

The background of the cover is a dark, almost black, space filled with fluid, wavy patterns in shades of teal and blue. These patterns resemble smoke or liquid in motion, creating a sense of depth and movement. The colors transition from a deep teal at the top to a vibrant blue at the bottom, with some lighter, ethereal highlights that suggest light reflecting off the moving surfaces.

Sparse Sensing for Statistical Inference

Theory, Algorithms, and Applications

Sundeep Prabhakar Chepuri

Sparse Sensing for Statistical Inference

Theory, Algorithms, and Applications

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op
maandag 25 januari 2016 om 12:30 uur

door

Sundeep Prabhakar CHEPURI

elektrotechnisch ingenieur, Technische Universiteit Delft
geboren te Bengaluru, India.

This dissertation has been approved by the
promotors: Prof. dr. ir. G.J.T. Leus and Prof. dr. ir. A.-J. van der Veen

Composition of the doctoral committee:

Rector Magnificus	chairman
Prof. dr. ir. G.J.T. Leus	Delft University of Technology
Prof. dr. ir. A.-J. van der Veen	Delft University of Technology

Independent members:

Prof. dr. ir. R.L. Lagendijk	Delft University of Technology
Prof. dr. G.B. Giannakis	University of Minnesota, USA
Prof. Dr.-Ing. A.M. Zoubir	Technische Universität Darmstadt, Germany
Prof. dr. ir. J.-P. Linnartz	Eindhoven University of Technology
Prof. dr. ir. A. Bertrand	Katholieke Universiteit Leuven, Belgium
Prof. dr. ir. B. De Schutter	Delft University of Technology, reserve member

ISBN 978-94-6186-570-0

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without the prior permission of the author.

Copyright © 2015 by S.P. Chepuri

To my parents, for making me who I am. To Kavya, for her ∞ love and patience. To my advisors, for what you are reading.

Acknowledgments

The scientist is not a person who gives the right answers, he's one who asks the right questions.

— Claude Lévi-Strauss

Geert and Alle-Jan, thank you so much for giving me the opportunity to write this dissertation. It was a great pleasure to be your student. I always had double the privileges all along this journey.

Geert, thank you for always supporting me and believing in me, for teaching me the trade (all the *secret* ingredients of research), and for the tremendous patience and a friendly supervision. You never fail to inspire me with your talent or amaze me with your biking records, and of course you don't cease to remind me that there is always someone working harder and more efficient than me. The numerous trips with you have been absolute fun, and needless to say, they were all with great memories (be it tuna in Spain or driving in India). Thank you for the amazing last 6 years and for all the years to come.

Alle-Jan, thank you for always encouraging me, for all the positive words that make even the very complicated things sound so simple, for teaching me the rare combination of precision, detail and the big picture, and for giving me all the freedom during my PhD. You are the most pragmatic and candid person I have met so far, and you will remain a great source of inspiration for me. Thank you again for everything.

It is a great pleasure and honor for me to have Prof. Georgios B. Gianakakis, Prof. Abdelhak M. Zoubir, Prof. Jean-Paul Linnartz, Prof. Alexander Bertrand, Prof. Inald Lagendijk, and Prof. Bart De Schutter on my doctoral committee. Thank you all for taking time to read and assess the first version of these pages.

Yorgos, thank you for hosting my visit to your group, and for the nice coffee at various places. It was a great pleasure spending time at SPiNCOM and working with you. Your energy and passion for teaching and research is truly amazing. Also I would like to thank all the group members of SPiNCOM for making my winter adventure in Minnesota a wonderful one.

I have learnt a lot from each one of my collaborators. I would like to thank them all, for all their efforts and for coauthoring papers with me.

Rafael Bidarra, thank you for being my graduate school mentor and a good friend. Seyran, Jos, Gerard, Iwan, Johan, Guido, and Relja, thank you all for your support and interest during the user meetings. Ruben, thank you for all the support since my MSc days at Holst Centre.

Minaksie, thank you for all your efforts in handling the administrative stuff. Antoon, thank you for keeping all the servers and our network in good shape. I would like to express my gratitude to all the CAS members (current and past). Thank you for making the work place really friendly and fun.

All my friends in the Netherlands — family away from home, thank you for being there with me during this journey and for making my life in Holland more awesome. Shahril, thank you for giving a professional touch to the cover page.

Last but not least, the most important acknowledgement goes to my mother Indira, father Prakash, sister Souju, brother-in-law Madhu, my wife Kavya, and my father- and mother-in-law. Thank you, for the trust you put in me, for the unconditional love and care, for the patience and persistence, and for a zillion other things. Our little champs, my nephew Sriram and niece Smriti, thank you for the smiles you guys bring in. I dedicate this thesis to my wonderful family.

Sundeep
Rijswijk, December 2015.

Contents

1	Introduction	1
1.1	Pervasive sensors and data deluge	1
1.2	Distributed sensing: synchronization and localization	5
1.3	Scope and context of this thesis	8
1.4	Summary of results and outline of the thesis	11
1.5	List of publications and other contributions	16
2	Sparse Sensing	21
2.1	Introduction	21
2.2	Sparse sensing model	22
2.3	Applications	26
2.4	Benefits	26
2.5	Contrast with compressive sensing	27
2.6	Sparse sensing design	28
2.7	Discussion	33
3	Sparse Sensing for Estimation	35
3.1	Introduction	36
3.2	Sensing nonlinear observations	38
3.3	$f(\mathbf{w})$ for estimation	39
3.4	Problem statement	44
3.5	Solvers	45
3.6	Numerical example: sensor placement	55
3.7	Dependent observations	62
3.8	Discussion	66

3.A	Performance thresholds	67
3.B	Projected Newton's method	69
3.C	Power iterations	69
4	Sparse Sensing for Filtering	71
4.1	Introduction	72
4.2	Sensing time-varying observations	74
4.3	$f(\mathbf{w})$ for filtering	75
4.4	Problem statement	78
4.5	Solvers	79
4.6	Numerical example: sensor scheduling	81
4.7	Dependent observations	85
4.8	Structured signals	86
4.9	Discussion	92
4.A	Extended Kalman filter	92
5	Sparse Sensing for Detection	95
5.1	Introduction	96
5.2	Sensing conditionally distributed observations	98
5.3	Problem statement	99
5.4	$f(\mathbf{w})$ for detection	100
5.5	Solvers	107
5.6	Illustrative examples	110
5.7	Dependent observations	117
5.8	Discussion	128
5.A	Proof of Proposition 5.1	128
5.B	Upper bound on P_m	129
5.C	Proof of Proposition 5.2	130
5.D	Expressions for $f_0(\mathbf{w})$ and $f_1(\mathbf{w})$	130
6	Continuous Sparse Sensing	133
6.1	Introduction	133
6.2	Sensitivity to gridding	135
6.3	Sensing model based on binning	137
6.4	Risk function for continuous sparse sensing	139

6.5	Solver	140
6.6	Discussion	144
7	Wireless Clock Synchronization	145
7.1	Introduction	145
7.2	System model	147
7.3	Passive listening protocol	148
7.4	Estimator	150
7.5	Cramér-Rao lower bound	154
7.6	Simulations	154
7.7	Discussion	156
8	Rigid Body Localization	157
8.1	Introduction	158
8.2	Problem formulation and modeling	161
8.3	Linear least squares estimators	166
8.4	Unitarily constrained Cramér-Rao bound	172
8.5	Unitarily constrained total least squares	175
8.6	Simulation results	178
8.7	Discussion	185
8.A	Derivation of the covariance matrix \mathbf{R}_n	185
8.B	Gauss-Newton iterations on the Stiefel manifold	187
8.C	Proof of Theorem 8.2	189
8.D	Proof of Theorem 8.3	189
8.E	Proof of Theorem 8.4	191
9	Conclusions and Future Research Directions	193
9.1	Concluding remarks	193
9.2	Directions for future research	196
	Glossary	201
	Bibliography	204
	Summary	219

Samenvatting	221
Biography	225

Chapter 1

Introduction

Contents

1.1	Pervasive sensors and data deluge	1
1.2	Distributed sensing: synchronization and localization	5
1.3	Scope and context of this thesis	8
1.4	Summary of results and outline of the thesis	11
1.5	List of publications and other contributions	16

1.1 Pervasive sensors and data deluge

Every day, we are generating data in the order of a billion gigabytes. This massive volume of data comes from omnipresent sensors used in medical imaging (e.g., breast or fetal ultrasound), seismic processing (e.g., from oil or gas field exploration), environmental monitoring (e.g., pollution, temperature, precipitation sensing), radio astronomy (e.g., from radio telescopes like the square kilometre array), power networks (e.g., to monitor wind farms or other distribution grids), smart infrastructures (e.g., to monitor the condition of railway tracks or bridges), localization and surveillance platforms (e.g., security cameras or drones, indoor navigation), and so on.

The acquired data samples are stored locally and then transported to a central location (e.g., a server or cloud) to extract meaningful information

(that is, for inference). Due to an unprecedented increase in the volume of the acquired data, it is becoming increasingly challenging to locally store and transport all the data samples to a central location for data/signal processing. This is because the amount of the sampled data quickly exceeds the storage and communication capacity by several orders of magnitude. Since the data processing is generally carried out at a central location with ample computing power, mainly the sensing, storage and transportation costs form the main bottleneck. To alleviate these bottlenecks, most of the data is blindly discarded without even being examined in order to limit the memory and communication requirements, causing a serious performance loss.

In this era of data deluge, it is of paramount importance to gather only the informative data needed for a specific task. If we had some prior knowledge about the task we want to perform on the data samples, then just a small portion of that data might be sufficient to reach a desired inference accuracy, thereby significantly reducing the amount of sampled and transported data. That is to say, if the inference task is known beforehand, less data needs to be acquired. Thus, the memory and bandwidth requirements can be seriously curtailed. In addition, the cost of data collection (or sensing) can be significantly reduced, where the major factors that determine the sensing costs are the number of physical sensors (and their economical and energy costs) and the physical space they occupy when installed. So, it is evident that there is an urgent need for developing unconventional and innovative sensing mechanisms tailored for specific inference tasks to extract as much information as possible yet collecting fewer data. This leads us to the first question:

Q1. *How can task-cognition be exploited to reduce the costs of sensing as well as the related storage and communications requirements?*

This is different from the classical big data setting in which the data is already available and the question is how to mine information from that large-scale data. Our problem has close similarities to sampling, and is on the other hand only related to model information, where the data is not yet available. Given the central role of sampling in engineering sciences, answering this question will impact a wide range of applications. The basic question of interest for such applications is, how the sensing systems should be designed to minimize the amount of data acquired yet reach a desired inference performance.




Figure 1.1: Illustration of an indoor localization setup. We show the floor plan of a building (e.g., museum) with candidate locations for installing the access points. The restriction on installing the access points in only certain areas might be for security or ambience purposes.

In particular, the design questions that should be answered are related to the optimal sensor placement in space and/or time, data rate, and sampling density to reduce the sensing cost as well as to reduce the storage and communications requirements. We next illustrate two specific examples of sensor placement for indoor localization and temperature sensing.

Example 1.1 (Target localization). *Indoor localization is becoming increasingly important in many applications. Some examples include: locating people inside a building for rescue operations, monitoring logistics in a production plant, lighting control, and so on. In such environments, global positioning system (GPS) signals are typically unavailable. Thus, other types of measurements such as visual, acoustic or radio waves revealing information about range, bearing, and/or Doppler are used. These measurements are gathered by access points, like cameras, microphones, radars, or wireless transceivers. One such scenario is illustrated in Figure 1.1, where we show an indoor localization setup for navigating a visitor inside a building. An interesting question is, instead of installing*

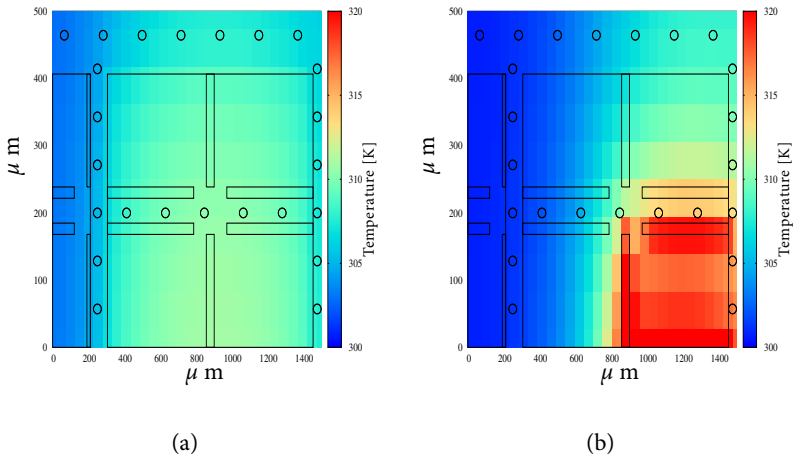


Figure 1.2: Heatmaps of a 32KB data cache (a) without and (b) with a hot spot. Black circles (\circ) denote the candidate temperature sensor locations— these are the areas with less or no active logic.

many such costly access points randomly, how can we minimize the number of access points (hence, the amount of data), by optimizing their characteristics (e.g., their spatial position, sampling rate) in such a way that a certain localization performance can be guaranteed.

Example 1.2 (Field detection). Consider a multi-core processor with a hot spot. A historical question of interest is to estimate the thermal distribution, for instance, by interpolating noisy measurements. In some applications, though, a precise estimation of the temperature field might not be required, instead, detecting the hot spots (i.e., the areas where the temperature exceeds a certain threshold) would be sufficient for subsequent control actions. Such a scenario is illustrated in Figure 1.2¹, where the image on the right (left) shows a 32 KB data cache with (no) hotspots. An important question of interest for such detection problems then is, how to design spatial samplers (i.e., sensor placement [Memik et al., 2008]) by exploiting the knowledge of the underlying model, physical space

¹We would like to thank Sumeet Kumar for the heatmaps [Kumar et al., 2015] .

and processing limitations.

Such optimally designed sensing systems can be used to perform a number of inference tasks. In the next section, we will introduce sensor networks as a spatial sampling device and discuss some related signal processing applications.

1.2 Distributed sensing: synchronization and localization

Over the past decade, advances in wireless sensor technology have enabled the usage of sensors to connect almost everything as a network. This so-called internet of things (IoT) is used for different purposes related to sensing, monitoring, and control. Such networks find applications ranging from monitoring natural ecosystems to buildings, industrial equipments, and vehicles, from military to civil localization applications, to name a few.

Due to the inherent discrete nature of the sensors—spatially localized objects, a sensor network performs sampling in space [Gastpar et al., 2006]. For many cases that we frequently encounter, a sensor network can be designed to faithfully represent distributed signals (e.g., a spatially varying phenomenon such as a temperature field). In addition, the distributed signals can be multi-dimensional, that is, they can exist in space and time. To acquire spatiotemporal variations of such distributed signals, we need to sample over both space and time, where the temporal sampling is achieved using analog-to-digital converters (ADCs) or time-to-digital converters (TDCs), for example. Each sensor has an independent sample clock, and its stability essentially determines the alignment of the temporal sampling grid across the sensors. This temporal sampling grid is perfectly aligned if all the sensors share a common clock. However, when the clocks are uncommon, the sample clocks drift from each other due to imperfections in the clock oscillator, aging and other environmental factors (e.g., temperature variation, vibration). We illustrate two sample clocks drifting from an ideal (or a reference) clock in Figure 1.3. This drift will result in the misalignment of the temporal sampling grid across the sensors. Therefore, we need to align them from time to time. In other words,

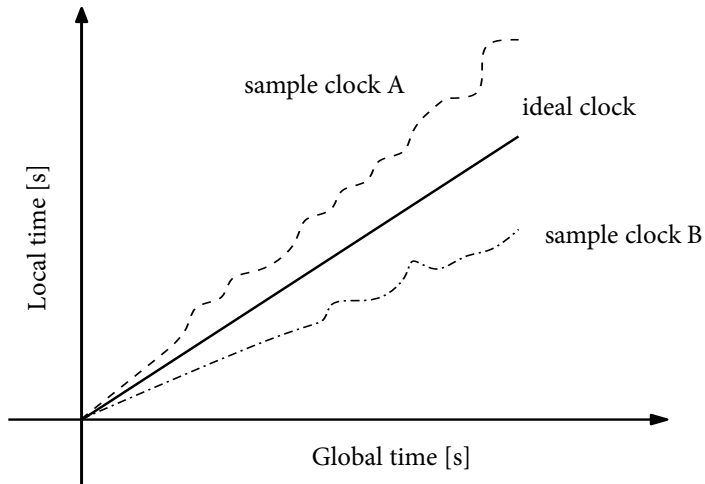


Figure 1.3: Illustration of two sample clocks drifting from each other with respect to an ideal clock.

we need to devise a mechanism to distribute the sample clock wirelessly. This brings us to the second question:

Q2. *How can wireless communications be exploited to synchronize spatially separated sample clocks?*

Answering this question impacts a range of other sensor network applications that demand for a common time frame for the entire network, such as sleep and wake-up coordination, time-based channel access, among others [Freris et al., 2010].

A vast number of the applications that use sensor networks rely on a fundamental aspect of either associating the location information to the data that is acquired by spatially distributed sensors (e.g., air quality measurements), or the acquired data is solely used to localize a target/source (e.g., indoor localization). One way to do this is to equip each sensor node with a GPS receiver, however, in many applications of interest the operating environment is harsh with GPS signals either being impaired or unavailable. Moreover, sensors are usually battery powered making GPS a less viable option. To facilitate

low-power and efficient localization solutions especially in GPS-denied environments, there exists a plethora of algorithms based on two localization paradigms: absolute or relative localization. In absolute localization, the aim is to estimate the absolute position of the sensor(s) using a few reference nodes whose absolute positions are known. Hence, these reference nodes are commonly referred to as *anchors*. Absolute localization problems are typically solved using measurements related to the propagation of radio or acoustic waves, e.g., time of arrival (TOA), time difference of arrival (TDOA), received signal strength (RSS), or angle of arrival (AOA), to list a few [Patwari et al., 2005, Gezici et al., 2005, Gustafsson and Gunnarsson, 2005]. Localization can also be relative, in which case the aim is to estimate the constellation of the sensors or the topology of the network, and determining the location of a sensor relative to the other sensors is sufficient. Classic solutions to relative localization are based on multi-dimensional scaling (MDS) using squared-range measurements [Cheung and So, 2005, Costa et al., 2006]. For relative localization, anchors are not needed.

The use of sensor networks to remotely monitor hazardous environments that are beyond human reach (e.g., leakage in oil pipes, surveillance of nuclear plants, or health of industrial machines) is gaining strong interest. Such tasks are generally performed using robots or drones (more generally, a sensing platform) with a number of sensors mounted on them. Consequently we now have to localize more than one sensor or even localize the whole sensing platform, and typically we know beforehand the sensor placement on the platform. Such a setting creates a need to solve the following problem:

Q3. *How can we extend the classical localization paradigm to localize a sensing platform by exploiting the knowledge of the sensor placement on the platform?*

We next illustrate the aforementioned problems (i.e., Q2 and Q3) with the following scenario.

Example 1.3 (Sensors on a platform). *Consider a number of sensors mounted on a moving platform as shown in Figure 1.4. These autonomous sensors collect data related to different physical phenomena, like temperature, vibration, pressure, and so on. This data has to be shipped to a central location (having several*




Figure 1.4: Illustration of distributed sensing with sensors mounted on a (moving) sensing platform. The sensor data has to be shipped to a central location with possibly multiple receivers mounted on a fixed world.

wireless transceivers —anchors) on a fixed platform. This is a typical setup in industrial machines or robots, where we generally know the sensor placement on the platform. However, the absolute position of the platform might not be known. Now the questions of interest are: can we use the available wireless links between the sensors and anchors to (A) synchronize the sample clocks and (B) localize such rigid platforms?

In the next section, we will discuss the context of this thesis and also pose the urgent questions that are addressed in this thesis.

1.3 Scope and context of this thesis

The research for this thesis was generously sponsored by the following two NWO/STW projects:

- VICI-SOWN: The VICI project on signal processing for self-organizing wireless networks (SOWN) aims at developing new mathematical sig-

nal processing tools for energy-efficient distributed information processing, spectral sensing, and localization in large sensor networks.

- FASTCOM: The reliable and fast wireless communication for lithography machines (FASTCOM) project aims to connect a sensor network on a moving platform to a control unit using high-speed links with low latency. To realize such a network, accurate sample clock synchronization and optimal sensing design to collect as little data as possible are needed.

Within the framework of the above two projects, we next pose the subquestions related to Q1, Q2, and Q3 that we have answered in this thesis. These subquestions are of general interest (impacts current hot topics like big data analytics, compressive sensing, internet of things) and goes far beyond the scope of VICI-SOWN and FASTCOM.

In order to reduce the sensing and other related costs, it is crucial to tailor the sensing mechanism for the specific inference task that will be performed on the acquired data samples. The tool that we will exploit in this thesis to reduce the cost of sensing is *sparse sensing*, which consists of an optimally designed structured and deterministic sparse (i.e., with many zeros and a few nonzeros) sensing function that is used to acquire the data in order to reach a desired inference performance. Here, the number of nonzeros determines the amount of data samples acquired (thus determines the amount of data reduction). This naturally leads to a number of questions related to the definition of the inference task and the related performance metric, which we pose as the following subquestions of Q1:

- Q1.1. How do we model sparse sensing functions to carry out fundamental signal processing tasks, like estimation, filtering, and detection?
- Q1.2. What are the reasonable inference performance metrics for the above tasks?
- Q1.3. Can we efficiently optimize (e.g., using a polynomial time algorithm) such inference performance metrics to obtain the sparse samplers of interest?

We will answer the above subquestions under the assumption that the data (i.e., actual measurements) is not yet available and that the model information is perfectly known.

Next, for a network comprising of several sensor nodes with independent clock oscillators, we aim to distribute the clock signal wirelessly. In other words, we address the problem of synchronizing the sample clocks of the sensor nodes in a network. The assumption is that there are several sensor units at known relative locations (absolute locations are, however, not known) and one sensor unit at an unknown location. The sensors have unreliable and uncommon clocks, except for one of them, which has a relatively stable clock. The goal is to estimate the clock deviations using time-of-flight measurements of messages. To this end, we pose the following subquestions of Q2:

- Q2.1. What is a reasonable parametric representation for the clock deviations?
- Q2.2. How can we fully exploit the broadcast nature of the wireless channel for clock synchronization?
- Q2.3. Is there an efficient estimator (e.g., unbiased and linear) to resolve the clock parameters and what are the theoretical limits on the variation of the estimates, that is, what is the Cramér-Rao bound (CRB)?

Finally, we consider the problem of localizing a sensing platform using sensor networks. The assumption is that the fixed world has several transceivers with known locations (anchors) and the sensing platform has several sensor units at known relative locations, and the platform is rigid; cf. Figure 1.4. The aim is to localize, that is, to estimate the position and orientation of the rigid platform using distance measurements. We refer to this problem as *rigid body localization*. The related subquestions of Q3 are:

- Q3.1. Is there a parametric representation for the rigid body localization problem?
- Q3.2. What are the theoretical limits (e.g., CRB) on the variation of the position and orientation estimates computed using distance measurements?

Q3.3. Can we solve the rigid body localization problem if the known sensor topology is perturbed, that is, if the body is not fully rigid?

In the next section, we will discuss the main results and organization of this thesis.

1.4 Summary of results and outline of the thesis

This thesis is organized into three parts. In the first part of this thesis (i.e., in Chapters 2—6), the theory and algorithms of sparse sensing are discussed in depth. In particular, we address subquestions Q1.1-Q1.3. In the second part of this thesis (i.e., in Chapter 7 and Chapter 8), applications of distributed sensing, more specifically, wireless clock synchronization and localization of a rigid platform are studied (here, we address subquestions Q2.1-Q2.3 and Q3.1-Q3.3). Finally, the thesis concludes with the third part (i.e., Chapter 9), where we pose some interesting open problems for future research. The content of Chapters 3—8 is published as papers, however with some new subtopics. The relation between these chapters and the publications is depicted in Table 1.1, while the list of publications is provided in §1.5.

Chapter 2 on *sparse sensing* fairly forms the backbone of the first part of this thesis. In this chapter, we will model the sparse sensing function as a linear projection operation, where the sensing function is parameterized by a sparse vector. This vector is basically a design parameter that is used as a handle to trade the amount of acquired data samples with the inference performance. We refer to this sparse sensing scheme as *discrete sparse sensing*, as the continuous observation domain is first discretized into grid points and we select (using the sparse vector) the best subset out of those grid points. To harness the full potential of sparse sensing, we need to sample in between the grid points and take samples anywhere in the continuous observation domain. We refer to such sensing mechanisms as *continuous sparse sensing*. We will discuss some applications of the proposed sparse sensing mechanisms and also list major differences with the state of the art in data reduction, that is, compressed sensing. Although the inference task is kept abstract in

	Chapter 3	Chapter 4	Chapter 5	Chapter 6	Chapter 7	Chapter 8
J1	•					
J2			•			
J3				•		
J4					•	
J5						•
C1	•	•	•			
C2	•					
C3	•					
C4		•				
C5		•				
C6			•			
C7						•
C8					•	
C9						•

Table 1.1: Connection between the papers and chapters.

this chapter, the obtained novel unifying view allows us to jointly treat sparse sensing mechanisms for different inference tasks considered in Chapters 3—6.

Chapter 3 focuses on discrete sparse sensing for a general *nonlinear estimation* problem. In particular, we solve the problem of choosing the best subset of observations that follow known nonlinear models with arbitrary yet independent distributions. We also extend this framework to nonlinear colored Gaussian observations as it is more suitable when the observations are subject to external noises or interference. The data is acquired using the discrete sparse sensing function, which is guided by a sparse vector. The CRB is used as an inference performance metric and we derive several functions of the CRB that include the sparse vector. To compute the sparse samplers, we propose convex relaxations of the derived inference performance metric and also develop low-complexity solvers. In sum, the discrete sparse samplers

for nonlinear inverse problems can be computed by solving a convex program.

Chapter 4 extends the theory developed in Chapter 3 to *nonlinear filtering* problems, that is, the focus will be on the design of discrete sparse sensing functions for systems that admit a known nonlinear state-space representation. In particular, we solve the problem of choosing the best subset of time-varying observations based on the entire history of measurements up to that point. The posterior CRB is used as the inference performance metric to decide on the best subset of observations. Although this framework is valid for independent observations that follow arbitrary distributions (e.g., non-Gaussian), we also extend it to colored Gaussian observations. Further, we introduce some additional constraints to obtain smooth sensing patterns over time. Finally, we devise sparse sensing mechanisms for *structured* time-varying observations (e.g., for time-varying sparse signals). In all these cases, the discrete sparse samplers can be designed by solving a convex program.

Chapter 5 is dedicated to discrete sparse sensing for *statistical detection*. Specifically, the aim is to choose the best subset of observations that are conditioned on the hypothesis, which belongs to a binary set. Naturally, the best subset of the observations is the one that results in a desired global error probability. Since the numerical optimization of the error probabilities is difficult, we adopt simpler costs related to distance measures between the conditional distributions of the sensor observations. We design sparse samplers for the Bayesian and Neyman-Pearson setting, where we respectively use the Bhattacharyya distance and Kullback-Leibler distance (and J-divergence) as the inference performance metric. For conditionally independent observations, we give an explicit solution, which is optimal in terms of the error exponents. More specifically, the best subset of observations is the one with the smallest local average root-likelihood ratio and largest local average log-likelihood ratio in the Bayesian and Neyman-Pearson setting, respectively. We supplement the proposed framework with a thorough analysis for Gaussian observations, including the case when the sensors are conditionally dependent, and also provide examples for other

observation distributions. One of the results shows that, for nonidentical Gaussian sensor observations with uncommon means and common covariances under both hypotheses, the number of sensors required to achieve a desired detection performance reduces significantly as the sensors become more coherent.

Chapter 6 contrasts with the discrete sparse sensing mechanisms that have been considered in Chapter 3 to Chapter 5, where the sparse sensing functions are parameterized by a discrete sparse vector that needs to be optimally designed. This basically means that the continuous observation domain is first discretized into grid points and we have to select the best subset out of those grid points. However, this discretization might be very coarse because of complexity reasons, preventing the system to achieve the best possible compression rates for the considered inference task. Therefore, in this chapter, we introduce *continuous sparse sensing* (or off-the-grid sparse sensing), where it is possible to sample in between the grid points and take samples anywhere in the continuous observation domain. The basic idea is to start from a discretized sampling space and to model every sampling point in the continuous sampling space as a discrete sampling point plus a perturbation. Then, the smallest set of possible discrete sampling points is searched for, along with the best possible perturbations, in order to reach the desired inference performance. We will demonstrate this approach for linear inverse problems, that is, for linear estimation problems with additive Gaussian noise.

Chapter 7 addresses subquestions of Q2 related to distributed sampling. In particular, this chapter is dedicated to *wireless clock synchronization*. To realize this, we assume an affine clock model, that is, we approximate the clock deviations using phase offset (or clock offset) and frequency offset (or clock skew), where we ignore the higher order terms like the frequency drift and so forth. In other words, we approximate the sample clock curves in Figure 1.3 with a piecewise straight line (within each observation interval) having slope and offset equal to the clock skew and clock offset, respectively. The assumption is that there are several anchor nodes with known relative locations and one sensor node with

an unknown position. Further, all the nodes have unreliable and uncommon clocks, except for one node that has a relatively stable clock (that is, one of the nodes has no clock offset and a clock skew equal to one). We estimate these clock parameters using time-of-flight measurements of messages. To fully harness the broadcast nature of the wireless medium, we allow all the nodes to passively listen to the messages and record time stamps. By doing so, we collect a significant amount of extra measurements, which we solve using a least squares estimator. Specifically, we solve for all the unknown clock skews and clock offsets along with the pairwise distances (i.e., ranges) of the sensor to each anchor. The proposed estimator is shown to be efficient, asymptotically meeting the theoretical CRB.

The proposed framework can be extended to jointly resolve unknown clock parameters and locations (instead of ranges) — a problem pertinent to time-based sensor network localization; see **C8** for the related results and the CRB. These results are, however, excluded in this thesis for the sake of conciseness.

Chapter 8 provides a framework for joint *position and orientation estimation* of a rigid platform. We consider a setup in which a few sensors are mounted on a rigid body. The absolute position of the rigid body is not known. However, we know how the sensors are mounted on the rigid body, i.e., the sensor topology is known. The rigid body is localized using noisy distance measurements between the sensors and a few anchors (nodes with known absolute positions), and without using any inertial measurements. We model the rigid body localization problem using an unknown rotation matrix and a translation vector that uniquely determine the orientation and position of the rigid platform, respectively. We propose a least squares, and a number of constrained least squares estimators, where the constrained estimators solve an optimization problem on the Stiefel manifold. As a benchmark, we derive a unitarily constrained CRB. Finally, the known topology of the sensors can be perturbed during fabrication or if the body is not entirely rigid. To take these perturbations into account, constrained total least squares estimators are also proposed.

We can further track the position and orientation of the rigid body using a state-space representation and a (constrained) Kalman filter; see **C9** for details. The results on tracking are, however, excluded here for the sake of conciseness.

Chapter 9 contains the conclusions and outlines a number of directions for future research along with some open problems.

1.5 List of publications and other contributions

The research work done for this thesis has resulted in the following journal papers, conference papers, and internal reports.

Included publications

Journal papers

- J1** S.P. Chepuri and G. Leus, “Sparsity-Promoting Sensor Selection for Non-linear Measurement Models”, *IEEE Trans. on Signal Processing*, vol. 63, no. 3, pp. 684–698, Feb. 2015.
- J2** S.P. Chepuri and G. Leus, “Sparse Sensing for Distributed Detection”, *IEEE Trans. on Signal Processing (To appear)*, Oct. 2015.
- J3** S.P. Chepuri and G. Leus, “Continuous Sensor Placement”, *IEEE Signal Processing Letters*, vol. 22, no. 5, pp. 544–548, May 2015.
- J4** S.P. Chepuri, R.T. Rajan, G. Leus, and A.-J. van der Veen, “Joint Clock Synchronization and Ranging: Asymmetrical Time-Stamping and Passive Listening”, *IEEE Signal Processing Letters*, vol. 20, no. 1, pp. 51 - 54, Jan. 2013.
- J5** S.P. Chepuri, G. Leus, and A.-J. van der Veen, “Rigid Body Localization Using Sensor Networks”, *IEEE Trans. on Signal Processing*, vol. 62, no. 18, pp. 4911 - 4924, Sep. 2013.

Conference papers

- C1 S.P. Chepuri** and G. Leus, “Sensor Selection for Estimation, Filtering, and Detection”, in *Proc. of the International Conference on Signal Processing and Communications (SPCOM 2014)*, Bangalore, India, Jul. 2014.
- C2 S.P. Chepuri**, G. Leus, and A.-J. van der Veen, “Sparsity-Exploiting Anchor Placement for Localization in Sensor Networks”, in *Proc. of the European Signal Processing Conference (Eusipco 2013)*, Marrakech, Morocco, Sep. 2013.
- C3 S.P. Chepuri** and G. Leus, “Sparse Sensing for Estimation with Correlated Observations”, in *Proc. of Asilomar Conf. Signals, systems, and Computers (Asilomar 2015)*, Pacific Grove, California, USA, Nov. 2015.
- C4 S.P. Chepuri** and G. Leus, “Compression schemes for time-varying sparse signals”, in *Proc. of Asilomar Conf. Signals, systems, and Computers (Asilomar 2014)*, Pacific Grove, California, USA, Nov. 2014.
- C5 S.P. Chepuri** and G. Leus, “Sparsity-Promoting Adaptive Sensor Selection for Non-linear Filtering”, in *Proc. of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2014)*, Florence, Italy, May 2014.
- C6 S.P. Chepuri** and G. Leus, “Sparse Sensing for Distributed Gaussian Detection”, in *Proc. of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2015)*, Brisbane, Australia, Apr. 2015, **(Best student paper award.)**
- C7 S.P. Chepuri**, G. Leus, and A.-J. van der Veen, “Position and orientation estimation of a rigid body: rigid body localization”, in *Proc. of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2013)*, Vancouver, Canada, May 2013.

Other related contributions

Journal papers

- J6** S. Khademi, **S.P. Chepuri**, Z. Irahhauten, G.J.M. Janssen, and A.-J. van der Veen, “Channel Measurements and Modeling for a 60 GHz Wireless Link Within a Metal Cabinet”, *IEEE Trans. on Wireless Communications*, vol. 14, no. 9, pp. 5098-5110, Sep. 2015.
- J7** G. Kail, **S.P. Chepuri**, and G. Leus. “Robust Censoring Using Metropolis-Hastings Sampling”, *IEEE Journal of Sel. Topics in Signal Processing (To appear)*, Mar. 2016.
- J8** S. Liu, **S.P. Chepuri**, M. Fardad, E. Masazade, G. Leus, and P.K. Varshney. “Sensor Selection for Estimation with Correlated Measurement Noise”, *IEEE Trans. on Signal Processing (Submitted)*, Aug. 2015.
- J9** S. Maleki, **S.P. Chepuri**, and G. Leus. “Optimization of Hard Fusion Based Spectrum Sensing for Energy-Constrained Cognitive Radio Networks”, *Elsevier Physical Communication*, vol. 9, pp. 193-198, Dec. 2013.

Conference papers

- C8** **S.P. Chepuri**, G. Leus, and A.-J. van der Veen, “Joint Localization and Clock Synchronization for Wireless Sensor Networks”, in *Proc. of Asilomar Conference on Signals, systems, and Computers (Asilomar 2012)*, Pacific Grove, California, USA, Nov. 2012.
- C9** **S.P. Chepuri**, A. Simonetto, G. Leus, and A.-J. van der Veen, “Tracking Position and Orientation of a Mobile Rigid Body”, in *Proc. of the IEEE Workshop on Comp. Adv. in Multi-Sensor Adaptive Proc. (CAMSAP 2013)*, St. Maarten, French Antilles, Dec. 2013.
- C10** K. Hu, **S.P. Chepuri**, and G. Leus, “Near-Field Source Localization Using Sparse Recovery Techniques”, in *Proc. of the International Conference on Signal Processing and Communications (SPCOM 2014)*, Bangalore, India, Jul. 2014.

- C11** K. Hu, **S.P. Chepuri**, and G. Leus, “Near-Field Source Localization: Sparse Recovery Techniques and Grid Matching”, in *Proc. of the Eighth IEEE Sensor Array and Multichannel Signal Processing Workshop (SAM 2014)*, A Corua, Spain, Jun. 2014, (**Finalist best student paper award contest.**)
- C12** S. Khademi, **S.P. Chepuri**, Z. Irahauten, G.J.M. Janssen, and A.-J. van der Veen, “Channel Characterization for Wideband 60 GHz Wireless Link Within a Metal Enclosure”, in *Proc. of the 8th European Conference on Antennas and Propagation (EuCAP 2014)*, The Hague, Netherlands, Apr. 2014.
- C13** V. Roy, **S.P. Chepuri**, and G. Leus, “Sparsity-Enforcing Sensor Selection for DOA Estimation”, in *Proc. of the IEEE Workshop on Comp. Adv. in Multi-Sensor Adaptive Proc. (CAMSAP 2013)*, St. Maarten, French Antilles, Dec. 2013.
- C14** S. Khademi, **S.P. Chepuri**, G. Leus, and A.-J. van der Veen, “Zero-forcing pre-equalization with transmit antenna selection in MIMO systems”, in *Proc. of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2013)*, Vancouver, Canada, May 2013.
- C15** G. Kail, **S.P. Chepuri**, and G. Leus, “Robust Censoring For Linear Inverse Problems”, in *Proc. of the IEEE 16th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC 2015)*, Stockholm, Sweden, Jun. 2015.
- C16** **S.P. Chepuri**, Y. Zhang, G. Leus, and G.B. Giannakis, “Big Data Sketching with Model Mismatch”, in *Proc. of the Asilomar Conf. Signals, systems, and Computers (Asilomar 2015)*, Pacific Grove, California, USA, Nov. 2015.
- C17** S. Rao, **S.P. Chepuri**, and G. Leus, “Greedy Sensor Selection for Non-Linear Models”, in *Proc. of the IEEE Workshop on Comp. Adv. in Multi-Sensor Adaptive Proc. (CAMSAP 2015)*, Cancun, Mexico, Dec. 2015.
- C18** S. Rao, **S.P. Chepuri**, and G. Leus, “DOA Estimation Using Sparse Vector Sensor Arrays”, in *Proc. of IEEE Workshop Workshop on Comp. Adv.*

in Multi-Sensor Adaptive Proc. (CAMSAP 2015), Cancun, Mexico, Dec. 2015.

- C19** A. Pizzo, **S.P. Chepuri**, and G. Leus, “Towards Multi-Rigid Body Localization”, Submitted to *International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2016)*, Shanghai, China, Mar. 2016.

Internal reports

- R1** **S.P. Chepuri** and A.-J. van der Veen, “Clock requirements and synchronization for the FASTCOM project”, March 25, 2014, 12 pages, (available on request.)

Chapter 2

Sparse Sensing

Contents

2.1	Introduction	21
2.2	Sparse sensing model	22
2.3	Applications	26
2.4	Benefits	26
2.5	Contrast with compressive sensing	27
2.6	Sparse sensing design	28
2.7	Discussion	33

2.1 Introduction

In this chapter we provide the theory of sparse sensing. This includes modeling as well as designing the sensing function that is used for gathering data (or sensing) to carry out a specific signal processing task. The sensing functions are designed knowing beforehand the inference task we want to perform on the data. The inference task could be as general as *estimation*, *filtering*, or *detection*, which are fundamental to statistical signal processing. The main aim of sparse sensing is therefore to exploit the knowledge of the inference task to




Figure 2.1: Continuous sparse sensing scheme. Here, we show an illustration, where we acquire $K = 3$ samples of a one-dimensional signal $x(t)$.

be performed on the acquired data in order to significantly reduce the sensing cost as well as reduce the storage and communications requirements yet assuring a desired inference quality.

We discuss two flavors of sparse sensing, namely, continuous and discrete sparse sensing in this chapter. We also present some applications and benefits of sparse sensing, and list the major differences of sparse sensing with compressed sensing—a popular tool for sensing cost reduction.

2.2 Sparse sensing model

Let $x(t)$ be a continuous-domain signal, where $t \in [0, T]$ denotes the sampling domain. The sampling domain can be space, time, or space-time and can be even of higher dimensions. For example, t could be p -dimensional, where $p = 4$ represents a (three-dimensional) spatio-temporal sampling domain.

We assume that the observation signal $x(t)$ follows a known model, which

relates the observation to the state of nature (e.g., through a parametric model or conditional distributions under different hypotheses) that we want to infer along with its statistical dependence on noise, if any.

In practice, the inference problem is typically solved in a digital fashion, that is, using a sampled version of the observation process. In this context, we are interested in the following fundamental question:

What are the best indices $\{\tau_m\}$ to optimally sample $x(t)$ to form $\{y_m = x(\tau_m)\}$ such that a desired inference performance is achieved?

This is tantamount to applying a continuous-domain *sparse sensing* function (hence the name) $w(t)$, modeled as a sum of Diracs

$$w(t) = \sum_{m=1}^K \delta(t - \tau_m) \quad (2.1)$$

that we apply on $x(t)$ to acquire

$$y(t) = w(t)x(t) = \sum_{m=1}^K x(\tau_m)\delta(t - \tau_m), \quad (2.2)$$

where we jointly design the *unknown* indices $\{\tau_m\}$, and the number of samples K , required in order to reach a desired inference performance. We label such a sensing mechanism as *continuous sparse sensing*; see the illustration in Figure 2.1. The sensing operation $w(t)$ is designed keeping in mind the known inference task that needs to be performed, and is related to it.

A way to design the continuous-domain sparse $w(t)$ is to discretize the sampling domain, and to assume that the indices $\{\tau_m\}_{m=1}^K$ lie on a discrete grid. In other words, we assume a set of $M \gg K$ candidate sampling locations $\{t_m\}_{m=1}^M$, and we alternatively model $y(t)$ as

$$y(t) = \sum_{m=1}^M w_m x(t_m)\delta(t - t_m), \quad (2.3)$$

where $w_m = (0)1$ indicates whether sample $x(t_m)$ is (not) selected.

The vector \mathbf{w} obtained by collecting $\{w_m\}_{m=1}^M$ as

$$\mathbf{w} = [w_1, w_2, \dots, w_M]^T \in \{0, 1\}^M$$

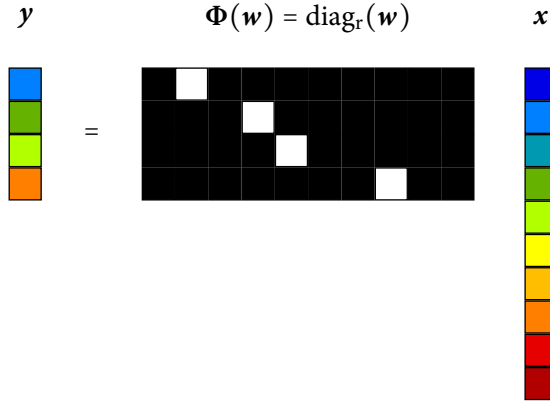


Figure 2.2: Discrete sparse sensing scheme. Here, a white (black) and colored square represents a one (zero) and an arbitrary value, respectively.

is used to construct a sensing matrix $\Phi(\mathbf{w}) = \text{diag}_r(\mathbf{w})$, which is then applied to the discrete signal $\mathbf{x} = [x(t_1), x(t_2), \dots, x(t_M)]^T$, to obtain a discrete-domain counterpart of (2.3) given as

$$\mathbf{y} = \Phi(\mathbf{w})\mathbf{x} = \text{diag}_r(\mathbf{w})\mathbf{x}, \quad (2.4)$$

where $\mathbf{y} = [y(t_1), y(t_2), \dots, y(t_M)]^T$. We term such a sensing mechanism as *discrete sparse sensing*; see the illustration in Figure 2.2. Thus, the design of a sparse function $w(t)$, can be simplified to the design of a sparse vector \mathbf{w} . Formally, we pose the question:

What is the sparsest \mathbf{w} to optimally sense \mathbf{x} to form \mathbf{y} such that a desired inference performance is achieved?

Sparse sensing does not necessarily mean that the sensing matrix should itself be sparse. That is, sparse sensing can be used to pick the best subset of rows of some matrix \mathbf{H} that are applied to the signal \mathbf{x} . In other words, the signal \mathbf{x} is acquired using a sensing function $\text{diag}_r(\mathbf{w})\mathbf{H}$. For example, the rows of \mathbf{H} can represent different receive beamformers, filters, or (sparsifying) basis functions that are selected using a sparse \mathbf{w} .

Sparse sensing guided by a sparse vector has several interesting properties:

1. **Linearity:** The compression is linear, however the sampling is typically nonuniform or irregular (as in [Marziliano and Vetterli, 2000, Marvasti, 2001, Vaidyanathan, 2001], for instance).
2. **Deterministic and structured samplers:** Sparse sensing is close to traditional sampling. The samplers are easier to implement as compared to sampling via random projections (as in [Drineas et al., 2006, Candès and Wakin, 2008], for instance), which is oftentimes not practical.
3. **Distributed sampling:** In contrast to a nonsparse linear compression, the construction of the sensing matrix $\Phi(\mathbf{w})$ enables a fully distributed sampling scheme, which is fundamental to distributed signal processing.
4. **Controllable:** Naturally, as with any subsampling scheme, sparse sensing also results in a reduction of the signal-to-noise ratio, by the compression factor, and leads to a loss in the inference performance. However, with sparse sensing, there exists a handle to trade this loss with the compression rate.

Other sub-optimal and trivial alternatives to the proposed sparse sensing mechanism are, for example,

1. **Uniform sensing:** A common practice is to use equally-spaced sampling indices denoted by $t_m = (m - 1)\Delta$, $m = 1, 2, \dots, K$, with a sampling interval of $\Delta = T/K$.
2. **Random sensing:** Another approach would be to instead pick K indices uniformly at random, i.e., $t_m \sim \mathcal{U}[0, T]$, $m = 1, 2, \dots, K$.

These sensing schemes are suboptimal as they ignore the inference task at hand and might not always guarantee a desired inference performance or limits the compression rate.

2.3 Applications

Many real-world applications like field (temperature, pollution, precipitation, sound) inference, target localization and tracking, radar and sonar systems, video surveillance, imaging, spectral sensing, seismology, control, to list a few, are carried out using sensor networks. For such applications, sensor placement, sensor selection, and sensor scheduling are some of the key design issues.

Choosing the best subset of sensors (or spatial locations) from a large set of candidate or available sensors (or spatial locations) such that a desired inference performance is achieved is referred to as *sensor selection (placement)*. Interpreting the entries of \mathbf{x} as the observations from different sensors (or spatial locations), sensor selection/placement can be achieved through sparse sensing. Sparse sensing can also be used to select temporal samples or to schedule space-time sensor activations (e.g., antenna thinning). Sparse sensing can be used for *source placement*—a dual problem of sensor placement. Source placement can be interpreted as the problem of choosing the best subset of source (e.g., heater) locations from a large set of candidate source locations in order to generate a desired field (e.g., temperature).

Applications such as radar (or indoor localization) systems can benefit from sparse sensing since the number of antennas (or access points) as well as their sampling rates can be significantly reduced for a given target detection probability or a given bearing/speed accuracy. Also radio astronomy systems can be improved in the sense that the antenna layouts can be optimized for the task at hand and the communications overhead in the network can be reduced. Similarly, seismic data acquisition systems can be improved by optimally placing the sensors for microseismic event detection/localization or for more general underground imaging applications. In sum, the proposed sparse sensing schemes can be used to optimally gather data for a wide range of signal processing problems.

2.4 Benefits

In this section, we will list the benefits of sparse sensing and answer the following questions: why and when is sparse sensing important?

1. **Economical constraints (hardware costs):** In many of the practical applications such as environmental monitoring, radio astronomy, localization services, the sensing devices (including communications and signal processing hardware; their maintenance) are expensive. In such cases, it is of paramount importance to minimize the number of sensing devices itself yet achieving the best possible inference performance.
2. **Storage and physical space:** If the data is not acquired smartly, then there might not be sufficient memory to store the acquired data for subsequent processing. In some applications, the physical space available for sensor placement might also be very restrictive and limited. For example, in thermal management of microprocessors there is not much physical space available for temperature sensor placement.
3. **Communications bandwidth:** Often the data acquired from the distributed sensors have to be transported to a central processing unit or a server. This consumes spectral resources, energy related to transmission and reception, and creates a need for a high data rate communication link. The communications requirements can be significantly reduced through sparse sensing.
4. **Processing and inference costs:** The data acquired has to be optimally processed to solve a specific inference task. Solving the inference problem becomes more and more difficult (e.g., increased latency, more computational capacity is required) as the data volume increases. Hence, through sparse sensing the processing requirements can be seriously diminished.

2.5 Contrast with compressive sensing

Sparse sensing differs from the broad research area of compressive sensing —state of the art in the field of sensing cost reduction [Donoho, 2006b, Candès and Wakin, 2008]. Compressive sensing is an elegant protocol for sensing and compressing data simultaneously. Although compressive sensing also aims at gathering fewer samples or measurements, there are a number of major differences with sparse sensing.

1. **Sparsity and signal processing task:** In compressive sensing, the signal of interest $x(t)$, is always considered sparse in some domain and the main goal is sparse signal reconstruction. On the other hand, for sparse sensing, the underlying signal does not necessarily have to be sparse and more general signal processing tasks can be considered. This can include sparse signal reconstruction, which we will discuss more in detail in Chapter 4, but is not limited to it.
2. **Samplers and compression:** The theory developed under the classical compressive sensing framework advocates random compression, which is essential to provide recovery algorithms, reconstruction guarantees, and performance analyses. Although random compression introduces robustness, it is difficult to realize in practice, particularly for applications requiring spatial sampling such as source localization, field estimation, imaging, and cognitive radio sensing, to list a few. Sparse sensing, on the other hand, is a deterministic type of data compression, where the sparse vector \mathbf{w} inside the sensing function gives a handle on the compression factor that can be used for optimally designing the sensing process.
3. **Inference quality:** The inference quality in compressive sensing, i.e., the reconstruction quality is generally characterized by a probabilistic measure on the space of random compression matrices. This means that the sensing function has to be constantly changed to achieve a desired result. In contrast, in sparse sensing we use a fixed sensing function, which is designed based on the probabilistic nature of the noise. Thus, it is practically more meaningful.

These differences are summarized in Table 2.1.

2.6 Sparse sensing design

To design the optimal sensing operator (characterized by the vector \mathbf{w}), we need to know the model of the physical world, and the definition of the inference task that we are trying to solve from the acquired data. These models describe the uncertainty about the state of nature through a probability measure

	Compressed sensing	Sparse sensing
Sparse $x(t)$	needed	not needed
Signal processing task	sparse signal reconstruction	any statistical inference task
Samplers	random	structured and deterministic
Compression	robust, but not always practical	practical and controllable

Table 2.1: Relation to compressive sensing.

for the noise. This uncertainty is captured by an inference performance metric or a (task-specific) statistical risk denoted by a function $f : \{0, 1\}^M \rightarrow \mathbb{R}$, which quantifies the inference performance. Depending on the definition of the inference task, the statistical risk can either be the estimation error, prediction error, or detection probability, for example. Further, the risk can be ensemble (thus data independent), where the average is computed under the noise pdf or the risk can be instantaneous (thus data dependent) leading to *model-driven* or *data-driven* sparse sensing, respectively.

In model-driven sparse sensing, the sensing function is designed to guarantee an average inference performance. Such risk functions can be computed offline. That is, the actual measurements (hence expensive hardware equipments) are not needed and only the model information is used. On the other hand, data-driven sparse sensing is appropriate for compressing already available data, e.g., sketching or censoring [Rago et al., 1996, Msechu and Giannakis, 2012]. This requires actual measurements (along with the model information) to compute the sensing function, thus, it incurs a sensing cost. In data-driven sparse sensing, the sensing function has to be designed for each data realization, which is more suitable for handling model mismatch and outliers. Since the assumption in this thesis is that the data is not yet available, we will restrict ourselves to model-driven sparse sensing throughout Chapters 3–6.

We are interested in the design of the lowest-cost sensing structures that guarantee a desired inference performance. Generally, this corresponds to a sparsest \mathbf{w} for a fixed statistical risk. Mathematically, it is a constrained cardinality minimization problem:

$$\begin{aligned} & \arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \\ & \text{s.to } f(\mathbf{w}) \leq \lambda, \end{aligned} \tag{P0}$$

where the ℓ_0 -(quasi) norm refers to the number of non-zero entries in \mathbf{w} , i.e., $\|\mathbf{w}\|_0 := |\{m : w_m \neq 0\}|$ and the threshold λ specifies the inference accuracy. Clearly, λ controls the sample size (and, hence the related sensing cost). Equivalently, the optimization problem in (P0) can also be formulated as

$$\begin{aligned} & \arg \min_{\mathbf{w} \in \{0,1\}^M} f(\mathbf{w}) \\ & \text{s.to } \|\mathbf{w}\|_0 = K, \end{aligned} \tag{P1}$$

where K is the desired number of samples. The design problems (P0) and (P1) are equivalent in the sense that with some threshold say λ^* , K samples can be selected. The problem of the form (P0) might be appropriate for certain designs where the number K , is not known, in which case λ should be known. In many applications, the number K might be known beforehand. This might happen, for example, when the sensors have already been purchased and we would want to use all of them. When K is *a priori* known, then the problem of the form (P1) is the obvious choice.

An optimal solution to (P0) and (P1), respectively, requires a combinatorial search over all the 2^M and $\binom{M}{K}$ possible combinations. This quickly becomes computationally intractable for modest values of M and K . For example, with $M = 100$ candidate sensors, there are in the order of 10^{30} possible choices whose direct enumeration is clearly impossible. This is essentially due to the Boolean constraint on the design variable \mathbf{w} . In addition, the cardinality function $\|\mathbf{w}\|_0$ is nonconvex in \mathbf{w} .

Depending on the shape of the statistical risk f , with respect to the selection variables, the above nonconvex Boolean optimization problem can be solved in the following two ways as discussed next.

2.6.1 Convex risk

The discrete combinatorial problems (P0) and (P1) can be approximately solved via convex optimization techniques assuming that there exists a risk $f(\cdot)$ that is a convex function of its argument. A convex function is formally defined as follows.

Definition 2.1 (Convex function). *Given a convex set \mathcal{W} , the function $f : \mathcal{W} \rightarrow \mathbb{R}$ is said to be convex, if it satisfies*

$$f(t\mathbf{w}_1 + (1-t)\mathbf{w}_2) \leq f(t\mathbf{w}_1) + (1-t)f(\mathbf{w}_2)$$

$\forall \mathbf{w}_1, \mathbf{w}_2 \in \mathcal{W}$ and $0 \leq t \leq 1$.

In order to solve (P0) and (P1) via convex optimization, we use standard convex relaxations: the discrete Boolean constraint $\mathbf{w} \in \{0,1\}^M$ is relaxed to a continuous set (also its convex hull) $0 \leq w_m \leq 1, m = 1, 2, \dots, M$, and the ℓ_0 -(quasi) norm can be approximated with the ℓ_1 -norm, its best convex approximation. By doing so, we can simplify the combinatorial problems (P0) and (P1) to the convex optimization problems

$$\begin{aligned} & \arg \min_{\mathbf{w}} \|\mathbf{w}\|_1 \\ \text{s.to } & f(\mathbf{w}) \leq \lambda, \\ & 0 \leq w_m \leq 1, m = 1, 2, \dots, M, \end{aligned} \tag{R0}$$

and

$$\begin{aligned} & \arg \min_{\mathbf{w}} f(\mathbf{w}) \\ \text{s.to } & \|\mathbf{w}\|_1 = K, \\ & 0 \leq w_m \leq 1, m = 1, 2, \dots, M, \end{aligned} \tag{R1}$$

respectively. An approximate Boolean solution can then be recovered from the solution of the above convex optimization problem either by simple thresholding or randomized rounding. Alternatively, the ℓ_0 -(quasi) norm can be approximated using the sum-of-logs $\sum_{m=1}^M \ln(w_m + \delta)$ with $\delta > 0$, which results in an iteratively weighted ℓ_1 -norm optimization problem. Typically, log-based heuristics result in a sparser solution, and thus better approximate the ℓ_0 -(quasi) norm.

Specific instances of relaxed problem (R1) have been proposed in [Joshi and Boyd, 2009] for estimation with linear and additive white Gaussian models. In this thesis, we provide a unifying theory and related algorithms for more complicated (nonlinear) inference tasks. More specifically, we will seek statistical risk functions $f(\mathbf{w})$ that are *convex* functions on $\mathbf{w} \in [0, 1]^M$ for fundamental statistical inference problems like estimation, filtering, and detection.

2.6.2 Submodular risk

An alternative way to solve (P0) and (P1) is to look for a risk function that is submodular in nature. The notion of submodularity is based on the property of diminishing returns. That is, for example, adding an observation to a set \mathcal{X} benefits less than or the same as adding the same observation to one of the subsets of \mathcal{X} . Mathematically, submodularity can be defined as follows.

Definition 2.2 (Submodular function). *Given two sets \mathcal{X} and \mathcal{Y} such that for every $\mathcal{X} \subseteq \mathcal{Y} \subseteq \mathcal{M}$ and $s \in \mathcal{M} \setminus \mathcal{Y}$, the set function $f : 2^{\mathcal{M}} \rightarrow \mathbb{R}$ defined on the subsets of \mathcal{M} is said to be submodular, if it satisfies*

$$f(\mathcal{X} \cup \{s\}) - f(\mathcal{X}) \geq f(\mathcal{Y} \cup \{s\}) - f(\mathcal{Y}).$$

Further, if the submodular function is monotone nondecreasing, that is, $f(\mathcal{X}) \leq f(\mathcal{Y})$ for all $\mathcal{X} \subseteq \mathcal{Y} \subseteq \mathcal{M}$ and normalized (i.e., $f(\emptyset) = 0$), then a greedy maximization of such a function is near-optimal and has a deterministic approximation factor of $(1 - 1/e)$, where e is the Euler's number [Nemhauser et al., 1978]. Hence, submodular risks are very useful to solve (discrete) combinatorial optimization problems using low-complexity greedy algorithms (see e.g., [Krause, 2008]). Some examples of submodular functions that are used in sensing optimization for estimation and filtering problems are frame potential [Ranieri et al., 2014], mutual information [Krause, 2008], and entropy [Krause, 2008].

Let us define an index set \mathcal{X} that is related to the sparse vector \mathbf{w} as

$$\mathcal{X} = \{m \mid w_m = 1, m = 1, 2, \dots, M\}.$$

Thus, the set \mathcal{X} is analogous (and maps uniquely) to the sparse vector \mathbf{w} . Assuming that there exists a submodular (task-specific) risk $f(\cdot)$, we can solve

Algorithm 2.1 Submodular sensing [Krause, 2008]

1. **Require** $\mathcal{X} = \emptyset, K$.
 2. **for** $k = 1$ to K
 3. $s^* = \arg \max_{s \notin \mathcal{X}} f(\mathcal{X} \cup \{s\})$
 4. $\mathcal{X} \leftarrow \mathcal{X} \cup \{s^*\}$
 5. **end**
 6. **Return** \mathcal{X}
-

the sparse sensing problem of the form (P1) using a greedy algorithm, which iteratively adds elements such that the uncertainty is reduced the most starting from an empty set; see Algorithm 2.1. Solving problems of the form (P0), i.e., minimizing the number of samples subject to a performance constraint is a straightforward adaptation of Algorithm 2.1, where the elements are added starting from an empty set until a desired performance is achieved.

Submodular sensing is useful for discrete optimization problems, thus only applicable for discrete sparse sensing and not for continuous sparse sensing. Moreover, many inference performance metrics are not readily monotone submodular, and their submodular surrogates (when available) do not always lead to a reasonable inference performance. On the other hand, convex optimization techniques can be used to design discrete as well as continuous sparse sensing mechanisms. For these reasons, the focus will be on sparse sensing design with convex risk functions.

2.7 Discussion

In this chapter we have kept the definition of the signal processing task and hence the risk function (i.e., the inference performance metric) abstract. We shall discuss more specifically different risk functions $f(\mathbf{w})$ for estimation, filtering, and detection in Chapter 3, Chapter 4, and Chapter 5, respectively. We will provide algorithms to solve the continuous sparse sensing problem in Chapter 6. In Chapters 3 till 6, the assumption is that the model information is perfectly known and the data is not yet available. We design sparse sensing

functions to acquire data in order to reach a desired average inference performance. Therefore, the sparse sensing functions can be designed offline (i.e., actual measurements are not needed) using only the available model information. Once the sensing functions are designed, solving the inference problem is not novel by itself and is based on classical signal processing tools.

Chapter 3

Sparse Sensing for Estimation

Contents

3.1	Introduction	36
3.2	Sensing nonlinear observations	38
3.3	$f(w)$ for estimation	39
3.4	Problem statement	44
3.5	Solvers	45
3.6	Numerical example: sensor placement	55
3.7	Dependent observations	62
3.8	Discussion	66
3.A	Performance thresholds	67
3.B	Projected Newton's method	69
3.C	Power iterations	69

Part of this chapter was published as: S.P. Chepuri and G. Leus. Sparsity-Promoting Sensor Selection for nonlinear Measurement Models. *IEEE Trans. on Signal Processing*, 63(3): 684-698, Feb. 2015.

3.1 Introduction

Discrete sparse sensing mechanisms enable the design of sparse space-time samplers that guarantee a desired estimation accuracy. Such problems are encountered, for example, in sensor placement (or selection), where the best subset of sensor locations (or measurements) are to be selected from a large set of candidate sensor locations (or measurements) subject to a specific performance constraint. Sensor selection/placement is pertinent to various sensor network and data analysis applications like network monitoring, location-aware services (e.g., target localization and tracking), field (e.g., heat, sound) estimation, and environmental (e.g., climate, precipitation) monitoring, to list a few. In such applications, we essentially solve an inverse problem, where the goal is to infer the parameters that describe the underlying physical phenomenon from a set of noisy measurements. These unknown parameters are related to the measurements through a model. In this chapter we are interested in designing sparse sensing mechanisms to gather only the most informative data being aware of the data model as well as the estimation task. In particular, the focus will be on nonlinear measurement models and developing risk functions that quantify the estimation accuracy.

3.1.1 Related prior works

A large volume of literature exists on sensor selection [Joshi and Boyd, 2009, and references therein]. The sensor selection problem is often formulated as an optimization problem based on some well-known performance measures from the optimal design of experiments [Ford et al., 1989, Pukelsheim, 1993]. For parameter estimation problems, the performance measures are related to the error covariance matrix denoted by $\mathbf{E} = \mathbb{E}\{(\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}})(\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}})^T\}$, and they are optimized with respect to the selection variables. Here, $\boldsymbol{\theta}$ and $\widehat{\boldsymbol{\theta}}$ denote the unknown parameter and its estimate, respectively. Some of the popular choices for the performance measures are:

1. *A-optimality*: sum of eigenvalues of \mathbf{E} , i.e., $\text{tr}\{\mathbf{E}\}$.
2. *E-optimality*: maximum eigenvalue of \mathbf{E} , i.e., $\lambda_{\max}\{\mathbf{E}\}$.
3. *D-optimality*: determinant of \mathbf{E} , i.e., $\det\{\mathbf{E}\}$.

All the above measures are equally reasonable, although neither of them completely characterizes the error covariance. There is no general answer to the question of how does one performance metric compare with the other.

Sensor selection for additive Gaussian linear models has been solved via convex relaxation techniques in [Joshi and Boyd, 2009], where the matrix E can be expressed in closed form (thus can be optimized). However, this is not true in general (e.g., for nonlinear or non-Gaussian measurement models). The solution from [Joshi and Boyd, 2009] has also been applied to sensor placement for power grid monitoring in [Kekatos et al., 2012].

Alternative greedy approaches exploiting the submodularity of the objective function [Krause et al., 2008b, Krause and Guestrin, 2007, Krause et al., 2008a, Shamaiah et al., 2010, Yao et al., 1993, Ranieri et al., 2014] are also proposed to solve the sensor selection for estimation [cf. §2.6.2 of Chapter 2].

Sensor selection for dynamical systems often referred to as sensor polling or scheduling, is studied in [Masazade et al., 2012, Carmi, 2010, Fu et al., 2012]. All the above literature (in general) deals with measurements that are related to additive Gaussian linear models. In [Kekatos and Giannakis, 2011], reliable sensor selection based on the actual measurements to identify the outliers is presented. A different problem, yet related to sensor selection, is the problem of identifying source-informative sensors, which is studied in [Schizas, 2013].

3.1.2 Contributions

We consider general scenarios where the measurements of the unknown parameter follow a nonlinear model (unlike [Joshi and Boyd, 2009] for instance). Nonlinear measurement models are frequently encountered in applications like source localization, field estimation, or phase retrieval, to list a few. The error covariance matrix for nonlinear models is not always available in closed form, and more importantly it depends on the unknown parameter. Our first contribution in the context of sensor selection is to leverage the additive property of the inverse Cramér-Rao bound (CRB) or the Fisher information matrix (FIM) for independent observations, and thus to express the performance requirement as a convex set. The CRB is a rigorous performance measure for optimality, and it generalizes very well for nonlinear measurement models (not necessarily in additive Gaussian noise). Although the first part of this

chapter focuses on independent observations, we also extend the framework to the case of nonlinear measurements in additive *correlated* Gaussian noise. In order to design the sensing mechanism of interest, we do not need the actual measurements, and hence, our framework is also well-suited for solving offline design problems (as it is data independent, but model driven).

The proposed sensor selection framework is generic and can be applied to any nonlinear estimation problem (linear being a special case). The selection problem is formulated as the design of a sparse selection vector, which is an ℓ_0 - (quasi) norm nonconvex Boolean optimization problem. The nonconvex sensor selection problem is relaxed using standard convex relaxation techniques that can be efficiently solved in polynomial time.

A sparsity-enhancing concave surrogate for the ℓ_0 - (quasi) norm is also proposed for sensor selection as an alternative to the traditional best convex relaxation. This is particularly advantageous when there are multiple (nearly) identical sensor measurements.

To cope with large-scale problems, we further present a projected subgradient algorithm. It is worth mentioning that the projected subgradient algorithm allows a very easy distributed implementation. In essence, we seek a sparse vector (i.e., a vector with many zeros and a few nonzero entries) that determines the sensing pattern. Sparse sensing leads to energy-efficient sampling schemes. We illustrate the sensor selection problem using examples of sensor placement for source localization.

3.2 Sensing nonlinear observations

We consider a generic nonlinear measurement model

$$x_m = h_m(\boldsymbol{\theta}, n_m), \quad m = 1, 2, \dots, M, \quad (3.1)$$

where x_m is the m th spatial or temporal sensor measurement, $\boldsymbol{\theta} \in \mathbb{R}^N$ is the unknown parameter, $\{n_m\}_{m=1}^M$ describe the noise components, and the regressors $\{h_m(\cdot, \cdot)\}_{m=1}^M$ are (in general) nonlinear functionals. Let the vector $\mathbf{x} = [x_1, x_2, \dots, x_M]^T \in \mathbb{R}^M$ collect the measurements. The likelihood of the measurements $p(\mathbf{x}; \boldsymbol{\theta})$ is the probability density function (pdf) of \mathbf{x} parameterized by the unknown vector $\boldsymbol{\theta}$. Similarly, the likelihood of the measurement $p_m(x; \boldsymbol{\theta})$ is the pdf of x_m parameterized by the unknown vector $\boldsymbol{\theta}$.

We acquire the data \mathbf{x} via the discrete sparse sensing mechanism that was introduced in Chapter 2. That is, we acquire data as

$$\mathbf{y} = \text{diag}_r(\mathbf{w})\mathbf{x} = \Phi(\mathbf{w})\mathbf{x},$$

where $\Phi(\mathbf{w}) = \text{diag}_r(\mathbf{w}) \in \{0, 1\}^{K \times M}$ is the sensing matrix characterized by the selection vector

$$\mathbf{w} = [w_1, w_2, \dots, w_M]^T \in \{0, 1\}^M.$$

Here, the variable $w_m = (0)1$ indicates whether the m th sensor is (not) selected. Note that we are interested in cases where $K \ll M$ and K is not known. The reduced dimension data vector $\mathbf{y} \in \mathbb{R}^K$ is used instead of $\mathbf{x} \in \mathbb{R}^M$ to solve the estimation problem.

Our goal is now to select the best subset ($\geq N$) of the M available (or candidate) sensors, that is, to design the entries of \mathbf{w} as sparse as possible, such that a certain accuracy on the estimate $\hat{\boldsymbol{\theta}}$ is guaranteed. For nonlinear inverse problems, the risk functions $f(\mathbf{w})$ that quantify the estimation accuracy are discussed next.

3.3 $f(\mathbf{w})$ for estimation

For nonlinear estimation problems, the error covariance matrix does not admit a closed-form expression or their expressions might not be suitable for numerical optimization, in general. Therefore, we will discuss a simpler and weaker surrogate, which can be optimized instead of the error covariance matrix. More specifically, we will use the CRB as a substitute for the error covariance matrix, however, we will not restrict ourselves to any specific estimator. The motivation behind using the CRB is twofold:

1. The CRB is a measure for the (local) identifiability of the problem. More specifically, a nonsingular FIM implies (local) solvability and a unique estimate of $\boldsymbol{\theta}$, however, the converse is not necessarily true [Rothenberg, 1971]. The sensor selection problem presented in this chapter seeks a subset of sensors for which the FIM has full rank in some domain such that the solvability of the problem in that domain is always ensured.

2. Typically, the subset of selected sensors that yields a lower CRB also yields a lower MSE, and thus improves the performance of any practical system.

The CRB also has a very attractive mathematical structure resulting in a selection problem that can be efficiently solved using numerical optimization techniques. Before we formally introduce the CRB, we make the following assumption:

Assumption 3.1 (Regularity condition). *The log-likelihood of the measurements satisfies the regularity condition, that is, $\mathbb{E}\{\partial \ln p(\mathbf{y}; \boldsymbol{\theta}) / \partial \boldsymbol{\theta}\} = \mathbf{0}$.*

The regularity condition in general holds for observations that belong to the family of exponential pdfs, and it already includes a large number of distributions. Thus, the proposed sparse sensing framework is valid as long as the above assumption is true.

Under Assumption 3.1 —a well-known condition for the CRB to exist [Kay, 1993], the covariance of any unbiased estimate $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^N$ of the unknown parameter satisfies the well-known inequality (also called the Cramér-Rao lower bound) [Kay, 1993]

$$\mathbb{E}\{(\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}})(\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}})^T\} \geq \mathbf{C}(\mathbf{w}, \boldsymbol{\theta}) = \mathbf{F}^{-1}(\mathbf{w}, \boldsymbol{\theta}),$$

where the Fisher information matrix (FIM) is given by [Kay, 1993]

$$\begin{aligned} \mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) &= -\mathbb{E}\left\{\frac{\partial^2}{\partial \boldsymbol{\theta}} \left(\frac{\ln p(\mathbf{y}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}\right)^T\right\} \\ &= \mathbb{E}\left\{\left(\frac{\partial \ln p(\mathbf{y}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}\right) \left(\frac{\partial \ln p(\mathbf{y}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}\right)^T\right\} \in \mathbb{R}^{N \times N}, \end{aligned} \quad (3.2)$$

and $\mathbf{C}(\mathbf{w}, \boldsymbol{\theta})$ is the CRB matrix.

We now introduce another assumption.

Assumption 3.2 (Independent observations). *The observations $\{x_m\}_{m=1}^M$ are a sequence of independent random variables, which depend on the unknown parameter $\boldsymbol{\theta}$, i.e., we have a class of pdfs that satisfy the relation $p(\mathbf{x}; \boldsymbol{\theta}) = \prod_{m=1}^M p_m(x; \boldsymbol{\theta})$.*

Under Assumption 3.2, the selection variable w_m modifies the log-likelihood of the selected measurements as

$$\ln p(\mathbf{y}; \boldsymbol{\theta}) = \ln \prod_{m=1}^M p_m(x; \boldsymbol{\theta})^{w_m} = \sum_{m=1}^M w_m \ln p_m(x; \boldsymbol{\theta}), \quad (3.3)$$

where the pdf of the selected measurements is of reduced dimension, i.e., it does not include the measurements that are set to zero. Using (3.3) in (3.2), the FIM $F(\mathbf{w}, \boldsymbol{\theta})$, can be explicitly expressed as a *linear function* of \mathbf{w} as

$$F(\mathbf{w}, \boldsymbol{\theta}) = \sum_{m=1}^M w_m F_m(\boldsymbol{\theta}), \quad (3.4)$$

where

$$\begin{aligned} F_m(\boldsymbol{\theta}) &= -\mathbb{E} \left\{ \frac{\partial^2}{\partial \boldsymbol{\theta}} \left(\frac{\ln p_m(x; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^T \right\} \\ &= \mathbb{E} \left\{ \left(\frac{\partial \ln p_m(x; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right) \left(\frac{\partial \ln p_m(x; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^T \right\}, \end{aligned} \quad (3.5)$$

is the $N \times N$ FIM of the m th (local) measurement. In other words, (3.4) means that every independent measurement contributes to the information measure and we use the Boolean selection parameter to choose the most informative sensors (or measurements). Note that the FIM for nonlinear models depends on the unknown vector $\boldsymbol{\theta}$.

One specific example that often occurs in practice is the case where the observations $\{x_m\}_{m=1}^M$, are related through the following additive Gaussian nonlinear model given by

$$x_m = h_m(\boldsymbol{\theta}) + n_m, \quad m = 1, 2, \dots, M, \quad (3.6)$$

with $n_m \sim \mathcal{N}(0, \sigma_m^2)$. It is easy to verify that for (3.6) the FIM (3.5) simplifies to

$$F_m(\boldsymbol{\theta}) = \frac{1}{\sigma_m^2} \left(\frac{\partial h_m(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right) \left(\frac{\partial h_m(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^T.$$

Remark 3.1 (Additive Gaussian linear model).

As a special case, when the measurement process is linear, we have $x_m = \mathbf{h}_m^T \boldsymbol{\theta} + n_m$, $m = 1, 2, \dots, M$, i.e., $h_m(\boldsymbol{\theta}, n_m) := \mathbf{h}_m^T \boldsymbol{\theta} + n_m$ with $\mathbf{h}_m \in \mathbb{R}^N$ being the regressor and $n_m \sim \mathcal{N}(0, \sigma_m^2)$. The computation of the FIM for a linear model is straightforward, and is given by

$$\mathbf{F}(\mathbf{w}) = \sum_{m=1}^M (w_m / \sigma_m^2) \mathbf{h}_m \mathbf{h}_m^T.$$

The CRB for linear models in additive Gaussian noise is equal to the mean squared error (MSE), and more importantly it is independent of the unknown vector $\boldsymbol{\theta}$.

In what follows, we will develop several scalar risk functions $f(\mathbf{w})$ together with the corresponding accuracy threshold λ that can be used in (P0) to design the sparse sensing operator as discussed in Chapter 2.

We constrain the estimation error $\boldsymbol{\varepsilon} = \widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}$ to be within an origin-centered circle of radius R_e with a probability higher than P_e , i.e.,

$$\Pr(\|\boldsymbol{\varepsilon}\|_2 \leq R_e) \geq P_e, \quad (3.7)$$

where the values of R_e and P_e define the accuracy required and are assumed to be known. A higher accuracy level is obtained by reducing R_e and/or increasing P_e . This metric is used in several occasions as an accuracy measure (e.g., see [Cover and Thomas, 2012, Gustafsson and Gunnarsson, 2005, Wang et al., 2009]). We next discuss two popular performance measures from the design of experiments that satisfy the above requirement.

Trace constraint

The risk function that satisfies the accuracy requirement in (3.7) is

$$f(\mathbf{w}) := \text{tr} \{ \mathbf{C}(\mathbf{w}, \boldsymbol{\theta}) \} = \text{tr} \left\{ \left(\sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) \right)^{-1} \right\}$$

with a sufficient condition (see Appendix 3.A)

$$\text{tr} \left\{ \left(\sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) \right)^{-1} \right\} \leq \lambda_{\text{tr}} = (1 - P_e) R_e^2. \quad (3.8)$$

This measure is related to the *A-optimality* or the average-variance criterion, which restricts the sum of the semi-axes of the confidence ellipsoid to λ_{tr} .

Minimum eigenvalue constraint

Another risk function that satisfies the accuracy requirement in (3.7) is

$$f(\mathbf{w}) := \lambda_{\max}\{\mathbf{C}(\mathbf{w}, \boldsymbol{\theta})\}$$

with a sufficient condition

$$\lambda_{\min}\{\mathbf{F}(\mathbf{w}, \boldsymbol{\theta})\} \geq \lambda_{\text{eig}} = \frac{N}{R_e^2} \left(\frac{1}{1 - P_e} \right),$$

where λ_{eig} is derived in [Wang et al., 2009] (see also Appendix 3.A). This measure is related to the *E-optimality* or the worst-case error, which restricts the semi-major axis of the confidence ellipsoid to λ_{eig} . The inequality constraint $\lambda_{\min}\{\mathbf{F}(\mathbf{w}, \boldsymbol{\theta})\} \geq \lambda_{\text{eig}}$ can be equivalently expressed as the following linear matrix inequality (LMI):

$$\sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) - \lambda_{\text{eig}} \mathbf{I}_N \succeq \mathbf{0}_N. \quad (3.9)$$

In other words, we put a lower bound on each eigenvalue of the matrix \mathbf{F} .

The above performance measures depend on the unknown parameter $\boldsymbol{\theta}$. In practice, the unknown parameter $\boldsymbol{\theta}$ has a physical meaning and takes values within a certain domain denoted by \mathcal{U} . For example, in the case of direction-of-arrival estimation, \mathcal{U} is the sector where the source is expected or for target localization it is the surveillance area where the target resides. Since the FIM for nonlinear models depends on the unknown $\boldsymbol{\theta}$, we propose to constrain every point within the domain \mathcal{U} .

Remark 3.2. *The trace constraint, which can also be represented by LMIs, has a larger feasible set as compared to the minimum eigenvalue constraint. However, the resulting sensor selection problem is computationally less attractive compared to the minimum eigenvalue constraint (as we show later on in §3.5.5).*

For the aforementioned reason, we focus on the minimum eigenvalue (LMI) constraint from now on. However, either one of the two performance constraints can be used.

3.4 Problem statement

Having introduced the risk functions we can now formally state the problem.

Problem 3.1 (Discrete sparse sensing for estimation).

Given the likelihoods of the measurements, $p_m(x; \boldsymbol{\theta})$, $m = 1, 2, \dots, M$, and a desired inference performance λ_{eig} , find a sparse vector $\mathbf{w} \in \{0, 1\}^M$ that selects the minimum number of most informative sensors satisfying the performance measure $\sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) - \lambda_{\text{eig}} \mathbf{I}_N \geq \mathbf{0}_N$, $\forall \boldsymbol{\theta} \in \mathcal{U}$.

Mathematically, the discrete sparse sensing problem (P0) introduced in Chapter 2 for estimation task specializes to

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \|\mathbf{w}\|_0 \quad (3.10a)$$

$$\text{s.to} \quad \sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) - \lambda_{\text{eig}} \mathbf{I}_N \geq \mathbf{0}_N, \forall \boldsymbol{\theta} \in \mathcal{U}, \quad (3.10b)$$

$$\mathbf{w} \in \{0, 1\}^M. \quad (3.10c)$$

The threshold λ_{eig} imposes the accuracy requirement. The threshold λ_{eig} is also the sparsity-inducing parameter, where $\lambda_{\text{eig}} \rightarrow 0$ implies a sparser solution. Alternatively, the sensor selection problem can also be expressed as (P1) described in Chapter 2 when K is known.

Suppose the domain \mathcal{U} consists of D points, obtained by gridding the entire domain (where the parameter is expected) at a certain resolution. The resulting multiple LMI constraints can then be stacked together as a single LMI constraint. Let us consider the domain $\mathcal{U} = \{\boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_D\}$ with $|\mathcal{U}| = D$. The constraints in (3.10b) can then be equivalently expressed as a single LMI constraint written as $\sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN} \geq \mathbf{0}_{DN}$, where

$$\mathbf{F}_m = \text{diag}(\mathbf{F}_m(\boldsymbol{\theta}_1), \mathbf{F}_m(\boldsymbol{\theta}_2), \dots, \mathbf{F}_m(\boldsymbol{\theta}_D)) \in \mathbb{S}^{DN}$$

for $m = 1, 2, \dots, M$. Note that the FIM after gridding is independent of $\boldsymbol{\theta}$. Henceforth, we denote this simply by \mathbf{F}_m (not explicitly as a function of $\boldsymbol{\theta}$). The computational complexity of the resulting solvers depends on the number of grid points, which is due to the fact that we do not exactly know where the true parameter is located and because we are dealing with a nonlinear system

model. We make the following remarks to indicate some scenarios where the computational burden due to gridding can be reduced.

Remark 3.3 (Worst-case constraints).

If for every $\boldsymbol{\theta} \in \mathcal{U}$, there exists some $\tilde{\boldsymbol{\theta}} \in \mathcal{U}_{\text{worst}} \subset \mathcal{U}$ such that

$$\lambda_{\min} \left\{ \mathbf{F}(\mathbf{w}, \tilde{\boldsymbol{\theta}}) \right\} \leq \lambda_{\min} \left\{ \mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) \right\}, \quad \forall \mathbf{w} \in \{0, 1\}^M,$$

then it is sufficient to constrain the performance for only the worst-case set $\mathcal{U}_{\text{worst}}$ instead of \mathcal{U} . This property can be used as a guideline for gridding.

Remark 3.4 (Bayesian CRB constraint).

In a Bayesian setting, when prior information of the unknown parameter $\boldsymbol{\theta}$ is available, this additional knowledge can be incorporated in the CRB. The related information matrix is often called the Bayesian information matrix (BIM) [Van Trees, 2004], and it is independent of the unknown parameter (hence, gridding is not needed). The BIM is given by

$$\mathbf{F}_B(\mathbf{w}) = \mathbf{J}_p + \mathbb{E}_{\boldsymbol{\theta}} \{ \mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) \},$$

where \mathbf{J}_p is the prior information matrix $\mathbf{J}_p = -\mathbb{E}_{\boldsymbol{\theta}} \left\{ \frac{\partial}{\partial \boldsymbol{\theta}} \left(\frac{\ln p(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^T \right\}$ with the (log) prior $\ln p(\boldsymbol{\theta})$, and the expectation $\mathbb{E}_{\boldsymbol{\theta}} \{ \cdot \}$ is under the pdf $p(\boldsymbol{\theta})$. The LMI constraint in (3.9) for the Bayesian setting will then be

$$\mathbf{J}_p + \sum_{m=1}^M w_m \mathbb{E}_{\boldsymbol{\theta}} \{ \mathbf{F}_m(\boldsymbol{\theta}) \} \geq \lambda_{\text{eig}} \mathbf{I}_N. \quad (3.11)$$

The prior information typically comes from the dynamics, previous measurements, or combining other available measurements.

In order to optimize the Bayesian CRB, we need to know the distribution of the unknown parameter.

3.5 Solvers

As discussed in Chapter 2, the optimization problem (3.10) that is of the form (P0) is nonconvex in \mathbf{w} . We next present a number of solvers based on the relaxed convex problem (R0) presented in Chapter 2, which can be solved efficiently in polynomial time.

3.5.1 Convex approximation based on ℓ_1 -norm

A computationally tractable (suboptimal) solution is to use the traditional best convex surrogate for the ℓ_0 -(quasi) norm, namely, the ℓ_1 -norm. The ℓ_1 -norm is known to represent an efficient heuristic for the ℓ_0 -(quasi) norm optimization with convex constraints especially when the solution is sparse [Polyak et al., 2013]. Such relaxations are well-studied for problems with linear constraints in the context of compressed sensing (CS) and sparse signal recovery [Donoho, 2006a]. The nonconvex Boolean constraint in (3.10c) is further relaxed to the convex box constraint $[0, 1]^M$. The constraint (3.10b) is convex on $\mathbf{w} \in [0, 1]^M$.

The relaxed optimization problem is given as the following problem

$$\widehat{\mathbf{w}} = \arg \min_{\mathbf{w} \in \mathbb{R}^M} \|\mathbf{w}\|_1 \quad (3.12a)$$

$$\text{s.to} \quad \sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN} \succeq \mathbf{0}_{DN}, \quad (3.12b)$$

$$0 \leq w_m \leq 1, \quad m = 1, 2, \dots, M. \quad (3.12c)$$

Due to the positivity constraint, the objective function $\|\mathbf{w}\|_1$ will simply be an affine function $\mathbf{1}_M^T \mathbf{w}$. The optimization problem in (3.12) is a standard semidefinite programming (SDP) problem in the inequality form, which can be efficiently solved in polynomial time using interior-point methods [Boyd and Vandenberghe, 2004], for instance. An implementation of the interior-point method for solving SDP problems in the inequality form is typically based on Newton's method using an approximating barrier function. A brief description of the *projected Newton's method* is provided in Appendix 3.B, which we use to analyze the computational complexity of the relaxed convex problem (3.12).

Remark 3.5 (Complexity per iteration).

The computational cost involved during each iteration is as follows [Boyd and Vandenberghe, 2004, p. 619]. The matrices $\{\mathbf{F}_m\}_{m=1}^M$ have a block-diagonal structure with D blocks. Forming the matrix $\mathbf{S} = \sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN}$ costs $\mathcal{O}(DMN^2)$ flops. Computing $\mathbf{S}^{-1} \mathbf{F}_i$ via Cholesky factorization costs $\mathcal{O}(MDN^3)$ flops (for each i); the Hessian matrix is computed via the inner product of the

matrices $\mathbf{S}^{-1}\mathbf{F}_i$ and $\mathbf{S}^{-1}\mathbf{F}_j$, which costs $\mathcal{O}(DM^2N^2)$ (for each i, j pair). Finally, the Newton step is computed via Cholesky factorization costing $\mathcal{O}(DM^3)$ flops, and the projection costs $\mathcal{O}(M)$ flops. Assuming that $M \gg N$, the overall computational complexity per iteration of the projected Newton's algorithm is then $\mathcal{O}(DM^3)$.

Implementations of the interior-point methods are easily available in the form of well-known toolboxes like YALMIP [Lofberg, 2004], SeDuMi [Sturm, 1999], and CVX [Grant and Boyd, 2012].

3.5.2 Projected subgradient algorithm

The second-order Newton's method [cf. Appendix 3.B] is computationally intensive when the number of candidate sensors is very large ($M \gg 1000$, for example). To circumvent this problem, we propose a subgradient algorithm. The projected subgradient algorithm is a first-order method that is attractive for large-scale problems as each iteration is much cheaper to process.

The subgradient method is typically used for optimizations involving non-differentiable functions [Boyd et al., 2003, Bertsekas, 1999]. The subgradient method is a generalization of the gradient method for nonsmooth and non-differentiable functions, such as, the ℓ_1 -norm and the minimum eigenvalue constraint functions. We next derive the projected subgradient algorithm.

The relaxed sensor selection problem in (3.12) can be equivalently expressed as

$$\arg \min_{\mathbf{w}} \mathbf{1}_M^T \mathbf{w} \quad (3.13a)$$

$$\text{s. to } f_{\text{eig}}(\mathbf{w}) \geq \lambda_{\text{eig}}, \quad (3.13b)$$

$$\mathbf{w} \in \mathcal{W}, \quad (3.13c)$$

where $f_{\text{eig}}(\mathbf{w}) := \lambda_{\min}\{\sum_{m=1}^M w_m \mathbf{F}_m\}$ is the constraint function in (3.12b), and the set

$$\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^M \mid 0 \leq w_m \leq 1, m = 1, 2, \dots, M\}$$

denotes the box constraints in (3.12c).

The objective $\mathbf{1}_M^T \mathbf{w}$ is affine, so a subgradient of the objective is the all-one vector $\mathbf{1}_M$. Let $\mathbf{g}^k \in \partial f_{\text{eig}}(\mathbf{w}^k)$ denote a subgradient of the constraint

function $f_{\text{eig}}(\mathbf{w})$ evaluated at $\mathbf{w} = \mathbf{w}^k$. Here, the set $\partial f_{\text{eig}}(\mathbf{w}^k)$ denotes the subdifferential of $f_{\text{eig}}(\mathbf{w})$ towards \mathbf{w} evaluated at $\mathbf{w} = \mathbf{w}^k$. To compute \mathbf{g}^k , we express the constraint function $f_{\text{eig}}(\mathbf{w}^k)$ as

$$f_{\text{eig}}(\mathbf{w}^k) = \inf_{\|\mathbf{v}\|_2=1} \mathbf{v}^T \left(\sum_{m=1}^M w_m^k \mathbf{F}_m \right) \mathbf{v}.$$

The computation of a subgradient is straightforward, and is given by

$$\mathbf{g}^k = [(\mathbf{v}_{\min}^k)^T \mathbf{F}_1 \mathbf{v}_{\min}^k, \dots, (\mathbf{v}_{\min}^k)^T \mathbf{F}_M \mathbf{v}_{\min}^k]^T \in \partial f_{\text{eig}}(\mathbf{w}^k), \quad (3.14)$$

where \mathbf{v}_{\min}^k is the eigenvector corresponding to the minimum eigenvalue of $\sum_{m=1}^M w_m^k \mathbf{F}_m$. The minimum eigenvalue and the corresponding eigenvector can be computed using a low-complexity iterative algorithm, for example, the power method (see Appendix 3.C) or using the standard eigenvalue decomposition [Golub and Van Loan, 1996]. Let the projection of a point onto the set \mathcal{W} be denoted by the operator $\mathcal{P}_{\mathcal{W}}(\cdot)$, which can be expressed element-wise as

$$[\mathcal{P}_{\mathcal{W}}(\mathbf{w})]_m = \begin{cases} 0 & \text{if } w_m \leq 0, \\ w_m & \text{if } 0 < w_m < 1, \\ 1 & \text{if } w_m \geq 1. \end{cases} \quad (3.15)$$

The projected subgradient algorithm then proceeds as follows:

$$\mathbf{w}^{k+1} = \begin{cases} \mathcal{P}_{\mathcal{W}}(\mathbf{w}^k - \alpha^k \mathbf{1}_M) & \text{if } f_{\text{eig}}(\mathbf{w}^k) \geq \lambda_{\text{eig}}, \\ \mathcal{P}_{\mathcal{W}}(\mathbf{w}^k + \alpha^k \mathbf{g}^k) & \text{if } f_{\text{eig}}(\mathbf{w}^k) < \lambda_{\text{eig}}. \end{cases} \quad (3.16)$$

In other words, if the current iterate \mathbf{w}^k is feasible (i.e., $f_{\text{eig}}(\mathbf{w}^k) \geq \lambda_{\text{eig}}$), we update \mathbf{w} in the direction of a negative objective subgradient, as if the LMI constraints were absent; if the current iterate \mathbf{w}^k is infeasible (i.e., $f_{\text{eig}}(\mathbf{w}^k) < \lambda_{\text{eig}}$), we update \mathbf{w} in the direction of a subgradient \mathbf{g}^k associated with the LMI constraints. After the update is computed, the iterate is projected onto the constraint set \mathcal{W} using $\mathcal{P}_{\mathcal{W}}(\cdot)$. When the k th iterate is feasible, a diminishing nonsummable step size $\alpha^k = 1/\sqrt{k}$ is used. When the iterate is not feasible Polyak's step size $\alpha^k = (f_{\text{eig}}(\mathbf{w}^k) + \alpha_0)/\|\mathbf{g}^k\|_2^2$ is used, where we adopt the optimal value for $\alpha_0 := \mathbf{1}_M^T \mathbf{w}^*$ when $\|\mathbf{w}\|_0$ is known (i.e., the number of sensors

Algorithm 3.1 Projected subgradient algorithm

1. **Initialize** iteration counter $k = 0$, $\mathbf{w}^k = \mathbf{1}_M$, $\mathbf{g}^k = \mathbf{0}$, k_{\max} , α_0 , and λ_{eig} .
 2. **for** $k = 0$ to k_{\max}
 3. **compute** $f_{\text{eig}}(\mathbf{w}^k) = \lambda_{\min}\{\sum_{m=1}^M w_m^k \mathbf{F}_m\}$
 4. **update**
 5. **if** $f_{\text{eig}}(\mathbf{w}^k) \geq \lambda_{\text{eig}}$
 6. $\mathbf{w}^{k+1} = \mathcal{P}_{\mathcal{W}}(\mathbf{w}^k - (1/\sqrt{k})\mathbf{1}_M)$
 7. **elseif** $f_{\text{eig}}(\mathbf{w}^k) < \lambda_{\text{eig}}$
 8. compute \mathbf{g}^k according to (3.14)
 9. $\mathbf{w}^{k+1} = \mathcal{P}_{\mathcal{W}}(\mathbf{w}^k + \frac{f_{\text{eig}}(\mathbf{w}^k) + \alpha_0}{\|\mathbf{g}^k\|_2^2} \mathbf{g}^k)$
 10. **end**
 11. **end**
 12. $\widehat{\mathbf{w}} = \mathbf{w}^{k_{\max}}$
-

to be selected is known). If this is not known, then we approximate it with $\alpha_0 := f_{\text{best}}^k + \gamma$, where $\gamma = 10/(10 + k)$, and $f_{\text{best}}^k = \min\{f_{\text{best}}^{k-1}, \mathbf{1}_M^T \mathbf{w}^k\}$ [Boyd et al., 2003]. The algorithm is terminated after a specified maximum number of iterations k_{\max} . Finally, the estimate is denoted by $\widehat{\mathbf{w}} = \mathbf{w}^{k_{\max}}$.

The convergence results of the subgradient method for the constrained optimization (i.e., without the projection step) are derived in [Boyd et al., 2003]. Since the projection onto a convex set is nonexpansive [Bertsekas, 1999], it does not affect the convergence. The projected subgradient algorithm is summarized as Algorithm 3.1.

Remark 3.6 (Complexity per iteration).

We first form the matrix $\sum_{m=1}^M w_m \mathbf{F}_m$, which costs $\mathcal{O}(DMN^2)$ flops. The minimum eigenvalue and the corresponding eigenvector can be computed using the power method at a cost of $\mathcal{O}(DN^2)$ flops [Golub and Van Loan, 1996]. Forming the vector \mathbf{g} costs $\mathcal{O}(DMN^2)$ flops, computing its norm costs $\mathcal{O}(M)$ flops, and the update and projection together cost $\mathcal{O}(M)$ flops. Assuming that $M \gg N$ as earlier, the computational cost of the projected subgradient algorithm is $\mathcal{O}(DMN^2)$, which is much lower than the complexity of the projected

Newton's method.

A distributed implementation of the projected subgradient algorithm is very easy. A simple distributed averaging algorithm (e.g., [Xiao and Boyd, 2004]) can be used to compute the sum of matrices $\sum_{m=1}^M w_m \mathbf{F}_m$. The minimum eigenvalue and the corresponding eigenvector can then be computed using power iterations at each node independently. The update step (3.16), the subgradient vector \mathbf{g} , and the projection are computed coordinatewise and are already distributed.

Subgradient methods are typically very slow compared to the interior-point method involving Newton iterations, and subgradient methods typically require a few hundred iterations. Newton's method typically requires in the order of ten steps. On the other hand, unlike the projected subgradient method, Newton's method cannot be easily distributed, and incurs a relatively high complexity per iteration due to the computation and storage of up to second-order derivatives. Depending on the scale of the problem and the resources available for processing one could choose between the subgradient or Newton's algorithm.

3.5.3 Concave surrogate based on sum-of-logarithms

The ℓ_1 -norm is customarily used as the best convex relaxation for the ℓ_0 - (quasi) norm. However, the intersection of the ℓ_1 -norm ball (or an affine subspace) with the positive semi-definite cone (i.e., the LMI constraint) is not always a unique point as shown in the following Theorem.

Theorem 3.1 (Uniqueness). *The projection of a point $\mathbf{w} \in [0, 1]^M$ onto a convex LMI constraint set $\sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN} \geq \mathbf{0}_{DN}$ under the ℓ_1 -norm is not always unique.*

Proof. The proof follows from the fact that the ℓ_1 -norm is not strictly convex, and from the linearity of the constraint set. Let us consider an example with $M = 2$ (w.l.o.g.), and $\mathbf{F}_1 = \mathbf{F}_2 \geq \lambda_{\text{eig}} \mathbf{I}_{DN}$. In other words, the observations are identical. In this case, the extreme points of the ℓ_1 -norm ball, i.e., $\widehat{\mathbf{w}}_1 = [1, 0]^T$ and $\widehat{\mathbf{w}}_2 = [0, 1]^T$ are two example solutions. Moreover, since the solution set of a convex minimization problem is convex, $\tau \widehat{\mathbf{w}}_1 + (1 - \tau) \widehat{\mathbf{w}}_2$ is also a solution for any $0 < \tau < 1$, which gives an infinite number of solutions to the relaxed

optimization problem (3.12). For such cases, the ℓ_1 -norm relaxation might not result in a sparse solution. \square

To improve upon the ℓ_1 -norm solution due to its nonuniqueness following from Theorem 3.1, we propose an alternative approximation to the ℓ_0 - (quasi) norm which also results in fewer selected sensors. Instead of relaxing the ℓ_0 - (quasi) norm with the ℓ_1 -norm, using a nonconvex surrogate function can yield a better approximation. It is motivated in [Candès et al., 2008] that the logarithm of the geometric mean of its elements can be used as an alternative surrogate function for linear inverse problems in CS. Adapting this to our sensor selection problem, we arrive at the optimization problem

$$\arg \min_{\mathbf{w} \in \mathbb{R}^M} \sum_{m=1}^M \ln(w_m + \delta) \quad (3.17a)$$

$$\text{s.to} \quad \sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN} \geq \mathbf{0}_{DN}, \quad (3.17b)$$

$$0 \leq w_m \leq 1, \quad m = 1, 2, \dots, M, \quad (3.17c)$$

where $\delta > 0$ is a small constant that prevents the cost from tending to $-\infty$. The cost (3.17a) is concave, but since it is smooth w.r.t. \mathbf{w} , iterative linearization can be performed to obtain a local minimum [Candès et al., 2008]. The first-order approximation of $\ln(w_m + \delta)$ around $w_m[i-1] + \delta$ results in

$$\ln(w_m + \delta) \leq \ln(w_m[i-1] + \delta) + \frac{w_m - w_m[i-1]}{w_m[i-1] + \delta}.$$

Here, i denotes the iteration index. Instead of minimizing the original cost, the majorizing cost (second term on the right-hand side of the above inequality) can be optimized to attain a local minimum. More specifically, the optimization problem (3.17) can be iteratively driven to a local minimum using the iterations

$$\widehat{\mathbf{w}}[i] = \arg \min_{\mathbf{w} \in \mathbb{R}^M} \sum_{m=1}^M \frac{w_m}{\widehat{w}_m[i-1] + \delta} \quad (3.19a)$$

$$\text{s.to} \quad \sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN} \geq \mathbf{0}_{DN}, \quad (3.19b)$$

$$0 \leq w_m \leq 1, \quad m = 1, 2, \dots, M. \quad (3.19c)$$

Algorithm 3.2 Sparsity-enhancing iterative algorithm

1. **Initialize** the iteration counter $i = 0$, the weight vector $\mathbf{u}[0] = [u_1[0], u_2[0], \dots, u_M[0]]^T = \mathbf{1}_M$, δ , and i_{\max} .
2. **for** $i = 0$ to i_{\max}
3. **solve**

$$\widehat{\mathbf{w}}[i] = \arg \min_{\mathbf{w} \in \mathbb{R}^M} \mathbf{u}[i]^T \mathbf{w} \quad (3.18a)$$

$$\text{s.to } \sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN} \geq \mathbf{0}_{DN}, \quad (3.18b)$$

$$0 \leq w_m \leq 1, m = 1, 2, \dots, M. \quad (3.18c)$$

4. **update** $u_m[i+1] = (\delta + \widehat{w}_m[i])^{-1}$, for each $m = 1, 2, \dots, M$.
 5. **end**
 6. $\widehat{\mathbf{w}} = \widehat{\mathbf{w}}[i_{\max}]$.
-

The iterative algorithm is summarized as Algorithm 3.2. Each iteration in (3.19) solves a weighted ℓ_1 -norm optimization problem. The weight updates force the small entries of the vector $\widehat{\mathbf{w}}[i]$ to zero and avoid inappropriate suppression of larger entries. The parameter δ provides stability, and guarantees that the zero-valued entries of $\widehat{\mathbf{w}}[i]$ do not strictly prohibit a nonzero estimate at the next step. Finally, the estimate is given by $\widehat{\mathbf{w}} = \widehat{\mathbf{w}}[i_{\max}]$, where i_{\max} is the specified maximum number of iterations.

Remark 3.7 (Sparsity-enhancing projected subgradient algorithm).

The projected subgradient algorithm can be adapted to fit into the sparsity-enhancing iterative algorithm as well. The optimization problem (3.18) is then replaced with the following update equations:

$$\mathbf{w}^{k+1}[i] = \begin{cases} \mathcal{P}_{\mathcal{W}}(\mathbf{w}^k[i] - \alpha^k \mathbf{u}[i]) & \text{if } f_{\text{eig}}(\mathbf{w}^k[i]) \geq \lambda_{\text{eig}}, \\ \mathcal{P}_{\mathcal{W}}(\mathbf{w}^k[i] + \alpha^k \mathbf{g}^k[i]) & \text{if } f_{\text{eig}}(\mathbf{w}^k[i]) < \lambda_{\text{eig}}, \end{cases}$$

where we solve a number of iterations (inner loop) of the projected subgradient

Algorithm 3.3 Randomized rounding algorithm

1. Generate L candidate estimates of the form $w_{m,l} = 1$ ($l = 1, 2, \dots, L$) with a probability \widehat{w}_m (or $w_{m,l} = 0$ with a probability $1 - \widehat{w}_m$) for $m = 1, 2, \dots, M$.
2. Define $\mathbf{w}_l = [w_{1,l}, \dots, w_{M,l}]^T$ and the index set of the candidate estimates satisfying the constraints as

$$\Omega \triangleq \{l \mid \lambda_{\min}\{\mathbf{F}(\mathbf{w}_l, \boldsymbol{\theta})\} \geq \lambda_{\text{eig}}, \forall \boldsymbol{\theta} \in \mathcal{U}, l = 1, 2, \dots, L\}.$$

3. If the set Ω is empty, go back to step 1.
4. The suboptimal Boolean estimate is the solution to the optimization problem

$$\widehat{\mathbf{w}}_{\text{bp}} = \arg \min_{l \in \Omega} \|\mathbf{w}_l\|_1.$$

algorithm within the i th iteration (outer loop) of Algorithm 3.2. Here, the k th iterate of the inner loop in the i th outer loop is denoted as $(\cdot)^k[i]$.

From the solution of the relaxed optimization problem, the approximate Boolean solution can be obtained using randomization techniques, as described next.

3.5.4 Randomized rounding

The solution of the relaxed optimization problem is used to compute the suboptimal Boolean solution for the selection problem. A straightforward technique that is often used is the simple rounding technique, in which the Boolean estimate is given by $\text{round}(\widehat{w}_m)$, $m = 1, 2, \dots, M$, where we define $\widehat{\mathbf{w}} = [\widehat{w}_1, \widehat{w}_2, \dots, \widehat{w}_M]^T$, and the $\text{round}(\cdot)$ operator rounds its arguments towards the nearest integer. However, there is no guarantee that the Boolean estimates obtained from the rounding technique always satisfy the performance constraints. Hence, we propose a randomized rounding technique, where the suboptimal Boolean estimates are computed based on random experiments

guided by the solution from the SDP problem (3.12) or the iterative version in (3.19). The randomized rounding technique is summarized as Algorithm 3.3.

3.5.5 Trace and determinant constraints

In this section, we will discuss the relaxed sensor selection problem based on the optimization criteria related to A-optimality and D-optimality.

Trace constraint

The relaxed sensor selection problem with the scalar trace constraint is given as follows

$$\begin{aligned} & \arg \min_{\mathbf{w} \in \mathbb{R}^M} \|\mathbf{w}\|_1 \\ & \text{s.to} \quad \text{tr} \left\{ \left(\sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) \right)^{-1} \right\} \leq \lambda_{\text{tr}}, \quad \forall \boldsymbol{\theta} \in \mathcal{U}, \\ & \quad 0 \leq w_m \leq 1, \quad m = 1, 2, \dots, M. \end{aligned} \quad (3.20)$$

The trace constraint in (3.20) is convex on $\mathbf{w} \in [0, 1]^M$; this is easier to verify when the above trace constraint is expressed as an LMI [Boyd and Vandenberghe, 2004, p. 387]. The optimization problem (3.20) is a convex problem, and can be cast as an SDP:

$$\begin{aligned} & \arg \min_{\mathbf{w} \in \mathbb{R}^M, \mathbf{x} \in \mathbb{R}^N} \|\mathbf{w}\|_1 \\ & \text{s.to} \quad \begin{bmatrix} \sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) & \mathbf{e}_n \\ \mathbf{e}_n^T & x_n \end{bmatrix} \geq \mathbf{0}_{N+1}, \quad n = 1, 2, \dots, N, \quad \forall \boldsymbol{\theta} \in \mathcal{U}, \\ & \quad \mathbf{1}_N^T \mathbf{x} \leq \lambda_{\text{tr}}, \\ & \quad x_n \geq 0, \quad n = 1, 2, \dots, N, \\ & \quad 0 \leq w_m \leq 1, \quad m = 1, 2, \dots, M, \end{aligned} \quad (3.21)$$

where $\mathbf{x} = [x_1, x_2, \dots, x_N]^T \in \mathbb{R}^N$ is the auxiliary variable and \mathbf{e}_n is the n th unit vector in \mathbb{R}^N .

In addition to the box constraints, the optimization problem (3.21) has N LMI constraints (of size $N + 1$) for every point in \mathcal{U} and $N + 1$ inequality constraints, while (3.12) has only one LMI constraint (of size N) for every point in \mathcal{U} . Hence, solving (3.21) is computationally more intense than solving (3.12).

Determinant constraint

Another popular scalar performance measure that quantifies the estimation accuracy is the determinant (product of eigenvalues) constraint. This measure is related to the D-optimality. The relaxed sensor selection problem with the determinant constraint is given as follows

$$\begin{aligned} \arg \min_{\mathbf{w} \in \mathbb{R}^M} \quad & \|\mathbf{w}\|_1 \\ \text{s.to} \quad & \ln \det \left\{ \sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) \right\} \geq \lambda_{\text{det}}, \forall \boldsymbol{\theta} \in \mathcal{U}, \\ & 0 \leq w_m \leq 1, m = 1, 2, \dots, M, \end{aligned} \quad (3.22)$$

where the threshold λ_{det} specifies the desired mean radius of the confidence ellipsoid (see Appendix 3.A). The log-determinant constraint is a concave function of $\mathbf{w} \in [0, 1]^M$. Note that although the constraint

$$\ln \det \left\{ \sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta}) \right\} \geq \lambda_{\text{det}}$$

is an indication of the performance of the estimator, it is not a sufficient condition for (3.7).

The relaxed sensor selection problem with the scalar (trace or determinant) constraints can be solved with either one of the two proposed approximations of the cardinality cost, i.e., the ℓ_1 -norm or log-based concave surrogate.

3.6 Numerical example: sensor placement

Localization is an important and extensively studied topic in wireless sensor networks (WSNs). Target localization can be performed using a plethora

of algorithms [Gustafsson and Gunnarsson, 2005, Gezici et al., 2005, Patwari et al., 2005] (and references therein), which exploit inter-sensor measurements like time of arrival (TOA), time difference of arrival (TDOA), angle of arrival (AOA), or received signal strength (RSS). The performance of any location estimator depends not only on the algorithm but also on the placement of the anchors (sensors with known locations). Sensor placement is a key challenge in localization system design, as certain sensor constellations not only deteriorate the performance, but also result in ambiguity or identifiability issues [Chepuri et al., 2013b].

The sensor placement problem can be interpreted as the problem where we divide a specific sensor area \mathcal{S} into M grid points and select the best subset of locations from these grid points. Here, the selected sensors are deemed the best, if they guarantee a certain prescribed accuracy on the location estimates for a target within a specific target area \mathcal{U} . We consider a two-dimensional network with one target located in the target area \mathcal{U} and M possible sensors located at the M grid points.

The absolute positions of the sensor grid points are known, hence, the considered sensors are commonly referred to as anchor nodes. Let the coordinates of the target and the m th anchor be denoted by the 2×1 vectors $\boldsymbol{\theta} = [\theta_1, \theta_2]^T$ and $\mathbf{a}_m = [a_{m,1}, a_{m,2}]^T$, respectively. Here, $\boldsymbol{\theta}$ is assumed to be unknown, but known to be within \mathcal{U} . We next illustrate the developed theory with an example.

Let the pairwise distance between the target and the m th anchor be denoted by $d_m = \|\boldsymbol{\theta} - \mathbf{a}_m\|_2$. In practice, the pairwise distances are obtained using time of arrival measurements of the ranging signals and they are noisy. The range measurements typically follow an additive Gaussian nonlinear model (see [Wang et al., 2009, § II-B], for example), as given by

$$y_m = d_m + n_m, \quad m = 1, 2, \dots, M, \quad (3.23)$$

where $n_m \sim \mathcal{N}(0, \sigma_m^2)$ is the noise with $\sigma_m^2 = \sigma^2/d_m^{-\eta}$. Here, σ^2 is the nominal noise variance, and η is the path-loss exponent. We can now write the FIM for the localization problem as $\mathbf{C}^{-1} = \mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) = \sum_{m=1}^M w_m \mathbf{F}_m(\boldsymbol{\theta})$, where using (3.5) we can compute

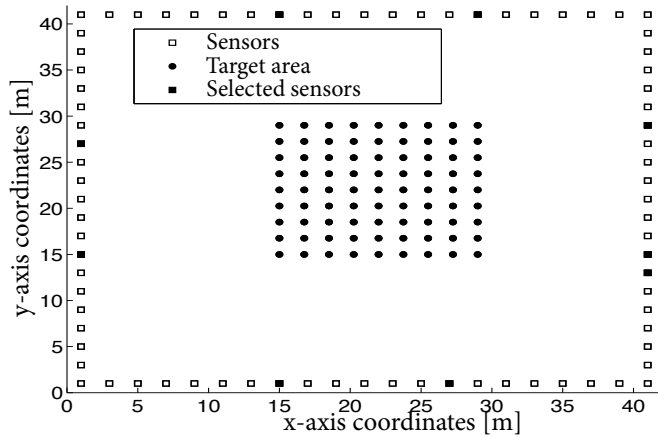
$$\mathbf{F}_m(\boldsymbol{\theta}) := \frac{(\boldsymbol{\theta} - \mathbf{a}_m)(\boldsymbol{\theta} - \mathbf{a}_m)^T}{\sigma_m^2 \|\boldsymbol{\theta} - \mathbf{a}_m\|_2^2}.$$

We apply the proposed sensor selection problem to sensor placement design for target localization. To test the proposed algorithms, we use CVX [Grant and Boyd, 2012]. CVX internally calls SeDuMi [Sturm, 1999], a MATLAB implementation of the second-order interior-point methods.

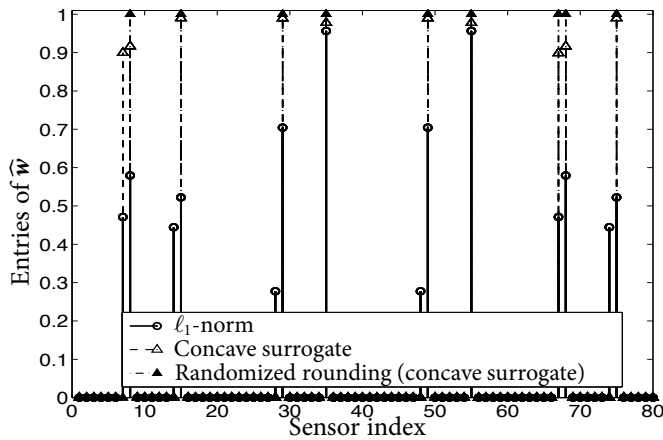
We consider the scenario shown in Figure 3.1(a) with $M = 80$ sensors to illustrate the sensor placement problem. Recall that the problem here is to choose the best sensor positions out of $M = 80$ available ones, such that a certain specified localization accuracy is achieved. We grid the target (or surveillance) area of $15 \times 15 \text{ m}^2$ uniformly with a resolution of 1.75 m along both horizontal and vertical directions as shown in Figure 3.1(a) to obtain the set \mathcal{U} containing $D = 81$ points.

The original nonconvex optimization problem is relaxed to an ℓ_1 -norm optimization problem. Alternatively, a concave surrogate function can be used to further enhance the sparsity. The optimization problem with the concave surrogate cost function is iteratively solved by affinely scaling the objective based on the solution from the previous iteration. For the sparsity-enhancing iterative Algorithm 3.2, we use $i_{\max} = 10$ and $\delta = 10^{-8}$. The number of randomizations used in the randomized rounding Algorithm 3.3 is $L = 100$. As observed in the simulations, a solution is typically found in the first batch itself, and a few tens of candidate entries are sufficient. We use the following parameters for the simulations: $\eta = 2$, $\sigma^2 = 1.78 \times 10^{-5}$, and $P_e = 0.9$.

Figure 3.1 shows the sensor selection for the distance (range) measurement model. The thresholds are computed with $R_e = 20 \text{ cm}$ and $P_e = 0.9$. The selection shown in Figure 3.1(a) is based on Algorithm 3.2 with randomized rounding to recover the approximate Boolean solution. The selection results based on the ℓ_1 -norm cost with the minimum eigenvalue constraint is shown in Figure 3.1(b). Figure 3.1(b) also shows that the solution based on the concave surrogate cost function with the minimum eigenvalue constraint leads to a sparser solution. The selection results based on the trace constraint obtained by solving (3.20) are illustrated in Figure 3.2. The sensors from the same region (close to the filled boxes in Figure 3.1(a)) are selected with either one of the two constraints. The sufficient trace constraint has a larger feasible set compared to the stronger sufficient minimum eigenvalue constraint. As a result, for the considered scenario, the minimum eigenvalue constraint leads



(a)



(b)

Figure 3.1: Sensor placement for target localization with $M = 80$ candidate sensor positions. The thresholds are computed using $R_e = 20$ cm and $P_e = 0.9$. (a) Selection based on sparsity-enhancing iterations with minimum eigenvalue constraints. The Boolean solution is recovered using randomized rounding. (b) Minimum eigenvalue constraints with ℓ_1 -norm and concave surrogate based relaxations. Randomized rounding is applied on the concave surrogate based solution.

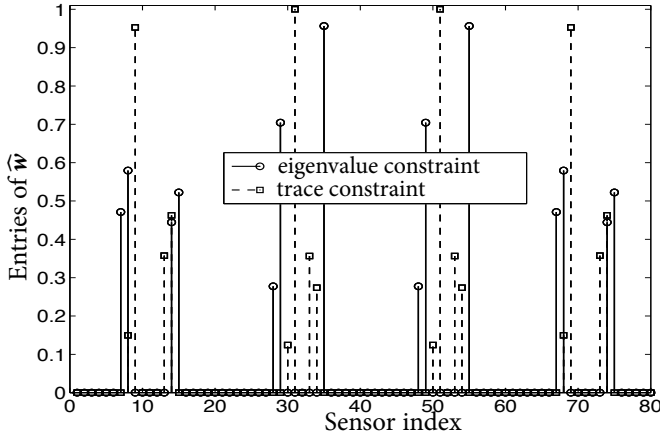


Figure 3.2: ℓ_1 -norm based selection with the trace and eigenvalue constraints. The thresholds are computed using $R_e = 20$ cm and $P_e = 0.9$.

to a slightly larger ℓ_1 -norm compared to the trace constraint. Nevertheless, for the considered example, the obtained solutions (i.e., samplers) are sparse in nature. Note that solutions from [Joshi and Boyd, 2009, Masazade et al., 2012] cannot be directly applied to such nonlinear models without an explicit linearization.

The optimization problem (3.12) is also solved using the projected subgradient method summarized in Algorithm 3.1 with $k_{\max} = 1000$ iterations. The solution of the projected subgradient is shown in Figure 3.3(a).

The convergence of the projected subgradient algorithm with respect to the solution from SeDuMi denoted by f_{opt} is shown in Figure 3.3(b). Even though the convergence of the projected subgradient algorithm is very slow, the estimated support after a few hundred iterations can be used along with randomized rounding to further refine the solution. The computation time on the same computer for the projected subgradient algorithm that solves (3.12) is around 8.84 seconds for 1000 iterations while SeDuMi takes around 4.03 seconds to solve the SDP problem in (3.12).

A practical estimator does not meet the CRB in some cases (for instance at low signal-to-noise ratios or finite data records). Therefore, the sensors selected with a specific R_e would lead to an underestimate of the desired MSE.

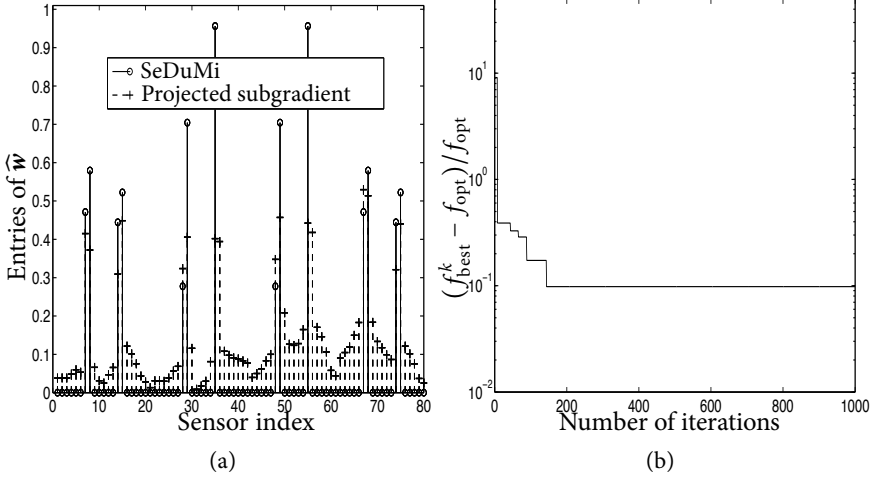


Figure 3.3: (a) The projected subgradient algorithm used to solve (3.12). (b) Performance of the projected subgradient algorithm. The thresholds are computed using $R_e = 20$ cm and $P_e = 0.9$.

We can account for this gap by choosing R_e appropriately. To this end, we give the entire solution path of the selected sensors for different values of R_e in Figure 3.4(a). The solution path can be efficiently computed by increasing R_e . The sensors corresponding to some R_e can then be used to meet the desired MSE requirement. The nonlinear model in (3.23) is solved in the least-squares sense iteratively using Gauss-Newton's method with 10 iterations [Kay, 1993]. The maximum root-MSE (RMSE), maximum root-CRB, average RMSE, and average root-CRB of the location estimates of a target within the target area using the selected sensors (as shown in the solution path) for different values of R_e are shown in Figure 3.4(b). For the considered scenario, both the maximum and average root-CRB satisfy the performance constraint which is given by the inequality in (3.31). The performance constraint is shown as a solid line in Figure 3.4(b). The maximum RMSE does not satisfy the accuracy requirement specified by a certain R_e , and this can be corrected by using an appropriate (lower) R_e . Moreover, for the considered scenario, the gap between the average RMSE and the performance constraint is still reasonable. We also show the maximum RMSE on top of Figure 3.4(a).

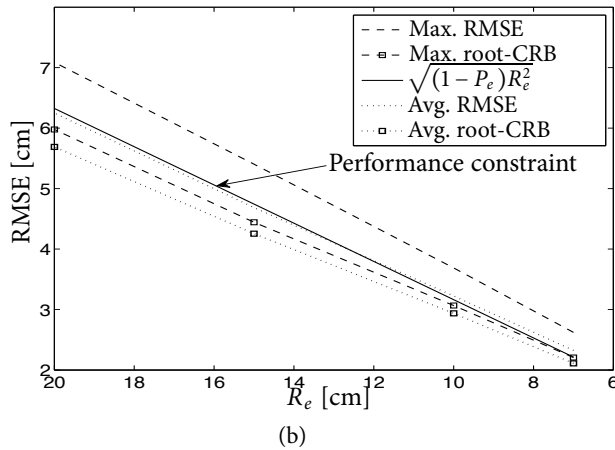
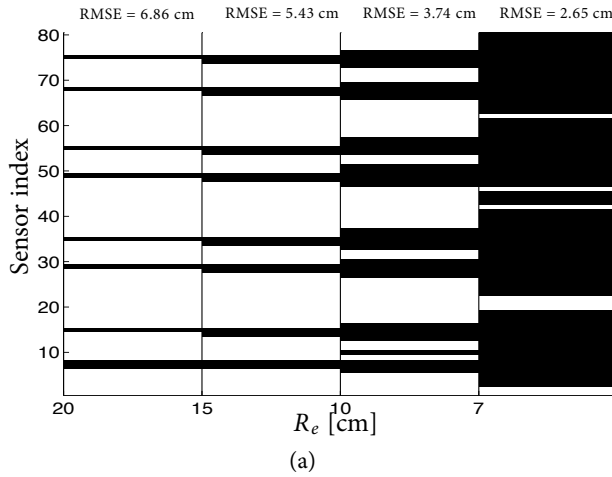


Figure 3.4: (a) Solution path of the sensors selected for different values of R_e and $P_e = 0.9$. Maximum RMSE based on selected sensors can be seen on the top of this plot. (b) Maximum and average RMSE of the location estimates based on Gauss-Newton's method, the corresponding maximum and average root-CRB, and the performance constraint in (3.31) for different values of R_e , and $P_e = 0.9$.

The proposed framework can be applied to a variety of signal processing problems (e.g., direction-of-arrival, frequency and delay estimation) as long as Assumption 3.1 and Assumption 3.2 are valid and all the model parameters are perfectly known. This is illustrated in [Chepuri and Leus, 2015b] for different target localization models based on bearing measurements, received signal strength, and energy measurements.

Although the illustrated example is related to passive sensing, we underline that the proposed sensor selection framework is not limited to passive sensors.

Remark 3.8 (Active sensor selection).

The sensor selection problem can also be formulated for active sensing. In active sensing, the sensors transmit probing signals (e.g., radar, sonar, remote sensing). The selection parameter w_m for active sensing is a soft parameter used for joint selection and resource allocation [Chepuri et al., 2013b], i.e., $w_m \in [0, 1]$ is a resource (e.g., transmit energy) normalized to the maximum prescribed value, and hence, it is dimensionless. The relaxed active sensor selection problem takes the same form as in (3.12). In fact, minimizing the ℓ_1 -norm in active sensor selection minimizes the overall network resources (e.g., overall transmit energy).

3.7 Dependent observations

Throughout this chapter so far, we have assumed that the observations are a sequence of independent random variables. This assumption is reasonable if the sensors are solely responsible for the noisy observations, for example, due to the internal thermal noise. If the observation signal itself is stochastic in nature or if the observations are subject to external noises or interference, then Assumption 3.2 will be too idealistic. As a consequence of relaxing Assumption 3.2, the FIM will not be linear (or convex) in \mathbf{w} anymore. Nevertheless, the FIM will still be a function of \mathbf{w} and sparse samplers can be designed using nonlinear and often nonconvex optimization techniques. In some cases, the solution can be computed using convex optimization techniques as illustrated in the following example.

Suppose the unknown vector $\boldsymbol{\theta} \in \mathbb{R}^N$ is related to the observations ac-

ording to

$$\mathbf{x} \sim \mathcal{N}(\mathbf{h}(\boldsymbol{\theta}), \boldsymbol{\Sigma}),$$

where $\mathbf{h}(\cdot) : \mathbb{R}^N \rightarrow \mathbb{R}^M$ is a nonlinear function and $\boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$ is the noise covariance matrix. As before, we acquire the data using the sparse sensing function $\Phi(\mathbf{w})$ to obtain $\mathbf{y} = \Phi(\mathbf{w})\mathbf{x}$.

Using (3.2), we can compute the FIM, and it is given by

$$\mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) = [\Phi(\mathbf{w})\mathbf{J}(\boldsymbol{\theta})]^T \boldsymbol{\Sigma}^{-1}(\mathbf{w}) [\Phi(\mathbf{w})\mathbf{J}(\boldsymbol{\theta})], \quad (3.24)$$

where $\mathbf{J}(\boldsymbol{\theta}) = \partial \mathbf{h}(\boldsymbol{\theta}) / \partial \boldsymbol{\theta}^T \in \mathbb{R}^{M \times N}$ is the Jacobian matrix and

$$\boldsymbol{\Sigma}(\mathbf{w}) = \Phi(\mathbf{w})\boldsymbol{\Sigma}\Phi^T(\mathbf{w}) \in \mathbb{R}^{K \times K}$$

is the submatrix of $\boldsymbol{\Sigma}$, which includes only the entries corresponding to the selected measurements. Clearly, the FIM (3.24) is no more additive or linear in \mathbf{w} . Consequently, the constraint $\lambda_{\min}\{\mathbf{F}(\mathbf{w}, \boldsymbol{\theta})\} \geq \lambda_{\text{eig}}$ in its current form is not convex in \mathbf{w} . This is also true for the trace and determinant constraints.

In what follows, we will provide some steps to express the minimum eigenvalue constraint as a convex constraint on \mathbf{w} . Firstly, we write the noise covariance matrix $\boldsymbol{\Sigma}$ as

$$\boldsymbol{\Sigma} = a\mathbf{I} + \mathbf{S}, \quad (3.25)$$

where a nonzero $a \in \mathbb{R}$ is chosen such that $\mathbf{S} \in \mathbb{R}^{M \times M}$ is invertible and well conditioned. Using (3.25) in (3.24), we obtain

$$\mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) = \mathbf{J}^T(\boldsymbol{\theta})\Phi^T(\mathbf{w})(a\mathbf{I} + \Phi(\mathbf{w})\mathbf{S}\Phi^T(\mathbf{w}))^{-1}\Phi(\mathbf{w})\mathbf{J}(\boldsymbol{\theta}).$$

We now state the following property.

Property 3.1. *Using the fact that $\Phi^T(\mathbf{w})\Phi(\mathbf{w}) = \text{diag}(\mathbf{w})$, we have*

$$\begin{aligned} \Phi^T(\mathbf{w})(a\mathbf{I} + \Phi(\mathbf{w})\mathbf{S}\Phi^T(\mathbf{w}))^{-1}\Phi(\mathbf{w}) &= \mathbf{S}^{-1} \\ &- \mathbf{S}^{-1}[\mathbf{S}^{-1} + a^{-1}\text{diag}(\mathbf{w})]^{-1}\mathbf{S}^{-1}. \end{aligned} \quad (3.26)$$

Proof. Applying the matrix inversion lemma [Kay, 1993]

$$\mathbf{C}(\mathbf{B}^{-1} + \mathbf{C}^T\mathbf{A}^{-1}\mathbf{C})^{-1}\mathbf{C}^T = \mathbf{A} - \mathbf{A}(\mathbf{A} + \mathbf{C}\mathbf{B}\mathbf{C}^T)^{-1}\mathbf{A},$$

with $\mathbf{C} = \Phi^T(\mathbf{w})$, $\mathbf{B}^{-1} = a\mathbf{I}$, and $\mathbf{A} = \mathbf{S}^{-1}$, it is easy to verify (3.26). \square

Therefore, from Property 3.1, we can equivalently express $F(\mathbf{w}, \boldsymbol{\theta})$ as

$$\begin{aligned} F(\mathbf{w}, \boldsymbol{\theta}) &= \mathbf{J}^T(\boldsymbol{\theta})\mathbf{S}^{-1}\mathbf{J}(\boldsymbol{\theta}) \\ &\quad - \mathbf{J}^T(\boldsymbol{\theta})\mathbf{S}^{-1}[\mathbf{S}^{-1} + a^{-1}\text{diag}(\mathbf{w})]^{-1}\mathbf{S}^{-1}\mathbf{J}^T(\boldsymbol{\theta}), \end{aligned} \quad (3.27)$$

where by definition $\boldsymbol{\Phi}^T(\mathbf{w})\boldsymbol{\Phi}(\mathbf{w}) = \text{diag}(\mathbf{w})$. In contrast to (3.24), the design parameter \mathbf{w} appears only once in (3.27), which makes the problem much easier. Using the Schur complement, the constraint $\lambda_{\min}\{F(\mathbf{w}, \boldsymbol{\theta})\} \geq \lambda_{\text{eig}}$ can now be equivalently expressed as an LMI:

$$\begin{bmatrix} \mathbf{S}^{-1} + a^{-1}\text{diag}(\mathbf{w}) & \mathbf{S}^{-1}\mathbf{J}(\boldsymbol{\theta}) \\ \mathbf{J}^T(\boldsymbol{\theta})\mathbf{S}^{-1} & \mathbf{J}^T(\boldsymbol{\theta})\mathbf{S}^{-1}\mathbf{J}(\boldsymbol{\theta}) - \lambda_{\text{eig}}\mathbf{I}_N \end{bmatrix} \succeq \mathbf{0}_{M+N}, \quad (3.28)$$

which is linear (thus convex) in \mathbf{w} . The above LMI is of size $M + N$ and it is larger than the size- N LMI (3.9), which is related to the independent observation case. Note that the constraint (3.28) also depends on the unknown parameter vector $\boldsymbol{\theta}$. We remark here that for linear measurement models, the above constraint is independent of the unknown parameter vector $\boldsymbol{\theta}$. In other words, in that case, $\mathbf{J}(\boldsymbol{\theta})$ will be independent of $\boldsymbol{\theta}$ and will simply be the regression matrix itself. Furthermore, because of (3.28) the matrix $\mathbf{S}^{-1} + a^{-1}\text{diag}(\mathbf{w})$ should be positive definite. This can be achieved, for example, by choosing a such that it satisfies the condition $0 < a < \lambda_{\min}\{\boldsymbol{\Sigma}\}$, since $w_m \geq 0$ for $m = 1, 2, \dots, M$.

3.7.1 Convex relaxation

The discrete sparse sensing design problem for nonlinear dependent Gaussian observations is obtained by replacing the LMI constraint (3.10b) with (3.28). The resulting nonconvex optimization problem can then be relaxed to a convex optimization problem along similar lines as explained in §3.5.1 or §3.5.3. For example, the sparsity-enhancing iterations, i.e., the re-weighted ℓ_1 -norm optimization problem [cf. (3.19) and Algorithm 3.2] for the dependent case is

given by

$$\begin{aligned}
 \widehat{\mathbf{w}}[i] &= \arg \min_{\mathbf{w} \in \mathbb{R}^M} \sum_{m=1}^M \frac{w_m}{\widehat{w}_m[i-1] + \delta} \\
 \text{s.to} \quad &\left[\begin{array}{cc} \mathbf{S}^{-1} + a^{-1} \text{diag}(\mathbf{w}) & \mathbf{S}^{-1} \mathbf{J}(\boldsymbol{\theta}) \\ \mathbf{J}^T(\boldsymbol{\theta}) \mathbf{S}^{-1} & \mathbf{J}^T(\boldsymbol{\theta}) \mathbf{S}^{-1} \mathbf{J}(\boldsymbol{\theta}) - \lambda_{\text{eig}} \mathbf{I}_N \end{array} \right] \geq \mathbf{0}_{M+N}, \\
 &\quad \forall \boldsymbol{\theta} \in \mathcal{U}, \\
 &0 \leq w_m \leq 1, \quad m = 1, \dots, M.
 \end{aligned} \tag{3.29}$$

In what follows, we will illustrate the sparse sensing design for dependent Gaussian observations applied to sensor placement for source localization.

3.7.2 Numerical example

We consider a similar setup as in §3.6, but with another popular measurement model as detailed next. In applications related to field estimation, (active/passive) radar, and sonar, it is important to estimate the location of a point source that emits or reflects energy. Suppose that there are M candidate sensors that can be placed at locations $\{\mathbf{a}_m \in \mathbb{R}^2\}_{m=1}^M$, and that measure the energy generated by a point source at location $\boldsymbol{\theta} \in \mathbb{R}^2$. The measurements are given as

$$y_m = h_m(\boldsymbol{\theta}) + n_m, \quad m = 1, 2, \dots, M, \tag{3.30}$$

where $h_m(\cdot)$ is an isotropic exponential attenuation function that is given by $h_m(\boldsymbol{\theta}) = \sqrt{e} \beta / (\beta + d_m^2)$ with $d_m = \|\boldsymbol{\theta} - \mathbf{a}_m\|_2$. Here, e is the known field intensity emitted or reflected by the point source, $\beta \geq 0$ is the gain, and n_m is the noise. In this case, we have

$$\mathbf{h}(\boldsymbol{\theta}) = [h_1(\boldsymbol{\theta}), h_2(\boldsymbol{\theta}), \dots, h_M(\boldsymbol{\theta})]^T$$

and $\mathbf{n} = [n_1, n_2, \dots, n_M]^T$ with $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$. The noise covariance matrix $\boldsymbol{\Sigma}$ might not be diagonal due to the multi-path effects, for example. Using the above model parameters, we have to choose the best subset of sensor locations out of M available locations such that a desired estimation accuracy for estimating $\boldsymbol{\theta}$ is achieved.

We consider the scenario shown in Figure 3.5 with $M = 80$ sensors. We assume that the noise correlation matrix is of the form

$$\Sigma = \begin{bmatrix} \Sigma_{\text{horz}} & \mathbf{0} \\ \mathbf{0} & \Sigma_{\text{vert}} \end{bmatrix} \in \mathbb{R}^{M \times M},$$

where Σ_{horz} is the noise correlation matrix corresponding to the horizontally located candidate sensors (denoted by (\diamond) in Figure 3.5) and Σ_{vert} is the noise correlation matrix corresponding to the vertically located candidate sensors (indicated by (\square) in Figure 3.5). We further assume that

$$\Sigma_{\text{horz}} = \sigma_e^2 [(1 - \rho)\mathbf{I} + \rho\mathbf{1}\mathbf{1}^T]$$

with correlation coefficient ρ and nominal noise variance σ_e^2 , and $\Sigma_{\text{vert}} = \sigma_e^2\mathbf{I}$. That is, the vertically located candidate sensors are uncorrelated while the horizontally located candidate sensors are equally correlated.

We use the following simulation parameters: $e = 1$, $\beta = 1$ and $\sigma_e^2 = 2 \times 10^{-5}$, $R_e = 20$ m, $\rho = 0.5$, and $P_e = 0.9$. Figure 3.5(a) shows the sensor placement obtained by solving Algorithm 3.2, but with (3.18) replaced with (3.29) and with $i_{\text{max}} = 10$ and $\delta = 10^{-6}$. Here, we use deterministic rounding. Figure 3.5(b) shows the selected sensor index.

We underline the following observations. Firstly, the sensors from the same region are selected as in Figure 3.1(a). This is due to the structural similarity of the FIMs corresponding to the models (3.23) and (3.30). Secondly, to achieve the desired performance requirement we see that fewer number of correlated sensors are selected as compared to the number of selected uncorrelated sensors. Finally, as we observed in the simulations (not shown here), for this particular numerical example, the ℓ_1 -norm based solution does not result in a sparse selection.

3.8 Discussion

In this chapter we discussed discrete sparse sensing for nonlinear parameter estimation problems. In particular, we focussed on observations having arbitrary yet independent distributions and colored Gaussian observations. The assumption is that the model parameters are perfectly known. We used

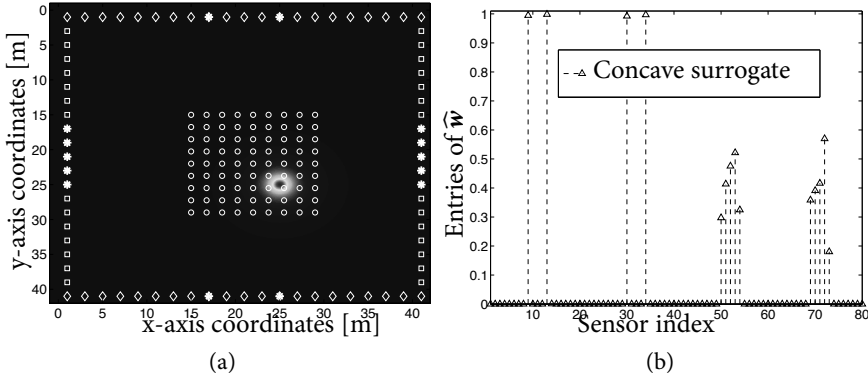


Figure 3.5: Sensor placement for field estimation. Here, the uncorrelated and correlated sensors are denoted by squares (\square) and diamonds (\diamond), respectively. The source domain is indicated by circles (\circ), while the selected sensors are indicated by ($*$). (a) Illustration of a field generated by a unit amplitude point source at location $\theta = [25, 25]^T$ m according to (3.30). Out of $M = 80$ available sensors 14 sensors are selected. (b) Sensor selection is solved using log-based heuristics. The thresholds are computed using $R_e = 20$ m, and $P_e = 0.9$. We use $e = 1$, $\beta = 1$ and $\sigma_e^2 = 2 \times 10^{-5}$.

a number of scalar functions related to the Cramér-Rao bound to design the sparse sensing function. The Cramér-Rao bound was used as a weaker surrogate for the error covariance matrix, which for nonlinear models does not admit a closed-form expression. The original nonconvex optimization problem is relaxed using convex relaxation techniques, which can then be efficiently solved in polynomial time. To handle large-scale problems, we have also presented a projected subgradient algorithm. The proposed framework is applied to sensor placement design for localization.

3.A Performance thresholds

Trace and minimum eigenvalue constraints

We can relate the accuracy requirement and CRB using Chebyshev's inequality [Cover and Thomas, 2012]

$$\Pr(\|\boldsymbol{\epsilon}\|_2 \geq R_e) \leq \text{tr}\{\mathbf{C}\}/R_e^2$$

which can be equivalently expressed as $\Pr(\|\boldsymbol{\varepsilon}\|_2 \leq R_e) \leq 1 - \text{tr}\{\mathbf{C}\}/R_e^2$. Combining this inequality together with $\Pr(\|\boldsymbol{\varepsilon}\|_2 \leq R_e) \geq P_e$ in (3.7) results in the following sufficient condition

$$\text{tr}\{\mathbf{C}\} \leq \lambda_{\text{tr}} = (1 - P_e)R_e^2. \quad (3.31)$$

Each eigenvalue of \mathbf{C}^{-1} is greater than $\lambda_{\min}(\mathbf{F})$, and as a result, $\text{tr}\{\mathbf{C}\} \leq N\lambda_{\min}^{-1}(\mathbf{F})$. Hence, a stronger sufficient condition (with a smaller feasible set) is $N\lambda_{\min}^{-1}(\mathbf{F}) \leq (1 - P_e)R_e^2$, or equivalently [Wang et al., 2009]

$$\lambda_{\min}(\mathbf{F}) \geq \lambda_{\text{eig}} = \frac{N}{R_e^2} \left(\frac{1}{1 - P_e} \right).$$

Determinant constraint

The determinant constraint is related to the volume or the mean radius of the confidence ellipsoid that contains $\boldsymbol{\varepsilon} = \boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}$ with probability P_e . Such a confidence ellipsoid can be expressed as

$$\mathcal{E} = \{\boldsymbol{\varepsilon} \mid \boldsymbol{\varepsilon}^T \mathbf{F}^{-1} \boldsymbol{\varepsilon} \leq \xi\},$$

where ξ is a constant that depends on P_e . Assuming \mathbf{F} has ordered eigenvalues $\lambda_{\max} \geq \lambda_2 \cdots \geq \lambda_{\min}$, the length of the n th semi-axis of the ellipsoid \mathcal{E} will be $\sqrt{\xi/\lambda_n}$. The geometric mean radius of the confidence ellipsoid \mathcal{E} is given by

$$\bar{R}_e = \sqrt{\xi}/(\det\{\mathbf{F}\})^{1/2N},$$

which gives a quantitative measure of how informative the observations are. For the estimates to be within the confidence ellipsoid \mathcal{E} , we use the constraint

$$\ln \det\{\mathbf{F}\} \geq 2N \ln \frac{\sqrt{\xi}}{\bar{R}_e} = \lambda_{\text{det}},$$

where \bar{R}_e and $\sqrt{\xi}$ specify the required accuracy, and are assumed to be known. A typical choice for ξ is $\xi = F_{\chi_N^2}^{-1}(P_e)$. Here, $F_{\chi_N^2}^{-1}$ is the cumulative distribution function of a chi-squared random variable with N degrees of freedom. This performance measure is related to the *D-optimality*.

3.B Projected Newton's method

In order to analyze the complexity of the interior point methods, we briefly describe the projected Newton's method. The Newton's method for an SDP problem in the inequality form is adapted to suit our problem [Boyd and Vandenberghe, 2004, p. 619].

The optimization problem in (3.12) can be approximated using the log-determinant barrier function which is given as

$$\arg \min_{\mathbf{w} \in [0,1]^N} \psi(\mathbf{w}) = t \mathbf{1}_M^T \mathbf{w} - \ln \det \left\{ \sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN} \right\},$$

where $t > 0$ is a parameter to tune the approximation. The projected Newton's update equation is given by

$$\mathbf{w}^{k+1} = \mathcal{P}_{\mathcal{W}} \left(\mathbf{w}^k - \alpha^k \mathbf{H}_{\psi}^{-1}(\mathbf{w}^k) \mathbf{g}_{\psi}(\mathbf{w}^k) \right), \quad (3.32)$$

where the (i, j) th entry of the Hessian matrix is given by

$$[\mathbf{H}_{\psi}(\mathbf{w}^k)]_{i,j} = \left. \frac{\partial^2 \psi(\mathbf{w})}{\partial w_i \partial w_j} \right|_{\mathbf{w}=\mathbf{w}^k} = \text{tr} \{ \mathbf{S}^{-1} \mathbf{F}_i \mathbf{S}^{-1} \mathbf{F}_j \},$$

and the i th entry of the gradient vector is given by

$$[\mathbf{g}_{\psi}(\mathbf{w}^k)]_i = \left. \frac{\partial \psi(\mathbf{w})}{\partial w_i} \right|_{\mathbf{w}=\mathbf{w}^k} = t + \text{tr} \{ \mathbf{S}^{-1} \mathbf{F}_i \}.$$

Here, we have introduced the matrix $\mathbf{S} = \sum_{m=1}^M w_m \mathbf{F}_m - \lambda_{\text{eig}} \mathbf{I}_{DN}$, and recall the projection operator $\mathcal{P}_{\mathcal{W}}(\cdot)$ defined in (3.15). The step length α^k is chosen by line search.

3.C Power iterations

We briefly describe the power iterations [Golub and Van Loan, 1996] to compute the minimum eigenvalue of a matrix $\mathbf{F} \in \mathbb{S}^N$. Assuming \mathbf{F} has ordered eigenvalues $\lambda_{\max} \geq \lambda_2 \cdots \geq \lambda_{\min}$, the power iterations

$$\mathbf{v}^{k+1} = \frac{\mathbf{F} \mathbf{v}^k}{\|\mathbf{F} \mathbf{v}^k\|_2}, \text{ and } \lambda^{k+1} = \frac{(\mathbf{v}^{k+1})^T \mathbf{F} \mathbf{v}^{k+1}}{\|\mathbf{v}^{k+1}\|_2},$$

converge to the maximum eigenvector \mathbf{v}_{\max} and maximum eigenvalue λ_{\max} , respectively, as $k \rightarrow \infty$. Here, we use $\mathbf{v}^0 = [1, \mathbf{0}_{N-1}^T]^T$. By forming a matrix $\tilde{\mathbf{F}} = \lambda_{\max} \mathbf{I}_N - \mathbf{F}$ which has the dominant eigenvalue $\lambda_{\max} - \lambda_{\min}$, we can apply the above power iterations on $\tilde{\mathbf{F}}$ to compute $\lambda_{\max} - \lambda_{\min}$ and \mathbf{v}_{\min} , and thus the minimum eigenvalue of \mathbf{F} and its corresponding eigenvector.

Chapter 4

Sparse Sensing for Filtering

Contents

4.1	Introduction	72
4.2	Sensing time-varying observations	74
4.3	$f(w)$ for filtering	75
4.4	Problem statement	78
4.5	Solvers	79
4.6	Numerical example: sensor scheduling	81
4.7	Dependent observations	85
4.8	Structured signals	86
4.9	Discussion	92
4.A	Extended Kalman filter	92

Part of this chapter was published as: S.P. Chepuri and G. Leus, “Sparsity-Promoting Adaptive Sensor Selection for nonlinear Filtering”, in *Proc. of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2014)*, Florence, Italy, May 2014, and S.P. Chepuri and G. Leus, “Compression schemes for time-varying sparse signals”, in *Proc. of Asilomar Conf. Signals, systems, and Computers (Asilomar 2014)*, Pacific Grove, California, USA, Nov. 2014.

4.1 Introduction

In this chapter we will extend the discrete sparse sensing framework for estimation introduced in Chapter 3 to nonlinear filtering problems. The proposed framework is useful for a variety of applications related to target or bearing tracking, dynamic field estimation, and nonlinear filtering in general, where the sensors (or measurements) are scheduled or activated parsimoniously to increase their battery lifetime as well as to reduce the communications and inference costs. The time-varying state parameters are described through a state-space model, which is assumed to be known. Given the state-space model, we will design sparse sensing mechanisms and develop the associated risk functions that quantify the filtering accuracy.

Over the past decade, sensor scheduling/polling problems for state estimation of linear dynamical systems (i.e., linear filtering problems) have been extensively studied with different flavors, such as different optimization criteria from experiment design, myopic (one-step ahead) and nonmyopic (a longer time horizon) scheduling, heuristic (submodular, convex, among others) algorithms, and budget constraints; see [Hernandez et al., 2004, Zuo et al., 2007, Krause, 2008, Joshi and Boyd, 2009, Fu et al., 2012, Jiang et al., 2013, Liu et al., 2014, Zhan et al., 2010]. For example, [Joshi and Boyd, 2009] proposed convex relaxation techniques, whereas [Krause, 2008] proposed greedy algorithms using submodular cost functions for sensor scheduling. In [Liu et al., 2014], the design of optimal periodic sensor scheduling (over an infinite time horizon) with restrictions on the total number of sensor activations was studied for additive linear Gaussian models. The optimal experiment design problems are well-studied for observations that follow a linear model with uncorrelated noise components for which the error covariance matrix has a known closed form. However, it is in general difficult to compute the error covariance matrix in closed form for nonlinear dynamical systems. Therefore, the above methods cannot be used directly.

We organize this chapter into three parts discussing sparse sensing mechanisms for: independent observations, dependent observations, and structured state sequences. To begin with, we focus on nonlinear state-space models with independent observations. In [Masazade et al., 2012], sensor selection for target tracking based on extended Kalman filtering has been pro-

posed, in which the selection is performed by designing a sparse gain matrix. Moreover, [Masazade et al., 2012] focuses on an additive Gaussian nonlinear model, that is linearized around the (noisy) past state estimate. This chapter, on the other hand, deals with general nonlinear models, without an explicit linearization. For a general nonlinear filtering problem, we use a number of scalar measures (from experiment design) of the posterior Cramér-Rao bound matrix. We saw in Chapter 3 that for nonlinear estimation problems the risk functions based on the Cramér-Rao bound depend on the true parameter. Similarly, for nonlinear filtering problems the risk functions based on the posterior Cramér-Rao bound depend on the current as well as the previous true state parameters. The sensing patterns are time-varying, and we do not need actual measurements to design them (i.e., they are data independent, but model driven). However, running a nonlinear filter in parallel and incorporating the entire history of (actual) measurements up to that point significantly simplifies the design problem. We further model the evolution of the sensing pattern in time to control its smoothness. This is beneficial for mobile sensing (e.g., sensing with sensors mounted on bicycles or automobiles), for instance. Next, we extend the framework to include nonlinear additive Gaussian models with correlated measurements.

We also study discrete sparse sensing for time-varying structured signals. In particular, we will illustrate this framework with time-varying sparse signals with possibly time-varying sparsity patterns and/or order. This has received a lot of attention in the recent past through *compressive sensing* (CS) [Baraniuk, 2007]. Time-varying sparse signal reconstruction has been studied in the past leading to various forms of sparsity-aware filters [Vaswani, 2008, Angelosante et al., 2009, Carmi et al., 2010, Malioutov et al., 2010], and are applied to problems like visual surveillance [Warnell et al., 2012] and target localization [Farahmand et al., 2014]. We study the design of sensing matrices for such problems; however, the focus will not be on the signal recovery itself.

Sensing matrix design for sparse recovery has been studied in various forms. For example, in [Eldar and Kutyniok, 2012, Ch. 6], [Haupt et al., 2009] the variance of the distribution from which the (random) sensing matrices are generated is designed such that the average information gain is maximized. The Bayesian CS framework [Ji et al., 2008] allows to quantify the sparse reconstruction error through the so-called error bars, which again allows to

adaptively design the sensing matrices. Both [Haupt et al., 2009] and [Ji et al., 2008] use experiment design techniques with performance measures like differential entropy to adaptively learn the sensing matrix starting from a random matrix. In [El Badawy et al., 2014], a greedy algorithm based on a sub-modular performance measure has been proposed for sensing operator design for a signal lying in the union of subspaces. However, the sensing design schemes discussed above are mostly limited to time-invariant signals and/or systems without any state-space representation. Hence, they are not adaptive in the true sense, and are not meant for tracking the signal variation over space and/or time. On the other hand, the proposed sensing matrix is designed at each time step based on the entire history of measurements and known dynamics described through a state-space model. Towards the end, we will also discuss a few extensions of the proposed framework to include general structured signals, such as group sparse and smooth signals.

4.2 Sensing time-varying observations

We consider a nonlinear measurement model for observing an unknown dynamic parameter $\boldsymbol{\theta}_k \in \mathbb{R}^{N \times 1}$ at time k :

$$\mathbf{x}_{k,m} = h_{k,m}(\boldsymbol{\theta}_k, n_{k,m}), m = 1, 2, \dots, M, \quad (4.1)$$

where $h_{k,m}(\cdot, \cdot)$ is a nonlinear functional of the unknown vector $\boldsymbol{\theta}_k$ and the noise component $n_{k,m}$. The spatial (and/or temporal) measurements at temporal block k , $\{\mathbf{x}_{k,m}\}_{m=1}^M$, are stacked in the measurement vector $\mathbf{x}_k = [x_{k,1}, x_{k,2}, \dots, x_{k,M}]^T \in \mathbb{R}^{M \times 1}$.

The unknown parameter is assumed to obey the dynamical model:

$$\boldsymbol{\theta}_{k+1} = \mathbf{A}_k \boldsymbol{\theta}_k + \mathbf{u}_k, \quad (4.2)$$

where $\mathbf{A}_k \in \mathbb{R}^{N \times N}$ is the state transition matrix and $\mathbf{u}_k \in \mathbb{R}^{N \times 1}$ is the process noise that accounts for any unmodeled dynamics. Here, we model $\mathbf{u}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_u)$, where $\boldsymbol{\Sigma}_u \in \mathbb{R}^{N \times N}$ represents the state noise covariance matrix.

At each time instance k , we acquire the data \mathbf{x}_k via the discrete sparse sensing mechanism that was introduced in Chapter 2. Specifically, we acquire data as

$$\mathbf{y}_k = \text{diag}_r(\mathbf{w}_k) \mathbf{x}_k = \boldsymbol{\Phi}_k(\mathbf{w}_k) \mathbf{x}_k$$

where $\Phi_k(\mathbf{w}_k) = \text{diag}_r(\mathbf{w}_k) \in \{0, 1\}^{K \times M}$ is the sparse sensing matrix whose entries are determined by the time-varying selection vector

$$\mathbf{w}_k = [w_{k,1}, w_{k,2}, \dots, w_{k,M}]^T \in \{0, 1\}^M.$$

Here, the time-varying variable $w_{k,m} = (0)1$ indicates whether $x_{k,m}$ is (not) selected.

We are interested in cases where $K \ll M$ and K is not known. The reduced dimension data vector $\mathbf{y}_k \in \mathbb{R}^K$ is used instead of $\mathbf{x}_k \in \mathbb{R}^M$ together with (4.2) to solve the nonlinear filtering problem. Therefore, our aim is to design, for each k , a sparsest \mathbf{w}_k (and, hence a sequence of matrices $\{\Phi_k(\mathbf{w}_k), k \in \mathbb{N}\}$) based on the entire history of measurements up to that point $\{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{k-1}\}$, such that a desired accuracy on the (*a posteriori*) estimate of $\boldsymbol{\theta}_k$ is guaranteed. For such nonlinear filtering problems, the risk function $f(\mathbf{w}_k)$ that quantifies the filtering accuracy is discussed next.

4.3 $f(\mathbf{w})$ for filtering

For general nonlinear filtering problems, it is difficult to compute the posterior error covariance matrix in closed form that is suitable for optimization. Therefore, along the lines of §3.3 of Chapter 3, we will discuss a weaker surrogate that can be optimized instead of the posterior error covariance.

We now recall Assumption 3.1 from Chapter 3 as:

Assumption 4.1 (Regularity condition). *The log-likelihood of the measurements satisfies the regularity condition, i.e., $\mathbb{E}\{\partial \ln p_k(\mathbf{y}; \boldsymbol{\theta}_k) / \partial \boldsymbol{\theta}\} = \mathbf{0}$, for all k , where $p_k(\mathbf{y}; \boldsymbol{\theta}_k)$ is the probability density function (pdf) of \mathbf{y}_k at time instance k parameterized by the unknown vector $\boldsymbol{\theta}_k$.*

Under the above assumption, the *a posteriori* estimate of $\boldsymbol{\theta}_k$ denoted by $\widehat{\boldsymbol{\theta}}_{k|k}$ satisfies the well-known posterior Cramér-Rao bound (PCRB) inequality given by

$$\mathbb{E}\{(\widehat{\boldsymbol{\theta}}_{k|k} - \boldsymbol{\theta}_k)(\widehat{\boldsymbol{\theta}}_{k|k} - \boldsymbol{\theta}_k)^T\} \geq \mathbf{C}_k = \mathbf{F}_k^{-1}, \quad (4.3)$$

where \mathbf{C}_k is the PCRB matrix, \mathbf{F}_k is the posterior Fisher information matrix (FIM), and the notation $\widehat{\boldsymbol{\theta}}_{k|l}$ denotes the estimate of $\boldsymbol{\theta}_k$ based on data up to

point l . The posterior FIM is given by the following recursion [Tichavsky et al., 1998]:

$$\mathbf{F}_k(\mathbf{w}_k, \boldsymbol{\theta}_k, \{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) = \overbrace{(\boldsymbol{\Sigma}_u + \mathbf{A}_k \mathbf{F}_{k-1}^{-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) \mathbf{A}_k^T)^{-1}}^{\mathbf{F}_{\text{prior},k-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k)} + \mathbf{F}_{\text{obs},k}(\mathbf{w}_k, \boldsymbol{\theta}_k) \in \mathbb{R}^{N \times N}, \quad (4.4)$$

with the FIM related to the observations at time instance k given by

$$\mathbf{F}_{\text{obs},k}(\mathbf{w}_k, \boldsymbol{\theta}_k) = \mathbb{E} \left\{ \left(\frac{\partial \ln p_k(\mathbf{y}; \boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}_k} \right) \left(\frac{\partial \ln p_k(\mathbf{y}; \boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}_k} \right)^T \right\} \in \mathbb{R}^{N \times N}. \quad (4.5)$$

The first term of (4.4) is the prior information related to the history up to time instance $k-1$.

We now introduce the independence assumption.

Assumption 4.2 (Independent observations). *The observations $\{x_{k,m}\}_{m=1}^M$ at time k are a sequence of independent random variables, which depend on the unknown parameter $\boldsymbol{\theta}_k$.*

Under Assumption 4.2, the information measure from each observation is *additive* in the variables $\{w_{k,m}\}_{m=1}^M$, which is intuitive as each independent measurement contributes some information (we have seen this property before in Chapter 3). Using this property, we can further simplify (4.4) to arrive at

$$\mathbf{F}_k(\mathbf{w}_k, \boldsymbol{\theta}_k, \{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) = \mathbf{F}_{\text{prior},k-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) + \sum_{m=1}^M w_{k,m} \mathbf{F}_{k,m}(\boldsymbol{\theta}_k), \quad (4.6)$$

where

$$\mathbf{F}_{k,m}(\boldsymbol{\theta}_k) = \mathbb{E} \left\{ \left(\frac{\partial \ln p_{k,m}(x; \boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}_k} \right) \left(\frac{\partial \ln p_{k,m}(x; \boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}_k} \right)^T \right\} \in \mathbb{R}^{N \times N}$$

with $p_{k,m}(x; \boldsymbol{\theta}_k)$ being the pdf of $x_{k,m}$ at time instance k parameterized by the unknown vector $\boldsymbol{\theta}_k$. The posterior FIM (4.6) depends on the true state $\boldsymbol{\theta}_k$ at time k as well as all the previous states.

Remark 4.1 (Additive Gaussian linear state-space models).

For a linear measurement model: $x_{k,m} = \mathbf{h}_{k,m}^T \boldsymbol{\theta}_k + n_{k,m}$, $m = 1, 2, \dots, M$, (i.e., $h_{k,m}(\boldsymbol{\theta}_k, n_{k,m}) := \mathbf{h}_{k,m}^T \boldsymbol{\theta}_k + n_{k,m}$) where $\mathbf{h}_{k,m} \in \mathbb{R}^N$ is the regressor. Assuming that the noise components $\{n_{k,m}\}_{m=1}^M$ are i.i.d. Gaussian with variance σ^2 , the posterior FIM for linear filtering is given by

$$\mathbf{F}_k(\mathbf{w}_k) = (\boldsymbol{\Sigma}_u + \mathbf{A}_k \mathbf{F}_{k-1}^{-1} \mathbf{A}_k^T)^{-1} + \frac{1}{\sigma^2} \sum_{m=1}^M w_{k,m} \mathbf{h}_{k,m} \mathbf{h}_{k,m}^T. \quad (4.7)$$

The PCRB for linear state-space models in additive Gaussian noise is equal to the posterior error covariance, and it is independent of the unknown state vectors.

As a performance measure, we constrain the posterior estimation error $\boldsymbol{\epsilon}_k = \widehat{\boldsymbol{\theta}}_{k|k} - \boldsymbol{\theta}_k$ to be within an origin centered circle of radius R_e with a probability higher than P_e , i.e.,

$$\Pr(\|\boldsymbol{\epsilon}_k\|_2 \leq R_e) \geq P_e, \quad (4.8)$$

where R_e and P_e are specified to achieve a desired accuracy. This accuracy constraint is satisfied by the two popular experiment design criteria [cf. §3.3 of Chapter 3]:

1. Trace constraint, which is related to the A-optimality measure. A sufficient condition for (4.8) is (see Appendix 3.A)

$$f(\mathbf{w}_k) := \text{tr} \left\{ \left(\mathbf{F}_k(\mathbf{w}_k, \boldsymbol{\theta}_k, \{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) \right)^{-1} \right\} \leq \lambda_{\text{tr}} = (1 - P_e) R_e^2.$$

2. Minimum eigenvalue constraint, which is related to the E-optimality measure. A sufficient condition for (4.8) related to the eigenvalue constraint is (see Appendix 3.A)

$$f(\mathbf{w}_k) := \lambda_{\min} \left\{ \mathbf{F}_k(\mathbf{w}_k, \boldsymbol{\theta}_k, \{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) \right\} \geq \lambda_{\text{eig}} = \frac{N}{R_e^2} \left(\frac{1}{1 - P_e} \right).$$

The above inequality can be alternatively expressed as the following linear matrix inequality (LMI):

$$\mathbf{F}_{\text{prior},k-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) + \sum_{m=1}^M w_{k,m} \mathbf{F}_{k,m}(\boldsymbol{\theta}_k) \geq \lambda_{\text{eig}} \mathbf{I}_N. \quad (4.9)$$

We underline here that the determinant constraint similar to §3.5.5 of Chapter 3 can also be used as a reasonable performance measure. Unless the process noise is zero (e.g., a deterministic trajectory), the true states are not known. However, in practice $\boldsymbol{\theta}_k$ takes values within a certain domain denoted by the set \mathcal{U}_k . Therefore, we constrain (4.9) for all $\boldsymbol{\theta}_\kappa \in \mathcal{U}_\kappa, \kappa = 0, 1, 2, \dots, k$.

As discussed in Chapter 3, the trace constraint, which can also be represented by LMIs, has a larger feasible set as compared to the minimum eigenvalue constraint. However, the resulting optimization problem is computationally less attractive compared to the minimum eigenvalue constraint (as we show later on in §3.5.5). For this reason, we focus on the minimum eigenvalue (LMI) constraint from now on.

4.4 Problem statement

The *adaptive discrete sparse sensing* (or adaptive sensor selection) problem can be interpreted as the problem of choosing the best subset of sensors out of the M available sensors to acquire measurements for time step k such that a certain accuracy on the estimate $\widehat{\boldsymbol{\theta}}_{k|k}$ is guaranteed. Thus, the adaptive sensor selection problem can be formally stated as follows.

Problem 4.1 (Adaptive discrete sparse sensing). *Given the state-space model (4.1) and (4.2), at each time step k , based on the entire history of measurements up to that point find a sparsest $\mathbf{w}_k \in \{0, 1\}^M$, which satisfies the accuracy constraint $\mathbf{F}_{\text{prior}, k-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) + \sum_{m=1}^M w_{k,m} \mathbf{F}_{k,m}(\boldsymbol{\theta}_k) \geq \lambda_{\text{eig}} \mathbf{I}_N, \forall \boldsymbol{\theta}_\kappa \in \mathcal{U}_\kappa, \kappa = 0, 1, 2, \dots, k$.*

At each time step, the above design problem is a specialization of the discrete sparse sensing problem (P0) introduced in Chapter 2, that is,

$$\arg \min_{\mathbf{w}_k} \|\mathbf{w}_k\|_0 \quad (4.10a)$$

$$\text{s.to } \mathbf{F}_{\text{prior}, k-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k) + \sum_{m=1}^M w_{k,m} \mathbf{F}_{k,m}(\boldsymbol{\theta}_k) \geq \lambda_{\text{eig}} \mathbf{I}_N, \quad (4.10b)$$

$$\forall \boldsymbol{\theta}_\kappa \in \mathcal{U}_\kappa, \kappa = 0, 1, 2, \dots, k,$$

$$\mathbf{w}_k \in \{0, 1\}^M. \quad (4.10c)$$

This is a nonconvex Boolean optimization problem. Clearly, the number of LMI constraints (4.10b) depends on $|\mathcal{U}_k|$ and this increases with k , that is, the number of LMI constraints is $\sum_{\kappa=0}^k |\mathcal{U}_\kappa|$. For example, if $|\mathcal{U}_\kappa| = D, \forall \kappa$, then at time step k , the optimization problem (4.10) will have D^k size- N LMI constraints. To reduce the computational complexity, the prior Fisher $F_{\text{prior},k-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k)$ can be simply evaluated at the past estimates

$$\tilde{\boldsymbol{\theta}}_{k-1} := \widehat{\boldsymbol{\theta}}_{k-1|k-1},$$

obtained by solving any nonlinear filter (e.g., extended or unscented Kalman filter, or particle filters). That is, we approximate $F_k(\mathbf{w}_k, \boldsymbol{\theta}_k, \{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^k)$ as

$$F_k(\mathbf{w}_k, \boldsymbol{\theta}_k) \approx \underbrace{F_{\text{prior},k-1}(\tilde{\boldsymbol{\theta}}_{k-1})}_{\text{prior}} + F_{\text{obs},k}(\mathbf{w}_k, \boldsymbol{\theta}_k) \in \mathbb{R}^{N \times N},$$

Therefore, running an extended Kalman filter in parallel (and incorporating the history of actual measurements) significantly reduces the number of constraints. Henceforth, we will not write the prior FIM $F_{\text{prior},k-1}(\tilde{\boldsymbol{\theta}}_{k-1})$ as an explicit function of $\tilde{\boldsymbol{\theta}}_{k-1}$, i.e., we will simply write it as $F_{\text{prior},k-1}$. As a consequence, (4.10) simplifies to

$$\arg \min_{\mathbf{w}_k} \|\mathbf{w}_k\|_0 \quad (4.11a)$$

$$\text{s.to } F_{\text{prior},k-1} + \sum_{m=1}^M w_{k,m} F_{k,m}(\boldsymbol{\theta}_k) \succeq \lambda_{\text{eig}} \mathbf{I}_N, \forall \boldsymbol{\theta}_k \in \mathcal{U}_k, \quad (4.11b)$$

$$\mathbf{w}_k \in \{0, 1\}^M. \quad (4.11c)$$

The posterior error covariance matrix and the prediction computed in the extended Kalman filter algorithm (see Appendix 4.A) can be used as a guideline to determine \mathcal{U}_k (we will discuss this with an example in §4.6).

4.5 Solvers

In this section we provide an algorithm to solve the proposed optimization problem. In addition, we also model the evolution of \mathbf{w}_k in time, which is useful to control the transient nature of the sensing patterns.

4.5.1 Convex approximation based on ℓ_1 -norm

The optimization problem in (4.10) is nonconvex due to the ℓ_0 (-quasi) norm cost function and the Boolean constraint. We use the traditional best convex surrogate for the ℓ_0 (-quasi) norm based on the ℓ_1 -norm heuristic, and the Boolean constraint is relaxed to the convex box constraint $[0, 1]^M$. Due to this box constraint, the ℓ_1 -norm will simply be the affine function $\mathbf{1}_M^T \mathbf{w}_k$. The relaxed optimization problem is of the form (R0) introduced in Chapter 2, and is given as

$$\widehat{\mathbf{w}}_k = \arg \min_{\mathbf{w}_k} \mathbf{1}_M^T \mathbf{w}_k \quad (4.12a)$$

$$\text{s.to } \mathbf{F}_{\text{prior}, k-1} + \sum_{m=1}^M w_{k,m} \mathbf{F}_{k,m}(\boldsymbol{\theta}_k) \geq \lambda_{\text{eig}} \mathbf{I}_N, \quad \forall \boldsymbol{\theta}_k \in \mathcal{U}_k, \quad (4.12b)$$

$$0 \leq w_{k,m} \leq 1, \quad m = 1, 2, \dots, M, \quad (4.12c)$$

where the LMI constraint (4.12b) is convex on $\mathbf{w}_k \in [0, 1]^M$.

The relaxed optimization problem is a standard semidefinite programming problem that can be solved efficiently in polynomial time using off-the-shelf solvers like SeDuMi [Sturm, 1999] or the projected subgradient algorithm developed in §3.5.2 of Chapter 3. Further, a concave surrogate based on the sum-of-logarithms can be used to approximate the ℓ_0 (-quasi) norm to obtain sparsity enhancing iterations as discussed in §3.5.3 of Chapter 3. An approximate Boolean solution can then be recovered by randomization rounding as explained in §3.5.4 of Chapter 3.

4.5.2 Smooth sensing

We now model the evolution of \mathbf{w}_k in time. A smooth evolution of the selection vector is important for mobile sensing to control the transient nature of a mobile sensor—a spatial sampling device. In other words, smoothness in the selection vector ensures an easy hand-off between the selected sensors.

The evolution of the selection vector is modeled as a linear recursion

$$\mathbf{w}_{k+1} = \mathbf{B}_k \mathbf{w}_k + \mathbf{v}_k, \quad (4.13)$$

where $\mathbf{B}_k \in \mathbb{R}^{M \times M}$ could be a banded matrix and $\mathbf{v}_k \in \mathbb{R}^M$ is the process noise vector. The smoothness depends on the construction of the matrix \mathbf{B}_k .

In order to incorporate the smoothing effect on the sensing pattern between subsequent time instances, we use the sparse estimate $\widehat{\mathbf{w}}_{k-1,\text{sm}}$ instead of \mathbf{w}_{k-1} (here, the subscript sm denotes smoothness). The optimization problem taking into account the smoothness is given as

$$\widehat{\mathbf{w}}_{k,\text{sm}} = \arg \min_{\mathbf{w}_k} \mathbf{1}_M^T \mathbf{w}_k + \mu \|\mathbf{w}_k - \mathbf{B}_{k-1} \widehat{\mathbf{w}}_{k-1,\text{sm}}\|_2^2 \quad (4.14a)$$

$$\text{s.to } \mathbf{F}_{\text{prior},k-1} + \sum_{m=1}^M w_{k,m} \mathbf{F}_{k,m}(\boldsymbol{\theta}_k) \geq \lambda \mathbf{I}_N, \forall \boldsymbol{\theta}_k \in \mathcal{U}_k, \quad (4.14b)$$

$$0 \leq w_{k,m} \leq 1, m = 1, 2, \dots, M, \quad (4.14c)$$

where μ is the smoothness controlling parameter.

4.6 Numerical example: sensor scheduling

We now apply the developed theory of adaptive sparse sensing to *sensor scheduling* for target tracking based on distance measurements. The sensor scheduling problem can be interpreted as the problem of polling (or activating) the best subset of sensors from a large pool of available sensors, such that a desired tracking accuracy is achieved. Sensor scheduling is typically used to increase the lifetime of battery powered sensor nodes, but also to reduce the communications and inferring costs.

At each time step k , the selected sensors are used to estimate the state vector $\boldsymbol{\theta}_k = [\mathbf{p}_k^T, \dot{\mathbf{p}}_k^T]^T \in \mathbb{R}^{4 \times 1}$, where $\mathbf{p}_k \in \mathbb{R}^{2 \times 1}$ is the target position vector and $\dot{\mathbf{p}}_k \in \mathbb{R}^{2 \times 1}$ is the velocity vector. We assume that M sensors are, respectively, located at known two-dimensional positions $\{\mathbf{a}_m\}_{m=1}^M$. They are each capable of measuring the distance to the target (e.g., from time of arrival measurements of the ranging signals). That is, we assume the measurement model

$$\begin{aligned} x_{k,m} &= \|\mathbf{p}_k - \mathbf{a}_m\|_2 + n_{k,m} \\ &= d_{k,m} + n_{k,m}, \quad m = 1, 2, \dots, M, \end{aligned}$$

where $n_{k,m} \sim \mathcal{N}(0, \sigma_{k,m}^2)$ with $\sigma_{k,m}^2 = \sigma^2/d_{k,m}^{-2}$ and σ^2 is the nominal noise variance. We use $\sigma^2 = 2 \times 10^{-5}$.

We consider an area of 60×60 square meter with $M = 49$ equally spaced sensors as shown in Figure 4.1(a). We use the following parameters for simu-

lations [Moon and Stirling, 2000]:

$$\mathbf{A}_k = \begin{bmatrix} 1 & 0 & \tau_s & 0 \\ 0 & 1 & 0 & \tau_s \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad \boldsymbol{\Sigma}_u = 10^{-2} \begin{bmatrix} \frac{\tau_s^3}{3} & 0 & \frac{\tau_s^2}{2} & 0 \\ 0 & \frac{\tau_s^3}{3} & 0 & \frac{\tau_s^2}{2} \\ \frac{\tau_s^2}{2} & 0 & \tau_s & 0 \\ 0 & \frac{\tau_s^2}{2} & 0 & \tau_s \end{bmatrix},$$

with sampling time $\tau_s = 2.5$ s. The dynamic model is initialized with $\mathbf{p}_0 \sim \mathcal{N}(\mathbf{1}_2, 2.778\mathbf{I}_2)$ and $\dot{\mathbf{p}}_0 \sim \mathcal{N}(2\mathbf{1}_2, 0.01\mathbf{I}_2)$ to emulate a target heading towards the north-east direction. We also run an extended Kalman filter in parallel and it is initialized with $\widehat{\boldsymbol{\theta}}_{0|0} = \mathbf{0}$, and $\mathbf{P}_{0|0} = 1000\mathbf{I}_2$. For the sake of completeness, we have summarized the extended Kalman filter algorithm in Appendix 4.A. The stochastic matrix \mathbf{B}_k is designed such that the transition to the one-hop sensor grid points and staying in the current state takes equal probabilities. In other words, in Figure 4.1(a), the corner most grid point has 3 one-hop neighbors, hence, it can move to any of these one-hop neighbors each with a probability of 1/4 or stay in the current state with a probability of 1/4. The parameters determining the accuracy are set to $R_e = 25$ cm and $P_e = 0.90$ to compute λ_{eig} .

We do not make velocity measurements, hence, we constrain only the FIM related to the distance measurements. Assuming that the FIM is composed of the following submatrices

$$\mathbf{F}_k(\mathbf{w}_k, \boldsymbol{\theta}_k) = \begin{bmatrix} \mathbf{F}_{k,pp} & \mathbf{F}_{k,p\dot{p}} \\ \mathbf{F}_{k,\dot{p}p} & \mathbf{F}_{k,\dot{p}\dot{p}} \end{bmatrix},$$

then using the Schur complement, the *a posteriori* position estimate $\widehat{\mathbf{p}}_{k|k}$ satisfies the following PCRBI inequality $\mathbb{E}\{(\widehat{\mathbf{p}}_{k|k} - \mathbf{p}_k)(\widehat{\mathbf{p}}_{k|k} - \mathbf{p}_k)^T\} \geq \tilde{\mathbf{F}}_k^{-1}(\mathbf{w}_k, \boldsymbol{\theta}_k)$, where

$$\tilde{\mathbf{F}}_k(\mathbf{w}_k, \boldsymbol{\theta}_k) = \mathbf{F}_{k,pp} - \mathbf{F}_{k,p\dot{p}}\mathbf{F}_{k,\dot{p}\dot{p}}^{-1}\mathbf{F}_{k,\dot{p}p}$$

assuming $\mathbf{F}_{k,\dot{p}\dot{p}}^{-1}$ exists. Similarly, assume that the prior FIM is also composed of the following submatrices

$$\mathbf{F}_{\text{prior},k}(\boldsymbol{\theta}_k) = \begin{bmatrix} \mathbf{F}_{\text{prior},k,pp} & \mathbf{F}_{\text{prior},k,p\dot{p}} \\ \mathbf{F}_{\text{prior},k,\dot{p}p} & \mathbf{F}_{\text{prior},k,\dot{p}\dot{p}} \end{bmatrix}.$$

Using the Schur complement and some straightforward matrix properties, we can show that

$$\tilde{\mathbf{F}}_k(\mathbf{w}_k, \boldsymbol{\theta}_k) = \tilde{\mathbf{F}}_{\text{prior},k-1} + \sum_{m=1}^M w_{k,m} \mathbf{F}_{k,m}(\mathbf{p}_k) \in \mathbb{R}^{2 \times 2},$$

where we define

$$\tilde{\mathbf{F}}_{\text{prior},k} = \mathbf{F}_{\text{prior},k,\mathbf{p}\mathbf{p}} - \mathbf{F}_{\text{prior},k,\mathbf{p}\hat{\mathbf{p}}} \mathbf{F}_{\text{prior},k,\hat{\mathbf{p}}\hat{\mathbf{p}}}^{-1} \mathbf{F}_{\text{prior},k,\hat{\mathbf{p}}\mathbf{p}}.$$

Therefore, for the case when only a part of the state parameters are measured, the LMI constraints in (4.10b) should be modified to

$$\tilde{\mathbf{F}}_{\text{prior},k-1} + \sum_{m=1}^M w_{k,m} \mathbf{F}_{k,m}(\mathbf{p}_k) \geq \lambda_{\text{eig}} \mathbf{I}_2, \quad \forall \boldsymbol{\theta}_k \in \mathcal{U}_k.$$

For a practical implementation of the algorithm, we construct \mathcal{U}_k based on the predicted estimate $\hat{\boldsymbol{\theta}}_{k|k-1}$ and the covariance matrix $\mathbf{P}_{k|k-1}$ (see Appendix 4.A). More specifically, \mathcal{U}_k is designed to have points within a circle of radius $5\sqrt{\text{tr}\{\mathbf{P}_{k|k-1}\}}$ centered around $\hat{\boldsymbol{\theta}}_{k|k-1}$. Since $\hat{\boldsymbol{\theta}}_{k|k-1} \sim \mathcal{N}(\boldsymbol{\theta}_k, \mathbf{P}_{k|k-1})$, the *true state* lies within a circle of radius $5\sqrt{\text{tr}\{\mathbf{P}_{k|k-1}\}}$ with an overwhelming probability.

Here, we discretize \mathcal{U}_k with 25 points as shown in Figure 4.1(a) and Figure 4.1(c) (indicated as the target area). For the sake of easy visibility, we plot the results in the time interval $(3\tau_s, 10\tau_s)$, as the target area is very large for initial estimates. Even though the predicted estimates are not necessarily on top of the true state, the true location will be within the target area with an overwhelming probability. Due to the assumed path-loss model, the sensors close to the target area are selected. The sensor activation time pattern without ($\mu = 0$) and with ($\mu > 0$) smooth sensing is shown in Figure 4.1(b) and Figure 4.1(d), respectively. An approximate Boolean solution is obtained by simply rounding all the nonzero entries of $\hat{\mathbf{w}}$ to one. The number of selected sensors with $\mu = 0.5$ is larger as compared to the case with $\mu = 0$. However, the sensors stay active for a longer duration ensuring a smooth hand-off between the selected sensors.

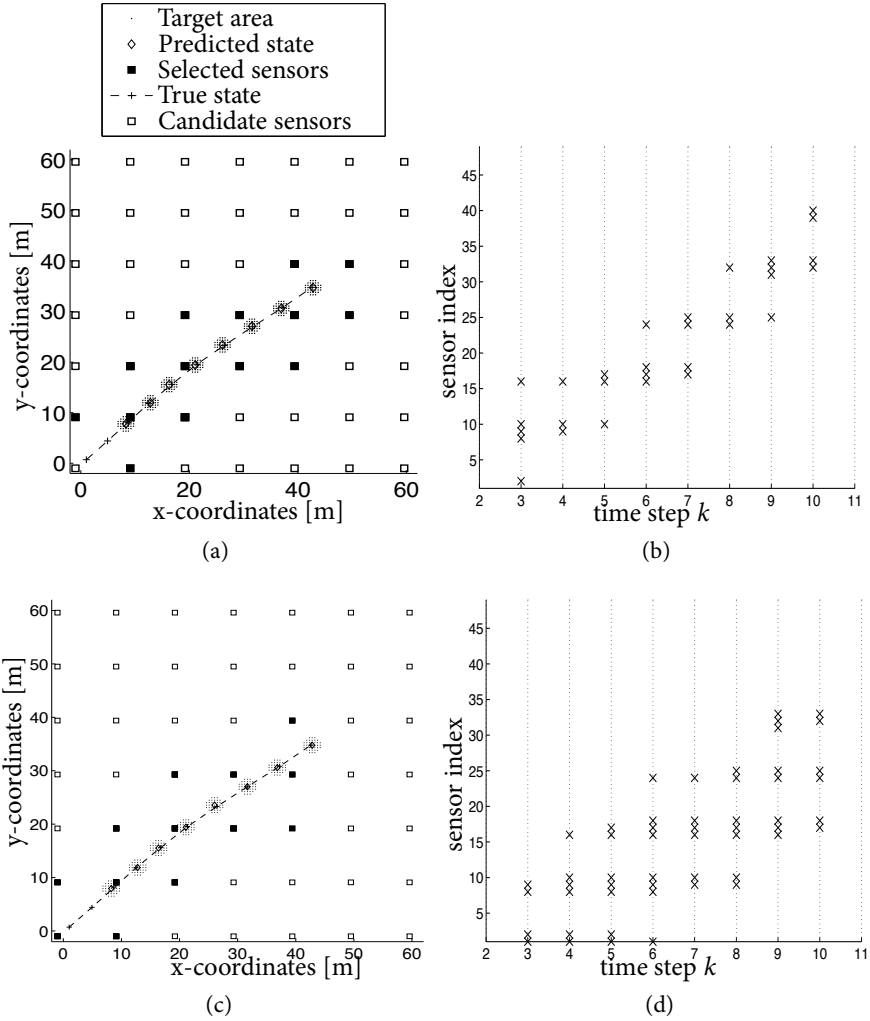


Figure 4.1: Sensor scheduling for target tracking based on range measurements for the time interval $(3\tau_s, 10\tau_s)$. (a) and (c) trajectory of the true state for a certain realization, predicted estimate from the extended Kalman filter, and the target area \mathcal{U}_k . (b) and (d) sensor activation time pattern for $\mu = 0$ (without smoothing) and $\mu = 0.5$ (with smoothing), respectively.

4.7 Dependent observations

So far we have assumed that the observations are mutually independent random variables. If we relax Assumption 4.1, then we can accommodate a much larger class of observation models, for example, to include ambient noises or stochastic state variables. However, by relaxing Assumption 4.1, the FIM information matrix related to the observations $\mathbf{F}_{\text{obs},k}(\mathbf{w}_k, \boldsymbol{\theta}_k)$, will not be linear in \mathbf{w}_k , which makes the optimization over \mathbf{w}_k much harder (i.e., nonconvex and nonlinear, in general). In this section, we will extend the results developed for sparse sensing for estimation with dependent Gaussian nonlinear measurements in §3.7 of Chapter 3 to filtering tasks.

Suppose the measurements (4.1) at time instance k are of the form

$$\mathbf{x}_k = \mathbf{h}_k(\boldsymbol{\theta}_k) + \mathbf{n}_k,$$

where $\mathbf{h}_k : \mathbb{R}^N \rightarrow \mathbb{R}^M$ is a nonlinear functional of the unknown vector $\boldsymbol{\theta}_k$ and \mathbf{n}_k is zero-mean Gaussian noise with variance $\boldsymbol{\Sigma}_k \in \mathbb{R}^{M \times M}$. For the above model, the FIM related to the observations [cf. (4.5)] will be

$$\mathbf{F}_{\text{obs},k}(\mathbf{w}_k, \boldsymbol{\theta}_k) = [\boldsymbol{\Phi}_k(\mathbf{w}_k)\mathbf{J}_k(\boldsymbol{\theta}_k)]^T \boldsymbol{\Sigma}_k^{-1}(\mathbf{w}_k) [\boldsymbol{\Phi}_k(\mathbf{w}_k)\mathbf{J}_k(\boldsymbol{\theta}_k)], \quad (4.15)$$

where $\mathbf{J}_k(\boldsymbol{\theta}_k) = \partial \mathbf{h}_k(\boldsymbol{\theta}_k) / \partial \boldsymbol{\theta}_k^T \in \mathbb{R}^{M \times N}$ and

$$\boldsymbol{\Sigma}_k(\mathbf{w}_k) = \boldsymbol{\Phi}_k(\mathbf{w}_k)\boldsymbol{\Sigma}_k\boldsymbol{\Phi}_k^T(\mathbf{w}_k) \in \mathbb{R}^{K \times K}$$

is a submatrix of $\boldsymbol{\Sigma}_k$ that includes only the entries corresponding to the selected measurements. The FIM (4.15) is no more additive or linear in \mathbf{w}_k .

Let us express the noise covariance matrix $\boldsymbol{\Sigma}_k$ as

$$\boldsymbol{\Sigma}_k = a_k \mathbf{I}_M + \mathbf{S}_k, \quad (4.16)$$

where any scalar $a_k \in \mathbb{R}$ is chosen such that $\mathbf{S}_k \in \mathbb{R}^{M \times M}$ is invertible and well conditioned. Using Property 3.1 from Chapter 3, we can equivalently express $\mathbf{F}_{\text{obs},k}(\mathbf{w}_k, \boldsymbol{\theta}_k)$ as

$$\begin{aligned} \mathbf{F}_{\text{obs},k}(\mathbf{w}_k, \boldsymbol{\theta}_k) &= \mathbf{J}_k^T(\boldsymbol{\theta}_k) \mathbf{S}_k^{-1} \mathbf{J}_k(\boldsymbol{\theta}_k) \\ &\quad - \mathbf{J}_k^T(\boldsymbol{\theta}_k) \mathbf{S}_k^{-1} [\mathbf{S}_k^{-1} + a_k^{-1} \text{diag}(\mathbf{w}_k)]^{-1} \mathbf{S}_k^{-1} \mathbf{J}_k^T(\boldsymbol{\theta}_k), \end{aligned} \quad (4.17)$$

where in contrast to (4.15), the design parameter \mathbf{w}_k appears only once in (4.17), which makes the problem much easier to solve.

Finally, using the Schur complement, the constraint $\lambda_{\min}\{\mathbf{F}_k(\mathbf{w}_k, \boldsymbol{\theta}_k)\} \geq \lambda_{\text{eig}}$ can now be equivalently expressed as a size- $M + N$ LMI:

$$\begin{bmatrix} \mathbf{S}_k^{-1} + a_k^{-1} \text{diag}(\mathbf{w}_k) & \mathbf{S}_k^{-1} \mathbf{J}_k(\boldsymbol{\theta}_k) \\ \mathbf{J}_k^T(\boldsymbol{\theta}_k) \mathbf{S}_k^{-1} & \mathbf{F}_{\text{prior},k-1} + \mathbf{J}_k^T(\boldsymbol{\theta}) \mathbf{S}_k^{-1} \mathbf{J}_k(\boldsymbol{\theta}_k) - \lambda_{\text{eig}} \mathbf{I}_N \end{bmatrix} \geq \mathbf{0}_{M+N}. \quad (4.18)$$

The constraint (4.18) also depends on the unknown parameter vector $\boldsymbol{\theta}_k$, which otherwise would be independent of $\boldsymbol{\theta}_k$ for additive linear Gaussian measurement models. For the dependent noise case, we also solve (4.10), but by replacing the size- N LMI in (4.10b) with the size- $M + N$ LMI (4.18). Furthermore, because of (4.18) the matrix $\mathbf{S}_k^{-1} + a_k^{-1} \text{diag}(\mathbf{w}_k)$ should be positive definite. This can be achieved, for example, by choosing a_k such that it satisfies the condition $0 < a_k < \lambda_{\min}\{\boldsymbol{\Sigma}_k\}$, since $w_{k,m} \geq 0$ for $m = 1, 2, \dots, M$.

4.8 Structured signals

In this section, we will discuss sparse sensing mechanisms for filtering problems involving structured signals (more generally, filtering with equality constraints on the state variables). In particular, we will illustrate sparse sensing for state sequences that are sparse in nature.

Suppose the time-varying vector of interest $\boldsymbol{\theta}_k \in \mathbb{R}^M$ at time k follows a linear model corrupted by additive noise:

$$\mathbf{x}_k = \mathbf{H}_k \boldsymbol{\theta}_k + \mathbf{n}_k, \quad (4.19)$$

where $\boldsymbol{\theta}_k \in \mathbb{R}^M$ (in contrast to the rest of this chapter, without loss of generality, here we assume $M = N$) has just a few nonzero coefficients, i.e., $\|\boldsymbol{\theta}_k\|_0 \ll M$ and \mathbf{H}_k denotes some known linear basis of size $M \times M$. The additive noise vector $\mathbf{n}_k \in \mathbb{R}^M$ is assumed to be zero-mean Gaussian with covariance matrix $\boldsymbol{\Sigma} = \sigma^2 \mathbf{I}_M$.

Under the assumption that the parameter vector $\boldsymbol{\theta}_k$ is sparse, CS theory asserts an exact recovery of $\boldsymbol{\theta}_k$ from observations which are typically much smaller than M , i.e, signals acquired via a linear compression matrix. We

now demonstrate that discrete sparse sensing mechanisms can be used for designing time-varying compression matrices as well as determining the optimal compression rate to reach a desired information gain or mean squared error. More specifically, we design $\Phi_k(\mathbf{w}_k) \in \{0, 1\}^{K \times M}$ to acquire the data as before, i.e.,

$$\mathbf{y}_k = \Phi_k(\mathbf{w}_k)\mathbf{x}_k. \quad (4.20)$$

For $K \ll M$, the sampling matrix will be a compression matrix, and the measurement vector will be much shorter than \mathbf{x}_k . Note that the sampling matrix Φ_k and sparsity pattern (including the sparsity order) of the vector $\boldsymbol{\theta}_k$ can both be time-varying.

One way to model the evolution of time-varying sparse sequences (more generally, states with structural constraints) is through the so-called pseudo-measurement formulation [Julier and LaViola, 2007], [Carmi et al., 2010], [Farahmand et al., 2014], where it is assumed that $\boldsymbol{\theta}_k$ evolves according to the following model

$$\text{dynamics: } \boldsymbol{\theta}_{k+1} = \mathbf{A}_k\boldsymbol{\theta}_k + \mathbf{u}_k; \quad (4.21a)$$

$$\text{pseudo-measurement: } 0 = g(\boldsymbol{\theta}_k) + e_k, \quad (4.21b)$$

where \mathbf{A}_k is an $M \times M$ state-transition matrix, $\mathbf{u}_k \in \mathbb{R}^M$ is the process noise, $g(\boldsymbol{\theta}_k)$ is a structure-constraining function, and e_k is (without loss of generality) zero-mean unit-variance noise. The basic idea behind such a formulation is to view the structural constraints as additional measurements whose outputs are always zero, and the pseudo-measurement noise variance specifies the amount of constraint violation. In this way, the state estimate (e.g., of the extended Kalman filter) conforms to the structural constraints as well.

To model sparse states, any one of the well-known approximations of the ℓ_0 -(-quasi) norm can be used for $g(\cdot)$, e.g., the ℓ_1 -norm function $(\gamma_k \|\boldsymbol{\theta}_k\|_1)^{1/2}$, the inverse Gaussian function $(\gamma_k \sum_{m=1}^M (1 - \exp(-\theta_{k,m}^2/2\sigma_g^2)))^{1/2}$ with shape parameter σ_g^2 , or the sum-of-logarithms function $(\gamma_k \sum_{m=1}^M \log(|\theta_{k,m}| + \delta))^{1/2}$ with $\delta > 0$. Here, γ_k is the tuning parameter and $\theta_{k,m}$ denotes the m th entry of $\boldsymbol{\theta}_k$. With the square root in each of these functions, the extra least-squares error term for the pseudo-measurement will be equivalent to the sparsity-inducing regularizer. Note that this formulation also accommodates a much richer class of structured signals as discussed later on in §4.8.3.

In the considered adaptive sparse sensing problem, we are basically replacing the random measurement operation traditionally used in the CS framework with a deterministic and structured sensing operation, which is more favorable for practical implementation. In what follows, we will develop risk functions that accommodate such structural constraints.

4.8.1 Risk function

The linear system (4.20) and (4.21a) can be solved using the celebrated Kalman filter algorithm [Kay, 1993]. We now recall Remark 4.1, where the posterior FIM (equal to the inverse of the posterior error covariance matrix) for the state-space equations (4.20) and (4.21a) without the pseudo-measurement was given by (4.7), that is,

$$\mathbf{F}_k(\mathbf{w}_k) = (\boldsymbol{\Sigma}_u + \mathbf{A}_k \mathbf{F}_{k-1}^{-1} \mathbf{A}_k^T)^{-1} + \frac{1}{\sigma^2} \sum_{m=1}^M w_{k,m} \mathbf{h}_{k,m} \mathbf{h}_{k,m}^T,$$

where $\mathbf{h}_{k,m} \in \mathbb{R}^{M \times 1}$, $m = 1, 2, \dots, M$, are the rows of \mathbf{H}_k .

Due to the compression ($K \ll M$), the conventional Kalman filter is less meaningful (especially for sparse recovery), unless the inherent sparsity of the state sequence is taken into account. We do this through an (independent) extra pseudo-measurement (4.21b). For the state-space equations (4.20) and (4.21), the additional measurement modifies the posterior FIM to

$$\begin{aligned} \mathbf{F}_k(\boldsymbol{\theta}_k, \mathbf{w}_k) &= (\boldsymbol{\Sigma}_u + \mathbf{A}_k \mathbf{F}_{k-1}^{-1} \mathbf{A}_k^T)^{-1} + \partial g(\boldsymbol{\theta}_k) \partial g(\boldsymbol{\theta}_k)^T \\ &\quad + \frac{1}{\sigma^2} \sum_{m=1}^M w_{k,m} \mathbf{h}_{k,m} \mathbf{h}_{k,m}^T, \end{aligned} \tag{4.22}$$

where $\partial g(\boldsymbol{\theta}) \in \mathbb{R}^M$ is the (sub)gradient of $g(\boldsymbol{\theta})$ towards $\boldsymbol{\theta}$ evaluated at $\boldsymbol{\theta}_k$. In fact, the above expression is the inverse of the posterior error covariance of the extended Kalman filter, when the true state $\boldsymbol{\theta}_k$ in (4.22) is replaced with the prediction $\widehat{\boldsymbol{\theta}}_{k|k-1} = \mathbf{A}_k \widehat{\boldsymbol{\theta}}_{k-1|k-1}$, where the past estimate can be computed using any of the favorite sparse recovery algorithms or using an iterative extended Kalman filter [Carmi et al., 2010, Farahmand et al., 2014, Vaswani, 2008].

As before, we could use either one of the constraints discussed in §4.3, i.e., use

$$f(\mathbf{w}_k) := \text{tr}\{(\mathbf{F}_k(\boldsymbol{\theta}_k, \mathbf{w}_k))^{-1}\} \leq \lambda_{\text{tr}}$$

or

$$f(\mathbf{w}_k) := \lambda_{\min}\{\mathbf{F}_k(\boldsymbol{\theta}_k, \mathbf{w}_k)\} \geq \lambda_{\text{eig}}$$

with $\mathbf{F}_k(\boldsymbol{\theta}_k, \mathbf{w}_k)$ in (4.22) in place of (4.12b) to solve (4.12).

4.8.2 Example: CS-based target tracking

In this section, we illustrate the developed theory with the following target tracking example. Let $\mathbf{p}_k = [p_{k,x}, p_{k,y}]^T \in \mathbb{R}^2$ denote the position of the target at time instance k and $\mathbf{a}_m \in \mathbb{R}^2$ denote the position of the m th sensor. Let us assume that there are M such locations where we can place these sensors. The sensors are capable of measuring the signal strength as

$$h_{k,m}(\mathbf{p}_k) = \frac{\beta s}{\beta + \|\mathbf{p}_k - \mathbf{a}_m\|_2^2}, m = 1, 2, \dots, M, \quad (4.23)$$

with a constant $\beta > 0$. Here, s denotes the signal strength. We linearize (4.23) around M grid points $\{\mathbf{g}_m\}_{m=1}^M$, where the target could be potentially located. As a result, we arrive at the linear grid-based model given by

$$x_{k,m} = \mathbf{h}_{k,m}^T \boldsymbol{\theta}_k + n_{k,m}, m = 1, 2, \dots, M,$$

where $\mathbf{h}_{k,m} = [h_{k,m}(\mathbf{g}_1), h_{k,m}(\mathbf{g}_2), \dots, h_{k,m}(\mathbf{g}_M)]^T \in \mathbb{R}^M$ is time-invariant, but $\boldsymbol{\theta}_k$ is time-varying. All the entries of the vector $\boldsymbol{\theta}_k \in \mathbb{R}^M$ are equal to zero except for the m th entry, $\theta_{k,m}$, which is equal to the target signal strength s at time k if and only if the target is located at the m th grid point, i.e., for $\mathbf{p}_k = \mathbf{g}_m$. Note that the number of grid points can be much larger than the number of (selected) sensors.

We are interested in tracking a target moving along the grid points $\mathbf{g}_m = [m, m]^T$ for $m = 0, 1, \dots, 29$ as shown in Figure 4.2(a). This can be modeled as

$$\boldsymbol{\theta}_{k+1} = \mathbf{A}\boldsymbol{\theta}_k + \mathbf{u}_k,$$

where the entries of the initial vector $\boldsymbol{\theta}_0$ are all zero except for the first entry $\theta_{1,1} = s$. Here, the state-transition matrix is a shift matrix, i.e.,

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 1 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 1 & 0 \end{bmatrix} \in \mathbb{R}^{M \times M}.$$

In scenarios like the one considered here, compression via random linear projections would still need all the M sensors with no reduction in the sensing and communications cost. On the contrary, sparse sensing enables a completely decentralized sensing, and it needs only $K \ll M$ sensors.

In this example, the sparsity pattern is time-varying, but the sparsity order is fixed. We stress here that the proposed framework is not limited to signals with a fixed sparsity order. We use the following parameters in the simulations: The number of grid points/candidate sensors $M = 30$, $K = 5$, $\beta = 100$, $s = 10$, and $\sigma = 10^{-3}$. The sensors are deployed uniformly at random within a 30×30 m² surveillance area as shown in Figure 4.2(a).

We use $f_k(\mathbf{w}_k) := \text{tr}\{(\mathbf{F}_k(\boldsymbol{\theta}_k, \mathbf{w}_k))^{-1}\}$ and λ_{tr} is chosen such that $N = 5$ sensors are selected at each time step. We evaluate $\mathbf{F}_k(\boldsymbol{\theta}_k, \mathbf{w}_k)$ at the prediction $\widehat{\boldsymbol{\theta}}_{k|k-1} = \mathbf{A}\widehat{\boldsymbol{\theta}}_{k-1|k-1}$, where the past estimate is obtained by solving the following ℓ_1 -norm regularized least squares:

$$\begin{aligned} \widehat{\boldsymbol{\theta}}_{k|k} = \arg \min_{\boldsymbol{\theta}_k} & \quad \|\widehat{\boldsymbol{\theta}}_{k|k-1} - \boldsymbol{\theta}_k\|_{\mathbf{P}_{k|k-1}^{-1}}^2 \\ & + \|\mathbf{y}_k - \boldsymbol{\Phi}_k(\mathbf{w}_k)\mathbf{H}\boldsymbol{\theta}_k\|_{\boldsymbol{\Sigma}^{-1}}^2 + g^2(\boldsymbol{\theta}_k), \end{aligned} \quad (4.24)$$

with $g(\boldsymbol{\theta}_k) = \sqrt{2\gamma_k\|\boldsymbol{\theta}_k\|_1}$ (here, $\|\boldsymbol{\theta}_k\|_1 = \mathbf{1}^T\boldsymbol{\theta}_k$ as the entries of $\boldsymbol{\theta}_k$ are all nonnegative), $\mathbf{P}_{k|k-1} = \mathbf{A}\mathbf{P}_{k-1|k-1}\mathbf{A}^T + \boldsymbol{\Sigma}_u$, and

$$\begin{aligned} \mathbf{P}_{k|k} = & \left(\mathbf{P}_{k|k-1}^{-1} + \frac{\gamma_k}{2\|\boldsymbol{\theta}_k\|_1} \mathbf{1}\mathbf{1}^T \right. \\ & \left. + \frac{1}{\sigma^2} \sum_{m=1}^M \mathbf{h}_{k,m} \mathbf{h}_{k,m}^T \right)^{-1}. \end{aligned}$$

Here, γ_k is a tuning parameter. We initialize $\widehat{\boldsymbol{\theta}}_{0|0} = \mathbf{1}_M$, $\mathbf{P}_{0|0} = \mathbf{I}_M$, $\boldsymbol{\Sigma}_u = 0.01\mathbf{I}_M$, and we compute γ_k using the method described in [Farahmand et al., 2014].

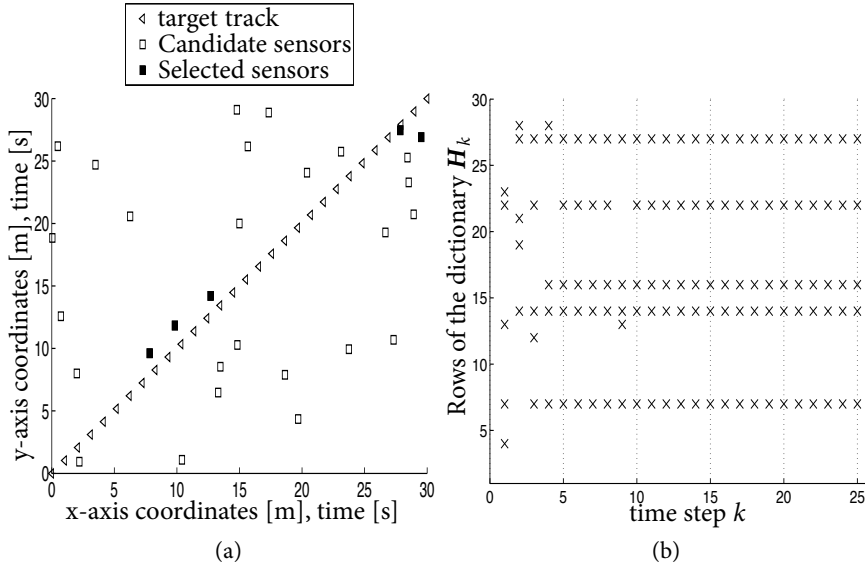


Figure 4.2: Tracking a target using a grid-based model with $M = 30$ and $K = 5$. (a) A target is moving along the straight line $p_{x,k} = p_{y,k} = k$, i.e., it moves with a constant velocity of $\sqrt{2}$ m/s. The selected sensors shown correspond to $k = 25$ s. (b) The solution path illustrating the selected rows of the dictionary \mathbf{H}_k for $k = 1, 2, \dots, 25$ s.

The relaxed optimization problem is solved using SeDuMi [Sturm, 1999]. Figure 4.2(b) illustrates the solution path for $k = 1, 2, \dots, 25$ s. The Boolean solution is recovered using deterministic rounding. The sensors selected for time step $k = 25$ s are also shown in Figure 4.2(a). In this example, the same subset of sensors are selected for $k > 10$ because the matrices \mathbf{H} and \mathbf{A} , and the sparsity order are not changing with time. In other words, the selected sensors optimize the average error, averaged over the entire track.

4.8.3 Extensions to other structured signals

In this section, we highlight some important generalizations of the proposed framework for structured signals, which are often studied together with the CS framework. The sparsity prior can be extended to a much broader class of structured signals, including structured sparse signals (or block-sparse sig-

nals) [Friedman et al., 2010], smooth (i.e., sparsity of the coefficients and also sparsity of their differences) [Tibshirani et al., 2005], to list a few. Depending on the structure of the state, the $g(\boldsymbol{\theta}_k)$ has to be modified accordingly. More specifically, for structured sparse signals we use a regularizer that accounts for block sparsity, i.e.,

$$g(\boldsymbol{\theta}_k) := \left(\gamma_k \sum_{i=1}^G \|\boldsymbol{\theta}_{k,i}\|_2 \right)^{1/2},$$

where the state vector $\boldsymbol{\theta}_k$ is grouped into G subvectors each of length N/G as

$$\boldsymbol{\theta}_k = [\boldsymbol{\theta}_{k,1}^T, \boldsymbol{\theta}_{k,2}^T, \dots, \boldsymbol{\theta}_{k,G}^T]^T.$$

Similarly, for signal smoothness, we use the regularizer

$$g(\boldsymbol{\theta}_k) := \left(\gamma_{k,1} \|\boldsymbol{\theta}_k\|_1 + \gamma_{k,2} \sum_{m=1}^{M-1} |\theta_{k,m} - \theta_{k,m-1}| \right)^{1/2},$$

where $\gamma_{k,1}$ and $\gamma_{k,2}$ are tuning parameters.

4.9 Discussion

In this chapter we have discussed discrete sparse sensing for state estimation in nonlinear dynamical systems. In particular, we have focused on filtering problems of three types: independent observations, additive Gaussian dependent observations, and structured state sequences. For designing sparse sensing patterns, we have used scalar measures related to the posterior Cramér-Rao bound, which depend on the previous as well as the current true states. The sensing patterns are designed one step ahead in time, hence can be computed offline.

4.A Extended Kalman filter

For the sake of completeness, we provide the steps involved in the Extended Kalman filter, which is used for state estimation with nonlinear additive Gaussian models; see [Kay, 1993] for more details.

Given the state-space equations

$$\begin{aligned}\mathbf{x}_k &= \mathbf{h}_k(\boldsymbol{\theta}_k) + \mathbf{n}_k, \\ \boldsymbol{\theta}_{k+1} &= \mathbf{A}_k \boldsymbol{\theta}_k + \mathbf{u}_k,\end{aligned}$$

where $\mathbf{h}_k(\cdot) : \mathbb{R}^N \rightarrow \mathbb{R}^M$ is a nonlinear function of the unknown state parameter vector $\boldsymbol{\theta}_k \in \mathbb{R}^N$, $\mathbf{A}_k \in \mathbb{R}^{N \times N}$ is the state-transition matrix, and $\mathbf{n}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ and $\mathbf{u}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_u)$ are the measurement and process noise vectors of length $M \times 1$, respectively.

Suppose that the estimate $\widehat{\boldsymbol{\theta}}_{k-1|k-1}$ and its covariance matrix $\mathbf{P}_{k-1|k-1}$ are available from the previous time step. At current time step k , the prediction and its covariance are given as

$$\begin{aligned}\widehat{\boldsymbol{\theta}}_{k|k-1} &= \mathbf{A}_k \widehat{\boldsymbol{\theta}}_{k-1|k-1}, \\ \mathbf{P}_{k|k-1} &= \mathbf{A}_k \mathbf{P}_{k-1|k-1} \mathbf{A}_k^T + \boldsymbol{\Sigma}_u.\end{aligned}$$

The update equations include the posterior estimate and posterior covariance matrices, which are computed using:

$$\begin{aligned}\widehat{\boldsymbol{\theta}}_{k|k} &= \widehat{\boldsymbol{\theta}}_{k|k-1} + \mathbf{P}_{k|k-1} \mathbf{H}_{k|k-1}^T [\mathbf{H}_{k|k-1} \mathbf{P}_{k|k-1} \mathbf{H}_{k|k-1}^T + \boldsymbol{\Sigma}]^{-1} [\mathbf{x}_k - \mathbf{h}_k(\widehat{\boldsymbol{\theta}}_{k|k-1})], \\ \mathbf{P}_{k|k} &= \mathbf{P}_{k|k-1} - \mathbf{P}_{k|k-1} \mathbf{H}_{k|k-1}^T (\mathbf{H}_{k|k-1} \mathbf{P}_{k|k-1} \mathbf{H}_{k|k-1}^T + \boldsymbol{\Sigma})^{-1} \mathbf{H}_{k|k-1} \mathbf{P}_{k|k-1},\end{aligned}$$

where $\mathbf{h}_k(\cdot)$ is linearized around the prediction to obtain

$$\mathbf{H}_{k|k-1} = \left. \frac{\partial \mathbf{h}_k(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T} \right|_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}_{k|k-1}} \in \mathbb{R}^{M \times N}.$$

Chapter 5

Sparse Sensing for Detection

Contents

5.1	Introduction	96
5.2	Sensing conditionally distributed observations	98
5.3	Problem statement	99
5.4	$f(\mathbf{w})$ for detection	100
5.5	Solvers	107
5.6	Illustrative examples	110
5.7	Dependent observations	117
5.8	Discussion	128
5.A	Proof of Proposition 5.1	128
5.B	Upper bound on P_m	129
5.C	Proof of Proposition 5.2	130
5.D	Expressions for $f_0(\mathbf{w})$ and $f_1(\mathbf{w})$	130

Part of this chapter was published as: S.P. Chepuri and G. Leus. Sparse Sensing for Distributed Detection. *IEEE Trans. on Signal Processing (To appear)*, Oct. 2015.

5.1 Introduction

In this chapter we will study the discrete sparse sensing framework for another important statistical inference problem, that is, detection. Statistical detection is pertinent to applications in sensor networks, radar and sonar systems, wireless cognitive radio networks, biometrics, social networks, imaging platforms, to list a few. We assume that the field is sampled by (spatially or temporally) distributed sensors, and these samples are delivered to a central unit. The central unit then makes a single global decision as to the true hypothesis using binary hypothesis testing. More specifically, the observations at each sensor are related to the state of nature \mathcal{H} , where the random variable \mathcal{H} is drawn from a binary alphabet set $\{\mathcal{H}_0, \mathcal{H}_1\}$. In the Bayesian setting, we assume that the prior probabilities $\pi_0 = \Pr(\mathcal{H}_0)$ and $\pi_1 = \Pr(\mathcal{H}_1)$ are known, whereas in the Neyman-Pearson setting, the prior probabilities are not known.

5.1.1 Related earlier works

The minimum error probability criterion is a standard performance measure for design problems related to statistical detection such as signal design [Grettenberg, 1963, Kadota and Shepp, 1967, Kailath, 1967], censoring [Rago et al., 1996], sampling design [Yu and Varshney, 1997], and so on. However, in most cases, optimizing the error probabilities is very difficult. This may be because these error probabilities do not admit a known closed form or their expression is too complicated for numerical optimization. Therefore, weaker performance criteria that are easier to evaluate and optimize are often used. A number of measures related to the distance between the conditional probabilities are widely used in the design of experiments as proxies for the error probability [Kailath, 1967, Rago et al., 1996, Yu and Varshney, 1997, Bajovic et al., 2011, Chamberland and Veeravalli, 2007]. Some of the prominent distance measures that are often used are the Kullback-Leibler distance, J-divergence, Chernoff information, and Bhattacharyya distance.

A related topic in the context of energy-efficient distributed detection is *data censoring*, wherein the uninformative sensor observations are not transmitted to the central unit [Rago et al., 1996, Appadwedula et al., 2008, Blum

and Sadler, 2008]. However, in censoring, data still has to be acquired in order to choose informative sensors, thus, it incurs a sensing cost. That is, censoring schemes are data dependent as opposed to the proposed data-independent sparse sensing schemes that can be designed offline. In other words, the actual measurements are not needed and only model information is used.

5.1.2 Main results

We focus on both the Bayesian as well as the Neyman-Pearson setting for binary hypothesis testing. The sparse sensing operation is designed based on a number of distance measures that belong to the general class of Ali-Silvey distances [Ali and Silvey, 1966].

The main question addressed in the chapter is similar to that of [Cambanis and Masry, 1983, Bahr and Bucklew, 1990, Yu and Varshney, 1997, Bajovic et al., 2011, Quan et al., 2009, Sung et al., 2005], but with the following differences. Firstly, the proposed framework is *general*, that is, it is not limited to Gaussian observations, especially for conditionally independent observations. Secondly, we propose a sparsity-promoting cost function to design structured samplers to achieve the lowest sensing cost as compared to the previously adopted periodic, regular, or random samplers. The main contributions of this chapter that broaden the existing literature are listed below.

- For conditionally independent observations, the best subset of sensors is the one with the *smallest local average root-likelihood ratio* and *largest local average log-likelihood ratio* in the Bayesian and Neyman-Pearson setting, respectively. This leads to an explicit solution for the sensing design problem that is optimal in terms of the error exponents. As a special case, for Gaussian observations with common covariances and uncommon means under both hypotheses, the selected sensors are also optimal in terms of the error probabilities (initial results for the Gaussian case were reported in [Chepuri and Leus, 2015a]). The computational complexity of the proposed solvers is independent of the number of candidate sensors, and is as low as $O(K)$, where K is the number of selected sensors (or sampling locations).

- For conditionally dependent observations, we focus on the Gaussian setting. When the mean vectors are uncommon and the covariance structure is common under both hypotheses, the sensing design problem can be relaxed to a convex optimization problem. Although this leads to a suboptimal solution, we propose a randomized rounding technique that further improves the solution. Moreover, in this case, for nonidentical sensor observations, we show that the number of sensors required to achieve a prescribed detection performance decreases significantly as the correlation among them increases (i.e., when the sensors become more coherent), which is in complete contrast to the case of identical sensor observations. When the covariances are uncommon and the means are common under both hypotheses, the sensing design problem remains nonconvex, except for the J-divergence optimization (this also holds for a more general case with uncommon means).

5.2 Sensing conditionally distributed observations

Consider a network with M candidate sensors. These candidate sensors might represent temporal, spatial, or even spatio-temporal samples. The observations are related to the following model

$$\mathcal{H}_0 : x_m \sim p_m(x|\mathcal{H}_0), m = 1, 2, \dots, M, \quad (5.1a)$$

$$\mathcal{H}_1 : x_m \sim p_m(x|\mathcal{H}_1), m = 1, 2, \dots, M, \quad (5.1b)$$

where the probability density function (pdf) of the observation at the m th sensor, x_m , conditioned on the state of nature \mathcal{H} is denoted by $p_m(x|\mathcal{H}_i)$ for $i = 0, 1$. Further, the observations are collected in $\mathbf{x} = [x_1, x_2, \dots, x_M]^T \in \mathbb{R}^M$. The pdf of \mathbf{x} under \mathcal{H}_0 and \mathcal{H}_1 is denoted by $p(\mathbf{x}|\mathcal{H}_0)$ and $p(\mathbf{x}|\mathcal{H}_1)$, respectively.

We acquire the data \mathbf{x} via a linear sensing operation, where the sensing task is modeled through a vector whose entries belong to a binary alphabet, i.e., through

$$\mathbf{w} = [w_1, w_2, \dots, w_M]^T \in \{0, 1\}^M,$$

where the variable $w_m = (0)1$ indicates whether the m th sensor is (not) selected. More specifically, we define the sensing matrix $\Phi(\mathbf{w}) = \text{diag}_r(\mathbf{w}) \in \{0, 1\}^{K \times M}$, to acquire the data as

$$\mathbf{y} = \text{diag}_r(\mathbf{w})\mathbf{x} = \Phi(\mathbf{w})\mathbf{x},$$

where K is not assumed to be known. Note that we are interested in cases where $K \ll M$. The reduced dimension data vector $\mathbf{y} \in \mathbb{R}^K$ is used instead of $\mathbf{x} \in \mathbb{R}^M$ to solve the detection problem. In this chapter we seek a sparsest \mathbf{w} , i.e., a vector with many zeros and just a few nonzero entries, such that a prescribed global detection performance is achieved.

5.3 Problem statement

Let $\widehat{\mathcal{H}}$ denote an estimate of the state of nature \mathcal{H} , based on a certain decision rule. In the Neyman-Pearson setting, the optimal detector minimizes the probability of miss detection (type II error),

$$P_m = \Pr(\widehat{\mathcal{H}} \neq \mathcal{H}_1 | \mathcal{H}_1)$$

for a fixed probability of false alarm (type I error),

$$P_f = \Pr(\widehat{\mathcal{H}} = \mathcal{H}_1 | \mathcal{H}_0).$$

This is the well-known Neyman-Pearson detector. In the Bayesian setting, given the prior probabilities $\pi_i = \Pr(\mathcal{H}_i)$ for $i = 0, 1$, the optimal detector minimizes the Bayesian error probability,

$$P_e = \Pr(\widehat{\mathcal{H}} \neq \mathcal{H}) = \pi_0 P_f + \pi_1 P_m,$$

or more generally, the detector minimizes the Bayes' risk. Having introduced the data model, we now formally state the design problem of interest.

Problem 5.1 (Sparse sampler design). *Given the data model (5.1), design a sparsest Boolean vector \mathbf{w} that results in a prescribed*

(i) *Bayesian probability of error, P_e , in the Bayesian setting, or*

- (ii) *probability of miss detection, P_m , for a fixed probability of false alarm, P_f , in the Neyman-Pearson setting.*

Mathematically, the sparse sensing problem for distributed detection can be formulated as

$$\begin{aligned} \text{P-B : } & \arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \\ & \text{s.to } P_e(\mathbf{w}) \leq e; \end{aligned} \quad (5.2a)$$

$$\begin{aligned} \text{P-N : } & \arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \\ & \text{s.to } P_f(\mathbf{w}) \leq \alpha, \text{ and } P_m(\mathbf{w}) \leq \beta, \end{aligned} \quad (5.2b)$$

where e , α and β are, respectively, the desired Bayesian probability of error, maximum false-alarm rate and maximum miss-detection rate. Here, $P_e(\mathbf{w})$, $P_f(\mathbf{w})$, and $P_m(\mathbf{w})$ denote the error probabilities due to the selected sensor subset indicated by the nonzero entries of \mathbf{w} . When prior probabilities are available, we solve P-B (P denotes problem and B denotes Bayesian), otherwise in the Neyman-Pearson setting we solve P-N (N denotes Neyman-Pearson).

In order to ease the design, we next discuss some performance measures that can substitute the error probabilities in the above optimization problems.

5.4 $f(\mathbf{w})$ for detection

The error probabilities P_e , P_m or P_f might not admit a known closed-form expression or their expressions might not be favorable for numerical optimization. In this section, we will discuss several weaker and simpler substitutes, which can be optimized instead of the error probabilities. These substitutes are based on the notion of distance (closeness or divergence) between the two distributions of the observations under test. They lead to tractable, if not always optimal (in terms of the error probabilities) design procedures for sampler design. Nevertheless, optimizing the distance measures improves the performance of any practical system.

Let the likelihood ratio of the two hypotheses under test be defined as

$$l(\mathbf{y}) = \frac{p(\mathbf{y}|\mathcal{H}_1)}{p(\mathbf{y}|\mathcal{H}_0)}.$$

In what follows, we consider a number of distance measures that belong to the general class of Ali-Silvey distances [Ali and Silvey, 1966], which are of the form

$$\psi(\mathbb{E}_{|\mathcal{H}_i} \{\phi[l(\mathbf{y})]\}),$$

where $\psi(\cdot)$ is an increasing real-valued function, $\phi[\cdot]$ is a continuous convex function on $(0, \infty)$, and the notation $\mathbb{E}_{|\mathcal{H}_i} \{\phi[l(\mathbf{y})]\}$ indicates that $\phi[l(\mathbf{y})]$ is averaged under the pdf $p(\mathbf{y}|\mathcal{H}_i)$ for either $i = 0$ or $i = 1$.

5.4.1 The Bayesian setting

The Bayes detector minimizes P_e , and makes a decision based on comparing the optimal statistic to a threshold:

$$\log l(\mathbf{y}) = \log \frac{p(\mathbf{y}|\mathcal{H}_1)}{p(\mathbf{y}|\mathcal{H}_0)} \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\gtrless}} \log \frac{\pi_0}{\pi_1}.$$

In the Bayesian setting, our goal is to choose the best subset of sensors that results in a prescribed Bayesian probability of error P_e . The best achievable exponent in the Bayesian probability of error is parameterized by the Chernoff information (sometimes also referred to as the Chernoff distance) [Cover and Thomas, 2012, Chernoff's theorem], and it is given by

$$\begin{aligned} \mathcal{C}(\mathcal{H}_1 \parallel \mathcal{H}_0) &= -\log \min_{0 \leq n \leq 1} \int [p(\mathbf{y}|\mathcal{H}_1)]^n [p(\mathbf{y}|\mathcal{H}_0)]^{1-n} d\mathbf{y} \\ &= -\log \min_{0 \leq n \leq 1} \mathbb{E}_{|\mathcal{H}_0} \{[l(\mathbf{y})]^n\}. \end{aligned} \quad (5.3)$$

Due to the involved minimization over n , the Chernoff information in (5.3) is difficult to optimize over \mathbf{w} . Therefore, we use a special case of the Chernoff information called the *Bhattacharyya distance* as the optimization criterion,

where the Bhattacharyya distance is obtained by fixing $n = 0.5$ in (5.3). The Bhattacharyya distance is given by

$$\mathcal{B}(\mathcal{H}_1\|\mathcal{H}_0) = -\log \rho, \quad (5.4)$$

where the *Bhattacharyya coefficient* [Kailath, 1967] or the *Hellinger integral* [Kadota and Shepp, 1967], ρ , is given by

$$\begin{aligned} \rho &= \int \sqrt{p(\mathbf{y}|\mathcal{H}_1)p(\mathbf{y}|\mathcal{H}_0)} d\mathbf{y} = \int p(\mathbf{y}|\mathcal{H}_0) \sqrt{\frac{p(\mathbf{y}|\mathcal{H}_1)}{p(\mathbf{y}|\mathcal{H}_0)}} d\mathbf{y} \\ &= \mathbb{E}_{|\mathcal{H}_0} \left\{ \sqrt{l(\mathbf{y})} \right\}. \end{aligned} \quad (5.5)$$

It is easy to verify from (5.5) that the Bhattacharyya distance is symmetric, which means $\mathcal{B}(\mathcal{H}_1\|\mathcal{H}_0) = \mathcal{B}(\mathcal{H}_0\|\mathcal{H}_1)$. More importantly, the upper and lower bounds for the Bayesian probability of error can be obtained using the Bhattacharyya coefficient. The bounds are given as follows [Kadota and Shepp, 1967, Appendix A], [Kailath, 1967]:

$$\frac{1}{2} \min(\pi_0, \pi_1) \rho^2 \leq P_e \leq \sqrt{\pi_0 \pi_1} \rho. \quad (5.6)$$

Therefore, in place of the Bayesian error probability, we minimize the Hellinger integral, or equivalently, maximize the Bhattacharyya distance.

Furthermore, when $\int [p(\mathbf{y}|\mathcal{H}_1)]^n [p(\mathbf{y}|\mathcal{H}_0)]^{1-n} d\mathbf{y}$ is symmetric in n and the observations are independent and identically distributed, the Bhattacharyya distance is exponentially the best [Kailath, 1967], i.e.,

$$P_e \stackrel{as.}{\approx} \exp(-\mathcal{B}(\mathcal{H}_1\|\mathcal{H}_0)) \quad \text{for } P_e \rightarrow 0.$$

We now introduce the following assumption:

Assumption 5.1 (Conditional independence). *The sensor observations are statistically independent, conditioned on the hypothesis \mathcal{H} .*

Under Assumption 5.1, the likelihood ratio simplifies to

$$l(\mathbf{y}) = \frac{p(\mathbf{y}|\mathcal{H}_1)}{p(\mathbf{y}|\mathcal{H}_0)} = \prod_{m=1}^M [l_m(x)]^{w_m}$$

where $l_m(x) = p_m(x|\mathcal{H}_1)/p_m(x|\mathcal{H}_0)$ is the local likelihood ratio related to the m th sensor, and $p_m(x|\mathcal{H}_i)$ for $i = 0, 1$ are the conditional pdfs of x for the m th sensor. Here, the conditional pdf of the selected sensors is of reduced dimension, i.e., it does not include the measurements that are set to zero.

Besides being a reasonable measure, the Bhattacharyya distance is much simpler to optimize under Assumption 5.1 because of the following result:

Proposition 5.1 (Linearity of the Bhattacharyya distance). *The considered sparse sensing mechanism preserves the additivity of the Bhattacharyya distance under Assumption 5.1, i.e., we can express*

$$f_B(\mathbf{w}) := \mathcal{B}(\mathcal{H}_1\|\mathcal{H}_0) = \sum_{m=1}^M w_m \mathcal{B}_m(\mathcal{H}_1\|\mathcal{H}_0), \quad (5.7)$$

where

$$\mathcal{B}_m(\mathcal{H}_1\|\mathcal{H}_0) = -\log \mathbb{E}_{|\mathcal{H}_0} \{\sqrt{l_m(x)}\}. \quad (5.8)$$

Proof. See Appendix 5.A □

Thus, Proposition 5.1 enables us to optimize $f_B(\mathbf{w})$ over \mathbf{w} (subscript B denotes Bayesian). We underline here that $f_B(\mathbf{w})$ assumes only the knowledge of the data model and does not need actual measurements, hence the sensing operation can be designed offline. We also remark that the Chernoff information (5.3) is not additive for conditionally independent observations, unlike the Bhattacharyya distance.

Before discussing the optimization criterion for the Neyman-Pearson setting, we end this subsection with the following remark that generalizes the sampling design in the Bayesian setting.

Remark 5.1 (Bayes risk). *Let C_{ij} be the cost if we decide \mathcal{H}_i when \mathcal{H}_j is true. A generalization of the minimum P_e detector, is to minimize the Bayes risk*

$$\mathcal{R} = \sum_{i=0}^1 \sum_{j=0}^1 C_{ij} \Pr(\mathcal{H}_i|\mathcal{H}_j) \Pr(\mathcal{H}_j),$$

where we arrive at a special case of $\mathcal{R} = P_e$ for $C_{00} = C_{11} = 0, C_{10} = C_{01} = 1$. This results in the sensing design problem

$$\arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \quad \text{s.to} \quad \mathcal{R}(\mathbf{w}) \leq e_r,$$

where $\mathcal{R}(\mathbf{w})$ denotes the Bayes risk due to the selected sensor subset indicated by the nonzero entries of \mathbf{w} , and e_r is the desired Bayes risk.

The bounds in (5.6) can be generalized to [Kobayashi and Thomas, 1967]

$$\mathcal{R}_0 + \mathcal{R}_2 \rho^2 \leq \mathcal{R} \leq \mathcal{R}_0 + \sqrt{\mathcal{R}_1} \rho,$$

where $\mathcal{R}_0 = \pi_0 C_{00} + \pi_1 C_{11}$, $\mathcal{R}_1 = \pi_0 \pi_1 (C_{11} - C_{01})(C_{00} - C_{10})$, and $\mathcal{R}_2 = \mathcal{R}_1 / (\pi_0 (C_{00} - C_{10}) + \pi_1 (C_{11} - C_{01}))$. Therefore, maximizing the Bhattacharyya distance (or minimizing the Hellinger integral) is a reasonable optimality criterion also for a more general minimum Bayes risk detector.

5.4.2 The Neyman-Pearson setting

When the prior probabilities are not known, we solve the Neyman-Pearson problem, where one of the error probabilities (P_f , for example) is fixed while the second error probability, P_m is minimized. More specifically, the decision is based upon the log-likelihood ratio test

$$\log l(\mathbf{y}) = \log \frac{p(\mathbf{y}|\mathcal{H}_1)}{p(\mathbf{y}|\mathcal{H}_0)} \underset{\mathcal{H}_1}{\overset{\mathcal{H}_0}{\leq}} \gamma, \quad (5.9)$$

where γ is the threshold obtained by setting $P_f = \alpha$. In what follows, we discuss two distance measures that we can optimize in the Neyman-Pearson setting.

Kullback-Leibler distance

For a Neyman-Pearson problem, the best achievable error exponent in the probability of error (P_m , for example) is given by the relative entropy or *Kullback-Leibler distance* $\mathcal{K}(\mathcal{H}_1 \|\mathcal{H}_0)$ [Cover and Thomas, 2012, Stein's lemma]. That is, for a fixed value of P_f ,

$$\log P_m \stackrel{as.}{\approx} -\mathcal{K}(\mathcal{H}_1 \|\mathcal{H}_0) \quad \text{for } P_m \rightarrow 0.$$

The Kullback-Leibler distance is the average log-likelihood ratio, and is given by [Kullback, 1959]

$$\begin{aligned} \mathcal{K}(\mathcal{H}_1 \|\mathcal{H}_0) &= \mathbb{E}_{|\mathcal{H}_1} \{\log l(\mathbf{y})\} \\ &= \int \log l(\mathbf{y}) p(\mathbf{y}|\mathcal{H}_1) d\mathbf{y}. \end{aligned} \quad (5.10)$$

A lower bound on P_m for a fixed P_f , say α ($0 \leq \alpha \leq 1$) can be obtained using [Kullback, 1959, pp. 74-75 and tables in pp. 378-379]

$$\begin{aligned} \mathcal{K}(\mathcal{H}_1 \parallel \mathcal{H}_0) &\geq \alpha \log \left(\frac{\alpha}{1 - P_m} \right) \\ &+ (1 - \alpha) \log \left(\frac{1 - \alpha}{P_m} \right) = g(P_m). \end{aligned} \quad (5.11)$$

Since $g(P_m)$ is a strictly monotonic function (for values of α that are of practical interest), we can write

$$P_m \geq g^{-1}(\mathcal{K}(\mathcal{H}_1 \parallel \mathcal{H}_0)). \quad (5.12)$$

For example, a very small (close to zero) α simplifies (5.12) to $P_m \geq \exp(-\mathcal{K}(\mathcal{H}_1 \parallel \mathcal{H}_0))$. The following theorem gives an upper bound on P_m .

Theorem 5.1 (Upper bound on P_m). *If the variance of the log-likelihood ratio is v^2 , then*

$$P_m \leq \frac{1}{1 + \frac{(\mathcal{K}(\mathcal{H}_1 \parallel \mathcal{H}_0) - \log \gamma)^2}{v^2}}, \quad (5.13)$$

where the threshold γ corresponds to a desired $P_f = \alpha$.

Proof. See Appendix 5.B. □

The bounds in (5.12) and (5.13) make the maximization of $\mathcal{K}(\mathcal{H}_1 \parallel \mathcal{H}_0)$ a reasonable optimality criterion. We stress here that the above bounds (5.12) and (5.13) are valid even when Assumption 5.1 is not true.

The following property of the Kullback-Leibler distance further allows its easy numerical optimization.

Proposition 5.2 (Linearity of the Kullback-Leibler distance). *The considered sparse sensing mechanism preserves the additivity of the Kullback-Leibler distance under Assumption 5.1, i.e., we can express*

$$f_{N,1}(\mathbf{w}) := \mathcal{K}(\mathcal{H}_1 \parallel \mathcal{H}_0) = \sum_{m=1}^M w_m \mathcal{K}_m(\mathcal{H}_1 \parallel \mathcal{H}_0), \quad (5.14)$$

where

$$\begin{aligned}\mathcal{K}_m(\mathcal{H}_1\|\mathcal{H}_0) &= \int \log l_m(x) p_m(x|\mathcal{H}_1) dx \\ &= \mathbb{E}_{|\mathcal{H}_1} \{\log l_m(x)\}\end{aligned}\quad (5.15)$$

with $l_m(x) = p_m(x|\mathcal{H}_1)/p_m(x|\mathcal{H}_0)$ being the local likelihood ratio that was defined earlier.

Proof. See Appendix 5.C □

Therefore, Proposition 5.2 allows us to maximize $f_{N,1}(\mathbf{w})$ over \mathbf{w} (subscript N denotes Neyman-Pearson).

Remark 5.2. For the problem that minimizes the probability of false alarm P_f for a fixed probability of miss detection P_m , the Kullback-Leibler distance

$$\begin{aligned}\mathcal{K}(\mathcal{H}_0\|\mathcal{H}_1) &= -\mathbb{E}_{|\mathcal{H}_0} \{\log l(\mathbf{y})\} \\ &= -\int \log l(\mathbf{y}) p(\mathbf{y}|\mathcal{H}_0) d\mathbf{y}\end{aligned}\quad (5.16)$$

has to be optimized. Note that the Kullback-Leibler distance is not symmetric, i.e., $\mathcal{K}(\mathcal{H}_0\|\mathcal{H}_1) \neq \mathcal{K}(\mathcal{H}_1\|\mathcal{H}_0)$. Furthermore, Proposition 5.2 holds with the 0 and 1 subscripts interchanged in equations (5.14) and (5.15), which leads to the cost function

$$f_{N,2}(\mathbf{w}) := \sum_{m=1}^M w_m \mathcal{K}_m(\mathcal{H}_0\|\mathcal{H}_1). \quad (5.17)$$

J-divergence

The symmetric form of the Kullback-Leibler distance, *J-divergence*, is another frequently used criterion in the design of experiments. The J-divergence is defined as

$$\mathcal{D}(\mathcal{H}_1\|\mathcal{H}_0) = \mathcal{K}(\mathcal{H}_1\|\mathcal{H}_0) + \mathcal{K}(\mathcal{H}_0\|\mathcal{H}_1). \quad (5.18)$$

A lower bound on $(P_f + P_m)/2$ can be obtained using [Kullback, 1959]

$$\begin{aligned}\mathcal{D}(\mathcal{H}_1\|\mathcal{H}_0) &\geq 2 \left[\frac{P_f + P_m}{2} \log \left(\frac{(P_f + P_m)/2}{1 - (P_f + P_m)/2} \right) \right. \\ &\quad \left. + \left(1 - \frac{P_f + P_m}{2} \right) \log \left(\frac{1 - (P_f + P_m)/2}{(P_f + P_m)/2} \right) \right],\end{aligned}$$

and the results from Theorem 5.1 can be generalized to arrive at an upper bound.

Remark 5.3. *The J-divergence is also a reasonable measure in the Bayesian setting with $\pi_0 = 0.5$ as the Bayesian error probability $P_e = (P_f + P_m)/2$ can be both upper and lower bounded by $\mathcal{D}(\mathcal{H}_1\|\mathcal{H}_0)$. However, for other prior probabilities an upper bound on P_e can be obtained in terms of $\mathcal{D}(\mathcal{H}_1\|\mathcal{H}_0)$ only for Gaussian observations [Kadota and Shepp, 1967].*

The J-divergence is also additive for conditionally independent observations, i.e.,

$$f_{N,3}(\mathbf{w}) := \mathcal{D}(\mathcal{H}_1\|\mathcal{H}_0) = \sum_{m=1}^M w_m \mathcal{D}_m(\mathcal{H}_1\|\mathcal{H}_0),$$

where

$$\mathcal{D}_m(\mathcal{H}_1\|\mathcal{H}_0) = \mathcal{K}_m(\mathcal{H}_1\|\mathcal{H}_0) + \mathcal{K}_m(\mathcal{H}_0\|\mathcal{H}_1). \quad (5.19)$$

The additive property of the J-divergence is straightforward to verify, and it follows directly from Proposition 5.2.

Note that all the distance measures introduced in this section admit a closed-form expression irrespective of the observation distributions. The solvers for designing the sensing operation based on the developed performance measures are presented next.

5.5 Solvers

The performance measures derived in §5.4 greatly simplify the sensing design problems P-B and P-N, which are otherwise difficult to solve. The simplified problem is stated as follows.

Problem 5.2 (Simplified sparse sensing design). *Under Assumption 5.1, given M candidate sensors characterized by the conditional pdfs $\{p_m(x|\mathcal{H}_i)\}_{m=1}^M$ for $i = 0, 1$, design a sparsest vector \mathbf{w} such that a desired*

- (i) *Bhattacharyya distance in the Bayesian setting, or*

(ii) *Kullback-Leibler distance (or J-divergence) in the Neyman-Pearson setting,*

is achieved.

These sampling design problems are, respectively, expressed as the following cardinality minimization problems

$$\text{S-B : } \arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \quad \text{s.to} \quad f_{\text{B}}(\mathbf{w}) \geq \lambda_{\text{B}}; \quad (5.20\text{a})$$

$$\text{S-N : } \arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \quad \text{s.to} \quad f_{\text{N}}(\mathbf{w}) \geq \lambda_{\text{N}}, \quad (5.20\text{b})$$

where λ_{B} and λ_{N} specify the required Bhattacharyya distance and Kullback-Leibler distance (or J-divergence), respectively. The optimization problems S-B and S-N (S denotes simplified problem) are Boolean linear programming problems. In place of $f_{\text{N}}(\mathbf{w})$ in (5.20b), either one of the three performance measures $f_{\text{N},1}(\mathbf{w})$, $f_{\text{N},2}(\mathbf{w})$, or $f_{\text{N},3}(\mathbf{w})$ can be used; however, there is no general answer to the question of how does one performance metric compare with the other. These problems are of the form (P0) introduced in Chapter 2.

For the sake of brevity, we collect the local distances $\{\mathcal{B}_m(\mathcal{H}_1\|\mathcal{H}_0)\}_{m=1}^M$, $\{\mathcal{K}_m(\mathcal{H}_1\|\mathcal{H}_0)\}_{m=1}^M$, $\{\mathcal{K}_m(\mathcal{H}_0\|\mathcal{H}_1)\}_{m=1}^M$, or $\{\mathcal{D}_m(\mathcal{H}_1\|\mathcal{H}_0)\}_{m=1}^M$ in a common vector denoted by $\mathbf{d} \in \mathbb{R}^M$. The optimization problems in (5.20) can then be expressed in a general form as

$$\arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \quad \text{s.to} \quad \mathbf{d}^T \mathbf{w} \geq \lambda, \quad (5.21)$$

where the threshold corresponds to $\lambda := \lambda_{\text{B}}$ or $\lambda := \lambda_{\text{N}}$ for the Bayesian or Neyman-Pearson setting, respectively, with $0 \leq \lambda \leq \mathbf{1}^T \mathbf{d}$. Boolean linear programming problems are in general hard to solve. However, S-B and S-N are some of the few special cases of a Boolean linear program that have an explicit solution. We give the solution to the considered offline sampling design problem in the following theorem.

Theorem 5.2 (Sparse sampler for distributed detection). *Assuming the entries of \mathbf{d} are (pre-)sorted in descending order and the entries of \mathbf{w} are sorted accordingly. The optimal solution \mathbf{w} to (5.21) has entries equal to 1 at the first \widehat{K} entries*

corresponding to the largest entries in \mathbf{d} , where

$$\widehat{K} = \min\{i \in \{1, 2, \dots, M\} | d_1 + d_2 + \dots + d_i \geq \lambda\}. \quad (5.22)$$

Proof. The proof is straightforward, thus, not detailed. \square

In essence, the integer program (5.21) has an explicit solution and it is optimal for (5.21). The solution can be interpreted as follows: recalling \widehat{K} from (5.22), the best subset of sensors out of the M candidate sensors are those \widehat{K} sensors having the smallest local average root-likelihood ratio and largest local average log-likelihood ratio in the Bayesian and Neyman-Pearson setting, respectively. This result on ordering the sensors based on the local Bhattacharyya distance also holds for a general Bayes detector that minimizes the Bayes risk [cf. Remark 5.1].

The appeal of the proposed solution lies in its simplicity. Computationally, the proposed solver is very attractive, for example, with a complexity of $\mathcal{O}(M \log M)$, which is essentially the complexity of the involved sorting algorithm [Papadimitriou, 2003]. A parallel implementation on different processors (i.e., still in an offline centralized setting) of the ordering algorithm further reduces the complexity to $\mathcal{O}(\widehat{K})$ using a back-off mechanism as detailed next: The distance measure d_m is made available to the central unit after a time c/d_m , where c is a known positive constant, and the central unit computes the sum of the received values. If the accumulated sum exceeds the desired threshold λ , the central unit declares a transmission stop¹. Thus, only the \widehat{K} largest distance values are gathered at the central unit.

In many applications, we might know the number of sensors to select (e.g., we might have already purchased the hardware and we want to use all of them). That is, for a fixed sampler size K , the sensing design problem can

¹If more than one distance is made available at the same time, we randomly pick as many as we need.

be expressed as

$$\text{E-B : } \arg \min_{\mathbf{w} \in \{0,1\}^M} P_e(\mathbf{w}) \quad \text{s.to} \quad \|\mathbf{w}\|_0 = K; \quad (5.23a)$$

$$\begin{aligned} \text{E-N : } \arg \min_{\mathbf{w} \in \{0,1\}^M} P_m(\mathbf{w}) & \quad \arg \min_{\mathbf{w} \in \{0,1\}^M} P_f(\mathbf{w}) \\ \text{s.to } P_f(\mathbf{w}) \leq \alpha, & \quad \text{or} \quad \text{s.to } P_m(\mathbf{w}) \leq \beta, \\ \|\mathbf{w}\|_0 = K, & \quad \|\mathbf{w}\|_0 = K, \end{aligned} \quad (5.23b)$$

where E-B (E-N) represents the equivalent Bayesian (equivalent Neyman-Pearson) problem, and α and β are, respectively, the maximum false-alarm rate and maximum miss-detection rate to be satisfied. By appropriately choosing the thresholds e , α and β in (5.2), we can obtain the optimal objective value of (5.2) equal to K , for which P-B (P-N) and E-B (E-N) are equivalent.

We can also simplify E-B and E-N using the Bhattacharyya and Kullback-Leibler distance (or J-divergence) as proxies for the error probabilities, respectively, to arrive at a general form given by

$$\arg \max_{\mathbf{w} \in \{0,1\}^M} \mathbf{d}^T \mathbf{w} \quad \text{s.to} \quad \|\mathbf{w}\|_0 = K, \quad (5.24)$$

where it is straightforward to verify that the optimal objective value is given by the sum of the K largest entries of \mathbf{d} .

We underline that the proposed solver is valid as long as Assumption 5.1 holds, and the observations need not necessarily be Gaussian distributed.

5.6 Illustrative examples

In this section, we illustrate the developed theory of offline sampling design for binary hypothesis testing with a few examples. The sensing operation is designed such that a desired detection performance determined by the Bhattacharyya distance, Kullback-Leibler distance, or J-divergence is achieved. We begin with some examples of Gaussian observations and later on extend it to exponential observation distributions.

5.6.1 Gaussian observations

Uncommon means and common covariances

Detecting signals in Gaussian noise is a well-studied problem in detection theory. In particular, it finds applications in spectrum sensing, target detection, and communications, to list a few. For binary signals in Gaussian noise, that is, observations with uncommon means and common covariance structure under both hypotheses, the conditional distributions are given by

$$\begin{aligned}\mathcal{H}_0: \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\theta}_0, \boldsymbol{\Sigma}) \\ \mathcal{H}_1: \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\theta}_1, \boldsymbol{\Sigma}),\end{aligned}\tag{5.25}$$

where $\mathcal{N}(\boldsymbol{\theta}, \boldsymbol{\Sigma})$ denotes a Gaussian distribution with mean vector $\boldsymbol{\theta}$ and covariance matrix $\boldsymbol{\Sigma}$, the mean vectors $\boldsymbol{\theta}_i = [\theta_{i,1}, \theta_{i,2}, \dots, \theta_{i,M}]^T \in \mathbb{R}^M$ for $i = 0, 1$ as well as the covariance matrix $\boldsymbol{\Sigma} = \text{diag}(\sigma_1^2, \sigma_2^2, \dots, \sigma_M^2) \in \mathbb{R}^{M \times M}$ are assumed to be perfectly known. The error probabilities admit the following expressions [Moon and Stirling, 2000, pg. 475]

$$\begin{aligned}P_f(\mathbf{w}) &= \mathcal{Q}\left(\frac{\gamma + s(\mathbf{w})/2}{\sqrt{s(\mathbf{w})}}\right); \\ P_m(\mathbf{w}) &= 1 - \mathcal{Q}\left(\frac{\gamma - s(\mathbf{w})/2}{\sqrt{s(\mathbf{w})}}\right),\end{aligned}\tag{5.26}$$

where γ is the threshold defined in (5.9),

$$s(\mathbf{w}) = (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0)^T \text{diag}(\mathbf{w}) \boldsymbol{\Sigma}^{-1} (\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0)\tag{5.27}$$

is the *signal-to-noise ratio* (sometimes referred to as the deflection coefficient), and \mathcal{Q} is the complementary Gaussian cumulative distribution function

$$\mathcal{Q}(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} \exp(-y^2/2) dy.$$

Note that the signal-to-noise ratio (5.27) is also linear in \mathbf{w} . The Bayesian error probability is given by [Moon and Stirling, 2000, pg. 494]

$$P_e(\mathbf{w}) = \pi_0 \mathcal{Q}\left(\frac{\gamma' + s(\mathbf{w})/2}{\sqrt{s(\mathbf{w})}}\right) + \pi_1 \left[1 - \mathcal{Q}\left(\frac{\gamma' - s(\mathbf{w})/2}{\sqrt{s(\mathbf{w})}}\right)\right], \quad (5.28)$$

where $\gamma' = \log(\pi_0/\pi_1)$ is the threshold in the Bayesian setting.

For the detection problem (5.25), the local Bhattacharyya distance, Kullback-Leibler distance, and J-Divergence can be computed respectively as

$$\begin{aligned} \mathcal{B}_m(\mathcal{H}_1\|\mathcal{H}_0) &= \frac{1}{8\sigma_m^2}(\theta_{0,m} - \theta_{1,m})^2, \\ \mathcal{K}_m(\mathcal{H}_1\|\mathcal{H}_0) &= \mathcal{K}_m(\mathcal{H}_0\|\mathcal{H}_1) = \frac{1}{2\sigma_m^2}(\theta_{0,m} - \theta_{1,m})^2, \\ \mathcal{D}_m(\mathcal{H}_1\|\mathcal{H}_0) &= \frac{1}{\sigma_m^2}(\theta_{0,m} - \theta_{1,m})^2. \end{aligned}$$

We next remark the following interesting observation. All the three distance measures are equal to the signal-to-noise ratio up to a constant. That is, $\mathcal{B}(\mathcal{H}_1\|\mathcal{H}_0) = s(\mathbf{w})/8$, $\mathcal{K}(\mathcal{H}_1\|\mathcal{H}_0) = s(\mathbf{w})/2$, and $\mathcal{D}(\mathcal{H}_1\|\mathcal{H}_0) = s(\mathbf{w})$. However, these relations are not universal (e.g., they do not hold for non-Gaussian observations). This fact allows us to state the following fundamental result in sampling design for Gaussian observations with common covariance.

Theorem 5.3. *For Gaussian observations with uncommon means and common covariance structure under both hypotheses, maximizing the signal-to-noise ratio over all the possible sampler choices is optimal for P-B and P-N.*

Proof. The proof is straightforward. It can be derived based on results from [Cambanis and Masry, 1983] and the monotonicity of the \mathcal{Q} function. Thus, it is omitted. \square

As an example, consider the sinusoidal detection problem with $M = 15$ candidate sensors. The means are $\theta_{0,m} = 0$ and $\theta_{1,m} = \cos 2\pi f m$ with $f := 0.33$

for $m = 1, 2, \dots, M$. Furthermore, we use $\Sigma = \mathbf{I}_M$, $\pi_0 = 0.3$, $\pi_1 = 0.7$, and $\alpha = 0.01$. In this example, we use a smaller dimension for M to compare the results with the optimal solution of (5.23). Nevertheless, the proposed solvers based on ordering easily scale to higher dimensional problems. We solve (5.23) using exhaustive search over all the $\binom{M}{K}$ combinations for different values of K such that the error probabilities (5.26) and (5.28) are optimized. This is labelled as “Neyman-Pearson/Bayesian optimal” in Figure. 5.1. For this particular example, due to Theorem 5.3, the simplified sensing design problem can be solved optimally also in terms of error probabilities. This is evident from Figure. 5.1, where the solution based on ordering the distance measures (labelled as “Neyman-Pearson/Bayesian simplified, sorting”) is on top of the optimal solution obtained from exhaustive search. The shaded regions in Figure. 5.1 indicate the error probabilities with the worst to best subset of K sensors (including any possible subset of K sensors) for different numbers of selected sensors. In particular, the error probabilities with *random sampling* (or any other sub-optimal sampling), for example, [Cambanis and Masry, 1983, Sung et al., 2005], would span the shaded region.

Remark 5.4 (Choosing λ). For a desired P_m , say β , and fixed P_f , say α , the threshold $\lambda := \lambda_N$ (for a desired signal-to-noise ratio) can be computed using (5.26). Specifically,

$$\lambda_N = \left(\mathcal{Q}^{-1}(\alpha) - \mathcal{Q}^{-1}(1 - \beta) \right)^2.$$

When λ does not admit a closed form (e.g., with other distributions), the solution path can be used as a guideline to choose λ that results in a desired error probability (often needs to be computed numerically); for example, see Figure. 5.2 to compute $\lambda := \lambda_B$, where we solve (5.21) with the same simulation parameters as before.

Uncommon covariances and common means

Detecting a change in variance is also frequently encountered in practice, for example, while measuring a physical phenomenon with different sensors each characterized with different noise levels both across the sensors and under

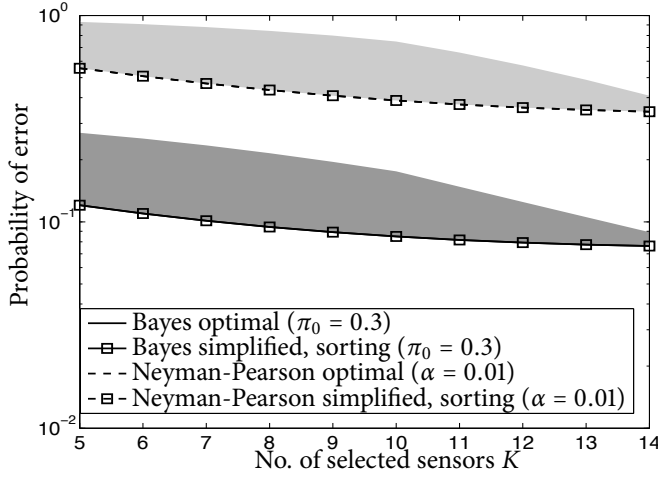


Figure 5.1: The probability of error (i.e., P_e or P_m) for (5.25) with different numbers of selected sensors K out of $M = 15$ sensors for independent observations. The shaded regions indicate the performance with the worst to best subset of K sensors.

both hypotheses. The conditional distributions in this case are given by

$$\begin{aligned} \mathcal{H}_0 : \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\theta}, \boldsymbol{\Sigma}_0) \\ \mathcal{H}_1 : \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\theta}, \boldsymbol{\Sigma}_1), \end{aligned} \quad (5.29)$$

where $\boldsymbol{\theta}$ is the known mean vector and $\boldsymbol{\Sigma}_i = \text{diag}(\sigma_{i,1}^2, \sigma_{i,2}^2, \dots, \sigma_{i,M}^2)$ for $i = 0, 1$ is the known diagonal covariance matrix. The local log-likelihood ratio is

$$\log l_m(x) = \frac{1}{2} \log \frac{\sigma_{0,m}^2}{\sigma_{1,m}^2} + x^2 \left(\frac{1}{2\sigma_{0,m}^2} - \frac{1}{2\sigma_{1,m}^2} \right).$$

Quantifying the performance of the detector, i.e., expressing P_m , P_f , and P_e in a known closed form is more difficult than before, as the pdf of $l(\mathbf{x})$ can be obtained only by numerical integration [Moon and Stirling, 2000]. However, the proposed performance measures admit known expressions as given next. The local Bhattacharyya distance between the conditional distributions

in (5.29) is given by

$$\mathcal{B}_m(\mathcal{H}_1\|\mathcal{H}_0) = \frac{1}{2} \log \left(\frac{\sigma_{0,m}^2 + \sigma_{1,m}^2}{2\sigma_{0,m}\sigma_{1,m}} \right), \quad (5.30)$$

the local Kullback-Leibler distance is given by

$$\mathcal{K}_m(\mathcal{H}_1\|\mathcal{H}_0) = \frac{1}{2} \left(\frac{\sigma_{1,m}^2}{\sigma_{0,m}^2} - 1 - \log \frac{\sigma_{1,m}^2}{\sigma_{0,m}^2} \right), \quad (5.31)$$

and $\mathcal{K}_m(\mathcal{H}_0\|\mathcal{H}_1)$ is obtained by interchanging the subscripts 0 and 1 in the above equation. Finally, the J-divergence is given by

$$\mathcal{D}_m(\mathcal{H}_1\|\mathcal{H}_0) = \frac{1}{2} \left(\frac{\sigma_{1,m}^2}{\sigma_{0,m}^2} + \frac{\sigma_{0,m}^2}{\sigma_{1,m}^2} - 2 \right). \quad (5.32)$$

Assume that

$$\Sigma_0 = \begin{bmatrix} 0.1 & 0 \\ 0 & 0.01 \end{bmatrix} \quad \text{and} \quad \Sigma_1 = \begin{bmatrix} 0.5 & 0 \\ 0 & 0.25 \end{bmatrix}$$

and that we want to find the best sensor out of $M = 2$ candidate sensors ($K = 1$). A quick calculation shows that $d_2 > d_1$ for all distances (i.e., the local distance measure of the second sensor is larger than that of the first sensor). Thus, the solution to the S-B (and S-N) will be $\mathbf{w} = [0, 1]^T$. This is intuitive as the conditional variance of the second sensor has a larger gap as compared to that of the first sensor, hence the second sensor is more informative.

5.6.2 Exponential observations

Exponentially distributed observations occur while detecting a complex Gaussian signal at the output of a noncoherent receiver. The conditional distributions for exponentially distributed observations for $m = 1, 2, \dots, M$, are given by

$$\begin{aligned} \mathcal{H}_0 : x_m &\sim \mu_{0,m} \exp(-\mu_{0,m}x) \\ \mathcal{H}_1 : x_m &\sim \mu_{1,m} \exp(-\mu_{1,m}x), \end{aligned} \quad (5.33)$$

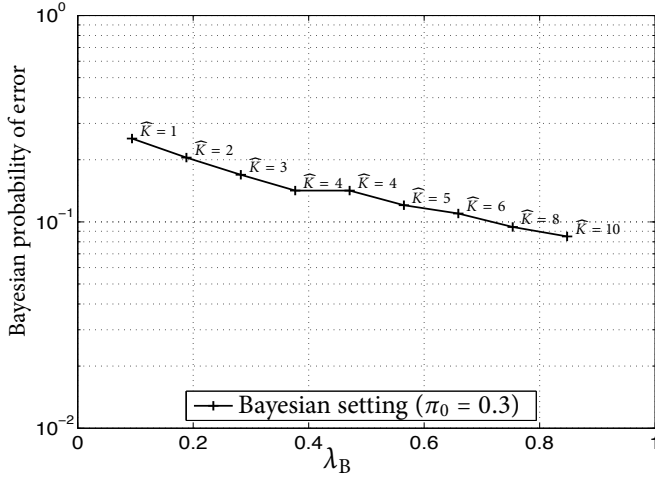


Figure 5.2: The solution path illustrates the Bayesian error probability for different values of the threshold λ_B . We use $M = 15$. The number of selected sensors \widehat{K} for a specific value of the threshold is also shown.

where $x \in [0, \infty)$. The local log-likelihood ratio is

$$\log l_m(x) = \log \frac{\mu_{1,m}}{\mu_{0,m}} + x(\mu_{0,m} - \mu_{1,m}).$$

Using (5.8), the local Bhattacharyya distance can be computed as

$$\mathcal{B}_m(\mathcal{H}_1 \parallel \mathcal{H}_0) = -\log \frac{\sqrt{4\mu_{0,m}\mu_{1,m}}}{\mu_{0,m} + \mu_{1,m}}.$$

Similarly, the local Kullback-Leibler distance $\mathcal{K}_m(\mathcal{H}_1 \parallel \mathcal{H}_0)$ can be computed as

$$\mathcal{K}_m(\mathcal{H}_1 \parallel \mathcal{H}_0) = \log \frac{\mu_{0,m}}{\mu_{1,m}} + \frac{\mu_{1,m}}{\mu_{0,m}} - 1,$$

the local Kullback-Leibler distance $\mathcal{K}_m(\mathcal{H}_0 \parallel \mathcal{H}_1)$ is obtained by interchanging the subscripts 0 and 1 in the above equation, and the local J-divergence is given as

$$\mathcal{D}_m(\mathcal{H}_1 \parallel \mathcal{H}_0) = 2 \log \frac{\mu_{1,m}}{\mu_{0,m}} + \frac{\mu_{0,m}^2 - \mu_{1,m}^2}{\mu_{0,m}\mu_{1,m}}.$$

These measures can be directly used in the proposed solvers to design sparse samplers.

5.7 Dependent observations

Throughout most of this chapter so far, we have assumed that the observations are conditionally independent. This assumption is generally valid if the sensors are responsible for the noise in the observations (i.e., receiver noise). However, if the sensors are subject to external noise or if the signal itself is stochastic in nature, then Assumption 5.1 might not be reasonable anymore. Consequently, the additive property of the considered distance measures is also no more valid.

The simplified design problem for this general case (i.e., without any independence assumption), again consists of finding a sparsest \mathbf{w} that results in a prescribed distance measure, where we express the Bhattacharyya distance, Kullback-Leibler distance, or J-divergence in terms of \mathbf{w} . The solution to the above generic problem is hard, nevertheless, we can solve it using standard nonlinear and often nonconvex optimization techniques for a given problem instance (see the example in §5.7.2). However, in some cases, a solution can be computed efficiently. As an example, the Gaussian observation case with uncommon means is detailed next.

5.7.1 Gaussian observations with uncommon means

Let us consider the case of binary signal detection in Gaussian noise, and assume the related conditional distributions are given by

$$\begin{aligned} \mathcal{H}_0: \quad \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\theta}_0, \boldsymbol{\Sigma}) \\ \mathcal{H}_1: \quad \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\theta}_1, \boldsymbol{\Sigma}), \end{aligned} \tag{5.34}$$

where the mean vectors $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$ as well as the covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$ are assumed to be perfectly known. Note that this model is a generalization of (5.25) with a nondiagonal covariance matrix. The results from Theorem 5.3 generalize to dependent observations. Thus, the error probabilities in (5.2) [or (5.23)] can without loss of optimality be replaced with the signal-to-noise

ratio (which is also related to the considered distance measures up to a constant)

$$s(\mathbf{w}) := [\Phi(\mathbf{w})\mathbf{m}]^T \Sigma^{-1}(\mathbf{w}) [\Phi(\mathbf{w})\mathbf{m}], \quad (5.35)$$

where we use $\mathbf{m} = \boldsymbol{\theta}_1 - \boldsymbol{\theta}_0$ and

$$\Sigma(\mathbf{w}) = \Phi(\mathbf{w})\Sigma\Phi^T(\mathbf{w}) \in \mathbb{R}^{K \times K}$$

is a submatrix of Σ that includes only the entries corresponding to the selected measurements. More specifically, we want to solve the problem

$$\arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \quad \text{s.to} \quad s(\mathbf{w}) \geq \lambda, \quad (5.36)$$

where λ is the desired signal-to-noise ratio (or distance measure, or error probability). However, in this case, the simplified problem does not admit an explicit solution. The optimal sampling scheme maximizes $s(\mathbf{w})$ in (5.35) over all possible $\mathbf{w} \in \{0,1\}^M$ such that \mathbf{w} is as sparse as possible. This incurs a combinatorial search over all the 2^M possible combinations. For example, with $M = 100$ candidate sensors, a performance evaluation of about 10^{30} possible choices is needed whose direct enumeration is clearly impossible.

The sampling design \mathbf{w} for (5.34) depends on the first and second order moments of the observations. In particular, it depends on $\boldsymbol{\theta}_0$, $\boldsymbol{\theta}_1$, and Σ .

We next propose some simplifications to solve this problem sub-optimally in polynomial time, yet with a performance that is comparable to the optimal one. Firstly, we write the covariance matrix Σ as

$$\Sigma = a\mathbf{I} + \mathbf{S}, \quad (5.37)$$

where a nonzero $a \in \mathbb{R}$ is chosen such that $\mathbf{S} \in \mathbb{R}^{M \times M}$ is invertible and well conditioned. Using (5.37) in (5.35), we obtain

$$s(\mathbf{w}) = \mathbf{m}^T \Phi^T(\mathbf{w}) [a\mathbf{I} + \Phi(\mathbf{w})\mathbf{S}\Phi^T(\mathbf{w})]^{-1} \Phi(\mathbf{w})\mathbf{m}. \quad (5.38)$$

Using Property 3.1 from Chapter 3, we can equivalently express $s(\mathbf{w})$ as

$$\begin{aligned} s(\mathbf{w}) &= \mathbf{m}^T \mathbf{S}^{-1} \mathbf{m} \\ &\quad - \mathbf{m}^T \mathbf{S}^{-1} [\mathbf{S}^{-1} + a^{-1} \text{diag}(\mathbf{w})]^{-1} \mathbf{S}^{-1} \mathbf{m}. \end{aligned} \quad (5.39)$$

Note that in contrast to (5.38), the design parameter \mathbf{w} only shows up at one place in (5.39), which makes the problem much easier. Using the Schur complement, the performance constraint in (5.36), i.e.,

$$\mathbf{m}^T \mathbf{S}^{-1} [\mathbf{S}^{-1} + a^{-1} \text{diag}(\mathbf{w})]^{-1} \mathbf{S}^{-1} \mathbf{m} \leq \lambda'$$

with $\lambda' := \lambda - \mathbf{m}^T \mathbf{S}^{-1} \mathbf{m}$ can be equivalently expressed as a linear matrix inequality in \mathbf{w} , i.e.,

$$\begin{bmatrix} \mathbf{S}^{-1} + a^{-1} \text{diag}(\mathbf{w}) & \mathbf{S}^{-1} \mathbf{m} \\ \mathbf{m}^T \mathbf{S}^{-1} & \lambda' \end{bmatrix} \succeq \mathbf{0}, \quad (5.40)$$

and therefore, it is convex in \mathbf{w} . The parameter a should be chosen such that \mathbf{S} is invertible and well-conditioned. Furthermore, because of (5.40) the matrix $\mathbf{S}^{-1} + a^{-1} \text{diag}(\mathbf{w})$ should be positive definite. This can be achieved, for example, by choosing a such that it satisfies the condition $0 < a < \lambda_{\min}\{\boldsymbol{\Sigma}\}$, since $w_m \geq 0$ for $m = 1, 2, \dots, M$. Although the constraint (5.40) is convex on \mathbf{w} , the optimization problem (5.36) is still not a convex problem due to the ℓ_0 -(quasi) norm cost function and the Boolean constraint.

Convex relaxation

The Boolean constraint set is relaxed to its convex hull, i.e., $0 \leq w_m \leq 1$, $m = 1, 2, \dots, M$, and we also relax the $\|\mathbf{w}\|_0$ constraint in (5.36) to its best convex approximate $\mathbf{1}^T \mathbf{w}$. Thus, the relaxed convex problem, more specifically, a semidefinite programming problem, is given as [cf. (R0) from Chapter 2]

$$\begin{aligned} & \arg \min_{\mathbf{w}} \mathbf{1}^T \mathbf{w} \\ & \text{s.to} \quad \begin{bmatrix} \mathbf{S}^{-1} + a^{-1} \text{diag}(\mathbf{w}) & \mathbf{S}^{-1} \mathbf{m} \\ \mathbf{m}^T \mathbf{S}^{-1} & \lambda' \end{bmatrix} \succeq \mathbf{0}, \\ & \quad 0 \leq w_m \leq 1, m = 1, 2, \dots, M. \end{aligned} \quad (5.41)$$

For a fixed K , the equivalent problem of the form (5.24) can be relaxed to

$$\begin{aligned} & \arg \max_{\mathbf{w}} s(\mathbf{w}) \\ & \text{s.to} \quad \mathbf{1}^T \mathbf{w} = K, \\ & \quad 0 \leq w_m \leq 1, m = 1, 2, \dots, M, \end{aligned}$$

which simplifies to

$$\begin{aligned} \arg \min_{\mathbf{w}} \quad & \mathbf{m}^T \mathbf{S}^{-1} [\mathbf{S}^{-1} + a^{-1} \text{diag}(\mathbf{w})]^{-1} \mathbf{S}^{-1} \mathbf{m} \\ \text{s.to} \quad & \mathbf{1}^T \mathbf{w} = K, \\ & 0 \leq w_m \leq 1, m = 1, 2, \dots, M. \end{aligned} \quad (5.42)$$

Here, only the second term of (5.39), which depends on \mathbf{w} is optimized (minimization is due to its negative sign). Writing (5.42) in the epigraph form [Boyd and Vandenberghe, 2004], we obtain

$$\begin{aligned} \arg \min_{\mathbf{w}, t} \quad & t \\ \text{s.to} \quad & \mathbf{1}^T \mathbf{w} = K, \\ & \begin{bmatrix} \mathbf{S}^{-1} + a^{-1} \text{diag}(\mathbf{w}) & \mathbf{S}^{-1} \mathbf{m} \\ \mathbf{m}^T \mathbf{S}^{-1} & t \end{bmatrix} \geq \mathbf{0}, \\ & 0 \leq w_m \leq 1, m = 1, 2, \dots, M, \end{aligned} \quad (5.43)$$

with auxiliary variable $t \in \mathbb{R}$.

Subsequently, the selected sensors (i.e., an approximate Boolean solution) can be computed using randomization techniques based on the solution from (5.41) or (5.43) as described in Algorithm 5.1. The relaxed convex problem can be solved using off-the-shelf software, for example, SeDuMi [Sturm, 1999].

Numerical example

To illustrate sparse sensing with dependent observations, we recall the simulation parameters from §5.6.1, but instead of independent noise, we use an autoregressive correlation matrix Σ , which is a Toeplitz matrix of the form

$$\Sigma = \sigma^2 \begin{bmatrix} 1 & \rho & \rho^2 & \dots & \rho^{M-1} \\ \rho & 1 & \rho & & \\ \rho^2 & \rho & 1 & & \vdots \\ \vdots & & & \ddots & \\ \rho^{M-1} & & & \dots & 1 \end{bmatrix}, \quad (5.44)$$

Algorithm 5.1 Randomized Rounding

1. **Given** the solution \mathbf{w}^* of (5.41) or (5.43) and a number of randomizations L .
2. **for** $l = 1$ to L
3. **generate** $w_{m,l} = 1$ with a probability w_m^*
 (or $w_{m,l} = 0$ with a probability $1 - w_m^*$)
 for $m = 1, 2, \dots, M$, where $\mathbf{w}_m^* = [\mathbf{w}^*]_m$.
4. **end**
5. **define** $\mathbf{w}_l = [w_{1,l}, \dots, w_{M,l}]^T$ and the index set of the candidate estimates satisfying the constraints as

$$\Omega \triangleq \begin{cases} \{l \mid s(\mathbf{w}_l) \geq \lambda, l = 1, 2, \dots, L\}, & \text{for (5.41)} \\ \{l \mid \|\mathbf{w}_l\|_0 = K, l = 1, 2, \dots, L\}, & \text{for (5.43)}. \end{cases}$$

6. If the set Ω is empty, go back to step 2.
7. **output** approximate solution $\mathbf{w}_{\text{round}}^* = \mathbf{w}_{l^*}$, where

$$l^* = \begin{cases} \arg \min_{l \in \Omega} \|\mathbf{w}_l\|_0, & \text{for (5.41)} \\ \arg \max_{l \in \Omega} s(\mathbf{w}_l), & \text{for (5.43)}. \end{cases}$$

with a known correlation coefficient $\rho \in [0, 1]$ and variance $\sigma^2 = 1$. Such a Σ is useful for modeling correlations between distributed sensors; for example, it can represent a spatially decaying correlation function. The convex relaxed problem (5.43) is solved using SeDuMi [Sturm, 1999].

The probability of error, i.e., P_m in the Neyman-Pearson setting and P_e in the Bayesian setting for different numbers of selected sensors is shown in Figure. 5.3. We underline the following observations. The solution with randomized rounding ($L = 50$) is shown in Figure. 5.3 for $\rho = \{0.25, 0.75\}$ with $a = 0.11$ in (5.37). For low values of the correlation coefficient, ρ , the convex relaxation with deterministic rounding is very close to optimal. For larger values of ρ , the solution of the relaxed problem with randomized rounding

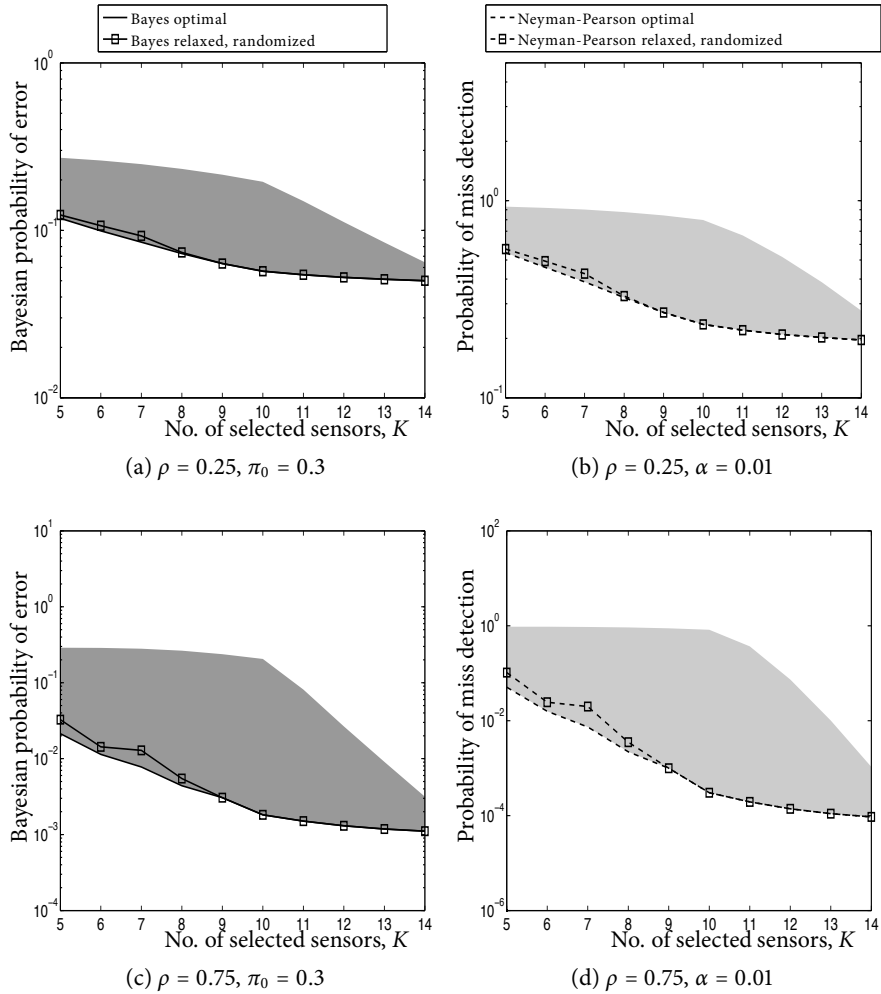


Figure 5.3: Error probabilities for (5.34) with different numbers of selected sensors K out of $M = 15$ sensors. The shaded regions indicate the performance with the worst to best subset of K sensors.

is still very close to optimal for large values of K , but less optimal for small values of K . As observed in the simulations, for $L = 50 \ll 2^{15}$, the sensing

design with randomization is near-optimal in terms of the error probability.

Correlation versus number of selected sensors

In this subsection, we focus on the number of sensors required to achieve a certain detection performance when the sensors become more *coherent*, i.e., as the correlation coefficient ρ approaches 1. To illustrate this, let us consider the numerical example introduced in §5.6.1 with $f \in \{0, 0.33\}$, but with an equi-correlated covariance matrix of the form

$$\boldsymbol{\Sigma} := \begin{bmatrix} 1 & \rho & \cdots & \rho \\ \rho & 1 & \cdots & \rho \\ \vdots & \vdots & \ddots & \vdots \\ \rho & \rho & \cdots & 1 \end{bmatrix} = (1 - \rho)\mathbf{I}_M + \rho\mathbf{1}_M\mathbf{1}_M^T, \quad (5.45)$$

with a known correlation coefficient $\rho \in [0, 1]$. Note that for such a covariance matrix, any $a \neq 1 - \rho$ leads to an invertible \mathbf{S} in (5.37) that can be used in the solver (5.41).

We first consider the case when $f = 0$, where all the M sensors have the same mean value, i.e., \mathbf{m} is the all-one vector up to a constant scaling. We refer to them as *identical sensors*. In this case, any subset of sensors is also the best subset of sensors, hence, random sensing is optimal. As the correlation coefficient ρ approaches 1, the amount of information (Kullback-Liebler distance/Bhattacharyya distance/J-divergence/signal-to-noise ratio) contributed by any random subset of $K > 1$ sensors is the same as that of the contribution from $K = 1$ sensor; see Figure 5.4(a). Thus, even with all the sensors selected the detection performance is limited to that of the performance with one sensor. This is a well-known result from distributed detection that extends to sampling design problems [Chamberland and Veeravalli, 2007].

A more interesting case, in particular for sensing design problems, is when the sensors are not identical ($f = 0.33$), i.e., \mathbf{m} has all different entries. When the sensors are not identical, as the correlation coefficient ρ approaches 1, the amount of information contained in the best subset of $K > 1$ sensors increases significantly; see Figure 5.4(b). More specifically, with equi-correlated yet different observations, to achieve a certain detection performance, the number of sensors required decreases significantly as the correlation coefficient




Figure 5.4: The signal-to-noise ratio for different values of the correlation coefficient ρ . (a) Identical sensors ($f = 0$). (b) Nonidentical sensors ($f = 0.33$).

ρ increases. The maximum achievable signal-to-noise ratio is proportional to the inverse of the minimum eigenvalue of $\Sigma(\mathbf{w})$, which is $\lambda_{\min}^{-1}\{\Sigma(\mathbf{w})\} = 1/(1 - \rho)$, for any sampler size $K \neq 0$. The optimal sparse sampler would choose the entries of \mathbf{m} that are most aligned to the eigenvector corresponding to the minimum eigenvalue of $\Sigma(\mathbf{w})$ (hence, as $\rho \rightarrow 1$ the signal-to-noise ratio is large). Similarly, if the entries of \mathbf{m} are parallel to the eigenvector corresponding to the maximum eigenvalue of $\Sigma(\mathbf{w})$, that is, the all-one vector, then the signal-to-noise ratio is minimized; this is the case in Figure 5.4(a).

5.7.2 Gaussian Observation with Uncommon Covariances

We now provide some extensions and offer guidelines for determining sparse sensing mechanisms for testing between two covariance matrices. That is, when the covariance structures are different under both hypotheses. Suppose the conditional distributions are given by

$$\begin{aligned} \mathcal{H}_0 : \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\theta}, \Sigma_0) \\ \mathcal{H}_1 : \mathbf{x} &\sim \mathcal{N}(\boldsymbol{\theta}, \Sigma_1), \end{aligned} \tag{5.46}$$

where the mean vector $\boldsymbol{\theta} \in \mathbb{R}^N$ as well as the $N \times N$ covariance matrices $\boldsymbol{\Sigma}_0$ and $\boldsymbol{\Sigma}_1$ are assumed to be perfectly known. This model is a generalization of (5.29) with nondiagonal covariance matrices.

As with (5.29), the distance measures are not equal to each other. Using (5.4), the Bhattacharyya distance for the observations of the form $\mathbf{y} = \boldsymbol{\Phi}(\mathbf{w})\mathbf{x}$ can be computed as

$$\begin{aligned} \mathcal{B}(\mathcal{H}_1 \parallel \mathcal{H}_0) &= \frac{1}{2} \log \det\{\boldsymbol{\Sigma}_{01}(\mathbf{w})\} \\ &\quad - \frac{1}{4} (\log \det\{\boldsymbol{\Sigma}_0(\mathbf{w})\} + \log \det\{\boldsymbol{\Sigma}_1(\mathbf{w})\}), \end{aligned} \quad (5.47)$$

where $\boldsymbol{\Sigma}_i(\mathbf{w}) = \boldsymbol{\Phi}(\mathbf{w})\boldsymbol{\Sigma}_i\boldsymbol{\Phi}^T(\mathbf{w})$ for $i = 0, 1$, with $2\boldsymbol{\Sigma}_{01} = \boldsymbol{\Sigma}_0 + \boldsymbol{\Sigma}_1$. Similarly, using (5.10), we can show that the Kullback-Leibler distance is given by

$$\begin{aligned} \mathcal{K}(\mathcal{H}_1 \parallel \mathcal{H}_0) &= \frac{1}{2} (\text{tr}\{\boldsymbol{\Sigma}_0^{-1}(\mathbf{w})\boldsymbol{\Sigma}_1(\mathbf{w})\} - \|\mathbf{w}\|_0 \\ &\quad - \log \det\{\boldsymbol{\Sigma}_0^{-1}(\mathbf{w})\boldsymbol{\Sigma}_1(\mathbf{w})\}). \end{aligned} \quad (5.48)$$

Here, $\boldsymbol{\Sigma}_0^{-1}(\mathbf{w})\boldsymbol{\Sigma}_1(\mathbf{w})$ is the *signal-to-noise ratio* matrix.

We can express the Bhattacharyya and Kullback-Leibler distance as a difference of concave functions by relaxing $\mathbf{w} \in \{0, 1\}^M$ to $[0, 1]^M$. That is, we can express (5.47) and (5.48) as

$$f(\mathbf{w}) = f_0(\mathbf{w}) - f_1(\mathbf{w}),$$

where $f_0(\mathbf{w})$ and $f_1(\mathbf{w})$ are concave functions of its arguments; see Appendix 5.D for the explicit expressions of $f_0(\mathbf{w})$ and $f_1(\mathbf{w})$. As a consequence, the relaxed problem (for fixed K)

$$\arg \min_{\mathbf{w} \in [0, 1]^M} f_1(\mathbf{w}) - f_0(\mathbf{w}) \quad \text{s.to} \quad \mathbf{1}^T \mathbf{w} = K,$$

is not a convex problem as the cost is not a convex function of its argument and has to be solved using nonconvex optimization techniques.

One such heuristic to solve the difference of convex problems is the convex-concave procedure [Yuille and Rangarajan, 2003], where the concave term (here, $f_1(\mathbf{w})$) is replaced with its affine approximation (more generally, any

reasonable convex approximation of $f_1(\mathbf{w})$ while the convex portion, i.e., $-f_0(\mathbf{w})$ is retained. The resulting convex problem is iteratively solved to obtain a local optimum.

The J-divergence can be computed using (5.18) as

$$\begin{aligned} \mathcal{D}(\mathcal{H}_1||\mathcal{H}_0) &= \frac{1}{2}\text{tr}\{\boldsymbol{\Sigma}_0^{-1}(\mathbf{w})\boldsymbol{\Sigma}_1(\mathbf{w})\} \\ &\quad + \frac{1}{2}\text{tr}\{\boldsymbol{\Sigma}_1^{-1}(\mathbf{w})\boldsymbol{\Sigma}_0(\mathbf{w})\} - \|\mathbf{w}\|_0. \end{aligned} \quad (5.49)$$

We next show that maximizing the J-divergence over \mathbf{w} can be cast as a convex problem.

Let the covariance matrices $\boldsymbol{\Sigma}_0 = \boldsymbol{\Sigma}_0^{1/2}\boldsymbol{\Sigma}_0^{T/2}$ and $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_1^{1/2}\boldsymbol{\Sigma}_1^{T/2}$, respectively, admit the decomposition

$$\boldsymbol{\Sigma}_0 = a_0\mathbf{I} + \mathbf{S}_0,$$

and

$$\boldsymbol{\Sigma}_1 = a_1\mathbf{I} + \mathbf{S}_1,$$

with scalars a_0 and a_1 chosen such that \mathbf{S}_0 and \mathbf{S}_1 are invertible. Using Property 3.1 from Chapter 3, we can show that the J-divergence (5.49) is equivalent to

$$\begin{aligned} \mathcal{D}(\mathcal{H}_1||\mathcal{H}_0) &= \frac{1}{2}\text{tr}\{\mathbf{S}_0^{-1}\boldsymbol{\Sigma}_1 \\ &\quad - \mathbf{S}_0^{-1}[\mathbf{S}_0^{-1} + a_0^{-1}\text{diag}(\mathbf{w})]^{-1}\mathbf{S}_0^{-1}\boldsymbol{\Sigma}_1\} \\ &\quad + \frac{1}{2}\text{tr}\{\mathbf{S}_1^{-1}\boldsymbol{\Sigma}_0 \\ &\quad - \mathbf{S}_1^{-1}[\mathbf{S}_1^{-1} + a_1^{-1}\text{diag}(\mathbf{w})]^{-1}\mathbf{S}_1^{-1}\boldsymbol{\Sigma}_0\} - \|\mathbf{w}\|_0. \end{aligned}$$

Thus, maximizing the J-divergence over \mathbf{w} for a fixed K is the same as minimizing

$$\begin{aligned} &\frac{1}{2}\text{tr}\{\mathbf{S}_0^{-1}[\mathbf{S}_0^{-1} + a_0^{-1}\text{diag}(\mathbf{w})]^{-1}\mathbf{S}_0^{-1}\boldsymbol{\Sigma}_1\} \\ &\quad + \frac{1}{2}\text{tr}\{\mathbf{S}_1^{-1}[\mathbf{S}_1^{-1} + a_1^{-1}\text{diag}(\mathbf{w})]^{-1}\mathbf{S}_1^{-1}\boldsymbol{\Sigma}_0\} \end{aligned}$$

over \mathbf{w} . To cast this as a convex problem, we introduce two variables

$$\begin{aligned}\mathbf{Z}_0 &= \boldsymbol{\Sigma}_1^{T/2} \mathbf{S}_0^{-1} [\mathbf{S}_0^{-1} + a_0^{-1} \text{diag}(\mathbf{w})]^{-1} \mathbf{S}_0^{-1} \boldsymbol{\Sigma}_1^{1/2}; \\ \mathbf{Z}_1 &= \boldsymbol{\Sigma}_0^{T/2} \mathbf{S}_1^{-1} [\mathbf{S}_1^{-1} + a_1^{-1} \text{diag}(\mathbf{w})]^{-1} \mathbf{S}_1^{-1} \boldsymbol{\Sigma}_0^{1/2},\end{aligned}$$

and obtain

$$\begin{aligned}\arg \min_{\mathbf{w}, \mathbf{Z}_0, \mathbf{Z}_1} & \frac{1}{2} \text{tr}\{\mathbf{Z}_0\} + \frac{1}{2} \text{tr}\{\mathbf{Z}_1\} \\ \text{s.to} & \quad \mathbf{1}^T \mathbf{w} = K, \\ & \quad \boldsymbol{\Sigma}_1^{T/2} \mathbf{S}_0^{-1} [\mathbf{S}_0^{-1} + a_0^{-1} \text{diag}(\mathbf{w})]^{-1} \mathbf{S}_0^{-1} \boldsymbol{\Sigma}_1^{1/2} \leq \mathbf{Z}_0 \\ & \quad \boldsymbol{\Sigma}_0^{T/2} \mathbf{S}_1^{-1} [\mathbf{S}_1^{-1} + a_1^{-1} \text{diag}(\mathbf{w})]^{-1} \mathbf{S}_1^{-1} \boldsymbol{\Sigma}_0^{1/2} \leq \mathbf{Z}_1 \\ & \quad 0 \leq w_m \leq 1, m = 1, 2, \dots, M.\end{aligned}\tag{5.50}$$

The second and the third constraint can be, respectively, expressed as an LMI in \mathbf{w} , i.e.,

$$\begin{aligned}\begin{bmatrix} \mathbf{Z}_0 & \mathbf{S}_0^{-1} \boldsymbol{\Sigma}_1^{1/2} \\ \boldsymbol{\Sigma}_1^{T/2} \mathbf{S}_0^{-1} & \mathbf{S}_0^{-1} + a_0^{-1} \text{diag}(\mathbf{w}) \end{bmatrix} & \geq \mathbf{0}, \\ \begin{bmatrix} \mathbf{Z}_1 & \mathbf{S}_1^{-1} \boldsymbol{\Sigma}_0^{1/2} \\ \boldsymbol{\Sigma}_0^{T/2} \mathbf{S}_1^{-1} & \mathbf{S}_1^{-1} + a_1^{-1} \text{diag}(\mathbf{w}) \end{bmatrix} & \geq \mathbf{0}.\end{aligned}$$

An approximate Boolean solution has to be subsequently computed using randomized rounding.

The optimization problem of the form (5.20) with unknown K can be derived along similar lines by relaxing the $\|\mathbf{w}\|_0$ in the cost function. Before we end this section, we make the following remarks.

- For Gaussian observations, we recall that an upper bound on P_e and P_m can be obtained in terms of J-divergence. Hence, optimizing J-divergence is reasonable under the Bayesian and Neyman-Pearson setting.

- For general Gaussian dependent observations (with uncommon means and uncommon covariances under both hypotheses), the design problems are straightforward combinations of the problems derived in §5.7.1 and §5.7.2.

5.8 Discussion

In this chapter we have developed a framework for structured and sparse sampler design for distributed detection problems. In particular, we have addressed binary hypothesis testing in both the Bayesian and Neyman-Pearson setting. The proposed framework can be directly applied to sensor placement/selection, sample selection, and fully-decentralized data compression, where we seek the best subset of sensor/sampling locations or data samples that results in a desired detection probability. To simplify the design problem, we have used a number of distance measures that quantify the closeness or divergence between the conditional distributions of the observations. We give an explicit solution for the sampling design problem with conditionally independent observations and the results are summarized as follows. The best sensors are the ones with the smallest local average root-likelihood ratio and largest local average log-likelihood ratio in the Bayesian and Neyman-Pearson setting, respectively. The framework has also been generalized to conditionally dependent observations with a thorough analysis for the Gaussian case. In that context, we have shown that, for uncommon means and common covariances under both hypotheses, the number of nonidentical Gaussian sensors required to achieve a desired detection performance reduces significantly as the sensors become more coherent.

5.A Proof of Proposition 5.1

In this section, we prove that the additivity of the Bhattacharyya distance is preserved with compression using $\Phi(\mathbf{w})$. Using the conditional independence assumption, i.e., Assumption 5.1, the Bhattacharyya distance in (5.4)

can be expressed as

$$\begin{aligned}\mathcal{B}(\mathcal{H}_1\|\mathcal{H}_0) &= -\log \mathbb{E}_{|\mathcal{H}_0} \left\{ \sqrt{l(\mathbf{y})} \right\} \\ &= -\log \mathbb{E}_{|\mathcal{H}_0} \left\{ \prod_{m=1}^M [l_m(x)]^{w_m/2} \right\} \\ &= -\log \prod_{m=1}^M \mathbb{E}_{|\mathcal{H}_0} \left\{ [l_m(x)]^{w_m/2} \right\},\end{aligned}$$

where $l_m(x)$ is the local likelihood ratio at the m th sensor. Since $w_m \in \{0, 1\}$, we can further simplify $\mathcal{B}(\mathcal{H}_1\|\mathcal{H}_0)$ to

$$\begin{aligned}\mathcal{B}(\mathcal{H}_1\|\mathcal{H}_0) &= -\log \prod_{m=1}^M \left(\mathbb{E}_{|\mathcal{H}_0} \left\{ \sqrt{l_m(x)} \right\} \right)^{w_m} \\ &= \sum_{m=1}^M -w_m \log \mathbb{E}_{|\mathcal{H}_0} \left\{ \sqrt{l_m(x)} \right\} \\ &= \sum_{m=1}^M w_m \mathcal{B}_m(\mathcal{H}_1\|\mathcal{H}_0).\end{aligned}$$

5.B Upper bound on P_m

To derive the upper bound on P_m stated in Theorem 5.1, we use Chebyshev's inequality [Hoeffding, 1963]

$$\Pr(X - \mathbb{E}\{X\} \geq t) \leq \frac{1}{1 + \frac{t^2}{v^2}}, \quad (5.51)$$

where X is a random variable with variance v^2 and t is a constant. Then, P_m simplifies to

$$\begin{aligned}P_m &= \Pr(\log l(\mathbf{y}) \leq \log \gamma | \mathcal{H}_1) \\ &= \Pr(\log l(\mathbf{y}) - \mathbb{E}_{|\mathcal{H}_1} \{\log l(\mathbf{y})\} \\ &\quad \leq \log \gamma - \mathbb{E}_{|\mathcal{H}_1} \{\log l(\mathbf{y})\} | \mathcal{H}_1) \\ &= \Pr(\log l(\mathbf{y}) - \mathcal{K}(\mathcal{H}_1\|\mathcal{H}_0) \geq \mathcal{K}(\mathcal{H}_1\|\mathcal{H}_0) - \log \gamma | \mathcal{H}_1),\end{aligned}$$

where the last equation has the same form as the inequality (5.51) with $t = \log \gamma - \mathcal{K}(\mathcal{H}_1 \| \mathcal{H}_0)$.

If the variance of $\log l(\mathbf{y})$ is v^2 , then, from (5.51), we have

$$P_m \leq \frac{1}{1 + \frac{(\mathcal{K}(\mathcal{H}_1 \| \mathcal{H}_0) - \log \gamma)^2}{v^2}}$$

This completes the proof.

5.C Proof of Proposition 5.2

In this section, we prove that the additivity of the Kullback-Leibler distance for independent observations is preserved with compression using $\Phi(\mathbf{w})$. Assuming Assumption 5.1 holds, then the Kullback-Leibler distance in (5.10) can be expressed as

$$\begin{aligned} \mathcal{K}(\mathcal{H}_1 \| \mathcal{H}_0) &= \mathbb{E}_{|\mathcal{H}_1} \{ \log l(\mathbf{y}) \} \\ &= \mathbb{E}_{|\mathcal{H}_1} \left\{ \log \prod_{m=1}^M [l_m(x)]^{w_m} \right\} \\ &= \mathbb{E}_{|\mathcal{H}_1} \left\{ \sum_{m=1}^M w_m \log l_m(x) \right\} \\ &= \sum_{m=1}^M w_m \mathbb{E}_{|\mathcal{H}_1} \{ \log l_m(x) \} \\ &= \sum_{m=1}^M w_m \mathcal{K}_m(\mathcal{H}_1 \| \mathcal{H}_0), \end{aligned}$$

where $l_m(x)$ is the local likelihood ratio at the m th sensor.

5.D Expressions for $f_0(\mathbf{w})$ and $f_1(\mathbf{w})$

Let the covariance matrices Σ_{01} , Σ_0 and Σ_1 , respectively, admit a decomposition of the form $\Sigma_{01} = a_{01}\mathbf{I} + \mathbf{S}_{01}$, $\Sigma_0 = a_0\mathbf{I} + \mathbf{S}_0$, and $\Sigma_1 = a_1\mathbf{I} + \mathbf{S}_1$. Here, the

scalars a_{01} , a_0 , and a_1 are, respectively, chosen such that the matrices \mathbf{S}_{01} , \mathbf{S}_0 , and \mathbf{S}_1 are invertible.

Using the Sylvester's determinant identity

$$\det\{\mathbf{A} + \mathbf{BC}\} = \det\{\mathbf{A}\} \det\{\mathbf{I} + \mathbf{CA}^{-1}\mathbf{B}\}, \quad (5.52)$$

we can express, for example,

$$\begin{aligned} \det\{\Phi \Sigma_0 \Phi^T\} &= \det\{a_0 \mathbf{I} + \Phi \mathbf{S}_0 \Phi^T\} \\ &= a_0^M \det\{\mathbf{I} + a_0^{-1} \Phi^T \Phi \mathbf{S}_0\} \\ &= a_0^M \det\{\mathbf{I} + a_0^{-1} \text{diag}(\mathbf{w}) \mathbf{S}_0\}. \end{aligned}$$

Bhattacharyya distance

Ignoring the terms that are independent of the optimization variable \mathbf{w} , we can express the Bhattacharyya distance (5.47) as

$$f(\mathbf{w}) = f_0(\mathbf{w}) - f_1(\mathbf{w}),$$

where

$$f_0(\mathbf{w}) := \frac{1}{2} \log \det\{\mathbf{I} + a_{01}^{-1} \text{diag}(\mathbf{w}) \mathbf{S}_{01}\}$$

and

$$\begin{aligned} f_1(\mathbf{w}) &:= \frac{1}{4} \left(\log \det\{\mathbf{I} + a_0^{-1} \text{diag}(\mathbf{w}) \mathbf{S}_0\} \right. \\ &\quad \left. + \log \det\{\mathbf{I} + a_1^{-1} \text{diag}(\mathbf{w}) \mathbf{S}_1\} \right), \end{aligned}$$

are concave functions on $\mathbf{w} \in [0, 1]^M$.

Kullback-Leibler distance

Using Property 3.1 from Chapter 3, we can write the first term of (5.48), that is,

$$\text{tr}\{\Sigma_0^{-1}(\mathbf{w}) \Phi \Sigma_1 \Phi^T\} = \text{tr}\{\Phi^T \Sigma_0^{-1}(\mathbf{w}) \Phi \Sigma_1\}$$

as

$$\text{tr}\{\mathbf{S}_0^{-1}\boldsymbol{\Sigma}_1 - \mathbf{S}_0^{-1}[\mathbf{S}_0^{-1} + a_0^{-1}\text{diag}(\mathbf{w})]^{-1}\mathbf{S}_0^{-1}\boldsymbol{\Sigma}_1\}.$$

The above function can be expressed as a convex function in \mathbf{w} (e.g., using the epigraph form). The second term of (5.48) can be relaxed to a convex function $\mathbf{1}^T \mathbf{w}$. The last term of (5.48), that is, $\log \det\{\boldsymbol{\Sigma}_0^{-1}(\mathbf{w})\boldsymbol{\Sigma}_1(\mathbf{w})\}$ is equivalent to

$$\begin{aligned} & \log \det\{\boldsymbol{\Phi}\boldsymbol{\Sigma}_1\boldsymbol{\Phi}^T\} - \log \det\{\boldsymbol{\Phi}\boldsymbol{\Sigma}_0\boldsymbol{\Phi}^T\} = \\ & \log \det\{\mathbf{I} + a_1^{-1}\text{diag}(\mathbf{w})\mathbf{S}_1\} - \log \det\{\mathbf{I} + a_0^{-1}\text{diag}(\mathbf{w})\mathbf{S}_0\}. \end{aligned}$$

Thus, we can equivalently express (5.48) as $f(\mathbf{w}) = f_0(\mathbf{w}) - f_1(\mathbf{w})$ with

$$\begin{aligned} f_0(\mathbf{w}) := & -\text{tr}\{\mathbf{S}_0^{-1}[\mathbf{S}_0^{-1} + a_0^{-1}\text{diag}(\mathbf{w})]^{-1}\mathbf{S}_0^{-1}\boldsymbol{\Sigma}_1\} - \mathbf{1}^T \mathbf{w} \\ & + \log \det\{\mathbf{I} + a_1^{-1}\text{diag}(\mathbf{w})\mathbf{S}_1\} \end{aligned}$$

and

$$f_1(\mathbf{w}) := \log \det\{\mathbf{I} + a_0^{-1}\text{diag}(\mathbf{w})\mathbf{S}_0\},$$

which are concave in \mathbf{w} .

Chapter 6

Continuous Sparse Sensing

Contents

6.1	Introduction	133
6.2	Sensitivity to gridding	135
6.3	Sensing model based on binning	137
6.4	Risk function for continuous sparse sensing	139
6.5	Solver	140
6.6	Discussion	144

6.1 Introduction

In the sparse sensing schemes considered so far in this thesis, the focus was on discrete sparse sensing. That is, we were selecting sparse sensing patterns from a discrete set of candidates, e.g., temporal samples, sensor positions, which were obtained by gridding the output space, in order to reach a desired inference performance. In this chapter we will discuss continuous (or off-the-grid) sparse sensing, where we can take samples anywhere in the continuous output space (i.e., we can sample in between the grid points). Continuous

Part of this chapter was published as: S.P. Chepuri and G. Leus. Continuous Sensor Placement. *IEEE Signal Processing Letters*, 22(5): 544-548, May 2015.

sparse sensing will reduce the sensing costs for a given inference performance, or in other words, it will improve the inference performance for a fixed sampler size.

To realize continuous sparse sensing, we start with a discretized output space and model every sampling point in the continuous sampling space as a discrete sampling point plus a continuous perturbation. Then, we solve for the smallest set of possible discrete points as well as the best possible perturbations that result in a desired inference performance. Although we focus on the estimation task with linear models in this chapter, the proposed approach can be generalized to nonlinear models and also other inference tasks (e.g., filtering and detection) discussed in this thesis. Depending on the nature of the inference task, the risk function that quantifies the inference quality will depend on the discrete sampling points as well as their perturbations. In sum, the main contribution of this chapter is a framework of continuous sparse sensing, which allows for off-the-grid sensor placement.

Let $x(t)$ denote the observation signal with a continuous-domain argument, where without loss of generality $t \in [0, T]$ denotes the sampling domain. We will restrict ourselves to the one-dimensional spatial domain, but the ideas presented can be applied directly to higher dimensions and even to temporal or spatio-temporal domains. Assume that $x(t)$ represents the measured physical field over a continuous one-dimensional space t , and it satisfies the linear model

$$x(t) = \mathbf{h}^H(t)\boldsymbol{\theta} + n(t) \quad (6.1)$$

where $\boldsymbol{\theta} \in \mathbb{R}^N$ collects the parameters to be estimated, $\mathbf{h}(t) \in \mathbb{C}^N$ is the known linear model representing the mapping between the parameters and the measurements, and $n(t)$ is the noise. Furthermore, we assume $\mathbf{h}(t) = 0$ for $t < 0$ and $t > T$. In other words, $\mathbf{h}(t)$ is completely described by its variation in the interval $t \in [0, T]$.

The fundamental question of interest is—where to sample $x(t)$ in order to reach a desired inference performance? We next state the problem more precisely.

Problem 6.1. *Given the model (6.1) and a desired estimation accuracy, find the optimal sampling locations in the range $[0, T]$ such that the number of samples is minimum and the desired estimation accuracy is achieved.*

6.2 Sensitivity to gridding

We now recall the sensing model (2.3) that we introduced in Chapter 2, where we discretize (e.g., by regular sampling) the output space with M points denoted by $\{t_m\}_{m=1}^M$. The inference performance (estimation accuracy, in this case) is limited by the choice of the initial grid points $\{t_m\}_{m=1}^M$ as the resolution might be too low, especially when $\mathbf{h}(t)$ is fast varying compared to the chosen grid.

Let $x_m = x(t_m)$ be the discrete-domain observations, $\mathbf{h}_m = \mathbf{h}(t_m)$ denote the discretized model, and $n_m = n(t_m)$ represent the noise. We assume that the noise is white Gaussian with variance σ^2 (for spatial sampling, the sensor noise variance is independent of the sampling density). Using the above notations, we can write the discrete-domain version of (6.1) as

$$x_m = \mathbf{h}_m^H \boldsymbol{\theta} + n_m, m = 1, 2, \dots, M. \quad (6.2)$$

The data $\mathbf{x} = [x(t_1), x(t_2), \dots, x(t_M)]^T$ is acquired through the discrete sparse sensing operation as

$$\mathbf{y} = \Phi(\mathbf{w})\mathbf{x} = \text{diag}_r(\mathbf{w})\mathbf{x},$$

where we design a sparsest vector

$$\mathbf{w} = [w_1, w_2, \dots, w_M] \in \{0, 1\}^M$$

by solving (P0) introduced in Chapter 2 with $f(\mathbf{w}) = \sigma^2 \text{tr}\{(\sum_{m=1}^M w_m \mathbf{h}_m \mathbf{h}_m^H)^{-1}\}$. That is, we solve

$$\begin{aligned} & \arg \min_{\mathbf{w} \in \{0,1\}^M} \|\mathbf{w}\|_0 \\ & \text{s.to } \sigma^2 \text{tr}\left\{\left(\sum_{m=1}^M w_m \mathbf{h}_m \mathbf{h}_m^H\right)^{-1}\right\} \leq \lambda. \end{aligned} \quad (6.3)$$

Recall that $f(\mathbf{w}) = \sigma^2 \text{tr}\{(\sum_{m=1}^M w_m \mathbf{h}_m \mathbf{h}_m^H)^{-1}\}$ is the Cramér-Rao bound (equal to the mean squared error of the least squares estimate) for linear models in Gaussian noise; see Remark 3.1 from Chapter 3.

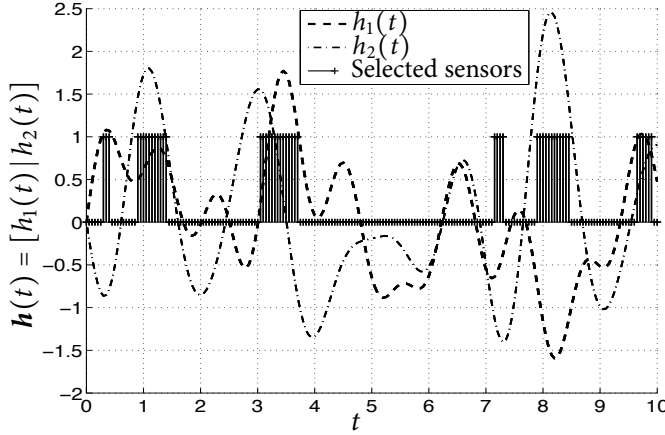


Figure 6.1: The field is measured at $M = 201$ discrete locations obtained by uniformly gridding the interval $[0, 10]$ m. The mean squared error bound is set to $\lambda = 0.026$, leading to $K = 50$ selected sensors.

We now illustrate the sensitivity to gridding with the following numerical example. Consider the linear model in (6.1) with the following specifications. Let the parameter vector $\boldsymbol{\theta}$ be of length 2×1 . Consider a sum of sinusoids model for $\mathbf{h}(t) = [h_1(t), h_2(t)]^T$ with $h_i(t) = \sum_{p=1}^{P_i} \alpha_{p,i} \sin(2\pi f_{p,i} t)$ for $i = 1, 2$. Let $\mathbf{f}_i = [f_{1,i}, \dots, f_{P_i,i}]^T$ and $\boldsymbol{\alpha}_i = [\alpha_{1,i}, \dots, \alpha_{P_i,i}]^T$ for $i = 1, 2$. We use the following parameters: $P_1 = 5, P_2 = 5, T = 10, \sigma^2 = 1$,

$$\begin{aligned} \mathbf{f}_1 &= [0.1, 0.33, 0.67, 0.78, 0.95]^T, \quad \mathbf{f}_2 = [0.15, 0.7, 0.4, 0.58, 0.85]^T, \\ \boldsymbol{\alpha}_1 &= [0.5, 0.65, 0.3, -0.15, 0.45]^T, \quad \boldsymbol{\alpha}_2 = [-0.25, -0.33, -0.6, 0.95, -0.25]^T. \end{aligned}$$

6.2.1 Coarse gridding

Assume that the field can be measured at $M = 5$ potential locations. Let us consider the following case of candidate sampling locations $t_m := \{1, 3, 5, 7, 9\}$. The mean squared error resulting from the samples at these locations is 1.47. The optimal mean squared error using 5 samples, on the other hand, is around 0.16.

We remark the following two observations. Firstly, the mean squared error resulting from the samples at the locations $t_m := \{1, 3, 5, 7, 9\}$ is much higher than the optimal mean squared error. Secondly, any subset of these 5 samples will naturally also result in a mean squared error larger than 1.47. Hence, due to the involved discretization, coarse gridding might not lead to the desired mean squared error even if all candidate samples/sensors are selected.

6.2.2 Fine gridding

Alternatively, the initial grid can be very dense with the candidate sampling locations at infinitesimal distance apart. Then (6.3) would choose many sensors within one or more virtual sampling bins.

For the setup discussed in the previous subsection, for a fine grid with $M = 201$ candidate sampling points the solution to the relaxed version of (6.3) is illustrated in Figure. 6.1. The sensors within the most informative bin are selected first before going to the next informative bin, and so on, till the desired mean squared error is achieved.

The fine gridding has two main drawbacks. Firstly, it might not be practically feasible to sample so close to each other. In addition, the reason why more samples are selected within a certain bin is to improve the signal-to-noise ratio. In case of spatial sampling, it might be desirable to restrict the number of spatial samples. That is, instead of placing additional (expensive) sensors within a certain bin, the signal-to-noise ratio can be compensated by other (cheaper) means, e.g., temporal averaging using a single sensor. Finally, note that the solvers based on convex optimization techniques incur a cubic complexity making fine gridding also computationally less viable.

6.3 Sensing model based on binning

The motivation behind a coarse discretization in discrete sparse sensing was computational tractability, but its performance is limited by the choice of the initial grid. On the contrary, fine gridding suffers from a high computational complexity and multiple closely spaced sensors for signal-to-noise ratio improvement. In this section, we present the sparse sensing model based on

binning, which allows to sample anywhere in the continuous sampling domain. Specifically, we take one sample per bin. To realize this, we augment the discrete model by including additional variables that account for the continuous nature of the sampling domain.

We now recall the sparse sensing model that we introduced in Chapter 2, where we acquire $x(t)$ through a continuous-domain sparse sensing function $w(t) = \sum_{m=1}^K \delta(t - \tau_m)$ with unknown indices $\{\tau_m\}_{m=1}^K$ and unknown number of samples K , as

$$y(t) = w(t)x(t) = \sum_{m=1}^K x(\tau_m)\delta(t - \tau_m).$$

By discretizing (e.g., regular sampling) the output space with $M \gg K$ points denoted by $\{t_m\}_{m=1}^M$ we alternatively modeled $y(t)$ as [cf. (2.3)]

$$y(t) = \sum_{m=1}^M w_m x(t_m) \delta(t - t_m)$$

where $w_m = (0)1$ indicates whether sample $x(t_m)$ is (not) selected. Here, the assumption was that the $\{\tau_m\}_{m=1}^K$ lie on the discrete grid.

When they do not lie on the discrete grid we follow a binning approach. If $x(t)$ is sufficiently smooth (i.e., its first-order derivative exists and is continuous), then local shifts of $x(t)$ can be approximated using its derivative based on a first-order Taylor expansion:

$$x(t_m + p_m) \approx x_m + p_m x'_m, \quad m = 1, 2, \dots, M, \quad (6.4)$$

where $\{t_m\}_{m=1}^M$ are the discrete sampling points, p_m represents the continuous perturbation around t_m with $|p_m| < 0.5\delta$ (δ denotes the bin size), and x'_m is the derivative of $x(t)$ towards t evaluated at t_m , i.e., $x'_m = \left. \frac{\partial x(t)}{\partial t} \right|_{t=t_m}$. Using this approximation, we arrive at an off-the-grid sensing model, i.e., we can model $y(t)$ as

$$y(t) = \sum_{m=1}^M w_m (x_m + p_m x'_m) \delta(t - t_m - p_m). \quad (6.5)$$

Such first-order interpolations have also been used in the context of sparse signal recovery in continuous compressive sensing to overcome problems due

to gridding, but at the input grid [Ekanadham et al., 2011, Zhu et al., 2011]. As a remark, alternative interpolation techniques (e.g., polar interpolation [Ekanadham et al., 2011]) can be considered.

Stacking $\{w_m\}_{m=1}^M$ in vector \mathbf{w} , $\{p_m\}_{m=1}^M$ in vector \mathbf{p} , and $\{x'_m\}_{m=1}^M$ in vector \mathbf{x}' , we can write the discrete-domain counterpart of (6.5) as

$$\mathbf{y} = \Phi(\mathbf{w}) [\mathbf{x} + \text{diag}(\mathbf{x}')\mathbf{p}] = \text{diag}_r(\mathbf{w}) [\mathbf{x} + \text{diag}(\mathbf{x}')\mathbf{p}], \quad (6.6)$$

where \mathbf{y} stacks $\{y_m\}_{m=1}^M$.

For the linear model (6.1), we can then represent the off-the-grid samples as

$$y_m = w_m [(\mathbf{h}_m + p_m \mathbf{h}'_m)^H \boldsymbol{\theta} + n_m], \quad m = 1, 2, \dots, M, \quad (6.7)$$

where $\mathbf{h}'_m = \left. \frac{\partial \mathbf{h}(t)}{\partial t} \right|_{t=t_m}$. Note that by using this idea, gridding actually results in binning, where we take at most one sample per bin. In what follows, we will derive the risk function that depends on \mathbf{w} and \mathbf{p} for the linear inverse problem.

6.4 Risk function for continuous sparse sensing

As discussed in Chapter 3, for statistical estimation problems, we use scalar measures of the Fisher information matrix to quantify the estimation accuracy. The computation of the FIM for a linear model (6.7) is straightforward [cf. (3.5) from Chapter 3]. It is given by

$$\begin{aligned} \mathbf{F}(\mathbf{w}, \mathbf{p}) &= \sum_{m=1}^M w_m \mathbf{h}_m \mathbf{h}_m^H + w_m p_m^2 \mathbf{h}'_m \mathbf{h}'_m{}^H \\ &\quad + w_m p_m (\mathbf{h}'_m \mathbf{h}_m^H + \mathbf{h}_m \mathbf{h}'_m{}^H). \end{aligned}$$

The A-optimality criterion [cf. §3.1.1 from Chapter 3], which is the trace of the inverse FIM, for linear Gaussian models corresponds to the mean squared error of the least squares estimate. Introducing variables $\mathbf{u} = [u_1, u_2, \dots, u_M]^T$ with $u_m = w_m p_m^2$ (which can also be written as $u_m = w_m^2 p_m^2$), and $\mathbf{v} = [v_1, v_2, \dots, v_M]^T$ with $v_m = w_m p_m$, we can write the risk function that we

optimize as

$$f(\mathbf{u}, \mathbf{v}, \mathbf{w}) = \sigma^2 \text{tr} \left\{ \left(\sum_{m=1}^M w_m \mathbf{h}_m \mathbf{h}_m^H + u_m \mathbf{h}'_m \mathbf{h}'_m{}^H + v_m (\mathbf{h}'_m \mathbf{h}_m^H + \mathbf{h}_m \mathbf{h}'_m{}^H) \right)^{-1} \right\}, \quad (6.8)$$

where $\mathbf{u} = \mathbf{v}^{\odot 2}$.

In contrast to discrete sparse sensing, the risk function in this case, depends on additional parameters that are related to the perturbation of the discrete sampling points.

6.5 Solver

The optimization variables $(\mathbf{u}, \mathbf{v}, \mathbf{w})$ in (6.8) are related through a structure. Recall that $w_m = 1$ indicates that the m th sample is selected. Only when w_m is nonzero, the corresponding continuous variables u_m and v_m are nonzero¹. In other words, the vectors \mathbf{u} , \mathbf{v} , and \mathbf{w} all share the same support set. Hence, instead of simply minimizing the cardinality of \mathbf{w} as in (P0) that was introduced in Chapter 2, we can exploit the structure and jointly optimize their cardinality to minimize the number of samples, and thus selecting the smallest set of perturbations associated with the discrete sampling points. Defining the matrix $\mathbf{Z} = [\mathbf{u}, \mathbf{v}, \mathbf{w}] \in \mathbb{R}^{M \times 3}$, the proposed continuous sparse sensing problem can be formulated as

$$\arg \min_{\mathbf{Z}} \|\mathbf{Z}\|_{2,0} \quad (6.9a)$$

$$\text{s.to } f(\mathbf{u}, \mathbf{v}, \mathbf{w}) \leq \lambda_c, \quad (6.9b)$$

$$\mathbf{Z} = [\mathbf{u}, \mathbf{v}, \mathbf{w}],$$

$$\mathbf{u} = \mathbf{v}^{\odot 2}, \quad (6.9c)$$

$$w_m \in \{0, 1\}, m = 1, 2, \dots, M, \quad (6.9d)$$

$$-0.5\delta < v_m < 0.5\delta, m = 1, 2, \dots, M, \quad (6.9e)$$

$$0 < u_m < 0.25\delta^2, m = 1, 2, \dots, M, \quad (6.9f)$$

¹We are not interested in a nonzero u_m or v_m when $w_m = 0$.

where the ℓ_2/ℓ_0 -(quasi) norm counts the number of nonzero rows of \mathbf{Z} as follows $\|\mathbf{Z}\|_{2,0} := |\{m : \sqrt{u_m^2 + v_m^2 + w_m^2} \neq 0\}|$, and the convex constraint (6.9b) specifies the estimation accuracy through a threshold λ_c (c denotes continuous). Since the continuous variable p_m takes values in the range $[-0.5\delta, 0.5\delta]$, we obtain the convex box constraints (6.9e) and (6.9f). The optimization problem (6.9) is nonconvex due to: (a) cardinality cost, (b) Boolean constraint (6.9d), and (c) quadratic equality (6.9c). Therefore, it is (in general) difficult to solve (6.9) optimally.

We now use some standard convex relaxation techniques to simplify (6.9) and solve it sub-optimally.

6.5.1 Convex relaxation

The ℓ_2/ℓ_0 -(quasi) norm is relaxed with its best convex approximation, i.e., an ℓ_2/ℓ_1 -mixed norm defined as $\|\mathbf{Z}\|_{2,1} := \sum_{m=1}^M \sqrt{u_m^2 + v_m^2 + w_m^2}$. The Boolean $w_m \in \{0, 1\}$ constraint is replaced with a convex set $w_m \in [0, 1]$. The constraint (6.9c) is equivalently expressed as $\mathbf{u} = \text{diag}(\mathbf{U})$, where $\mathbf{U} = \mathbf{v}\mathbf{v}^H$ is a rank-1 matrix with $[\mathbf{U}]_{i,j \neq i} = 0, \forall i, j$. Dropping the rank constraint on \mathbf{U} and replacing the equality with an inequality as

$$\mathbf{U} \geq \mathbf{v}\mathbf{v}^H \Leftrightarrow \begin{bmatrix} \mathbf{U} & \mathbf{v} \\ \mathbf{v}^H & 1 \end{bmatrix} \geq 0,$$

we arrive at the relaxed continuous sparse sensing problem:

$$\begin{aligned} & \arg \min_{\mathbf{Z}, \mathbf{U}} \|\mathbf{Z}\|_{2,1} \\ & \text{s.to } f(\mathbf{u}, \mathbf{v}, \mathbf{w}) \leq \lambda_c, \\ & \quad \mathbf{Z} = [\mathbf{u}, \mathbf{v}, \mathbf{w}], \\ & \quad \begin{bmatrix} \mathbf{U} & \mathbf{v} \\ \mathbf{v}^H & 1 \end{bmatrix} \geq 0, \\ & \quad \text{diag}(\mathbf{U}) = \mathbf{u}, [\mathbf{U}]_{i,j \neq i} = 0, \forall i, j, \\ & \quad 0 \leq w_m \leq 1, m = 1, 2, \dots, M, \\ & \quad -0.5\delta < v_m < 0.5\delta, m = 1, 2, \dots, M, \\ & \quad 0 < u_m < 0.25\delta^2, m = 1, 2, \dots, M. \end{aligned} \tag{6.10}$$

Subsequently, an approximate Boolean solution for w has to be recovered either by deterministic or randomized rounding as discussed in Chapter 3. Finally, the sensor placements are given by shifting the locations of the selected sensors according to v . The relaxed optimization problem can be solved using off-the-shelf solvers like SeDuMi [Sturm, 1999]. We underline here that the proposed sensor placement is not limited to the initial chosen grid points, and we basically replace grid points with bins allowing one sensor per bin. However, this feature comes at an additional complexity compared to that of solving the sensor selection problem with a fixed discrete grid. The increase in complexity is due to the additional variables like the continuous perturbation parameter and the associated box constraints.

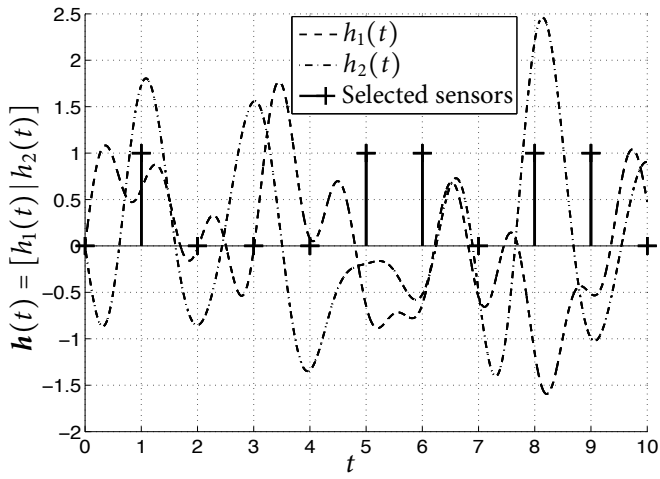
Furthermore, the proposed algorithm is a single step approach. However, it is also possible to compute the initial points from discrete sparse sensing in step-1, and based on these initial points, iteratively solve for the sampling locations using gradient descent techniques to find good local solutions in step-2. But, such a two-step approach can result in a local optimum. Moreover, it will be a complex iterative approach, which can be completely different from the solution of the proposed approach.

6.5.2 Numerical example

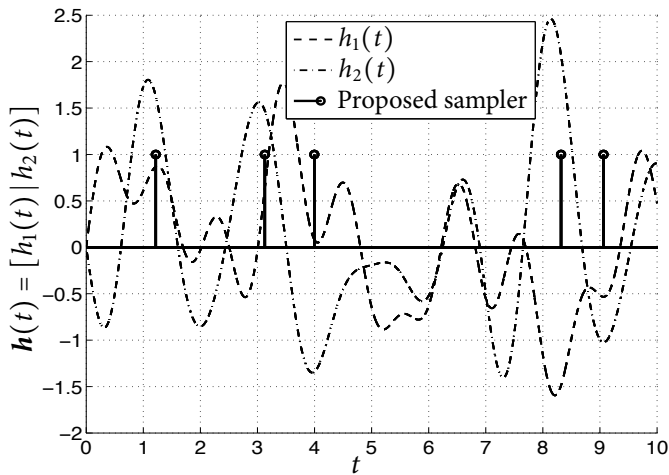
To validate the proposed continuous sparse sensing approach, we refer to the sum of sinusoids example introduced earlier in §6.2. Let the initial coarse grid include $M = 11$ sampling locations $\{t_m = (m - 1)\delta \mid m \in \{1, 2, \dots, 11\}\}$ with $\delta = 1$. Note that the proposed framework is not limited to the sum of sinusoids model, but is valid for any general known model.

Figure. 6.2(a) illustrates the sensor placement via discrete sparse sensing (specifically, sensor selection [Joshi and Boyd, 2009]). The best subset of sensors is computed by solving the relaxed version of (6.3). For sensor selection, we choose $\lambda := 0.47$ such that 5 sensors are selected out of 11 available sensors. For the considered scenario, the mean squared error achieved with such a sensor placement is ≈ 0.47 .

Figure. 6.2(b) illustrates the results from the proposed continuous sensor placement obtained by solving the relaxed optimization problem (6.10). We use $\lambda_c = 0.064$, which has also been chosen such that 5 sensors are selected.



(a)



(b)

Figure 6.2: The field is measured at $M = 11$ discrete locations with $\delta = 1$ and $T = 10$. (a.) Sensor placement via sensor selection. (b.) Proposed continuous sparse sensing. A different threshold λ is used for (a) and (b), such that 5 sensors are selected.

	$M = 5$	$M = 11$	$M = 21$	$M = 41$	$M = 81$
Sensor selection	1.47	0.47	0.28	0.20	0.18
Proposed sampler	1.32	0.36	0.22	0.18	0.17

Table 6.1: Mean squared error with 5 selected sensors for different grid densities.

The mean squared error achieved with the proposed placement is ≈ 0.36 , which is lower than the mean squared error obtained by the sensor placement through sensor selection. The threshold λ_c is an underestimate of the mean squared error (unlike λ), and this is due to the approximation in (6.4). The threshold corresponding to a certain mean squared error can be chosen by computing the solution path for different λ_c values.

Finally, in Table 6.1 we evaluate the mean squared error with 5 selected sensors obtained by solving the sensor selection problem and the proposed continuous placement for different grid densities. The optimization problems are solved in MATLAB using SeDuMi [Sturm, 1999]. We consider different grid densities $M = \{5, 11, 21, 41, 81\}$ and in each case we use a threshold that selects 5 sensors. The continuous sensor placement offers better mean squared error with a reasonable increase in complexity.

6.6 Discussion

We have proposed a framework of continuous sparse sensing in this chapter, where we select sparse sensing patterns from a continuous domain instead of a discrete one. We model an off-the-grid sampling point as an on-the-grid sampling point plus a perturbation assuming that the continuous-domain function is sufficiently smooth. In other words, we can take samples in between the grid points. Expressing the inference quality determining risk as a function of the discrete sampling points and their perturbations, we have designed a continuous sparse sensing operator by solving a convex program.

Chapter 7

Wireless Clock Synchronization

Contents

7.1	Introduction	145
7.2	System model	147
7.3	Passive listening protocol	148
7.4	Estimator	150
7.5	Cramér-Rao lower bound	154
7.6	Simulations	154
7.7	Discussion	156

7.1 Introduction

Sensor networks can be designed to faithfully represent distributed signals (e.g., a spatially varying phenomenon such as a temperature field). In other words, a sensor network can be used as a spatial sampling device. Further, to acquire multidimensional distributed signals that exist in space and time, we also need to perform sampling in time. The temporal sampling is achieved

Part of this chapter was published as: S.P. Chepuri *et al.* Joint Clock Synchronization and Ranging: Asymmetrical Time-Stamping and Passive Listening. *IEEE Signal Processing Letters*, 20 (1): 51 - 54, Jan. 2013.

using analog-to-digital converters or time-to-digital converters, for example. Each sensor has an independent sample clock (i.e., oscillator), and its stability basically determines the alignment of the temporal sampling grid across the sensors. This temporal sampling grid is perfectly aligned if all the sensors share a common clock. When uncommon, the individual clocks drift from each other due to imperfections in the oscillator, aging, and other environmental variations. This drift will result in the misalignment of the temporal sampling grid across the sensors. Therefore, clock synchronization among different nodes each having its own autonomous clock forms a key component of a sensor network.

A plethora of clock synchronization algorithms based on the time-of-flight measurements of the messages have been proposed [Noh et al., 2008, Freris et al., 2010, Wu et al., 2011, Rajan and van der Veen, 2011], which could operate via a two-way time stamp exchange [Wu et al., 2011] or pairwise broadcast synchronization (PBS) [Noh et al., 2008]. Assuming an affine (i.e., first order) clock model and one of the nodes as a reference, and using least squares the unknown clock skews and clock offsets of the remaining nodes in a network can be estimated from time stamps recorded with the two-way time stamp exchange protocol [Rajan and van der Veen, 2011]. Sensor nodes are usually battery powered. Thus, all the tasks of a sensor network, including synchronization, should be carefully performed to ensure longer operating lifetime. For synchronization, this means to minimize the number of transmissions between nodes during which the time stamps are recorded.

In this chapter, we extend the joint clock synchronization and ranging algorithm in [Rajan and van der Veen, 2011] to fully harness the broadcast nature of the wireless medium. By doing so, the number of active transmissions between the nodes can be significantly reduced for a fixed synchronization accuracy. In other words, the synchronization accuracy can be improved for a fixed number of active transmissions. To realize this, we propose an asymmetrical time stamping and passive listening (ATPL) protocol. The ATPL protocol is based on the protocols proposed in [Wang et al., 2011] and [Noh et al., 2008]. The main goal of [Noh et al., 2008] was synchronization and did not focus on ranging (distance between the nodes was assumed to be known). The algorithm proposed in [Wang et al., 2011] also exploits the broadcast property and focused on localization of a target node in an asynchronous network,




Figure 7.1: The ATPL protocol with the i th anchor transmitting. Solid (dotted) lines refer to the active (passive) links. Dark (light) shaded lines refer to the forward (reverse) link.

however, estimation of the clock parameters was not explicitly considered. In the ATPL protocol, during communication between a pair of nodes, time stamps are recorded and exchanged. Besides this, the remaining nodes in the network also passively listen and record the time stamps with their respective clocks, in a cooperative way. However, they do not respond back to either of the active node pair. In addition, the protocol does not put any constraint on the sequence of transmissions, and this together with passive listening results in asymmetrical links, and hence, asymmetrical time stamps. The ATPL protocol is energy efficient in the sense that we obtain more information just by passive listening, and reception usually consumes less power than transmission. For a fully asynchronous network with one sensor and M anchors, we propose a least squares estimator for jointly estimating all the unknown clock skews, clock offsets, and pairwise distances of the sensor to each anchor using the time stamps recorded using the ATPL protocol.

7.2 System model

We consider a fully asynchronous network with M anchors (nodes with known relative locations) and one sensor (*node 0*). We assume that one of the nodes

has a relatively stable clock oscillator and is used as a clock reference. All the other nodes suffer from clock skews and clock offsets. The network model considered here is a special case of the model in [Rajan and van der Veen, 2011], as the pairwise distances of certain nodes (anchors) are now assumed to be known.

The distance between the i th and the j th node is denoted by $d_{i,j} = d_{j,i}$. The distance between the sensor and the i th anchor is denoted by $d_{0,i} = d_{i,0}$, and is unknown. Let t_i be the local time at the i th node and t be the reference time. We assume that the relation between the local time and the reference time can be given by a first order affine clock model [Rajan and van der Veen, 2011]:

$$t_i = \omega_i t + \phi_i \quad \Leftrightarrow \quad t = \alpha_i t_i + \beta_i \quad (7.1)$$

where $\omega_i \in \mathbb{R}_+$ is the clock skew, $\phi_i \in \mathbb{R}$ is the clock offset, $\alpha_i = \omega_i^{-1}$ and $\beta_i = -\omega_i^{-1}\phi_i$ are the synchronization parameters of the i th node. Without loss of generality, we assume that *node M* has a stable clock, i.e., $[\omega_M, \phi_M] = [1, 0]$. The unknown synchronization parameters are collected in

$$\boldsymbol{\alpha} = [\alpha_0, \alpha_1, \dots, \alpha_{M-1}]^T \quad \text{and} \quad \boldsymbol{\beta} = [\beta_0, \beta_1, \dots, \beta_{M-1}]^T.$$

The unknown clock skews and clock offsets are, respectively, given by

$$\boldsymbol{\omega} = \mathbf{1}_M \otimes \boldsymbol{\alpha} \quad \text{and} \quad \boldsymbol{\phi} = -\boldsymbol{\beta} \otimes \boldsymbol{\alpha}. \quad (7.2)$$

The transmission and reception time stamps are recorded both during the forward link (i th active anchor to the sensor) and the reverse link (sensor to the i th active anchor). The time stamp recorded at the i th node when the k th iteration message departs is denoted by $T_i^{(k)}$, and on arrival of the corresponding message, the j th node records the time stamp $R_{i,j}^{(k)}$. Note that the time stamps recorded at the sensor will be either $T_0^{(k)}$ or $R_{i,0}^{(k)}$.

7.3 Passive listening protocol

In the two-way time stamp exchange protocol between the i th anchor and the sensor, the remaining nodes of the network are idle. On the other hand, in the




Figure 7.2: An example sequence of the recorded time stamps. Solid (dotted) lines refer to the active (passive) links. Dark (light) shaded lines refer to the forward (reverse) link.

ATPL protocol, we propose that all the remaining $M - 1$ anchors passively listen to the communication between the i th anchor and the sensor, and record the time stamps $R_{i,j \neq 0}^{(k)}$ of their respective local clocks. By doing so, we obtain more information with extra equations corresponding to transmissions between a) active anchor and other passive anchors, and b) sensor and remaining passive anchors. This is additional to the equations corresponding to the active anchor sensor pair as compared to the two-way time stamp exchange. The ATPL protocol initiated by the i th anchor is illustrated in Figure 7.1. An illustration of the sequence of time stamps recorded during the ATPL protocol is shown in Figure 7.2.

In the proposed protocol, we do not put any constraints on the sequence of forward links and reverse links [Rajan and van der Veen, 2011], i.e., the reverse link need not always follow the forward link as in [Freris et al., 2010, Wu et al., 2011, Noh et al., 2008]. This means that the sensor need not respond to the request from the anchor immediately. Therefore the processing time at the sensor or the network delay typically considered in clock synchronization algorithms [Freris et al., 2010, Wu et al., 2011, Wang et al., 2011, Noh et al., 2008] need not be taken into account as long as the clock parameters are stable within certain reasonable limits.

Remark 7.1. (*Protocol modes*): Possible ways of executing the ATPL protocol

are **mode a**) Anchor node i makes K_i transmissions and the sensor replies back with K_0 messages, with K_0 not necessarily equal to K_i and the transmissions need not be sequential. This is repeated by all the remaining anchors; **mode b**) each anchor node i makes K_i transmissions, and, in the end the sensor replies only once with K_0 messages; and **mode c**) m -out-of- M anchors ($m \leq M$) make transmissions, and the sensor replies as described in either mode a or mode b.

A suitable protocol mode can be adopted depending on the performance requirement and the energy constraint per node.

Remark 7.2. (With or without a central unit): The computation can be done in a centralized way in a fusion center (FC). However, there is an involved communication load in transmitting the time stamps recorded at each node to a FC.

An FC based approach can be avoided by including the time stamps $R_{i,0}^{(k)}$, $k = 1, 2, \dots, K_i$ in the payload when the sensor responds to the i th anchor. However, additional broadcast messages to distribute the time stamps a) $R_{0,i}^{(k)}$, b) $R_{i,j}^{(k)}$ and $R_{0,j}^{(k)}$ are still required. This approach would avoid transmission of the computed unknown parameters to the nodes, that is required with an FC based approach. Moreover, it allows each node to independently perform computations without a FC.

7.4 Estimator

The time-of-flight for a line-of-sight (LOS) transmission between the i th and the j th node can be defined as $\tau_{i,j} = v^{-1}d_{i,j}$, where v denotes the speed of a wave (electromagnetic or acoustic) in a medium. Using (7.1), $\tau_{i,j}$ can be written in terms of the time stamps recorded using respective local clocks of the i th and j th node as

$$\tau_{i,j} = (\alpha_j R_{i,j}^{(k)} + \beta_j) - (\alpha_i T_i^{(k)} + \beta_i) + n_{i,j}^{(k)} \quad (7.3)$$

where $n_{i,j}^{(k)}$ denotes the aggregate measurement error on the time stamps.

The transmission and reception time stamps recorded at the i th and the j th node are, respectively, collected in vectors

$$\mathbf{t}_i = [T_i^{(1)}, T_i^{(2)}, \dots, T_i^{(K_i)}]^T \in \mathbb{R}^{K_i \times 1}$$

and

$$\mathbf{r}_{i,j} = [R_{i,j}^{(1)}, R_{i,j}^{(2)}, \dots, R_{i,j}^{(K_i)}]^T \in \mathbb{R}^{K_i \times 1},$$

where K_i is the number of transmissions made by the i th node. The error vector is denoted by $\mathbf{n}_{i,j} = [n_{i,j}^{(1)}, n_{i,j}^{(2)}, \dots, n_{i,j}^{(K_i)}]^T \in \mathbb{R}^{K_i \times 1}$.

For the sake of exposition, we consider a network with one sensor (*node 0*) and $M = 2$ anchors (*node 1* and *node 2*) and the following example protocol: (i) *node 1* makes K_1 transmissions, *node 0* and *node 2* passively listen, (ii) *node 2* makes K_2 transmissions, *node 0* and *node 1* passively listen, and finally (iii) sensor *node 0* responds with K_0 messages and *node 1* and *node 2* passively listen. This is an example of protocol *mode b* that was described earlier.

Collecting the clock parameters of the i th node in a vector $\mathbf{b}_i = [\alpha_i, \beta_i]^T$, we can now write the equations of the form given in (7.3), obtained for all the $K = K_0 + K_1 + K_2$ time stamps recorded in a matrix-vector form as

$$\begin{array}{l} \text{Node 0 responds} \quad j = 1 \\ \quad (i = 0) \quad \quad \quad j = 2 \\ \text{Node 1 transmits} \quad j = 0 \\ \quad (i = 1) \quad \quad \quad j = 2 \\ \text{Node 2 transmits} \quad j = 0 \\ \quad (i = 2) \quad \quad \quad j = 1 \end{array} \begin{array}{c} \overbrace{\left[\begin{array}{cccc|cccc|cccc} \mathbf{t}_0 & \mathbf{1}_{K_0} & -\mathbf{r}_{0,1} & -\mathbf{1}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} & \mathbf{1}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} \\ \mathbf{t}_0 & \mathbf{1}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} & -\mathbf{r}_{0,2} & -\mathbf{1}_{K_0} & \mathbf{0}_{K_0} & \mathbf{1}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} \\ -\mathbf{r}_{1,0} & -\mathbf{1}_{K_1} & \mathbf{t}_1 & \mathbf{1}_{K_1} & \mathbf{0}_{K_1} & \mathbf{0}_{K_1} & \mathbf{1}_{K_1} & \mathbf{0}_{K_1} & \mathbf{0}_{K_1} & \mathbf{0}_{K_1} \\ \mathbf{0}_{K_1} & \mathbf{0}_{K_1} & \mathbf{t}_1 & \mathbf{1}_{K_1} & -\mathbf{r}_{1,2} & -\mathbf{1}_{K_1} & \mathbf{0}_{K_1} & \mathbf{0}_{K_1} & \mathbf{0}_{K_1} & \mathbf{1}_{K_1} \\ -\mathbf{r}_{2,0} & -\mathbf{1}_{K_2} & \mathbf{0}_{K_2} & \mathbf{0}_{K_2} & \mathbf{t}_2 & \mathbf{1}_{K_2} & \mathbf{0}_{K_2} & \mathbf{1}_{K_2} & \mathbf{0}_{K_2} & \mathbf{0}_{K_2} \\ \mathbf{0}_{K_2} & \mathbf{0}_{K_2} & -\mathbf{r}_{2,1} & -\mathbf{1}_{K_2} & \mathbf{t}_2 & \mathbf{1}_{K_2} & \mathbf{0}_{K_2} & \mathbf{0}_{K_2} & \mathbf{0}_{K_2} & \mathbf{1}_{K_2} \end{array} \right]}^{A \in \mathbb{R}^{2K \times 9}, K=K_0+K_1+K_2} \\ \underbrace{\left[\begin{array}{c} \mathbf{b}_0 \\ \mathbf{b}_1 \\ \mathbf{b}_2 \\ \tau_{0,1} \\ \tau_{0,2} \\ \tau_{1,2} \end{array} \right]}_{\mathbf{n} \in \mathbb{R}^{2K \times 1}} = \left\{ \begin{array}{l} \left[\begin{array}{c} \mathbf{n}_{0,1} \\ \mathbf{n}_{0,2} \end{array} \right] \\ \left[\begin{array}{c} \mathbf{n}_{1,0} \\ \mathbf{n}_{1,2} \end{array} \right] \\ \left[\begin{array}{c} \mathbf{n}_{2,0} \\ \mathbf{n}_{2,1} \end{array} \right] \end{array} \right\} \begin{array}{l} \mathbf{n}_0 \\ \mathbf{n}_1 \\ \mathbf{n}_2 \end{array} \end{array} \quad (7.4)$$

The ordering of the rows of the system matrix \mathbf{A} is arbitrary and does not imply the order of transmission. The columns of \mathbf{A} corresponding to $\tau_{0,1}$, $\tau_{0,2}$, and $\tau_{1,2}$ have two nonzero subvectors each as $\tau_{i,j} = \tau_{j,i}$.

Remark 7.3. (*Rank-deficiency*): The linear model in (7.4) does not have a unique solution, unless we impose certain constraints. Here, we do that by assigning *node 2* as the clock reference, i.e., $\mathbf{b}_2 = [1, 0]^T$.

We define the vector $\boldsymbol{\tau}_0 = [\tau_{0,1}, \tau_{0,2}, \dots, \tau_{0,M}]^T \in \mathbb{R}^{M \times 1}$, where the entries of $\boldsymbol{\tau}_0$ are not known. Note that $\tau_{1,2} = v^{-1}d_{1,2}$ corresponds to the distance between the nodes 1 and 2, and is known. Moving all the knowns (columns corresponding to \mathbf{b}_2 and $\tau_{1,2}$) to one side, (7.4) simplifies to the following linear model given as

$$\begin{array}{l}
 \text{Node 0 responds } j = 1 \\
 \text{--- } (i = 0) \text{--- } j = 2 \\
 \text{Node 1 transmits } j = 0 \\
 \text{--- } (i = 1) \text{--- } j = 2 \\
 \text{Node 2 transmits } j = 0 \\
 \text{--- } (i = 2) \text{--- } j = 1
 \end{array}
 \begin{array}{c}
 \overbrace{\left[\begin{array}{cc|cc|cc}
 \mathbf{t}_0 & \mathbf{1}_{K_0} & -\mathbf{r}_{0,1} & -\mathbf{1}_{K_0} & \mathbf{1}_{K_0} & \mathbf{0}_{K_0} \\
 \mathbf{t}_0 & \mathbf{1}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} & \mathbf{1}_{K_0} \\
 -\mathbf{r}_{1,0} & -\mathbf{1}_{K_1} & \mathbf{t}_1 & \mathbf{1}_{K_1} & \mathbf{1}_{K_1} & \mathbf{0}_{K_1} \\
 \mathbf{0}_{K_1} & \mathbf{0}_{K_1} & \mathbf{t}_1 & \mathbf{1}_{K_1} & \mathbf{0}_{K_1} & \mathbf{0}_{K_1} \\
 -\mathbf{r}_{2,0} & -\mathbf{1}_{K_2} & \mathbf{0}_{K_2} & \mathbf{0}_{K_2} & \mathbf{0}_{K_2} & \mathbf{1}_{K_2} \\
 \mathbf{0}_{K_2} & \mathbf{0}_{K_2} & -\mathbf{r}_{2,1} & -\mathbf{1}_{K_2} & \mathbf{0}_{K_2} & \mathbf{0}_{K_2}
 \end{array} \right]}^{\bar{\mathbf{A}} \in \mathbb{R}^{2K \times 6}}
 \end{array}
 \begin{array}{c}
 \theta \in \mathbb{R}^{6 \times 1} \\
 \left[\begin{array}{c} \mathbf{b}_0 \\ \mathbf{b}_1 \\ \boldsymbol{\tau}_0 \end{array} \right]
 \end{array}
 \\
 \\
 = - \underbrace{\left[\begin{array}{cc|c}
 \mathbf{0}_{K_0} & \mathbf{0}_{K_0} & \mathbf{0}_{K_0} \\
 -\mathbf{r}_{0,2} & -\mathbf{1}_{K_0} & \mathbf{0}_{K_0} \\
 \mathbf{0}_{K_1} & \mathbf{0}_{K_1} & \mathbf{0}_{K_1} \\
 -\mathbf{r}_{1,2} & -\mathbf{1}_{K_1} & \mathbf{1}_{K_1} \\
 \mathbf{t}_2 & \mathbf{1}_{K_2} & \mathbf{0}_{K_2} \\
 \mathbf{t}_2 & \mathbf{1}_{K_2} & \mathbf{1}_{K_2}
 \end{array} \right]}_{\mathbf{x} \in \mathbb{R}^{2K \times 1}}
 \begin{array}{c}
 \left[\begin{array}{c} \mathbf{b}_2 \\ \tau_{1,2} \end{array} \right]
 \end{array}
 + \mathbf{n}.
 \end{array}
 \tag{7.5}$$

The generalization of the data model (7.5) for any $M > 2$ is straightforward and can be easily derived along the same lines. The generalized linear model based on the ATPL protocol is given by

$$\bar{\mathbf{A}}\boldsymbol{\theta} = \mathbf{x} + \mathbf{n} \tag{7.6}$$

where $\bar{\mathbf{A}} \in \mathbb{R}^{KM \times 3M}$, $\boldsymbol{\theta} \in \mathbb{R}^{3M \times 1}$, $\mathbf{x} \in \mathbb{R}^{KM \times 1}$ and $\mathbf{n} \in \mathbb{R}^{KM \times 1}$, all having a similar structure as that of (7.5).

Remark 7.4. (*Correlated error vector*): In case of broadcasting, the entries of the error vector \mathbf{n} are not uncorrelated due to a common error on the transmit time stamp $T_i^{(k)}$.

We assume that the aggregate error $n_{i,j}^{(k)}$ in (7.3) is due to the additive stochastic noise components on the time stamps, $T_i^{(k)}$ denoted by $e_i^{(k)}$ and

the time stamps, $R_{i,j}^{(k)}$ denoted by $e_{i,j}^{(k)}$. We model the aggregate error in (7.3) as

$$\mathbf{n}_{i,j}^{(k)} = \mathbf{e}_i^{(k)} + \mathbf{e}_{i,j}^{(k)} \quad (7.7)$$

where both $e_i^{(k)}$ and $e_{i,j}^{(k)}$ are modeled as zero mean i.i.d. Gaussian [Patwari et al., 2005] with variance $0.5\sigma^2$, such that, $\mathbb{E}\{e_i^{(k)} e_{i,j}^{(k)}\} = 0$ for $i \neq j$. (This is a simplified noise model and more accurate models could be considered.)

We can compute the noise covariance matrix as $\mathbf{\Sigma} = \text{diag}(\mathbf{\Sigma}_0, \mathbf{\Sigma}_1, \dots, \mathbf{\Sigma}_M) \in \mathbb{R}^{MK \times MK}$, where $\mathbf{\Sigma}_i = \mathbb{E}\{\mathbf{n}_i \mathbf{n}_i^T\}$. For $M = 2$, we find

$$\mathbf{\Sigma}_i = \begin{bmatrix} \sigma^2 \mathbf{I}_{K_i} & 0.5\sigma^2 \mathbf{I}_{K_i} \\ 0.5\sigma^2 \mathbf{I}_{K_i} & \sigma^2 \mathbf{I}_{K_i} \end{bmatrix} \in \mathbb{R}^{2K_i \times 2K_i} \quad (7.8)$$

The structure of $\mathbf{\Sigma}$ can be generalized for any $M > 2$ in a similar way, leading to $\mathbf{\Sigma}_i = 0.5\sigma^2(\mathbf{1}_M \mathbf{1}_M^T + \mathbf{I}_M) \otimes \mathbf{I}_{K_i} \in \mathbb{R}^{MK_i \times MK_i}$.

We can now prewhiten the observation model in (7.5) by forming $\mathbf{\Sigma}^{-1/2} \bar{\mathbf{A}}$ and $\mathbf{\Sigma}^{-1/2} \mathbf{x}$. For $K \geq 3$, $\bar{\mathbf{A}}$ is tall and is left-invertible. Hence, the unknown parameters in $\boldsymbol{\theta}$ can be estimated using least squares, i.e.,

$$\hat{\boldsymbol{\theta}} = (\bar{\mathbf{A}}^T \mathbf{\Sigma}^{-1} \bar{\mathbf{A}})^{-1} \bar{\mathbf{A}}^T \mathbf{\Sigma}^{-1} \mathbf{x}. \quad (7.9)$$

Subsequently, the clock skews $\boldsymbol{\omega}$, clock offsets $\boldsymbol{\phi}$ can be obtained using the relation in (7.2), and the pairwise distances of the sensor to each anchor using the relation $\mathbf{d}_0 = \nu \boldsymbol{\tau}_0$.

Remark 7.5. (*Sensor does not respond*):

When only one of the nodes transmits, say node 1, $\bar{\mathbf{A}}$ in (7.5) will not have rows corresponding to transmissions of node 0 and node 2, and it is rank-deficient as the columns two and five are dependent. This also holds when only either node 2 or node 0 transmits.

If only anchor nodes transmit, and the sensor does not respond, then $\bar{\mathbf{A}}$ will not have rows corresponding to transmissions of node 0. In that case, $\bar{\mathbf{A}}$ will be again rank-deficient, as column two is a linear combination of columns five and six. Therefore, for (7.6) to have a unique solution, the sensor should respond at least once, i.e., $K_0 = 1$ with $K \geq 3$.

The possibility that the sensor node responds with only one message in the end makes the protocol energy efficient.

In sum, the proposed algorithm exploits: (1) known distances between anchors, and (2) the broadcast property, which results in additional observations and a correlated error as compared to the pairwise transmissions without passive listening.

7.5 Cramér-Rao lower bound

It follows from the Cramér-Rao lower bound (CRB) theorem that the covariance of any unbiased estimate $\widehat{\bar{\boldsymbol{\theta}}}$ of the unknown parameter $\bar{\boldsymbol{\theta}}$ satisfies the well-known inequality [Kay, 1993]

$$\mathbb{E} \left\{ (\bar{\boldsymbol{\theta}} - \widehat{\bar{\boldsymbol{\theta}}})(\bar{\boldsymbol{\theta}} - \widehat{\bar{\boldsymbol{\theta}}})^T \right\} \geq \mathbf{F}^{-1},$$

where $\bar{\boldsymbol{\theta}} = [\boldsymbol{\omega}^T, \boldsymbol{\phi}^T, \mathbf{d}_0^T]^T$ and $\mathbf{F} \in \mathbb{R}^{3M \times 3M}$ is the Fisher information matrix. If the error \mathbf{n} in (7.6) is zero-mean Gaussian with covariance matrix $\boldsymbol{\Sigma}$, then \mathbf{F} can be computed as $\mathbf{F} = \mathbf{J}^T \boldsymbol{\Sigma}^{-1} \mathbf{J}$, where \mathbf{J} is a Jacobian matrix. The Jacobian matrix is given by

$$\mathbf{J} = \frac{\partial(\bar{\mathbf{A}}\bar{\boldsymbol{\theta}} - \mathbf{x})}{\partial\bar{\boldsymbol{\theta}}^T} = [\mathbf{J}_\omega \quad \mathbf{J}_\phi \quad \mathbf{J}_{d_0}] \in \mathbb{R}^{KM \times 3M} \quad (7.10)$$

with sub-blocks

$$\begin{aligned} \mathbf{J}_\omega &= \frac{\partial(\bar{\mathbf{A}}\bar{\boldsymbol{\theta}} - \mathbf{x})}{\partial\boldsymbol{\omega}^T} = -(\bar{\mathbf{A}}\mathbf{S}_\alpha - \bar{\mathbf{A}}\mathbf{S}_\beta \odot \mathbf{1}_{KM}\boldsymbol{\phi}^T) \oslash (\mathbf{1}_{KM}\boldsymbol{\omega}^T)^{\odot 2}, \\ \mathbf{J}_\phi &= \frac{\partial(\bar{\mathbf{A}}\bar{\boldsymbol{\theta}} - \mathbf{x})}{\partial\boldsymbol{\phi}^T} = -\bar{\mathbf{A}}\mathbf{S}_\beta \oslash \mathbf{1}_{KM}\boldsymbol{\omega}^T, \\ \mathbf{J}_{d_0} &= \frac{\partial(\bar{\mathbf{A}}\bar{\boldsymbol{\theta}} - \mathbf{x})}{\partial\mathbf{d}_0^T} = v^{-1}\bar{\mathbf{A}}\mathbf{S}_{\tau_0} \end{aligned}$$

where \mathbf{S}_α , \mathbf{S}_β , and \mathbf{S}_{τ_0} are selection matrices to select the columns of $\bar{\mathbf{A}}$ corresponding to $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\tau}_0$, respectively.

7.6 Simulations

A network with one sensor and 10 anchors is considered for simulations. Both the sensor and the anchor nodes are deployed randomly within a range of




(a) clock offsets ϕ .(b) clock skews ω .(c) pairwise distances d_0 .

Figure 7.3: RMSE of the estimated unknown parameters.

100 m. Clock skews ω and clock offsets ϕ are uniformly distributed in the range $[1 - 100 \text{ ppm}, 1 + 100 \text{ ppm}]$ and $[-1 \text{ s}, 1 \text{ s}]$, respectively. We use an observation interval of 100 s during which the clock parameters are fixed and $v = 3 \times 10^8 \text{ m/s}$. The time stamps are corrupted with an i.i.d. Gaussian process having a standard deviation $\sigma = 1 \text{ ns}$ [Patwari et al., 2005].

The proposed estimator based on the ATPL protocol is compared with the global least squares (GLS) algorithm proposed in [Rajan and van der Veen, 2011, Figure 3(c)], as it is already shown to outperform other existing synchronization algorithms. We apply the GLS algorithm based on two-way communication between each sensor-anchor pair.

Figure 7.3 shows the root mean square error (RMSE) of the estimates ϕ and ω , and d_0 for different number of messages, K . We show simulations for *mode a* and *mode c* of the ATPL protocol described in §7.3. It can be seen from the figures that the proposed algorithm performs better than GLS in both the considered scenarios due the additional links obtained from passive listening. The proposed algorithm also achieves the theoretical root CRB (RCRB).

7.7 Discussion

In this chapter, we have considered a fully asynchronous network with one sensor and M anchors. We have proposed a least squares estimator to synchronize the sample clocks of wireless sensors based on the ATPL protocol that fully exploits the broadcast nature of the wireless medium. We estimate all the unknown clock skews and clock offsets along with the pairwise distances of the sensor to each anchor. The proposed estimator is shown to be efficient, asymptotically meeting the theoretical CRB, and it outperforms available algorithms. Pairwise distances form a major input to any localization scheme as we see in the next chapter.

Chapter 8

Rigid Body Localization

Contents

8.1	Introduction	158
8.2	Problem formulation and modeling	161
8.3	Linear least squares estimators	166
8.4	Unitarily constrained Cramér-Rao bound	172
8.5	Unitarily constrained total least squares	175
8.6	Simulation results	178
8.7	Discussion	185
8.A	Derivation of the covariance matrix R_n	185
8.B	Gauss-Newton iterations on the Stiefel manifold	187
8.C	Proof of Theorem 8.2	189
8.D	Proof of Theorem 8.3	189
8.E	Proof of Theorem 8.4	191

Part of this chapter was published as: S.P. Chepuri *et al.* Rigid Body Localization Using Sensor Networks. *IEEE Trans. on Signal Processing*, 62(18): 4911-4924, Sept. 2014.

8.1 Introduction

In this chapter, we provide a different flavor of localization, called *rigid body localization*. In rigid body localization, we use a few sensors on a rigid sensing platform and exploit the knowledge of how the sensors are mounted on the body (i.e., sensor placement on the platform) to jointly estimate the position as well as the orientation of the rigid body based on distance measurements.

8.1.1 Applications and prior works

Rigid body localization has potential applications in a variety of fields. To list a few, it is useful for location services involving underwater (or in-liquid) systems, orbiting satellites, mechatronic systems, aircrafts, underwater vehicles, ships, robotic systems, or automobiles. In such applications, classical localization of the node(s) is not sufficient. For example, in an autonomous underwater vehicle (AUV), or an orbiting satellite, the sensing platform is not only subject to motion but also to rotation. Hence, next to position, determining the orientation of the body also forms a key component, and is essential for controlling, maneuvering, and monitoring purposes.

Traditionally, position and orientation are treated separately even though they are closely related. The orientation of a body is usually measured using inertial measurement units (IMUs) comprised of accelerometers [Salhuana, 2012] and gyroscopes. However, IMUs generally suffer from accumulated errors often referred to as drift errors. The drift calibration is typically achieved using different sensor technologies including vision, magnetometers, ultra wide band (UWB), or GPS [Hol et al., 2009, Hol, 2011], leading to dependencies between these technologies. Sometimes these different sensors cannot be coherently fused, for instance magnetometer based calibration needs an undistorted magnetic environment, which is typically difficult to guarantee.

GPS-based attitude determination [Wahba, 1965, Cohen, 1992, Juang and Huang, 1997] is closely related to our work, in which multiple antennas on a platform are used. Here, the attitude is estimated from GPS carrier phase measurements which involves a complicated integer problem with no unique solution in general.

8.1.2 Contributions

We propose a framework for joint position and orientation estimation of a rigid body in a three-dimensional space by borrowing techniques from classical sensor localization, i.e., using only *range* measurements between all the sensor-anchor pairs. We consider a rigid body on which a few sensor nodes are mounted. The absolute position of the rigid body is not known. However, the topology of how the sensors are mounted on the rigid body is known up to a certain accuracy. The orientation of the rigid body is expressed as a *rotation matrix* and the absolute position of the rigid body (instead of the absolute position of each individual sensor) as a *translation vector*. In other words, the absolute position of the sensors is expressed as an affine function of the Stiefel manifold.

The maximum likelihood (ML) estimators for the original problem involve solving a constrained (nonconvex) nonlinear least-squares (NLS) problem, which is in general difficult to solve. In order to simplify this problem, we linearize the problem by squaring the measurements. We use the linearized model in a least-squares (LS) estimator to jointly estimate the rotation matrix (to begin with, its structure is ignored) and the translation vector. Since rotation matrices are unitary matrices, we also propose a unitarily constrained least-squares (CLS) estimator and a simplified unitarily constrained least-squares (SCLS) estimator, both of which solve an optimization problem on the Stiefel manifold. The solutions from the proposed estimators can be used as an initialization to solve the maximum-likelihood estimators or the original nonlinear LS problem if needed. We also derive a unitarily constrained Cramér-Rao bound (CCRB), which is used as a benchmark for the proposed estimators.

In many applications, the sensor topology might not be accurately known, i.e., the known topology can be noisy. These perturbations are typically introduced while mounting the sensors during fabrication or if the body is not entirely rigid. To account for such perturbations, we propose a unitarily constrained total least squares (CTLS) estimator and a simplified unitarily constrained total least squares (SCTLS) estimator. The performance of the proposed estimators is analyzed using simulations. Using a sensor array with a known geometry not only enables orientation estimation, but also yields a




Figure 8.1: An illustration of the sensors on a rigid body undergoing a rotation and translation.

better localization performance. The initial results on rigid body localization using range measurements, viz., SCLS and SCTLs were proposed in [Chepuri et al., 2013a].

The proposed framework of rigid body localization can also be used as an add-on to the existing IMU based systems to correct the drift errors, or in environments where inertial measurements and/or positioning via GPS is not feasible. The proposed framework is based on a static position and orientation, unlike most of the orientation estimators which are based on inertial measurements and a certain dynamical state-space model (e.g., see [Hol et al., 2009]). Hence, our approach is useful when there is no dynamic model available. We should stress, however, that the proposed framework is also suitable for the estimation (tracking) of dynamic position and orientation using a state-constrained Kalman filter for instance, and some initial results on this topic can be found in [Chepuri et al., 2013d].

8.2 Problem formulation and modeling

8.2.1 Problem formulation

Consider a network with M anchors (nodes with known absolute locations) and N sensors in a 3-dimensional space. The sensors are mounted on a rigid body as illustrated in Figure 8.1. The wireless sensors are mounted on the rigid body (e.g., at the factory), and the topology of how these sensors are mounted is known up to a certain accuracy. In other words, we connect a so-called *reference frame* to the rigid body denoted by the axes $\mathbf{q}_1 = [1, 0, 0]^T$, $\mathbf{q}_2 = [0, 1, 0]^T$, and $\mathbf{q}_3 = [0, 0, 1]^T$ in Figure 8.1. In that reference frame, the coordinates of the n th sensor are given by the known 3×1 vector $\mathbf{c}_n = [c_{n,1}, c_{n,2}, c_{n,3}]^T$. The sensor topology is basically determined by the matrix $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_N] \in \mathbb{R}^{3 \times N}$. Let the absolute coordinates of the m th anchor and the n th sensor be denoted by a 3×1 vector \mathbf{a}_m and \mathbf{s}_n , respectively, where \mathbf{s}_n is not known. The absolute positions of the anchors and the sensors are collected in the matrices $\mathbf{A} = [\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_M] \in \mathbb{R}^{3 \times M}$ and $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N] \in \mathbb{R}^{3 \times N}$, respectively.

Rigid body transformation

A Stiefel manifold [Eldén and Park, 1999] in three dimensions, denoted by $\mathcal{V}_{3,3}$, is the set of all 3×3 unitary matrices $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3] \in \mathbb{R}^{3 \times 3}$, i.e.,

$$\mathcal{V}_{3,3} = \{\mathbf{Q} \in \mathbb{R}^{3 \times 3} : \mathbf{Q}^T \mathbf{Q} = \mathbf{Q} \mathbf{Q}^T = \mathbf{I}_3\}.$$

The absolute position of the n th sensor can be written as an affine function of a point on the Stiefel manifold, i.e.,

$$\begin{aligned} \mathbf{s}_n &= c_{n,1} \mathbf{q}_1 + c_{n,2} \mathbf{q}_2 + c_{n,3} \mathbf{q}_3 + \mathbf{t} \\ &= \mathbf{Q} \mathbf{c}_n + \mathbf{t}, \end{aligned} \tag{8.1}$$

where $\mathbf{t} \in \mathbb{R}^{3 \times 1}$ denotes the unknown translation. More specifically, the parameter vector \mathbf{t} refers to the unknown position of the rigid body. The combining weights \mathbf{c}_n are the *known* coordinates of the n th sensor in the reference frame. This means that the unknown unitary matrix \mathbf{Q} actually tells us how

the rigid body has rotated in the reference frame. When there is no rotation, then we have $\mathbf{Q} = \mathbf{I}_3$. The relation in (8.1) is sometimes also referred to as the *rigid body transformation*. The rotation matrices can uniquely represent the orientation of a rigid body unlike Euler angles or unit quaternions (see [Chaturvedi et al., 2011] for more details). The rigid body transformation is also used in computer vision applications [Arun, 1992, Horn et al., 1988, Horn, 1987].

With (8.1), the absolute position of all the sensors can be written as

$$\mathbf{S} = \mathbf{Q}\mathbf{C} + \mathbf{t}\mathbf{1}_N^T = \overbrace{\left[\mathbf{Q} \mid \mathbf{t} \right]}^{\Theta} \overbrace{\begin{bmatrix} \mathbf{C} \\ \mathbf{1}_N^T \end{bmatrix}}^{\mathbf{c}_e}, \quad (8.2)$$

where $\Theta \in \mathbb{R}^{3 \times 4}$ is the unknown transformation matrix.

Range measurements

The proposed framework is general and can be applied to range estimates obtained from any one of the standard ranging techniques (e.g., two-way ranging based on TOA measurements [Patwari et al., 2005, Chepuri et al., 2013c]). The framework is valid as long as the range estimates between all the sensor-anchor pairs are available. Further, we assume that the body position is nearly static during the ranging process, i.e., the linear and angular velocities are negligible compared to the propagation speed.

Let the range (or the Euclidean distance) between the m th anchor and the n th sensor be denoted by $\rho_{mn} = \|\mathbf{a}_m - \mathbf{s}_n\|_2$. The noisy range measurement between the m th anchor and the n th sensor can be expressed as

$$\begin{aligned} y_{mn} &= \|\mathbf{a}_m - \mathbf{s}_n\|_2 + v_{mn} \\ &= \|\mathbf{a}_m - (\mathbf{Q}\mathbf{c}_n + \mathbf{t})\|_2 + v_{mn}, \end{aligned} \quad (8.3)$$

where $v_{mn} \sim \mathcal{N}(0, \sigma_{mn}^2)$ is the stochastic noise resulting from the ranging process. The ranging noise v_{mn} , $m = 1, 2, \dots, M$; $n = 1, 2, \dots, N$, is a sequence of independent random variables whose variance σ_{mn}^2 is assumed to be known or easily estimated.

Problem statement

Having introduced the rigid body transformation in (8.1) and the measurement model (8.3), we can now formally state the rigid body localization problem as follows.

Problem 8.1 (Rigid body localization). *Given one range measurement between each sensor-anchor pair, i.e., y_{mn} as in (8.3), the ranging noise variance σ_{mn}^2 , for $m = 1, 2, \dots, M$ and $n = 1, 2, \dots, N$, the positions of the anchors \mathbf{A} , and the topology of the sensors on the rigid body determined by the matrix \mathbf{C} , jointly estimate the position $\mathbf{t} \in \mathbb{R}^{3 \times 1}$ and orientation $\mathbf{Q} \in \mathcal{V}_{3,3}$ of the rigid body.*

The ML estimator for jointly estimating the orientation and translation is to solve the following optimization problem

$$\begin{aligned} \arg \min_{\mathbf{Q}, \mathbf{t}} \quad & \sum_{m=1}^M \sum_{n=1}^N \sigma_{mn}^{-2} (y_{mn} - \|\mathbf{a}_m - (\mathbf{Q}\mathbf{c}_n + \mathbf{t})\|_2)^2 \\ \text{s.to} \quad & \mathbf{Q}^T \mathbf{Q} = \mathbf{I}_3. \end{aligned} \quad (8.4)$$

The above problem is a nonlinear and a nonconvex optimization problem, and is in general difficult to solve. To simplify this problem, we next linearize the model in (8.3), which can then be solved using linear LS based estimators. The solution from the proposed estimators can then be used as an initialization to solve the above NLS problem if needed.

8.2.2 Squared-range measurements

The model in (8.3) is nonlinear in \mathbf{s}_n (or \mathbf{Q} and \mathbf{t}). Therefore, we linearize the nonlinear model in (8.3) by squaring it. Squaring the measurements in (8.3) results in a noise term with a nonnegative known mean¹ σ_{mn}^2 . Subtracting that mean σ_{mn}^2 from the squared-range measurements between the m th anchor and the n th sensor, we obtain

$$\begin{aligned} d_{mn} &= y_{mn}^2 - \sigma_{mn}^2 \\ &= \|\mathbf{a}_m\|_2^2 - 2\mathbf{a}_m^T \mathbf{s}_n + \|\mathbf{s}_n\|_2^2 + n_{mn}, \end{aligned} \quad (8.5)$$

¹For low noise levels, this nonnegative mean which is simply the variance of the range error in (8.3) can be ignored.

where

$$n_{mn} = 2\rho_{mn}v_{mn} + v_{mn}^2 - \sigma_{mn}^2 \quad (8.6)$$

is the new zero-mean noise term due to squaring. Collecting these new squared-range measurements between the n th sensor and all the anchors in

$$\mathbf{d}_n = [d_{1n}, d_{2n}, \dots, d_{Mn}]^T \in \mathbb{R}^{M \times 1},$$

we can write (8.5) in a vector form as

$$\mathbf{d}_n = \boldsymbol{\alpha} - 2\mathbf{A}^T \mathbf{s}_n + \|\mathbf{s}_n\|_2^2 \mathbf{1}_M + \mathbf{n}_n, \quad (8.7)$$

where

$$\boldsymbol{\alpha} = [\|\mathbf{a}_1\|_2^2, \|\mathbf{a}_2\|_2^2, \dots, \|\mathbf{a}_M\|_2^2]^T \in \mathbb{R}^{M \times 1},$$

is known, and

$$\mathbf{n}_n = [n_{1n}, n_{2n}, \dots, n_{Mn}]^T \in \mathbb{R}^{M \times 1}.$$

Subtracting the knowns in (8.7) from the measurements, we arrive at

$$\mathbf{d}_n - \boldsymbol{\alpha} = -2\mathbf{A}^T \mathbf{s}_n + \|\mathbf{s}_n\|_2^2 \mathbf{1}_M + \mathbf{n}_n. \quad (8.8)$$

We next eliminate the vector $\|\mathbf{s}_n\|_2^2 \mathbf{1}_M$ in (8.8) using an isometry decomposition of the projection matrix

$$\mathbf{P}_M = \mathbf{I}_M - \frac{1}{M} \mathbf{1}_M \mathbf{1}_M^T = \mathbf{U}_M \mathbf{U}_M^T \in \mathbb{R}^{M \times M},$$

where \mathbf{U}_M is an $M \times (M-1)$ matrix obtained by collecting orthonormal basis vectors of the null-space of $\mathbf{1}_M^T$ such that $\mathbf{U}_M^T \mathbf{1}_M = \mathbf{0}_{M-1}$. Pre-multiplying both sides of (8.8) with \mathbf{U}_M^T , we arrive at

$$\mathbf{U}_M^T (\mathbf{d}_n - \boldsymbol{\alpha}) = -2\mathbf{U}_M^T \mathbf{A}^T \mathbf{s}_n + \mathbf{U}_M^T \mathbf{n}_n. \quad (8.9)$$

Stacking (8.9) for all the N sensors, we obtain

$$\mathbf{U}_M^T \mathbf{D} = -2\mathbf{U}_M^T \mathbf{A}^T \mathbf{S} + \mathbf{U}_M^T \mathbf{N}, \quad (8.10)$$

where we define the following $M \times N$ matrices:

$$\mathbf{D} = [\mathbf{d}_1, \mathbf{d}_2, \dots, \mathbf{d}_N] - \boldsymbol{\alpha} \mathbf{1}_N^T,$$

and $\mathbf{N} = [\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_N].$

The linear model in (8.10) can then be compactly expressed as

$$\bar{\mathbf{D}} = \bar{\mathbf{A}}\mathbf{S} + \bar{\mathbf{N}}, \quad (8.11)$$

where we have introduced the following matrices:

$$\begin{aligned} \bar{\mathbf{D}} &= \mathbf{U}_M^T \mathbf{D} \in \mathbb{R}^{(M-1) \times N}, \\ \bar{\mathbf{A}} &= -2\mathbf{U}_M^T \mathbf{A}^T \in \mathbb{R}^{(M-1) \times 3}, \\ \text{and } \bar{\mathbf{N}} &= \mathbf{U}_M^T \mathbf{N} \in \mathbb{R}^{(M-1) \times N}. \end{aligned}$$

Vectorizing¹ (8.11) leads to

$$\bar{\mathbf{d}} = \overbrace{(\mathbf{I}_N \otimes \bar{\mathbf{A}})}^{(M-1)N \times 3N} \mathbf{s} + \bar{\mathbf{n}}, \quad (8.12)$$

where

$$\begin{aligned} \mathbf{s} &= \text{vec}(\mathbf{S}) \in \mathbb{R}^{3N}, \\ \bar{\mathbf{d}} &= \text{vec}(\bar{\mathbf{D}}) \in \mathbb{R}^{(M-1)N}, \\ \text{and } \bar{\mathbf{n}} &= \text{vec}(\bar{\mathbf{N}}) = (\mathbf{I}_N \otimes \mathbf{U}_M^T) \text{vec}(\mathbf{N}) \in \mathbb{R}^{(M-1)N}. \end{aligned}$$

Using the rigid body transformation in (8.2), we can relate the measurements $\bar{\mathbf{D}}$ in (8.11) and the transformation matrix Θ . Substituting (8.2) in (8.11), we arrive at the following linear model

$$\bar{\mathbf{D}} = \bar{\mathbf{A}}\Theta\mathbf{C}_e + \bar{\mathbf{N}}, \quad (8.13)$$

which can then be vectorized to

$$\bar{\mathbf{d}} = \overbrace{(\mathbf{C}_e^T \otimes \bar{\mathbf{A}})}^{(M-1)N \times 9} \boldsymbol{\theta} + \bar{\mathbf{n}}, \quad (8.14)$$

where

$$\boldsymbol{\theta} = \text{vec}(\Theta) = [\mathbf{q}_1^T, \mathbf{q}_2^T, \mathbf{q}_3^T, \mathbf{t}^T]^T \in \mathbb{R}^{12 \times 1}$$

is the unknown parameter vector that has to be estimated.

¹We use the matrix property $\text{vec}(\mathbf{ABC}) = (\mathbf{C}^T \otimes \mathbf{A})\text{vec}(\mathbf{B})$.

The covariance matrix of the noise $\bar{\mathbf{n}}$ in (8.14) is denoted by

$$\mathbf{R}_{\bar{\mathbf{n}}} = (\mathbf{I}_N \otimes \mathbf{U}_M^T) \mathbf{R}_{\mathbf{n}} (\mathbf{I}_N \otimes \mathbf{U}_M) \in \mathbb{R}^{(M-1)N \times (M-1)N},$$

where $\mathbf{R}_{\mathbf{n}}$ is the covariance matrix of $\mathbf{n} = \text{vec}(\mathbf{N})$, and is developed in Appendix 8.A. To whiten¹ the noise, the vectorized model in (8.14) (equivalently the model in (8.12)) is transformed to

$$\begin{aligned} \bar{\mathbf{d}}' &= \mathbf{R}_{\bar{\mathbf{n}}}^{-1/2} \bar{\mathbf{d}} \\ &= \mathbf{R}_{\bar{\mathbf{n}}}^{-1/2} ((\mathbf{I}_N \otimes \bar{\mathbf{A}}) \mathbf{s} + \bar{\mathbf{n}}) \end{aligned} \quad (8.15)$$

$$= \mathbf{R}_{\bar{\mathbf{n}}}^{-1/2} ((\mathbf{C}_e^T \otimes \bar{\mathbf{A}}) \boldsymbol{\theta} + \bar{\mathbf{n}}). \quad (8.16)$$

Here, the notation $\mathbf{R}_{\bar{\mathbf{n}}}^{1/2}$ is defined from the Cholesky decomposition $\mathbf{R}_{\bar{\mathbf{n}}} := \mathbf{R}_{\bar{\mathbf{n}}}^{1/2} \mathbf{R}_{\bar{\mathbf{n}}}^{T/2}$.

In the next section, we propose several estimators of $\boldsymbol{\theta}$ from the processed squared-range measurements $\bar{\mathbf{d}}'$.

8.3 Linear least squares estimators

To begin with, we first look at the topology-agnostic classical LS-based location estimator (i.e., ignoring the prior sensor placement information).

8.3.1 Classical LS-based localization (topology-agnostic)

We use the classical (weighted) LS estimator of \mathbf{s} from $\bar{\mathbf{d}}'$ in (8.15) to estimate the absolute position of the sensors as

$$\begin{aligned} \widehat{\mathbf{s}}_{\text{LS}} &= \arg \min_{\mathbf{s} \in \mathbb{R}^{3N}} \|\bar{\mathbf{d}}' - \mathbf{R}_{\bar{\mathbf{n}}}^{-1/2} (\mathbf{I}_N \otimes \bar{\mathbf{A}}) \mathbf{s}\|_2^2 \\ &= (\mathbf{R}_{\bar{\mathbf{n}}}^{-1/2} (\mathbf{I}_N \otimes \bar{\mathbf{A}}))^\dagger \bar{\mathbf{d}}', \end{aligned} \quad (8.17)$$

¹The noise covariance is parameter dependent, and hence, for whitening it we use an estimated covariance matrix $\widehat{\mathbf{R}}_{\bar{\mathbf{n}}}$ as discussed in §8.6.

which is unique if $\mathbf{I}_N \otimes \bar{\mathbf{A}}$ has full column-rank. This requires $M \geq 4$. Finally, we have

$$\widehat{\mathbf{S}}_{\text{LS}} = \text{unvec}(\widehat{\mathbf{s}}_{\text{LS}}) \in \mathbb{R}^{3 \times N}.$$

In this classical LS-based localization, the knowledge about the known sensor topology is not exploited, and the absolute position of each sensor is estimated separately.

8.3.2 Unconstrained LS estimator

The unknown parameter vector $\boldsymbol{\theta}$ has a structure because $\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \mathbf{q}_3]$ is a unitary matrix. To begin with, we propose to estimate $\boldsymbol{\theta}$, ignoring its structure, from $\bar{\mathbf{d}}'$ in (8.16) using the following (weighted) LS estimator

$$\begin{aligned} \widehat{\boldsymbol{\theta}}_{\text{LS}} &= \arg \min_{\boldsymbol{\theta}} \|\bar{\mathbf{d}}' - \mathbf{R}_{\bar{\mathbf{n}}}^{-1/2} (\mathbf{C}_e^T \otimes \bar{\mathbf{A}}) \boldsymbol{\theta}\|_2^2 \\ &= (\mathbf{R}_{\bar{\mathbf{n}}}^{-1/2} (\mathbf{C}_e^T \otimes \bar{\mathbf{A}}))^\dagger \bar{\mathbf{d}}'. \end{aligned} \quad (8.18)$$

The estimator in (8.18) will have a unique solution if the matrix $\mathbf{C}_e^T \otimes \bar{\mathbf{A}}$ has full column-rank, i.e., \mathbf{C}_e^T and $\bar{\mathbf{A}}$ are both full-column rank, and this requires $(M-1)N \geq 12$. Finally, we have

$$\widehat{\boldsymbol{\Theta}}_{\text{LS}} = \text{unvec}(\widehat{\boldsymbol{\theta}}_{\text{LS}}) = \left[\widehat{\mathbf{Q}}_{\text{LS}} \mid \widehat{\mathbf{t}}_{\text{LS}} \right]. \quad (8.19)$$

8.3.3 Unitarily constrained estimators

The LS estimate $\widehat{\mathbf{Q}}_{\text{LS}}$ obtained in (8.19) is typically (in presence of noise) not a rotation matrix. Hence, we next propose two LS estimators with a unitary constraint on \mathbf{Q} . Both these estimators solve an optimization problem on the Stiefel manifold.

For this purpose, we decouple the rotations and the translations in (8.2) by eliminating the all-one vector $\mathbf{1}_N^T$, and hence the matrix $\mathbf{t}\mathbf{1}_N^T$. In order to eliminate $\mathbf{t}\mathbf{1}_N^T$, we use an isometry matrix \mathbf{U}_N , and as earlier, this matrix is obtained by the isometry decomposition of \mathbf{P}_N , given by

$$\mathbf{P}_N = \mathbf{I}_N - \frac{1}{N} \mathbf{1}_N \mathbf{1}_N^T = \mathbf{U}_N \mathbf{U}_N^T,$$

where \mathbf{U}_N is an $N \times (N - 1)$ matrix obtained by collecting orthonormal basis vectors of the null-space of $\mathbf{1}_N^T$ such that $\mathbf{1}_N^T \mathbf{U}_N = \mathbf{0}_{N-1}^T$. Right-multiplying both sides of (8.2) with \mathbf{U}_N leads to

$$\mathbf{S}\mathbf{U}_N = \mathbf{Q}\mathbf{C}\mathbf{U}_N. \quad (8.20)$$

Combining (8.11) and (8.20) we get the following linear model

$$\tilde{\mathbf{D}}\mathbf{U}_N = \tilde{\mathbf{A}}\mathbf{Q}\mathbf{C}\mathbf{U}_N + \tilde{\mathbf{N}}\mathbf{U}_N,$$

which can be further simplified as

$$\tilde{\mathbf{D}} = \tilde{\mathbf{A}}\mathbf{Q}\tilde{\mathbf{C}} + \tilde{\mathbf{N}}, \quad (8.21)$$

where we have introduced the following matrices:

$$\begin{aligned} \tilde{\mathbf{D}} &= \mathbf{U}_M^T \mathbf{D}\mathbf{U}_N \in \mathbb{R}^{(M-1) \times (N-1)}, \\ \tilde{\mathbf{C}} &= \mathbf{C}\mathbf{U}_N \in \mathbb{R}^{3 \times (N-1)}, \\ \text{and } \tilde{\mathbf{N}} &= \mathbf{U}_M^T \mathbf{N}\mathbf{U}_N \in \mathbb{R}^{(M-1) \times (N-1)}. \end{aligned}$$

Vectorizing (8.21), we obtain

$$\tilde{\mathbf{d}} = \overbrace{(\tilde{\mathbf{C}}^T \otimes \tilde{\mathbf{A}})}^{K \times 9} \mathbf{q} + \tilde{\mathbf{n}}, \quad (8.22)$$

where $K = (M - 1)(N - 1)$, $\tilde{\mathbf{d}} = \text{vec}(\tilde{\mathbf{D}})$, $\mathbf{q} = \text{vec}(\mathbf{Q})$, and $\tilde{\mathbf{n}} = \text{vec}(\tilde{\mathbf{N}})$. The covariance matrix of the noise $\tilde{\mathbf{n}}$ in (8.22) is denoted by

$$\mathbf{R}_{\tilde{\mathbf{n}}} = (\mathbf{U}_N^T \otimes \mathbf{U}_M^T) \mathbf{R}_n (\mathbf{U}_N \otimes \mathbf{U}_M) \in \mathbb{R}^{K \times K}.$$

We will estimate \mathbf{Q} based on a (weighted) LS formulation with a *unitary* constraint, as given by

$$\arg \min_{\mathbf{Q}} \|\mathbf{R}_{\tilde{\mathbf{n}}}^{-1/2} (\tilde{\mathbf{d}} - (\tilde{\mathbf{C}}^T \otimes \tilde{\mathbf{A}}) \mathbf{q})\|_2^2 \quad (8.23a)$$

$$\text{s.to } \mathbf{q} = \text{vec}(\mathbf{Q}), \mathbf{Q}^T \mathbf{Q} = \mathbf{I}_3. \quad (8.23b)$$

The optimization problem in (8.23) is *nonconvex* due to the quadratic equality constraint, and does not generally admit a known closed-form solution. However, such optimization problems can be solved iteratively as will be discussed later on. Before presenting the iterative algorithm, we will first look at a simplified version of (8.23).

Simplified unitarily constrained LS (SCLS) estimator

The optimization problem in (8.23) can be simplified and brought to the standard form of an *orthogonal Procrustes problem* (OPP) with a noniterative known solution. The OPP is generally used to compute rotations between subspaces.

Assuming that $\bar{\mathbf{A}}$ has full column-rank (this can be ensured with optimal anchor placement as discussed in Chapter 3 and [Chepuri et al., 2013b]), and multiplying both sides of (8.21) with $\bar{\mathbf{A}}^\dagger$, we obtain

$$\check{\mathbf{D}} = \mathbf{Q}\bar{\mathbf{C}} + \check{\mathbf{N}}, \quad (8.24)$$

where $\check{\mathbf{D}} := \bar{\mathbf{A}}^\dagger \tilde{\mathbf{D}}$ and $\check{\mathbf{N}} := \bar{\mathbf{A}}^\dagger \tilde{\mathbf{N}}$. The simplified unitarily constrained LS (SCLS) problem is then given as

$$\begin{aligned} \widehat{\mathbf{Q}}_{\text{SCLS}} &= \arg \min_{\mathbf{Q}} \|\check{\mathbf{D}} - \mathbf{Q}\bar{\mathbf{C}}\|_F^2 \\ \text{s.to } &\mathbf{Q}^T \mathbf{Q} = \mathbf{I}_3. \end{aligned} \quad (8.25)$$

The SCLS estimator in (8.25) is *suboptimal* for the problem in (8.23) due to the colored noise $\check{\mathbf{N}}$ in (8.24).

Theorem 8.1 (Solution to the SCLS problem). *The constrained LS problem in (8.25) admits a noniterative known solution given by $\widehat{\mathbf{Q}}_{\text{SCLS}} = \mathbf{V}\mathbf{U}^T$, where \mathbf{V} and \mathbf{U} are obtained from the singular value decomposition (SVD) of $\check{\mathbf{D}}\bar{\mathbf{C}}^T =: \mathbf{V}\boldsymbol{\Sigma}\mathbf{U}^T$ in which matrices $\mathbf{V} \in \mathbb{R}^{3 \times 3}$, $\mathbf{U} \in \mathbb{R}^{3 \times 3}$ are unitary, and $\boldsymbol{\Sigma} \in \mathbb{R}^{3 \times 3}$ is diagonal. The obtained solution is unique if and only if $\check{\mathbf{D}}\bar{\mathbf{C}}^T$ is nonsingular.*

Proof. See [Golub and Van Loan, 1996, pg. 601]. □

Remark 8.1 (Alternative SCLS formulation). *Instead of pseudo inverting $\bar{\mathbf{A}}$ in (8.21) to arrive at (8.24), we can alternatively pseudo-invert $\bar{\mathbf{C}}$ in (8.21) to arrive at another OPP given by*

$$\begin{aligned} \widehat{\mathbf{Q}}_{\text{A-SCLS}} &= \arg \min_{\mathbf{Q}} \|\check{\mathbf{D}} - \bar{\mathbf{A}}\mathbf{Q}\|_F^2 \\ \text{s.to } &\mathbf{Q}^T \mathbf{Q} = \mathbf{I}_3, \end{aligned} \quad (8.26)$$

where $\check{\mathbf{D}} := \check{\mathbf{D}}\check{\mathbf{C}}^\dagger$. The OPP in (8.26) has a closed-form solution $\widehat{\mathbf{Q}}_{\text{A-SCLS}} = \mathbf{U}\mathbf{V}^T$, where the unitary matrices $\mathbf{U} \in \mathbb{R}^{3 \times 3}$ and $\mathbf{V} \in \mathbb{R}^{3 \times 3}$ are obtained from the SVD of $\check{\mathbf{A}}^T \check{\mathbf{D}} =: \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$.

Pseudo inverting $\check{\mathbf{C}}$ can often assure better conditioning as the topology matrix is usually designed at the factory. However, the alternative SCLS formulation in (8.26) cannot be used in case of perturbations on the sensor positions, which is discussed in §8.5. Hence, from now on we will not consider the approach in Remark 8.1.

Subsequently, the SCLS estimate of the translation \mathbf{t} can be computed using $\widehat{\mathbf{Q}}_{\text{SCLS}}$ obtained by solving (8.25). We can write (8.14) equivalently as

$$\check{\mathbf{d}} = \left[(\mathbf{C}^T \otimes \check{\mathbf{A}} \mid (\mathbf{1}_N \otimes \check{\mathbf{A}}) \right] \begin{bmatrix} \mathbf{q} \\ \mathbf{t} \end{bmatrix} + \check{\mathbf{n}}, \quad (8.27)$$

where $(\mathbf{C}^T \otimes \check{\mathbf{A}}) \in \mathbb{R}^{(M-1)N \times 9}$, and $(\mathbf{1}_N \otimes \check{\mathbf{A}}) \in \mathbb{R}^{(M-1)N \times 3}$. Substituting $\widehat{\mathbf{q}}_{\text{SCLS}} := \text{vec}(\widehat{\mathbf{Q}}_{\text{SCLS}})$ in the above model, and moving the knowns to the left side, we get the following linear model

$$\check{\mathbf{d}} - (\mathbf{C}^T \otimes \check{\mathbf{A}})\widehat{\mathbf{q}}_{\text{SCLS}} = (\mathbf{1}_N \otimes \check{\mathbf{A}})\mathbf{t} + \check{\mathbf{n}}.$$

The SCLS estimate for the translations is given by the following LS estimate

$$\begin{aligned} \widehat{\mathbf{t}}_{\text{SCLS}} &= \arg \min_{\mathbf{t}} \left\| \check{\mathbf{d}} - (\mathbf{C}^T \otimes \check{\mathbf{A}})\widehat{\mathbf{q}}_{\text{SCLS}} - (\mathbf{1}_N \otimes \check{\mathbf{A}})\mathbf{t} \right\|_2^2 \\ &= (\mathbf{1}_N \otimes \check{\mathbf{A}})^\dagger (\check{\mathbf{d}} - (\mathbf{C}^T \otimes \check{\mathbf{A}})\widehat{\mathbf{q}}_{\text{SCLS}}). \end{aligned} \quad (8.28)$$

Since we solve the SCLS estimates in an unweighted LS sense, we need not compute the related noise covariance estimates.

Unitarily constrained LS (CLS) estimator

The CLS estimates are obtained by solving the optimization problem that was introduced earlier in (8.23), which we recall as

$$\widehat{\mathbf{Q}}_{\text{CLS}} = \arg \min_{\mathbf{Q}} \left\| \mathbf{R}_{\check{\mathbf{n}}}^{-1/2} (\check{\mathbf{d}} - (\check{\mathbf{C}}^T \otimes \check{\mathbf{A}})\mathbf{q}) \right\|_2^2 \quad (8.29a)$$

$$\text{s.to } \mathbf{q} = \text{vec}(\mathbf{Q}), \mathbf{Q}^T \mathbf{Q} = \mathbf{I}_3. \quad (8.29b)$$

The optimization problem in (8.29) is a generalization of the OPP, and is sometimes also referred to as the *weighted orthogonal Procrustes problem* (WOPP) [Viklands, 2008]. Unlike the OPP of (8.25), which has a closed-form analytical solution, the optimization problem (8.29) does not admit a closed-form solution. However, it can be solved using iterative methods based on either Gauss-Newton's method, Newton's method [Viklands, 2008], or steepest descent [Manton, 2002]. In this chapter, we restrict ourselves to Gauss-Newton's method for solving (8.29) because of the availability of a good initial value (e.g., the closed-form solution from Theorem 8.1) for the iterative algorithm.

The optimization problem in (8.29) is a LS problem on the Stiefel manifold. To simplify the notations we write (8.29) in a more general form:

$$\begin{aligned} \arg \min_{\mathbf{Q}} \|\mathbf{z} - \text{Lvec}(\mathbf{Q})\|_2^2 \\ \text{s.to } \mathbf{Q} \in \mathcal{V}_{3,3}, \end{aligned} \quad (8.30)$$

Algorithm 8.1 CLS based on Gauss-Newton's method.

1. **Compute** initial value $\mathbf{Q}_0 := \widehat{\mathbf{Q}}_{\text{SCLS}}$,
2. **initialize** \mathbf{L} , \mathbf{z} , iteration counter $k = 0$, ε , and ε_0 .
3. **while** $\varepsilon_k > \varepsilon$
4. **compute** \mathbf{J} using (8.48)
5. **compute** a *Gauss-Newton* step

$$\mathbf{x}_k = (\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T (\mathbf{z} - \text{Lvec}(\mathbf{Q}_k))$$

6. **compute** the optimal step-length α_k using (8.50).
 7. **update** $\mathbf{Q}_{k+1} = \mathbf{Q}_k \mathbf{\Omega}(\alpha_k \mathbf{x}_k)$,
 where $\mathbf{\Omega}(\cdot)$ is defined in (8.46).
 8. **increment** $k = k + 1$.
 9. **compute** $\varepsilon_{k+1} = \frac{\|\mathbf{J}^T (\mathbf{z} - \text{Lvec}(\mathbf{Q}_k))\|_2}{\|\mathbf{J}\|_F \|\mathbf{z} - \text{Lvec}(\mathbf{Q}_k)\|_2}$.
 10. **end while**.
-

and for solving (8.29) we use

$$\begin{aligned} \mathbf{z} &:= \mathbf{R}_{\tilde{\mathbf{n}}}^{-1/2} \tilde{\mathbf{d}} \in \mathbb{R}^{K \times 1}, \\ \text{and } \mathbf{L} &:= \mathbf{R}_{\tilde{\mathbf{n}}}^{-1/2} (\tilde{\mathbf{C}}^T \otimes \tilde{\mathbf{A}}) \in \mathbb{R}^{K \times 9}. \end{aligned} \quad (8.31)$$

The algorithm is derived in Appendix 8.B, and it is summarized as Algorithm 8.1. A more profound treatment on WOPP is found in [Viklands, 2008].

As before, the estimate for the translation \mathbf{t} can then be computed using $\widehat{\mathbf{q}}_{\text{CLS}} := \text{vec}(\widehat{\mathbf{Q}}_{\text{CLS}})$ in (8.27). The CLS estimate for the translation is given by the following (weighted) LS estimate

$$\begin{aligned} \widehat{\mathbf{t}}_{\text{CLS}} &= \arg \min_{\mathbf{t}} \|\tilde{\mathbf{d}}' - \mathbf{R}_{\tilde{\mathbf{n}}}^{-1/2} ((\mathbf{C}^T \otimes \tilde{\mathbf{A}}) \widehat{\mathbf{q}}_{\text{CLS}} + (\mathbf{1}_N \otimes \tilde{\mathbf{A}}) \mathbf{t})\|_2^2 \\ &= (\mathbf{R}_{\tilde{\mathbf{n}}}^{-1/2} (\mathbf{1}_N \otimes \tilde{\mathbf{A}}))^\dagger (\tilde{\mathbf{d}}' - \mathbf{R}_{\tilde{\mathbf{n}}}^{-1/2} (\mathbf{C}^T \otimes \tilde{\mathbf{A}}) \widehat{\mathbf{q}}_{\text{CLS}}). \end{aligned}$$

8.3.4 Topology-aware (TA) localization

A complementary by-product of the rigid body localization is the *topology-aware* localization. In this case, the position and orientation estimation is not the main interest, but the absolute position of each sensor node has to be estimated, given that the sensors follow a certain topology. This latter information can be used as a constraint for estimating the sensor positions rather than estimating them separately. For the rigidity constraint, using $\widehat{\mathbf{Q}}$ and $\widehat{\mathbf{t}}$ obtained from either the SCLS or CLS estimator, we can then compute the absolute positions of each sensor on the rigid body as

$$\widehat{\mathbf{S}}_{\text{TA}} = \widehat{\mathbf{Q}}\mathbf{C} + \widehat{\mathbf{t}}\mathbf{1}_N^T. \quad (8.32)$$

8.4 Unitarily constrained Cramér-Rao bound

Suppose we want to estimate the unknown vector $\boldsymbol{\theta} = [\mathbf{q}_1^T, \mathbf{q}_2^T, \mathbf{q}_3^T, \mathbf{t}^T]^T \in \mathbb{R}^{12 \times 1}$ from the range measurements y_{mn} corrupted by independent noise $v_{mn} \sim \mathcal{N}(0, \sigma_{mn}^2)$ for $n = 1, 2, \dots, N$, and $m = 1, 2, \dots, M$, where the observations follow the nonlinear model (8.3). We can compute the CRB for any unbiased estimator of $\boldsymbol{\theta}$ as described next.

8.4.1 Unconstrained CRB

The covariance matrix of any unbiased estimate of the parameter vector $\boldsymbol{\theta}$ satisfies [Kay, 1993]

$$\mathbb{E}\{(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})^T\} \geq \mathbf{C}_{\text{CRB}} = \mathbf{F}^{-1},$$

where the Fisher information matrix (FIM) \mathbf{F} is given by

$$\mathbf{F} = \sum_{m=1}^M \sum_{n=1}^N \mathbb{E} \left\{ \left(\frac{\partial \ln p(y_{mn}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right) \left(\frac{\partial \ln p(y_{mn}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^T \right\}.$$

This is the *Cramér-Rao bound theorem* and \mathbf{C}_{CRB} is the Cramér-Rao lower bound (CRB) matrix. Let us define $\mathbf{c}_{e,n} = [\mathbf{c}_n^T, 1]^T \in \mathbb{R}^4$ for $n = 1, 2, \dots, N$. The computation of the FIM for observations with Gaussian likelihoods is straightforward, and the 12×12 FIM is given as

$$\mathbf{F}(\boldsymbol{\theta}) = \sum_{m=1}^M \sum_{n=1}^N \frac{(\mathbf{c}_{e,n} \otimes \mathbf{I}_3)(\mathbf{a}_m - (\mathbf{Q}\mathbf{c}_n + \mathbf{t}))(\mathbf{a}_m - (\mathbf{Q}\mathbf{c}_n + \mathbf{t}))^T (\mathbf{c}_{e,n}^T \otimes \mathbf{I}_3)}{\sigma_{mn}^2 \|\mathbf{a}_m - (\mathbf{Q}\mathbf{c}_n + \mathbf{t})\|_2^2}. \quad (8.33)$$

8.4.2 Constrained CRB

The FIM in (8.33), does not take into account the unitary constraint on the matrix \mathbf{Q} , i.e., $\mathbf{Q}^T \mathbf{Q} = \mathbf{I}_3$. Generally, if the parameter vector $\boldsymbol{\theta}$ is subject to some P continuously differentiable (nonredundant) constraints $\mathbf{h}(\boldsymbol{\theta}) = \mathbf{0}$, then with these constraints, the resulting constrained CRB is lower than the unconstrained CRB. In [Stoica and Ng, 1998], it is shown that the constrained CRB (CCRB) has the form

$$\mathbb{E}\{(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})^T\} \geq \mathbf{C}_{\text{CCRB}}(\boldsymbol{\theta}) = \mathbf{M}(\mathbf{M}^T \mathbf{F} \mathbf{M})^{-1} \mathbf{M}^T, \quad (8.33)$$

where \mathbf{F} is the FIM for the unconstrained estimation problem as in (8.33), and an isometry matrix $\mathbf{M} \in \mathbb{R}^{12 \times (12-P)}$ is obtained by collecting orthonormal basis vectors of the null-space of the gradient matrix

$$\mathbf{G}(\boldsymbol{\theta}) = \frac{\partial \mathbf{h}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T} \in \mathbb{R}^{P \times 12}.$$

The nonredundant constraints ensure that the matrix $\mathbf{G}(\boldsymbol{\theta})$ is full row-rank, and implies

$$\mathbf{G}(\boldsymbol{\theta})\mathbf{M} = \mathbf{0}$$

while $\mathbf{M}^T\mathbf{M} = \mathbf{I}_{12-P}$. For the unitarily constrained CRB (CCRB) denoted by $\mathbf{C}_{\text{CCRB}}(\boldsymbol{\theta})$, we have to consider the unitary constraint $\mathbf{Q}^T\mathbf{Q} = \mathbf{I}_3$, which can be expressed by the following $P = 6$ nonredundant constraints

$$\mathbf{h}(\boldsymbol{\theta}) = [\mathbf{q}_1^T\mathbf{q}_1 - 1, \mathbf{q}_2^T\mathbf{q}_1, \mathbf{q}_3^T\mathbf{q}_1, \mathbf{q}_2^T\mathbf{q}_2 - 1, \mathbf{q}_3^T\mathbf{q}_2, \mathbf{q}_3^T\mathbf{q}_3 - 1]^T = \mathbf{0}_6 \in \mathbb{R}^{6 \times 1}, \quad (8.34)$$

where the symmetric redundant constraints are discarded. The gradient matrix for the constraints in (8.34) can be computed as follows

$$\mathbf{G}(\boldsymbol{\theta}) = \frac{\partial \mathbf{h}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}^T} = \begin{bmatrix} 2\mathbf{q}_1^T & \mathbf{0}_3^T & \mathbf{0}_3^T & \mathbf{0}_3^T \\ \mathbf{q}_2^T & \mathbf{q}_1^T & \mathbf{0}_3^T & \mathbf{0}_3^T \\ \mathbf{q}_3^T & \mathbf{0}_3^T & \mathbf{q}_1^T & \mathbf{0}_3^T \\ \mathbf{0}_3^T & 2\mathbf{q}_2^T & \mathbf{0}_3^T & \mathbf{0}_3^T \\ \mathbf{0}_3^T & \mathbf{q}_3^T & \mathbf{q}_2^T & \mathbf{0}_3^T \\ \mathbf{0}_3^T & \mathbf{0}_3^T & 2\mathbf{q}_3^T & \mathbf{0}_3^T \end{bmatrix} \in \mathbb{R}^{6 \times 12}.$$

An orthonormal basis of the null-space of the gradient matrix is finally given by

$$\mathbf{M} = \frac{1}{\sqrt{2}} \left[\begin{array}{ccc|c} -\mathbf{q}_3 & \mathbf{0}_3 & \mathbf{q}_2 & \mathbf{0}_{9 \times 3} \\ \mathbf{0}_3 & -\mathbf{q}_3 & -\mathbf{q}_1 & \\ \mathbf{q}_1 & \mathbf{q}_2 & \mathbf{0}_3 & \\ \hline & \mathbf{0}_{3 \times 3} & & \sqrt{2}\mathbf{I}_3 \end{array} \right].$$

The CCRB for TA-localization can be easily derived from $\mathbf{C}_{\text{CCRB}}(\boldsymbol{\theta})$ using the transformation of parameters [Kay, 1993].

Theorem 8.2 (Biased estimator). *An unbiased constrained estimator for \mathbf{Q} does not exist, except for the noiseless case.*

Proof. See Appendix 8.C. □

Due to Theorem 8.2, the mean-squared-error (MSE) of any unitarily constrained estimator will be lower than the CCRB for high noise levels. However, at low noise levels, the bias tends to zero, and the CCRB gives a good lower bound on the MSE of the unitarily constrained estimators.

8.5 Unitarily constrained total least squares

In the previous section, we assumed that the position of the sensors on the rigid body in the reference frame, i.e., the matrix \mathbf{C} , is accurately known. In practice, there is no reason to believe that errors are restricted only to the range measurements and there are no perturbations on the initial sensor positions. Such perturbations can be introduced for instance during fabrication or if the body is not entirely rigid (e.g., wing flexing of an aircraft).

Let us assume that the position of the n th sensor in the reference frame \mathbf{c}_n is noisy, and let us denote the perturbation on \mathbf{c}_n by \mathbf{e}_n , and the perturbation on $\mathbf{C} = [\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_N]$ by the matrix $\mathbf{E} = [\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_N] \in \mathbb{R}^{3 \times N}$. The covariance matrix of the perturbation $\mathbf{e} = \text{vec}(\mathbf{E})$ is denoted by $\mathbf{R}_e = \sigma_e^2 \mathbf{I}_{3N}$. In other words, we assume that the perturbations $\mathbf{e}_n, n = 1, 2, \dots, N$, are a sequence of independent and identically distributed (i.i.d.) random vectors. To account for such errors in the model, we next propose total least squares (TLS) estimators again with unitary constraints.

8.5.1 Simplified unitarily constrained TLS (SCTLS) estimator

Taking the perturbations on the known topology into account, the data model in (8.21) will be modified as

$$\tilde{\mathbf{D}} = \bar{\mathbf{A}}\mathbf{Q}(\bar{\mathbf{C}} + \bar{\mathbf{E}}) + \tilde{\mathbf{N}},$$

where $\bar{\mathbf{E}} = \mathbf{E}\mathbf{U}_N$. Multiplying both sides of the above equation with $\bar{\mathbf{A}}^\dagger$, we get

$$\check{\mathbf{D}} = \mathbf{Q}(\bar{\mathbf{C}} + \bar{\mathbf{E}}) + \check{\mathbf{N}}. \quad (8.35)$$

The solution to the data model in (8.35) admits a classical TLS formulation, but with a unitary constraint. The SCTLS optimization problem is given by

$$\begin{aligned} \widehat{\mathbf{Q}}_{\text{SCTLS}} &= \arg \min_{\mathbf{Q}} \|\bar{\mathbf{E}}\|_F^2 + \|\check{\mathbf{N}}\|_F^2 \\ \text{s.to } \check{\mathbf{D}} &= \mathbf{Q}(\bar{\mathbf{C}} + \bar{\mathbf{E}}) + \check{\mathbf{N}}, \text{ and } \mathbf{Q}^T \mathbf{Q} = \mathbf{I}_3. \end{aligned} \quad (8.36)$$

Similar to SCLS, the optimization problem in (8.36) admits a known closed-form solution which makes it simplified compared to the weighted problem discussed in the next section. Also, the noise $\check{\mathbf{N}}$ in (8.35) is not white leading to a suboptimal solution.

Theorem 8.3 (Solution to SCTLS). *The SCTLS problem in (8.36) has the same solution as the simplified unitarily constrained LS problem, i.e., $\widehat{\mathbf{Q}}_{\text{SCTLS}} = \widehat{\mathbf{Q}}_{\text{SCLS}}$.*

Proof. See Appendix 8.D. □

We next estimate the translation vector \mathbf{t} . Taking the perturbations into account, we can modify the model in (8.13) as

$$\check{\mathbf{D}} = \bar{\mathbf{A}}\mathbf{Q}\mathbf{C} + \bar{\mathbf{A}}\mathbf{t}\mathbf{1}_N^T + \bar{\mathbf{A}}\mathbf{Q}\mathbf{E} + \check{\mathbf{N}},$$

which can be vectorized and further simplified to

$$\bar{\mathbf{d}} - (\mathbf{C}^T \otimes \bar{\mathbf{A}})\mathbf{q} = (\mathbf{1}_N \otimes \bar{\mathbf{A}})\mathbf{t} + \overbrace{(\mathbf{I} \otimes \bar{\mathbf{A}}\mathbf{Q})\mathbf{e}}^{\mathbf{v} \in \mathbb{R}^{(M-1)N}} + \check{\mathbf{n}}. \quad (8.37)$$

Using $\widehat{\mathbf{q}}_{\text{SCLS}}$ in the above model, the SCTLS estimator for the translation taking perturbations into account, is then given by the following (unweighted) LS problem

$$\begin{aligned} \widehat{\mathbf{t}}_{\text{SCLS}} &= \arg \min_{\mathbf{t}} \|\bar{\mathbf{d}} - (\mathbf{C}^T \otimes \bar{\mathbf{A}})\widehat{\mathbf{q}}_{\text{SCLS}} - (\mathbf{1}_N \otimes \bar{\mathbf{A}})\mathbf{t}\|_2^2 \\ &= (\mathbf{1}_N \otimes \bar{\mathbf{A}})^\dagger (\bar{\mathbf{d}} - (\mathbf{C}^T \otimes \bar{\mathbf{A}})\widehat{\mathbf{q}}_{\text{SCLS}}). \end{aligned}$$

The algorithms to compute the SCTLS estimates are the same as that of the SCLS estimator. As before, we solve SCTLS in an unweighted LS sense, therefore, we need not estimate the related noise covariance matrix for whitening.

8.5.2 Unitarily constrained TLS (CTLS) estimator

Similar to the CLS formulation, the TLS estimator can be derived without pseudo inverting the matrix $\bar{\mathbf{A}}$. The data model in (8.21) taking into account the perturbations on the known sensor topology is then given by

$$\tilde{\mathbf{D}} = \bar{\mathbf{A}}\mathbf{Q}(\tilde{\mathbf{C}} + \tilde{\mathbf{E}}) + \tilde{\mathbf{N}}, \quad (8.38)$$

which can be further vectorized as

$$\tilde{\mathbf{d}} = (\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\mathbf{Q})\tilde{\mathbf{c}} + (\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\mathbf{Q})\tilde{\mathbf{e}} + \tilde{\mathbf{n}}, \quad (8.39)$$

where $\tilde{\mathbf{c}} = \text{vec}(\tilde{\mathbf{C}}) \in \mathbb{R}^{3(N-1)}$, and $\tilde{\mathbf{e}} = \text{vec}(\tilde{\mathbf{E}}) \in \mathbb{R}^{3(N-1)}$.

Assuming that the pre-whitening matrix takes the block diagonal form

$$\mathbf{R}_k^{-1/2} := \text{diag}(\sigma_e^{-1}\mathbf{I}_{3(N-1)}, \mathbf{R}_{\tilde{\mathbf{n}}}^{-1/2})$$

with $\mathbf{k} = [\tilde{\mathbf{e}}^T, \tilde{\mathbf{n}}^T]^T$, the CTLS optimization problem is given by

$$\begin{aligned} \arg \min_{\mathbf{Q}} \quad & \|\sigma_e^{-1}\tilde{\mathbf{e}}\|_2^2 + \|\mathbf{R}_{\tilde{\mathbf{n}}}^{-1/2}\tilde{\mathbf{n}}\|_2^2 \\ \text{s.to} \quad & \tilde{\mathbf{D}} = \bar{\mathbf{A}}\mathbf{Q}(\tilde{\mathbf{C}} + \tilde{\mathbf{E}}) + \tilde{\mathbf{N}}, \\ & \tilde{\mathbf{e}} = \text{vec}(\tilde{\mathbf{E}}), \quad \tilde{\mathbf{n}} = \text{vec}(\tilde{\mathbf{N}}), \\ & \mathbf{Q}^T\mathbf{Q} = \mathbf{I}_3. \end{aligned} \quad (8.40)$$

Theorem 8.4 (Solution to CTLS). *Assuming that the covariance matrix of the perturbation vector is a scaled identity matrix, the unitarily constrained TLS problem in (8.40) has the same solution as a specifically weighted CLS, i.e., it is the solution to*

$$\begin{aligned} \hat{\mathbf{Q}}_{\text{CTLS}} = \arg \min_{\mathbf{Q}} \quad & \|\Lambda^{1/2}(\tilde{\mathbf{d}} - (\tilde{\mathbf{C}}^T \otimes \bar{\mathbf{A}})\mathbf{q})\|_2^2 \\ \text{s.to} \quad & \mathbf{Q}^T\mathbf{Q} = \mathbf{I}_3, \end{aligned} \quad (8.41)$$

where matrix $(\tilde{\mathbf{C}}^T \otimes \bar{\mathbf{A}}) \in \mathbb{R}^{K \times 9}$ was defined earlier, and $\Lambda \in \mathbb{R}^{K \times K}$ is the specific weighting matrix defined in (8.57).

Proof. See Appendix 8.E. □

The optimization problem (8.41) does not have a closed-form solution, and has to be solved iteratively using for instance Gauss-Newton's method (summarized in Algorithm 8.1) with

$$\begin{aligned} \mathbf{z} &:= \mathbf{\Lambda}^{1/2} \tilde{\mathbf{d}} \in \mathbb{R}^{K \times 1}, \\ \text{and } \mathbf{L} &:= \mathbf{\Lambda}^{1/2} (\tilde{\mathbf{C}}^T \otimes \tilde{\mathbf{A}}) \in \mathbb{R}^{K \times 9}. \end{aligned} \quad (8.42)$$

The CTLS estimates for the translations can be computed similar to SCTLS as earlier, however, for CTLS we first pre-whiten the noise. The covariance of the noise \mathbf{v} in (8.37) is denoted by an $(M-1)N \times (M-1)N$ matrix $\mathbf{R}_v := \sigma_c^2 (\mathbf{I} \otimes \tilde{\mathbf{A}} \tilde{\mathbf{A}}^T) + \mathbf{R}_{\tilde{\mathbf{h}}}$. We can then use a weighted LS estimator to estimate the translations accounting for the perturbations. The CTLS translation estimates are given by

$$\begin{aligned} \hat{\mathbf{t}}_{\text{CTLS}} &= \arg \min_{\mathbf{t}} \|\mathbf{R}_v^{-1/2} (\tilde{\mathbf{d}} - (\mathbf{C}^T \otimes \tilde{\mathbf{A}}) \hat{\mathbf{q}}_{\text{CTLS}} - (\mathbf{1}_N \otimes \tilde{\mathbf{A}}) \mathbf{t})\|_2^2 \\ &= (\mathbf{R}_v^{-1/2} (\mathbf{1}_N \otimes \tilde{\mathbf{A}}))^\dagger \mathbf{R}_v^{-1/2} (\tilde{\mathbf{d}} - (\mathbf{C}^T \otimes \tilde{\mathbf{A}}) \hat{\mathbf{q}}_{\text{CTLS}}). \end{aligned}$$

8.6 Simulation results

We consider $N = 5$ sensors mounted on the vertices of a rigid body (rectangle based pyramid as in Figure 8.1) with

$$\mathbf{C} = \begin{bmatrix} 0.5 & 1.5 & 1.5 & 0.5 & 1 \\ 0 & 0 & 1.5 & 1.5 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \text{ m,}$$

and $M = 4$ anchors fixed at location

$$\mathbf{A} = \begin{bmatrix} 0 & 100 & 0 & 100 \\ 100 & 100 & 0 & 0 \\ 0 & 100 & 100 & 0 \end{bmatrix} \text{ m.}$$

Let us define a function $\mathcal{R}(\cdot) : \mathbb{R}^3 \rightarrow \mathcal{V}_{3,3}$ that maps angles in degrees along each dimension into a rotation matrix, and its inverse $\mathcal{R}^{-1}(\cdot) : \mathcal{V}_{3,3} \rightarrow \mathbb{R}^3$

which maps a rotation matrix into corresponding angles in degrees (for details on converting angles to a rotation matrix and vice versa, see [Diebel, 2006]). Collecting the angles in a 3×1 vector $\boldsymbol{\phi} = [20, -25, 10]^T$ deg, the rotation matrix is then generated according to $\mathbf{Q} = \mathcal{R}(\boldsymbol{\phi})$. We use a translation vector $\mathbf{t} = [15, 5, 10]^T$ m. The simulations are averaged over $N_{\text{exp}} = 1000$ independent Monte-Carlo experiments.

The performance of the proposed estimators is analyzed in terms of the root mean squared error (RMSE) of the estimates of \mathbf{Q} , $\boldsymbol{\phi}$, and \mathbf{t} , and are respectively given as

$$\begin{aligned} \text{RMSE}(\mathbf{Q}) &= \left(\frac{1}{N_{\text{exp}}} \sum_{n=1}^{N_{\text{exp}}} \|\mathbf{Q} - \widehat{\mathbf{Q}}^{(n)}\|_F^2 \right)^{1/2}, \\ \text{RMSE}(\boldsymbol{\phi}) &= \left(\frac{1}{N_{\text{exp}}} \sum_{n=1}^{N_{\text{exp}}} \|\boldsymbol{\phi} - \mathcal{R}^{-1}(\widehat{\mathbf{Q}}^{(n)})\|_2^2 \right)^{1/2} \text{ deg}, \\ \text{RMSE}(\mathbf{t}) &= \left(\frac{1}{N_{\text{exp}}} \sum_{n=1}^{N_{\text{exp}}} \|\mathbf{t} - \widehat{\mathbf{t}}^{(n)}\|_2^2 \right)^{1/2} \text{ m}, \end{aligned}$$

where $\widehat{\mathbf{Q}}^{(n)}$ and $\widehat{\mathbf{t}}^{(n)}$ denote the estimates during the n th Monte-Carlo experiment. Since the rotation matrix estimated using the unconstrained LS estimator is not a valid rotation matrix, we first project it onto $\mathcal{V}_{3,3}$ using (8.51) before converting it into corresponding angles.

We use the same noise variance for all the range measurements, i.e., $\sigma_{mn} = \sigma$ m for $m = 1, 2, \dots, M$, and $n = 1, 2, \dots, N$. The covariance matrix of the noise \mathbf{n} , i.e., (see Appendix 8.A)

$$\mathbf{R}_n = \text{diag}(2\rho)\mathbf{R}_v\text{diag}(2\rho) + \mathbf{R}_{v \odot 2} - \boldsymbol{\mu}\boldsymbol{\mu}^T,$$

depends on the unknown parameter

$$\boldsymbol{\rho} = [\rho_{11}, \rho_{12}, \dots, \rho_{M1}, \dots, \rho_{MN}]^T \in \mathbb{R}^{MN}.$$

Hence, to whiten it, we use the noisy range measurements

$$\mathbf{y} = [y_{11}, y_{12}, \dots, y_{M1}, \dots, y_{MN}]^T \in \mathbb{R}^{MN}$$

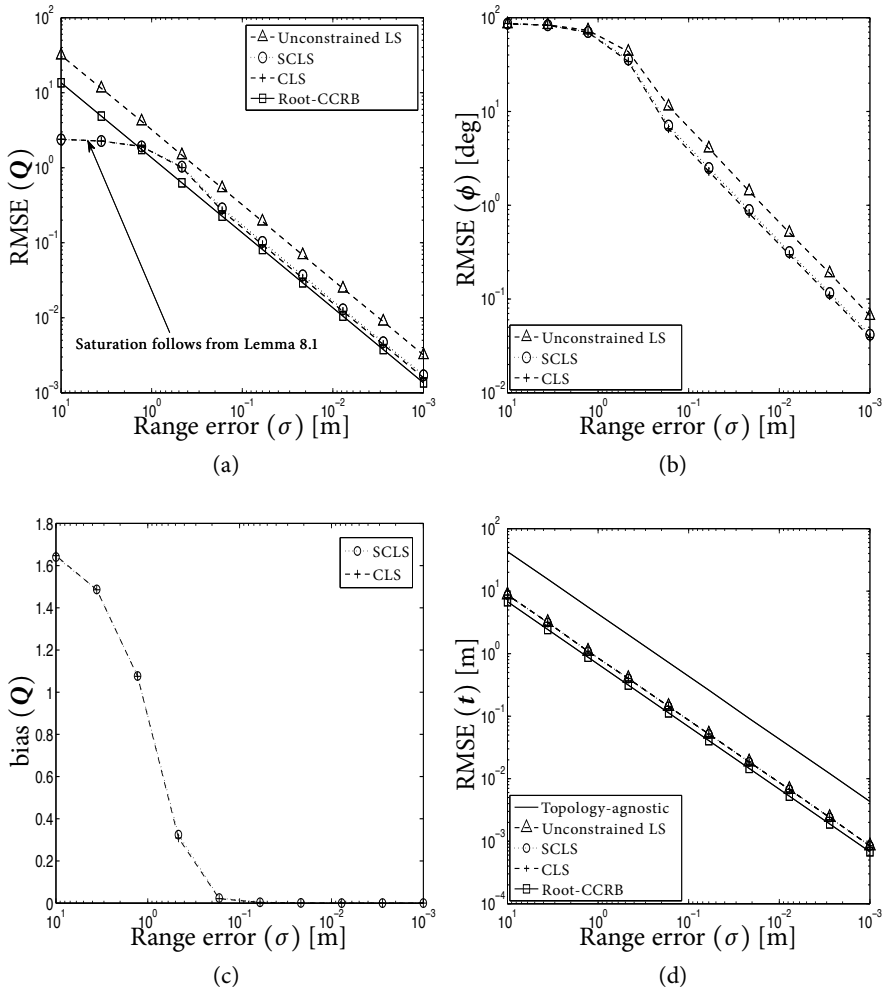


Figure 8.2: (a) RMSE of the estimated rotation matrix \mathbf{Q} . The RMSE of the constrained estimators is upper bounded as discussed in Remark 8.1. (b) RMSE in degrees of the estimated rotations. (c) Bias in the SCLS and CLS estimators for \mathbf{Q} . Bias tends to zero for low noise variance. (d) RMSE of the estimated translation vector \mathbf{t} along with the solution from the classical LS-based localization.

in (8.3) instead of $\boldsymbol{\rho}$ to compute the estimated covariance matrix

$$\widehat{\mathbf{R}}_n = \text{diag}(2\mathbf{y})\mathbf{R}_v\text{diag}(2\mathbf{y}) + \mathbf{R}_{v\odot 2} - \boldsymbol{\mu}\boldsymbol{\mu}^T.$$

We use this estimated covariance matrix for pre-whitening the noise. Simulations are provided for different values of the nominal ranging noise σ m.

In Figure 8.2, the RMSE of the estimates \mathbf{Q} , $\boldsymbol{\phi}$, and \mathbf{t} for different values of σ is shown. The estimators in Figure 8.2 are LS based where the topology of the sensors is assumed to be accurately known. Due to Theorem 8.2, the RMSE of \mathbf{Q} for the constrained estimators is lower than the CCRB at high noise levels. The saturation of the RMSE in Figure 8.2(a) for $\sigma > 1$ m follows from the following lemma.

Lemma 8.1 (Frobenius norm induced distance). *For any matrix \mathbf{Q}_i and \mathbf{Q}_j , such that, $\mathbf{Q}_i \in \mathcal{V}_{n,n}$ and $\mathbf{Q}_j \in \mathcal{V}_{n,n}$, the Frobenius norm induced distance is always upper bounded by $\sqrt{2n}$, i.e., $\|\mathbf{Q}_i - \mathbf{Q}_j\|_F \leq \sqrt{\|\mathbf{Q}_i\|_F + \|\mathbf{Q}_j\|_F} = \sqrt{2n}$.*

However, the CCRB computed using (8.33) does not saturate in this range as there exists no unbiased estimator for high noise values as discussed in Theorem 8.2. Figure 8.2(b) shows the RMSE in degrees, which gives an insight into how the error on the range measurements translates to the error on the estimated rotations. For the considered scenario, the range error that yields a positioning accuracy of the order of 10 cm also yields an orientation error accuracy of the order of 0.1 deg.

The bias of both the SCLS and CLS estimators is shown in Figure 8.2(c), and it can be seen that the bias tends to zero for $\sigma < 0.1$ m (as discussed in Theorem 8.2), whereas the unconstrained LS is an unbiased estimator. The bias is computed as

$$\text{bias}(\mathbf{Q}) = \left\| \frac{1}{N_{\text{exp}}} \sum_{n=1}^{N_{\text{exp}}} \text{vec}(\widehat{\mathbf{Q}}^{(n)}) - \text{vec}(\mathbf{Q}) \right\|_2.$$

We can see a significant (close to an order of magnitude) improvement in the performance of the location estimates when the knowledge of the topology is exploited as compared to the topology-agnostic case (see Figure 8.2(d)). The performance of the SCLS estimator is similar (slightly worse) to that of

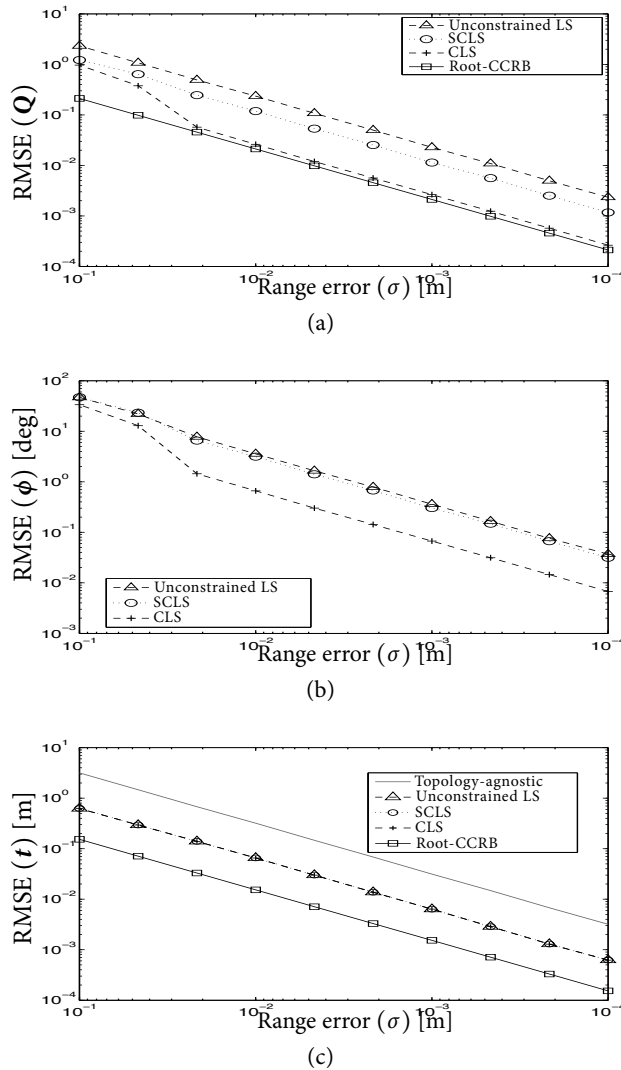


Figure 8.3: Effect of bad anchor geometry ($\text{cond}(\mathbf{A}) = 100$): (a) RMSE of the estimated rotation matrix \mathbf{Q} . (b) RMSE in degrees of the estimated rotations. (c) RMSE of the estimated translation vector \mathbf{t} .




Figure 8.4: Convergence of Gauss-Newton iterations for range error $\sigma = 10^{-2}$ m.

the iterative CLS estimator. The resulting gap between the RMSE and the root-CRB is reasonable, and thus the proposed estimators are robust to the linearization via squaring.

Figure 8.3 illustrates the effect of a bad anchor geometry, where we use an ill-conditioned matrix A with a condition number of 100. The performance of the SCLS estimators (based on pseudo inverting \bar{A}) deteriorates for scenarios with a bad anchor geometry, however, the performance of the CLS estimators is not affected. Nevertheless, if the topology is not subject to perturbations, the solution proposed in Remark 8.1 can be used for scenarios with a bad anchor geometry, in which a well-conditioned C can always be designed.

The convergence (i.e., $\epsilon_k := \frac{\|J^T(z - \text{Lvec}(\mathbf{Q}_k))\|_2}{\|J\|_F \|z - \text{Lvec}(\mathbf{Q}_k)\|_2}$, where J defined in (8.48)) of Gauss-Newton's method with $\sigma = 10^{-2}$ m for the optimal step size and a fixed step size is shown in Figure 8.4. We can see that it is sufficient to use a fixed step size $\alpha_k = 1$ at low noise levels. As observed from the simulations, the iterative algorithm requires typically ten or fewer iterations.

In order to analyze the performance of the estimators for the case when the sensor topology is perturbed, we corrupt the sensor coordinates in the reference frame with a zero-mean i.i.d. Gaussian random process of standard deviation $\sigma_e = 1$ cm. The RMSE of the estimates of \mathbf{Q} , ϕ , and \mathbf{t} using

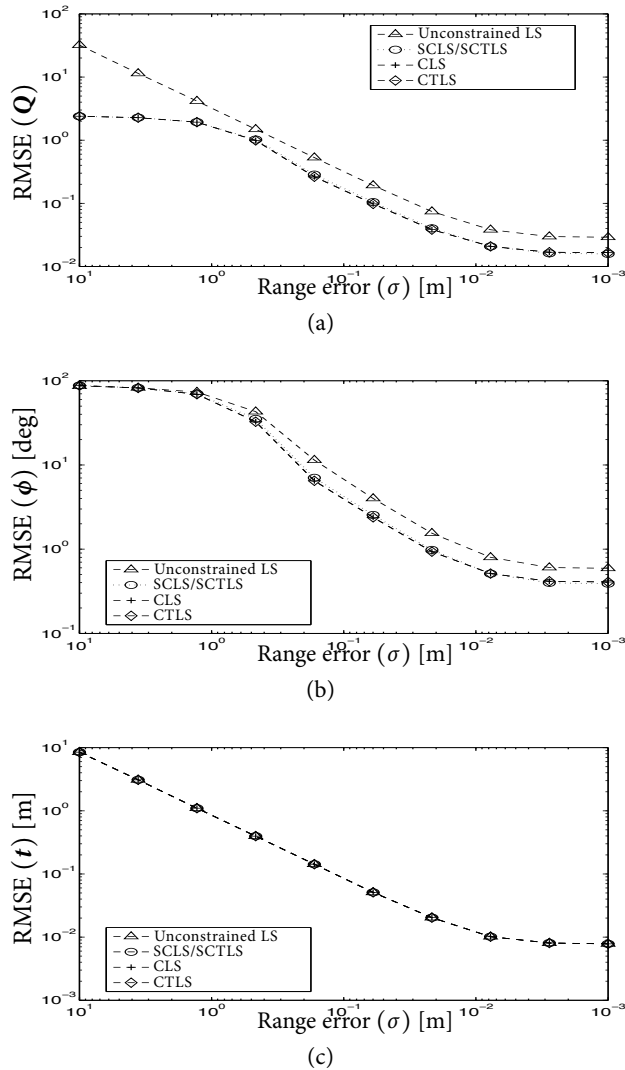


Figure 8.5: Results for a perturbed topology C . We use $\sigma_e = 1$ cm. (a) RMSE of the estimated rotation matrix Q . (b) RMSE in degrees of the estimated rotations. (c) RMSE of the estimated translation vector t .

the unconstrained LS, SCLS/SCTLS, CLS and CTLS estimators is shown in Figure 8.5. The performance of these estimators is similar to that of the LS-based estimators, except for the error floor, and this is due to the model error (perturbations on the sensor topology). The estimators hit the error floor for lower values of σ as σ_c^2 increases. With an ill-conditioned matrix \mathbf{A} , the performance of the SCLS/SCTLS (and the unconstrained LS) estimator is worse than the CLS and CTLS estimators.

8.7 Discussion

A framework for joint position and orientation estimation of a rigid body based on range measurements is proposed. We refer to this problem as rigid body localization. Sensor nodes can be mounted on the rigid bodies (e.g., satellites, robots) during fabrication at the factory, and the geometry of how these sensors are mounted is known *a priori* up to a certain accuracy. However, the absolute position of the rigid body is not known. The original non-linear problem is linearized via squaring of the range measurements. The squared-range measurements between the anchors and the sensors on the rigid body can then be used to estimate the position and the orientation of the body. The position and orientation of the rigid body is determined by a translation vector and a rotation matrix, respectively. Ignoring the fact that the rotation matrix is unitary, an unconstrained estimator is proposed. A number of unitarily constrained LS estimators is also proposed, all of which solve an optimization problem on the Stiefel manifold. All the proposed estimators are robust to linearization via squaring. For good anchor geometries, the performance of the SCLS estimator with a closed-form solution is already reasonable. The gap between the RMSE and root-CCRB of the SCLS and CLS estimates is negligible, however, the estimators do not achieve the CCRB. In addition to this, constrained TLS estimators that take into account the inaccuracies in the known sensor topology are also proposed.

8.A Derivation of the covariance matrix \mathbf{R}_n

In this section, we develop the covariance matrix \mathbf{R}_n used for pre-whitening. Modeling the noise on the range measurements v_{mn} as a zero-mean additive

white Gaussian process having a variance σ_{mn}^2 , we can compute the statistics (up to the second-order) of the zero-mean noise in (8.6), i.e.,

$$\mathbf{n}_{mn} = 2\rho_{mn}\mathbf{v}_{mn} + \mathbf{v}_{mn}^2 - \sigma_{mn}^2 \quad (8.43)$$

as

$$\begin{aligned} \mathbb{E}\{\mathbf{n}_{mn}\} &= \mathbf{0}, \\ \mathbb{E}\{(\mathbf{n}_{mn} - \mathbb{E}\{\mathbf{n}_{mn}\})^2\} &= 4\rho_{mn}^2\sigma_{mn}^2 + 2\sigma_{mn}^4, \\ \mathbb{E}\{(\mathbf{n}_{mn} - \mathbb{E}\{\mathbf{n}_{mn}\})(\mathbf{n}_{ln} - \mathbb{E}\{\mathbf{n}_{ln}\})\} &= \mathbf{0}, \quad m \neq l, \\ \text{and } \mathbb{E}\{(\mathbf{n}_{mn} - \mathbb{E}\{\mathbf{n}_{mn}\})(\mathbf{n}_{ml} - \mathbb{E}\{\mathbf{n}_{ml}\})\} &= \mathbf{0}, \quad n \neq l, \end{aligned}$$

where we use the fact that, if $\mathbf{v}_{mn} \sim \mathcal{N}(0, \sigma_{mn}^2)$ then $\mathbb{E}\{\mathbf{v}_{mn}^4\} = 3\sigma_{mn}^4$ and $\mathbb{E}\{\mathbf{v}_{mn}^3\} = 0$. Collecting ρ_{mn} , $m = 1, 2, \dots, M$; $n = 1, 2, \dots, N$, in the vector

$$\boldsymbol{\rho} = [\rho_{11}, \rho_{12}, \dots, \rho_{M1}, \dots, \rho_{MN}]^T \in \mathbb{R}^{MN},$$

\mathbf{v}_{mn} , $m = 1, 2, \dots, M$; $n = 1, 2, \dots, N$, in the vector

$$\mathbf{v} = [v_{11}, v_{12}, \dots, v_{M1}, \dots, v_{MN}]^T \in \mathbb{R}^{MN},$$

and σ_{mn}^2 , $m = 1, 2, \dots, M$; $n = 1, 2, \dots, N$, in the vector

$$\boldsymbol{\mu} = [\sigma_{11}^2, \sigma_{12}^2, \dots, \sigma_{M1}^2, \dots, \sigma_{MN}^2]^T \in \mathbb{R}^{MN},$$

we can write (8.43) compactly as

$$\mathbf{n} = \text{diag}(2\boldsymbol{\rho})\mathbf{v} + \mathbf{v}^{\odot 2} - \boldsymbol{\mu},$$

where

$$\mathbf{n} = [n_{11}, n_{12}, \dots, n_{M1}, \dots, n_{MN}]^T \in \mathbb{R}^{MN}.$$

We can then compute the mean of \mathbf{n} as $\mathbb{E}\{\mathbf{n}\} = \mathbf{0}_{MN}$, and the covariance matrix of \mathbf{n} as

$$\begin{aligned} \mathbf{R}_n &= \mathbb{E}\{(\mathbf{n} - \mathbb{E}\{\mathbf{n}\})(\mathbf{n} - \mathbb{E}\{\mathbf{n}\})^T\} \\ &= \text{diag}(2\boldsymbol{\rho})\mathbf{R}_v\text{diag}(2\boldsymbol{\rho}) + \mathbf{R}_{v^{\odot 2}} - \boldsymbol{\mu}\boldsymbol{\mu}^T \in \mathbb{R}^{MN \times MN}, \end{aligned}$$

where

$$\mathbf{R}_v = \text{diag}(\boldsymbol{\mu}) \in \mathbb{R}^{MN \times MN},$$

and

$$\mathbf{R}_{\mathbf{v}^{\otimes 2}} = \begin{bmatrix} 3\sigma_{11}^4 & \sigma_{11}^2\sigma_{12}^2 & \cdots & \sigma_{11}^2\sigma_{MN}^2 \\ \sigma_{11}^2\sigma_{12}^2 & 3\sigma_{12}^4 & \cdots & \vdots \\ \vdots & \cdots & \ddots & \vdots \\ \sigma_{11}^2\sigma_{MN}^2 & \cdots & \cdots & 3\sigma_{MN}^4 \end{bmatrix} \in \mathbb{R}^{MN \times MN}.$$

8.B Gauss-Newton iterations on the Stiefel manifold

The Newton method to solve the WOPP is developed in [Viklands, 2008], and it is adapted to suit our problem. In this section, we derive the Gauss-Newton iterations to solve the LS problems on the Stiefel manifold.

The CLS and CTLS problems solve an optimization problem on the Stiefel manifold of the form

$$\arg \min_{\mathbf{Q}} \|\mathbf{z} - \text{Lvec}(\mathbf{Q})\|_2^2 \quad \text{s.to} \quad \mathbf{Q} \in \mathcal{V}_{3,3}. \quad (8.44)$$

We can represent any unitary matrix \mathbf{Q} in the vicinity of a given unitary matrix \mathbf{Q}_k as

$$\mathbf{Q} = \mathbf{Q}_k \boldsymbol{\Omega}(\mathbf{x}), \quad (8.45)$$

where the operator $\boldsymbol{\Omega}(\cdot) : \mathbb{R}^3 \rightarrow \mathcal{V}_{3,3}$ is defined as

$$\boldsymbol{\Omega}(\mathbf{x}) = \exp(\mathbf{X}(\mathbf{x})), \quad (8.46)$$

and $\mathbf{X}(\mathbf{x})$ is a skew-symmetric matrix

$$\mathbf{X}(\mathbf{x}) = \begin{bmatrix} 0 & -x_1 & -x_2 \\ x_1 & 0 & -x_3 \\ x_2 & x_3 & 0 \end{bmatrix} \in \mathbb{R}^{3 \times 3},$$

with $\mathbf{x} = [x_1, x_2, x_3]^T \in \mathbb{R}^{3 \times 1}$. We use the matrix exponential $\exp(\mathbf{X})$ to map a point $\mathbf{x} \in \mathbb{R}^3$ onto the Stiefel manifold $\mathcal{V}_{3,3}$.

We linearize the matrix exponential by using a first-order expansion of the matrix exponential

$$\boldsymbol{\Omega}(\mathbf{x}) \approx \mathbf{I}_3 + \mathbf{X}.$$

Using this linearization¹ in (8.45) we get

$$\mathbf{Q} \approx \mathbf{Q}_k(\mathbf{I}_3 + \mathbf{X}).$$

We can then express $L\text{vec}(\mathbf{Q})$ as

$$L\text{vec}(\mathbf{Q}) \approx L\text{vec}(\mathbf{Q}_k) + L\text{vec}(\mathbf{Q}_k\mathbf{X}),$$

which is a function of \mathbf{x} , i.e.,

$$f_k(\mathbf{x}) = L\text{vec}(\mathbf{Q}_k) + \mathbf{J}\mathbf{x}, \quad (8.47)$$

where

$$\mathbf{J} = \frac{\partial L\text{vec}(\mathbf{Q}_k\mathbf{X})}{\partial \mathbf{x}^T} \in \mathbb{R}^{K \times 3}. \quad (8.48)$$

Next, we solve the optimization problem in (8.44) iteratively as follows. Using (8.47) in (8.44), we can transform the unitarily constrained optimization problem into an unconstrained minimization problem. More specifically, during the k -th iteration we compute the Gauss-Newton search direction by minimizing the following unconstrained LS problem

$$\begin{aligned} \mathbf{x}_k &= \arg \min_{\mathbf{x} \in \mathbb{R}^3} \psi(\mathbf{x}) = \|\mathbf{z} - f_k(\mathbf{x})\|_2^2 \\ &= (\mathbf{J}^T \mathbf{J})^{-1} \mathbf{J}^T (\mathbf{z} - L\text{vec}(\mathbf{Q}_k)), \end{aligned} \quad (8.49)$$

and subsequently compute the rotation update

$$\mathbf{Q}_{k+1} = \mathbf{Q}_k \mathbf{\Omega}(\alpha_k \mathbf{x}_k).$$

Here, α_k is the step size. The optimal step size is obtained by solving

$$\alpha_k = \arg \min_{\alpha \in (0,1]} \|\mathbf{z} - L\text{vec}(\mathbf{Q}_k \mathbf{\Omega}(\alpha \mathbf{x}_k))\|_2^2, \quad (8.50)$$

¹Instead of a matrix exponential, a Cayley transformation $\mathbf{\Omega}(\mathbf{x}) = (\mathbf{I}_3 + \mathbf{X})(\mathbf{I}_3 + \mathbf{X})^{-1}$ can be alternatively used, which can be then linearized by using a first-order expansion of $(\mathbf{I}_3 + \mathbf{X})^{-1} \approx \mathbf{I}_3 - \mathbf{X}$ (see [Viklands, 2008]). As a result, we get a similar expression $\mathbf{Q} \approx \mathbf{Q}_k(\mathbf{I}_3 + 2\mathbf{X})$.

whose solution is the root of the polynomial equation obtained by expanding the matrix exponential [Viklands, 2008], or can be computed simply by line search. With a good initial point and at low noise levels we can take $\alpha_k = 1$.

The solution $\widehat{\mathbf{Q}}_{\text{SCLS}}$ of the SCLS algorithm can be used as an initial point for the iterative algorithm. Alternatively, the initial point can be computed by orthonormalizing the solution of the unconstrained LS solution $\widehat{\mathbf{Q}}_{\text{LS}}$. The latter orthonormalization procedure solves a special case of OPP, and is given as

$$\begin{aligned} \mathbf{Q}_0 &:= \arg \min_{\mathbf{Q}} \|\mathbf{Q} - \widehat{\mathbf{Q}}_{\text{LS}}\|_F^2 \quad \text{s.to} \quad \mathbf{Q}^T \mathbf{Q} = \mathbf{I}_3 \\ &= (\widehat{\mathbf{Q}}_{\text{LS}} \widehat{\mathbf{Q}}_{\text{LS}}^T)^{-1/2} \widehat{\mathbf{Q}}_{\text{LS}}. \end{aligned} \quad (8.51)$$

8.C Proof of Theorem 8.2

We prove the claim of Theorem 8.2 by contradiction. Let there exist an unbiased constrained estimator $\widehat{\mathbf{Q}}$ such that $\widehat{\mathbf{Q}} \in \mathcal{V}_{3,3}$. Then $\widehat{\mathbf{Q}} = \mathbf{Q} + \boldsymbol{\Psi}$ where $\boldsymbol{\Psi}$ is the estimation error such that $\mathbb{E}\{\widehat{\mathbf{Q}}\} = \mathbf{Q}$ or $\mathbb{E}\{\boldsymbol{\Psi}\} = \mathbf{0}$. Since $\widehat{\mathbf{Q}} \in \mathcal{V}_{3,3}$, we have $\widehat{\mathbf{Q}}\widehat{\mathbf{Q}}^T = \mathbf{I}_3$, and hence

$$(\mathbf{Q} + \boldsymbol{\Psi})(\mathbf{Q} + \boldsymbol{\Psi})^T = \mathbf{I}_3. \quad (8.52)$$

Using $\mathbf{Q}\mathbf{Q}^T = \mathbf{I}_3$ and taking expectations on both sides, (8.52) can be further simplified to

$$\text{tr}\{\mathbb{E}\{\boldsymbol{\Psi}\}\mathbf{Q}^T\} + \text{tr}\{\mathbf{Q}\mathbb{E}\{\boldsymbol{\Psi}^T\}\} = -\mathbb{E}\{\|\boldsymbol{\Psi}\|_F^2\}. \quad (8.53)$$

Due to the assumption that $\mathbb{E}\{\boldsymbol{\Psi}\} = \mathbf{0}$, the left-hand side of (8.53) is zero, but, the right-hand side is strictly less than zero. Hence a contradiction occurs, unless the noise is zero.

8.D Proof of Theorem 8.3

The proof from [Arun, 1992] is provided here to aid the understanding of the proof of the next theorem. For any \mathbf{Q} , we can re-write the constraint in (8.36) as

$$\left[\mathbf{Q} \mid -\mathbf{I}_3 \right] \begin{bmatrix} \tilde{\mathbf{E}} \\ \check{\mathbf{N}} \end{bmatrix} = - \left[\mathbf{Q} \mid -\mathbf{I}_3 \right] \begin{bmatrix} \tilde{\mathbf{C}} \\ \check{\mathbf{D}} \end{bmatrix}.$$

Using the unitary constraint on \mathbf{Q} , and pseudo-inverting the wide matrix $\begin{bmatrix} \mathbf{Q} & | & -\mathbf{I}_{3(N-1)} \end{bmatrix}$ we get

$$\begin{aligned} \begin{bmatrix} \bar{\mathbf{E}} \\ \check{\mathbf{N}} \end{bmatrix} &= -\frac{1}{2} \begin{bmatrix} \mathbf{Q}^T \\ -\mathbf{I}_3 \end{bmatrix} \begin{bmatrix} \mathbf{Q} & | & -\mathbf{I}_3 \end{bmatrix} \begin{bmatrix} \bar{\mathbf{C}} \\ \check{\mathbf{D}} \end{bmatrix} \\ &= -\frac{1}{2} \begin{bmatrix} \mathbf{I}_3 & -\mathbf{Q}^T \\ -\mathbf{Q} & \mathbf{I}_3 \end{bmatrix} \begin{bmatrix} \bar{\mathbf{C}} \\ \check{\mathbf{D}} \end{bmatrix} \\ &= -\frac{1}{2} \begin{bmatrix} \bar{\mathbf{C}} - \mathbf{Q}^T \check{\mathbf{D}} \\ \check{\mathbf{D}} - \mathbf{Q} \bar{\mathbf{C}} \end{bmatrix}. \end{aligned}$$

We can now re-write the objective in (8.36) to compute the minimum-norm square solution

$$\begin{aligned} &\text{tr} \left\{ \begin{bmatrix} \bar{\mathbf{E}}^T & | & \check{\mathbf{N}}^T \end{bmatrix} \begin{bmatrix} \bar{\mathbf{E}} \\ \check{\mathbf{N}} \end{bmatrix} \right\} \\ &= \frac{1}{2} \text{tr}(\bar{\mathbf{C}}^T \bar{\mathbf{C}} - \check{\mathbf{D}}^T \mathbf{Q} \bar{\mathbf{C}} - \bar{\mathbf{C}}^T \mathbf{Q}^T \check{\mathbf{D}} + \check{\mathbf{D}}^T \check{\mathbf{D}}) \\ &= \frac{1}{2} \|\bar{\mathbf{C}}\|_F^2 - \text{tr}(\mathbf{Q} \bar{\mathbf{C}} \check{\mathbf{D}}^T) + \frac{1}{2} \|\check{\mathbf{D}}\|_F^2. \end{aligned}$$

The solution to the SCTLS problem is then obtained by optimizing the term depending only on \mathbf{Q} , i.e., by maximizing $\text{tr}\{\mathbf{Q} \bar{\mathbf{C}} \check{\mathbf{D}}^T\}$. This is the same cost as that of the SCLS problem (see [Golub and Van Loan, 1996, pg. 601]). Hence, the solution to the unitarily constrained TLS problem is

$$\widehat{\mathbf{Q}}_{\text{SCTLS}} = \mathbf{V} \mathbf{U}^T, \quad (8.54)$$

where the matrices \mathbf{V} and \mathbf{U} are obtained by computing the SVD of $\check{\mathbf{D}} \bar{\mathbf{C}}^T =: \mathbf{V} \mathbf{\Sigma} \mathbf{U}^T$.

8.E Proof of Theorem 8.4

The proof proceeds along the similar lines as the proof of Theorem 8.3. For any \mathbf{Q} the constraint in the optimization problem (8.39) can be written as

$$\begin{aligned} & \left[\left(\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\mathbf{Q} \right) \mid -\mathbf{I}_K \right] \begin{bmatrix} \bar{\mathbf{e}} \\ \bar{\mathbf{n}} \end{bmatrix} \\ & = - \left[\left(\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\mathbf{Q} \right) \mid -\mathbf{I}_K \right] \begin{bmatrix} \bar{\mathbf{c}} \\ \bar{\mathbf{d}} \end{bmatrix}. \end{aligned} \quad (8.55)$$

Multiplying both sides of (8.55) with the pseudo-inverse of the wide-matrix $\left[\left(\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\mathbf{Q} \right) \mid -\mathbf{I}_K \right]$ given by

$$\begin{aligned} & \left[\left(\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\mathbf{Q} \right) \mid -\mathbf{I}_K \right]^\dagger \\ & = \begin{bmatrix} \mathbf{I}_{3(N-1)} \otimes \mathbf{Q}^T \bar{\mathbf{A}}^T \\ -\mathbf{I} \end{bmatrix} \left(\left(\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\bar{\mathbf{A}}^T \right) + \mathbf{I}_K \right)^{-1}, \end{aligned}$$

we get the minimum-norm solution $\mathbf{k} = [\bar{\mathbf{e}}^T, \bar{\mathbf{n}}^T]^T$ to the system of equations in (8.55) which is given by

$$\begin{aligned} \begin{bmatrix} \bar{\mathbf{e}} \\ \bar{\mathbf{n}} \end{bmatrix} & = - \begin{bmatrix} \left(\mathbf{I}_{3(N-1)} \otimes \mathbf{Q}^T \bar{\mathbf{A}}^T \right) \\ -\mathbf{I}_K \end{bmatrix} \left(\left(\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\bar{\mathbf{A}}^T \right) + \mathbf{I}_K \right)^{-1} \\ & \quad \times \left[\left(\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\mathbf{Q} \right) \mid -\mathbf{I}_K \right] \begin{bmatrix} \bar{\mathbf{c}} \\ \bar{\mathbf{d}} \end{bmatrix}. \end{aligned} \quad (8.56)$$

Assuming that the covariance matrix of the perturbation vector is a scaled identity matrix, it is straightforward to verify that the objective in (8.40) using (8.56) simplifies to

$$\text{tr} \left\{ \left[\bar{\mathbf{e}}^T \mid \bar{\mathbf{n}}^T \right] \mathbf{R}_k \begin{bmatrix} \bar{\mathbf{e}} \\ \bar{\mathbf{n}} \end{bmatrix} \right\} = \|\Lambda^{1/2}(\bar{\mathbf{d}} - (\bar{\mathbf{C}}^T \otimes \bar{\mathbf{A}})\mathbf{q})\|_2^2,$$

where

$$\Lambda^{1/2} = \left(\sigma_e^{-2} (\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\bar{\mathbf{A}}^T) + \mathbf{R}_n^{-1} \right)^{1/2} \left(\left(\mathbf{I}_{3(N-1)} \otimes \bar{\mathbf{A}}\bar{\mathbf{A}}^T \right) + \mathbf{I}_K \right)^{-1}. \quad (8.57)$$

Hence, the solution to the optimization problem (8.40) is equivalent to a specifically weighted CLS.

Chapter 9

Conclusions and Future Research Directions

Contents

9.1	Concluding remarks	193
9.2	Directions for future research	196

9.1 Concluding remarks

Pervasive sensors collect massive amounts of data. As a consequence, it is becoming increasingly difficult to store and ship the acquired data to a central location for signal/data processing. Therefore, to overcome these problems, we have focused on developing sensing mechanisms to extract as much information as possible yet collecting fewer data. In particular, in the first part of this thesis, we have addressed the question “*How can task-cognition be exploited to reduce the costs of sensing as well as the related storage and communications requirements?*”. To answer this question, the tool that we have exploited for data reduction is sparse sensing. It consists of a deterministic sparse sensing function (guided by a sparse vector) that is optimally designed to achieve a desired inference performance with the reduced number of samples.

The sparse sensing model to acquire continuous-domain and discrete-domain observations, which we term as *continuous sparse sensing* and *discrete sparse sensing*, respectively, has been derived in Chapter 2. One of the fundamental differences with compressive sensing (a well-known data reduction mechanism used for sparse signal recovery) is that in sparse sensing the underlying signals need not be sparse. This allows us to focus on general signal processing tasks such as estimation, filtering, and detection. The developed theory can be applied to sensor placement, sensor/sample selection, data compression, with an impact on a range of signal processing applications.

Chapters 3, 4, and 5 have been, respectively, dedicated to discrete sparse sensing for estimation, filtering, and detection tasks. In all these chapters, under the assumption that the data is not yet available, the main questions that we have answered are: “*What are the reasonable inference performance metrics for estimation, filtering, and detection tasks?*” and “*Can we efficiently optimize (e.g., using a polynomial time algorithm) such inference performance metrics to obtain the sparse samplers of interest?*”. To this end, we have derived convex risk functions that include the sparse vector for specific inference tasks and thus, can be optimized via convex programming. Relying on the assumption that the data model is perfectly known, we have shown that measures such as the Cramér-Rao bound (CRB) for nonlinear estimation, posterior CRB for filtering, and Ali-Silvey distances (i.e., Bhattacharyya distance, Kullback-Leibler distance, or J-divergence) for binary hypothesis testing can be used as reasonable inference performance metrics. Furthermore, we have also shown that all these (ensemble) inference performance measures (A) can be computed in closed form and are independent of the observations, thus enabling offline (or data-independent) designs, and (B) can be optimized using a convex program, in many cases of interest. The optimization problem is relatively easier to solve when the observations are independent, irrespective of their distributions. By choosing appropriate instantaneous (instead of ensemble) inference performance metrics (that depend on data), the unifying approach presented in this thesis can be extended to *data-driven sparse sensing* for machine learning problems, like clustering or classification, which provides room for further research.

In Chapter 6, we have presented the *continuous sparse sensing* framework. Here, unlike the discrete case, we have shown that it is possible to sample in

between the grid points, that is, anywhere in the continuous observation domain. We have done this by modeling an off-the-grid sampling point as a discrete sampling point plus a perturbation. Then, the smallest set of possible discrete sampling points is searched for, along with the best possible perturbations, in order to reach the desired inference performance. We have solved the continuous sparse sensing problem via convex programming for linear inverse problems. There is still room for extending this framework (especially the solver) to other inference tasks, like hypothesis testing and nonlinear estimation/filtering with correlated errors.

Next, in the second part of this thesis, we have focused on some signal processing problems related to distributed sensing and sensor networks. More specifically, we have addressed the questions: “*How can wireless communications be exploited to synchronize spatially separated sample clocks?*”, and “*How can we extend the classical localization paradigm to localize a sensing platform by exploiting the knowledge of the sensor placement on the platform?*”.

In Chapter 7, we have provided a solution to synchronize the sample clocks of the sensors by fully exploiting the broadcast nature of the wireless medium, and using the time-of-flight measurements of the messages. We have assumed an affine clock model and that there exists at least one line-of-sight path between the nodes, which are reasonable in many cases of interest. Specifically, we have solved for all the unknown clock skews and clock offsets using least squares. The proposed estimator is shown to be efficient, asymptotically meeting the theoretical CRB. However, when there is no line-of-sight path, the wireless clock synchronization remains challenging.

Finally, in Chapter 8, we have solved the rigid body localization (i.e., position and orientation estimation) problem using only the distance measurements. To achieve this, we have exploited the prior information about the sensor placement on the platform. We have shown that the position and orientation of the rigid platform can be represented by a rotation matrix and translation vector. We have proposed a number of constrained least squares estimators, and we have also derived the theoretical limits (i.e., CRB) for the position and orientation estimation problem using distance measurements.

9.2 Directions for future research

The research for this thesis has generated a number of new challenging research questions, which remain open. In what follows, we conclude this thesis by listing some directions for future research.

9.2.1 Sparse sensing for constrained inverse problems

In Chapter 3, the focus was on designing sparse sensing mechanisms for nonlinear estimation problems. In particular, the goal was to sparsely sense the signal \mathbf{x} that is related to an unknown (nonrandom) parameter vector $\boldsymbol{\theta}$ through a known nonlinear model in order to achieve a desired estimation accuracy. We have used scalar functions of the CRB as the inference performance metric to quantify the estimation accuracy. In many applications, the parameter space may be restricted to a known subset of the Euclidean space through smooth functional constraints on the parameters. The question then is, how can we design sparse sensing functions that take into account such constraints on the unknown parameter. Note that this is different from the Bayesian approach that we have discussed in Remark 3.4 of Chapter 3, where the assumption was that the parameter is random and the prior information was related to its distribution.

If the parameter vector $\boldsymbol{\theta}$ is subject to some continuously differentiable constraints $\mathbf{g}(\boldsymbol{\theta}) = \mathbf{0}$, then with these constraints, the resulting constrained CRB (CCRB) is lower than the unconstrained CRB [cf. §8.4.2 from Chapter 8]. In [Stoica and Ng, 1998], it is shown that the CCRB has the form

$$\mathbb{E}\{(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})^T\} \geq \mathbf{C}_{\text{CCRB}}(\mathbf{w}, \boldsymbol{\theta}) = \mathbf{M}(\mathbf{M}^T \mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) \mathbf{M})^{-1} \mathbf{M}^T,$$

where $\widehat{\boldsymbol{\theta}}$ denotes the estimate of $\boldsymbol{\theta}$, $\mathbf{F}(\mathbf{w}, \boldsymbol{\theta})$ is the FIM for the unconstrained estimation problem as defined in (3.2), and \mathbf{M} is an isometry matrix whose columns form an orthonormal basis for the nullspace of the gradient matrix $\partial \mathbf{g}(\boldsymbol{\theta}) / \partial \boldsymbol{\theta}^T$. Here, \mathbf{w} is the sparse sensing vector to be designed. The question is, can we optimize the risk functions [cf. §3.3 of Chapter 3] $f(\mathbf{w}) := \text{tr}\{\mathbf{C}_{\text{CCRB}}(\mathbf{w}, \boldsymbol{\theta})\}$ or $f(\mathbf{w}) := \lambda_{\max}\{\mathbf{C}_{\text{CCRB}}(\mathbf{w}, \boldsymbol{\theta})\}$ using a convex program.

9.2.2 Sparse sensing for composite detection problems

In Chapter 5, the focus was on choosing the best subset of observations for a binary hypothesis testing problem. In particular, the assumption was that the model parameters were perfectly known. That is, our goal was to sparsely sense the signal \mathbf{x} that follows, for example, the conditional distributions:

$$\begin{aligned}\mathcal{H}_0: \quad \mathbf{x} &\sim h_0(\boldsymbol{\theta}_0, \mathbf{n}) \\ \mathcal{H}_1: \quad \mathbf{x} &\sim h_1(\boldsymbol{\theta}_1, \mathbf{n}),\end{aligned}$$

where we assumed that the parameter vectors $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$ of the functions h_0 and h_1 as well as the probability density function of the noise \mathbf{n} are perfectly known. When the parameter vectors $\boldsymbol{\theta}_0$ and $\boldsymbol{\theta}_1$ are unknown, we generally solve the so-called composite binary hypothesis testing problem [Scharf, 1991].

For composite hypothesis testing problems, the question is, how to design sparse sensing mechanisms. Clearly, the inference performance metrics that were discussed in §5.4 of Chapter 5 will now depend on the unknown parameters. Consequently, the risk function of the form $f(\mathbf{w}, \boldsymbol{\theta}_0, \boldsymbol{\theta}_1)$ has to be optimized over the sparse sensing vector \mathbf{w} , $\forall \boldsymbol{\theta}_0$ and $\forall \boldsymbol{\theta}_1$. This is reminiscent of the risk functions of the nonlinear estimation problem discussed in Chapter 3. Therefore, a way to solve the sparse sensing problem would be to grid the parameter space (every grid point would result in an additional performance constraint), or alternatively marginalize the unknown parameters using the prior statistical information, when available. This leads to a number of interesting questions, such as: “Is optimizing the signal-to-noise ratio optimal in terms of the error probabilities for Gaussian composite hypothesis testing?” and “How much data reduction can be achieved via sparse sensing for composite hypothesis testing as compared to (nonlinear) parameter estimation?”

9.2.3 Distributed sparse sensing

Throughout this thesis, we have assumed that the signals were available only at a single instance in space, frequency, or time. Suppose now that they are available at multiple instances, that is to say, the signals are distributed. For

example, imagine spatially separated sensors observing continuous-time phenomena. The question then is, how we can jointly design the different sparse sampling functions at different physical sensors to reach the best compression rate per sensor, subject to specific performance guarantees.

While deriving sparse sensing mechanisms for estimation, filtering, or detection, we always considered a single realization of the signal \mathbf{x} was available and designed the sparse sensing function $\Phi(\mathbf{w})$ to obtain the sparsely sensed signal \mathbf{y} . Now if multiple realizations are available, e.g., when \mathbf{x} is a spatial-domain signal, we could gather this signal at different time intervals or in different frequency bands. All these realizations can be sparsely sensed as before, by employing the same sensing function at different sensors. This might not be the best thing to do, as we do not fully exploit the available diversity to acquire the available multiple realizations using different sensing functions. In order to exploit the available diversity, clearly we should have the freedom to design different sparse sensing functions at different physical sensors. For example, let $\mathbf{x}_1 = [x_{1,1}, x_{2,1}, \dots, x_{M,1}]^T$ and $\mathbf{x}_2 = [x_{1,2}, x_{2,2}, \dots, x_{M,2}]^T$ be the two available temporal snapshots of the spatial signal $\mathbf{x} = [\mathbf{x}_1^T, \mathbf{x}_2^T]^T$. We acquire such distributed multidimensional signals using two sparse vectors \mathbf{w}_1 and \mathbf{w}_2 related to two temporal snapshots (instead of one sparse vector) as

$$\mathbf{y} = \Phi(\mathbf{w}_1, \mathbf{w}_2)\mathbf{x},$$

where $\Phi(\cdot, \cdot)$ denotes the sparse sensing function. Although the theory developed in this thesis can be used to design such sparse sensing functions, the interesting question is, to study the trade-off between spatial and temporal sampling densities.

9.2.4 Clock synchronization in non line-of-sight scenarios

The algorithms for wireless clock synchronization derived in Chapter 7 were based on time-of-flight measurements of messages between the sensor nodes. This relies on a fundamental assumption that there exists a line-of-sight path between the sensor nodes. The framework in Chapter 7 can be extended to multipath environments that include the line-of-sight path. This can, for example, be done by resolving the line-of-sight path using techniques like ray tracing or outlier rejection methods (i.e., to look for the consistent system

of equations). However, when there is no line-of-sight path, synchronizing sample clocks is a challenging problem.

One possible way to synchronize clocks is to use a phased-lock loop — a well-known phase and frequency tracking tool used in wireless communications. Let us define $a(t)$ as the output of an oscillator [Allan, 1987]

$$a(t) = a_0 \sin(2\pi f_0 t + \psi(t)),$$

where f_0 is the nominal frequency, a_0 is the constant amplitude, and $\psi(t)$ captures all the residual phase deviations. For an ideal clock, $\psi(t) = 0, \forall t$. In Chapter 7, we have modeled only the first order clock deviations and ignored the higher-order terms. That is, the assumption was

$$\psi(t) \approx \phi + \omega t = \phi + 2\pi f t,$$

where t is the *true* time, ϕ is the phase offset (or the clock offset) and $\omega = 2\pi f$ is the frequency offset (or the clock skew). The question is, can we use a phased-lock loop (or rather a dual phased-lock loop) discussed in [Johnson and Sethares, 2004, pp. 200-202] at each sensor to track the phase and frequency offsets with respect to a synchronization beacon transmitted from a sensor with a relatively stable clock. The second question is, what are the theoretical limits on synchronization based on such an approach.

9.2.5 Sensor, source placement, and closing the loop

In the first part of this thesis, the focus was on sensor placement, that is, to choose the best subset of sensor locations, subject to specific inference performance guarantees. The dual problem of sensor placement is the source placement problem (e.g., see [Ranieri and Vetterli, 2014]). The source placement problem can be interpreted as choosing the best subset of source locations in order to generate a desired field. Naturally, one would think of closing this loop through a joint sensor and source placement mechanism.

We will illustrate this with the rigid body localization problem that was discussed in Chapter 8. In particular, the question is, what are the optimal sensor/anchor placement (on the fixed world) and optimal source placement on the rigid platform to achieve the desired localization accuracy. To realize

this, we can express the Fisher information matrix [cf. (8.33)] in terms of two sparse vectors, where one vector is associated with the candidate anchor positions and the other one is associated with the candidate sensor positions. The resulting constrained CRB [cf. (8.33)] is then optimized to jointly minimize the number of sensors and sources. The question of interest is, how efficiently can we solve this optimization problem.

9.2.6 Robust sparse sensing

Data analytics (e.g., robust learning, classification, clustering) with large-scale data, such as data generated from social networks, financial markets, genomics, consumer behavior, and so on, is infeasible without dimensionality reduction due to limited computational capacity [Slavakis et al., 2014]. An interesting question is, how can sparse sensing be used to optimally compress already acquired data. This is different from the setup considered in this thesis as in data analytics the data is already available and stored. So, we need to study inference performance metrics for sketching or throwing away less informative samples. In this case, the design of sparse samplers will be data driven, where the samplers have to be designed for each data realization. Such a data dependent design is more appropriate for handling outliers or model mismatch (i.e., when the model information is not completely known), and are robust to them.

In addition to the above list, there is still room for studying greedy (e.g., submodular) algorithms for continuous and discrete sparse sensing (with dependent observations, in particular) problems that were considered in this thesis and for studying sparse sensing for multiple hypothesis testing. Developing a theory on relative position and orientation estimation between rigid bodies in an anchorless scenario, also remain open.

The ideas presented in this thesis with innovative perspectives on sparse sensing and distributed sensing are certainly pertinent to the design of modern sensing systems. We hope that these ideas will also open up new lines of research within the broad field of signal processing.

Glossary of Notation and Abbreviations

Notation

Sets

\mathbb{R}	Real numbers.
\mathbb{R}_+	Nonnegative real numbers.
\mathbb{C}	Complex numbers.
\mathbb{R}^N	Real length- N vectors.
$\mathbb{R}^{M \times N}$	Real $M \times N$ matrices.
\mathbb{C}^N	Complex length- N vectors.
$\mathbb{C}^{M \times N}$	Complex $M \times N$ matrices.
$\{0, 1\}^N$	Boolean length- N vectors.
$\{0, 1\}^{M \times N}$	Boolean $M \times N$ matrices.
\mathbb{Z}	Integers.
\mathbb{N}	Natural numbers.
\mathbb{S}^N	Symmetric $N \times N$ matrices.
\mathbb{S}_+^N	Symmetric positive semidefinite $N \times N$ matrices.
\mathbb{S}_{++}^N	Symmetric positive definite $N \times N$ matrices.
$ \mathcal{A} $	Cardinality of set \mathcal{A} .

Vectors and matrices

x	Scalar.
\mathbf{x}	Vector denoted by lower bold face letters.
\mathbf{X}	Matrix denoted by upper bold face letters.
\mathbf{X}^T	Transpose of matrix \mathbf{X} .
\mathbf{X}^H	Hermitian of matrix \mathbf{X} .
$\text{diag}(\cdot)$	Block diagonal matrix with the elements in its argument on the main diagonal.
$\text{diag}_r(\cdot)$	Diagonal matrix with the argument on its diagonal but with the all-zero rows removed.
$[\mathbf{X}]_{i,j}$	(i, j) th entry of matrix \mathbf{X} .
$[\mathbf{x}]_i$	i th entry of vector \mathbf{x} .
$\mathbf{1}_N$	$N \times 1$ vector of all ones.
$\mathbf{0}_N$	$N \times 1$ vector of all zeros.
\mathbf{I}_N	Identity matrix of size N .
$\mathbf{x}^{\odot 2}$	Element-wise squaring of vector \mathbf{x} .
$\text{tr}\{\mathbf{X}\}$	Trace of matrix \mathbf{X} .
$\det\{\mathbf{X}\}$	Determinant of matrix \mathbf{X} .
$\lambda_{\min}\{\mathbf{X}\}$	Minimum eigenvalue of a symmetric matrix \mathbf{X} .
$\lambda_{\max}\{\mathbf{X}\}$	Maximum eigenvalue of a symmetric matrix \mathbf{X} .
$\mathbf{X} \geq \mathbf{Y}$	$\mathbf{X} - \mathbf{Y}$ is a positive semidefinite matrix.
$\mathbf{X} \otimes \mathbf{Y}$	Kronecker product of matrix \mathbf{X} and matrix \mathbf{Y} .
$\mathbf{X}^\dagger = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T$	Pseudo inverse (or the left-inverse) of a full-column rank tall matrix \mathbf{X} .
$\mathbf{X}^\dagger = \mathbf{X}^T (\mathbf{X} \mathbf{X}^T)^{-1}$	Pseudo inverse (or the right-inverse) of a full-row rank wide matrix \mathbf{X} .
$\text{vec}(\mathbf{X})$	$MN \times 1$ vector formed by stacking the columns of an $M \times N$ matrix \mathbf{X} .
$\text{unvec}(\mathbf{X})$	$M \times N$ matrix \mathbf{X} formed by the inverse $\text{vec}(\mathbf{X})$ operation on an $MN \times 1$ vector.
$\text{cond}(\mathbf{X})$	Condition number of a matrix \mathbf{X} .

Norms and distances

$\ \mathbf{w}\ _0$	ℓ_0 -(quasi) norm, i.e., number of non-zero entries of vector \mathbf{w} .
$\ \mathbf{w}\ _1$	ℓ_1 -norm of vector \mathbf{w} .
$\ \mathbf{w}\ _2$	Euclidean (or ℓ_2 -)norm of vector \mathbf{w} .
$\ \mathbf{w}\ _A^2$	Weighted squared ℓ_2 -norm $\mathbf{w}^T \mathbf{A} \mathbf{w}$.
$\ \mathbf{W}\ _{2,0}$	ℓ_2/ℓ_0 -(quasi) norm counts the number of nonzero rows of matrix \mathbf{W} .

Stochastic processes

$\Pr(H)$	Probability of event H .
$\mathbb{E}\{\mathbf{x}\}$	Expected value of random vector \mathbf{x} .
$\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$	Gaussian distribution with mean vector $\boldsymbol{\mu}$, covariance matrix $\boldsymbol{\Sigma}$.
$\mathcal{U}(a, b)$	Uniform distribution with support $[a, b]$.
$p(x; \boldsymbol{\theta})$	Probability density function of x parameterized by $\boldsymbol{\theta}$.

Abbreviations

ADC	analog-to-digital converter
AOA	angle of arrival
CCRB	constrained Cramér-Rao bound
CLS	unitarily constrained least squares
CRB	Cramér-Rao bound
CS	compressive sensing
CTLS	unitarily constrained total least squares
FIM	Fisher information matrix
GPS	global positioning system
i.i.d.	independent and identically distributed
LS	least squares
LMI	linear matrix inequality
MDS	multi-dimensional scaling
MSE	mean squared error

OPP	orthogonal Procrustes problem
PCRB	posterior Cramér-Rao bound
RCRB	root Cramér-Rao bound
pdf	probability density function
RMSE	root mean squared error
RSS	received signal strength
SCLS	simplified unitarily constrained least squares
SCTLS	simplified unitarily constrained total least squares
SDP	semidefinite programming
SNR	signal-to-noise ratio
TDOA	time difference of arrival
TDC	time-to-digital converter
TOA	time of arrival
TLS	total least squares

Bibliography

- [Ali and Silvey, 1966] Ali, S. and Silvey, S. D. (1966). A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Stat. Society. Series B (Methodological)*, 28(1):131–142.
- [Allan, 1987] Allan, D. W. (1987). Time and frequency(time-domain) characterization, estimation, and prediction of precision clocks and oscillators. *IEEE Trans. on ultrasonics, ferroelectrics, and frequency control*, 34(6):647–654.
- [Angelosante et al., 2009] Angelosante, D., Giannakis, G. B., and Grossi, E. (2009). Compressed sensing of time-varying signals. In *Proc. of 16th International Conference on Digital Signal Processing*, Jul. 2009, Santorini, (Hellas) Greece.
- [Appadwedula et al., 2008] Appadwedula, S., Veeravalli, V. V., and Jones, D. L. (2008). Decentralized detection with censoring sensors. *IEEE Trans. Signal Process.*, 56(4):1362–1373.
- [Arun, 1992] Arun, K. (1992). A unitarily constrained total least squares problem in signal processing. *SIAM Journal on Matrix Analysis and Applications*, 13(3):729–745.
- [Bahr and Bucklew, 1990] Bahr, R. K. and Bucklew, J. A. (1990). Optimal sampling schemes for the Gaussian hypothesis testing problem. *IEEE Trans. Acoust., Speech, Signal Process.*, 38(10):1677–1686.

- [Bajovic et al., 2011] Bajovic, D., Sinopoli, B., and Xavier, J. (2011). Sensor selection for event detection in wireless sensor networks. *IEEE Trans. Signal Process.*, 59(10):4938–4953.
- [Baraniuk, 2007] Baraniuk, R. G. (2007). Compressive sensing. *IEEE Signal Process. Mag.*, 24(4):118–121.
- [Bertsekas, 1999] Bertsekas, D. P. (1999). *Nonlinear programming*. Athena Scientific optimization and computation series. Belmont, MA: Athena Scientific.
- [Blum and Sadler, 2008] Blum, R. S. and Sadler, B. M. (2008). Energy efficient signal detection in sensor networks using ordered transmissions. *IEEE Trans. Signal Process.*, 56(7):3229–3235.
- [Boyd and Vandenberghe, 2004] Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press.
- [Boyd et al., 2003] Boyd, S., Xiao, L., and Mutapcic, A. (2003). *Subgradient methods*. Lecture notes of EE364b. Stanford Univ., Stanford, CA, USA.
- [Cambanis and Masry, 1983] Cambanis, S. and Masry, E. (1983). Sampling designs for the detection of signals in noise. *IEEE Trans. Inf. Theory*, 29(1):83–104.
- [Candès and Wakin, 2008] Candès, E. and Wakin, M. (2008). An introduction to compressive sampling. *IEEE Signal Process. Mag.*, 25(2):21–30.
- [Candès et al., 2008] Candès, E., Wakin, M., and Boyd, S. (2008). Enhancing sparsity by reweighted ℓ_1 minimization. *Journal of Fourier Analysis and Applications*, 14(5):877–905.
- [Carmi, 2010] Carmi, A. (2010). Sensor scheduling via compressed sensing. In *Proc. of 13th Conference on Information Fusion*, Jul. 2010, Edinburgh, United Kingdom.
- [Carmi et al., 2010] Carmi, A., Gurfil, P., and Kanevsky, D. (2010). Methods for sparse signal recovery using Kalman filtering with embedded

- pseudo-measurement norms and quasi-norms. *IEEE Trans. Signal Process.*, 58(4):2405–2409.
- [Chamberland and Veeravalli, 2007] Chamberland, J. and Veeravalli, V. (2007). Wireless sensors in distributed detection applications. *IEEE Signal Process. Mag.*, 24(3):16–25.
- [Chaturvedi et al., 2011] Chaturvedi, N. A., Sanyal, A. K., and Mc Clamroch, N. H. (2011). Rigid-body attitude control. *IEEE Control Syst. Mag.*, 31(3):30–51.
- [Chepuri and Leus, 2015a] Chepuri, S. P. and Leus, G. (2015a). Sparse sensing for distributed Gaussian detection. In *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2015, Brisbane, Australia.
- [Chepuri and Leus, 2015b] Chepuri, S. P. and Leus, G. (2015b). Sparsity-promoting sensor selection for non-linear measurement models. *IEEE Trans. Signal Process.*, 63(3):684–698.
- [Chepuri et al., 2013a] Chepuri, S. P., Leus, G., and van der Veen, A.-J. (2013a). Position and orientation estimation of a rigid body: Rigid body localization. In *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2013, Vancouver, Canada.
- [Chepuri et al., 2013b] Chepuri, S. P., Leus, G., and van der Veen, A.-J. (2013b). Sparsity-exploiting anchor placement for localization in sensor networks. In *Proc. of the 21st European Signal Processing Conference*, Sept. 2013, Marrakech, Morocco.
- [Chepuri et al., 2013c] Chepuri, S. P., Rajan, R., Leus, G., and van der Veen, A.-J. (2013c). Joint clock synchronization and ranging: Asymmetrical time-stamping and passive listening. *IEEE Signal Process. Lett.*, 20(1):51–54.
- [Chepuri et al., 2013d] Chepuri, S. P., Simonetto, A., Leus, G., and van der Veen, A.-J. (2013d). Tracking position and orientation of a mobile rigid body. In *Proc. of of the IEEE Workshop on Comp. Adv. in Multi-Sensor Adaptive Processing*, Dec. 2013, St. Maarten, French Antilles.

- [Cheung and So, 2005] Cheung, K. W. and So, H. C. (2005). A multidimensional scaling framework for mobile location using time-of-arrival measurements. *IEEE Trans. Signal Process.*, 53(2):460–470.
- [Cohen, 1992] Cohen, C. (1992). *Attitude determination using GPS*. Ph.D. dissertation, Dept. of Aeronautics and Astronautics, Stanford Univ., Stanford, CA, USA.
- [Costa et al., 2006] Costa, J. A., Patwari, N., and Hero, III, A. O. (2006). Distributed weighted-multidimensional scaling for node localization in sensor networks. *ACM Trans. Sen. Netw.*, 2(1):39–64.
- [Cover and Thomas, 2012] Cover, T. M. and Thomas, J. A. (2012). *Elements of information theory*. Hoboken, NJ, USA: John Wiley & Sons.
- [Diebel, 2006] Diebel, J. (2006). *Representing attitude: Euler angles, unit quaternions, and rotation vectors*. Stanford Univ., Stanford, CA, USA.
- [Donoho, 2006a] Donoho, D. (2006a). For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution. *Communications on pure and applied mathematics*, 59(6):797–829.
- [Donoho, 2006b] Donoho, D. L. (2006b). Compressed sensing. *IEEE Trans. Inf. Theory*, 52(4):1289–1306.
- [Drineas et al., 2006] Drineas, P., Mahoney, M., and Muthukrishnan, S. (2006). Sampling algorithms for ℓ_2 regression and applications. In *Proc. of 17th annual ACM-SIAM symp. on Discrete algorithm*, Jan. 2006, Miami, FL, USA.
- [Ekanadham et al., 2011] Ekanadham, C., Tranchina, D., and Simoncelli, E. (2011). Recovery of sparse translation-invariant signals with continuous basis pursuit. *IEEE Trans. Signal Process.*, 59(10):4735–4744.
- [El Badawy et al., 2014] El Badawy, D., Ranieri, J., and Vetterli, M. (2014). Near-optimal sensor placement for signals lying in a union of subspaces. In *Proc. of 22nd European Signal Processing Conference*, Sept. 2014, Lisbon, Portugal.

- [Eldar and Kutyniok, 2012] Eldar, Y. C. and Kutyniok, G. (2012). *Compressed sensing: theory and applications*. Cambridge, U.K.: Cambridge Univ. Press.
- [Eldén and Park, 1999] Eldén, L. and Park, H. (1999). A Procrustes problem on the Stiefel manifold. *Numerische Mathematik*, 82(4):599–619.
- [Farahmand et al., 2014] Farahmand, S., Giannakis, G. B., Leus, G., and Tian, Z. (2014). Tracking target signal strengths on a grid using sparsity. *EURASIP Journal on Advances in Signal Processing*, 2014(1):7.
- [Ford et al., 1989] Ford, I., Titterington, D. M., and Kitsos, C. P. (1989). Recent advances in nonlinear experimental design. *Technometrics*, 31(1):49–60.
- [Freris et al., 2010] Freris, N., Kowshik, H., and Kumar, P. R. (2010). Fundamentals of large sensor networks: Connectivity, capacity, clocks, and computation. *Proc. of the IEEE*, 98(11):1828–1846.
- [Friedman et al., 2010] Friedman, J., Hastie, T., and Tibshirani, R. (2010). A note on the group LASSO and a sparse group LASSO. *arXiv preprint arXiv:1001.0736*.
- [Fu et al., 2012] Fu, Y., Ling, Q., and Tian, Z. (2012). Distributed sensor allocation for multi-target tracking in wireless sensor networks. *IEEE Trans. Aerosp. Electron. Syst.*, 48(4):3538–3553.
- [Gastpar et al., 2006] Gastpar, M., Vetterli, M., and Dragotti, P. L. (2006). Sensing reality and communicating bits: A dangerous liaison. *IEEE Signal Process. Mag.*, 23(4):70–83.
- [Gezici et al., 2005] Gezici, S., Tian, Z., Giannakis, G. B., Kobayashi, H., Molisch, A. F., Poor, H. V., and Sahinoglu, Z. (2005). Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *IEEE Signal Process. Mag.*, 22(4):70–84.
- [Golub and Van Loan, 1996] Golub, G. H. and Van Loan, C. F. (1996). *Matrix Computations*. Johns Hopkins Studies in the Mathematical Sciences. Baltimore, MD, USA: Johns Hopkins Univ. Press.

- [Grant and Boyd, 2012] Grant, M. and Boyd, S. (2012). CVX: Matlab software for disciplined convex programming, version 2.0 beta. <http://cvxr.com/cvx>.
- [Grettenberg, 1963] Grettenberg, T. (1963). Signal selection in communication and radar systems. *IEEE Trans. Inf. Theory*, 9(4):265–275.
- [Gustafsson and Gunnarsson, 2005] Gustafsson, F. and Gunnarsson, F. (2005). Mobile positioning using wireless networks: possibilities and fundamental limitations based on available wireless network measurements. *IEEE Signal Process. Mag.*, 22(4):41 – 53.
- [Haupt et al., 2009] Haupt, J., Nowak, R., and Castro, R. (2009). Adaptive sensing for sparse signal recovery. In *Proc. of IEEE 13th Digital Signal Processing Workshop and 5th IEEE Signal Processing Education Workshop*, Jan. 2009, Marco Island, Florida, USA.
- [Hernandez et al., 2004] Hernandez, M., Kirubarajan, T., and Bar-Shalom, Y. (2004). Multisensor resource deployment using Posterior Cramér-Rao bounds. *IEEE Trans. Aerosp. Electron. Syst.*, 40(2):399–416.
- [Hoeffding, 1963] Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30.
- [Hol, 2011] Hol, J. (2011). *Sensor fusion and calibration of inertial sensors, vision, Ultra-Wideband and GPS*. Ph.D. dissertation, Dept. of Electrical Engineering, Automatic Control. Linköping University, Sweden.
- [Hol et al., 2009] Hol, J., Dijkstra, F., Luinge, H., and Schon, T. (2009). Tightly coupled UWB/IMU pose estimation. In *Proc. of IEEE International Conference on Ultra-Wideband*, Sept. 2009, Vancouver, Canada.
- [Horn, 1987] Horn, B. K. P. (1987). Closed-form solution of absolute orientation using unit quaternions. *Journal of the Optical Society of America A*, 4(4):629–642.

- [Horn et al., 1988] Horn, B. K. P., Hilden, H. M., and Negahdaripour, S. (1988). Closed-form solution of absolute orientation using orthonormal matrices. *Journal of the Optical Society of America A*, 5(7):1127–1135.
- [Ji et al., 2008] Ji, S., Xue, Y., and Carin, L. (2008). Bayesian compressive sensing. *IEEE Trans. Signal Process.*, 56(6):2346–2356.
- [Jiang et al., 2013] Jiang, F., Chen, J., and Swindlehurst, A. L. (2013). Linearly reconfigurable Kalman filtering for a vector process. In *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2013, Vancouver, Canada.
- [Johnson and Sethares, 2004] Johnson, C. R. and Sethares, W. A. (2004). *Telecommunication breakdown: concepts of communication transmitted via software-defined radio*. Englewood Cliffs, NJ, USA: Prentice-Hall.
- [Joshi and Boyd, 2009] Joshi, S. and Boyd, S. (2009). Sensor selection via convex optimization. *IEEE Trans. Signal Process.*, 57(2):451–462.
- [Juang and Huang, 1997] Juang, J.-C. and Huang, G.-S. (1997). Development of GPS-based attitude determination algorithms. *IEEE Trans. Aerosp. Electron. Syst.*, 33(3):968–976.
- [Julier and LaViola, 2007] Julier, S. and LaViola, J. (2007). On Kalman filtering with nonlinear equality constraints. *IEEE Trans. Signal Process.*, 55(6):2774–2784.
- [Kadota and Shepp, 1967] Kadota, T. and Shepp, L. A. (1967). On the best finite set of linear observables for discriminating two Gaussian signals. *IEEE Trans. Inf. Theory*, 13(2):278–284.
- [Kailath, 1967] Kailath, T. (1967). The divergence and Bhattacharyya distance measures in signal selection. *IEEE Trans. Commun. Technol.*, 15(1):52–60.
- [Kay, 1993] Kay, S. M. (1993). *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice-Hall.
- [Kekatos and Giannakis, 2011] Kekatos, V. and Giannakis, G. B. (2011). From sparse signals to sparse residuals for robust sensing. *IEEE Trans. Signal Process.*, 59(7):3355–3368.

- [Kekatos et al., 2012] Kekatos, V., Giannakis, G. B., and Wollenberg, B. (2012). Optimal placement of phasor measurement units via convex relaxation. *IEEE Trans. Power Syst.*, 27(3):1521–1530.
- [Kobayashi and Thomas, 1967] Kobayashi, H. and Thomas, J. B. (1967). Distance measures and related criteria. In *Proc. 5th Annu. Allerton Conf. Circuit and System Theory*, Oct. 1967, Monticello, IL, USA.
- [Krause, 2008] Krause, A. (2008). *Optimizing sensing: Theory and applications*. Ph.D. dissertation, School of Comput. Sci. Carnegie Mellon Univ., Pittsburgh, PA, United States.
- [Krause and Guestrin, 2007] Krause, A. and Guestrin, C. (2007). Near-optimal observation selection using submodular functions. In *Proc. of Twenty-Second Conference on Artificial Intelligence*, Jul. 2007, Vancouver, Canada.
- [Krause et al., 2008a] Krause, A., McMahan, H. B., Guestrin, C., and Gupta, A. (2008a). Robust submodular observation selection. *The Journal of Machine Learning Research*, 9(12):2761–2801.
- [Krause et al., 2008b] Krause, A., Singh, A., and Guestrin, C. (2008b). Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies. *The Journal of Machine Learning Research*, 9(2):235–284.
- [Kullback, 1959] Kullback, S. (1959). *Information theory and statistics*. Hoboken, NJ, USA: John Wiley & Sons.
- [Kumar et al., 2015] Kumar, S., Zjajo, A., and van Leuken, R. (2015). Ctherm: An integrated framework for thermal-functional co-simulation of systems-on-chip. In *Proc. of the 23rd IEEE/EuroMicro Conference on Parallel, Distributed and Network-Based Processing*, Mar. 2015, Turku, Finland.
- [Liu et al., 2014] Liu, S., Fardad, M., Varshney, P. K., and Masazade, E. (2014). Optimal periodic sensor scheduling in networks of dynamical systems. *IEEE Trans. Signal Process.*, 62(12):3055–3068.

- [Lofberg, 2004] Lofberg, J. (2004). YALMIP : A toolbox for modeling and optimization in MATLAB. In *Proc. of International Symposium on Computer Aided Control System Design*, Sept. 2004, Taipei, Taiwan.
- [Malioutov et al., 2010] Malioutov, D. M., Sanghavi, S. R., and Willsky, A. S. (2010). Sequential compressed sensing. *IEEE J. Sel. Topics Signal Process.*, 4(2):435–444.
- [Manton, 2002] Manton, J. H. (2002). Optimization algorithms exploiting unitary constraints. *IEEE Trans. Signal Process.*, 50(3):635–650.
- [Marvasti, 2001] Marvasti, F. (2001). *Nonuniform sampling: theory and practice*. New York, USA: Kluwer Academic/Plenum Publishers.
- [Marziliano and Vetterli, 2000] Marziliano, P. and Vetterli, M. (2000). Reconstruction of irregularly sampled discrete-time bandlimited signals with unknown sampling locations. *IEEE Trans. Signal Process.*, 48(12):3462–3471.
- [Masazade et al., 2012] Masazade, E., Fardad, M., and Varshney, P. (2012). Sparsity-promoting extended Kalman filtering for target tracking in wireless sensor networks. *IEEE Signal Process. Lett.*, 19(12):845–848.
- [Memik et al., 2008] Memik, S. O., Mukherjee, R., Ni, M., and Long, J. (2008). Optimizing thermal sensor allocation for microprocessors. *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, 27(3):516–527.
- [Moon and Stirling, 2000] Moon, T. K. and Stirling, W. C. (2000). *Mathematical methods and algorithms for signal processing*. Englewood Cliffs, NJ, USA: Prentice-Hall.
- [Msechu and Giannakis, 2012] Msechu, E. J. and Giannakis, G. B. (2012). Sensor-centric data reduction for estimation with WSNs via censoring and quantization. *IEEE Trans. Signal Process.*, 60(1):400–414.
- [Nemhauser et al., 1978] Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions— I. *Mathematical Programming*, 14(1):265–294.

- [Noh et al., 2008] Noh, K.-L., Serpedin, E., and Qaraqe, K. (2008). A new approach for time synchronization in wireless sensor networks: Pairwise broadcast synchronization. *IEEE Trans. Wireless Commun.*, 7(9):3318 – 3322.
- [Papadimitriou, 2003] Papadimitriou, C. H. (2003). *Computational complexity*. Hoboken, NJ, USA: John Wiley & Sons.
- [Patwari et al., 2005] Patwari, N., Ash, J. N., Kyperountas, S., Hero III, A. O., Moses, R., and Correal, N. (2005). Locating the nodes: cooperative localization in wireless sensor networks. *IEEE Signal Process. Mag.*, 22(4):54 – 69.
- [Polyak et al., 2013] Polyak, B., Khlebnikov, M., and Shcherbakov, P. (2013). An LMI approach to structured sparse feedback design in linear control systems. In *Proc. of European Control Conference (ECC)*, Jul. 2013, Zurich, Switzerland.
- [Pukelsheim, 1993] Pukelsheim, F. (1993). *Optimal design of experiments*, volume 50. Philadelphia, PA: SIAM.
- [Quan et al., 2009] Quan, Z., Kaiser, W. J., and Sayed, A. H. (2009). Innovations diffusion: A spatial sampling scheme for distributed estimation and detection. *IEEE Trans. Signal Process.*, 57(2):738–751.
- [Rago et al., 1996] Rago, C., Willett, P., and Bar-Shalom, Y. (1996). Censoring sensors: A low-communication-rate scheme for distributed detection. *IEEE Trans. Aerosp. Electron. Syst.*, 32(2):554–568.
- [Rajan and van der Veen, 2011] Rajan, R. and van der Veen, A.-J. (2011). Joint ranging and clock synchronization for a wireless network. In *Proc. of 4th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, Dec. 2011, San Juan, Puerto Rico.
- [Ranieri et al., 2014] Ranieri, J., Chebira, A., and Vetterli, M. (2014). Near-optimal sensor placement for linear inverse problems. *IEEE Trans. Signal Process.*, 62(5):1135–1146.

- [Ranieri and Vetterli, 2014] Ranieri, J. and Vetterli, M. (2014). Near-optimal source placement for linear physical fields. In *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2014, Florence, Italy.
- [Rothenberg, 1971] Rothenberg, T. (1971). Identification in parametric models. *Econometrica: Journal of the Econometric Society*, 39(3):577–591.
- [Salhuana, 2012] Salhuana, L. (2012). Tilt sensing using linear accelerometers. In *Appl. note AN3461*, Feb. 2012, TX, USA. Freescale Semiconductor Inc.
- [Scharf, 1991] Scharf, L. L. (1991). *Statistical signal processing*. Reading, MA, USA: Addison-Wesley.
- [Schizas, 2013] Schizas, I. (2013). Distributed informative-sensor identification via sparsity-aware matrix decomposition. *IEEE Trans. Signal Process.*, 61(18):4610–4624.
- [Shamaiah et al., 2010] Shamaiah, M., Banerjee, S., and Vikalo, H. (2010). Greedy sensor selection: Leveraging submodularity. In *Proc. of 49th IEEE Conference on Decision and Control*, Dec. 2010, Atlanta, Georgia, USA.
- [Slavakis et al., 2014] Slavakis, K., Giannakis, G., and Mateos, G. (2014). Modeling and optimization for big data analytics:(statistical) learning tools for our era of data deluge. *IEEE Signal Process. Mag.*, 31(5):18–31.
- [Stoica and Ng, 1998] Stoica, P. and Ng, B. C. (1998). On the Cramér-Rao bound under parametric constraints. *IEEE Signal Process. Lett.*, 5(7):177–179.
- [Sturm, 1999] Sturm, J. F. (1999). Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization methods and software*, 11(1-4):625–653.
- [Sung et al., 2005] Sung, Y., Tong, L., and Poor, H. V. (2005). A large deviations approach to sensor scheduling for detection of correlated random fields. In *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing*, pages Mar. 2005, Philadelphia, PA, USA.

- [Tibshirani et al., 2005] Tibshirani, R., Saunders, M., Rosset, S., Zhu, J., and Knight, K. (2005). Sparsity and smoothness via the fused LASSO. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(1):91–108.
- [Tichavsky et al., 1998] Tichavsky, P., Muravchik, C., and Nehorai, A. (1998). Posterior Cramér–Rao bounds for discrete-time nonlinear filtering. *IEEE Trans. Signal Process.*, 46(5):1386–1396.
- [Vaidyanathan, 2001] Vaidyanathan, P. (2001). Generalizations of the sampling theorem: Seven decades after Nyquist. *IEEE Trans. Circuits Syst. I: Fundamental Theory and Applications*, 48(9):1094–1109.
- [Van Trees, 2004] Van Trees, H. L. (2004). *Detection, Estimation, and Modulation Theory, Optimum Array Processing*. Hoboken, NJ, USA: John Wiley & Sons.
- [Vaswani, 2008] Vaswani, N. (2008). Kalman filtered compressed sensing. In *Proc. of 15th IEEE International Conference on Image Processing*, Oct. 2008, San Diego, CA, USA.
- [Viklands, 2008] Viklands, T. (2008). *Algorithms for the Weighted Orthogonal Procrustes problem and Other Least Squares Problems*. Ph.D. dissertation, Dept. Comput. Sci., Umea Univ., Umea, Sweden.
- [Wahba, 1965] Wahba, G. (1965). A least squares estimate of satellite attitude. *SIAM review*, 7(3):409–409.
- [Wang et al., 2009] Wang, T., Leus, G., and Huang, L. (2009). Ranging energy optimization for robust sensor positioning based on semidefinite programming. *IEEE Trans. Signal Process.*, 57(12):4777 –4787.
- [Wang et al., 2011] Wang, Y., Ma, X., and Leus, G. (2011). Robust time-based localization for asynchronous networks. *IEEE Trans. Signal Process.*, 59(9):4397 –4410.
- [Warnell et al., 2012] Warnell, G., Reddy, D., and Chellappa, R. (2012). Adaptive rate compressive sensing for background subtraction. In *Proc. of IEEE*

- International Conference on Acoustics, Speech and Signal Processing*, Mar. 2012, Kyoto, Japan. IEEE.
- [Wu et al., 2011] Wu, Y.-C., Chaudhari, Q., and Serpedin, E. (2011). Clock synchronization of wireless sensor networks. *IEEE Signal Process. Mag.*, 28(1):124–138.
- [Xiao and Boyd, 2004] Xiao, L. and Boyd, S. (2004). Fast linear iterations for distributed averaging. *Systems & Control Letters*, 53(1):65–78.
- [Yao et al., 1993] Yao, L., Sethares, W., and Kammer, D. (1993). Sensor placement for on-orbit modal identification via a genetic algorithm. *The American Institute of Aeronautics and Astronautics Journal*, 31(10):1922–1928.
- [Yu and Varshney, 1997] Yu, C.-T. and Varshney, P. K. (1997). Sampling design for Gaussian detection problems. *IEEE Trans. Signal Process.*, 45(9):2328–2337.
- [Yuille and Rangarajan, 2003] Yuille, A. L. and Rangarajan, A. (2003). The concave-convex procedure. *Neural computation*, 15(4):915–936.
- [Zhan et al., 2010] Zhan, P., Casbeer, D., and Swindlehurst, A. (2010). Adaptive mobile sensor positioning for multi-static target tracking. *IEEE Trans. Aerosp. Electron. Syst.*, 46(1):120–132.
- [Zhu et al., 2011] Zhu, H., Leus, G., and Giannakis, G. B. (2011). Sparsity-cognizant total least-squares for perturbed compressive sampling. *IEEE Trans. Signal Process.*, 59(5):2002–2016.
- [Zuo et al., 2007] Zuo, L., Niu, R., and Varshney, P. (2007). Posterior CRLB based sensor selection for target tracking in sensor networks. In *Proc. of IEEE International Conference on Acoustics, Speech and Signal Processing*, Apr. 2007, Honolulu, Hawaii, USA.

Summary

Sparse Sensing for Statistical Inference Theory, Algorithms, and Applications

In today's society, we are flooded with massive volumes of data in the order of a billion gigabytes on a daily basis from pervasive sensors. It is becoming increasingly challenging to locally store and transport the acquired data to a central location for signal/data processing (i.e., for inference). To alleviate these problems, it is evident that there is an urgent need to significantly reduce the sensing cost (i.e., the number of expensive sensors) as well as the related memory and bandwidth requirements by developing unconventional sensing mechanisms to extract as much information as possible yet collecting fewer data.

The first aim of this thesis is to develop theory and algorithms for data reduction. We develop a data reduction tool called sparse sensing, which consists of a deterministic and structured sensing function (guided by a sparse vector) that is optimally designed to achieve a desired inference performance with the reduced number of data samples. The first part of this thesis is dedicated to the development of sparse sensing mechanisms and convex programs to efficiently design sparse sensing functions. We design sparse sensing functions under the assumption that the data is not yet available and the model information is perfectly known.

Sparse sensing offers a number of advantages over compressed sensing (a state-of-the-art data reduction method for sparse signal recovery). One of the major differences is that in sparse sensing the underlying signals need not be sparse. This allows us to consider general signal processing tasks (not

just sparse signal recovery) under the proposed sparse sensing framework. Specifically, we focus on fundamental statistical inference tasks, like estimation, filtering, and detection. In essence, we present topics that transform classical (e.g., random or uniform) sensing methods to low-cost data acquisition mechanisms tailored for specific inference tasks. The developed framework can be applied to sensor selection, sensor placement, or sensor scheduling, for example.

In the second part of this thesis, we focus on some applications related to distributed sampling using sensor networks. Sensor networks can be used as a spatial sampling device, that is, to faithfully represent distributed signals (e.g., a spatially varying phenomenon such as a temperature field). On top of that, the distributed signals can exist in space and time, where the temporal sampling is achieved using analog-to-digital converters, for example. Each sensor has an independent sample clock, and its stability essentially determines the alignment of the temporal sampling grid across the sensors. Due to imperfections in the oscillator, the sample clocks drift from each other, resulting in the misalignment of the temporal sampling grids. To overcome this issue, we devise a mechanism to distribute the sample clock wirelessly. Specifically, we perform wireless clock synchronization based on the time-of-flight measurements of broadcast messages. In addition, clock synchronization also plays a central role in other time-based sensor network applications such as localization.

Localization is increasingly gaining popularity in many applications, especially for monitoring environments beyond human reach, e.g., using robots or drones with several sensor units mounted on it. Consequently we now have to localize more than one sensor or even localize the whole sensing platform. Therefore, we extend the classical localization paradigm to localize a (rigid) sensing platform by exploiting the knowledge of the sensor placement on the platform. In particular, we develop algorithms for rigid body localization, i.e., for estimating the position and orientation of the rigid platform using distance measurements.

Given the central role of sensing and sensor networks, the results presented in this thesis impacts a wide range of applications.

Samenvatting

Schaarse Data-Acquisitie voor Statistische Inferentie Theorie, Algoritmes, en Toepassingen

De maatschappij van tegenwoordig wordt overspoeld met grote volumes data in de orde van biljoenen gigabytes per dag afkomstig van alomtegenwoordige sensoren. Het wordt alsnog moeilijker om die data lokaal op te slaan en door te sturen naar een centrale locatie voor signaal- en dataverwerking (m.a.w. voor inferentie). Vandaar dat onconventionele data-acquisitietechnieken dringend nodig zijn om zo veel mogelijk informatie te vergaren en dit met zo weinig data-acquisitie als nodig. Dit zal zowel de kost van data-acquisitie (aantal dure sensoren) als de benodigde geheugeninhoud en communicatiebandbreedte verminderen.

Het eerste doel van deze thesis is het ontwikkelen van een theorie en algoritmes voor datareductie. We ontwikkelen een datareductietechniek, genaamd schaarse data-acquisitie of “sparse sensing”, die bestaat uit een deterministische en gestructureerde data-acquisitiefunctie (geparametriseerd met een schaarse vector) die optimaal ontworpen is om een bepaalde inferentieperformantie te behalen met een gereduceerd aantal data. Het eerste deel van deze thesis is gewijd aan het ontwikkelen van schaarse data-acquisitietechnieken en convexe programmas om schaarse data-acquisitiefuncties te ontwerpen. We ontwerpen schaarse data-acquisitiefuncties in de veronderstelling dat de data nog niet beschikbaar is maar de modelinformatie perfect gekend is.

Schaarse data-acquisitie biedt een aantal voordelen t.o.v. gecomprimeerde data-acquisitie of “compressed sensing”(een van de nieuwste datareductie-

methodes voor de reconstructie van schaarse signalen). Eén van de grootste verschillen is dat de onderliggende signalen voor schaarse data-acquisitie niet schaars hoeven te zijn. Dit laat ons toe om in het voorgestelde schaarse data-acquisitiekader meer algemene signaalverwerkingstechnieken te behandelen (en niet enkel de reconstructie van schaarse signalen). Meer specifiek richten we ons op fundamentele statistische inferentietaken zoals schatting, filtering, en detectie. In essentie stellen we werktuigen voor die klassieke data-acquisitietechnieken (bijv. random of uniform) transformeren in nieuwe data-acquisitietechnieken met een lage kost en toegespitst op specifieke inferentietaken. Het ontwikkelde kader kan toegepast worden op sensorselectie, sensorpositionering, of sensorplanning.

In het tweede deel van deze thesis richten we ons op een aantal toepassingen gerelateerd aan bemonstering via sensornetwerken. Sensornetwerken kunnen gebruikt worden als een ruimtelijk bemonsteringssysteem, of met andere woorden, om op een betrouwbare manier gedistribueerde signalen voor te stellen (bijv. fenomenen die variëren in de ruimte zoals een temperatuursveld). Daarenboven kunnen gedistribueerde signalen bestaan in ruimte en tijd, waarbij de bemonstering in het tijdsdomein wordt uitgevoerd door de sensor zijn analoog-digitaalomzetter bijvoorbeeld. Iedere sensor heeft een onafhankelijke bemonsteringsklok en de stabiliteit van deze klok bepaalt of de sensoren synchroon lopen of niet. Onnauwkeurigheden in de klokoscillator kunnen drift veroorzaken en zorgen ervoor dat de sensoren niet synchroon bemonsteren. Om dit probleem op te lossen hebben we een techniek ontwikkeld om een kloksignaal draadloos uit te wisselen. Meer specifiek hebben we een draadloos synchronisatiealgoritme ontworpen gebaseerd op de propagatietijd van communicatiesignalen. Verder speelt kloksynchronisatie ook een belangrijke rol voor andere toepassingen gebaseerd op tijdsmetingen zoals lokalisatie.

Lokalisatie wint aan populariteit in vele geautomatiseerde observatietoepassingen, bijv. door middel van robots of drones waarop verschillende sensoren worden gemonteerd. Om zo een sensorstructuur te lokaliseren moeten we dus gelijktijdig verschillende sensoren lokaliseren. Daarom breiden we in deze thesis het klassieke lokalisatieparadigma uit naar het lokaliseren van een (rigide) sensorstructuur gebruik makende van de kennis van de locatie van de sensoren in die structuur. In het bijzonder ontwikkelen we algoritmes voor de

lokalisatie van rigide sensorstructuren, m.a.w. om de positie en de oriëntatie van de structuur te schatten op basis van afstandsmetingen.

Gegeven de centrale rol die data-acquisitie en sensornetwerken spelen hebben de resultaten die voorgesteld zijn in deze thesis een impact op een overvloed aan toepassingen.

Geert Leus

Delft, december 2015

Biography

S.P. (Sundeep Prabhakar) Chepuri was born in Bangalore, India in 1986. He received the Bachelors in Engineering degree (*with distinction*) from the PES Institute of Technology, Bangalore, India, in 2007, and the Master of Science degree (*cum laude*) from the Delft University of Technology, The Netherlands, in 2011. He is currently with the Circuits and Systems group at the Faculty of Electrical Engineering, Mathematics and Computer Science of the Delft University of Technology, The Netherlands. During his PhD thesis work (Sept. 2011-Sept. 2015), he was employed on the STW FASTCOM project and later on an NWO/STW VICI project. He has held positions at Robert Bosch, India, during 2007-2009, and Holst Centre/imec-nl, The Netherlands, during 2010-2011. He has received the Best Student Paper Award for his publication at ICASSP 2015 conference in Australia (ICASSP is the largest and a very important conference in signal processing receiving around 3000 submissions every year). His general research interest lies in the field of mathematical signal processing, statistical inference, sensor networks, and wireless communications.

