next derive (4.21) using $\mathbf{M}(n|\mathbf{b}_{1:n}):=\mathbb{E}\{[\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})][\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})]^T|\,\mathbf{b}_{1:n}\}$. It should be noted that given $\mathbf{b}_{1:n}$, as in most non-linear filtering, see e.g., [30], the conditional ECM defined above is different from the unconditional ECM $\bar{\mathbf{M}}(n|\mathbf{b}_{1:n}):=\mathbb{E}\{[\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})][\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})]^T\}$. In contrast, for the Gauss-Markov model in the clairvoyant KF [P2]-[C2] the ECM $\mathbf{M}(n|\mathbf{y}_{1:n})$ is independent of $\mathbf{y}_{1:n}$, i.e., $\bar{\mathbf{M}}(n|\mathbf{y}_{1:n})=\mathbf{M}(n|\mathbf{y}_{1:n})$. Derivation of $\mathbf{M}(n|\mathbf{b}_{1:n})$ uses again the iterated conditional expectation in (C.1). We first write $\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})$ using (C.6) and (C.8) as follows

$$\begin{aligned}\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})=&\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})+\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})\\ =&\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})+\mathbf{k}^c(n)[\tilde{y}(n)-\mathbb{E}\{\tilde{y}(n)|\mathbf{b}_{1:n-1},b(n)=i\}].\end{aligned} \tag{C.13}$$

With $g(\boldsymbol{\theta}(n))=[\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})][\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}(n|\mathbf{b}_{1:n})]^T$ in (C.1), the ECM $\mathbf{M}(n|\mathbf{b}_{1:n})$ can be written as

$$\begin{aligned}\mathbf{M}(n|\mathbf{b}_{1:n})=&\mathbb{E}\{g(\boldsymbol{\theta}(n))\,|\,\mathbf{b}_{1:n-1},\,b(n)=i\}\\ =&\mathbb{E}\left\{\mathbb{E}\{g(\boldsymbol{\theta}(n))\,|\mathbf{b}_{1:n-1},\,\tilde{y}(n)\}\,|\,\mathbf{b}_{1:n-1},\,b(n)=i\right\}.\end{aligned} \tag{C.14}$$

Considering first the inner expectation, we obtain upon substituting from (C.13)

$$\begin{aligned}
&\mathbb{E}\{g(\boldsymbol{\theta}(n))|\mathbf{b}_{1:n-1},\tilde{y}(n)\}\\
&=\mathbb{E}\{[\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})][\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})]^T|\mathbf{b}_{1:n-1},\tilde{y}(n)\}\\
&\quad+\mathbf{k}^c(n)\,[\tilde{y}(n)-\mathbb{E}\{\tilde{y}(n)|\mathbf{b}_{1:n-1},b(n)=i\}]\\
&\quad\times[\tilde{y}(n)-\mathbb{E}\{\tilde{y}(n)|\mathbf{b}_{1:n-1},b(n)=i\}]^T\,\mathbf{k}^{c\,T}(n). \qquad \text{(C.15)}
\end{aligned}$$

This expression follows since both $\tilde{y}(n)$ and $\mathbb{E}\{\tilde{y}(n)|\mathbf{b}_{1:n-1},b(n)\}$ are deterministic functions of the variables $\mathbf{b}_{1:n-1},\tilde{y}(n)$ and $\mathbb{E}\{[\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})]|\mathbf{b}_{1:n-1},\tilde{y}(n)\}=\mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1},\tilde{y}(n)\}-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})=\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})=\mathbf{0}$. Since $\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})=\mathbb{E}\{\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1},\tilde{y}(n)\}$, we observe that the first term of (C.15) is the covariance of $p[\boldsymbol{\theta}(n)|\mathbf{b}_{1:n-1},\tilde{y}(n)]$, which from $p[\boldsymbol{\theta}(n),\tilde{y}(n)|\mathbf{b}_{1:n-1}]$ in (C.5) is given by

$$\begin{aligned}
\mathbf{M}^c(n|\mathbf{b}_{1:n}) :=&\mathbb{E}\{[\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})][\boldsymbol{\theta}(n)-\hat{\boldsymbol{\theta}}^c(n|\mathbf{b}_{1:n})]^T\,|\mathbf{b}_{1:n-1},\,\tilde{y}(n)\}\\
=&\mathbf{M}(n|\mathbf{b}_{1:n-1})\;-\;\mathbf{k}^c(n)\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1})\,. \qquad \text{(C.16)}
\end{aligned}$$

Next, we pursue the outer expectation in (C.14) using (C.16) to obtain

$$\begin{aligned}
\mathbf{M}(n|\mathbf{b}_{1:n}) &= \mathbf{M}^c(n|\mathbf{b}_{1:n})+\mathbf{k}^c(n)\mathbb{E}\left\{[\tilde{y}(n)-\mathbb{E}\{\tilde{y}(n)|\mathbf{b}_{1:n}\}][\tilde{y}(n)-\mathbb{E}\{\tilde{y}(n)|\mathbf{b}_{1:n}\}]^T|\mathbf{b}_{1:n}\right\}\mathbf{k}^{c\,T}(n)\\
&= \mathbf{M}^c(n|\mathbf{b}_{1:n})\;+\;\mathbf{k}^c(n)\,\mathrm{var}\{\tilde{y}(n)|\,\mathbf{b}_{1:n}\}\,\mathbf{k}^{c\,T}(n)\,. \qquad \text{(C.17)}
\end{aligned}$$

The conditional variance term in (C.17) can be expressed using the pdf in (C.9) as

$$\begin{aligned}
\mathrm{var}\{\tilde{y}(n)|\,\mathbf{b}_{1:n-1},b(n)=i\} &= \mathbb{E}\{\tilde{y}^2(n)|\mathbf{b}_{1:n-1},b(n)=i\}-\mathbb{E}^2\{\tilde{y}(n)|\mathbf{b}_{1:n-1},b(n)=i\} \qquad \text{(C.18)}\\
&=\frac{\int_{\tau_i(n)}^{\tau_{i+1}(n)}\tilde{y}^2(n)p[\tilde{y}(n)|\mathbf{b}_{1:n-1}]d\tilde{y}(n)}{\mathrm{P}\{\tau_i(n)\le\tilde{y}(n)<\tau_{i+1}(n)|\mathbf{b}_{1:n-1}\}}-\mathbb{E}^2\{\tilde{y}(n)|\mathbf{b}_{1:n-1},b(n)=i\}.
\end{aligned}$$

Using integration by parts, the numerator of the first term of (C.18) is

$$\begin{aligned}
&\int_{\tau_i(n)}^{\tau_{i+1}(n)}\tilde{y}^2(n)\mathcal{N}[\tilde{y}(n);0,\sigma^2_{y_n}]d\tilde{y}(n)\\
&\quad=-\frac{\sigma_{y_n}}{\sqrt{2\pi}}\big[\tilde{y}(n)e^{-\frac{\tilde{y}^2(n)}{2\sigma^2_{y_n}}}\big]_{\tau_i(n)}^{\tau_{i+1}(n)}+\frac{\sigma_{y_n}}{\sqrt{2\pi}}\int_{\tau_i(n)}^{\tau_{i+1}(n)}e^{-\frac{\tilde{y}^2(n)}{2\sigma^2_{y_n}}}d\tilde{y}(n)\\
&=\frac{\sigma^2_{y_n}}{\sqrt{2\pi}}\left[\frac{\tau_i(n)}{\sigma_{y_n}}\exp\left(-\frac{\tau_i^2(n)}{2\,\sigma^2_{y_n}}\right)\;-\;\frac{\tau_{i+1}(n)}{\sigma_{y_n}}\exp\left(-\frac{\tau_{i+1}^2(n)}{2\,\sigma^2_{y_n}}\right)\right]\\
&\quad+\sigma^2_{y_n}\left[Q\left(\frac{\tau_i(n)}{\sigma_{y_n}}\right)-Q\left(\frac{\tau_{i+1}(n)}{\sigma_{y_n}}\right)\right]\\
&=\frac{\sigma^2_{y_n}}{\sqrt{2\pi}}\Big[\Delta_i(n)\exp\left(-\frac{\Delta_i^2(n)}{2}\right)-\Delta_{i+1}(n)\exp\left(-\frac{\Delta_{i+1}^2(n)}{2}\right)\Big]\\
&\quad+\;\sigma^2_{y_n}\left[Q\big(\Delta_i(n)\big)-Q\big(\Delta_{i+1}(n)\big)\right]
\end{aligned}$$

where $\Delta_i(n) := \tau_i(n)/\sigma_{y_n}$. The second term of (C.18) is the square of (C.11), i.e., $\sigma^2_{y_n}\alpha_i^2(n)$. Since $\Pr\{\,\tau_i(n) \le \tilde{y}(n) < \tau_{i+1}(n)\,|\,\mathbf{b}_{1:n-1}\} = Q\big(\Delta_i(n)\big) - Q\big(\Delta_{i+1}(n)\big)$, it follows that (C.18) becomes

$$\begin{aligned}\mathrm{var}\{\tilde{y}(n)|\,\mathbf{b}_{1:n-1}, b(n)=i\} =& \frac{\sigma^2_{y_n}}{\sqrt{2\pi}}\frac{\Delta_i(n)\exp\big(-\frac{1}{2}\Delta_i^2(n)\big) - \Delta_{i+1}(n)\exp\big(-\frac{1}{2}\Delta^2_{i+1}(n)\big)}{Q\big(\Delta_i(n)\big) - Q\left(\Delta_{i+1}(n)\right)} \\ &+ \sigma^2_{y_n} - \sigma^2_{y_n}\alpha_i^2(n) \\ :=& \sigma^2_{y_n}\left[1-\beta_i(n)\right] \end{aligned} \tag{C.19}$$

where the last equality follows from the definition of $\beta_i(n)$ in (4.19). Substituting $\mathbf{M}^c(n|\mathbf{b}_{1:n})$ from (C.16) and the variance in (C.19) into (C.17), we obtain

$$\begin{aligned}\mathbf{M}(n|\mathbf{b}_{1:n}) = \mathbf{M}(n|\mathbf{b}_{1:n-1}) &- \mathbf{k}^c(n)\mathbf{h}^T(n)\mathbf{M}(n|\mathbf{b}_{1:n-1}) \\ &+ \mathbf{k}^c(n)\sigma^2_{y_n}\left[1-\beta_i(n)\right]\mathbf{k}^{c\,T}(n)\,. \end{aligned} \tag{C.20}$$

Expanding $\mathbf{k}^{c\,T}(n)$ based on (C.7) we have $\mathbf{k}^c(n)\sigma^2_{y_n}\mathbf{k}^{c\,T}(n) = \mathbf{k}^c(n)\,\mathbf{h}^T(n)\,\mathbf{M}(n|\mathbf{b}_{1:n-1})$ which when substituted into (C.20), simplifies $\mathbf{M}(n|\mathbf{b}_{1:n})$ to

$$\begin{aligned}\mathbf{M}(n|\mathbf{b}_{1:n}) =& \mathbf{M}(n|\mathbf{b}_{1:n-1}) - \beta_i(n)\mathbf{k}^c(n)\mathbf{h}^T(n)\,\mathbf{M}(n|\mathbf{b}_{1:n-1}) \\ =& \mathbf{M}(n|\mathbf{b}_{1:n-1}) - \beta_i(n)\frac{\mathbf{M}(n|\mathbf{b}_{1:n-1})\,\mathbf{h}(n)\mathbf{h}^T(n)\,\mathbf{M}(n|\mathbf{b}_{1:n-1})}{\sigma^2_{y_n}} \end{aligned} \tag{C.21}$$

which is the ECM in (4.21). □

C.2 Proof of proposition 4.3: Derivation of IQKF

Proof: We first prove the predictor step [P4] of (4.44) and (4.45) using the state equation (4.39). Taking expectation with respect to $p[\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}]$, we obtain (4.44) as

$$\begin{aligned}\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) :&= \mathbb{E}\{\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}\} \\ &= \breve{\mathbf{A}}(n)\,\mathbb{E}\{\breve{\boldsymbol{\theta}}(n-1)|\mathbf{b}_{1:n-1}\} + \mathbb{E}\{\breve{\mathbf{u}}(n)|\mathbf{b}_{1:n-1}\} \\ &= \breve{\mathbf{A}}(n)\,\hat{\breve{\boldsymbol{\theta}}}(n-1|\mathbf{b}_{1:n-1}) \end{aligned}$$

where since $\mathbf{b}_{1:n-1}$ are quantized versions of $\{\breve{y}(1), \breve{y}(2), \ldots, \breve{y}(n-1)\}$, it follows that $\mathbb{E}\{\breve{\mathbf{u}}(n)|\mathbf{b}_{1:n-1}\} = \mathbb{E}\{\breve{\mathbf{u}}(n)\} = \mathbf{0}$; and also [cf. Proposition 4.3], $\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1}\}$ and $\hat{\breve{\boldsymbol{\theta}}}(n-1|\mathbf{b}_{1:n-1}) := \mathbb{E}\{\breve{\boldsymbol{\theta}}(n-1)|\mathbf{b}_{1:n-1}\}$.

To find an expression for $\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1}) := \mathbb{E}\{[\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) - \breve{\boldsymbol{\theta}}(n)][\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) - \breve{\boldsymbol{\theta}}(n)]^T\}$ we use $\hat{\breve{\boldsymbol{\theta}}}(n|\mathbf{b}_{1:n-1}) - \breve{\boldsymbol{\theta}}(n) = \breve{\mathbf{A}}(n)[\hat{\breve{\boldsymbol{\theta}}}(n-1|\mathbf{b}_{1:n-1}) - \breve{\boldsymbol{\theta}}(n-1)] - \breve{\mathbf{u}}(n)$ from (4.39). We thus

obtain (4.45) as

$$\begin{aligned}\breve{\mathbf{M}}(n|\mathbf{b}_{1:n-1}) &= \breve{\mathbf{A}}(n)\,\mathbb{E}\Big\{\big[\hat{\breve{\boldsymbol{\theta}}}(n-1|\mathbf{b}_{1:n-1})-\breve{\boldsymbol{\theta}}(n-1)\big]\big[\hat{\breve{\boldsymbol{\theta}}}(n-1|\mathbf{b}_{1:n-1})-\breve{\boldsymbol{\theta}}(n-1)\big]^T\Big\}\breve{\mathbf{A}}^T(n)\\ &\quad+\mathbb{E}\Big\{\breve{\mathbf{u}}(n)\breve{\mathbf{u}}^T(n)\Big\}\\ &= \breve{\mathbf{A}}(n)\breve{\mathbf{M}}(n-1|\mathbf{b}_{1:n-1})\breve{\mathbf{A}}^T(n)+\mathbf{C}_{\breve{\mathbf{u}}}(n)\end{aligned}$$

where $\mathbb{E}\Big\{[\hat{\breve{\boldsymbol{\theta}}}(n-1|\mathbf{b}_{1:n-1})-\breve{\boldsymbol{\theta}}(n-1)]\breve{\mathbf{u}}^T(n)\Big\}=\mathbf{0}$, since $\breve{\mathbf{u}}(n)$ is uncorrelated with $\breve{\boldsymbol{\theta}}(n-1)$ and $\mathbf{b}_{1:n-1}$.

We next derive the corrector step [C6] based on the approach used for deriving the corrector step of the more general batch quantized KF in Appendix A. The proof details the ith iterative step in (4.46)-(4.47) using $p[\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1},\mathbf{b}^{(1:i-1)}(n)] = \mathcal{N}[\breve{\boldsymbol{\theta}}(n);\hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1}),\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})]$ and the binary quantizer in (4.43). The quantization intervals are defined as $\mathcal{R}_j := \big[\tau_j(n),\tau_{j+1}(n)\big)$, $j\in\{1,2\}$ where $\tau_1(n)=-\infty$, $\tau_2(n)=0$, and $\tau_3(n)=+\infty$, similar to the quantizer of BQKF detailed in Section 4.2.

Let $\tilde{\breve{\boldsymbol{\theta}}}(n) := \breve{\boldsymbol{\theta}}(n)-\hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1})$, $\tilde{y}(n) := \breve{y}(n)-\breve{\mathbf{h}}^T(n)\hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{h}}^T(n)\tilde{\breve{\boldsymbol{\theta}}}(n)+\breve{v}(n)$, where $\breve{y}(n)$ is defined in (4.40). Since $\breve{\boldsymbol{\theta}}(n)$ and $\breve{v}(n)$ conditioned on $[\mathbf{b}_{1:n-1},\mathbf{b}^{(1:i-1)}(n)]$ are Gaussian and independent, by using $\sigma^2_{y_n} := \breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)+c_{\breve{v}}(n)$, we have

$$p[\breve{\boldsymbol{\theta}}(n),\,\tilde{y}(n)|\mathbf{b}_{1:n-1},\mathbf{b}^{(1:i-1)}(n)] = \mathcal{N}\left[\breve{\boldsymbol{\theta}}(n),\,\tilde{y}(n);\begin{pmatrix}\hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\\ 0\end{pmatrix},\right.\tag{C.22}$$
$$\left.\begin{pmatrix}\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) & \breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\\ \breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) & \sigma^2_{y_n}\end{pmatrix}\right].$$

From (C.22), it is clear that if $\tilde{y}(n)\in\mathcal{R}_j$ for $j\in\{1,2\}$, then $\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1})$ [cf. (4.46)] and $\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1})$ [cf. (4.47)] are the conditional mean and conditional covariance of $p[\breve{\boldsymbol{\theta}}(n)|\mathbf{b}_{1:n-1},\mathbf{b}^{(1:i-1)}(n),\tilde{y}(n)\in\mathcal{R}_j]$ and can be found using the iterated conditional means in the same way as for the BQKF in Appendix A. In fact, for $N=2$ (binary quantizer) with $\tau_1(n)=-\infty$, $\tau_2(n)=0$, $\tau_3=+\infty$, and $\Delta_j(n)=\tau_j(n)/\sigma_{y_n}$, we have $\Delta_1(n)=-\infty$, $\Delta_2(n)=0$ and $\Delta_3(n)=+\infty$.

Equation (C.12) then simplifies to

$$\hat{\breve{\boldsymbol{\theta}}}^{(i)}(n|\mathbf{b}_{1:n-1}) = \hat{\breve{\boldsymbol{\theta}}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) + \sqrt{\frac{2}{\pi}}\frac{\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}{\sigma_{y_n}}\,b^{(i)}(n)$$

where [cf. (4.18)] $\frac{\exp\left[-\Delta_j^2(n)/2\right]-\exp\left[-\Delta_{j+1}^2(n)/2\right]}{Q[\Delta_j(n)]-Q[\Delta_{j+1}(n)]}$ equals -2 $(+2)$ if $j=1$ $(j=2)$. This is the estimator in (4.46), since the noise variance $c_{\breve{v}}(n)=0$ in the augmented state formulation detailed in Section 4.3.1.

Likewise for the ECM $\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1})$ - see (C.21) and its derivation - we can use this binary quantizer (in place of the batch quantizer in Appendix A) and the resulting $\{\Delta_j(n)\}_{j=1}^3$ to show that $\beta_j(n) = 2/\pi$ [cf. (4.19)] which leads to the conditional ECM

$$\breve{\mathbf{M}}^{(i)}(n|\mathbf{b}_{1:n-1}) = \breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1}) - \frac{2}{\pi}\frac{\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})}{\breve{\mathbf{h}}^T(n)\breve{\mathbf{M}}^{(i-1)}(n|\mathbf{b}_{1:n-1})\breve{\mathbf{h}}(n)}$$

as given in (4.47). □

C.3 Proof of optimal binary quantizer threshold in (4.31)

Upon defining $x := \Delta(n)$ in (4.31) and $g(x) = e^{x^2}Q(x)Q(-x)$, it follows readily that

$$\arg \max_{\Delta(n)\in\mathbb{R}} \bar{\beta}(n) = \arg \min_{x\in\mathbb{R}} g(x)\,. \tag{C.23}$$

Note that $g(x)$ is symmetric about $x=0$, i.e., $g(x)=g(-x)$. To show that $x=0$ minimizes $g(x)$ and consequently $\Delta(n)=0$ maximizes $\bar{\beta}(n)$, we will prove that $g(x)$ is convex and symmetric. To this end, we use the following lemma:

Lemma C.1 *If a function $f(x)$ is convex and symmetric about $x=0$, then $f(0)$ is a minimum of $f(x)$.*

Proof: Letting $f(x)$ be convex and symmetric, we have that $f(\lambda x_1+(1-\lambda)x_2) \le \lambda f(x_1)+(1-\lambda)f(x_2)$, $\forall x_1, x_2 \in \mathbb{R}$, and $0 \le \lambda \le 1$. Also, by symmetry of $f(x)$, we have $f(x) = f(-x)$, $\forall x \in \mathbb{R}$. Setting $x_1=x$, $x_2=-x$, and $\lambda=1/2$, we obtain $f(0) \le f(x)$, $\forall x \in \mathbb{R}$. ■

Proof: We next prove that $g(x)$ is convex. From [73, pp.85, 88], $Q(x)=\frac{1}{\pi}\int_0^{\pi/2} e^{-\frac{x^2}{2\sin^2\theta}}d\theta$, and $Q^2(x)=\frac{1}{\pi}\int_0^{\pi/4} e^{-\frac{x^2}{2\sin^2\theta}}d\theta$ from which

$$\begin{aligned} g(x) =& e^{x^2}Q(x)[1-Q(x)] \\ =& e^{x^2}\left[Q(x)-Q^2(x)\right] \\ =& \frac{1}{\pi}\int_{\pi/4}^{\pi/2} e^{x^2(1-\frac{1}{2\sin^2\theta})}d\theta\,. \end{aligned}$$

Since $\forall\theta \in \left[\frac{\pi}{4}, \frac{\pi}{2}\right)$, we have $1-\frac{1}{2\sin^2\theta} \ge 0$, hence $e^{x^2(1-\frac{1}{2\sin^2\theta})}$ is convex w.r.t. x for $\theta \in \left[\frac{\pi}{4}, \frac{\pi}{2}\right)$. From [11, pp. 79], integration (or sum) of convex functions preserves convexity. Thus, $g(x)$ is convex and symmetric, and by Lemma C.1, $x=0$ minimizes $g(x)$. From (C.23), $\Delta(n)=0$ maximizes $\bar{\beta}(n)$ in (4.31). □