Sparsity-Cognizant Algorithms with Applications to Communications, Signal Processing, and the Smart Grid

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

Hao Zhu

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy

Georgios B. Giannakis, Advisor

August, 2012

Acknowledgements

First and foremost, my deepest gratitude goes to my advisor Professor Georgios B. Giannakis. I am very grateful to him for an exciting and fruitful experience during my graduate studies at the University of Minnesota, one of the best periods of my life so far. His advice has been extraordinary by all means, dedicating extensive amounts of time to help me through my aspiration to start thinking critically, improve technical presentation and writing skills, and overall become a good researcher. He has been a true inspiration for me, to the extent that I now have a better understanding of what it takes to succeed, and further develop my appreciation for the beauty of life.

Due thanks go also to Professors Mos Kaveh, Jarvis Haupt, Yousef Saad, Nikos Sidiropoulos, and Bruce Wollenberg, for agreeing to serve on my committee.

The work in one chapter of this Thesis is a result of a very fruitful side-by-side collaboration with Professor Geert Leus at the Delft University of Technology, Netherlands. Thanks to his visit to Minneapolis, I enjoyed a very pleasant and inspiring research collaboration experience. Other results in this Thesis further benefited from discussions with current and former members of the SPiNCOM group: Dr. Daniele Angelosante, Brian Baingana, Juan-Andrés Bazerque, Prof. Xiaodong Cai, Dr. Alfonso Cano, Dr. Emiliano Dall'Anese, Dr. Shahrokh Farahmand, Dr. Pedro Forero, Dr. Nikos Gatsis, Nitin Jain, Dr. Vassilis Kekatos, Prof. Seung-Jun Kim, Dr. Guobing Li, Morteza Mardani, Prof. Antonio Marques, Dr. Gonzalo Mateos, Dr. Eric Msechu, Dr. Ketan Rajawat, Prof. Alejandro Ribeiro, Prof. Yannis Schizas, Dr. Tairan Wang, Nasim Yahya Soltani, and Yu Zhang. At this point, I would also like to thank my colleagues at the NEC Labs, Princeton, during the summer of 2011; and especially Dr. Naranyan Prasad who shared with me many invaluable research and life experiences.

My sincere gratitude is also due to my family and friends, for their continuous

generosity in supporting me, for their precious time and sincerity to care about me, for their pleasant company and sharing of those exciting moments. In particular, I would like to thank my great friends from the Club of Amazing Mathematics & Engineering Ladies (CAMEL) at UMN, including Xingjie Li (Co-founder), Huiqiong Deng, Xiaoqing He, Hui Li, Zhihua Su, Yi Wang, and Lin Yang. Special thanks to my parents, your love is always the essential strength for me to face up everything coming to me and rejoice upon everything rejoicing me.

Hao Zhu, Minneapolis, August 3, 2012.

Dedication

To all my mentors, especially my parents Shuying and Yinhuai.

Abstract

Sparsity plays an instrumental role in a plethora of scientific fields, including statistical inference for variable selection, parsimonious signal representations, and solving under-determined systems of linear equations - what has led to the ground-breaking result of compressive sampling (CS). This Thesis leverages exciting ideas of sparse signal reconstruction to develop sparsity-cognizant algorithms, and analyze their performance. The vision is to devise tools exploiting the 'right' form of sparsity for the 'right' application domain of multiuser communication systems, array signal processing systems, and the emerging challenges in the smart power grid.

Two important power system monitoring tasks are addressed first by capitalizing on the hidden sparsity. To robustify power system state estimation, a sparse outlier model is leveraged to capture the possible corruption in every datum, while the problem nonconvexity due to nonlinear measurements is handled using the semidefinite relaxation technique. Different from existing iterative methods, the proposed algorithm approximates well the global optimum regardless of the initialization. In addition, for enhanced situational awareness, a novel sparse overcomplete representation is introduced to capture (possibly multiple) line outages, and develop real-time algorithms for solving the combinatorially complex identification problem. The proposed algorithms exhibit near-optimal performance while incurring only linear complexity in the number of lines, which makes it possible to quickly bring contingencies to attention.

This Thesis also accounts for two basic issues in CS, namely fully-perturbed models and the finite alphabet property. The sparse total least-squares (S-TLS) approach is proposed to furnish CS algorithms for fully-perturbed linear models, leading to statistically optimal and computationally efficient solvers. The S-TLS framework is well motivated for grid-based sensing applications and exhibits higher accuracy than existing sparse algorithms. On the other hand, exploiting the finite alphabet of unknown signals emerges naturally in communication systems, along with sparsity coming from the low activity of each user. Compared to approaches only accounting for either one of the two, joint exploitation of both leads to statistically optimal detectors with improved error performance.

Contents

Acknowledgements **i**

Dedication **iii**

Abstract **iv**

List of Tables **ix**

List of Figures **x**

1 Introduction **1**

1.1 Compressive Sampling and Sparse Signal Processing 3

1.1.1 Theoretical Advances 3

1.1.2 Algorithmic Developments 5

1.2 Motivation and Context 6

1.2.1 Power System Operation by Exploiting Hidden Sparsity 7

1.2.2 Fully-Perturbed Sparse Linear Regression 10

1.2.3 Exploiting Sparsity along with Finite-Alphabet Properties 11

1.3 Thesis Outline and Contributions 13

1.4 Notational conventions 15

2 Robust State Estimation by Exploiting Outlier Sparsity **17**

2.1 Modeling and Problem Statement 17

2.2 Solving R-SE using SDR 21

2.2.1 Convexifying R-SE via SDR 23

2.3 Dual Equivalence . . . 25
2.4 Numerical Tests . . . 27

3 Sparse Representations for Identifying Power Line Outages 31
3.1 Modeling and Problem Statement . . . 32
3.1.1 Linear DC Power Flow Model and its Graph Laplacian . . . 32
3.1.2 Unveiling Power Line Outages . . . 34
3.2 A Sparse Overcomplete Representation . . . 35
3.3 Sparse Line-Outage Identification . . . 38
3.3.1 Greedy Line-Outage Identification via OMP . . . 39
3.3.2 Lassoing Line Outages via CD . . . 40
3.3.3 Selection of Tuning Parameters . . . 43
3.4 Numerical Tests . . . 44
3.5 Generalizing the Overcomplete Representation . . . 49
3.5.1 Incorporating Information on Internal Noise . . . 49
3.5.2 From Line Outages to Line Changes . . . 50
3.5.3 Decoupled Power Flow Model . . . 51

4 Sparsity-Cognizant Total Least-Squares 52
4.1 Preliminaries and Problem Statement . . . 53
4.2 MAP Optimality of S-TLS for EIV Models . . . 55
4.3 S-TLS Solvers . . . 57
4.3.1 Bisection-based ε-Optimal Algorithm . . . 57
4.3.2 Alternating Descent Sub-Optimal Algorithm . . . 62
4.4 Weighted and Structured S-TLS . . . 65
4.5 S-TLS Applications . . . 71
4.5.1 Cognitive Radio Sensing . . . 71
4.5.2 DoA Estimation via Sparse Linear Regression . . . 73
4.6 Simulated Tests . . . 75

5 Exploiting Sparse User Activity in Multiuser Detection 81
5.1 Modeling and Problem Statement . . . 82
5.2 Sparsity-Exploiting MAP Detector . . . 84

5.3 Relaxed S-MAP Detectors . . . 86
5.3.1 Linear Ridge MUD . . . 87
5.3.2 Lasso-based MUD . . . 89
5.4 S-MAP Detectors with Lattice Search . . . 90
5.4.1 Sparsity-Exploiting Decision-Directed MUD . . . 91
5.4.2 Sparsity-Exploiting Sphere Decoding-based MUD . . . 92
5.5 Generalizations of S-MAP Detectors . . . 97
5.5.1 Exploiting user (in)activity across symbols . . . 97
5.5.2 Under-determined CDMA Systems . . . 99
5.5.3 Group Lassoing block activity . . . 101
5.6 Simulations . . . 102

6 Conclusions and Directions for Future Research 106
6.1 Thesis Summary . . . 106
6.2 Future research . . . 109
6.2.1 Multi-area (R-)SE using distributed SDP . . . 109
6.2.2 Recursive S-TLS for power system mode estimation . . . 110

References 113

Appendix A. Acronyms 125

Appendix B. Technical Details in Chapter 2 129
B.1 Proof of Lemma 2.1 . . . 129
B.2 Proof of Proposition 2.2 . . . 129
B.3 Proof of Proposition 2.3 . . . 130
B.4 PMU-aided SDR-based SE . . . 130

Appendix C. Technical Details in Chapter 4 133
C.1 Proof of Lemma 4.1 . . . 133
C.2 Proof of Proposition 4.2 . . . 133
C.3 Proof of Proposition 4.3 . . . 135
C.4 Proof of Proposition 4.4 . . . 135
C.5 Proof of Lemma 4.5 . . . 136

C.6 Proof of Proposition 4.7 136

Appendix D. Technical Details in Chapter 5 **137**

D.1 Proof of (5.21) 137

D.2 Proof of (5.22) 137

List of Tables

2.1 Estimation error with % of convergence for Test Case 2.1. 28
3.1 118-bus system with all bus phasor angle measurements. 45
3.2 300-bus system with all bus phasor angle measurements. 46
3.3 118-bus system with phasor angle measurements from $\mathcal{N}_I$. 47
3.4 Average running times in seconds. 49
A.1 Acronyms 125

List of Figures

1.1 Examples of sparsity in signals 2
1.2 The ℓ_1-norm and sparse signal reconstruction 4
1.3 Historic timeline of power and data networks. 8
2.1 Power system state estimation 18
2.2 Estimation errors for SE of 30-bus system 29
2.3 Estimation errors for R-SE of 30-bus system 30
3.1 Identification results for three line outages 48
4.1 Setup for cognitive radio spectrum sensing 72
4.2 Global versus local S-TLS algorithms 76
4.3 Comparing S-TLS with Lasso, TLS, and ℓ_2 regularized TLS 77
4.4 Spectrum sensing using Lasso and S-TLS 78
4.5 DoA Estimation using Lasso and S-TLS 80
5.1 Wireless sensors access a UAV. 83
5.2 Quantization thresholds for RD-MUD 102
5.3 Error performance for sub-/near-optimal S-MAP detectors 103
5.4 Exploiting user (in)activity across symbols 104
5.5 Exploiting sparsity in under-determined CDMA systems 105

Chapter 1

Introduction

Sparsity is an attribute characterizing a variety of natural and man-made signals. First of all, nature is inherently parsimonious. For example, sparse signals are naturally present due to the low occurrence of rare events, such as contingencies. The satellite images in Fig. 1.1(a)[1] depict the area affected by the 2003 U.S.-Canada Northeast blackout on both August 14, 2003 (right) and the day before (left), which started at Ohio and propagated to even some parts of Canada. This is a relatively small region when compared to the overall geographical area and also across the time horizon, especially at an early stage before the blackout cascades. But even the man-made radio spectrum is sparsely occupied across frequencies over the time and space domains. This can be illustrated by the Ofcom spectrum occupancy measurements depicted in Fig. 1.1(b)[2] over a rural area (top), near Heathrow airport (middle) and in central London (bottom).

Secondly, sparse signals could emerge due to parsimonious engineering designs. Data integrity for example, is a major challenge for the control center of various complex large-scale networks, such as the one present in a typical power system; see Fig. 1.1(c)[3]. As data are mostly collected from distant locations and telemetered to a central location, it is inevitable that a few of them may suffer from contamination or synchronicity issues. A fresh robust look at this problem is to construct an overcomplete outlier vector pertaining to all data, each entry of which is nonzero only if the corresponding

[1] Source: `http://spie.org/x16032.xml`

[2] Source: `http://stakeholders.ofcom.org.uk/market-data-research/other/technology-research/`

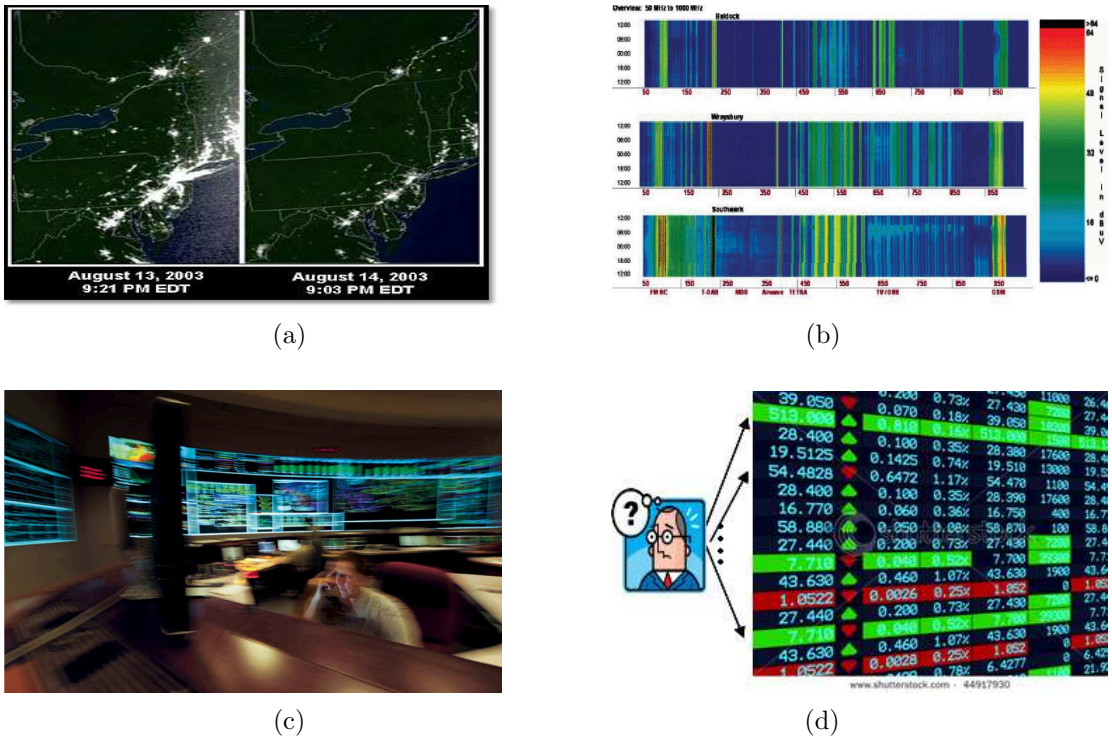

Figure 1.1: (a) Satellite images of U.S.-Canada Northeast on the night of August 13 (left) and August 14 (right), 2003; (b)Ofcom spectrum frequency measurements over frequency (horizonal axis) and time (vertical axis); (c) Control center of the PJM Interconnection LLC (PJM), with the world's largest competitive wholesale electricity market; (d) Portfolio selection on stock share investment.

datum is an outlier, and equals zero otherwise. As the number of outliers is indeed much lower than the total number of data points, the man-made outlier signal is indeed sparse.

Last but not least, sparse signals are motivated by practical constraints which encourage as few degrees of freedom as possible. Financial engineers are interested in the problem of portfolio selection, which entails possibly investing over the huge number of shares available in the market to optimize a "best-return" criterion; see e.g., Fig. 1.1(d). Constrained by the overhead fee associated with each share, it is of great importance

[3] Source: National Geographic, July 2010.

to select the most profitable portfolio with as few shares as possible.

1.1 Compressive Sampling and Sparse Signal Processing

To sum up, sparsity has emerged to play a critical role in a plethora of application areas. Reflecting back however, what sparked the growing popularity of sparsity-aware algorithms in recent years is perhaps the surge of theoretical research and development efforts toward reconstructing sparse signals. This includes the ground-breaking results on compressive sampling (CS), where sparsity has been proved instrumental to solving *under-determined* systems of linear equations.

1.1.1 Theoretical Advances

Consider the problem of reconstructing a sparse signal vector of interest $\mathbf{x}$ from an (under-determined) linear system of equations given by

$$\mathbf{y} = \mathbf{A}\mathbf{x} \tag{1.1}$$

where the data vector $\mathbf{y}$ and regression matrix $\mathbf{A}$ are given either through direct measurements or via system modeling assumptions. Provided that matrix $\mathbf{A}$ is fat and satisfies the so-termed restricted isometry property, the CS results [15, 16, 30] have established that the sparsest solution to (1.1) can be recovered exactly via minimizing the number of nonzero entries, as given by ($\mathbb{1}$ denotes the indicator function)

$$\min_{\mathbf{x}} \quad \|\mathbf{x}\|_0 := \sum_i \mathbb{1}(x_i \neq 0) \tag{1.2a}$$

$$\text{s. to} \quad \mathbf{y} = \mathbf{A}\mathbf{x}. \tag{1.2b}$$

This is a ground-breaking result, because it shows that it is possible to recover a number of unknowns with many less equations (observations), by exploiting the information on the sparsity of the unknown vector $\mathbf{x}$. However, the ℓ_0-norm minimization (1.2) is very challenging to solve. In fact, it is known to be NP-hard; see e.g., [16, 21]. Fortunately, by relaxing the ℓ_0-norm in (1.2a) to its best convex approximation, namely the ℓ_1-norm, it is still possible to find the sparsest solution to (1.1) with performance guarantees by

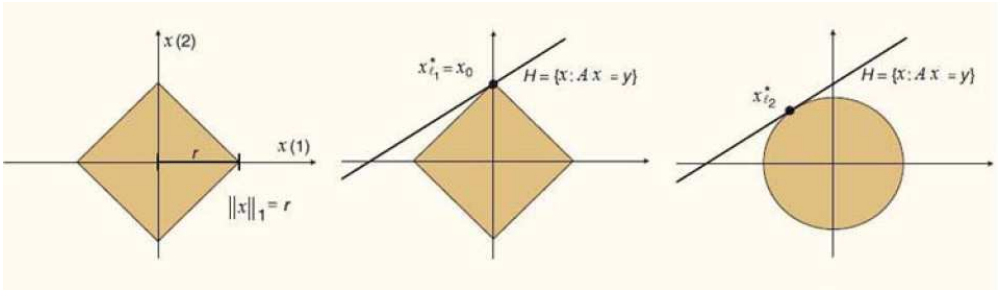

Figure 1.2: The ℓ_1 norm ball of dimensionality two (left) attains the sparse cover solution with only one nonzero entry (middle), as compared to the ℓ_2-norm ball (right).

solving the following convex problem [16, 30]

$$\min_{\mathbf{x}} \quad \|\mathbf{x}\|_1 := \sum_i |x_i| \tag{1.3a}$$

$$\text{s. to} \quad \mathbf{y} = \mathbf{A}\mathbf{x}. \tag{1.3b}$$

The ℓ_1-norm minimization (1.3) is a convex quadratic program (QP) [13, Sec. 4.4]. It can be solved efficiently using various off-the-shelf convex solvers, or, CS-tailored algorithms that will be outlined soon in Section 1.1.2. Before that, a simple example with $\mathbf{x}$ having length two is given in Fig. 1.2, to offer geometric insights on the power of the ℓ_1-norm in recovering sparse vectors. While the linear constraint in (1.3b) boils down to one single line in this setup, the ℓ_1-norm "ball" almost surely touches upon the line at its corner point on the vertical axis. This leads to a sparse reconstruction solution with only one nonzero entry, which is not typically the case with the ℓ_2-norm ball. This interesting example shows that the curvature of the ℓ_1-norm ball with acute corner points is instrumental for reconstructing sparse signals.

One practical concern is the additive noise present in data measurements. This could also be the case for the linear measurements $\mathbf{y}$ in (1.1). For the noisy setup, it is suggested to relax the equality constraint (1.3b) to a norm constraint on the reconstruction error, which is known as basis pursuit (BP) [21]. BP solves the optimization problem

$$\min_{\mathbf{x}} \quad \|\mathbf{x}\|_1 := \sum_i |x_i| \tag{1.4a}$$

$$\text{s. to} \quad \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 \leq \epsilon \tag{1.4b}$$

where $\epsilon > 0$ denotes a given tolerance. Casting BP to its equivalent Lagragian dual form leads to the well- known least-absolute shrinkage and selection operator (Lasso) [51, 83, 119], which is the solution of the unconstrained problem

$$\boxed{\min_{\mathbf{x}} \quad \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda \|\mathbf{x}\|_1} \tag{1.5}$$

for some positive penalty factor λ. While related BP ideas have leveraged sparsity to obtain parsimonious signal representations primarily in *deterministic* settings, the Lasso has been exploited vastly over the last dozen years in a broad range of *statistical* inference tasks concerned with the choice of most informative variables in linear regression models.

Collectively, CS, BP, and Lasso have been successfully applied to diverse applications involving sparse signals, ranging from the aforementioned portfolio selection problem in finance [14], to discovering interpretable biological structure in gene expression microarray data [51]. A huge literature is also available on these topics from various fields; see e.g., [5, 17, 84, 99].

1.1.2 Algorithmic Developments

Approaches to reconstructing sparse coefficient vectors of linear regression models will be briefly summarized in this subsection. They can be broadly grouped in two categories: those relying on greedy approximation schemes, and those minimizing the ℓ_1-norm of the sparse vector as outlined in Section 1.1.1.

The first category is rooted on the matching pursuit (MP) algorithm for signal approximation [66]. MP has undergone improvements which led to the orthogonal matching pursuit (OMP) algorithm [85], that draws its popularity from its computational simplicity and guaranteed performance. Albeit heuristic, greedy schemes such as the OMP have been shown capable of recovering the optimal representation of an exactly sparse vector provided that the regression matrix (a.k.a. dictionary) satisfies the so-termed 'coherence' conditions [85]. More recently, approximation guarantees for various greedy heuristic schemes including the OMP have been established using the notion of submodularity; see e.g., [27].

The second category relies on the ℓ_1-norm to yield the sparsest solution to an underdetermined linear system of equations. As mentioned earlier, popular solvers of the resultant optimization problems are variants of BP, or the Lasso. Albeit approximating

the optimal ℓ_0-norm, this relaxation has been shown to exhibit guaranteed performance of recovering sparse signal vectors; see e.g., [16]. One issue in these approaches is the tuning of regularization parameters, which can be handled using cross-validation (CV) [73] or other statistical methods. Aiming to solve a convex QP, the global minimizer of the ℓ_1-norm regularized least-squares (LS) reconstruction in Lasso can be efficiently obtained for a sequence of fixed sparsity levels using coordinate descent (CD) iterations with "warm starts" [86]. This has motivated fast solvers to obtain the regularization paths for recovering sparse solutions to general linear models using the Lasso; see e.g., [34].

1.2 Motivation and Context

Building upon both theoretical and algorithmic advances in this promising CS realm, sparse signal reconstruction has well-documented merits and enjoys increasing popularity primarily in applications, where data obey the classical linear regression model. Nonetheless, when applied to various practical problems encountered in communication, signal processing, and electrical power systems considered in the present Thesis, these state-of-the-art sparse signal recovery schemes may fall short in addressing several interesting questions, including the following ones:

(Q1) How can one uncover the sparsity of unknown signals when it is *hidden* in the problem structure?

(Q2) What if the adopted sparse linear regression models entail *perturbations* both in the data vector as well as in the regression matrix?

(Q3) Is it possible to jointly account for sparsity and *finite-alphabet* constraints on the unknowns?

While (Q1) mainly focuses on challenging problems related to smartly operating the power grid, both (Q2) and (Q3) aim at generic concerns underlying the framework of CS and sparse signal recovery, which find interesting applications in wireless communications and array signal processing problems. As we will see soon, these questions are

not only of theoretical interest, but are clearly motivated by real-life problems involving both sparsity in the unknowns and unique challenges (as well as opportunities for research), that arise as one attempts to address these three questions.

1.2.1 Power System Operation by Exploiting Hidden Sparsity

The power grid in North America has been lauded as the most important engineering achievement of the 20th century by the National Academy of Engineering [97]. A timeline comparison is depicted in Fig. 1.3 between the power grid, particularly the power transportation infrastructure, and the data transportation network. The latter has been successfully evolved to nowadays' backbone network, which is fast growing and remains for the most part freely accessible. On the contrary, the power grid today is still strictly administered by regional operators, such as the PJM in Fig. 1.1(c). Power outages are estimated to cost Americans $80 billions a year. A popular dictum from the power industry argues that "Alexander Graham Bell would not recognize today's telephone network, but Thomas Edison would feel right at home running our current electrical grid."

In addition, looking behind the curtain for outages of power lines, the report on the infamous 2003 U.S.-Canada Northeast blackout [87] identified these two major culprits, both related to power system operation concerns:

- Cause 1: "inadequate system understanding" – stated 20 times; and
- Cause 2: "inadequate situation awareness" – stated 14 times.

Besides outages, the modern power grid faces major challenges related to environmental, stability, security, and market diversity issues, which cannot be tackled solely by the domain expertise in power systems [89]. The smart grid *vision* aspires to build a transformative cyber-physical network that can address the challenges associated with the aging infrastructure by capitalizing upon state-of-the-art information technologies in sensing, control, communication, and machine learning. Important avenues where such technologies can make significant impact on power networks include: (a) agile monitoring techniques that can estimate the status of large-scale grids with precision; (b) intelligent contingency detection strategies that bring any anomalies quickly to the operators' attention, while at the same time (c) ensuring robustness and scalability

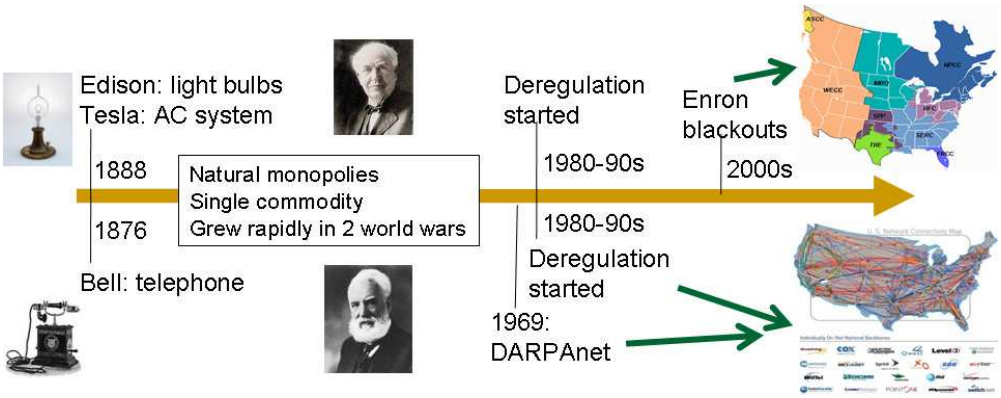

Figure 1.3: Historic timeline of power and data networks.

in the face of limited data and computational power available, in addition to possibly malicious attacks posed to this burgeoning key infrastructure.

This Thesis touches upon two intriguing challenges along these avenues, as detailed next.

Robust Power System State Estimation using Outlier Sparsity

The first challenge relates to *robust* power system *state estimation (SE)* for accurate monitoring. SE for power systems aims to obtain estimates of the voltage phasors at all buses[4] in the network [1]. In addition to outliers contaminating the measurements, the SE problem is inherently nonconvex giving rise to many local optima, due to the *nonlinear* coupling present among meter measurements and the states of interest. Therefore, Gauss-Newton iterations [11, Sec. 1.5] have been used extensively in practice, reckoned as the algorithmic foundation of (robust) SE; see e.g., [1,68]. Since this iterative procedure is in fact related to gradient descent solvers for nonconvex problems, it inevitably faces severe limitations associated with sensitivity to the choice of initialization and the resultant convergence issues [11, Sec. 1.5].

In addition to the long-time concern on noise robustness [1, Ch. 5-6], the growing interests on cyber-security issues [58,62], justify well the recent efforts devoted to account for various sources of outliers present in power meter measurements. Traditional robust SE (R-SE) techniques include the largest normalized residual (LNR) test, M-estimators,

[4] In power engineering terminology, bus refers to a node in the corresponding graph representation.

and the least-absolute value (LAV) estimation; see e.g., [1, Ch. 5-6] and [68]. However, all of them rely on the aforementioned iterative procedure, and thus suffer from convergence issues. Recently, a novel framework capitalizing on the *hidden sparsity* of the outliers has been advocated for R-SE [57, 98], but the nonlinearity therein still limits existing approaches to iterative Gauss-Newton solvers [98].

The latest trend for (robust) SE is to incorporate *linear* state measurements offered by synchronized phasor measurement units (PMUs); see e.g., [43, 100] and references therein. However, limited PMU deployment currently confines SE to mostly rely on the nonlinear meters, and its companion Gauss-Newton iterative methods. Recently, for other power system operation tasks, the semidefinite relaxation (SDR) technique [41] has been proposed for solving the optimal power flow problem [4]; see also [61], where a zero-duality gap result and guarantees for global optimality have been reported. Hence, the grand challenge for many years has been to develop an R-SE solver attaining the *global optimum* at *polynomial-time* complexity.

In Chapter 2, leveraging physical properties of nonlinear AC power systems, a novel convex SDR technique is pursued to develop polynomial-time R-SE algorithms, offering the potential to attain the globally optimum.

Sparse Representations for Identifying Power Line Outages

The second challenge considered relates to improved situational awareness for the power grid, by timely monitoring the status of transmission lines. By viewing the grid as a graph comprising interconnected power systems, PMUs provide voltage and power data per local (a.k.a. *internal*) system in real time. Likewise, real-time data are telemetered internally per system to offer topology-bearing information. On the other hand, power flow conservation across interconnected systems allows for identifying changes even in *external* system lines, which is critical for wide-area monitoring. This would have been a non-issue if inter-system data were available. Unfortunately, the system data exchange (SDX) module of the North-American Electric Reliability Corporation (NERC) can provide grid-wide interarea (a.k.a. *basecase*) topology information on an hourly basis [88], but the desiderata is near real-time monitoring of transmission lines. In a nutshell, the need arises for each internal system to identify in a computationally *efficient* manner line outages (and generally line changes) in its external counterpart relying only on

basecase topology information and local PMU data.

Existing approaches typically formulate external line-outage identification as a *combinatorially complex* (integer programming) problem, which can be computationally tractable only for single or at most double line outages [31,80,81]. However, it is becoming increasingly crucial to cope with multiple line outages in face of cascading failures that have caused recent blackouts. This has motivated several recent works [47, 48], which consider how to handle *multiple* line outage distribution factors (LODFs), as compared to coping with a single LODF. In addition, an alternative recent approach to line outage identification adopts a Gauss-Markov graphical model of the power network and can deal with multiple outages at affordable complexity, but assumes conditionally independent phasor angle measurements and requires inter-system PMU data to be available in real time *across* the grid [52]. It is challenging and important to develop a computationally efficient algorithm for near real-time identification of *multiple* external line outages (and generally changes), using only hourly basecase topology information, and *local* phasor angle data available by PMUs.

In this context, Chapter 3 proposes a overcomplete representation of the outage-induced innovation, which in turn enables casting outage identification as a *sparse* vector estimation problem, which can be solved at affordable complexity by adapting efficient CS algorithms.

1.2.2 Fully-Perturbed Sparse Linear Regression

One major hurdle to overcome by existing CS, BP, and Lasso-based approaches relates to perturbations present in the matrix $\mathbf{A}$ in (1.1), which in the BP (respectively Lasso) parlance is referred to as the representation basis or dictionary (correspondingly regression) matrix. Such perturbations appear when there is a *mismatch* between the adopted basis matrix and the actual but *unknown* one – a performance-critical issue in e.g., sparsity-exploiting approaches to localization, time delay, and Doppler estimation in communications, radar, and sonar applications [35], [65], [70], [9], [3]. In addition to model mismatch, entries of $\mathbf{A}$ in sparse portfolio selection pertain to returns of individual shares; hence, they are inevitably subject to perturbations due to any investor's limited information of the whole market. Performance analysis of CS and BP approaches for the partially-perturbed linear model with perturbations present only in the basis

matrix, as well as for the fully-perturbed one with perturbations present also in the measurements, was pursued recently in [53], [22], and [20]. But devising a systematic approach to reconstructing sparse vectors under either type of perturbed models was left open.

Interestingly, for *non-sparse* over-determined linear systems, such an approach is available within the framework of total least-squares (TLS), the basic generalization of LS tailored for fitting fully-perturbed linear models [90]. TLS and its variants involving regularization with the ℓ_2-norm of the unknown vector [77], have found widespread applications in diverse areas, including system identification with errors-in-variables (EIV), retrieval of spatial and temporal harmonics, reconstruction of medical images, and forecasting of financial data [67]. TLS was also utilized by [23] for dictionary learning, but the problem reduces to an over-determined linear system with a non-sparse unknown vector. Unfortunately, TLS approaches, with or without existing regularization terms, cannot yield consistent estimators when the linear model is under-determined, nor they account for sparsity present in the unknown vector of regression coefficients.

From a high-level vantage point, Chapter 4 develops what is termed *sparse (S-) TLS* framework for fitting sparse, perturbed, linear models, which provides CS algorithms suitable for fitting fully-perturbed linear models, as well as furnishes TLS with sparsity-cognizant regularized alternatives with improved consistency properties even for under-determined systems.

1.2.3 Exploiting Sparsity along with Finite-Alphabet Properties

Another limitation of the existing CS framework lies in the fact that it has not dealt with sparse signal vectors $\mathbf{x}$ whose entries adhere to *finite-alphabet (FA)* constraints.

Joint exploitation of sparsity and FA properties is motivated naturally in code-division multiple access (CDMA) systems, whereby user (in)activities are unknown and the activity factor (probability of each user being active) is low. Multiuser detection (MUD) algorithms play a major role for mitigating multi-access interference present in CDMA systems; see e.g., [93] and references therein. These well-appreciated MUD algorithms simultaneously detect the transmitted symbols of all active user terminals. However, they require knowledge of which terminals are active, and do not exploit possible user (in)activity. Even though sparsity alone has been used for parameter

estimation in communication systems, see e.g., [3,9,18,33], the setup therein still follows the reconstruction of real-valued unknown signals. On the other hand, the overloaded CDMA system, modeled by a under-determinant linear model (1.1), has recently become extremely popular due to the high congestion present in the modern mobile networks; see e.g., [72]. However, sparse user activity has not been considered under this setup.

Other examples for which the sparsity-aware detection is useful include selecting optimal sampling patterns in multi-coset sampling. To overcome the limitation of uniform sampling at or above the Nyquist rate for spectrally sparse multi-band signals, the so-termed 'multi-coset sampling' has been proposed at sub-Nyquist rates; see e.g., [91]. However, selecting the sampling pattern with minimal noise-induced error is left open. Challenged by combinatorial complexity, a greedy suboptimal algorithm can be applied to trade-off complexity for performance [12], while extensions to variable sampling rates has also been explored in [10]. Recently, [71] showed how to formulate this problem as a sparse detection task involving binary symbols, with greedy solvers. However, no near-optimal alternatives at affordable complexity have been offered under this setup.

The proposed framework can further be tailored for the optimal fault detection problem, arising in numerous computer-based engineering systems, including power systems; see e.g., [120] and references therein. A number of heuristics have been proposed to identify faults in linear dynamical systems, especially for detecting binary faults; see e.g., [37, 44]. Recently, convex relaxation for tackling the combinatorial complexity in identifying binary faults has been discussed in [120], using either *linear* relaxation of Boolean constraints, or the SDR technique as mentioned in Section 1.2.1. However, (near-) optimal solvers for the challenging case of under-determined linear systems remains unexplored.

In Chapter 5, a sparsity-exploiting MUD approach is introduced to incorporate user inactivity by augmenting the underlying alphabet with an extra zero point. Thus, low user activity suggests an optimal maximum *a posteriori* (MAP) criterion for detecting a sparse vector, for which various sub-/near-optimal schemes are offered by trading off error performance for computational complexity, even with under-determinant systems.

1.3 Thesis Outline and Contributions

Chapter 2 deals with robust estimation of the power system state, namely the complex voltage phasors. Challenged by the nonlinearity and gross contamination of meter measurements, the proposed approach systematically formulates the robust estimation problem as a nonlinear LS one regularized by the ℓ_1-norm of the outlier vector, and further relaxes it to a convex semidefinite program (SDP) using the SDR technique. SDR has been applied successfully to approximately solve nonconvex problems in signal processing and communications; see e.g., [64]; and more recently for convexifying the optimal power flow problem in power system operation [4, 61]. Leveraging the SDR technique for R-SE is the main novelty and contribution of Chapter 2.

Compared to existing R-SE methods relying on Gauss-Newton iterations [1, Ch. 5-6] and [98], the SDR-based R-SE formulation enables computationally efficient solvers with guaranteed convergence, regardless of the initial guess used. A dual equivalence result is provided to shed light on the approximation performance of this relaxation based algorithm. Several numerical tests are presented to corroborate the merits of the proposed SDR approaches over the Gauss-Newton method, in terms of (near-)optimal performance they exhibit. The SDR-based SE approach is further generalized to include linear synchrophasor measurements. The material of Chapter 2 has been reported in [106, 107] and [110].

To improve power system situational awareness, Chapter 3 contributes a computationally efficient algorithm for near real-time identification of *multiple* external line outages (and generally changes) using only hourly basecase topology information, and *local* voltage phasor angle data available by PMUs. It relies on the standard direct-current (DC) linear power flow model, but also applies readily to its alternating-current (AC) linearized counterpart per iteration. The novel approach views the topology-bearing basecase information as the weighted Laplacian matrix of the grid-induced graph. This leads to an overcomplete representation of the outage-induced innovation, which in turn enables casting line-outage identification as a *sparse* vector estimation problem. Solving the latter draws from recent advances in CS and sparse linear regression, and leverages the greedy OMP and the convex Lasso to identify line outages at affordable complexity. Interesting extensions of the sparse overcomplete representation are further offered.

Simulated tests using IEEE benchmark systems show that compared to the existing ones in [31, 80, 81], the proposed algorithms result in hundreds of time computational time savings, while exhibiting near-optimal identification performance. This corroborates the advantages of the proposed sparse overcomplete representation in tackling the combinatorial complexity for analyzing power system contingencies. The results of Chapter 3 have been reported in [108] and [109].

For the sparse fully-perturbed linear regression problem, Chapter 4 presents a unified framework to complement the existing CS algorithms for fitting partially-perturbed linear models, while furnishing TLS with sparsity-cognizant regularized alternatives with improved consistency. The novel framework does not require a priori information on the underlying perturbations, and in this sense S-TLS based algorithms enjoy universal applicability. The framework is flexible enough to accommodate both deterministic as well as probabilistic prior information that maybe available on the perturbed data. The practical impact is apparent to CS, BP/Lasso, and TLS-related applications involving reconstruction of a sparse vector based on data adhering to an over- or underdetermined, partially- or fully-perturbed, linear model.

Starting with an outline of the pertinent CS-BP-Lasso-TLS context, Chapter 4 introduces the S-TLS problem with two equivalent formulations, which are first used to establish the maximum *a posteriori* (MAP) optimality of S-TLS estimators, under a fully-perturbed errors-in-variables (EIV) model. Subsequently, the same formulations are utilized to develop near-optimum and reduced-complexity suboptimum S-TLS solvers with convergence guarantees. The scope of S-TLS is considerably broadened to incorporate a priori information on the deterministic structure of the data vector, the basis matrix, and/or the statistics of the perturbations, for developing weighted and structured (WS) S-TLS criteria along with associated algorithms. The impact of WSS-TLS is further demonstrated using two paradigms: cognitive radio sensing, and direction of arrival estimation with (possibly uncalibrated) antenna arrays. Simulated tests using both random settings and synthetic data illustrate the merits of the novel (WS)S-TLS framework relative to BP, Lasso, and TLS alternatives. The results of Chapter 4 are included in [114, 115] and [116].

Different from the CS settings of estimating real-valued signals, Chapter 5 develops

multiuser detectors accounting for sparsity but also for the finite alphabet of the desired sparse symbol vector. Taking into consideration the sparsity in MUD under FA constraints however, incurs complexity that is exponential in the vector length. To cope with this challenge, Chapter 5 introduces an optimal sparsity-exploiting MAP detector, which leverages sparsity to either relax or judiciously search over subsets of the alphabet for developing sub-/near-optimal alternatives. The resultant MUD algorithms trade off optimality in detection error performance for computational complexity. In a nutshell, the contribution of Chapter 5 is the development of efficient algorithms for MUD under sparsity and finite-alphabet constraints, as corroborated by simulated tests.

The sparse MUD framework is further generalized in several directions, among which the case of overloaded CDMA systems is of special interest, because it can cope with under-determined linear systems. In the CS literature, a number of works is available to analyze the performance of recovering sparse sequences of size exceeding the number of observations; see e.g., [16, 17] and references therein. These works establish the minimum number of observations needed to estimate sparse vectors with overwhelming probability. In the design of practical CDMA systems however, one is also interested in saving bandwidth and power resources. This becomes possible by reducing the size of the required spreading gain, which in turn reduces latency and energy consumption. Taking advantage of the low activity factor, enables designs with spreading gains smaller than the number of candidate users. Efficient multiuser detectors are also developed here for this kind of under-determined cases encountered in CDMA systems. The material in Chapter 5 has been reported [111] and [112].

The Thesis is summarized and interesting open problems are included in Chapter 6.

1.4 Notational conventions

The following notational conventions will be adopted throughout the subsequent chapters. Upper (lower) boldface letters will be used for matrices (column vectors); $(\cdot)^T$ denotes transposition; $(\cdot)^{\mathcal{H}}$ complex-conjugate transposition; $\mathrm{Tr}(\cdot)$ the matrix trace; $\mathrm{rank}(\cdot)$ the matrix rank; $(\cdot)^\dagger$ the matrix pseudo-inverse; $\mathrm{vec}(\cdot)$ the column-wise matrix vectorization; $\otimes$ the Kronecker product; $\mathbf{1}_{m\times n}$ the $m \times n$ matrix of all ones; $\mathbf{0}_{m\times n}$ the $m \times n$ matrix of all zeros; $\mathbf{I}$ the identity matrix; $\|\cdot\|_F$ the Frobenius norm; and $\|\cdot\|_p$ the

p-th vector norm for $p \geq 1$; $\mathrm{Re}(\cdot)$ the real part; $\mathrm{Im}(\cdot)$ the imaginary part; $\mathrm{sign}(\cdot)$ the sign operator; $\lceil \cdot \rceil$ the ceiling function; $|\cdot|$ the magnitude of a variable or the cardinality of a set; $\mathcal{S}_1 \bigcup \mathcal{S}_2$ the union of sets $\mathcal{S}_1$ and $\mathcal{S}_2$; $\mathcal{S}_1 \backslash \mathcal{S}_2$ the relative complement of the set $\mathcal{S}_2$ in the set $\mathcal{S}_1$; $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ the vector Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$; and $p[x = x'|y = y']$ the conditional probability density function (pdf) of the continuous random variable (r.v.) x taking the value x', given that the r.v. y took the value y'.

Chapter 2

Robust State Estimation by Exploiting Outlier Sparsity

An important monitoring task for power networks is to estimate accurately the underlying system state, which is useful for security-constrained dispatch and power system control. For nonlinear AC power systems, the state estimation (SE) problem is inherently nonconvex giving rise to many local optima, as well as challenged by the data integrity issue. As a result, existing robust estimators used extensively in practice rely on iterative optimization methods and restrain to two-step bad data analysis, which suffers from issues in convergence and lack of global optimality. To tackle the latter, a novel robust SE (R-SE) is formulated via capitalizing sparsity over an overcomplete outlier vector, while the *convex* semidefinite relaxation (SDR) technique is pursued to render the nonconvex problem efficiently solvable. A dual equivalence is also provided, which sheds light on the performance of the SDR-based R-SE solution of polynomial complexity. The resultant algorithms outperform existing iterative alternatives and exhibit near-optimal performance, as corroborated by numerical tests on the IEEE 30-bus benchmark system.

2.1 Modeling and Problem Statement

Consider a power transmission network with N buses denoted by the set of nodes $\mathcal{N} := \{1, \ldots, N\}$, and L transmission lines represented by the set of edges $\mathcal{E} := \{(n, m)\} \subseteq$

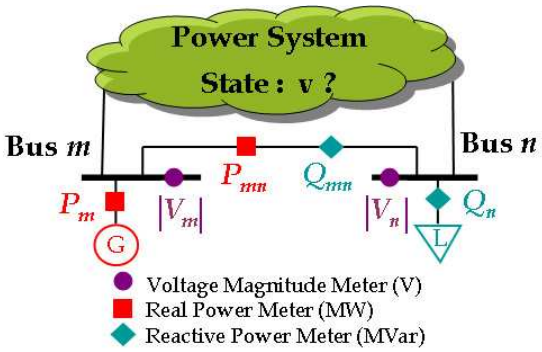

Figure 2.1: SE through measuring power system variables. (Here bus m connects to the generation, and bus n to the load.)

$\mathcal{N} \times \mathcal{N}$. To estimate the complex voltage V_n at each bus $n \in \mathcal{N}$, measurements are taken for a subset of the following system variables[1] (see also Fig. 2.1):

- $P_n(Q_n)$: the real (reactive) power injection at bus n (negative if bus n is connected to a load);
- $P_{mn}(Q_{mn})$: the real (reactive) power flow from bus m to bus n; and
- $|V_n|$: the voltage magnitude at bus n.

Compliant with the well-known AC power flow model [96, Ch. 4], these measurements are nonlinearly related with the power system state of interest, namely the voltage vector $\mathbf{v} := [V_1, \ldots, V_N]^T \in \mathbb{C}^N$. To specify this relationship, similarly collect the injected currents of all buses in $\mathbf{i} := [I_1, \ldots, I_N]^T \in \mathbb{C}^N$, and let $\mathbf{Y} \in \mathbb{C}^{N \times N}$ represent the grid's symmetric bus admittance matrix. Kirchoff's law in vector-matrix form dictates (see e.g., [8, Sec. 9.1])

$$\mathbf{i} = \mathbf{Y}\mathbf{v} \tag{2.1}$$

where the (m, n)-th entry of $\mathbf{Y}$ is given by

$$Y_{mn} := \begin{cases} -y_{mn}, & \text{if } (m,n) \in \mathcal{E} \\ y_{nn} + \sum_{\nu \in \mathcal{N}_n} y_{n\nu}, & \text{if } m = n \\ 0, & \text{otherwise} \end{cases} \tag{2.2}$$

[1] For the distribution-level SE, line current magnitude measurements may also be available [1, Sec. 2.6]; see also Remark 2.1.

with y_{mn} denoting the line admittance between buses m and n, y_{nn} bus n's admittance to the ground, and $\mathcal{N}_n$ the set of all buses linked to bus n through transmission lines.

Letting $\bar{y}_{mn}$ stand for the shunt admittance at bus m associated with the line (m,n), the current flowing from bus m to n is $I_{mn} = \bar{y}_{mn}V_m + y_{mn}(V_m - V_n)$. In addition to currents, the AC power flow model further asserts that the apparent power injection into bus n is given by $P_n + jQ_n = V_n I_n^{\mathcal{H}}$, while the apparent power flow from bus m to bus n by $P_{mn} + jQ_{mn} = V_m I_{mn}^{\mathcal{H}}$. Finally, expressing the squared bus voltage magnitude as $|V_n|^2 = V_n V_n^{\mathcal{H}}$, it is clear that all measurable quantities listed earlier are nonlinearly (in fact quadratically) related to the state $\mathbf{v}$.

Collect these (possibly noisy) measurements in the $M \times 1$ vector $\mathbf{z} := [\{\check{P}_n\}_{n\in\mathcal{N}_P}, \{\check{Q}_n\}_{n\in\mathcal{N}_Q}, \{\check{P}_{mn}\}_{(m,n)\in\mathcal{E}_P}, \{\check{Q}_{mn}\}_{(m,n)\in\mathcal{E}_Q}, \{|\check{V}_n|^2\}_{n\in\mathcal{N}_V}]^T$, where the check mark differentiates measured values from the noise-free variables[2] . The ℓ-th entry of $\mathbf{z}$ can be written as $z_\ell = h_\ell(\mathbf{v}) + \epsilon_\ell$, where $h_\ell(\cdot)$ denotes the nonlinear relationship specified according to the aforementioned AC power flow equations. The zero-mean Gaussian noise ϵ_ℓ at the ℓ-th meter is assumed independent across meters with variance σ_ℓ^2. Due to the independence among errors, the maximum-likelihood (ML) criterion for estimating $\mathbf{v}$ boils down to the weighted least-squares (WLS) one, yielding the optimal state estimator as

$$\hat{\mathbf{v}} := \arg\min_{\mathbf{v}} \sum_{\ell=1}^{M} w_\ell \, [z_\ell - h_\ell(\mathbf{v})]^2 \tag{2.3}$$

where $w_\ell := 1/\sigma_\ell^2 \; \forall \ell$. The nonlinear WLS SE formulation in (2.3) is clearly nonconvex.

The Gauss-Newton iterative solver for nonlinear WLS problems [11, Sec. 1.5] has been widely used for SE; see e.g., [1, Ch. 2] and [96, Ch. 12]. Using Taylor's expansion around a given starting point, the pure form of Gauss-Newton methods approximates the cost in (2.3) with a linear WLS one, and relies on its minimizer to initialize the subsequent iteration. This iterative procedure is closely related to gradient descent algorithms for solving nonconvex problems, see e.g., [11, Ch. 1], which are known to encounter two issues: i) sensitivity to the initial guess; and ii) convergence concerns. Typical WLS-based SE iterations start with a flat voltage profile, where all bus voltages

[2] For consistency with other measurements in Section 2.2, the squared magnitude $|V_n|^2$ is considered from now on. This is possible by adopting the model $|\check{V}_n| = |V_n| + \epsilon_V$, where ϵ_V is zero-mean Gaussian with small variance σ_V^2, to an approximate model for the squared magnitude; namely, $|\check{V}_n|^2 \approx |V_n|^2 + \epsilon_V'$, where ϵ_V' has variance $4|\check{V}_n|^2\sigma_V^2$.

are initialized with the same real number. Unfortunately, this does not guarantee convergence to the global optimum. Existing variants have asserted improved numerical stability [1, Ch. 3], but they are all limited to improving the approximate WLS cost per iteration. Recently, due to rapid developments in PMU technology, SE has benefited greatly by including the synchrophasor data, which adhere to *linear* measurement models with respect to (wrt) the unknown $\mathbf{v}$. Nonetheless, challenges emerging from the nonlinearity of legacy measurements must be addressed in the resultant PMU-aided SE methods. Apart from Chapter 2, Appendix B.4 will extend the SE framework to include the linear synchrophasor data offered by PMUs.

Apart from the nonlinearity present, another challenge present in the SE is due to grossly corrupted meter measurements (a.k.a. bad data). Statistical tests such as the largest normalized residuals of the weighted least-squares (WLS) estimation error are typically employed to reveal and remove bad data [68]. Alternatively, robust estimators, such as the least-absolute deviation, or Huber's M-estimators have also been considered; see e.g., [1, Ch. 6]. Motivated by recent advances in sparsity-aware robust statistical inference [40], a fresh look is to rewrite the ℓ-th measurement as

$$z_\ell = h_\ell(\mathbf{v}) + \epsilon_\ell + o_\ell \tag{2.4}$$

where the $M \times 1$ vector $\mathbf{o}$ is unknown with o_ℓ being nonzero if z_ℓ is a bad datum. Recovering both $\mathbf{v}$ and $\mathbf{o}$ essentially reveals the state and identifies faulty measurements. However, this casts an overcomplete model regarding $\mathbf{v}$ and $\mathbf{o}$, as the number of $\mathbf{z}$ is always less that of the unknowns. By capitalizing on the sparsity of $\mathbf{o}$ though, the goal of jointly estimating and identifying can be achieved by solving

$$\{\hat{\mathbf{v}}, \hat{\mathbf{o}}\} := \arg\min_{\mathbf{v},\mathbf{o}} \sum_{\ell=1}^{M} w_\ell \, [z_\ell - h_\ell(\mathbf{v}) - o_\ell]^2 + \lambda \|\mathbf{o}\|_1 \tag{2.5}$$

where the positive λ will be chosen later on. The ℓ_1-norm based R-SE has been considered recently with nonlineaer measurememnts in [98]. However, the pertinent performance analysis has been given for general nonlinear regression, but not for the specific nonlinear measurement models of AC power systems. Moreover, the tranditional iterative optimization framework has been adapted to solve (2.5), which has no guarantee regarding convergence or global optimality. In a nutshell, the grand challenge so far has

been to develop an R-SE solver specifically to tackle the nonlinearity in the measurement model, while attaining or approximating the *global optimum* at *polynomial-time* complexity, which is the subject of the ensuing section.

2.2 Solving R-SE using SDR

The main challenge in the R-SE problem lies in the nonlinearity of $h_\ell(\cdot)$, leading to a nonconvex WLS cost in the first summand of (2.5). This motivates us to consider how to reformulate it to a 'pseudo-linear' measurement model. Notice that all results developed in the follows are also applicable for the SE problem (2.3) without accounting for outliers. To this end, express each quadratic measurement in $\mathbf{z}$ linearly in terms of the outer-product matrix $\mathbf{V} := \mathbf{v}\mathbf{v}^{\mathcal{H}}$. Let $\{\mathbf{e}_n\}_{n=1}^N$ denote the canonical basis of $\mathbb{R}^N$, and define the following admittance-related matrices

$$\mathbf{Y}_n := \mathbf{e}_n\mathbf{e}_n^T\mathbf{Y} \tag{2.6a}$$

$$\mathbf{Y}_{mn} := (\bar{y}_{mn} + y_{mn})\mathbf{e}_m\mathbf{e}_m^T - y_{mn}\mathbf{e}_m\mathbf{e}_n^T \tag{2.6b}$$

and their related Hermitian counterparts

$$\mathbf{H}_{P,n} := \frac{1}{2}\left(\mathbf{Y}_n + \mathbf{Y}_n^{\mathcal{H}}\right), \quad \mathbf{H}_{Q,n} := \frac{j}{2}\left(\mathbf{Y}_n - \mathbf{Y}_n^{\mathcal{H}}\right) \tag{2.7a}$$

$$\mathbf{H}_{P,mn} := \frac{1}{2}\left(\mathbf{Y}_{mn} + \mathbf{Y}_{mn}^{\mathcal{H}}\right), \mathbf{H}_{Q,mn} := \frac{j}{2}\left(\mathbf{Y}_{mn} - \mathbf{Y}_{mn}^{\mathcal{H}}\right) \tag{2.7b}$$

$$\mathbf{H}_{V,n} := \mathbf{e}_n\mathbf{e}_n^T. \tag{2.7c}$$

Using these definitions, the following lemma is proved in Appendix B.1 to establish a linear model in the complex $\mathbf{V}$; its counterpart involving the real outer-product matrix of $[\mathrm{Re}(\mathbf{v})^T\ \mathrm{Im}(\mathbf{v})^T]^T$ was given in [61].

Lemma 2.1 *All error-free measurement variables are linearly related with the outer-product* $\mathbf{V}$ *as*

$$P_n = \mathrm{Tr}(\mathbf{H}_{P,n}\mathbf{V}), \qquad Q_n = \mathrm{Tr}(\mathbf{H}_{Q,n}\mathbf{V}) \tag{2.8a}$$

$$P_{mn} = \mathrm{Tr}(\mathbf{H}_{P,mn}\mathbf{V}),\ Q_{mn} = \mathrm{Tr}(\mathbf{H}_{Q,mn}\mathbf{V}) \tag{2.8b}$$

$$|V_n|^2 = \mathrm{Tr}(\mathbf{H}_{V,n}\mathbf{V}). \tag{2.8c}$$

Thus, the meter measurement z_ℓ in (2.4) *can be written as*

$$z_\ell = h_\ell(\mathbf{v}) + \epsilon_\ell + o_\ell = \mathrm{Tr}(\mathbf{H}_\ell \mathbf{V}) + \epsilon_\ell + o_\ell \tag{2.9}$$

where $\mathbf{H}_\ell$ is a Hermitian matrix specified in accordance with (2.7a)-(2.7c).

Lemma 2.1 implies the following equivalent reformulation of (2.5) [cf. (2.9)]

$$\left\{\hat{\mathbf{V}}_1, \hat{\mathbf{o}}_1\right\} := \arg\min_{\mathbf{V},\mathbf{o}} \sum_{\ell=1}^{M} w_\ell \left[z_\ell - \mathrm{Tr}(\mathbf{H}_\ell \mathbf{V}) - o_\ell\right]^2 + \lambda \|\mathbf{o}\|_1 \tag{2.10a}$$

$$\text{s.to } \mathbf{V} \in \mathbb{C}^{N\times N} \succeq \mathbf{0}, \text{and } \mathrm{rank}(\mathbf{V}) = 1 \tag{2.10b}$$

where the positive semi-definiteness and rank constraints jointly ensure that for any $\mathbf{V}$ admissible to (2.10b), there always exists a $\mathbf{v} \in \mathbb{C}^N$ such that $\mathbf{V} = \mathbf{v}\mathbf{v}^{\mathcal{H}}$.

Albeit the linearity between z_ℓ and $\mathbf{V}$ in the new formulation (2.10), nonconvexity is still present in two aspects: i) the cost in (2.10a) has degree 4 wrt the entries of $\mathbf{V}$; and ii) the rank constraint in (2.10b) is nonconvex. Aiming for an SDP formulation of (2.10), Schur's complement lemma, see e.g., [13, Appx. 5.5], can be leveraged to convert the summands in (2.10a) to a linear cost over an auxiliary vector $\boldsymbol{\chi} \in \mathbb{R}^M$. Specifically, with $\mathbf{w} := [w_1, \ldots, w_L]^T$ and likewise for $\boldsymbol{\chi}$, consider the another R-SE reformulation as

$$\left\{\hat{\mathbf{V}}_2, \hat{\mathbf{o}}_2, \hat{\boldsymbol{\chi}}_2\right\} := \arg\min_{\mathbf{V},\mathbf{o},\boldsymbol{\chi}} \sum_{\ell=1}^{M} w_\ell \chi_\ell + \lambda \|\mathbf{o}\|_1 = \arg\min_{\mathbf{V},\mathbf{o},\boldsymbol{\chi}} \mathbf{w}^T \boldsymbol{\chi} + \lambda \|\mathbf{o}\|_1 \tag{2.11a}$$

$$\text{s.to} \quad \mathbf{V} \succeq \mathbf{0}, \text{and } \mathrm{rank}(\mathbf{V}) = 1, \tag{2.11b}$$

$$\begin{bmatrix} -\chi_\ell & z_\ell - \mathrm{Tr}(\mathbf{H}_\ell \mathbf{V}) - o_\ell \\ z_\ell - \mathrm{Tr}(\mathbf{H}_\ell \mathbf{V}) - o_\ell & -1 \end{bmatrix} \preceq \mathbf{0} \ \forall \ell. \tag{2.11c}$$

The equivalence among all three R-SE formulations can be asserted in the following proposition, as proven in Appendix B.2.

Proposition 2.2 *All three nonconvex optimization problems in* (2.5), (2.10), *and* (2.11) *solve an equivalent R-SE problem under the AC flow model. For the optima of these problems, it holds that*

$$\hat{\mathbf{V}}_1 = \hat{\mathbf{V}}_2 = \hat{\mathbf{v}}\hat{\mathbf{v}}^{\mathcal{H}} \quad \text{and} \quad \hat{\chi}_{2,\ell} = \left[\hat{z}_\ell - \mathrm{Tr}(\mathbf{H}_\ell \hat{\mathbf{V}}_2)\right]^2 \ \forall \ell. \tag{2.12}$$

Remark 2.1 *(Line current magnitude measurements).* For certain distribution-level (R-)SE problems, line current magnitude measurements $|\check{I}_{mn}|$ are also available; see e.g., [1, Sec. 2.6]. It is worth pointing out that all (R-)SE reformulations can also handle this type of measurements. Since I_{mn} is linear in $\mathbf{v}$, its magnitude squared $|I_{mn}|^2$ will be quadratic in $\mathbf{v}$, and likewise linear in $\mathbf{V}$ as in (2.9). Hence, the equivalence in Proposition 2.2 and the ensuing analysis apply even with current magnitude measurements, which for brevity are henceforth omitted.

2.2.1 Convexifying R-SE via SDR

Proposition 2.2 established the relevance of the novel R-SE formulation (2.11), which is still nonconvex though, due to the rank-1 constraint. Fortunately, problem (2.11) is amenable to the SDR technique, which amounts to dropping the rank constraint and has well-appreciated merits as an optimization tool; see e.g., the seminal work of [41]. The SDR technique has also recently provided new perspectives for a number of nonconvex problems in the field of signal processing and communications, thanks to its provable performance guarantees and implementation advantages; see e.g., [64] for a tutorial treatment. The contribution here consists in permeating the benefits of this powerful optimization tool to estimating the state of AC power systems. In the spirit of SDR, relaxing the rank constraint in (2.11b) leads to the following SDP

$$\left\{\hat{\mathbf{V}}, \hat{\mathbf{o}}, \hat{\boldsymbol{\chi}}\right\} := \arg \min_{\mathbf{V},\mathbf{o},\boldsymbol{\chi}} \mathbf{w}^T\boldsymbol{\chi} + \lambda\|\mathbf{o}\|_1 \tag{2.13a}$$

$$\text{s.to} \quad \mathbf{V} \succeq \mathbf{0}, \tag{2.13b}$$

$$\begin{bmatrix} -\chi_\ell & z_\ell - \mathrm{Tr}(\mathbf{H}_\ell\mathbf{V}) - o_\ell \\ z_\ell - \mathrm{Tr}(\mathbf{H}_\ell\mathbf{V}) - o_\ell & -1 \end{bmatrix} \preceq \mathbf{0} \ \forall \ell. \tag{2.13c}$$

SDR endows R-SE with a convex SDP formulation for which efficient schemes are available to obtain the global optimum using, e.g., the interior-point solver SeDuMi [79]. The worst-case complexity of this SDP problem is $\mathcal{O}(M^4\sqrt{N}\log(1/\epsilon))$ for a given solution accuracy $\epsilon > 0$ [64]. For typical power networks, M is in the order of N, and thus the worst-case complexity becomes $\mathcal{O}(N^{4.5}\log(1/\epsilon))$. In addition, it is possible to leverage the special problem structure to further reduce complexity; see e.g., [28]. Indeed, all matrices in (2.7) are markedly sparse. For example, only the (n,n)-th entry

of $\mathbf{H}_{V,n}$ is non-zero. A closer look reveals that all $\{\mathbf{H}_\ell\}$ matrices have non-zero entries only at their diagonal entries, and those (m,n)-th off-diagonal ones that correspond to transmission lines in $\mathcal{E}$. This can be used to exploit the so-called "chordal" data structure of $\mathbf{V}$, which has led to major computational savings for solving the SDR-based OPF problem [54]. Another special structure relates to the low rank of $\{\mathbf{H}_\ell\}$ matrices, which could greatly simplify the Schur complement matrix computation step of interior-point methods [28].

Nonetheless, the SDP problem (2.13) is only a relaxed version of the equivalent R-SE in (2.11); hence, its solution $\hat{\mathbf{V}}$ may have rank greater than 1, which makes it necessary to recover a feasible estimate $\hat{\mathbf{v}}$ from $\hat{\mathbf{V}}$. This is possible by eigen-decomposing $\hat{\mathbf{V}} = \sum_{i=1}^{r} \lambda_i \mathbf{u}_i \mathbf{u}_i^{\mathcal{H}}$, where $r := \text{rank}(\hat{\mathbf{V}})$, $\lambda_1 \geq \cdots \geq \lambda_r > 0$ denote the positive ordered eigenvalues, and $\{\mathbf{u}_i \in \mathbb{C}^N\}_{i=1}^r$ are the corresponding eigenvectors. Since the best (in the minimum-norm sense) rank-one approximation of $\hat{\mathbf{V}}$ is $\lambda_1 \mathbf{u}_1 \mathbf{u}_1^{\mathcal{H}}$, the state estimate can be chosen equal to $\hat{\mathbf{v}}(\mathbf{u}_1) := \sqrt{\lambda_1}\mathbf{u}_1$.

Besides this eigenvector approach, *randomization* offers another way to extract an approximate R-SE vector from $\hat{\mathbf{V}}$, with quantifiable approximation accuracy; see e.g., [64]. The basic idea is to generate multiple Gaussian distributed random vectors $\boldsymbol{\nu} \sim \mathcal{CN}(\mathbf{0}, \hat{\mathbf{V}})$, and pick the one with the minimum error cost corresponding to the set of inlier meters $\mathcal{M}_i := \{\ell | 1 \leq \ell \leq M,\ \hat{o}_\ell \neq 0\}$. Note that although any vector $\boldsymbol{\nu}$ is feasible for (2.5), it is still possible to decrease the minimum achievable cost by rescaling to obtain $\hat{\mathbf{v}}(\boldsymbol{\nu}) = \hat{c}\boldsymbol{\nu}$, where the optimal weight can be chosen as the solution of the following convex problem as

$$\hat{c} = \arg\min_{c>0} \sum_{\ell \in \mathcal{M}_i} w_\ell \left(z_\ell - c^2 \boldsymbol{\nu}^{\mathcal{H}} \mathbf{H}_\ell \boldsymbol{\nu}\right)^2 = \sqrt{\frac{\sum_{\ell \in \mathcal{M}_i} w_\ell z_\ell \boldsymbol{\nu}^{\mathcal{H}} \mathbf{H}_\ell \boldsymbol{\nu}}{\sum_{\ell \in \mathcal{M}_i} w_\ell (\boldsymbol{\nu}^{\mathcal{H}} \mathbf{H}_\ell \boldsymbol{\nu})^2}}. \tag{2.14}$$

A couple of remarks are now in order.

Remark 2.2 *(Reference bus).* For power system SE, the reference bus convention is adopted, and the corresponding bus voltage angle is set to 0; see e.g., [96, pg. 76]. As all measurements in (2.9) are quadratically related to $\mathbf{v}$, the outer-product $(e^{j\theta}\mathbf{v})(e^{j\theta}\mathbf{v})^{\mathcal{H}} = \mathbf{v}\mathbf{v}^{\mathcal{H}}$ remains invariant to phase rotation $\theta \in [-\pi, \pi]$. To account for this, once an estimate $\hat{\mathbf{v}}$ is recovered, it can be rotated by multiplying with $\hat{V}_{\text{ref}}^{\mathcal{H}}/|\hat{V}_{\text{ref}}|$, where $\hat{V}_{\text{ref}}$ denotes the estimated reference-bus voltage.

Remark 2.3 *(Hierarchical multi-area R-SE).* Most existing (R-)SE algorithms are implemented at a central controller, namely the energy management system (EMS), to which the measurements are typically telemetered. Such central processing has two serious limitations: i) vulnerability to unreliable telemetry; and ii) high computational complexity at the EMS; see e.g., [42] and references therein. Taking advantage of the decomposable convex problem structure, it is possible to devise a *distributed* multi-area solver of the proposed SDR-based R-SE using only local communication in a consensus fashion; see also [104] for our related work in the context of distributed demodulation. Recently, efficient distributed implementations has been proposed for the optimal power flow (OPF) problem using SDP, which also suggests promising directions for computational savings in solving the SDR-based R-SE problem; see e.g., [57] for a distributed R-SE algorithm under linear measurement model.

2.3 Dual Equivalence

This section relates the relaxed and unrelaxed R-SE problems through the equivalence of their dual problems. Consider the dual of (2.13) by defining the Lagrange multiplier matrix associated with the ℓ-th inequality in (2.13c) as

$$\boldsymbol{\mu}_\ell := \begin{bmatrix} \mu_{\ell,0} & \mu_{\ell,1} \\ \mu_{\ell,1} & \mu_{\ell,2} \end{bmatrix} \succeq \mathbf{0}. \tag{2.15}$$

Using (2.15), the Lagrangian corresponding to (2.13) becomes

$$\begin{aligned} &\mathcal{L}(\mathbf{V}, \mathbf{o}, \boldsymbol{\chi}, \{\boldsymbol{\mu}_\ell\}) \\ :=&\mathbf{w}^T\boldsymbol{\chi} + \lambda\|\mathbf{o}\|_1 + \sum_{\ell=1}^{M} \mathrm{Tr}\left\{ \begin{bmatrix} -\chi_\ell & z_\ell - \mathrm{Tr}(\mathbf{H}_\ell\mathbf{V}) - o_\ell \\ z_\ell - \mathrm{Tr}(\mathbf{H}_\ell\mathbf{V}) - o_\ell & -1 \end{bmatrix} \boldsymbol{\mu}_\ell \right\} \\ =&\sum_{\ell=1}^{M} \{(2\mu_{\ell,1}z_\ell - \mu_{\ell,2}) + (w_\ell - \mu_{\ell,0})\chi_\ell - 2\mathrm{Tr}(\mu_{\ell,1}\mathbf{H}_\ell\mathbf{V}) + (\lambda|o_\ell| - 2\mu_{\ell,1}o_\ell)\}. \end{aligned} \tag{2.16}$$

The dual problem amounts to maximizing over $\boldsymbol{\mu}_\ell \succeq \mathbf{0}$, a cost equal to the minimum Lagrangian $\mathcal{L}(\mathbf{V}, \mathbf{o}, \boldsymbol{\chi}, \{\boldsymbol{\mu}_\ell\})$ over $\boldsymbol{\chi}$, $\mathbf{o}$, and $\mathbf{V} \succeq \mathbf{0}$. For the second and third summands in (2.16), at the minimum both equalities $w_\ell - \mu_{\ell,0} = 0$ and $-2\sum_{\ell=1}^{M} \mu_{\ell,1}\mathbf{H}_\ell \succeq \mathbf{0}$, have to be satisfied because of their linear relations wrt variables χ_ℓ and $\mathbf{V}$, see e.g., [13,

Sec. 4.6]. Regarding the last summand of (2.16) which involves absolute function, its minimum is attained at the point $\boldsymbol{\mu}$ with $|2\mu_{\ell,1}| \leq \lambda$, $\forall \ell$, which always takes the cost 0. This way, the dual problem can be formulated as

$$\{\hat{\boldsymbol{\mu}}_\ell\}_{\ell=1}^M := \arg\max_{\boldsymbol{\mu}_\ell} \sum_{\ell=1}^M (\mu_{\ell,2} - 2\mu_{\ell,1} z_\ell) \tag{2.17a}$$

$$\text{s.to} \quad \boldsymbol{\mu}_\ell = \begin{bmatrix} w_\ell & \mu_{\ell,1} \\ \mu_{\ell,1} & \mu_{\ell,2} \end{bmatrix} \succeq \mathbf{0}, \ |2\mu_{\ell,1}| \leq \lambda, \quad \forall \ell \tag{2.17b}$$

$$\mathbf{A} := 2\sum_{\ell=1}^M \mu_{\ell,1} \mathbf{H}_\ell \preceq \mathbf{0}. \tag{2.17c}$$

Proposition 2.3 *The SDP problem in* (2.17) *is the dual of both the R-SE in* (2.11) *and the SDR-based R-SE in* (2.13)*. Strong duality holds between* (2.17) *and* (2.13)*, while the primal variable* $\mathbf{V} \succeq \mathbf{0}$ *corresponds to the Lagrange multiplier associated with* (2.17c)*.*

The strong duality asserted by Proposition 2.3, as proven in Appendix B.3, implies that the SDR-based R-SE problem (2.13) is also the dual of (2.17), and thus Lagrangian bidual of the original R-SE (2.11); see e.g., [13, App. B]. Hence, apart from the rank relaxation interpretation in the primal domain, this provides another interpretation based on the Lagrangian dual equivalence. Additional complexity reduction is envisioned from this dualization; see e.g., [63]. The dual R-SE problem (2.17) entails $2M$ unknown variables, where for typical power networks the number of measurements M is in the order of the number of lines $|\mathcal{E}|$, which is of order N. Compared to the $N \times N$ complex matrix $\mathbf{V}$ in the primal SDP (2.13), the dual one has considerably less variables to optimize over.

Another insight of the dual R-SE comes from the relation between $\mu_{\ell,1}$ and the regularization parameter λ. The ℓ_1-norm of the outlier $\mathbf{o}$ leads to the inequality constraint for $\mu_{\ell,1}$ in (2.17b). Therefore, the larger λ is chosen, the more likely that strict inequality would hold. From the Lagrangian duality perspective, if only if $|2\mu_{\ell,1}| < \lambda$, then at the minimum of the Lagrangian $\mathcal{L}(\mathbf{V}, \mathbf{o}, \boldsymbol{\chi}, \{\boldsymbol{\mu}_\ell\})$ the corresponding outlier o_ℓ will take the value 0. This explains the effect of outlier controlling using the ℓ_1-norm, as a higher

regularization factor will lead to less number of outliers in the solution. From the limiting perspective, the R-SE problem boils down to the SE one in (2.3) if the parameter $\lambda \to \infty$. This dual insight sheds light on how to select the parameter λ, depending on the prior information regarding the inliers or the outliers (number or statistics). In general, this is possible through inspecting the solution and then adjusting the level of sparsity in the outlier vector; see [40] for more detailed discussions on this matter.

2.4 Numerical Tests

The novel SDR-based (R-)SE algorithms are tested in this section using the IEEE 30-bus system with 41 transmission lines [74], and compared to existing WLS methods that are based on Gauss-Newton iterations. The software toolbox MATPOWER [102] is used to generate the pertinent power flow and meter measurements. In addition, its SE function `doSE` has been adapted to realize the WLS Gauss-Newton iterations. The iterations terminate either upon convergence, or, once the condition number of the approximate linearization exceeds 10^8, which flags divergence of the iterates. To solve the (augmented) SDR-based SE problems, the MATLAB-based optimization modeling package `CVX` [46] is used, together with the interior-point method solver SeDuMi [79].

Test Case 2.1 : The pure WLS error criterion is first considered, to illustrate the effectiveness of SDR-based SE. The real and reactive power flows along all 41 lines are measured, together with voltage magnitudes at 30 buses. Independent Gaussian noise corrupts all measurements, with σ_ℓ equal to 0.02 at power meters, and 0.01 at voltage meters. Except for the reference bus phasor $V_{ref} = 1$, each bus has its voltage magnitude Gaussian distributed with mean 1 and variance 0.01, and its voltage angle uniformly distributed over $[-\theta,\ \theta]$. For three choices, namely $\theta = 0.3\pi$, 0.4π, and 0.5π, the empirical estimation errors $\|\mathbf{v} - \hat{\mathbf{v}}\|_2$ are averaged over 500 Monte-Carlo realizations for the SDR approach and the WLS one using three different initializations, as listed in Table 2.1. The percentage of realizations that the iterative WLS method converges is also given in parentheses. The SDR estimator is recovered from the SDR-based SE solution $\hat{\mathbf{V}}$, by picking the minimum-cost vector over the eigenvector solution and 50 randomization samples, as explained in Section 2.2.1. The first WLS estimator, termed WLS/FVP, corresponds to the WLS solution initialized by the flat-voltage profile (FVP)

Table 2.1: Estimation error with % of convergence for Test Case 2.1.

θ	SDR	WLS/FVP	WLS/DC	WLS/SDR
0.3π	0.070	0.097 (98.6%)	0.042 (100%)	0.042 (100%)
0.4π	0.081	0.593(88.6%)	0.255 (97.2%)	0.044 (100%)
0.5π	0.088	2.228 (68.6%)	1.161 (88.0%)	0.047 (100%)

point; that is, the one using the all-one vector as initial guess. For a better starting point, the second WLS/DC one is obtained by initializing the voltage angles using the DC model SE [1, Sec. 2.8], and the magnitudes using the corresponding meter measurements. To gauge the SDR approach's near-optimal performance wrt the global solution, the SDR estimator is further used to initialize the WLS iterations, and the abbreviation used for this estimator is WLS/SDR.

Table 2.1 clearly shows that the DC model based SE provides a much better initialization compared to the FVP one, in terms of smaller estimation error and higher probability of convergence. When the actual voltage angles are small ($\theta = 0.3\pi$), the WLS linear approximation is quite accurate with either the FVP or the DC model based initialization, and thus convergence to the global optimum can be guaranteed. Especially for the WLS/DC with $\theta = 0.3\pi$, the empirical error 0.042 can be considered as the benchmark estimation error achieved for such meter placements and noise levels. As θ increases however, the nonlinearity in the measurement model is responsible for the performance degradation exhibited by the WLS/FVP and WLS/DC estimators. Interestingly, estimation accuracy of the SDR estimator is still competitive to the benchmark and comes close to the global optimum. With any choice of θ, the WLS/SDR estimator is always convergent and attains the benchmark accuracy 0.042 within numerical accuracy. This suggests that the SDR-based estimator comes with numerically verifiable approximation bounds relative to the global optimum. Further evidence to this effect is provided by the empirical voltage angle and magnitude errors per bus, which are plotted in Fig. 2.2. With $\theta = 0.3\pi$, Fig. 2.2(a) demonstrates that the SDR estimator exhibits error variation similar to both WLS/DC and WLS/SDR, which is roughly twice that of these two optimal schemes. However, as θ increases to 0.4π, Fig. 2.2(b) illustrates that the WLS/DC estimator blows up due to possible divergence especially

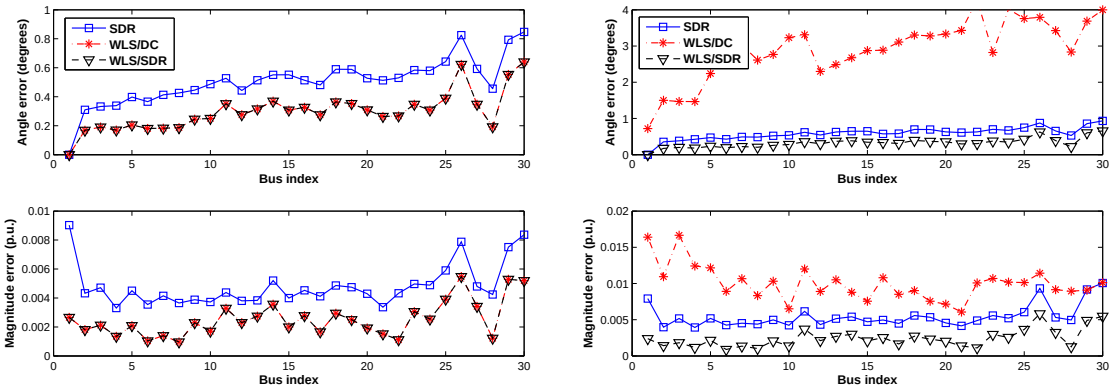

(a) (b)

Figure 2.2: Comparing estimation errors in voltage magnitudes and angles between SDR and WLS solvers at different buses for Test Case 2.1 with (a) $\theta = 0.3\pi$; and (b) $\theta = 0.4\pi$.

in the angle estimates, while both the SDR and WLS/SDR show comparable accuracy as well as analogous performance. This test case numerically supports the near-optimal performance of the proposed SDR-based SE algorithm.

Test Case 2.2 : Some preliminary results are provided to illustrate the usefulness of robust estimation. The system settings follow from Test Case 2.1 with $\theta = 0.5\pi$, and each time one power flow meter measurement is randomly chosen as a bad datum, via multiplying the meter reading by 1.2. The iterative reweighed procedure in [98] is performed to compare with the proposed SDR-based SE in (2.13), and the empirical estimation errors are plotted in Fig. 2.3. Clearly, the propose algorithm greatly reduces the effects of bad data in the estimation error in voltage phase angles, which is a more important SE performance metric than the magnitude one.

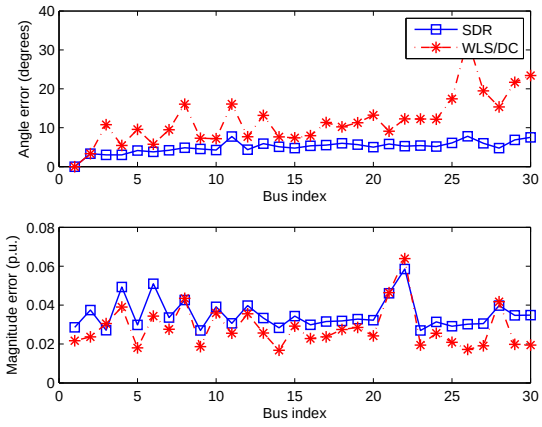

Figure 2.3: Comparing estimation errors in voltage magnitudes and angles between SDR and WLS solvers at different buses for Test Case 2.2.

Chapter 3

Sparse Representations for Identifying Power Line Outages

Fast and accurate unveiling of power line outages is of paramount importance not only for preventing faults that may lead to blackouts, but also for routine monitoring and control tasks of the smart grid, including state estimation and optimal power flow. Existing approaches are either challenged by the combinatorial complexity issues involved, and are thus limited to identifying single and double line-outages; or, they invoke less pragmatic assumptions such as conditionally independent phasor angle measurements available across the grid. Using only a subset of voltage phasor angle data, this chapter develops a near real-time algorithm for identifying multiple line outages at the affordable complexity of solving a sparse signal reconstruction problem via either greedy steps or coordinate descent iterations. Recognizing that the number of line outages is a small fraction of the total number of lines, the novel approach relies on reformulating the DC linear power flow model as a sparse overcomplete expansion, and leveraging contemporary advances in compressive sampling and variable selection. This sparse representation can also be extended to incorporate available information on the internal system, and more general line parameter faults. Analysis and simulated tests on 118-, 300-, and 2,383-bus systems confirm the effectiveness of identifying sparse power line outages.

3.1 Modeling and Problem Statement

Consider a power transmission network consisting of N buses denoted by the set of nodes $\mathcal{N} := \{1, \ldots, N\}$, and L transmission lines represented by the set of edges $\mathcal{E} := \{(m, n)\} \subseteq \mathcal{N} \times \mathcal{N}$. Collect all the voltage phasor angles $\{\theta_n\}$, one per bus $n \in \mathcal{N}$, in the vector $\boldsymbol{\theta} \in \mathbb{R}^N$; and correspondingly the injected power variables $\{P_n\}$ in $\mathbf{p} \in \mathbb{R}^N$. The network buses comprise the union $\mathcal{N} = \mathcal{N}_I \bigcup \mathcal{N}_E$ with $\mathcal{N}_I \bigcap \mathcal{N}_E = \varnothing$, where $\mathcal{N}_I$ denotes the subset of observable buses in the internal system, and $\mathcal{N}_E$ stands for the unobservable buses of the external system. Accordingly, $\boldsymbol{\theta}$ comprises two sub-vectors $\boldsymbol{\theta}_I$ and $\boldsymbol{\theta}_E$, collecting phasor angles of voltages at $\mathcal{N}_I$ and $\mathcal{N}_E$ buses, respectively; and similarly for the $\mathbf{p}_I$ and $\mathbf{p}_E$ partitions of $\mathbf{p}$. Supposing that the network-wide injected-power vector $\mathbf{p}$ only changes gradually across time, the present work aims to unveil line changes, including (possibly multiple) line outages anywhere in the network, using data $\boldsymbol{\theta}_I$ acquired in real time from PMUs. Before the problem formulation, it is important to understand how the network topology dictates the relationship between $\mathbf{p}$ and $\boldsymbol{\theta}$.

3.1.1 Linear DC Power Flow Model and its Graph Laplacian

Power flow models are useful for determining how injected power $\mathbf{p}$ flows along all transmission lines. To cope with the nonlinear AC power-voltage relationships, the DC power flow model offers a simple linear analysis tool, assuming constant voltage magnitudes per bus and negligible transmission losses [96, Sec. 4.1]. Providing a linear approximation of the actual AC system, the DC power flow model turns out to facilitate a variety of power system monitoring tasks under normal operating conditions, including security-constrained contingency analysis, and system state estimation; see e.g., [96, Ch. 11-12].

In the linear DC model, the power injected to bus n balances all the line flows originating from it; that is,

$$P_m = \sum_{n \in \mathcal{N}_m} P_{mn} = \sum_{n \in \mathcal{N}_m} \frac{1}{x_{mn}} (\theta_m - \theta_n) \tag{3.1}$$

where $x_{mn} = x_{nm}$ denotes the reactance along line (m, n), and $\mathcal{N}_m$ the set of neighboring buses linked to bus m. Writing (3.1) in vector-matrix form yields

$$\mathbf{p} = \mathbf{B}\boldsymbol{\theta} \tag{3.2}$$

where the $N \times N$ matrix $\mathbf{B}$ has its (m,n)-th entry given by

$$B_{mn} = \begin{cases} -1/(x_{mn}) & \text{if } (m,n) \in \mathcal{E} \\ \sum_{\nu \in \mathcal{N}_m} 1/(x_\nu), & \text{if } m = n \\ 0, & \text{otherwise.} \end{cases} \tag{3.3}$$

The DC power flow model (3.2) is an approximation of the AC one in (2.1), under the assumptions that the power system has minimal transmission losses and voltage amplitudes are uniform, which are very reasonable for the transmission networks.

Matrix $\mathbf{B}$, relating the voltage-phasor angle $\boldsymbol{\theta}$ to the injected power $\mathbf{p}$ as in (3.2), is uniquely determined by the line reactance parameters $\{x_{mn}\}$ of the network $(\mathcal{N}, \mathcal{E})$, and topology-bearing information, namely $\mathcal{E}$, provided by the SDX. At this point, it is worth noting that each x_{mn} reactance is present in only four entries of $\mathbf{B}$, namely, $(m,m), (m,n), (n,m)$, and (n,n). This observation about line's (m,n) presence in $\mathbf{B}$ intuitively describes the selective effect topology changes due to line outages can have on the angle vector $\boldsymbol{\theta}$ in (3.2). To formulate this topological effect concretely, $\mathbf{B}$ can be viewed as the weighted Laplacian matrix of the graph $(\mathcal{N}, \mathcal{E})$. To this end, consider the $N \times L$ bus-line incidence matrix $\mathbf{M}$, see e.g., [95, pg. 56], formed by columns $\{\mathbf{m}_\ell\}_{\ell=1}^L$ of length N. With subscript ℓ corresponding to the line (m,n), the column $\mathbf{m}_\ell$ has all its entries equal to 0 except the m-th and n-th, which take the value 1 and -1, respectively. Clearly, the rank-one matrix $\mathbf{m}_\ell \mathbf{m}_\ell^T$ has nonzero entries on its diagonal positions (m,m) and (n,n) which equal 1, and also on its (m,n) and (n,m) positions which equal -1. Thus, when summing such rank-one matrices for distinct lines, those having common m or n are superimposed on the diagonal, while those off the diagonal still have values equal to -1. This argument and (3.3) establish readily the following representation of the network topology matrix

$$\sum_{\ell=1}^{L} \frac{1}{x_\ell} \mathbf{m}_\ell \mathbf{m}_\ell^T = \mathbf{M} \mathbf{D}_x \mathbf{M}^T = \mathbf{B} \tag{3.4}$$

where the diagonal matrix $\mathbf{D}_x$ has its ℓ-th diagonal entry $1/x_\ell$ equal to the inverse reactance $1/x_{mn}$, if ℓ corresponds to the line (m,n). In addition, matrix $\mathbf{B}$ is symmetric, and has rank $N-1$ if the power network is connected, since its null space is only spanned by $\mathbf{1}$; see e.g., [95, pg. 469]. Rank deficiency of $\mathbf{B}$ gives rise to multiple solutions for $\boldsymbol{\theta}$ in (3.2). To fix this ambiguity, one generation bus is typically chosen as reference with

its phasor angle set to zero, case in which phasor angles of all other buses denote their differences relative to the reference bus; see e.g., [96, pg. 76]. Clearly, with the reference bus convention, the $(N-1)\times(N-1)$ matrix $\mathbf{B}$ has full rank, and can be formed by $\mathbf{M}$ as in (3.4) after removing the corresponding row of the incidence matrix.

3.1.2 Unveiling Power Line Outages

Given the linear DC power flow model (3.2), and the aforementioned reference bus convention, the pre-event phasor angles $\boldsymbol{\theta}$ can be uniquely characterized by the injected power vector $\mathbf{p}$ and the topology-dependent matrix $\mathbf{B}$. Suppose that due to changes in the grid, e.g., cascading failures in an early stage, several outages occur on lines, collected in the subset $\tilde{\mathcal{E}} \subseteq \mathcal{E}$. Line outages on the transmission network yield the post-event graph $(\mathcal{N}, \mathcal{E}')$, with[1] $\mathcal{E}' := \mathcal{E} \backslash \tilde{\mathcal{E}}$.

As in [80], it is assumed that fast system dynamics are well damped, and that the system settles down to a quasi-stable state following the line outages. The possibility of poor system damping can be accounted for by low-pass filtering of the measured sequential phasor angle data to smooth out any system oscillations, as detailed in [80]. Furthermore, as assumed in [81], the outaged lines in $\tilde{\mathcal{E}}$ will not result in islanding of the post-event system; that is, the underlying graph will not become disconnected. This precludes considerable changes between pre- and post-event bus power injections, such as those caused by generator outages or radial line outages. Under these considerations, the linear DC model for the quasi-stable post-event network is given by [cf. (3.2)]

$$\mathbf{p}' = \mathbf{p} + \boldsymbol{\eta} = \mathbf{B}'\boldsymbol{\theta}' \tag{3.5}$$

where $\mathbf{B}'$ is the post-event weighted Laplacian of $(\mathcal{N}, \mathcal{E}')$, and noise vector $\boldsymbol{\eta}$ accounts for the small perturbations between $\mathbf{p}'$ and $\mathbf{p}$ due to, e.g., variations in the bus loads, which can usually be modeled as zero-mean (possibly Gaussian) vector with covariance matrix $\sigma_\eta^2\mathbf{I}$; see e.g., [76].

Voltage phasor angles are only available at the subset of internal buses $\mathcal{N}_I$. Thus, unveiling outages amounts to identifying lines in $\tilde{\mathcal{E}}$, given the pre-event network-wide topology matrix $\mathbf{B}$, as well as pre- and post-event internal phasor angle vectors $\boldsymbol{\theta}_I$ and

[1] As a mnemonic, post-event quantities are denoted with prime, and the differences relative to their pre-event counterparts are denoted with tilde.

$\boldsymbol{\theta}'_I$. It will turn out that solving this problem incurs combinatorial complexity. Fortunately, using (3.2) and (3.5), the ensuing section leverages the expression (3.3) of $\mathbf{B}$ as the weighted Laplacian of the underlying graph to derive an overcomplete representation for line outages, allowing for tractable solvers of the identification problem at hand. But before this, it is important to point out the difference between the linear DC power flow model here and the probabilistic dependence graph model in [52]. The Gauss-Markov random field (GMRF) approach pursued in [52] requires the less pragmatic assumptions on the *conditional independence* among bus phasor angles, as well as the availability of $\boldsymbol{\theta}$ across the network. In contrast, these assumptions are not required for the DC model adopted here, and also in [31, 80, 81].

Remark 3.1 *(Non-random errors due to the DC power flow model approximation).* Being an approximation of the real, AC physical system behavior, the linear DC flow model may introduce sizeable non-random errors on top of the random errors captured by $\boldsymbol{\eta}$, which are due to power injection variations. Bounds on the non-random errors have been derived in [56], and tested extensively in [78]. If sizeable, such errors will render the DC model less accurate, and thus challenge *all* approaches to line outage identification that rely on the DC model. Fortunately, effective remedies are available to decrease these errors, such as using "hot-start" modeling along with unequal branch sending and receiving flows [78]. In lieu of analytic means, the effect of DC approximation errors will be assessed in Section 3.4 via numerical tests.

3.2 A Sparse Overcomplete Representation

In this section, the line outage identification task is formulated as that of recovering the coefficients of a sparse overcomplete representation of the pre- and post-event phase difference of the internal system's PMU data. To this end, the difference $\tilde{\mathbf{B}} := \mathbf{B} - \mathbf{B}'$ denoting the weighted Laplacian for the outage lines in $\tilde{\mathcal{E}}$, is written as [cf. (3.4)]

$$\tilde{\mathbf{B}} = \sum_{\ell \in \tilde{\mathcal{E}}} \frac{1}{x_\ell} \mathbf{m}_\ell \mathbf{m}_\ell^T. \tag{3.6}$$

Substituting the pre-event power flow model (3.2) into the post-event one in (3.5) leads to

$$\mathbf{B}\boldsymbol{\theta} + \boldsymbol{\eta} = \mathbf{B}'\boldsymbol{\theta}' = \mathbf{B}\boldsymbol{\theta}' - \tilde{\mathbf{B}}\boldsymbol{\theta}'. \tag{3.7}$$

Consider now the phasor angle change vector $\tilde{\boldsymbol{\theta}} := \boldsymbol{\theta}' - \boldsymbol{\theta}$, and partition it into the available $\tilde{\boldsymbol{\theta}}_I$ and the unavailable $\tilde{\boldsymbol{\theta}}_E$, each corresponding to buses in $\mathcal{N}_I$ and $\mathcal{N}_E$. With these notational conventions, substituting (3.6) into (3.7), yields

$$\begin{aligned} \mathbf{B}\tilde{\boldsymbol{\theta}} &= \tilde{\mathbf{B}}\boldsymbol{\theta}' + \boldsymbol{\eta} \\ &= \sum_{\ell \in \tilde{\mathcal{E}}} s_\ell \, \mathbf{m}_\ell + \boldsymbol{\eta} \end{aligned} \tag{3.8}$$

where $s_\ell := \mathbf{m}_\ell^T \boldsymbol{\theta}'/x_\ell$, $\forall \ell \in \tilde{\mathcal{E}}$. Solving (3.8) with respect to (wrt) $\tilde{\boldsymbol{\theta}}$, and extracting the rows corresponding to the buses in $\mathcal{N}_I$ leads to the following expression for the internal post- minus pre-event phase difference vector

$$\tilde{\boldsymbol{\theta}}_I = [\mathbf{B}^{-1}]_I \left(\sum_{\ell \in \tilde{\mathcal{E}}} s_\ell \, \mathbf{m}_\ell \right) + [\mathbf{B}^{-1}]_I \boldsymbol{\eta} \tag{3.9}$$

where $[\mathbf{B}^{-1}]_I$ is constructed from the rows corresponding to buses in $\mathcal{N}_I$, of the inverse matrix $\mathbf{B}^{-1}$.

Suppose that the number of line outages, namely $L_o := |\tilde{\mathcal{E}}|$, is given. The number of possible topologies, each having L_o outages, is clearly the number of combinations $C := \binom{L}{L_o}$. Let $\{\tilde{\mathcal{E}}^{(c)}\}_{c=1}^{C}$ denote the sets of line outages corresponding to these candidate topologies. Existing approaches exhaustively check the ℓ_2-norm of the least-squares (LS) error in the internal phase difference vector for each candidate $\tilde{\mathcal{E}}^{(c)}$, and select the minimum. That is, given $\tilde{\boldsymbol{\theta}}_I$ and $\mathbf{B}$, the chosen topology index is

$$\hat{c} := \underset{c=1,\ldots,C}{\arg\min} \left\{ \min_{\{s_\ell\}} \left\| \tilde{\boldsymbol{\theta}}_I - \left(\sum_{\ell \in \tilde{\mathcal{E}}^{(c)}} [\mathbf{B}^{-1}]_I \mathbf{m}_\ell s_\ell \right) \right\|_2^2 \right\}. \tag{3.10}$$

Such an approach incurs combinatorial complexity, and has limited existing methods based on exhaustive enumeration of all combinations, to identifying single line-outage ($L_o = 1$) [80], or at most double line-outages ($L_o = 2$) [81]. Also, it will become clear soon that the covariance information of $\boldsymbol{\eta}$ is not exploited optimally in (3.10).

To bypass this combinatorial complexity, the fresh idea here is to consider an overcomplete representation capturing all possible line outages. However, notice that the inversion in (3.9) introduces colored perturbation. To account for this, consider the compact singular value decomposition (SVD) of the fat matrix $[\mathbf{B}^{-1}]_I = \mathbf{U}_I \mathbf{\Sigma}_I \mathbf{V}_I^T$, where the square diagonal matrix $\mathbf{\Sigma}_I$ comprises all its $|\mathcal{N}_I|$ nonzero singular values. Upon defining $\mathbf{y} := \mathbf{\Sigma}_I^{-1} \mathbf{U}_I^T \tilde{\boldsymbol{\theta}}_I$, the data model in (3.9) can be reduced to a sparse linear regression one with unknown regression coefficients contained in the $L \times 1$ vector $\mathbf{s}$, whose ℓ-th entry equals s_ℓ, if $\ell \in \tilde{\mathcal{E}}$, and 0 otherwise. Thus, (3.8) can be written as

$$\mathbf{y} = \mathbf{V}_I^T \left(\sum_{\ell \in \tilde{\mathcal{E}}} s_\ell \, \mathbf{m}_\ell \right) + \mathbf{V}_I^T \boldsymbol{\eta} = \mathbf{V}_I^T \mathbf{M} \mathbf{s} + \mathbf{V}_I^T \boldsymbol{\eta} \tag{3.11}$$

$$= \mathbf{A}\mathbf{s} + \mathbf{V}_I^T \boldsymbol{\eta} \tag{3.12}$$

where the transformed incidence (a.k.a. adjacency) matrix $\mathbf{A} := \mathbf{V}_I^T \mathbf{M}$, now viewed as a regression matrix, captures all the pre-event transmission lines.

Remark 3.2 *(Perturbation whitening).* After transforming $\tilde{\boldsymbol{\theta}}_I$ to $\mathbf{y}$, the effective perturbation $\mathbf{V}_I^T \boldsymbol{\eta}$ in (3.12) is still zero-mean with covariance $\sigma_\eta^2 \mathbf{I}$. However, the enumeration based approaches in [80,81], or the mixed-integer programming one in [31] compute the ℓ_2-norm LS error for each candidate topology with line outages relying on the pre-whitened inverted linear system, and do not account for the non-white perturbation therein. Since optimal estimation (in the unbiased minimum variance sense) of linear regression models with non-white noise calls for *weighted* LS solvers, ignoring the noise color leads to suboptimal performance of the LS-based enumeration approaches that rely on (3.10), especially when the power perturbation is large. This will be illustrated further in the simulation results of Section 3.4.

In essence, the overcomplete representation in (3.12) allows one to cast the problem of recovering the subset $\tilde{\mathcal{E}}$ as one of estimating the *sparse* vector $\mathbf{s}$. Different from (3.9), the line outage set $\tilde{\mathcal{E}}$ is no longer present in (3.12), which eliminates the need for exhaustive search. The key *premise* is that the number of line outages is a small fraction of the total number of lines; i.e., $L_o := |\tilde{\mathcal{E}}| \ll L$. This holds even for multiple line outages such as those occurring during cascading failures, at least in the early stage when only a small number of lines start to fail. The same premise is adopted

for contingency analysis, where usually single and double line-outages are of primary concern; see e.g., [96, Ch. 11]. Under this sparsity constraint, the lines in outage are few, meaning that the vector $\mathbf{s}$ has only a few nonzero entries. In turn, this suggests recovering line outages by exploiting the sparsity in $\mathbf{s}$, which falls under the realm of sparse signal reconstruction approaches. Reconstruction of sparse signals has become very popular recently, especially after the emergence of compressive sampling (CS) theory; see e.g., [16,21,85] and references therein, wherein several efficient algorithms have been proposed. A related sparse overcomplete representation has been used in [45] to address the fault estimation problem for electrical circuits, under a completely different system model based on Kirchhoff's laws.

Compared to previous works constrained to at most double line-outages [31, 80, 81], the proposed formulation overcomes the inherent combinatorial complexity with a virtual "outage" at every possible transmission line. By leveraging the sparse signal representation framework, the next section will focus on algorithms offering the potential to achieve high accuracy in recovering faults even with more than two line outages, while bypassing the exhaustive search which is computationally prohibitive.

3.3 Sparse Line-Outage Identification

Approaches to reconstructing sparse coefficient vectors of linear regression models can be broadly grouped into two categories: those relying on greedy approximation schemes, and those minimizing the ℓ_1-norm of the sparse vector.

The first category is rooted on the matching pursuit (MP) algorithm for signal approximation [66]. MP has undergone improvements which led to the orthogonal matching pursuit (OMP) algorithm [85], that draws its popularity from its computational simplicity and guaranteed performance. The second category relies on the ℓ_1-norm of the wanted vector offering a convex relaxation of the ℓ_0-norm based minimization problems, which are known to be NP-hard but can yield the sparsest solution to an under-determined linear system of equations; see e.g., [16, 21]. Popular solvers of the resultant optimization problems are variants of the basis pursuit (BP) algorithm [21], or the Lasso [83]. Albeit approximating the optimal ℓ_0-norm, this relaxation has been shown to exhibit guaranteed performance of recovering sparse signal vectors; see e.g., [16]

and references therein. As a convex quadratic program (QP), the global minimizer of the ℓ_1-norm regularized Lasso problem can be efficiently obtained under various sparsity levels using coordinate descent (CD) iterations [86].

3.3.1 Greedy Line-Outage Identification via OMP

Given vector $\mathbf{y}$ obtained from $\tilde{\boldsymbol{\theta}}_I$ and $\mathbf{A}$, the greedy OMP algorithm will be outlined next for finding sparse solutions to (3.12) wrt a fixed number of nonzero entries (a.k.a. sparsity level) κ of the sought solution $\hat{\mathbf{s}}^\kappa$ to (3.12). The number κ can be directly related to the number of lines in outage, namely L_o, and its choice will be discussed in Section 3.3.3.

Greedy schemes generally aim to approximate the vector $\mathbf{y}$ by successively selecting columns $\{\mathbf{a}_\ell\}_{\ell=1}^L$ of $\mathbf{A}$ to superimpose with weights given by the wanted nonzero entries of $\mathbf{s}$. The selection criteria may be different but all share the common idea of choosing the next column as the one that correlates "best" with the current approximation error. Updating the sparsity level κ one by one, the approximation error vector per step is given by

$$\mathbf{r}^\kappa := \mathbf{y} - \mathbf{A}\hat{\mathbf{s}}^\kappa, \quad \kappa \geq 1\,; \quad \mathbf{r}^0 = \mathbf{y}\,. \tag{3.13}$$

Furthermore, with $\mathcal{L}^\kappa$ denoting the subset of indices corresponding to nonzero entries in $\hat{\mathbf{s}}^\kappa$, once a new column ℓ^κ is chosen at step κ for inclusion in $\mathcal{L}^\kappa$, the $\ell^\kappa-$th entry of $\hat{\mathbf{s}}^\kappa$ will be nonzero compared to $\hat{\mathbf{s}}^{\kappa-1}$. The orthogonality in OMP is manifested by updating the estimate $\hat{\mathbf{s}}^\kappa$ via LS fitting of $\mathbf{y}$ using all the columns in $\mathcal{L}^\kappa$, such that $\mathbf{r}^\kappa$ in (3.13) becomes orthogonal to all the chosen columns in $\mathcal{L}^\kappa$. The OMP based line-outage solution to (3.12) with maximum sparsity level $\kappa_{\max}$ is tabulated as Algorithm 3.1.

Since Algorithm 3.1 chooses the column by comparing its correlation with the residual error vector, it is useful to normalize all columns first, so that comparison is in essence performed based on the correlation coefficient. Also, due to its greedy update, the sequence of solutions $\{\hat{\mathbf{s}}^\kappa\}_{\kappa=1}^{\kappa_{\max}}$ is nested in terms of their nonzero entries, but not necessarily in the exact entry values after the orthogonal LS fitting step. In addition to computational efficiency, the simulations in Section 3.4 will demonstrate that Algorithm 3.1 exhibits also quite competitive reconstruction performance. Analytically, the latter

Algorithm 3.1 (OMP): Input $\mathbf{y}$, $\mathbf{A}$, and $\kappa_{\max}$. Output $\{\hat{\mathbf{s}}^{\kappa}\}_{\kappa=1}^{\kappa_{\max}}$.

Initialize $\mathbf{r}^0 = \mathbf{y}$, and the column subset $\mathcal{L}^0 = \emptyset$.
for $\kappa = 1, \ldots, \kappa_{\max}$ **do**
  Choose $\ell^{\kappa} \in \arg\max_{\ell} |\mathbf{a}_{\ell}^T \mathbf{r}^{\kappa-1}|$ (arbitrarily breaking ties), and update $\mathcal{L}^{\kappa} := \mathcal{L}^{\kappa-1} \bigcup \{\ell^{\kappa}\}$.
  Estimate $\hat{\mathbf{s}}^{\kappa} := \arg\min_{\mathbf{s}} \|\mathbf{y} - \mathbf{A}\mathbf{s}\|_2^2$, by setting its $\ell-$th entry $s_{\ell} = 0$ for all $\ell \notin \mathcal{L}^{\kappa}$.
  Update $\mathbf{r}^{\kappa} := \mathbf{y} - \mathbf{A}\hat{\mathbf{s}}^{\kappa}$.
end for

depends on the so-termed coherence of matrix $\mathbf{A}$ that is defined as

$$\mu := \max_{\ell \neq \ell'} \left| \mathbf{a}_{\ell}^T \mathbf{a}_{\ell'} \right|. \tag{3.14}$$

By adapting [85, Thm. B], the following proposition can be established for the present context of identifying line outages.

Proposition 3.1 *(Exact recovery of OMP). Consider the noise-free case, i.e.,* $\boldsymbol{\eta} = \mathbf{0}$ *in* (3.12)*, and suppose that the actual vector* $\mathbf{s}_o$ *has its number of nonzero entries* L_o *satisfying*

$$L_o < (\mu^{-1} + 1)/2. \tag{3.15}$$

Given the vector $\mathbf{y} = \mathbf{A}\mathbf{s}_o$ *and matrix* $\mathbf{A}$*, the OMP in Algorithm 3.1 is capable of recovering the unknown* $\mathbf{s}_o$ *exactly, if the greedy steps stop once* $\|\mathbf{r}^{\kappa}\|_2 = 0$*.*

3.3.2 Lassoing Line Outages via CD

Lasso promotes sparse solutions by regularizing the LS error criterion with the ℓ_1-norm of $\mathbf{s}$. Specifically, the coefficient vector whose nonzero entries identify line outages is found as

$$\hat{\mathbf{s}}^{\lambda} := \arg\min_{\mathbf{s}} \|\mathbf{y} - \mathbf{A}\mathbf{s}\|_2^2 + \lambda \|\mathbf{s}\|_1 \tag{3.16}$$

where λ is a tuning parameter whose choice is detailed in Section 3.3.3. Thanks to the overcomplete representation (3.12) and the ℓ_1-norm relaxation, the line outage identification is reduced to the convex QP in (3.16). Hence, its global minimizer can be efficiently obtained using general-purpose convex solvers, such as interior-point algorithms [13, Ch. 11]. Instead of these off-the-shelf solvers, the CD iterative solver [86]

will be adapted, because it has been shown to exploit the specific problem structure and can yield the so-termed regularization path of Lasso-based solutions as a function of λ [34]. This approach is advocated especially for sparse unknown vectors of high dimension; that is, for a large number of lines L.

To allow for a sequence of λ values, consider first (3.16) with a fixed λ. CD optimizes the regularized LS cost (3.16) by cyclically minimizing over the coordinates, namely, the scalar entries of $\mathbf{s}$. It yields successive estimates of each coordinate, while keeping the rest fixed. Suppose that the ℓ-th entry $s_\ell(i)$ is to be found. Precursor entries $\{s_1(i),\ldots,s_{\ell-1}(i)\}$ have been already obtained in the i-th iteration along with postcursor entries $\{s_{\ell+1}(i-1),\ldots,s_L(i-1)\}$ are also available from the previous $(i-1)$-st iteration. Thus, the effect of these given entries can be removed from $\mathbf{y}$ by forming

$$\mathbf{e}_\ell(i) := \mathbf{y} - \sum_{j=1}^{\ell-1} \mathbf{a}_j s_j(i) - \sum_{j=\ell+1}^{L} \mathbf{a}_j s_j(i-1) \tag{3.17}$$

where $\mathbf{a}_j$ denotes the $j-$th column of matrix $\mathbf{A}$. Using (3.17), the vector optimization problem in (3.16) reduces to the following scalar one with $s_\ell(i)$ as unknown: $s_\ell(i) = \arg\min_{s_\ell}[\|\mathbf{e}_\ell(i)-\mathbf{a}_\ell s_\ell\|_2^2+\lambda|s_\ell|]$. This is the popular *scalar* Lasso problem, which admits a closed-form solution expressed via a soft thresholding operator as (see e.g., [34])

$$s_\ell(i)=\text{sign}\big(\mathbf{a}_\ell^T\mathbf{e}_\ell(i)\big)\left[\frac{|\mathbf{a}_\ell^T\mathbf{e}_\ell(i)|}{\|\mathbf{a}_\ell\|_2^2}-\frac{\lambda}{2\|\mathbf{a}_\ell\|_2^2}\right]_+, \quad \forall \ell \tag{3.18}$$

where $[\chi]_+ := \chi$, if $\chi > 0$, and zero otherwise.

Cycling through the closed forms in (3.17) and (3.18) explains why CD here is faster than, and thus preferable over general-purpose convex solvers. The residual update in (3.17) is feasible iteratively, while the soft thresholding operation in (3.18) is also very fast. Following the basic convergence results in [86], the CD iteration is provably convergent to the global optimum $\hat{\mathbf{s}}^\lambda$ of (3.16), as asserted in the following proposition. The argument relies on the basic convergence result in [86], and can be referred to Appendix C.4.

Proposition 3.2 *(Convergence of CD). Given the parameter λ and arbitrary initialization, the iterates $\{\mathbf{s}(i)\}$ given by* (3.18) *converge monotonically to the global optimum $\hat{\mathbf{s}}^\lambda$ of the line outage identification problem in* (3.16).

Algorithm 3.2 (CD): Input $\mathbf{y}$, $\mathbf{A}$, and a decreasing sequence of λ values. Output $\{\hat{\mathbf{s}}^{\lambda}\}$ in (3.16) for each λ.

Initialize with $i = -1$ and $\mathbf{s}(-1) = \mathbf{0}$.
for each λ value from the decreasing sequence **do**
  repeat
    Set $i := i + 1$.
    for $\ell = 1, \ldots, L$ **do**
      Compute the residual $\mathbf{e}_{\ell}(i)$ as in (3.17).
      Update the scalar $s_{\ell}(i)$ via (3.18).
    end for
  until CD convergence is achieved.
  Save $\hat{\mathbf{s}}^{\lambda} = \mathbf{s}(i)$ for the current λ value.
  Initialize with $i = -1$ and $\mathbf{s}(-1) = \mathbf{s}(i)$ as the warm start for the next λ.
end for

The identification path is further obtained by applying the CD method for a decreasing sequence of λ values. Larger λ's in (3.18) force more entries of $\mathbf{s}(i)$ to be nulled. Hence, if a large enough parameter λ is picked, the corresponding $\hat{\mathbf{s}}^{\lambda}$ will eventually become $\mathbf{0}$. With a decreasing sequence of λ's, $\hat{\mathbf{s}}^{\lambda}$ for a large λ can be used as a *warm start* for solving (3.16) with the second largest λ. This way, the CD based line-outage solution path of (3.16), as tabulated in Algorithm 3.2, exploits both the efficient scalar solution in (3.18) as well as warm starts, to ensure reduced complexity and algorithmic stability.

Remark 3.3 *(Computational complexity).* As mentioned in Section 3.2, enumeration incurs combinatorial complexity while checking all C possible topologies $\{\tilde{\mathcal{E}}^{(c)}\}_{c=1}^{C}$ with a fixed L_o. Hence, its complexity is $\mathcal{O}(L^{L_o})$, which grows exponentially with L_o. In contrast, the recursive OMP Algorithm 3.1 has complexity $\mathcal{O}(\kappa_{\max} L)$, which is bi-linear in L and $\kappa_{\max}$. For Algorithm 3.2, each CD iteration entails only L scalar operations, while further acceleration is possible by exploiting sparsity and warm starts across CD iterations. Numerical running time comparisons will also be given in Section 3.4 to corroborate the linear complexity of the proposed algorithms.

3.3.3 Selection of Tuning Parameters

Algorithms 3.1 and 3.2 both yield a sequence of line-outage identification solutions. Selection of the parameter κ or λ however, becomes a critical issue, as it either directly or indirectly relates to the underlying sparsity level, and thus the (generally unknown) L_o. With additional prior information, existing statistical tests can be adopted in order to choose the actual number of lines in outage.

Number of line outages L_o is fixed or upper bounded by $\kappa_{\max}$

If $|\tilde{\mathcal{E}}| = L_o$ is given, the greedy OMP algorithm can automatically find the solution $\hat{\mathbf{s}}^\kappa$ with $\kappa = L_o$. As for the Lasso Algorithm 3.2, the parameter λ can also be determined if L_o is known. In fact, by directly inspecting the regularization path one can determine λ, so that the solution $\hat{\mathbf{s}}^\lambda$ has the number of nonzero entries equal to the given L_o. The indices of nonzero entries in $\hat{\mathbf{s}}^\lambda$ also correspond to the identified lines in $\tilde{\mathcal{E}}$. Note that the number of line outages L_o is assumed known and fixed in all existing works [31,52,80,81].

If the maximum number of line outages is prescribed, i.e., $L_o \leq \kappa_{\max}$, Algorithm 3.1 can output all the greedy solutions $\{\hat{\mathbf{s}}^\kappa\}$ from $\kappa = 1$ to $\kappa_{\max}$. Likewise, Algorithm 3.2 can yield solution $\hat{\mathbf{s}}^\lambda$ along the identification path with the number of nonzero entries not exceeding $\kappa_{\max}$. In order to determine the actual cardinality, one can adopt a minimum description length (MDL) type test; see e.g., [49]. Considering the OMP solutions $\{\hat{\mathbf{s}}^\kappa\}$ for example, the MDL test will specify a rating score per sparsity level κ, given by

$$\rho(\kappa) = \tfrac{1}{\sigma_\eta^2}\|\mathbf{y} - \mathbf{A}\hat{\mathbf{s}}^\kappa\|_2^2 + \kappa \log(N_I), \qquad \kappa = 1, \ldots, \kappa_{\max} \tag{3.19}$$

where $N_I := |\mathcal{N}_I|$. Based on (3.19), the chosen number of line outages will be $\hat{\kappa} := \arg\min_\kappa \rho(\kappa)$, thus yielding the identified lines in outage in the set $\mathcal{L}^{\hat{\kappa}}$ obtained from $\hat{\mathbf{s}}^{\hat{\kappa}}$.

When computing the MDL score for the Lasso solutions along the path, the LS fitting error in the first summand of (3.19) is obtained not by simply replacing $\hat{\mathbf{s}}^\kappa$ by $\hat{\mathbf{s}}^\lambda$. Instead, the zero entries of $\hat{\mathbf{s}}^\lambda$ are retained and the minimum LS cost corresponding to the nonzero entries is used in the ℓ_2-norm term in (3.19). This is because the ℓ_1-norm in the Lasso cost (3.16) introduces bias in the estimate $\hat{\mathbf{s}}^\lambda$, and this bias can be removed using the aforementioned LS fitting.

Variance of injected power noise is known

If only the variance σ_η^2 of entries in $\boldsymbol{\eta}$ is known, one can handle even an unknown number of outage lines as follows. Consider again $\hat{\mathbf{s}}^\kappa$ for each κ, and find the equivalent LS reconstruction error to obtain the sample variance of $\boldsymbol{\eta}$ given by

$$\hat{\sigma}_\eta^2(\kappa) = \frac{1}{N_I}\|\mathbf{y} - \mathbf{A}\hat{\mathbf{s}}^\kappa\|_2^2,\ \kappa = 1, \ldots, \kappa_{\max}. \tag{3.20}$$

This test selects the desired sparsity level κ as the one yielding the most accurate variance estimate as $\hat{\kappa} := \arg\min_\kappa |\hat{\sigma}_\eta^2(\kappa) - \sigma_\eta^2|$. The latter is used to identify the outage lines in $\mathcal{L}^{\hat{\kappa}}$ as described in Section 3.3.1. A similar method also applies to the Lasso approach after removing the bias in estimates $\hat{\mathbf{s}}^\lambda$, as explained earlier.

3.4 Numerical Tests

The proposed sparsity-exploiting fault diagnosis algorithms are tested in this section using first the IEEE 118-bus and 300-bus benchmark systems [74]. Both systems are initially tested with phasor angles available at all buses, in order to illustrate the merits of the proposed overcomplete representations in comparison with the optimal exhaustive search. Subsequently, line outage identification relying on a subset of phasor angles is tested on the 118-bus system partitioned as in [31]. Specifically, the internal system contains buses with indices in the set $\mathcal{N}_I := \{1-45, 113, 114, 115, 117\}$, and the external one with those in $\mathcal{N}_E := \{46-112, 116, 118\}$. The software toolbox MATPOWER [102] is used throughout to generate the phasor angle measurements as well as the pertinent power flows.

Test Case 3.1 : The 118-bus system with the complete information on phasor angles available is considered first. AC power flows are generated for both pre- and post-event systems. All possible candidate topologies with a single line-outage ($L_o = 1$), and 500 randomly chosen topologies with double line-outages ($L_o = 2$) are tested. For each line-outage topology tested, 10 realizations of the perturbation noise $\boldsymbol{\eta}$ in (3.5) are generated. Three methods are compared: Algorithms 3.1 and 3.2, and a statistically optimal exhaustive search (ES) using (3.8). Specifically, instead of solving for $\tilde{\boldsymbol{\theta}}_I$ in (3.9) and using the criterion (3.10) as proposed in [80] and [81], the optimal ES here

Table 3.1: 118-bus system with all bus phasor angle measurements.

	ES	Alg. 3.1	Alg. 3.2
Single (0%)	**95.5**	**95.5**	**95.5**
Single (1%)	92.1	92.1	**93.0**
Single (2%)	88.0	88.0	**89.8**
Single (5%)	76.8	76.8	**76.9**
Double (0%)	**95.0**	**95.0**	94.9
Double (1%)	**91.8**	91.7	91.7
Double (2%)	87.7	87.7	**88.3**
Double (5%)	76.7	**76.8**	**76.8**

utilizes the complete information of $\tilde{\boldsymbol{\theta}}$ and picks

$$\hat{c}_o := \arg \min_{c=1,\ldots,C} \left\{ \min_{\{s_\ell\}} \left\| \mathbf{B}\tilde{\boldsymbol{\theta}} - \left(\sum_{\ell \in \tilde{\mathcal{E}}^{(c)}} \mathbf{m}_\ell \; s_\ell \right) \right\|_2^2 \right\} \tag{3.21}$$

which is optimal because the noise $\boldsymbol{\eta}$ in (3.8) is white. (Contrast this with the enumeration in (3.10) used by [80, 81], which is suboptimum because the color introduced by the matrix $[\mathbf{B}]_I^{-1}$ multiplying $\boldsymbol{\eta}$ is not accounted for.) The number L_o is assumed known to all three algorithms. For both single and double line-outage tests, the standard deviation of $\boldsymbol{\eta}$ is set either equal to zero, corresponding to a noise-free case, or, equal to 1%, 2%, or 5%, of the average pre-event power injection. The percentage of correctly identified line outages is listed[2]in Table 3.1. Clearly, even under the AC flow model, Algorithms 3.1 and 3.2 exhibit performance comparable to the benchmark enumeration based schemes of (3.21), while both incur complexity that grows linearly in L, as opposed to the combinatorial search, which for $L_o = 2$ is $\mathcal{O}(L^2)$. This complexity comparison will be numerically demonstrated soon in Test Case 3.4.

For any fixed noise level, the performance of double line-outage tests does not degrade appreciably when compared to that of single line-outage tests. As the injection noise level increases, all three approaches experience performance degradation, which is expected. The OMP algorithm's performance resembles the optimal ES one thanks to its exact recovery asserted in Proposition 3.1. Interestingly, the ES method is not

[2] In all tables, "ES" stands for the exhaustive search algorithm, and the number in boldface font indicates the best performance per row.

Table 3.2: 300-bus system with all bus phasor angle measurements.

	ES	Alg. 3.1	Alg. 3.2
Single (0%)	**96.6**	**96.6**	**96.6**
Single (1%)	**82.4**	**82.4**	82.0
Single (2%)	39.6	39.6	**40.3**
Double (0%)	92.4	**92.6**	90.2
Double (1%)	79.1	**79.2**	74.2
Double (2%)	36.5	**36.7**	33.9

always the best one, while the Lasso solution of Algorithm 3.2 in certain cases exhibits improved noise resilience. This can be attributed to two reasons. First, the *empirical* identification rates in Table 3.1 depend on the random line-outage topology and injection noise realizations, and may not be identical with the exact rates. This explains why results for all three approaches are numerically comparable. Second, statistical optimality of the ES criterion (3.21) implicitly assumes that the variable s_ℓ is uniformly distributed. However, it corresponds to the post-outage power flow for any $\ell \in \mathcal{E}$ [cf. (3.8)], and differs from line to line. Hence, it is possible that the ℓ_1-norm in the Lasso cost (3.16) aids discerning the wrong non-zero entries induced by the noise.

Test Case 3.2 : The IEEE 300-bus system is tested here, keeping all other parameters identical to those of Test Case 3.1. The percentage of correctly identified line outages is listed in Table 3.2, which also shows that the proposed algorithms are quite competitive wrt the enumeration based benchmark. Similar observations are also made when comparing different noise levels and the results of single and double line-outage cases. This effect of injection noises to performance degradation becomes more pronounced, and should be attributed to the approximation induced by the linear DC model for larger systems (cf. Remark 3.1).

Test Case 3.3 : Based on the aforementioned internal and external grid partition, the 118-bus system is tested here using only part of the phasor angles in $\boldsymbol{\theta}_I$. All other settings are as in Test Case 3.1, and the results are listed in Table 3.3. The suboptimum exhaustive search in (3.10) without whitening the colored noise is also included for

Table 3.3: 118-bus system with phasor angle measurements from $\mathcal{N}_I$.

	ES without whitening	ES with whitening	Alg. 3.1	Alg. 3.2
Single (0%)	64.8	**67.2**	**67.2**	64.0
Single (1%)	34.9	**60.8**	**60.8**	57.7
Single (2%)	25.3	**56.9**	**56.9**	55.6
Single (5%)	12.5	**52.1**	**52.1**	51.2
Double (0%)	59.1	**63.9**	63.8	56.1
Double (1%)	45.2	**60.4**	**60.4**	52.3
Double (2%)	36.5	53.1	**53.2**	49.0
Double (5%)	27.3	**50.7**	**50.7**	45.6

comparison, and performs much worse than other (near-) optimal algorithms, especially when the noise level is high (5%). On the contrary, the exhaustive search based on the whitened system (3.12), given by

$$\hat{c}_o := \arg \min_{c=1,\ldots,C} \left\{ \min_{\{s_\ell\}} \left\| \mathbf{y} - \left(\sum_{\ell \in \tilde{\mathcal{E}}^{(c)}} \mathbf{a}_\ell s_\ell \right) \right\|_2^2 \right\} \tag{3.22}$$

is optimal in the unbiased minimum variance sense, and for this reason it is used to benchmark the performance of others. Table 3.3 speaks for the importance of the whitening step (3.11). With partial phasor angle measurements, performance of the three (near-) optimal algorithms degrades relative to the ones in Test Case 3.1. This is expected because observability for the 118-bus system is compromised if only angle measurements from buses in $\mathcal{N}_I$ are available. It is also worth stressing that the proposed schemes, especially the OMP Algorithm 3.1, still perform comparably to the benchmark scheme based on (3.22).

To further illustrate the performance of the proposed algorithms, the lines in outage were considered drawn from the set $\tilde{\mathcal{E}} := \{(4, 11), (42, 49), (65, 68)\}$, corresponding to lines indexed by $\ell = 10, 66, 103$. Notice that two of these three lines are linked to buses in the unobservable external system $\mathcal{N}_E$ at both sides. The noise level is set equal to 2% of the average pre-event power injection. Algorithm 3.1 is simulated with the maximum $\kappa_{\max} := 6$, and outputs the sequence of identified lines in $\mathcal{L}^6 = \{66, 10, 103, 42, 110, 34\}$. Clearly, the set $\tilde{\mathcal{E}}$ is identified exactly from the first three elements of $\mathcal{L}^6$. Subsequently,

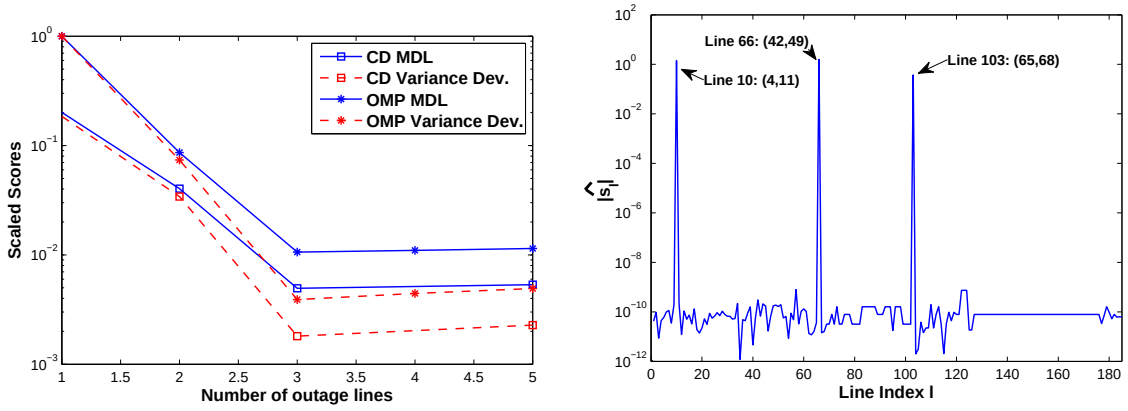

(a) (b)

Figure 3.1: (a) Scaled scores versus a variable L_o for MDL and variance tests; (b) Absolute value of entries of $\hat{s}_\ell^\lambda$ versus the index ℓ, for the λ value yielding three line outages..

Algorithm 3.2 is applied to obtain the line outage identification path for $\hat{\mathbf{s}}^\lambda$ with an exponentially decreasing sequence of 20 values of λ.

When L_o is unknown, the tests based on the MDL and the sample variance deviation criterion are implemented with the number of outaged lines ranging from 1 to 5. The test scores are obtained by scaling relative to their respective maximum ones, and are plotted in Fig. 3.1(a). The plot demonstrates that the minimum scores are achieved at $|\tilde{\mathcal{E}}| = 3$. Clearly, this is also the right κ for the OMP, which confirms that $\mathcal{L}^\kappa$ correctly identifies $\tilde{\mathcal{E}}$. For Algorithm 3.2, Fig. 3.1(b) depicts the absolute value of the $\hat{\mathbf{s}}^\lambda$ entries, corresponding to a λ value that yields three line outages. All entries in $\hat{\mathbf{s}}^\lambda$ are numerically equal to zero, except for three positions, coinciding exactly with lines in $\tilde{\mathcal{E}}$. This clearly demonstrates the effectiveness of the proposed algorithms.

Test Case 3.4 : To gauge computational complexity of the proposed algorithms, their running times are provided for the 118-, 300-, and 2,383-bus systems. While the first two are the aforementioned IEEE benchmark systems from [74], the last one is from the Polish power system provided by MATPOWER 'case2383wp' casefile. All algorithms are

Table 3.4: Average running times in seconds.

	ES	Alg. 3.1	Alg. 3.2
Single, 118-bus	5.9e-2	1.3e-4	4.9e-2
Double, 118-bus	4.1	2.7e-4	9.7e-2
Single, 300-bus	0.15	4.0e-4	0.39
Double, 300-bus	22.5	1.1e-3	0.97
Single, 2383-bus	1.2	1.4e-2	2.6
Double, 2383-bus	$\sim$	2.9e-2	4.2

run using the MATLAB® R2011a software, on a typical Windows XP computer with a 2.8GHz CPU. The running times for both single and double line-outages averaged over 500 Monte-Carlo iterations are listed in Table 3.4. The OMP Algorithm 3.1 clearly outperforms the rest in running time, and scales well with the number of buses. For both Algorithms 3.1 and 3.2, running time for double-line outages is about twice that for single-line outages[3] . However, the combinatorially complex ES approach does not scale well with the system size and the number of line outages, resulting in hundreds of times more computational time for double line-outage tests with the 118- or the 300-bus system. As this increase grows to be in the order of thousands for the 2,383-bus system, it is omitted here.

3.5 Generalizing the Overcomplete Representation

This section generalizes the novel notion of sparse overcomplete representation for identifying line outages, along with Algorithms 3.1 and 3.2, to three directions of practical interest.

3.5.1 Incorporating Information on Internal Noise

It is often the case that the internal system in $\mathcal{N}_I$ can acquire $\mathbf{p}_I$ in real time, using PMU data and also state estimates obtained from legacy meters; see e.g., [31]. As the

[3] Although Algorithm 3.2 may not appear as attractive, there are tailored programs in [34] to compute the entire identification path within 1 second for systems involving as many as 5,000 unknowns.

noise vector $\boldsymbol{\eta}$ accounts for the perturbations present in the post-event power vector relative to the pre-event power vector, the internal system is able to recover exactly the part $\boldsymbol{\eta}_I$ of this noise vector, since it can measure both $\mathbf{p}_I$ and $\mathbf{p}'_I$ at the buses in $\mathcal{N}_I$. Upon forming $\boldsymbol{\eta}_I := \mathbf{p}'_I - \mathbf{p}_I$, this internal noise vector can be removed from both sides of (3.8). After this "internal noise" cancelation, $\boldsymbol{\eta}$ is no longer white, and likewise $\mathbf{V}_I^T\boldsymbol{\eta}$ in (3.12) is colored too. While the internal noise is absent ($\boldsymbol{\eta}_I = \mathbf{0}$), the external one ($\boldsymbol{\eta}_E$) is still zero-mean white with covariance $\sigma_\eta^2\mathbf{I}$.

To account for the absence of internal noise, let $[\mathbf{B}^{-1}]_{IE}$ denote the submatrix of $\mathbf{B}^{-1}$ formed by rows and columns from buses in $\mathcal{N}_I$ and $\mathcal{N}_E$, respectively. Using the latter, consider replacing the transformed noise $[\mathbf{B}^{-1}]_I\boldsymbol{\eta}$ in (3.9) with $[\mathbf{B}^{-1}]_{IE}\boldsymbol{\eta}_E$. It now becomes possible to employ a similar whitening procedure using the SVD $[\mathbf{B}^{-1}]_{IE} = \mathbf{U}_{IE}\boldsymbol{\Sigma}_{IE}\mathbf{V}_{IE}^T$ to obtain the data vector

$$\check{\mathbf{y}} := \boldsymbol{\Sigma}_{IE}^{-1}\mathbf{U}_{IE}^T\tilde{\boldsymbol{\theta}}_I = \boldsymbol{\Sigma}_{IE}^{-1}\mathbf{U}_{IE}^T[\mathbf{B}^{-1}]_I\mathbf{M}\mathbf{s} + \mathbf{V}_{IE}^T\boldsymbol{\eta}_E\,. \tag{3.23}$$

Apart from a different regression matrix, the modified data model (3.23) entails again a sparse overcomplete representation, and thus Algorithms 3.1 and 3.2 can be readily applied. It is expected that by incorporating the extra information on the internal system noise will markedly enhance accuracy without sacrificing computational efficiency.

3.5.2 From Line Outages to Line Changes

A second extension of interest is to allow for more general *line faults* including transmission line parameter changes. As suggested by the expression of $\tilde{\mathbf{B}}$ in (3.6), changes in the inverse reactance $\{1/x_\ell\}$ of lines in outage manifest themselves in changes of the difference matrix $\tilde{\mathbf{B}}$ in the pre- and post-event network Laplacians. Suppose that instead of having only line outages, a subset of lines experiences change in reactance parameters, giving rise to a modified set of lines $\tilde{\mathcal{E}}_{\text{mod}}$. Clearly, $\tilde{\mathcal{E}}_{\text{mod}}$ can also capture line outages which correspond to line reactance parameters going to infinity. Suppose that the ℓ-th line experiences a sudden change in its reactance from x_ℓ to x'_ℓ, for all $\ell \in \tilde{\mathcal{E}}'$. This can be also captured by the matrix

$$\tilde{\mathbf{B}} := \mathbf{B} - \mathbf{B}' = \sum_{\ell\in\tilde{\mathcal{E}}_{\text{mod}}} \delta_\ell \mathbf{m}_\ell\mathbf{m}_\ell^T \tag{3.24}$$

where the difference in the inverse reactance $\delta_\ell := \frac{1}{x_\ell} - \frac{1}{x'_\ell}$. Clearly, a line outage is a special type of a line change, since if the ℓ-th line is in outage it holds that $x'_\ell \to \infty$, and thus $\delta_\ell = 1/x_\ell$. However, it is worth noticing that the system model (3.12) still holds with $s_\ell := \delta_\ell \mathbf{m}_\ell^T \boldsymbol{\theta}'$ for $\ell \in \tilde{\mathcal{E}}_{\text{mod}}$. Thus, the general line change identification problem can again be solved under the sparse signal reconstruction framework [cf. (3.16)].

3.5.3 Decoupled Power Flow Model

The DC power flow model discussed so far actually coincides with the system obtained by linearizing the decoupled *nonlinear* AC power flow model obtained after the appropriate parameter simplifications, see e.g., [8, Sec. 10.7]. To handle nonlinear power flows, iterative approximations are typically employed. Among the most popular approximation methods, the fast decoupled power flow model is preferred as an efficient linearization method as it only needs four to seven iterations to converge in practice; see e.g., [8, Sec. 10.7]. Compared to the DC power flow involving phasor angles and real power injection, the fast decoupled model includes an additional system of equations, relating the phasor amplitudes to the reactive power injection. Interestingly, both transformation matrices in the fast decoupled model are also related to the Laplacian $\mathbf{B}$, and thus they can be used to infer changes in the grid topology.

Chapter 4

Sparsity-Cognizant Total Least-Squares

Solving linear regression problems based on the total least-squares (TLS) criterion has well-documented merits in various applications, where perturbations appear both in the data vector as well as in the regression matrix. However, existing TLS approaches do not account for sparsity possibly present in the unknown vector of regression coefficients. On the other hand, sparsity is the key attribute exploited by modern compressive sampling and variable selection approaches to linear regression, which include noise in the data, but do not account for perturbations in the regression matrix. The present chapter fills this gap by formulating and solving (regularized) TLS optimization problems under sparsity constraints. Near-optimum and reduced-complexity suboptimum sparse (S-) TLS algorithms are developed to address the perturbed compressive sampling (and the related dictionary learning) challenge, when there is a mismatch between the true and adopted bases over which the unknown vector is sparse. The novel S-TLS schemes also allow for perturbations in the regression matrix of the least-absolute selection and shrinkage selection operator (Lasso), and endow TLS approaches with ability to cope with sparse, under-determined "errors-in-variables" models. Interesting generalizations can further exploit prior knowledge on the perturbations to obtain novel weighted and structured S-TLS solvers. Analysis and simulations demonstrate the practical impact of S-TLS in calibrating the mismatch effects of contemporary grid-based approaches to

cognitive radio sensing, and robust direction-of-arrival estimation using antenna arrays.

4.1 Preliminaries and Problem Statement

Consider the *under-determined* linear system of equations, $\mathbf{y} = \mathbf{C}\boldsymbol{\theta}_o$, where the unknown $n \times 1$ vector $\boldsymbol{\theta}_o$ is to be recovered from the given $m \times 1$ data vector $\mathbf{y}$ and the $m \times n$ matrix $\mathbf{C}$. With $m < n$ and no further assumption, only approximations of $\boldsymbol{\theta}_o$ are possible using the minimum-norm solution; or, the least-squares (LS) regularized by the ℓ_2-norm, which solves in closed form the quadratic problem: $\min_{\boldsymbol{\theta}} \|\mathbf{y} - \mathbf{C}\boldsymbol{\theta}\|_2^2 + \gamma\|\boldsymbol{\theta}\|_2^2$ for some chosen $\gamma > 0$. Suppose instead that over a known basis matrix $\mathbf{B}$, the unknown vector satisfies $\boldsymbol{\theta}_o = \mathbf{B}\mathbf{x}_o$ with $\mathbf{x}_o$ being *sparse*, meaning that:

(as4.0) *The $n \times 1$ vector $\mathbf{x}_o$ contains more than $n-m$ zero elements at unknown entries.*

Under (as4.0) and certain conditions on the matrix $\mathbf{A} := \mathbf{CB}$, compressive sampling (CS) theory asserts that exact recovery of $\mathbf{x}_o$ can be guaranteed by solving the non-convex, combinatorially complex problem: $\min_{\mathbf{x}} \|\mathbf{x}\|_0$ subject to (s.to) $\mathbf{y} = \mathbf{A}\mathbf{x}$. More interestingly, the same assertion holds with quantifiable chances if one relaxes the ℓ_0- via the ℓ_1-norm, and solves efficiently the convex problem: $\min_{\mathbf{x}} \|\mathbf{x}\|_1$ s.to $\mathbf{y} = \mathbf{A}\mathbf{x}$ [5,15,21].

Suppose now that due to data perturbations the available vector $\mathbf{y}$ adheres only approximately to the linear model $\mathbf{A}\mathbf{x}_o$. The ℓ_1-norm based formulation accounting for the said perturbations is known as basis pursuit (BP) [21], and the corresponding convex problem written in its Lagrangian form is: $\min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda_1\|\mathbf{x}\|_1$, where $\lambda_1 > 0$ is a sparsity-tuning parameter. (For large λ_1, the solution is driven toward the all-zero vector; whereas for small λ_1 it tends to the LS solution.) This form of BP coincides with the Lasso approach developed for variable selection in linear regression problems [51, 83]. For uniformity with related problems, the BP/Lasso solvers can be equivalently written as

$$\{\hat{\mathbf{x}}_{Lasso}, \hat{\mathbf{e}}_{Lasso}\} := \arg\min_{\mathbf{x},\mathbf{e}} \ \|\mathbf{e}\|_2^2 + \lambda_1\|\mathbf{x}\|_1 \tag{4.1a}$$

$$\text{s. to} \ \ \mathbf{y} + \mathbf{e} = \mathbf{A}\mathbf{x}\,. \tag{4.1b}$$

Two interesting questions arise at this point: i) How is the performance of CS and BP/Lasso based reconstruction affected if perturbations appear also in $\mathbf{A}$? and ii)

How can sparse vectors be efficiently reconstructed from over- and especially under-determined linear regression models while accounting for perturbations present in $\mathbf{y}$ and/or $\mathbf{A}$?

In the context of CS, perturbations in $\mathbf{A}$ can be due to disturbances in the compressing matrix $\mathbf{C}$, in the basis matrix $\mathbf{B}$, or in both. Those in $\mathbf{C}$ can be due to non-idealities in the analog implementation of CS; while those in $\mathbf{B}$ can also emerge because of mismatch between the adopted basis $\mathbf{B}$ and the actual one, which being unknown, is modeled as $\mathbf{B}+\mathbf{E}_B$. This mismatch emerges with grid-based approaches to localization, time delay, and spatio-temporal frequency or Doppler estimation [3, 6, 9, 19, 35, 36]. In these applications, the entries of $\boldsymbol{\theta}_o$ have e.g., a sparse discrete-time Fourier transform with peaks off the frequency grid $\{2\pi k/n\}_{k=0}^{n-1}$, but the postulated $\mathbf{B}$ is the fast Fourier transform (FFT) matrix built from this canonical grid. In this case, the actual linear relationship is $\boldsymbol{\theta}_o = (\mathbf{B}+\mathbf{E}_B)\mathbf{x}_o$ with $\mathbf{x}_o$ sparse. Bounds on the CS reconstruction error under basis mismatch are provided in [22]; see also [20], where the mismatch-induced error was reduced by increasing the grid density. Performance of BP/Lasso approaches for the under-determined, fully-perturbed (in both $\mathbf{y}$ and $\mathbf{A}$) linear model was analyzed in [53] by bounding the reconstruction error, and comparing it against its counterpart derived for the partially-perturbed (only in $\mathbf{y}$) model derived in [15]. Collectively, [22] and [53] address the performance question i), but provide no algorithms to address the open research issue ii).

The overarching theme of the present chapter is to address this issue by developing a sparse total least-squares (S-TLS) framework. Without exploiting sparsity, TLS has well-documented impact in applications as broad as linear prediction, system identification with errors-in-variables (EIV), spectral analysis, image reconstruction, speech, and audio processing, to name a few; see [90] and references therein. For over-determined models with unknown vectors $\mathbf{x}_o$ not abiding with (as4.0), TLS estimates are given by

$$\left\{\hat{\mathbf{x}}_{TLS}, \hat{\mathbf{E}}_{TLS}, \hat{\mathbf{e}}_{TLS}\right\} := \arg\min_{\mathbf{x},\mathbf{E},\mathbf{e}} \ \|[\mathbf{E}\ \mathbf{e}]\|_F^2 \tag{4.2a}$$

$$\text{s. to} \ \ \mathbf{y}+\mathbf{e} = (\mathbf{A}+\mathbf{E})\mathbf{x}. \tag{4.2b}$$

To cope with ill-conditioned matrices $\mathbf{A}$, an extra constraint bounding $\|\boldsymbol{\Gamma}\mathbf{x}\|_2$ is typically added in (4.2) to obtain different regularized TLS estimates depending on the choice of matrix $\boldsymbol{\Gamma}$ [7, 77].

The distinct objective of S-TLS relative to (regularized) TLS is twofold: account for sparsity as per (as4.0), and develop S-TLS solvers especially for under-determined, fully-perturbed linear models. To accomplish these goals, one must solve the S-TLS problem formulated (for $\lambda > 0$) as [cf. (4.1), (4.2)]

$$\left\{\hat{\mathbf{x}}_{S-TLS}, \hat{\mathbf{E}}_{S-TLS}, \hat{\mathbf{e}}_{S-TLS}\right\} := \arg\min_{\mathbf{x},\mathbf{e},\mathbf{E}} \ \|[\mathbf{E}\ \mathbf{e}]\|_F^2 + \lambda\|\mathbf{x}\|_1 \tag{4.3a}$$

$$\text{s. to} \ \ \mathbf{y} + \mathbf{e} = (\mathbf{A} + \mathbf{E})\mathbf{x}\,. \tag{4.3b}$$

The main goal is to develop efficient algorithms attaining at least the local and hopefully the global optimum of (4.3) – a challenging task since presence of the product $\mathbf{Ex}$ reveals that the problem is generally nonconvex. Similar to LS, BP, Lasso, and TLS, it is also worth stressing that the S-TLS estimates sought in (4.3) are universal in the sense that perturbations in $\mathbf{y}$ and $\mathbf{A}$ can be random or deterministic with or without a priori known structure.

But if prior knowledge is available on the perturbations, can weighted and structured S-TLS problems be formulated and solved? Can the scope of S-TLS be generalized (e.g., to recover a sparse matrix $\mathbf{X}_o$ using $\mathbf{A}$ and a data matrix $\mathbf{Y}$), and thus have impact in classical applications such as calibration of antenna arrays, or contemporary ones, such as cognitive radio sensing? Can S-TLS estimates be (e.g., Bayes) optimal if additional modeling assumptions are invoked? These questions will be addressed in the ensuing sections, starting from the last one.

4.2 MAP Optimality of S-TLS for EIV Models

Consider the EIV model with perturbed input ($\mathbf{A}$) and perturbed output ($\mathbf{y}$) obeying the relationship

$$\mathbf{y} = \mathbf{A}_o\mathbf{x}_o + (-\mathbf{e}_y)\,, \quad \mathbf{A} = \mathbf{A}_o + (-\mathbf{E}_A) \tag{4.4}$$

where the notation of the model perturbations $\mathbf{e}_y$ and $\mathbf{E}_A$ stresses their difference with $\mathbf{e}$ and $\mathbf{E}$, which are variables selected to yield the optimal S-TLS fit in (4.3). In a system identification setting, $\mathbf{e}_y$ and $\mathbf{E}_A$ are random perturbations giving rise to noisy output/input data $\mathbf{y}/\mathbf{A}$, based on which the task is to estimate the system vector $\mathbf{x}_o$

(comprising e.g., impulse response or pole-zero parameters), and possibly the inaccessible input matrix $\mathbf{A}_o$. To assess statistical optimality of the resultant estimators, collect the model perturbations in a column-vector form as vec($[\mathbf{E}_A \ \mathbf{e}_y]$), and further assume that:

(as4.1) *Perturbations of the EIV model in* (4.4) *are independent identically distributed (i.i.d.), Gaussian r.v.s, i.e., vec*$([\mathbf{E}_A \ \mathbf{e}_y]) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$*, independent from* $\mathbf{A}_o$ *and* $\mathbf{x}_o$*. Entries of* $\mathbf{x}_o$ *are zero-mean, i.i.d., according to a common Laplace distribution. In addition, either* **(a)** *the entries of* $\mathbf{x}_o$ *have common Laplacian parameter* $2/\lambda$*, and are independent from* $\mathbf{A}_o$*, which has i.i.d. entries drawn from a zero-mean uniform (i.e., non-informative) prior pdf; or,* **(b)** *the common Laplacian parameter of* $\mathbf{x}_o$ *entries is* $2(\sigma^2+1)/(\lambda\sigma^2)$*, and* $\mathbf{A}_o$ *conditioned on* $\mathbf{x}_o$ *has i.i.d. rows with pdf* $\mathcal{N}(\mathbf{0}, \sigma^2[\mathbf{I} - (1 + \|\mathbf{x}_o\|_2^2)^{-1}\mathbf{x}_o\mathbf{x}_o^T])$.

Note that the heavy-tailed Laplacian prior on $\mathbf{x}_o$ under (as4.1) is in par with the "non-probabilistic" sparsity attribute in (as4.0). It has been used to establish that the Lasso estimator in (4.1) is optimal, in the maximum *a posteriori* (MAP) sense, when $\mathbf{E}_A \equiv \mathbf{0}$ [83]. If on the other hand, $\mathbf{x}_o$ is viewed as non-sparse, deterministic and $\mathbf{A}_o$ as deterministic or as adhering to (as4.1b), it is known that the TLS estimator in (4.2) is optimum in the maximum likelihood (ML) sense for the EIV model in (4.4); see [67] and [69].

Aiming to establish optimality of S-TLS under (as4.1), it is useful to re-cast (4.3) as described in the following lemma, as proven in Appendix C.1. (This lemma will be used also in developing S-TLS solvers in Section 4.3.)

Lemma 4.1 *The constrained S-TLS formulation in* (4.3) *is equivalent to two unconstrained (also nonconvex) optimization problems: (a) one involving* $\mathbf{x}$ *and* $\mathbf{E}$ *variables, namely*

$$\boxed{\left\{\hat{\mathbf{x}}_{S-TLS}, \hat{\mathbf{E}}_{S-TLS}\right\} = \arg\min_{\mathbf{x},\mathbf{E}} \ \left[\|\mathbf{y} - (\mathbf{A}+\mathbf{E})\mathbf{x}\|_2^2 + \|\mathbf{E}\|_F^2 + \lambda\|\mathbf{x}\|_1\right]} \tag{4.5}$$

and (b) one of fractional form involving only the variable $\mathbf{x}$*, expressed as*

$$\boxed{\hat{\mathbf{x}}_{S-TLS} := \arg\min_{\mathbf{x}} \frac{\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2}{1 + \|\mathbf{x}\|_2^2} + \lambda\|\mathbf{x}\|_1.} \tag{4.6}$$

Using Lemma 4.1, it is possible to establish MAP optimality of the S-TLS estimator as follows, with the proof given in Appendix C.2.

Proposition 4.2 *(MAP optimality). Under (as4.1), the S-TLS estimator in* (4.3) *is MAP optimal for the EIV model in* (4.4)*. Specifically,* (4.5) *is MAP optimal for estimating both* $\mathbf{x}_o$ *and* $\mathbf{A}_o$ *under (as4.1a), while* (4.6) *is MAP optimal for estimating only* $\mathbf{x}_o$ *under (as4.1b).*

Proposition 4.2 will be generalized in Section 4.4 to account for structured and correlated perturbations with known covariance matrix. But before pursuing these generalizations, S-TLS solvers of the problem in (4.3) are in order.

4.3 S-TLS Solvers

Two iterative algorithms are developed in this section to solve the S-TLS problem in (4.3), which was equivalently re-formulated as in (4.5) and (4.6). The first algorithm can approach the *global optimum* but is computationally demanding; while the second one guarantees convergence to a *local optimum* but is computationally efficient. Thus, in addition to being attractive on its own, the second algorithm can serve as initialization to speed up convergence (and thus reduce computational burden) of the first one. To appreciate the challenge and the associated performance-complexity tradeoffs in developing algorithms for optimizing S-TLS criteria, it is useful to recall that all S-TLS problems are nonconvex; hence, unlike ordinary TLS that can be globally optimized (e.g., via SVD [67]), no efficient convex optimization solver is available with guaranteed convergence to the global optimum of (4.3), (4.5), or (4.6).

4.3.1 Bisection-based ε-Optimal Algorithm

Viewing the cost in (4.6) as a Lagrangian function, allows casting this unconstrained minimization problem as a constrained one. Indeed, sufficiency of the Lagrange multiplier theory implies that [11, Sec. 3.3.4]: using the solution $\hat{\mathbf{x}}_{S-TLS}$ of (4.6) for a given multiplier $\lambda > 0$ and letting $\mu := \|\hat{\mathbf{x}}_{S-TLS}\|_1$, the pertinent constraint is

$\mathcal{X}_1(\mu) := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_1 \leq \mu\}$; and the equivalent constrained minimization problem is [cf. (4.6)]

$$\hat{\mathbf{x}}_{S-TLS} := \arg \min_{\mathbf{x} \in \mathcal{X}_1(\mu)} f(\mathbf{x}) \,, \quad f(\mathbf{x}) := \frac{\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2}{1 + \|\mathbf{x}\|_2^2} \,. \tag{4.7}$$

There is no need to solve (4.6) in order to specify μ, because a cross-validation scheme can be implemented to specify μ in the stand-alone problem (4.7), along the lines used by e.g., [73] to determine λ in (4.6). The remainder of this subsection will thus develop an iterative scheme converging to the global optimum of (4.7), bearing in mind that this equivalently solves (4.6), (4.5) and (4.3) too.

From a high-level view, the novel scheme comprises an *outer iteration* loop based on the bisection method [29], and an *inner iteration* loop that relies on a variant of the branch-and-bound (BB) method [2]. A related approach was pursued in [7] to solve the clairvoyant TLS problem (4.2) under ℓ_2-norm regularization constraints. The challenging difference with the S-TLS here is precisely the non-differentiable ℓ_1-norm constraint in $\mathcal{X}_1(\mu)$. The outer iteration "squeezes" the minimum cost $f(\mathbf{x})$ in (4.7) between successively shrinking lower and upper bounds expressible through a parameter a. Per outer iteration, these bounds are obtained via inner iterations equivalently minimizing a surrogate quadratic function $g(\mathbf{x}, a)$, which does not have fractional form, and is thus more convenient to optimize than $f(\mathbf{x})$.

Given an upper bound a on $f(\mathbf{x})$, the link between $f(\mathbf{x})$ and $g(\mathbf{x}, a)$ follows if ones notes that

$$0 \leq a^\star := \min_{\mathbf{x} \in \mathcal{X}_1(\mu)} f(\mathbf{x}) = \min_{\mathbf{x} \in \mathcal{X}_1(\mu)} \frac{\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2}{1 + \|\mathbf{x}\|_2^2} \leq a \tag{4.8}$$

is equivalent to

$$g^\star(a) := \min_{\mathbf{x} \in \mathcal{X}_1(\mu)} g(\mathbf{x}, a) = \min_{\mathbf{x} \in \mathcal{X}_1(\mu)} \left\{\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 - a(1 + \|\mathbf{x}\|_2^2)\right\} \leq 0. \tag{4.9}$$

Suppose that after outer iteration i the optimum $a^\star$ in (4.8) belongs to a known interval $\mathcal{I}_i := [l_i, u_i]$. Suppose further that the inner loop yields the global optimum in (4.9) for $a = (l_i + u_i)/2$, and consider evaluating the sign of $g^\star(a)$ at this middle point $a = (l_i + u_i)/2$ of the interval $\mathcal{I}_i$. If $g^\star((l_i + u_i)/2) > 0$, the equivalence between (4.9) and (4.8) implies that $a^\star > (l_i + u_i)/2 > l_i$; and hence, $a^\star \in \mathcal{I}_{i+1} := [(l_i + u_i)/2, u_i]$, which

yields a reduced-size interval $\mathcal{I}_{i+1}$ by shrinking $\mathcal{I}_i$ from the left. On the other hand, if $g^\star((l_i+u_i)/2) < 0$, the said equivalence will imply that $a^\star \in \mathcal{I}_{i+1} := [l_i, (l_i+u_i)/2]$, which shrinks the $\mathcal{I}_i$ interval from the right. This successive shrinkage through bisection explains how the outer iteration converges to the global optimum of (4.7).

What is left before asserting rigorously this convergence, is to develop the inner iteration which ensures that the global optimum in (4.9) can be approached for any given a specified by the outer bisection-based iteration. To appreciate the difficulty here note that the Hessian of $g(\mathbf{x}, a)$ is given by $\mathbf{H} := 2(\mathbf{A}^T\mathbf{A} - a\mathbf{I})$. Clearly, $\mathbf{H}$ is not guaranteed to be positive or negative definite since a is positive. As a result, the cost $g(\mathbf{x}, a)$ in (4.9) bypasses the fractional form of $f(\mathbf{x})$ but it is still an indefinite quadratic, and hence nonconvex. Nonetheless, the quadratic form of $g(\mathbf{x}, a)$ allows adapting the BB iteration of [2], which can yield a feasible and δ-optimum solution $\mathbf{x}_g^\star$ satisfying: a) $\mathbf{x}_g^\star \in \mathcal{X}_1(\mu)$; and b) $g^\star(a) \leq g(\mathbf{x}_g^\star, a) \leq g^\star(a) + \delta$, where δ denotes a pre-specified margin.

In the present context, the BB algorithm finds successive upper and lower bounds of the function

$$g_{\text{box}}^\star(a) := \min_{\mathbf{x}\in\mathcal{X}_1(\mu), \mathbf{x}_L \leq \mathbf{x} \leq \mathbf{x}_U} g(\mathbf{x}, a) \tag{4.10}$$

where the constraint $\mathbf{x}_L \leq \mathbf{x} \leq \mathbf{x}_U$ represents a box that shrinks as iterations progress. Upon converting the constraints of (4.10) to linear ones, upper bounds $\mathcal{U}$ on the function $g_{\text{box}}^\star(a)$ in (4.10) can be readily obtained via suboptimum solvers of the constrained optimization of the indefinite quadratic cost $g(\mathbf{x}, a)$; see e.g., [11, Chp. 2]. Lower bounds on $g_{\text{box}}^\star(a)$ can be obtained by minimizing a convex function $g_L(\mathbf{x}, a)$, which under-approximates $g(\mathbf{x}, a)$ over the interval $\mathbf{x}_L \leq \mathbf{x} \leq \mathbf{x}_U$. This convex approximant is given by

$$g_L(\mathbf{x}, a) = g(\mathbf{x}, a) + (\mathbf{x} - \mathbf{x}_L)^T\mathbf{D}(\mathbf{x} - \mathbf{x}_U) \tag{4.11}$$

where $\mathbf{D}$ is a diagonal positive semi-definite matrix chosen to ensure that $g_L(\mathbf{x}, a)$ is convex, and stays as close as possible below $g(\mathbf{x}, a)$. Such a matrix $\mathbf{D}$ can be found by minimizing the maximum distance between $g_L(\mathbf{x}, a)$ and $g(\mathbf{x}, a)$, and comes out as the solution of the following minimization problem:

$$\min_{\mathbf{D}} \ (\mathbf{x}_U - \mathbf{x}_L)^T\mathbf{D}(\mathbf{x}_U - \mathbf{x}_L) \ \text{ s. to } \ \mathbf{H} + 2\mathbf{D} \succeq \mathbf{0} \tag{4.12}$$

Algorithm 4.1-a (BB): Input $\mathbf{y}$, $\mathbf{A}$, a, and δ. Output a δ-optimal solution $\mathbf{x}_g^\star$ of (4.9)

Set $\mathbf{x}_L = -\mu\mathbf{1}$, $\mathbf{x}_U = \mu\mathbf{1}$, $\Omega := \{(\mathbf{x}_L, \mathbf{x}_U, -\infty\}$, and initialize with $\mathcal{U} = \infty$.
repeat
  Let $(\mathbf{x}_L, \mathbf{x}_U, c)$ be one triplet of Ω with the smallest c; and set $\Omega = \Omega\backslash(\mathbf{x}_L, \mathbf{x}_U, c)$.
  Solve (4.10) locally to obtain $\hat{\mathbf{x}}_g^\star$.
  if $g(\hat{\mathbf{x}}_g^\star, a) < \mathcal{U}$ **then**
    Set $\mathcal{U} = g(\hat{\mathbf{x}}_g^\star, a)$ and $\mathbf{x}_g^\star = \hat{\mathbf{x}}_g^\star$. // update the minimum
  end if
  Minimize globally the convex $g_L(\mathbf{x}, a)$ in (4.11) with the optimum $\mathbf{D}$ in (4.12), to obtain $\check{\mathbf{x}}_g^\star$ and $\mathcal{L} := g_L(\check{\mathbf{x}}_g^\star, a)$.
  if $\mathcal{U} - \mathcal{L} > \delta$ // need to split **then**
    Find $i = \arg\max_n([\mathbf{x}_U]_n - [\mathbf{x}_L]_n)$.
    Set $\mathbf{x}_{L,1}$ ($\mathbf{x}_{U,1}$) and $\mathbf{x}_{L,2}$ ($\mathbf{x}_{U,2}$) equal to $\mathbf{x}_L$ ($\mathbf{x}_U$) except for the i-th entry. // split the maximum separation
    Set $[\mathbf{x}_{L,1}]_i = [\mathbf{x}_L]_i$, $[\mathbf{x}_{U,1}]_i = ([\mathbf{x}_U]_i - [\mathbf{x}_L]_i)/2$, $[\mathbf{x}_{L,2}]_i = ([\mathbf{x}_U]_i - [\mathbf{x}_L]_i)/2$, and $[\mathbf{x}_{U,2}]_i = [\mathbf{x}_U]_i$.
    Augment the set of unsolved boxes $\Omega = \Omega \bigcup \{(\mathbf{x}_{L,1}, \mathbf{x}_{U,1}, \mathcal{L}), (\mathbf{x}_{L,2}, \mathbf{x}_{U,2}, \mathcal{L})\}$.
  end if
until $\Omega = \emptyset$

where the constraint on the Hessian ensures that $g_L(\mathbf{x}, a)$ remains convex. Since (4.12) is a semi-definite program, it can be solved efficiently using available convex optimization software; e.g., the interior point optimization routine in SeDuMi [79]. Having selected $\mathbf{D}$ as in (4.12), $\min_{\mathbf{x}\in\mathcal{X}_1(\mu), \mathbf{x}_L\le\mathbf{x}\le\mathbf{x}_U} g_L(\mathbf{x}, a)$ is a convex problem (quadratic cost under linear constraints); thus, similar to the upper bound $\mathcal{U}$, the lower bound $\mathcal{L}$ on $g_{\text{box}}^\star(a)$ can be obtained efficiently.

The detailed inner loop (BB scheme) is tabulated as Algorithm 4.1-a. It amounts to successively splitting the initial box $-\mu\mathbf{1} \le \mathbf{x} \le \mu\mathbf{1}$, which is the smallest one containing $\mathcal{X}_1(\mu)$. Per inner iteration i, variable $\mathcal{U}$ keeps track of the upper bound on $g_{\text{box}}^\star(a)$, which at the end outputs to the outer loop the nearest estimate of $g^\star(a)$. Concurrently, the lower bound $\mathcal{L}$ on $g_{\text{box}}^\star(a)$ determines whether the current box needs to be further split, or discarded, if the difference $\mathcal{U} - \mathcal{L}$ is smaller than the pre-selected margin δ. This iterative splitting leads to a decreasing $\mathcal{U}$ and a tighter $\mathcal{L}$, both of which prevent further splitting.

Recapitulating, the outer bisection-based iteration tabulated as Algorithm 4.1-b calls Algorithm 4.1-a to find a feasible δ-optimal solution $\mathbf{x}_g^\star$ to evaluate the sign of $g^\star(a)$ in

Algorithm 4.1-b (Bisection): Input $\mathbf{y}$, $\mathbf{A}$, and tolerances ε and δ. Output an ε-optimal solution $\mathbf{x}_\varepsilon^\star$ to (4.7)

Set $l_0 = 0$, $u_0 = \|\mathbf{y}\|_2^2$, iteration index $i = 0$, and initialize the achievable cost $f_m = u_0$ with $\mathbf{x}_\varepsilon^\star = \mathbf{0}_{n\times 1}$.
repeat
  Let $a = (l_i + u_i)/2$ and call Algorithm 4.1-a to find a feasible δ-optimal solution $\mathbf{x}_g^\star$ to (4.9).
  Calculate $f_g = f(\mathbf{x}_g^\star)$, and update the iteration $i = i + 1$.
  if $f_g < f_m$ **then**
    Set $f_m = f_g$ and $\mathbf{x}_\varepsilon^\star = \mathbf{x}_g^\star$. // update the minimum
  end if
  if $g(\mathbf{x}_g^\star, a) \leq 0$ **then**
    Update $u_i = a$ and $l_i = l_{i-1}$.
  else if $g(\mathbf{x}_g^\star, a) \geq \delta$ **then**
    Update $l_i = a$ and $u_i = u_{i-1}$.
  else
    Update $l_i = a - \delta$ and $u_i = u_{i-1}$.
  end if
  Set $u_i = \min(u_i, f_g)$.
until $u_i - l_i \leq \varepsilon$

(4.9). Since $\mathbf{x}_g^\star$ is not the exact global minimum of (4.9), positivity of $g(\mathbf{x}_g^\star, a)$ does not necessarily imply $g^\star(a) > 0$. But $\mathbf{x}_g^\star$ is δ-optimal, meaning that $g^\star(a) \geq g(\mathbf{x}_g^\star, a) - \delta$; thus, $g(\mathbf{x}_g^\star, a) > \delta$, in which case the lower bound l_{i+1} is updated to $(l_i+u_i)/2$; otherwise, if $g(\mathbf{x}_g^\star, a) \in (0, \delta)$, then l_{i+1} should be set to $(l_i + u_i)/2 - \delta$.

As far as convergence is concerned, the following result can be proven as given in Appendix C.3.

Proposition 4.3 *(ε-optimal convergence) After at most* $\left\lceil \ln\left(\frac{u_0}{\varepsilon - 2\delta}\right) / \ln(2) \right\rceil$ *iterations, Algorithm 4.1-b outputs an ε-optimal solution $\mathbf{x}_\varepsilon^\star$ to* (4.7)*; that is,*

$$\mathbf{x}_\varepsilon^\star \in \mathcal{X}_1(\mu), \quad \textit{and} \quad a^\star \leq f(\mathbf{x}_\varepsilon^\star) \leq a^\star + \varepsilon. \tag{4.13}$$

Proposition 4.3 quantifies the number of outer iterations needed by the bisection-based Algorithm 4.1-b to approach within ε the global optimum of (4.7). In addition, the inner (BB) iterations bounding $g_{\text{box}}^\star(a)$ are expected to be fast converging because the box function in (4.10) is tailored for the box constraints induced by the ℓ_1-norm regularization. Nonetheless, similar to all BB algorithms, the complexity of Algorithm

4.1-a does not have guaranteed polynomial complexity on average. The latter necessitates as few calls of Algorithm 4.1-a, which means as few outer iterations. Proposition 4.3 reveals that critical to this end is the initial upper bound u_0 (Algorithm 4.1-b simply initializes with $u_0 = f(\mathbf{0})$).

This motivates the efficient suboptimal S-TLS solver of the next subsection, which is of paramount importance not only on its own, but also for initializing the ε-optimal algorithm.

4.3.2 Alternating Descent Sub-Optimal Algorithm

The starting point for a computationally efficient S-TLS solver is the formulation in (4.5). Given $\mathbf{E}$, the cost in (4.5) has the form of the Lasso problem in (4.1); while given $\mathbf{x}$, it reduces to a quadratic form, which admits closed-form solution wrt $\mathbf{E}$. These observations suggest an iterative block coordinate descent algorithm yielding successive estimates of $\mathbf{x}$ with $\mathbf{E}$ fixed, and alternately of $\mathbf{E}$ with $\mathbf{x}$ fixed. Specifically, with the iterate $\mathbf{E}(i)$ given per iteration $i \geq 0$, the iterate $\mathbf{x}(i)$ is obtained by solving the Lasso-like convex problem as [cf. (4.1)]

$$\mathbf{x}(i) = \arg\min_{\mathbf{x}} \|\mathbf{y} - [\mathbf{A} + \mathbf{E}(i)]\mathbf{x}\|_2^2 + \lambda\|\mathbf{x}\|_1 \ . \tag{4.14}$$

With $\mathbf{x}(i)$ available, $\mathbf{E}(i+1)$ for the ensuing iteration is found as

$$\mathbf{E}(i+1) = \arg\min_{\mathbf{E}} \|\mathbf{y} - \mathbf{A}\mathbf{x}(i) - \mathbf{E}\mathbf{x}(i)\|_2^2 + \|\mathbf{E}\|_F^2. \tag{4.15}$$

By setting the first-order derivative of the cost wrt $\mathbf{E}$ equal to zero, the optimal solution to the quadratic problem (4.15) is obtained in closed form as

$$\mathbf{E}(i+1) = (1 + \|\mathbf{x}(i)\|_2^2)^{-1}[\mathbf{y} - \mathbf{A}\mathbf{x}(i)]\mathbf{x}^T(i). \tag{4.16}$$

The iterations are initialized at $i = 0$ by setting $\mathbf{E}(0) = \mathbf{0}_{m\times n}$. Substituting the latter into (4.14), yields $\mathbf{x}(0) = \hat{\mathbf{x}}_{Lasso}$ in (4.1). That this is a good initial estimate is corroborated by the result in [53], which shows that even with perturbations present in both $\mathbf{A}$ and $\mathbf{y}$, the CS (and thus Lasso) estimators yield accurate reconstruction. In view of the fact that the block coordinate descent iterations ensure that the cost in (4.5) is non-increasing, the final estimates upon convergence will be at least as accurate.

The block coordinate descent algorithm is provably convergent to a stationary point of the S-TLS cost in (4.5), and thus to its equivalent forms in (4.3), (4.6) and (4.7), as asserted in the following proposition and proven in Appendix C.4.

Proposition 4.4 *(Convergence of alternating descent) Given arbitrary initialization, the iterates* $\{\mathbf{E}(i), \mathbf{x}(i)\}$ *given by* (4.14) *and* (4.16) *converge monotonically at least to a stationary point of the S-TLS problem* (4.3)*.*

Proposition 4.4 solidifies the merits of the alternating descent S-TLS solver. Simulated tests will further demonstrate that the local optimum guaranteed by this computationally efficient scheme is very close to the global optimum attained by the more complex scheme of the previous subsection.

Since estimating $\mathbf{E}$ is simple using the closed form in (4.15), it is useful at this point to explore modifications, extensions and tailored solvers for the problem in (4.14) by adapting to the present setup existing results from the Lasso literature dealing with problem (4.1). From the plethora of available options to solve (4.14), it is worth mentioning two computationally efficient ones: the least-angle regression (LARS), and the coordinate descent (CD); see e.g., [51]. LARS provides the entire "solution path" of (4.14) for all $\lambda > 0$ at complexity comparable to LS. On the other hand, if a single "best" value of λ is fixed using the cross-validation scheme [73], then CD is the state-of-the-art choice for solving (4.14).

CD in the present context cycles between iterates $\mathbf{E}(i)$, and scalar iterates of the $\mathbf{x}(i)$ entries. Suppose that the ν-th entry $x_\nu(i)$ is to be found. Precursor entries $\{x_1(i), \ldots, x_{\nu-1}(i)\}$ have been already obtained in the i-th iteration, and postcursor entries $\{x_{\nu+1}(i-1), \ldots, x_n(i-1)\}$ are also available from the previous $(i-1)$-st iteration along with $\mathbf{E}(i)$ obtained in closed form as in (4.16). If $\boldsymbol{\alpha}_\nu(i)$ denotes the ν-th column of $[\mathbf{A}+\mathbf{E}(i)]$, the effect of these known entries can be removed from $\mathbf{y}$ by forming

$$\mathbf{e}_\nu(i) := \mathbf{y} - \sum_{j=1}^{\nu-1} \boldsymbol{\alpha}_j(i) x_j(i) - \sum_{j=\nu+1}^{n} \boldsymbol{\alpha}_j(i) x_j(i-1) \,. \tag{4.17}$$

Using (4.17), the vector optimization problem in (4.14) reduces to the following scalar one with $x_\nu(i)$ as unknown: $x_\nu(i) = \arg\min_{x_\nu} [\|\mathbf{e}_\nu(i) - \boldsymbol{\alpha}_\nu(i) x_\nu\|_2^2 + \lambda |x_\nu|]$. This *scalar* Lasso problem is known to admit a closed-form solution expressed in terms of a soft

Algorithm 4.2 (CD): Input $\mathbf{y}$, $\mathbf{A}$, and coefficient λ. Output the iterates $\mathbf{E}(i)$ and $\mathbf{x}(i)$ upon convergence.

Initialize with $\mathbf{E}(0) = \mathbf{0}_{m\times n}$ and $\mathbf{x}(-1) = \mathbf{0}_{n\times 1}$
for $i = 0, 1, \ldots$ **do**
  for $\nu = 1, \ldots, n$ **do**
    Compute the residual $\mathbf{e}_\nu(i)$ as in (4.17).
    Update the scalar $x_\nu(i)$ via (4.18).
  end for
  Update the iterate $\mathbf{E}(i+1)$ as in (4.16).
end for

thresholding operator (see e.g., [51])

$$x_\nu(i) = \text{sign}\left(\mathbf{e}_\nu^T(i)\boldsymbol{\alpha}_\nu(i)\right)\left[\frac{|\mathbf{e}_\nu^T(i)\boldsymbol{\alpha}_\nu(i)|}{\|\boldsymbol{\alpha}_\nu(i)\|_2^2} - \frac{\lambda}{2\|\boldsymbol{\alpha}_\nu(i)\|_2^2}\right]_+, \quad \nu = 1, \ldots, n \tag{4.18}$$

where sign$(\cdot)$ denotes the sign operator, and $[\chi]_+ := \chi$, if $\chi > 0$, and zero otherwise.

Cycling through the closed forms (4.16)-(4.18) explains why CD here is faster than, and thus preferable over general-purpose convex optimization solvers of (4.14). Another factor contributing to its speed is the sparsity of $\mathbf{x}(i)$, which implies that starting up with the all-zero vector, namely $\mathbf{x}(-1) = \mathbf{0}_{n\times 1}$, offers initialization close to a stationary point of the cost in (4.5). Convergence to this stationary point is guaranteed by using the results in [86], along the lines of Proposition 4.4. Note also that larger values of λ in (4.18) force more entries of $\mathbf{x}(i)$ to be shrunk to zero, which corroborates the role of λ as a sparsity-tuning parameter. The CD based S-TLS solver is tabulated as Algorithm 4.2.

Remark 4.1 *(Regularization options for S-TLS).* Lasso estimators are known to be biased, but modifications are available to remedy bias effects. One such modification is the weighted Lasso, which replaces the ℓ_1-norm in (4.3) by its weighted version, namely $\sum_{\nu=1}^n w_\nu|x_\nu|$, where the weights $\{w_\nu\}$ are chosen using the LS solution [119]. An alternative popular choice is to replace the ℓ_1-norm with concave regularization terms [32], such as $\sum_{\nu=1}^n \log(x_\nu+\delta_1)$, where δ_1 is a small positive constant introduced to avoid numerical instability. In addition to mitigating bias effects, concave regularization terms provide tighter approximations to the ℓ_0-(pseudo)norm, and although they render the cost in (4.3) nonconvex, they are known to converge very fast to an improved

estimate of $\mathbf{x}$, when initialized with the Lasso solution [32].

Remark 4.2 *(Group Lasso and Matrix S-TLS).* When groups $\{\mathbf{x}_g\}_{g=1}^G$ of $\mathbf{x}$ entries are a priori known to be zero or nonzero (as4. a group), the ℓ_1-norm in (4.3) must be replaced by the sum of ℓ_2-norms, namely $\sum_{g=1}^G \|\mathbf{x}_g\|_2$. The resulting group S-TLS estimate can be obtained using the group-Lasso solver [51]. In the present context, this is further useful if one considers the matrix counterpart of the S-TLS problem in (4.3), which in its unconstrained form can be written as [cf. (4.5)]

$$\left\{\hat{\mathbf{X}}_{S-TLS}, \hat{\mathbf{E}}_{S-TLS}\right\} = \arg\min_{\mathbf{X},\mathbf{E}} \left[\|\mathbf{Y} - (\mathbf{A}+\mathbf{E})\mathbf{X}\|_F^2 + \|\mathbf{E}\|_F^2 + \lambda \sum_{\nu=1}^{n} \|\mathbf{x}_\nu^T\|_2\right] \quad (4.19)$$

where $\mathbf{x}_\nu^T$ denotes the ν-th row of the $n \times L$ unknown matrix $\mathbf{X}$, which is sparse in the sense that a number of its rows are zero, and has to be estimated using an $m \times L$ data matrix $\mathbf{Y}$ along with the regression matrix $\mathbf{A}$, both with perturbations present. Problem (4.19) can be solved using block coordinate descent cycling between iterates $\mathbf{E}(i)$ and rows $\mathbf{x}_\nu^T(i)$ as opposed to scalar entries as in (4.18).

4.4 Weighted and Structured S-TLS

Apart from the optimality links established in Proposition 4.2 under (as4.1), the S-TLS criteria in (4.3), (4.5), and (4.6) make no assumption on the perturbations $[\mathbf{E}\ \mathbf{e}]$. In this sense, the S-TLS solvers of the previous section find universal applicability. However, one expects that exploiting prior information on $[\mathbf{E}\ \mathbf{e}]$, can only lead to improved performance. Thinking for instance along the lines of weighted LS, one is motivated to *weight* $\|\mathbf{E}\|_F^2$ and $\|\mathbf{e}\|_2^2$ in (4.5) by the inverse covariance matrix of $\mathbf{E}$ and $\mathbf{e}$, respectively, whenever those are known and are not both equal to $\mathbf{I}$. As a second motivating example, normal equations, involved in e.g., linear prediction, entail *structure* in $\mathbf{E}$ and $\mathbf{e}$ that capture sample estimation errors present in the matrix $[\mathbf{A}\ \mathbf{y}]$, which is Toeplitz. Prompted by these examples, this section is about broadening the scope of S-TLS with weighted and structured forms capitalizing on prior information available about the matrix $[\mathbf{E}\ \mathbf{e}]$. To this end, it is prudent to quantify first the notion of structure.

Definition 4.1 *The $m \times (n+1)$ data matrix $[\mathbf{A}\ \mathbf{y}](\mathbf{p})$ has structure characterized by an $n_p \times 1$ parameter vector $\mathbf{p}$, if and only if there is a mapping such that $\mathbf{p} \in \mathbb{R}^{n_p} \to [\mathbf{A}\ \mathbf{y}](\mathbf{p}) := \mathbf{S}(\mathbf{p}) \in \mathbb{R}^{m \times (n+1)}$.*

Definition 4.1 is general enough to encompass any (even unstructured) matrix $[\mathbf{A}\mathbf{y}](\mathbf{p})$, by simply letting $\mathbf{p} := \text{vec}([\mathbf{A}\ \mathbf{y}]) \in \mathbb{R}^{m(n+1)}$ comprise all entries of $[\mathbf{A}\ \mathbf{y}]$. However, it becomes more relevant when $n_p \ll m(n+1)$, the case in which $\mathbf{p}$ characterizes $[\mathbf{A}\ \mathbf{y}]$ parsimoniously. Application examples are abundant: structure in Toeplitz and Hankel matrices encountered with system identification, deconvolution, and linear prediction; as well as in circulant and Vandermonde matrices showing up in spatio-temporal harmonic retrieval problems [67]. Structured matrices $\mathbf{A}$ and sparse vectors $\mathbf{x}_o$ emerge also in contemporary CS gridding-based applications e.g., for spectral analysis and estimation of time-varying channels, where rows of the FFT matrix are selected at random. (This last setting appears when training orthogonal frequency-division multiplexing (OFDM) input symbols are used to estimate communication links exhibiting variations due to mobility-induced Doppler effects [9].)

Consider now re-casting the S-TLS criteria in terms of $\mathbf{p}$, and its associated perturbation vector denoted by $\boldsymbol{\epsilon} \in \mathbb{R}^{n_p}$. The Frobenius norm in the cost of (4.3a) is mapped to the ℓ_2-norm of $\boldsymbol{\epsilon}$; and to allow for weighting the structured perturbation vector using a symmetric positive definite matrix $\mathbf{W} \in \mathbb{R}^{n_p \times n_p}$, the weighted counterpart of $\|[\mathbf{E}\,\mathbf{e}]\|_F^2$ becomes $\boldsymbol{\epsilon}^T\mathbf{W}\boldsymbol{\epsilon}$. With regards to the constraint, recall first from Definition 4.1 that $\mathbf{S}(\mathbf{p}) = [\mathbf{A}\ \mathbf{y}]$, which implies $\mathbf{S}(\mathbf{p}+\boldsymbol{\epsilon}) = [\mathbf{A}+\mathbf{E}\ \mathbf{y}+\mathbf{e}]$; hence, re-writing (4.3b) as $[\mathbf{A}+\mathbf{E}\ \mathbf{y}+\mathbf{e}]\left[\mathbf{x}^T, -1\right]^T = \mathbf{0}$, yields the structured constraint as $\mathbf{S}(\mathbf{p}+\boldsymbol{\epsilon})\left[\mathbf{x}^T, -1\right]^T = \mathbf{0}$. Putting things together, leads to the combined weighted-structured S-TLS version of (4.3) as

$$\min_{\mathbf{x},\boldsymbol{\epsilon}} \quad \boldsymbol{\epsilon}^T\mathbf{W}\boldsymbol{\epsilon} + \lambda\|\mathbf{x}\|_1 \tag{4.20a}$$

$$\text{s. to} \quad \mathbf{S}(\mathbf{p}+\boldsymbol{\epsilon}) \begin{bmatrix} \mathbf{x} \\ -1 \end{bmatrix} = \mathbf{0} \tag{4.20b}$$

which clearly subsumes the structure-only form as a special case corresponding to $\mathbf{W} = \mathbf{I}$.

To confine the structure quantified in Definition 4.1, two conditions will be imposed, which are commonly adopted by TLS approaches [67], and are satisfied by most applications mentioned so far.

(as4.2) *The structure mapping in Definition 4.1 is separable, meaning that with* $\mathbf{p} = [(\mathbf{p}^A)^T \ (\mathbf{p}^y)^T]^T$*, where* $\mathbf{p}^A \in \mathbb{R}^{n_A}$ *and* $\mathbf{p}^y \in \mathbb{R}^{n_y}$*, it holds that* $\mathbf{S}(\mathbf{p}) := [\mathbf{A} \ \mathbf{y}](\mathbf{p}) = [\mathbf{A}(\mathbf{p}^A) \ \mathbf{y}(\mathbf{p}^y)]$*. In addition, the separable structure mapping is linear (more precisely affine), if and only if the* $\mathbf{S}(\mathbf{p})$ *matrix is composed of known structural elements, namely "matrix atoms"* $\mathbf{S}_0$*,* $\{\mathbf{S}_k^A\}_{k=1}^{n_A}$ *and "vector atoms"* $\{\mathbf{s}_k^y\}_{k=1}^{n_y}$*, so that*

$$\mathbf{S}(\mathbf{p}) = \mathbf{S}_0 + \left[\sum_{k=1}^{n_A} p_k^A \mathbf{S}_k^A \qquad \sum_{k=1}^{n_y} p_k^y \mathbf{s}_k^y\right] \tag{4.21}$$

where p_k^A *(*p_k^y*) denotes the* k*-th entry of* $\mathbf{p}^A$*(*$\mathbf{p}^y$*).*

Similar to Definition 4.1, (4.21) is general enough to encompass even unstructured matrices $\mathbf{S}(\mathbf{p}) := [\mathbf{A} \ \mathbf{y}]$, by setting $\mathbf{S}_0 = \mathbf{0}$, $\mathbf{p} := \text{vec}([\mathbf{A} \ \mathbf{y}]) \in \mathbb{R}^{m(n+1)}$, $n_p := n_A + n_y = mn + m$, and selecting the m vector atoms (mn matrix atoms) as the canonical vectors (matrices), each with one entry equal to 1 and all others equal to 0. Again, interesting structures are those with $n_A \ll mn$ and/or $n_y \ll m$. (Consider for instance a circulant $m \times n$ matrix $\mathbf{A}$, which can be represented as in (4.21) using $n_A = m$ matrix atoms.)

Separability and linearity will turn out to simplify the constraint in (4.20b) for some given matrix atoms and vector atoms collected for notational brevity in the matrices

$$\boldsymbol{S}^A := [\mathbf{S}_1^A \cdots \mathbf{S}_{n_A}^A] \quad \text{and} \quad \boldsymbol{S}^y := [\mathbf{s}_1^y \ \cdots \ \mathbf{s}_{n_y}^y] \in \mathbb{R}^{m \times n_y} \ . \tag{4.22}$$

Indeed, linearity in (as4.2) allows one to write $\mathbf{S}(\mathbf{p}+\boldsymbol{\epsilon}) = \mathbf{S}(\mathbf{p}) + \mathbf{S}(\boldsymbol{\epsilon})$, and the constraint (4.20b) as: $\mathbf{S}(\boldsymbol{\epsilon})[\mathbf{x}^T, -1]^T = -\mathbf{S}(\mathbf{p})[\mathbf{x}^T, -1]^T = \mathbf{y} - \mathbf{A}\mathbf{x}$; while separability implies that $\mathbf{S}(\boldsymbol{\epsilon})[\mathbf{x}^T, -1]^T = [\sum_{k=1}^{n_A} \epsilon_k^A \mathbf{S}_k^A \quad \sum_{k=1}^{n_y} \epsilon_k^y \mathbf{s}_k^y][\mathbf{x}^T, -1]^T = \boldsymbol{S}^A(\mathbf{I} \otimes \mathbf{x})\boldsymbol{\epsilon}^A - \boldsymbol{S}^y\boldsymbol{\epsilon}^y$, where the definitions $\boldsymbol{\epsilon} := [(\boldsymbol{\epsilon}^A)^T \ (\boldsymbol{\epsilon}^y)^T]^T$ and (4.22) were used in the last equality along with the identity $\sum_{k=1}^{n_A} \epsilon_k^A \mathbf{S}_k^A \mathbf{x} = \boldsymbol{S}^A(\mathbf{I} \otimes \mathbf{x})\boldsymbol{\epsilon}^A$. In a nutshell, (4.20b) under (as4.2) becomes $\boldsymbol{S}^A(\mathbf{I} \otimes \mathbf{x})\boldsymbol{\epsilon}^A - \boldsymbol{S}^y\boldsymbol{\epsilon}^y = \mathbf{y} - \mathbf{A}\mathbf{x}$, in which $\boldsymbol{\epsilon}^A$ is decoupled from $\boldsymbol{\epsilon}^y$.

Therefore, the weighted and structured (WS)S-TLS problem in (4.20) reduces to [cf. (4.3)]

$$\min_{\mathbf{x},\boldsymbol{\epsilon}^A,\boldsymbol{\epsilon}^y} \quad \begin{bmatrix} \boldsymbol{\epsilon}^A \\ \boldsymbol{\epsilon}^y \end{bmatrix}^T \mathbf{W} \begin{bmatrix} \boldsymbol{\epsilon}^A \\ \boldsymbol{\epsilon}^y \end{bmatrix} + \lambda\|\mathbf{x}\|_1 \tag{4.23a}$$

$$\text{s. to} \quad [\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x}) \; -\boldsymbol{S}^y] \begin{bmatrix} \boldsymbol{\epsilon}^A \\ \boldsymbol{\epsilon}^y \end{bmatrix} = \mathbf{y} - \mathbf{A}\mathbf{x} \tag{4.23b}$$

or in a more compact form as: $\min_{\mathbf{x},\boldsymbol{\epsilon}}\{\boldsymbol{\epsilon}^T\mathbf{W}\boldsymbol{\epsilon}+\lambda\|\mathbf{x}\|_1\}$ s.to $\boldsymbol{G}(\mathbf{x})\boldsymbol{\epsilon} = \mathbf{r}(\mathbf{x})$, after defining

$$\boldsymbol{G}(\mathbf{x}) := [\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x}) \;\; \boldsymbol{S}^y] \quad \text{and} \quad \mathbf{r}(\mathbf{x}) := \mathbf{y} - \mathbf{A}\mathbf{x}\,. \tag{4.24}$$

Comparing (4.3) with (4.23) allows one to draw apparent analogies: both involve three sets of optimization variables, and both are nonconvex because two of these sets enter the corresponding constraints in a bilinear fashion [cf. product of $\mathbf{E}$ with $\mathbf{x}$ in (4.3b), and $\boldsymbol{\epsilon}^A$ with $\mathbf{x}$ in (4.23b)].

Building on these analogies, the following lemma shows how to formulate WSS-TLS criteria, paralleling those of Lemma 4.1, where one or two sets of variables were eliminated to obtain efficient, provably convergent solvers, and establish statistical optimality links within the EIV model in (4.4).

Lemma 4.5 *The constrained WSS-TLS form in* (4.20) *is equivalent to two unconstrained nonconvex optimization problems: (a) one involving* $\mathbf{x}$ *and* $\boldsymbol{\epsilon}^A$ *variables, namely*

$$\{\hat{\mathbf{x}},\hat{\boldsymbol{\epsilon}}^A\}_{WSS-TLS} = \arg\min_{\mathbf{x},\boldsymbol{\epsilon}^A} \begin{bmatrix} \boldsymbol{\epsilon}^A \\ (\boldsymbol{S}^y)^\dagger[\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x})\boldsymbol{\epsilon}^A - \mathbf{r}(\mathbf{x})] \end{bmatrix}^T \mathbf{W}\times \begin{bmatrix} \boldsymbol{\epsilon}^A \\ (\boldsymbol{S}^y)^\dagger[\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x})\boldsymbol{\epsilon}^A - \mathbf{r}(\mathbf{x})] \end{bmatrix} + \lambda\|\mathbf{x}\|_1 \tag{4.25}$$

where $\boldsymbol{S}^y$ *is assumed full rank and square*[1]*, i.e.,* $m = n_y$ *in* (4.24)*; and also (b) one involving only the variable* $\mathbf{x}$*, expressed using the definitions in* (4.24)*, as*

$$\hat{\mathbf{x}}_{WSS-TLS} = \arg\min_{\mathbf{x}} \; \{\mathbf{r}^T(\mathbf{x}) \; \left[\boldsymbol{G}(\mathbf{x})\,\mathbf{W}^{-1}\,\boldsymbol{G}^T(\mathbf{x})\right]^\dagger \; \mathbf{r}(\mathbf{x}) \; + \; \lambda\|\mathbf{x}\|_1\}\,. \tag{4.26}$$

[1] Tall $\boldsymbol{S}^y$ matrices with full column rank can be handled too for block diagonal weight matrices $\mathbf{W}$ typically adopted with separable structures; see also [114]. This explains why the pseudo-inverse of $\boldsymbol{S}^y$ is used in this section instead of its inverse; but exposition of the proof simplifies considerably for the square case. Note also that the full rank assumption is not practically restrictive because data matrices perturbed by noise of absolutely continuous pdf have full rank almost surely.

The formulation in (4.25) suggests directly an iterative WSS-TLS solver based on the block coordinate descent method. Specifically, suppose that the estimate $\boldsymbol{\epsilon}^A(i)$ of $\boldsymbol{\epsilon}^A$ is available at iteration i. Substituting $\boldsymbol{\epsilon}^A(i)$ into (4.25), allows estimating $\mathbf{x}$ as

$$\mathbf{x}(i)=\arg\min_{\mathbf{x}}\begin{bmatrix}\boldsymbol{\epsilon}^A(i)\\ (\boldsymbol{S}^y)^\dagger[\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x})\boldsymbol{\epsilon}^A(i)-\mathbf{r}(\mathbf{x})\end{bmatrix}^T\mathbf{W}\begin{bmatrix}\boldsymbol{\epsilon}^A(i)\\ (\boldsymbol{S}^y)^\dagger[\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x})\boldsymbol{\epsilon}^A(i)-\mathbf{r}(\mathbf{x})]\end{bmatrix}+\lambda\|\mathbf{x}\|_1. \tag{4.27}$$

Since $\mathbf{r}(\mathbf{x})$ is linear in $\mathbf{x}$ [cf. (4.24)], the cost in (4.27) is convex (quadratic regularized by the ℓ_1-norm as in the Lasso cost in (4.1)); thus, it can be minimized efficiently. Likewise, given $\mathbf{x}(i)$ the perturbation vector for the ensuing iteration can be found in closed form since the pertinent cost is quadratic; that is,

$$\begin{aligned}\boldsymbol{\epsilon}^A(i+1)=\arg\min_{\boldsymbol{\epsilon}^A}&\begin{bmatrix}\boldsymbol{\epsilon}^A\\ (\boldsymbol{S}^y)^\dagger[\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x}(i))\boldsymbol{\epsilon}^A-\mathbf{r}(\mathbf{x}(i))]\end{bmatrix}^T\mathbf{W}\times\\ &\begin{bmatrix}\boldsymbol{\epsilon}^A\\ (\boldsymbol{S}^y)^\dagger[\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x}(i))\boldsymbol{\epsilon}^A-\mathbf{r}(\mathbf{x}(i))]\end{bmatrix}.\end{aligned} \tag{4.28}$$

To express $\boldsymbol{\epsilon}^A(i+1)$ compactly, partition $\mathbf{W}$ in accordance with $\mathbf{p}=[(\mathbf{p}^A)^T,\ (\mathbf{p}^y)^T]^T$; i.e., let

$$\mathbf{W}=\begin{bmatrix}\mathbf{W}_{AA} & \mathbf{W}_{Ay}\\ \mathbf{W}_{Ay}^T & \mathbf{W}_{yy}\end{bmatrix}. \tag{4.29}$$

Using (4.29), and equating to zero the gradient (wrt $\boldsymbol{\epsilon}^A$) of the cost in (4.28), yields the closed form

$$\boldsymbol{\epsilon}^A(i+1)=\left\{\check{\boldsymbol{S}}(\mathbf{x}(i))\mathbf{W}\check{\boldsymbol{S}}^T(\mathbf{x}(i))\right\}^\dagger\check{\boldsymbol{S}}(\mathbf{x}(i))[\mathbf{W}_{Ay}^T,\ \mathbf{W}_{yy}^T]^T(\boldsymbol{S}^y)^\dagger\mathbf{r}(\mathbf{x}(i)) \tag{4.30}$$

where $\check{\boldsymbol{S}}(\mathbf{x}(i)):=\left[\mathbf{I},\ \left[(\boldsymbol{S}^y)^\dagger\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x}(i))\right]^T\right]$.

Initialized with $\boldsymbol{\epsilon}^A(0)=\mathbf{0}_{n_A\times 1}$, the algorithm cycles between iterations (4.27) and (4.30). Mimicking the steps of Appendix C.4, it is easy to show that these iterations are convergent as asserted in the following.

Proposition 4.6 *(Convergence). The iterates in* (4.27) *and* (4.30) *converge monotonically at least to a stationary point of the cost in* (4.20)*, provided that* $\boldsymbol{S}^y$ *in* (4.24) *has full column rank.*

As with the solver of Section 4.3.2, CD is also applicable to the WSS-TLS solver, by cycling between $\boldsymbol{\epsilon}^A(i)$ and scalar iterates of the $\mathbf{x}(i)$ entries. To update the ν-th entry $x_\nu(i)$, suppose precursor entries $\{x_1(i),\ldots,x_{\nu-1}(i)\}$ have been already obtained in the i-th iteration, and postcursor entries $\{x_{\nu+1}(i-1),\ldots,x_n(i-1)\}$ are also available from the previous $(i-1)$-st iteration along with $\boldsymbol{\epsilon}^A(i)$, found in closed form as in (4.30). Letting $\boldsymbol{\alpha}_\nu(i)$ denote the ν-th column of $\left[(\boldsymbol{S}^y)^\dagger\left(\mathbf{A}+\sum_{k=1}\epsilon_k^A\mathbf{S}_k^A\right)\right]$, the effect of these known entries can be removed from $\mathbf{y}$ by forming [cf. (4.17)]

$$\mathbf{e}_\nu(i) := (\boldsymbol{S}^y)^\dagger\mathbf{y} - \sum_{j=1}^{\nu-1}\boldsymbol{\alpha}_j(i)x_j(i) - \sum_{j=\nu+1}^{n}\boldsymbol{\alpha}_j(i)x_j(i-1)\,. \tag{4.31}$$

Using (4.31), the vector optimization problem in (4.27) now reduces to the following scalar one with $x_\nu(i)$ as unknown: $x_\nu(i)=\arg\min_{x_\nu}\{\|\boldsymbol{\alpha}_\nu(i)x_\nu-\mathbf{e}_\nu(i)\|^2_{W_{yy}}+2[\boldsymbol{\epsilon}^A(i)]^T\mathbf{W}_{Ay}\boldsymbol{\alpha}_\nu(i)x_\nu+\lambda|x_\nu|\}$, where $\|\cdot\|_{W_{yy}}$ denotes the ℓ_2-norm weighted by $\mathbf{W}_{yy}$. The solution of this scalar Lasso problem can be expressed using the same soft-thresholding form as in (4.18), and is given by

$$x_\nu(i)=\text{sign}\left([\mathbf{e}_\nu(i)\mathbf{W}_{yy}-\boldsymbol{\epsilon}^A(i)\mathbf{W}_{Ay}^T]^T\boldsymbol{\alpha}_\nu(i)\right)\left[\frac{\left|[\mathbf{e}_\nu(i)\mathbf{W}_{yy}-\boldsymbol{\epsilon}^A(i)\mathbf{W}_{Ay}^T]^T\boldsymbol{\alpha}_\nu(i)\right|}{\|\boldsymbol{\alpha}_\nu(i)\|^2_{W_{yy}}}-\frac{\lambda}{2\|\boldsymbol{\alpha}_\nu(i)\|^2_{W_{yy}}}\right]_+\,. \tag{4.32}$$

This block CD algorithm enjoys fast convergence (at least) to a stationary point, thanks both to the simplicity of (4.32), and the sparsity of $\mathbf{x}(i)$, as explained in Section IV-B.

The WSS-TLS criterion in (4.25) is also useful to establish its statistical optimality under a structured EIV model, with output-input data obeying the relationships

$$\mathbf{y}=\mathbf{A}(\mathbf{p}_o^A)\mathbf{x}_o+(-\boldsymbol{S}^y\boldsymbol{\epsilon}_y)\,,\quad \mathbf{A}=\mathbf{A}(\mathbf{p}_o^A)+\left(-\sum_{k=1}^{n_A}\epsilon_{A,k}\mathbf{S}_k^A\right) \tag{4.33}$$

where perturbation vectors $\boldsymbol{\epsilon}_A$ and $\boldsymbol{\epsilon}_y$ play the role of $\mathbf{E}_A$ and $\mathbf{e}_y$ in (4.4), and differ from the optimization variables $\boldsymbol{\epsilon}^A$ and $\boldsymbol{\epsilon}^y$ in (4.23). Unknown are the vector $\mathbf{x}_o$, and the inaccessible input matrix $\mathbf{A}(\mathbf{p}_o^A)$, characterized by the vector $\mathbf{p}_o^A$. The model in (4.33) obeys the following structured counterpart of (as4.1a).

(as4.1′) *Perturbations in* (4.33) *are jointly Gaussian, i.e.,* $[\boldsymbol{\epsilon}_A\ \boldsymbol{\epsilon}_y]\sim\mathcal{N}(\mathbf{0},\mathbf{W}^{-1})$*, as well as independent from* $\mathbf{p}_o^A$ *and* $\mathbf{x}_o$*. Vector* $\mathbf{x}_o$ *has i.i.d. entries with the same prior*

as in (as4.1a); and it is independent from $\mathbf{p}_o^A$*, which has i.i.d. entries drawn from a zero-mean uniform (i.e., non-informative) prior pdf.*

The following optimality claim holds for the WSS-TLS estimator in (4.25), assured to be equivalent to the solution of problem (4.23) by Lemma 4.5.

Proposition 4.7 *(MAP optimality of WSS-TLS). Under (as4.1′) and (as4.2), the equivalent WSS-TLS problem in* (4.25) *yields the MAP optimal estimator of* $\mathbf{x}_o$ *and* $\mathbf{p}_A$ *in the structured EIV model* (4.33)*.*

4.5 S-TLS Applications

In this section, the practical impact of accounting for perturbations present in the data matrix $[\mathbf{A}\ \mathbf{y}]$ will be demonstrated via two sensing applications involving reconstruction of sparse vectors. In both, the perturbation $\mathbf{E}_A$ comes from inaccurate modeling of the underlying actual matrix $\mathbf{A}_o$, while $\mathbf{e}_y$ is due to measurement noise.

4.5.1 Cognitive Radio Sensing

Consider N_s sources located at unknown positions, each transmitting a radio frequency (RF) signal with power spectral density (PSD) $\Phi_s(f)$ that is well approximated by a basis expansion model: $\Phi_s(f) = \sum_{\nu=1}^{N_b} x_{s\nu} b_\nu(f)$, where $\{b_\nu(f)\}_{\nu=1}^{N_b}$ are known (e.g., rectangular) basis functions, and $\{x_{s\nu}\}_{s=1}^{N_s}$ are unknown power coefficients. As source positions are also unknown, a Cartesian grid of known points $\{\boldsymbol{l}_g\}_{g=1}^{N_g}$ is adopted to describe *candidate* locations that transmitting radios could be positioned [6, 19]; see also Fig. 4.1.

The task is to estimate the locations and powers of active sources based on PSD samples measured at N_r cognitive radios (CRs) at known locations $\{\boldsymbol{\ell}_r\}_{r=1}^{N_r}$. Per frequency f_k, these samples obey the model

$$\begin{aligned}\hat{\Phi}_r(f_k) &= \sum_{g=1}^{N_g} \gamma_{gr}\Phi_g(f_k) + \sigma_r^2 + e_r(f_k) \\ &= \left[\sum_{g=1}^{N_g}\sum_{\nu=1}^{N_b} \gamma_{gr} b_\nu(f_k)\right] x_{g\nu} + \sigma_r^2 + e_r(f_k) \ = \mathbf{a}_r^T(f_k)\mathbf{x}_o + e_r(f_k) \end{aligned} \tag{4.34}$$

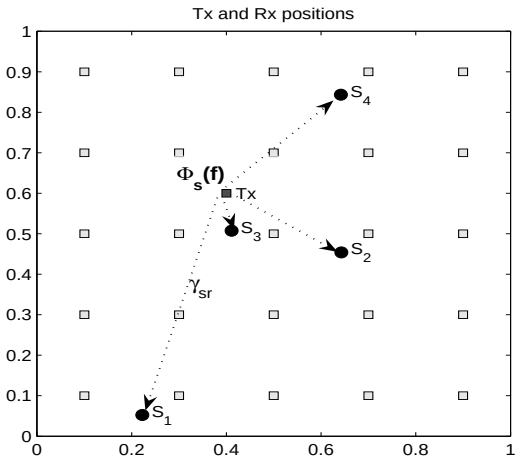

Figure 4.1: Grid topology with $N_g = 25$ candidate locations, $N_s = 1$ transmitting source, and $N_r = 4$ receiving CRs.

where the PSD $\Phi_g(f)$ is nonzero only if a transmitting source is present at $\boldsymbol{l}_g$; γ_{gr} represents the channel gain from the candidate source at $\boldsymbol{l}_g$ to the CR at $\boldsymbol{\ell}_r$ that is assumed to follow a known pathloss function of the distance $||\boldsymbol{l}_g - \boldsymbol{\ell}_r||$; σ_r^2 denotes the known noise variance at receiver r; the $N_b N_g \times 1$ vector $\mathbf{a}_r(f_k)$ collects products $\gamma_{gr} b_\nu(f_k)$; vector $\mathbf{x}_o$ contains the $N_b N_g$ unknown power coefficients $x_{g\nu}$; and $e_r(f_k)$ captures the error between the true, $\Phi_r(f_k)$, and estimated, $\hat{\Phi}_r(f_k)$, PSDs.

Estimated PSD samples at K frequencies from all N_r receivers are first compensated by subtracting the corresponding noise variances, and subsequently collected to form the data vector $\mathbf{y}$ of length $m = KN_r$. Noise terms $-e_r(f_k)$ are similarly collected to build the perturbation vector $\boldsymbol{\epsilon}_y$. Likewise, row vectors $\mathbf{a}_r^T(f_k)$ of length $n = N_b N_g$ are concatenated to form the $m \times n$ matrix $\mathbf{A}$. The latter is perturbed (relative to the inaccessible $\mathbf{A}_o$) by a matrix $\mathbf{E}_A$, which accounts for the mismatch between grid location vectors, $\{\boldsymbol{l}_g\}_{g=1}^{N_g}$, and those of the actual sources, $\{\boldsymbol{\kappa}_s\}_{s=1}^{N_s}$. To specify $\mathbf{E}_A$, let $\epsilon_{gr}^s := \gamma_{sr} - \gamma_{gr}$ for the source at $\boldsymbol{\kappa}_s$ closest to $\boldsymbol{l}_g$, and substitute $\gamma_{gr} = \gamma_{sr} - \epsilon_{gr}^s$ into the double sum inside the square brackets of (4.34). This allows writing $\mathbf{A} = \mathbf{A}_o - \mathbf{E}_A$, where $\mathbf{E}_A$ is affine structured with coefficients $\{\epsilon_{gr}^s\}$ and matrix atoms formed by $\{b_\nu(f_k)\}$. All in all, the setup fits nicely the structured EIV model in (4.33).

Together with $\{\epsilon_{gr}^s\}$, the support of $\mathbf{x}_o$ estimates the locations of sources, and the

nonzero entries of $\mathbf{x}_o$ their transmit powers. Remarkably, this grid-based approach reduces localization – traditionally a nonlinear estimation task – to a linear one, by increasing the problem dimensionality ($N_g \gg N_s$). What is more, $\mathbf{x}_o$ is sparse for two reasons: a) relative to the swath of available bandwidth, the transmitted PSDs are narrowband; hence, the number of nonzero $x_{g\nu}$s is small relative to N_b; and (b) the number of actual sources (N_s) is much smaller than the number of grid points (N_g) that is chosen large enough to localize sources with sufficiently high resolution. Existing sparsity-exploiting approaches to CR sensing rely on BP/Lasso, and do not take into account the mismatch arising due to griding [6, 19]. Simulations in Section 4.6 will demonstrate that sensing accuracy improves considerably if one accounts for grid-induced errors through the EIV model, and compensates for them via the novel WSS-TLS estimators.

4.5.2 DoA Estimation via Sparse Linear Regression

The setup here is the classical one in sensor array processing: plane waves from N_s far-field, narrowband sources impinge on a uniformly-spaced linear array (ULA) of N_r (possibly uncalibrated) antenna elements. Based on as few $N_r \times 1$ vectors of spatial samples collected across the ULA per time instant t (snapshot), the task is to localize sources by estimating their directions-of-arrival (DoA) denoted as $\{\vartheta_s\}_{s=1}^{N_s}$. High-resolution, (weighted) subspace-based DoA estimators are nonlinear, and rely on the sample covariance matrix of these spatio-temporal samples, which requires a relatively large number of snapshots for reliable estimation especially when the array is not calibrated; see e.g., [55]. This has prompted recent DoA estimators based on sparse linear regression, which rely on a uniform polar grid of N_g points describing *candidate* DoAs $\{\theta_g\}_{g=1}^{N_g}$ [36, 65, 70]. Similar to the CR sensing problem, the g-th entry x_g of the $N_g \times 1$ unknown vector of regression coefficients, $\mathbf{x}_{o,t}$, is nonzero and equal to the transmit-source signal power, if a source is impinging at angle θ_g, and zero otherwise.

The $N_r \times 1$ array response vector to a candidate source at DoA θ_g is

$$\mathbf{a}(\theta_g) = [1 \; e^{-j\alpha_g} \; \cdots \; e^{-j\alpha_g(N_r-1)}]^T$$

where $\alpha_g := 2\pi d \sin(\theta_g)$ denotes the phase shift relative to the source signal wavelength between neighboring ULA elements separated by distance d. The per-snapshot received

data vector $\mathbf{y}_t$ of length $m = N_r$ obeys the EIV model: $\mathbf{y}_t = (\mathbf{A} + \mathbf{E}_A)\ \mathbf{x}_{o,t} + (-\mathbf{e}_{y,t})$, where $-\mathbf{e}_{y,t}$ represents the additive noise across the array elements; the $m \times n$ matrix $\mathbf{A} := [\mathbf{a}(\theta_1)\ \cdots\ \mathbf{a}(\theta_{N_g})]$ denotes the grid angle scanning matrix of $n = N_g$ columns; and $\mathbf{E}_A$ represents perturbations arising because DoAs from actual sources do not necessarily lie on the postulated grid points. Matrix $\mathbf{E}_A$ can also account for gain, phase, and position errors of antenna elements when the array is uncalibrated.

To demonstrate how a structured S-TLS approach applies to the DoA estimation problem at hand, consider for simplicity one source from direction ϑ_s, whose nearest grid angle is θ_g; and let $\epsilon_g^s := \vartheta_s - \theta_g$ be the corresponding error that vanishes as the grid density grows large. For small ϵ_g^s, the actual source-array phase shift $\alpha_s := 2\pi d \sin(\theta_g + \epsilon_g^s)$ can be safely approximated as $2\pi d[\sin(\theta_g)\cos(\epsilon_g^s) + \cos(\theta_g)\sin(\epsilon_g^s)] \approx 2\pi d[\sin(\theta_g) + \epsilon_g^s \cos(\theta_g)]$; or, more compactly as $\alpha_s \approx \alpha_g + \epsilon_g \beta_g$, where $\beta_g := 2\pi d \cos(\theta_g)$. As a result, using the approximation $\exp(-j\alpha_s) \approx \exp[-j(\alpha_g + \epsilon_g^s \beta_g)] \approx (1 - j\epsilon_g^s \beta_g)\exp(-j\alpha_g)$, the actual array response vector can be approximated as a linear function of ϵ_g^s; thus, it be expressed as

$$\begin{aligned} &\mathbf{a}(\vartheta_s) := \mathbf{a}(\theta_g) + \epsilon_g^s \boldsymbol{\phi}(\theta_g), \\ &\boldsymbol{\phi}(\theta_g) := [0\ \ -j\beta_g e^{-j\alpha_g}, \ldots, -j\beta_g (N_r - 1) e^{-j\alpha_g (N_r - 1)}]^T . \end{aligned} \tag{4.35}$$

With columns obeying (4.35), the actual array manifold is modeled as $\mathbf{A}_o = \mathbf{A} + \mathbf{E}_A$, where the perturbation matrix is structured as $\mathbf{E}_A = \sum_{g=1}^{N_g} \epsilon_g^s \mathbf{S}_g^A$, with the $N_r \times N_g$ matrix $\mathbf{S}_g^A$ having all zero entries, except for the g-th column that equals $\boldsymbol{\phi}(\theta_g)$. With such an array manifold and $\boldsymbol{S}^y = \mathbf{I}$, the grid-based DoA setup matches precisely the structured EIV model in (4.33). The simulated tests in the ensuing section will illustrate, among other things, the merits of employing WSS-TLS solvers to estimate $\mathbf{x}_{o,t}$ and ϵ_g^s based on data collected by possibly antenna arrays. But before this, a final remark is in order.

Remark 4.3 *(Relationships with [35] and [65]).* Although $\mathbf{E}_A$ is not explicitly included in the model of existing grid-based approaches, this mismatch has been mitigated either by iteratively refining the grid around the region where sources are present [65], or, by invoking the minimum description length (MDL) test to estimate the number of actual sources N_s, followed by spatial interpolation to estimate their DoAs [35]. These remedies require post-processing the initial estimates obtained by sparse linear regression. In

contrast, the proposed *structured* S-TLS based approach *jointly* estimates the nonzero support of $\mathbf{x}_{o,t}$ along with grid-induced perturbations. This allows for direct compensation of the angle errors to obtain high-resolution DoA estimates in a single step, and in certain cases without requiring multiple snapshots. Of course, multiple snapshots are expected to improve estimation performance using the matrix S-TLS solver mentioned in Remark 4.2.

4.6 Simulated Tests

Four simulated tests are presented in this section to illustrate the merits of the S-TLS approach[2] , starting from the algorithms of Section 4.3.

Test Case 4.1 (Optimum vs. suboptimum S-TLS): The EIV model in (4.4) is simulated here with a 6×10 matrix $\mathbf{A}$, whose entries are i.i.d. Gaussian having variance $1/6$, so that the expected ℓ_2-norm of each column equals 1. The entries of $\mathbf{E}_A$ and $\mathbf{e}_y$ are also i.i.d. Gaussian with variance $0.0025/6$ corresponding to entry-wise signal-to-noise ratio (SNR) of 26dB. Vector $\mathbf{x}_o$ has only nonzero elements in the two first entries: $x_{o,1} = -1.3$ and $x_{o,2} = 5$. Algorithm 1 is tested with $\mu = 5$ against Algorithm 4.2 implemented with different values of λ to obtain a solution satisfying $\|\hat{\mathbf{x}}_{S-TLS}\|_1 = \mu$. For variable tolerance values ε in Algorithm 4.1-b, the attained minimum cost $f(\mathbf{x})$ in (4.8) is plotted in Fig. 4.2. To serve as a benchmark, a genie-aided globally optimum scheme is also tested with the support of $\mathbf{x}$ known and equal to that of $\mathbf{x}_o$. Specifically, the genie-aided scheme minimizes $f(\mathbf{x})$ over all points with ℓ_1-norm equal to μ, and all entries being 0 except for the first two. Using the equivalence between (4.8) and (4.9), the genie-aided scheme per iteration amounts to minimizing a scalar quadratic program under linear constrains, which is solved efficiently using the interior-point optimization routine in [79].

Fig. 4.2 shows that as ε becomes smaller, the minimum achieved value $f(\mathbf{x})$ decreases monotonically, and drops sharply to the global minimum attained by the genie-aided bisection scheme. Interestingly, the alternating descent algorithm that guarantees convergence to a stationary point, exhibits performance comparable to the global algorithm.

[2] Matlab® code for the algorithms in this chapter is available at `http://www.tc.umn.edu/~zhuh`

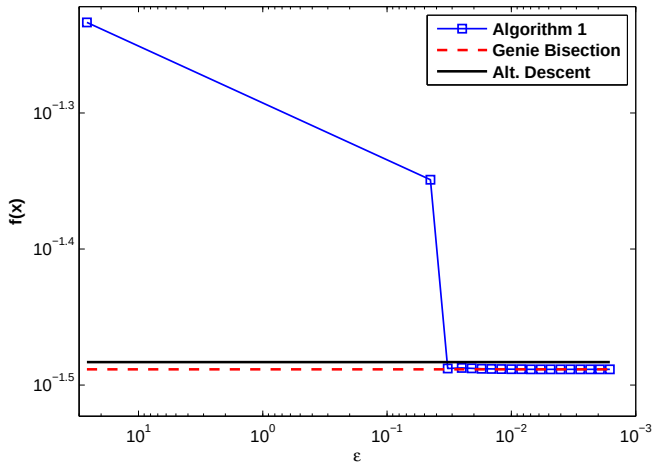

Figure 4.2: Attained $f(\mathbf{x})$ for variable tolerance values ε by the global Algorithm 4.1-b, compared to the alternating descent local algorithm, and the genie-aided global solver.

For this reason, only the alternating descent algorithm is used in all subsequent tests. Next, S-TLS estimates are compared with those obtained via BP/Lasso and (regularized) TLS in the context of the CR sensing and array processing applications outlined in Section 4.5.

Test Case 4.2 (S-TLS vs. Lasso vs. TLS): The setup here is also based on the EIV model (4.4), with $\mathbf{A}$ of size 20×40 having i.i.d. Gaussian entries; and $\mathbf{x}_o$ having 5 nonzero i.i.d. standardized Gaussian entries. All other parameters are as in Test Case 1 adapted to the different problem size here. By averaging results over 200 Monte Carlo runs, the S-TLS solution is compared against the Lasso one for 20 values of λ (uniformly spaced in log-scale), based on the ℓ_2, ℓ_1, and ℓ_0 errors of the estimated vectors relative to $\mathbf{x}_o$. (The ℓ_0 error equals the percentage of entries for which the support of the two vectors is different.) Fig. 4.3 corroborates the improvement of S-TLS over Lasso, especially in the ℓ_0 norm. Fig. 4.3(c) further demonstrates that over a range of moderate λ values, S-TLS consistently outperforms Lasso in recovering the true support of $\mathbf{x}_o$. For high λ's, both estimates come close to the all-zero vector, so that the ℓ_0 errors become approximately the same, even though the ℓ_2 and ℓ_1 errors are smaller for Lasso. However, for both error norms S-TLS has a slight edge over moderate values of λ.

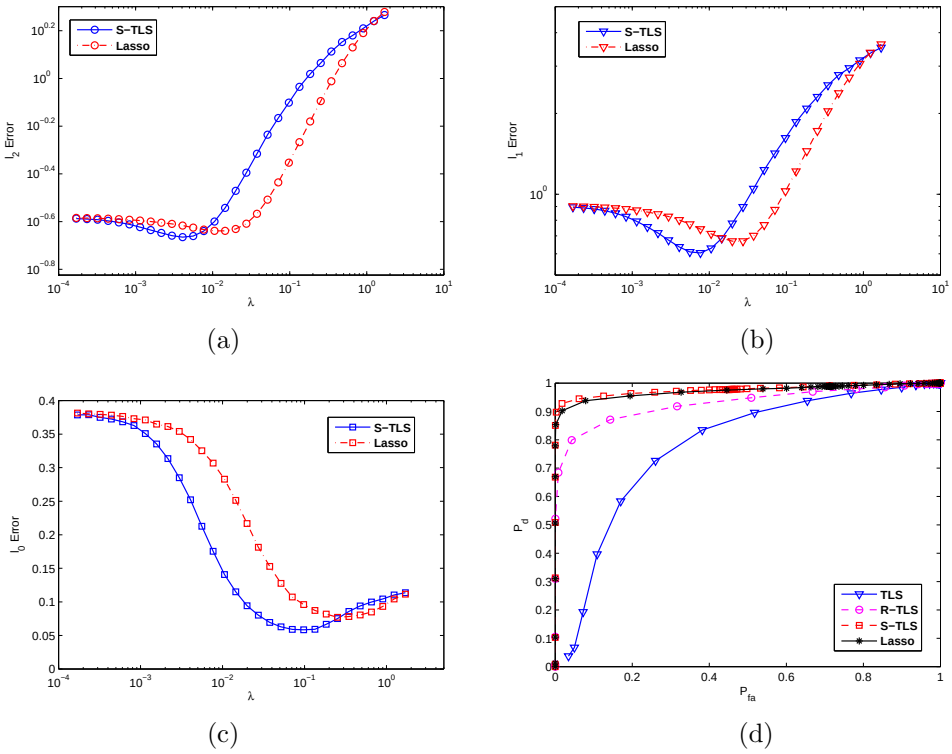

Figure 4.3: Comparison between S-TLS and Lasso in terms of: (a) ℓ_2-norm, (b) ℓ_1-norm, and (c) ℓ_0-norm of the estimation errors; (d) Prob. of detection versus prob. of false alarms for TLS, ℓ_2 regularized (R-)TLS, S-TLS and Lasso.

Receiver operating characteristic (ROC) curves are plotted in Fig. 4.3(d) to illustrate the merits of S-TLS and Lasso over (regularized) TLS in recovering the correct support. The "best" λ for the S-TLS and Lasso algorithms is chosen using cross-validation [73]. As TLS cannot be applied to under-determined systems, a 40×40 matrix $\mathbf{A}$ is selected. Since TLS and LS under an ℓ_2-norm constraint $\|\mathbf{x}\|_2 \leq \mu$ are known to be equivalent when μ is small [77], the regularized TLS is tested using the function `lsqi` for regularized LS from [50]. The probability of correct detection, P_d, is calculated as the probability of identifying correctly the support over nonzero entries of $\mathbf{x}_o$, and the probability of false alarms, P_{fa}, as that of incorrectly deciding zero entries to be nonzero. The ROC curves in Fig. 4.3(d) demonstrate the advantage of Lasso, and more clearly that of S-TLS, in

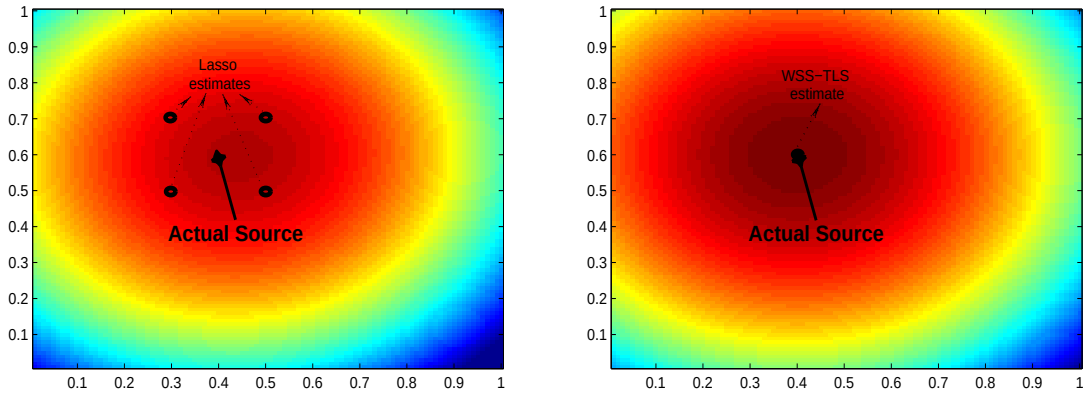

(a) (b)

Figure 4.4: Comparison between PSD maps estimated by (a) Lasso, and (b) WSS-TLS for the CR network in Fig. 4.1.

recovering the correct support.

Test Case 4.3 (CR spectrum sensing): This simulation is performed with reference to the CR network in the region $[0\ 1] \times [0\ 1]$ in Fig. 4.1. The setup includes $N_r = 4$ CRs deployed to estimate the power and location of a single source with position vector [0.4 0.6], located at the center of four neighboring grid points. The CRs scan $K = 128$ frequencies from 15MHz to 30MHz, and adopt the basis expansion model in Section 4.5-A with $N_b = 16$ rectangular $b_\nu(f)$ functions, each of bandwidth 1MHz. The actual source only transmits over the $\nu = 6$-th band. The channel gains are exponentially decaying in distances with exponent $-1/2$. The received data are generated using the transmit PSD described earlier, a regular Rayleigh fading channel with 6 taps, and additive white Gaussian receiver noise at SNR=0dB. Receive-PSDs are obtained using exponentially weighted periodograms (with weight 0.99) averaged over 1,000 coherence blocks; see also [6] for more details of a related simulation. The WSS-TLS approach is used to account for perturbations ϵ_{gr}^s in the channel gains. A diagonal matrix $\mathbf{W}$ is used with each diagonal entry equal to $\hat{\sigma}_\epsilon^{-2}$ (inversely proportional to the average of of sample variances of ϵ_{gr}^s).

With λ chosen as in [21], both Lasso and WSS-TLS identify the active frequency band correctly (only the entries $\{x_{g6}\}_{g=1}^{16}$ were estimated as nonzero). However, Lasso identifies four transmitting sources at positions [0.3(0.5) 0.5(0.7)], the four grid points closest to [0.4 0.6]. WSS-TLS returns only one source at position [0.5 0.5], along with the estimated $\hat{\epsilon}_{gr}^s$ that yields $\hat{\gamma}_{sr} = \hat{\epsilon}_{gr}^s + \gamma_{gr}$. Concatenate the latter to form $\hat{\boldsymbol{\gamma}}_s$ of length $N_r = 4 \ll m$. Using a refined grid of 25 points uniformly spaced over the "zoom-in" region [0.3 0.7] $\times$ [0.3 0.7] centered at [0.5 0.5], correlation coefficients between $\hat{\boldsymbol{\gamma}}_s$ and those of each candidate point are evaluated. The source position is estimated as the point with maximum correlation coefficient, which for WSS-TLS occurs at the true location [0.4 0.6]. To illustrate graphically the two alternatives, the estimated maps of the spatial PSDs at the 6th frequency band are plotted in Fig. 4.4(a) using the Lasso, and in Fig. 4.4(b) using WSS-TLS. The marked point indicates the actual source location [0.4 0.6] in both maps. Unlike Lasso, the WSS-TLS identifies correctly the true position of the source.

Test Case 4.4 (DoA estimation): The setup here entails a ULA consisting of $N_r = 8$ antenna elements with inter-element spacing $d = 1/2$, and a grid of $N_g = 90$ scanning angles from -90° to 90° wrt the array boresight. Two sources ($N_s = 2$) of unit amplitude impinge from angles $\vartheta_s = 1^\circ$ and -9°, both 1° off their nearest grid DoAs. As in the single-snapshot test in [65], the SNR is set to 20dB. The variance of ϵ_g^s in (4.35) is obtained from the uniform distribution in $[-1^\circ, 1^\circ]$. Selecting λ according to the noise level as in [65], Lasso returns four nonzero entries, two around each source at $\vartheta_s \pm 1^\circ$; while WSS-TLS gives two nonzero θ_g estimates at -10° ($g = 40$) and 0° ($g = 45$), along with perturbation estimates $\hat{\epsilon}_{40}^s$ and $\hat{\epsilon}_{45}^s$. Using the latter, the DoAs are estimated as $\hat{\vartheta}_s := \hat{\theta}_g + \hat{\epsilon}_g^s$ for $g = 40,\ 45$. The angle spectra using Lasso, and WSS-TLS with estimated $\hat{\vartheta}_s$, are compared in Fig. 4.5(a). The two black arrows depict the actual source angles, and benchmark the true angular spectrum.

To further illustrate the merits of WSS-TLS in estimating correctly the closest grid point and subsequently each DoA, the sample variance of a DoA estimate is plotted versus SNR in Fig. 4.5(b) using Monte Carlo runs, each with a single source randomly placed over $[-1^\circ, 1^\circ]$. Both WSS-TLS and Lasso are post-processed by interpolating peaks in the obtained spectra from two nearest grid points, linearly weighted by the

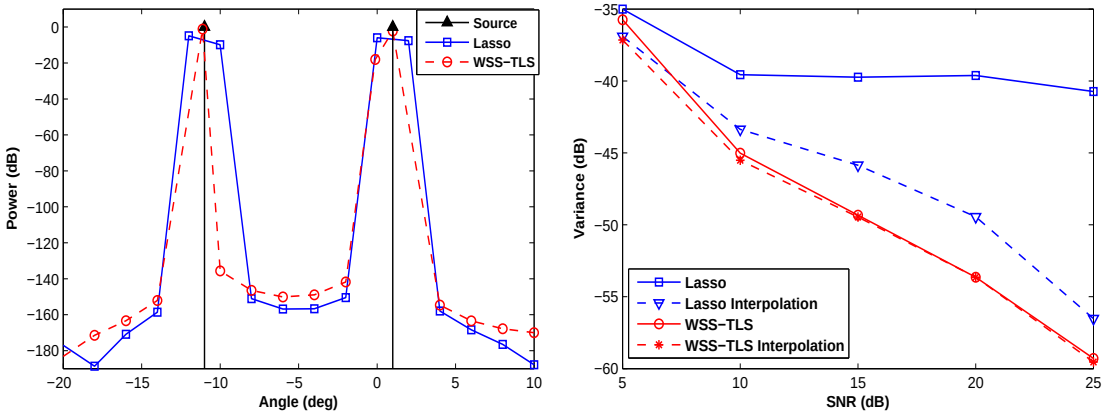

(a) (b)

Figure 4.5: (a) Angular spectra estimated using Lasso and WSS-TLS as compared to the actual transmission pattern; (b) Comparison of angle estimation variances of Lasso, WSS-TLS, without and with interpolation.

estimated amplitudes as in [36]. Both curves confirm that WSS-TLS outperforms the Lasso. More interestingly, the two WSS-TLS curves almost coincide, which further corroborates that WSS-TLS manages in a single step to identify correctly the support of $\mathbf{x}_{o,t}$ without requiring post processing.

Chapter 5

Exploiting Sparse User Activity in Multiuser Detection

The number of *active* users in code-division multiple access (CDMA) systems is often much lower than the spreading gain. This chapter exploits fruitfully this *a priori* information to improve performance of multiuser detectors. A low-activity factor manifests itself in a sparse symbol vector with entries drawn from a finite alphabet that is augmented by the zero symbol to capture user inactivity. The non-equiprobable symbols of the augmented alphabet motivate a sparsity-exploiting maximum *a posteriori* (S-MAP) probability criterion, which is shown to yield a cost comprising the ℓ_2 least-squares error penalized by the p-th norm of the wanted symbol vector ($p = 0, 1, 2$). Related optimization problems appear in variable selection (shrinkage) schemes developed for linear regression, as well as in the emerging field of compressive sampling (CS). The contribution of this work to such *sparse* CDMA systems is a gamut of sparsity-exploiting multiuser detectors trading off performance for complexity requirements. From the vantage point of CS and the least-absolute shrinkage selection operator (Lasso) spectrum of applications, the contribution amounts to sparsity-exploiting algorithms when the entries of the wanted signal vector adhere to finite-alphabet constraints.

5.1 Modeling and Problem Statement

Consider the uplink of a CDMA system with K user terminals and spreading gain N. Assume first that $N \geq K$. The under-determined case ($N < K$) will also be addressed later on in Sec. 5.5.2. Suppose the system has a relatively low activity factor, which analytically means that each terminal is active with probability (w.p.) $p_a < 1/2$ per symbol, and the events "*active*" are independent across symbols and across users. The case of correlated (in)activity of users across symbols will be dealt with in Sec. 5.5.1. Let $b_k \in \mathcal{A}$ denote the symbol, drawn from a finite alphabet by the k-th user, when active; otherwise, $b_k = 0$. Incorporating possible (in)activity per user is equivalent to having b_k take values from an *augmented* alphabet $\mathcal{A}_a := \mathcal{A} \bigcup \{0\}$.

The access point (AP) receives the superimposed modulated (quasi-) synchronous signature waveforms through (possibly frequency-selective) fading channels in the presence of additive white Gaussian noise (AWGN); and projects on the orthonormal space spanned by the aggregate waveforms to obtain the received chip samples collected in the $N \times 1$ vector $\mathbf{y}$. With the $K \times 1$ vector $\mathbf{b}$ containing the symbols of all (active and inactive) users, the canonical input-output relationship is, see e.g., [93, Sec. 2.9]

$$\mathbf{y} = \mathbf{H}\mathbf{b} + \mathbf{w} \tag{5.1}$$

where $\mathbf{H}$ is an $N \times K$ matrix capturing transmit-receive filters, spreading sequences, channel impulse responses, and timing offsets; and the $N \times 1$ vector $\mathbf{w}$ is the AWGN. Without loss of generality (w.l.o.g.), $\mathbf{w}$ can be scaled to have unit variance. Note that (5.1) holds for quasi-synchronous systems too, where relative user asynchronism is bounded to a few chips per symbol, provided that: either i) user transmissions include guard bands to eliminate inter-symbol interference (ISI); or ii) the received vector $\mathbf{y}$ collects only the chips of each user that belong to the part of the common symbol interval under consideration.

The low activity factor implies that $\mathbf{b}$ is a sparse vector. However, the AP is neither aware of the positions nor the number of zero entries in $\mathbf{b}$. In order to perform multiuser detection (MUD) needed to determine the optimal $\hat{\mathbf{b}}$, the AP must account for the augmented alphabet $\mathcal{A}_a$, i.e., consider all the possible candidate vectors $\mathbf{b} \in \mathcal{A}_a^K$. This way, the MUD also determines the k-th user's activity captured by the extra constellation point $b_k = 0$. Supposing that the AP has the channel matrix $\mathbf{H}$ available (e.g., via

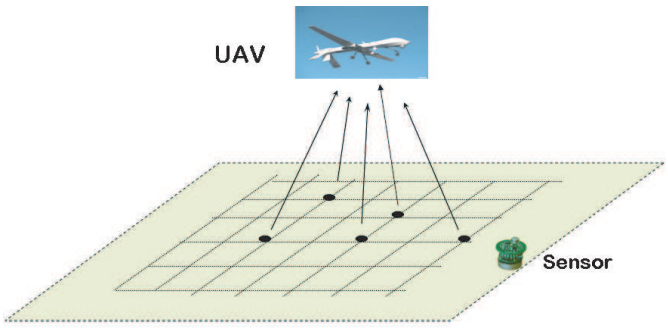

Figure 5.1: Wireless sensors access a UAV.

training), the goal of this chapter is to detect the optimal $\hat{\mathbf{b}}$ given the received vector $\mathbf{y}$ by exploiting the sparsity of active users.

To motivate this sparsity-exploiting MUD setup in the CDMA context, consider a set of terminals wishing to link with a common AP. Suppose that the AP acquires the full matrix $\mathbf{H}$ (with all terminals active) during a training phase. Those channels may include either non-dispersive or multipath fading, and are assumed invariant during the coherence time, which is typically larger compared to the symbol period. Each terminal accesses the channel randomly, and the AP receives the superposition of signals from the active terminals only. The AP is interested in determining both the active terminals and the symbols transmitted.

Another scenario where $\mathbf{H}$ is known and sparsity-exploiting MUD is well motivated, entails an unmanned aerial vehicle (UAV) collecting information from a ground wireless sensor network (WSN) placed over a grid, as depicted in Fig. 5.1. As the UAV flies over the grid of sensors, it collects the signals from a random subset of them. If the channel fading is predominantly affected by path loss, the UAV can acquire $\mathbf{H}$ based on the relative AP-to-sensor distances. Again, the UAV faces the problem of determining both identities of active sensors, and the symbols each active sensor transmits.

With this problem setup in mind, we will develop different MUD strategies, which account for the low activity factor. First, we will look at the maximum *a posteriori* probability (MAP) optimum MUD that exploits the sparsity present.

5.2 Sparsity-Exploiting MAP Detector

The goal is to detect $\mathbf{b}$ in (5.1), given a prescribed activity factor, the received vector $\mathbf{y}$, and the matrix $\mathbf{H}$. Recall though that the low activity factor leads to a sparse $\mathbf{b}$, i.e., each entry b_k is more likely to take the value 0 from the alphabet. Because entries $\{b_k\}_{k=1}^K$ are non-equiprobable, the optimal detector in the sense of minimizing the detection error rate is the MAP one.

Aiming at a sparsity-aware MAP criterion, consider first the prior probability for $\mathbf{b}$. For simplicity in exposition, suppose for now that each terminal transmits binary symbols when active, i.e., $\mathcal{A} = \{\pm 1\}$. (It will become clear later on that all sparsity-cognizant MUD schemes are applicable to finite alphabets of general constellations, not necessarily binary.) If b_k takes values from $\{-1, 0, 1\}$, with corresponding probabilities $\{p_a/2, 1-p_a, p_a/2\}$, and since each entry b_k is independent from $b_{k'}$ for $k \neq k'$, the prior probability for $\mathbf{b}$ can be expressed as

$$\Pr(\mathbf{b}) = \prod_{k=1}^{K} \Pr(b_k) = (1-p_a)^{K-\|\mathbf{b}\|_0} (p_a/2)^{\|\mathbf{b}\|_0} \tag{5.2}$$

where $\|\mathbf{b}\|_0$ denotes the ℓ_0 (pseudo) norm that is by definition equal to the number of non-zero entries in the vector $\mathbf{b}$. Upon taking logarithms, (5.2) yields

$$\ln \Pr(\mathbf{b}) = -\lambda \|\mathbf{b}\|_0 + K \ln(1-p_a) \tag{5.3}$$

where

$$\lambda := \ln \frac{1-p_a}{p_a/2}. \tag{5.4}$$

Since low activity factor means $p_a < 1/2$, it follows readily from (5.4) that $\lambda > 0$. With the prior distribution of $\mathbf{b}$ in (5.3), the sparsity-aware MAP (S-MAP) detector is

$$\begin{aligned} \hat{\mathbf{b}}^{\mathrm{MAP}} &= \arg \max_{\mathbf{b} \in \mathcal{A}_a^K} \Pr(\mathbf{b}|\mathbf{y}) = \arg \min_{\mathbf{b} \in \mathcal{A}_a^K} -\ln p(\mathbf{y}|\mathbf{b}) - \ln \Pr(\mathbf{b}) \\ &= \arg \min_{\mathbf{b} \in \mathcal{A}_a^K} \frac{1}{2} \|\mathbf{y} - \mathbf{H}\mathbf{b}\|_2^2 + \lambda \|\mathbf{b}\|_0 \end{aligned} \tag{5.5}$$

where the last equality follows from (5.3) and the Gaussianity of $\mathbf{w}$. Hence, the S-MAP detection task amounts to finding the vector in the constraint set $\mathcal{A}_a^K$, which minimizes the cost in (5.5).

Two remarks are now in order.

Remark 5.1 *(General constellations).* Beyond binary alphabets, it is easy to see why the S-MAP detector in (5.5) applies to more general constellations, including pulse amplitude modulation (PAM), phase-shift keying (PSK), and quadrature amplitude modulation (QAM). Specifically, for general M-ary constellations with $M \geq 2$ it suffices to adjust accordingly the prior probability as a function of $\|\mathbf{b}\|_0$ in (5.2). This will render the S-MAP MUD in (5.5) applicable to general M-ary constellations, provided that λ in (5.4) is replaced by $\lambda := \ln \frac{1-p_a}{p_a/M}$.

Remark 5.2 *(Scale λ as a function of p_a).* The definition in (5.4) reveals the explicit relationship of λ with the activity factor p_a. Different from CS and VS approaches, where λ is a tuning parameter often chosen with cross-validation techniques as a function of the data size N and K, here it is directly coupled with the user activity factor p_a. Such a coupling carries over even when users have distinct activity factors. Specifically, if the user k is active w.p. $p_{a,k}$, then the term $\lambda\|\mathbf{b}\|_0 := \lambda \sum_{k=1}^{K} |b_k|$ in (5.5) should change to $\sum_{k=1}^{K} \lambda_k |b_k|$, where λ_k is defined as in (5.4) with $p_{a,k}$ substituting p_a. This user-specific regularization will be used in Sec. 5.5.1 to adaptively estimate user activity factors on-the-fly, and thus enable sparsity-aware MUD which accounts for correlated user (in)activity across the time slots.

With the variable b_k only taking values from $\{\pm 1, 0\}$, it holds for $p \geq 1$ that

$$\boxed{\|\mathbf{b}\|_0 = \sum_{k=1}^{K} |b_k|^p = \|\mathbf{b}\|_p^p, \quad \forall\, \mathbf{b} \in \mathcal{A}_a^K.} \tag{5.6}$$

Hence, the S-MAP detector (5.5) for binary transmissions is equivalent to

$$\boxed{\hat{\mathbf{b}}^{\text{MAP}} = \arg \min_{\mathbf{b} \in \mathcal{A}_a^K} \frac{1}{2}\|\mathbf{y} - \mathbf{H}\mathbf{b}\|_2^2 + \lambda \|\mathbf{b}\|_p^p \ , \ \ \forall\, p \geq 1.} \tag{5.7}$$

Notice that the equivalence between (5.5) and (5.7) is based on the norm equivalence in (5.6), which holds only for constant modulus constellations. Although the cost in (5.7) will turn out to be of interest on its own, it is not an S-MAP detector for non-constant modulus constellations.

Interestingly, since low-activity factor implies a positive λ, the problem (5.7) entails a convex cost function, compared to the non-convex one in (5.5). In lieu of the

finite-alphabet constraint, the criterion of (5.7) consists of the least-squares (LS) cost regularized by the ℓ_p norm, which in the context of linear regression has been adopted to mitigate data overfitting. In contrast, the LS estimator only considers the goodness-of-fit, and thus tends to overfit the data. Shrinking the LS estimator, by penalizing its size through the ℓ_p norm, typically outperforms LS in practice. For example, the Lasso adopts the ℓ_1 norm through which it effects sparsity [83]. In the MUD context for CDMA systems with low-activity factor, the vector $\mathbf{b}$ has a sparse structure, which motivates well this regularizing strategy. What is distinct and interesting here is that this penalty-augmented LS approach under finite-alphabet constraints emerges naturally as the logarithm of the prior in the S-MAP detector.

However, the finite-alphabet constraint renders the solution of (5.7) combinatorially complex. For general $\mathbf{H}$ and $\mathbf{y}$, the solution of (5.7) requires exhaustive search over all the 3^K feasible points, with the complexity growing exponentially in the problem dimension K. Likewise, for general M-ary alphabets the complexity incurred by (5.5) is $\mathcal{O}((M+1)^K)$. On the other hand, many (sub-) optimal alternatives are available in the MUD literature; see e.g., [93, Ch. 5-7]. Similarly here, we will develop different (sub-) optimal algorithms to trade off complexity for probability of error performance in sparsity-exploiting MUD alternatives.

Since the exponential complexity of MUD stems from the finite-alphabet constraint, one reduced-complexity approach is to solve the unconstrained convex problem, and then quantize the resultant soft decision to the nearest point in the alphabet. This approach includes the sub-optimal linear MUD algorithms (decorrelating and minimum mean-square error (MMSE) detectors). Another approach is to search over (possibly a subset of) the alphabet lattice directly as in the decision-directed detectors or the sphere decoders [94]. Likewise, it is possible to devise (sub-) optimal algorithms for solving the S-MAP MUD problem (5.7) along these two categories. First, we will present the sparsity-exploiting MUD algorithms by relaxing the finite-alphabet constraint.

5.3 Relaxed S-MAP Detectors

In addition to offering an S-MAP detector for constant modulus constellations, the cost in (5.7) is convex. Thus, by relaxing the combinatorial constraint, the optimization

problem (5.7) can be solved efficiently by capitalizing on convexity. As mentioned earlier, this problem is similar to the penalty-augmented LS criterion that is used for VS in linear regression, where the choice of p is important for controlling the shrinking effect, that is the degree of sparsity in the solution. Next, we will develop detectors for two choices of p, and compare them in terms of complexity and performance.

5.3.1 Linear Ridge MUD

The choice $p = 2$ is a popular one in statistics, well-known as *Ridge regression*. Its popularity is mainly due to the fact that it can regularize the LS solution while retaining its closed-form expression as a linear function of the data $\mathbf{y}$. A relaxed detection algorithm for S-MAP MUD can be developed accordingly with $p = 2$, what we term Ridge detector (RD). Ignoring the finite-alphabet constraint, the optimal solution of (5.7) for $p = 2$ takes a linear form

$$\begin{aligned}\mathbf{b}^{\mathrm{RD}} &= \arg\min_{\mathbf{b}} \frac{1}{2}\|\mathbf{y}-\mathbf{H}\mathbf{b}\|_2^2 + \lambda\|\mathbf{b}\|_2^2 \\ &= (\mathbf{H}^T\mathbf{H} + 2\lambda\mathbf{I})^{-1}\mathbf{H}^T\mathbf{y}. \end{aligned} \tag{5.8}$$

In addition to its simplicity, and different from LS, the existence of the inverse in (5.8) is ensured even for ill-posed or under-determined problems; i.e., when $\mathbf{H}$ is rank deficient or fat ($N < K$). Notice that $\mathbf{b}^{\mathrm{RD}}$ takes a form similar to the linear MMSE multiuser detector, with the parameter λ replacing the noise variance, and connecting the activity factor with the degree of regularization applied to the matrix $\mathbf{H}^T\mathbf{H}$.

Upon quantizing each entry of the soft decision $\mathbf{b}^{\mathrm{RD}}$ with the operator

$$\mathcal{Q}_\theta(x) := \mathrm{sign}(x)\mathbb{1}(|x| \geq \theta) \tag{5.9}$$

where $\theta > 0$, $\mathrm{sign}(x) = 1(-1)$ with $x > (<)0$, and $\mathbb{1}$ denoting the indicator function, the hard RD is

$$\hat{\mathbf{b}}^{\mathrm{RD}} = \mathcal{Q}_\theta(\mathbf{b}^{\mathrm{RD}}) = \mathcal{Q}_\theta\left((\mathbf{H}^T\mathbf{H} + 2\lambda\mathbf{I})^{-1}\mathbf{H}^T\mathbf{y}\right). \tag{5.10}$$

Because the detector in (5.8) is linear, it is possible to express linearly its soft output $\mathbf{b}^{\mathrm{RD}}$ with respect to (wrt) the input symbol vector $\mathbf{b}$. Based on this relationship, one can subsequently derive the symbol error rate (SER) of the hard detected symbols in $\hat{\mathbf{b}}^{\mathrm{RD}}$ as a function of the quantization threshold θ. These steps will be followed next to obtain the performance of the RD.

Performance Analysis

Letting $\check{\mathbf{b}}$ denote the vector transmitted, and substituting $\mathbf{y} = \mathbf{H}\check{\mathbf{b}} + \mathbf{w}$ into (5.8) yields

$$\mathbf{b}^{\mathrm{RD}} = (\mathbf{H}^T\mathbf{H} + 2\lambda\mathbf{I})^{-1}\mathbf{H}^T(\mathbf{H}\check{\mathbf{b}} + \mathbf{w}) = \mathbf{G}\check{\mathbf{b}} + \mathbf{w}' \tag{5.11}$$

where $\mathbf{G} := \mathbf{I} - 2\lambda(\mathbf{H}^T\mathbf{H} + 2\lambda\mathbf{I})^{-1}$, and the colored noise $\mathbf{w}' := (\mathbf{H}^T\mathbf{H} + 2\lambda\mathbf{I})^{-1}\mathbf{H}^T\mathbf{w}$ is zero-mean Gaussian with covariance matrix $\mathbf{\Sigma}_{w'} := E\{\mathbf{w}'(\mathbf{w}')^T\} = (\mathbf{H}^T\mathbf{H}+2\lambda\mathbf{I})^{-2}\mathbf{H}^T\mathbf{H}$.

It follows readily from (5.11) that the k-th entry of $\mathbf{b}^{\mathrm{RD}}$ satisfies

$$b_k^{\mathrm{RD}} = G_{kk}\check{b}_k + \sum_{\ell\neq k} G_{k\ell}\check{b}_\ell + w'_k \tag{5.12}$$

where $G_{k\ell}$ and w'_k are the (k,ℓ)-th and k-th entries of $\mathbf{G}$ and $\mathbf{w}'$, respectively. The last two terms in the right-hand side of (5.12) capture the multiuser interference-plus-noise effect, which has variance

$$\sigma_k^2 = \mathrm{var}\left\{\sum_{\ell\neq k} G_{k\ell}\check{b}_\ell + w'_k\right\} = \sum_{\ell\neq k} G_{k\ell}^2 p_a + \Sigma_{w',kk} \tag{5.13}$$

where $\Sigma_{w',kk}$ denotes the (k,k)-th entry of $\mathbf{\Sigma}_{w'}$.

With the interference-plus-noise term being approximately Gaussian distributed, deciphering $\check{b}_k$ from (5.12) amounts to detecting a ternary deterministic signal in the presence of zero-mean, Gaussian noise of known variance. Hence, the symbol error rate (SER) for the k-th entry using the quantization rule in (5.10) entry-wise can be analytically obtained as

$$P_{e,k}^{\mathrm{RD}} = 2(1-p_a)Q\left(\frac{\theta}{\sigma_k}\right) + p_a Q\left(\frac{G_{kk}-\theta}{\sigma_k}\right) \tag{5.14}$$

where $Q(\mu) := (1/\sqrt{2\pi})\int_\mu^\infty \exp(-\nu^2/2)d\nu$ denotes the Gaussian tail function.

The SER in (5.14) is a convex function of the threshold θ. Thus, taking the first-order derivative wrt θ and setting it equal to zero yields the optimal threshold for the k-th entry as

$$\hat{\theta}_k = \frac{G_{kk}}{2} + \frac{\sigma_k^2}{G_{kk}}\lambda\,. \tag{5.15}$$

The corresponding minimum SER becomes [cf. (5.14) and (5.15)]

$$\hat{P}_{e,k}^{\mathrm{RD}} = 2(1-p_a)Q\left(\frac{G_{kk}}{2\sigma_k} + \frac{\lambda\sigma_k}{G_{kk}}\right) + p_a Q\left(\frac{G_{kk}}{2\sigma_k} - \frac{\lambda\sigma_k}{G_{kk}}\right). \tag{5.16}$$

As the CDMA system signal-to-noise ratio (SNR) goes to infinity, asymptotically we have $\mathbf{G} \to \mathbf{I}$ and $\mathbf{\Sigma}_{w'} \to \mathbf{0}$, so the optimal threshold in (5.15) approaches 0.5. The numerical tests in Sec. 5.6 will also confirm that selecting $\theta_k = 0.5$ approaches the minimum SER $\hat{P}_{e,k}^{\mathrm{RD}}$ over the range of SNR values encountered in most practical settings.

The clear advantage of RD-MUD is its simplicity as a linear detector. However, using the ℓ_2 norm for regularization, the RD-MUD inherently introduces a Gaussian prior for the unconstrained symbol vector and is thus not effecting sparsity in $\mathbf{b}^{\mathrm{RD}}$; see also [83]. This renders the performance of RD dependent on the quantization threshold θ – a fact also corroborated by the simulations in Sec. 5.6. These considerations motivate the ensuing development of an alternative relaxed S-MAP algorithm, which accounts for the sparsity present in $\mathbf{b}$.

5.3.2 Lasso-based MUD

Another popular regression method is the Lasso one, which regularizes the LS cost with the ℓ_1 norm. In the Bayesian formulation, regularization with the ℓ_1 norm corresponds to adopting a Laplacian prior for $\mathbf{b}$ [83]. The nice feature of Lasso-based regression is that it ensures sparsity in the resultant estimates. The degree of sparsity depends on the value of λ, which is selected here using the *a priori* information available on the activity factor [cf. (5.4)]. The optimal solution of (5.7) for $p = 1$ without the finite-alphabet constraint yields the Lasso detector (LD) as

$$\mathbf{b}^{\mathrm{LD}} = \arg\min_{\mathbf{b}} \frac{1}{2}\|\mathbf{y} - \mathbf{H}\mathbf{b}\|_2^2 + \lambda\|\mathbf{b}\|_1. \tag{5.17}$$

While a closed-form solution is impossible for general $\mathbf{H}$, the minimization in (5.17) is a quadratic programming (QP) problem that can be readily accomplished using available QP solvers, such as SeDuMi [79]. Upon slicing the solution in (5.17), we obtain the detection result as

$$\hat{\mathbf{b}}^{\mathrm{LD}} = \mathcal{Q}_{\theta}(\mathbf{b}^{\mathrm{LD}}). \tag{5.18}$$

The larger λ is, the more sparse $\mathbf{b}^{\mathrm{LD}}$ becomes [cf. (5.4)]. This is intuitively reasonable, because λ is inversely proportional to the activity factor p_a. Since the Lasso approach (5.17) yields sparse estimates systematically, and can be obtained via QP solvers in

polynomial time, LD is a competitive MUD alternative. Lack of a closed-form solution prevents analytical evaluation of the SER, which will be tested using simulations in Sec. 5.6.

Remark 5.3 *(Relaxed detectors for general constellations).* So far, we assumed $\mathcal{A} = \{\pm 1\}$ to ensure equivalence of the ℓ_p-norm regularized S-MAP detector (5.7) with the more general one in (5.5). However, the sub-optimal algorithms of this section ignore the finite-alphabet constraint, and just rely on the convexity of the cost function in (5.7) to offer MUD schemes that can be implemented efficiently, either in linear closed-form or through quadratic programming. In fact, starting from (5.7) and forgoing its equivalence with (5.5), the RD and LD relaxations of (5.7) apply also for general M-ary alphabets for any $M > 2$. Of course, the quantization thresholds required for slicing the soft symbol estimates in order to obtain hard symbol estimates must be modified in accordance with the corresponding constellation. For non-constant modulus transmissions, the cost in (5.7) favors low-energy (close to the origin) constellation points, but this effect is mitigated by the judicious selection of the quantization thresholds.

Note that forgoing the equivalence of (5.7) with (5.5) is less of an issue for RD and LD because the major limitation of these simple relaxation-based algorithms is their sub-optimality, which emerges because they do not account for the finite-alphabet symbol constraints. Next, MUD algorithms are developed to minimize the S-MAP cost while adhering to the constraint in (5.5) explicitly.

5.4 S-MAP Detectors with Lattice Search

User symbols in this section are drawn from an M-ary PAM alphabet $\mathcal{A} = \{\pm 1, \pm 3, \ldots, \pm(M-1)\}$, with M even. Consider also reformulating the S-MAP problem in (5.5) using the QR decomposition of the matrix $\mathbf{H}$ (assumed here to be square or tall with full column rank) as $\mathbf{H} = \mathbf{QR}$, where $\mathbf{R}$ is a $K \times K$ upper triangular matrix, and $\mathbf{Q}$ is an $N \times K$ unitary matrix. Substituting this QR decomposition into (5.5), and left multiplying with the unitary $\mathbf{Q}$ inside the LS cost, the S-MAP detector becomes: $\hat{\mathbf{b}}^{\text{MAP}} = \arg\min_{\mathbf{b}\in\mathcal{A}_a^k} \frac{1}{2}\left\|\mathbf{Q}^T\mathbf{y} - \mathbf{Q}^T(\mathbf{QR})\mathbf{b}\right\|_2^2 + \lambda\|\mathbf{b}\|_0 = \arg\min_{\mathbf{b}\in\mathcal{A}_a^K} \frac{1}{2}\|\mathbf{y}' - \mathbf{Rb}\|_2^2 +$

$\lambda\|\mathbf{b}\|_0$, where $\mathbf{y}' := \mathbf{Q}^T\mathbf{y}$; or, after using the definitions of the norms,

$$\boxed{\hat{\mathbf{b}}^{\mathrm{MAP}} = \arg\min_{\mathbf{b}\in\mathcal{A}_a^K} \sum_{k=1}^{K} \left\{ \frac{1}{2}\left(y'_k - \sum_{\ell=k}^{K} R_{k\ell} b_\ell \right)^2 + \lambda |b_k|_0 \right\} .} \tag{5.19}$$

Although the optimal solution of (5.19) still incurs exponential complexity, the upper triangular form of $\mathbf{R}$ enables decomposition of (5.19) into sub-problems involving only scalar variables. As it will be seen later in Sec. 5.4.1, the S-MAP problem accepts a neat closed-form solution in the scalar case ($K = 1$). This is instrumental for the development of efficient (near-) optimal algorithms searching over the finite-alphabet induced lattice. One such sub-optimal MUD algorithm is described next.

5.4.1 Sparsity-Exploiting Decision-Directed MUD

Close look at (5.19) reveals that once the estimates $\{\hat{b}_\ell\}_{\ell=k+1}^{K}$ are available, the optimal $\hat{b}_k$ can be obtained by minimizing the cost corresponding to the k-th summand of (5.19). This leads to the per-symbol optimal decision-directed detector (DDD), which following related schemes in different contexts, could be also called *successive interference cancellation* or *decision-feedback decoding*, see e.g., [39, Sec. 9.4] and [93, Ch. 7]. The main difference here is that $\mathbf{b}$ is sparse.

The DDD algorithm relies on back substitution to decompose the overall S-MAP cost into K sub-costs each dependent on a single scalar variable and accepting a closed-form solution. Specifically, supposing the symbols $\{\hat{b}_\ell^{\mathrm{DDD}}\}_{\ell=k+1}^{K}$ have been already detected, the DDD algorithm detects the k-th symbol as

$$\hat{b}_k^{\mathrm{DDD}} = \arg\min_{b_k\in\mathcal{A}_a} \frac{1}{2}\left(y'_k - \sum_{\ell=k+1}^{K} R_{k\ell}\hat{b}_\ell^{\mathrm{DDD}} - R_{kk} b_k \right)^2 + \lambda |b_k|_0 \, . \tag{5.20}$$

This minimization problem entails only one scalar variable taking one of $(M+1)$ possible points in $\mathcal{A}_a$. Thus, the minimum is found after comparing the costs corresponding to these $(M+1)$ values. Appendix D.1 proves that this optimal solution can be found in closed form as

$$\hat{b}_k^{\mathrm{DDD}} = \lfloor b_k^{\mathrm{LS}} \rceil \mathbb{1}\left(2b_k^{\mathrm{LS}} \lfloor b_k^{\mathrm{LS}} \rceil - \lfloor b_k^{\mathrm{LS}} \rceil^2 - 2\lambda/R_{kk}^2 > 0 \right) \tag{5.21}$$

Algorithm 5.1 (DDD): Input $\mathbf{y}'$, $\mathbf{R}$, λ, and M even. Output $\hat{\mathbf{b}}^{\mathrm{DDD}}$.

for $k = K, K-1, \ldots, 1$ **do**
(Back substitution) Compute the unconstrained LS solution $b_k^{\mathrm{LS}} := (y'_k - \sum_{\ell=k+1}^{K} R_{k\ell}\hat{b}_\ell^{\mathrm{DDD}})/R_{kk}$.
(Quantize to $\mathcal{A}$) Set $\hat{b}_k^{\mathrm{DDD}} := \lfloor b_k^{\mathrm{LS}} \rceil$.
(Compare with 0) If $2b_k^{\mathrm{LS}}\lfloor b_k^{\mathrm{LS}} \rceil - \lfloor b_k^{\mathrm{LS}} \rceil^2 - 2\lambda/R_{kk}^2 \leq 0$, then set $\hat{b}_k^{\mathrm{DDD}} := 0$.
end for

where $b_k^{\mathrm{LS}} := (y'_k - \sum_{\ell=k+1}^{K} R_{k\ell}\hat{b}_\ell^{\mathrm{DDD}})/R_{kk}$, and $\lfloor \cdot \rceil$ quantizes to the nearest point in $\mathcal{A}$. The simple implementation steps are tabulated as Algorithm 5.1.

When $\mathbf{R}$ is diagonal, Algorithm 1 yields the optimal S-MAP detection result; i.e., $\hat{\mathbf{b}}^{\mathrm{MAP}} = \hat{\mathbf{b}}^{\mathrm{DDD}}$ for this case. However, since the DDD detects symbols sequentially, it is prone to error propagation, especially at low SNR values. The error propagation can be mitigated by preprocessing and ordering methods [93, Ch. 7]. Also similar to all related detectors that rely on back substitution, error performance of the sparse DDD will be analyzed assuming there is no error propagation. For the special case of $M = 2$, Appendix D.2 shows that under this assumption, the SER for (5.21) becomes

$$P_{e,k}^{\mathrm{DDD}} = 2(1-p_a)Q\left(\frac{|R_{kk}|}{2} + \frac{\lambda}{|R_{kk}|}\right) + p_a Q\left(\frac{|R_{kk}|}{2} - \frac{\lambda}{|R_{kk}|}\right). \tag{5.22}$$

For a general M-ary constellation, it is also possible to approximate the SER using the union bound.

Because it accounts for the finite-alphabet constraint, sparse DDD outperforms the relaxed detectors of the previous section – a fact that will be confirmed also by simulations. However, the error propagation emerging at medium-low SNR degrades the sparse DDD performance when compared to the optimal but computationally complex S-MAP detector. As a compromise, a branch-and-bound type of MUD algorithm is developed next, to attain (near-) optimal performance by exploiting the finite-alphabet and sparsity constraints, at the price of increased complexity compared to DDD.

5.4.2 Sparsity-Exploiting Sphere Decoding-based MUD

Sphere decoding (SD) algorithms have been widely used for maximum-likelihood (ML) demodulation of multiuser and/or multiple-input multiple-output (MIMO) transmissions. Given a PAM or QAM alphabet, SD yields (near-) ML performance at polynomial

average complexity; see e.g., [39, Sec. 5.2]. However, different from the ML-optimal SD that minimizes an LS cost, the S-MAP problem (5.19) entails also a regularization term to account for sparsity in $\mathbf{b}$. Although the resultant algorithm will be termed sparse sphere decoder (SSD), it searches in fact within an "ℓ_0-*norm regularized* sphere," which is not a sphere but a hyper-solid.

The goal is to find the unknown $K \times 1$ vector $\mathbf{b} \in \mathcal{A}_a^K$, which minimizes the distance metric [cf. (5.19)]

$$D_1^K(\mathbf{b}) := \sum_{k=1}^{K} \left\{ \frac{1}{2} \left(y'_k - \sum_{\ell=k}^{K} R_{k\ell} b_\ell \right)^2 + \lambda |b_k|_0 \right\}. \tag{5.23}$$

For a large enough threshold τ, candidate vectors (and thus the minimizer of D_1^K too) satisfy[1]

$$D_1^K(\mathbf{b}) < \tau \tag{5.24}$$

that specifies a hyper-solid inside which the wanted minimizer must lie. Define now [cf. (5.21)]

$$\rho_k := \left(y'_k - \sum_{\ell=k+1}^{K} R_{k\ell} b_\ell \right) \Bigg/ R_{kk}, \quad k = K, K-1 \cdots 1 \tag{5.25}$$

where $\rho_K := y'_K / R_{KK}$. Note that ρ_k depends on $\{b_\ell\}_{\ell=k+1}^K$.

Using (5.25), the hyper-solid in (5.24) can be expressed as

$$D_1^K(\mathbf{b}) = \sum_{k=1}^{K} \left\{ \frac{R_{kk}^2}{2} (\rho_k - b_k)^2 + \lambda |b_k|_0 \right\} < \tau \tag{5.26}$$

or, in a more compact form as $D_1^K(\mathbf{b}) = \sum_{k=1}^K d_k(b_k) < \tau$, where $d_k(b_k) := (R_{kk}^2/2)(\rho_k - b_k)^2 + \lambda |b_k|_0$. In addition to the overall metric D_1^K assessing a candidate vector $\mathbf{b} \in \mathcal{A}_a^K$, as well as the per entry metric d_k for each candidate symbol $b_k \in \mathcal{A}_a$, it will be useful to define the *accumulated* metric $D_k^K := \sum_{\ell=k}^K d_\ell(b_\ell)$ corresponding to the $K-k+1$ candidate symbols from entry K down to entry k.

Per entry k, which subsequently will be referred to as level k, eq. (5.26) implies a set of inequalities:

$$\text{Level } k: \quad d_k(b_k) < \tau - D_{k+1}^K, \qquad \text{for} \quad k = K, K-1 \ldots 1, \tag{5.27}$$

[1] At initialization, τ is set equal to ∞ so that (5.24) is always satisfied.

with $D_{K+1}^K := 0$. SSD relies on the Schnorr-Euchner (SE) enumeration, see e.g., [26], properly adapted here to account for the ℓ_0-norm regularization. SE capitalizes on the inequalities (5.27) to search efficiently over all possible vectors $\mathbf{b}$ with entries belonging to $\mathcal{A}_a$. Any candidate $\mathbf{b} \in \mathcal{A}_a^K$ obeying the K inequalities in (5.27) for a given τ, will be termed *admissible*. Threshold τ is reduced after each admissible $\mathbf{b}$ is identified, as will be detailed soon. The SE-based SSD amounts to a *depth-first* tree search, which seeks and checks candidate vectors starting from entry (level) K and working downwards to entry 1 per candidate vector.

At level K, SE search chooses $b_K = \lfloor \rho_K \rceil \mathbb{1}(2\rho_K \lfloor \rho_K \rceil - \lfloor \rho_K \rceil^2 - 2\lambda/R_{KK}^2 > 0)$, which we know from (5.21) is the constellation point yielding the smallest d_K. If this choice of b_K does not satisfy the inequality (5.27) with $k = K$, no other constellation point will satisfy it either, and the minimizer of D_1^K in (5.23) must lie outside[2] the hyper-solid postulated by (5.24). If this choice of b_K satisfies (5.27), SE proceeds to level $K-1$ in which (5.25) is used first with $k = K-1$ to find ρ_{K-1} that depends on the chosen b_K from level K; subsequently, b_{K-1} is selected as $b_{K-1} = \lfloor \rho_{K-1} \rceil \mathbb{1}(2\rho_{K-1} \lfloor \rho_{K-1} \rceil - \lfloor \rho_{K-1} \rceil^2 - 2\lambda/R_{K-1,K-1}^2 > 0)$. If this choice of b_{K-1} does not satisfy (5.27) with $k = K-1$, then we move back to level K, and select b_K equal to the constellation point yielding the *second* smallest d_K, and so on; otherwise, we proceed to level $K-2$. Continuing this procedure down to level 1, yields the first candidate vector $\hat{\mathbf{b}}$, which is deemed admissible since it has entries belonging to $\mathcal{A}_a$ and also satisfying (5.24). This candidate is stored, and the threshold is updated to $\tau := D_1^K(\hat{\mathbf{b}})$.

Then, the search proceeds looking for a better candidate. Now at level 1, we move up to level 2 and choose b_2 equal to the constellation point yielding the *second* smallest cost d_2. If this b_2 satisfies (5.27) at level 2 with the current τ, we move down to level 1 to update the value of b_1 (note that b_2 has just been updated and $\{b_\ell\}_{\ell=3}^K$ are equal to the corresponding entries in $\hat{\mathbf{b}}$). If (5.27) at level 2 is not satisfied with the current τ, we move up to level 3 to update the value of b_3, and so on.

Finally, when it fails to find any other admissible candidate satisfying (5.27), the search stops, and the latest admissible candidate $\hat{\mathbf{b}}$ is the optimal $\hat{\mathbf{b}}^{\text{MAP}}$ solution sought. With $\tau = \infty$, the first found admissible candidate $\hat{\mathbf{b}}$ is the $\hat{\mathbf{b}}^{\text{DDD}}$ solution of Sec. 5.4.1.

[2] This will never happen with $\tau = \infty$ in (5.24).

Before summarizing the SSD steps, it is prudent to elaborate on the *ordered* enumeration of the constellation points *per level*, which in fact constitutes the main difference of SSD relative to SD. In lieu of the 0 constellation point and the ℓ_0 norm, SE in SD enumerates the PAM symbols per level k in the order of increasing cost as: $\{b_k, b_k + 2\Delta_k, b_k - 2\Delta_k, b_k + 4\Delta_k, b_k - 4\Delta_k, \cdots\} \bigcap \mathcal{A}$, with $b_k := \lfloor \rho_k \rceil$ and $\Delta_k := \text{sign}(\rho_k - b_k)$. (With $\lfloor \rho_k \rceil$ yielding the smallest d_k, if $\Delta_k = 1$, then $\lfloor \rho_k \rceil + 2$ yields the second smallest d_k and $\lfloor \rho_k \rceil - 2$ the third; and the other way around, if $\Delta_k = -1$.) SD effects such an ordered enumeration by alternately updating $b_k = b_k + 2\Delta_k$ and $\Delta_k = -\Delta_k - \text{sign}(\Delta_k)$, [26]. To demonstrate how SSD further accounts for the ℓ_0-norm regularization and the augmented alphabet of S-MAP which includes 0, let $b_k^{(i)} \in \mathcal{A}_a$ denote the symbol for level k incurring the i-th smallest ($i = 1, 2, \ldots, M+1$) cost d_k. If $b_k^{(i)} \in \mathcal{A}$, then $b_k^{(i+1)}$ will be either 0 or $b_k^{(i)} + 2\Delta_k$. If $d_k(0) < d_k(b_k^{(i)} + 2\Delta_k)$, then the next symbol in the ordered enumeration should be $b_k^{(i+1)} = 0$, and an auxiliary variable $b_k^{(c)}$ is used to cache the subsequent symbol in the order as $b_k^{(i+2)} = b_k^{(i)} + 2\Delta_k$. With $b_k^{(i+1)} = 0$, the auxiliary variable allows the wanted $b_k^{(i+2)}$ at the next enumeration step to be retrieved from $b_k^{(c)}$.

Similar to SD, the ordered enumeration pursued by SSD per level implies a corresponding order in considering all $\mathbf{b} \in \mathcal{A}_a^K$, which leads to a repetition-free and exhaustive search of all admissible candidate vectors. At the same time, the hyper-solid postulated by (5.24) shrinks as τ decreases, until no other admissible vector can be found. This guarantees that the SSD outputs the vector with the smallest D_1^K, and thus the optimal solution to (5.19). The SSD algorithm can be summarized in the following six steps 1–6 tabulated as Algorithm 5.2.

Remark 5.4 *(SSD computational complexity).* SSD inherits all the attractive features of SD [26]. Specifically, during the search one basically needs to store D_k^K per level k. Its *in place update* for each b_k candidate implies that SSD *memory requirements* are only linear in K. In addition, the *computational efficiency* of SSD (relative to that of ML which is $\mathcal{O}((M+1)^K)$) stems from four factors: (i) the DDD solution provides an admissible initialization reducing the search space at the outset; (ii) the recursive search enabled by the QR decomposition gives rise to the causally dependent inequalities (5.27), which restrict admissible candidate entries to choices that even decrease over successive

Algorithm 5.2 (SSD): Input $\tau = \infty$, $\mathbf{y}'$, $\mathbf{R}$, λ, and M even. Output the solution $\hat{\mathbf{b}}^{\text{MAP}} := \hat{\mathbf{b}}^{\text{SSD}}$ to (5.19).

1: (Initialization) Set $k := K$, $D_{k+1}^K := 0$.

2: Compute ρ_k as in (5.25), set $b_k := \lfloor \rho_k \rceil$, $b_k^{(c)} := 0$, and $\Delta_k := \text{sign}(\rho_k - b_k)$.
If $2\rho_k b_k - b_k^2 - 2\lambda/R_{kk}^2 < 0$, then // `symbol 0 yields smaller` d_k `than` $\lfloor \rho_k \rceil$
Set $b_k^{(c)} := b_k$, and $b_k := 0$.
End if and go to Step 3.

3: **If** $d_k(b_k) := (R_{kk}^2/2)(\rho_k - b_k)^2 + \lambda |b_k|_0 \geq \tau - D_{k+1}^K$, then
go to Step 4. // `outside hyper-solid in` (5.24)
Else if $|b_k| > M - 1$, then
go to Step 6. // `inside hyper-solid in` (5.24)`, but outside` $\mathcal{A}_a$
Else if $k > 1$, then
set $D_k^K := D_{k+1}^K + d_k(b_k)$ and $k := k - 1$, go to Step 2. // `deeper in the tree`
Else go to Step 5. // $k = 1$`, at the tree's bottom`
End if

4: **If** $k = K$, then terminate
Else set $k := k + 1$, go to Step 6.
End if

5: (An admissible $\mathbf{b}$ is found)
Set $\tau := D_2^K + d_1(b_1)$, $\hat{\mathbf{b}}^{\text{SSD}} := \mathbf{b}$, and $k := k + 1$; then go to Step 6.

6: (Enumeration at level k proceeds to the candidate symbol next in the order)
If $b_k = 0$, then
Retrieve the next (based on d_k ordering) $b_k := b_k^{(c)}$, and flag $b_k^{(c)} := FLAG$.
Else set $b_k := b_k + 2\Delta_k$, and $\Delta_k := -\Delta_k - \text{sign}(\Delta_k)$.
If $b_k^{(c)} \neq FLAG$ and $2\rho_k b_k - b_k^2 - 2\lambda/R_{kk}^2 < 0$, then
Set $b_k^{(c)} := b_k$, and $b_k := 0$.
End if
End if and go to Step 3.

depth-first passes of the search; (iii) ordering per level increases the likelihood of finding "near-optimal admissible" candidates early, which means quick and sizeable shrinkage of the hyper-solid, and thus fast convergence to the S-MAP optimal solution; and (iv) metrics involved in the search can be efficiently reused since children of the same level in the tree share the already computed accumulated metric of the "partial path" from this level to the root.

Compared to other sub-optimal detection schemes proposed in previous sections, the SSD algorithm can return the S-MAP optimal solution possibly at exponential complexity, unless one stops the search at the affordable complexity – case in which the solution is only ensured to be near-optimal. Fortunately, at medium-high SNR, both

SD and SSD return the optimal solution at average complexity which is polynomial (typically cubic). Moreover, SSD can be generalized to provide symbol-by-symbol soft output with approximate *a posteriori* probabilities, as is the case with the SD; see e.g., [39, Ch. 5].

5.5 Generalizations of S-MAP Detectors

Up to now, four sparsity-exploiting MUD algorithms have been developed to solve the integer program associated with the linear model in (5.1). The present section will present interesting generalizations to account for correlated user (in)activity across symbols, and under-determined CDMA systems.

5.5.1 Exploiting user (in)activity across symbols

Sparsity-aware detectors for the linear model in (5.1) were so far developed on a symbol-by-symbol basis, which does not account for the fact that user (in)activity typically persists across multiple symbols. To this end, user activity across time can for instance be thought of as a Markov chain with two states (active-inactive). Once a user terminal starts transmitting to the AP, it becomes more likely to stay active for the next symbol slot too; and likewise, inactive once it stops transmitting. In this model, the state transition probability from either one state to the other is relatively much smaller than that of staying unchanged, and this manifests itself to the said dependence of user (in)activities across time.

Admittedly, MUD schemes accounting for this dependence must process the aggregation of data vectors $\mathbf{y}$ obeying (5.1) across slots. With N_s denoting the number of slots, the number of unknowns (KN_s) can grow prohibitively with N_s. One approach to cope with this "curse of dimensionality" is via dynamic programming, which can take advantage of the fact that correlation is only present between two consecutive slots; see e.g., [93, Sec. 4.2]. However, for M-ary alphabets, the resultant sequential detector requires evaluating per symbol slot the path weights of all $(M+1)^K$ possible symbol vectors. This high computational burden is impractical for real-time implementations.

The proposed alternative to bypass this challenge stems from the observation that for a given slot t the influence of all previous slots $\{t'\}_{t'=0}^{t-1}$ on the S-MAP detection rule

is reflected only in the prior probability of each user being active at time t; i.e., user k's current (and time-varying) activity factor $p_{a,k}(t)$. The natural means to capture this influence online is to track each user's activity factor using the recursive LS (RLS) estimator [75, Ch. 12], based on activity factors from previous slots; that is,

$$\hat{p}_{a,k}(t) = \arg\min_{p} \sum_{t'=0}^{t-1} \beta_k^{t,t'} (p - |\hat{b}_k(t')|_0)^2, \quad t = 1, \ldots \tag{5.28}$$

where $\hat{b}_k(t')$ denotes user k's detected symbol at time t', and the so-called "forgetting-factor" $\beta_k^{t,t'}$ describes the effective memory (data windowing). A popular choice is the exponentially decaying window, for which $\beta_k^{t,t'} := \beta_k^{t-t'}$ for some $0 \ll \beta_k < 1$. Accordingly, (5.28) can expressed in closed form, recursively as

$$\begin{aligned} \hat{p}_{a,k}(t) &= \frac{1-\beta_k}{\beta_k} \left(\sum_{t'=0}^{t-1} \beta_k^{t-t'} |\hat{b}_k(t')|_0 \right) / (1-\beta_k^t) \\ &= \frac{\beta_k - \beta_k^t}{1-\beta_k^t} \hat{p}_{a,k}(t-1) + \frac{1-\beta_k}{1-\beta_k^t} |\hat{b}_k(t-1)|_0, \quad t = 1, \ldots \end{aligned} \tag{5.29}$$

where the last equality comes from back substitution of $\hat{p}_{a,k}(t-1)$. The choice of β_k critically depends on the (in)activity correlation between consecutive slots. In the extreme case where user (in)activities across slots are independent, the infinite-memory window ($\beta_k = 1$) is optimal, and (5.29) reduces to the simple online time-averaging estimate $\hat{p}_{a,k}(t) = \frac{1}{t} \sum_{t'=0}^{t-1} |\hat{b}_k(t')|_0$.

Adapting the user activity factors allows one to weigh entries of ℓ_0-norm regularization which in turn affects the prior probability in the S-MAP detector (5.5) through the coefficient $\lambda_k(t)$ corresponding to $\hat{p}_{a,k}(t)$ (cf. Remark 5.2). Note that when the correlation across time is strong, it is possible that $\hat{p}_{a,k}(t)$ can approach 1, case in which $\lambda_k(t)$ is not guaranteed to stay positive. This will cause problems to the relaxed S-MAP detectors of Sec. 5.3, as those schemes rely on the convexity of the cost function in (5.7). However, the DDD and SSD algorithms will remain operational, because they rely on enumeration per symbol (in DDD) or per group of symbols within a sphere (in SSD). Note also that with $\lambda < 0$ the regularization term in the minimization of (5.19) is non-positive; hence, $\lambda < 0$ encourages searching over the non-zero constellation points (and thus discourages sparsity), whereas $\lambda > 0$ promotes sparsity.

5.5.2 Under-determined CDMA Systems

Minimal-size spreading sequences, even smaller than the number of users, is well motivated for bandwidth and power savings. Without finite-alphabet constraints on the wanted vector, results available in the CS literature guarantee recovery of sparse signals from a few observations; see e.g., [16,17] and references therein. Specifically, [16] shows that if the vector of interest is sparse (or compressible) over a known basis, then it is possible to reconstruct it with very high accuracy from a small number of random linear projections at least in the ideal noise-free case. For non-ideal observations corrupted with unknown noise of bounded perturbation, [17] provides an upper bound on the reconstruction error, which is proportional to the noise variance for a sufficiently sparse signal. However, CS theory pertains to sparse analog-valued signals. Moreover, the noise considered in a practical communication system is typically Gaussian, or generally drawn from a distribution having possibly unbounded support. Therefore, existing results from the CS literature do not carry over to the present context.

Nevertheless, it is still interesting to consider extensions of all the (sub-)optimal sparsity-exploiting MUD algorithms to an under-determined CDMA system with $N < K$, where the observation matrix $\mathbf{H}$ becomes fat. Consider first the two types of relaxed S-MAP detectors. The RD-MUD in (5.10) clearly works when $N < K$, since the $2\lambda\mathbf{I}$ term inside the inversion renders the overall matrix full rank. However, since the RD is a linear detector, it is expected to lose identifiability in the under-determined case, similar to the MMSE detectors for the sparsity-agnostic MUD schemes. Similarly, the LD problem (5.17) is also solvable for a fat $\mathbf{H}$ matrix, as the Lasso problem in CS. However, neither of them accounts for the augmented finite-alphabet constraint present in the original S-MAP problem (5.7).

The S-MAP detectors with lattice search are challenging to implement when $N < K$. The main obstacle is the QR decomposition of the fat matrix $\mathbf{H}$, which yields the upper triangular matrix $\mathbf{R}$ of the same dimension $N \times K$. Instead of a single unknown symbol, now the sparse DDD must optimize over the last $(N-K+1)$ symbols in $\mathbf{b}$. However, apart from exhaustive search there is no low-complexity method to solve the aforementioned problem involving $(N-K+1)$ variables, because sub-optimal alternatives introduce severe error propagation.

The same problem appears also with the SSD. To tackle the under-determined case,

the generalized SD in [26] fixes the last $(N-K)$ symbols of $\mathbf{b}$ and relies on the standard SSD to detect the remaining K symbols that minimize a cost similar to the one in (5.23). Repeating this search for every choice of the last $(N-K)$ symbols, yields eventually the overall optimum vector. The complexity of the latter is exponential in $(N-K)$, regardless of the SNR. Recently, an alternative SD approach to avoid this exponential complexity has been developed for the under-determined case [24]. This algorithm takes advantage of the fact that for constant-modulus constellations the usual LS cost can be modified without affecting optimality, by adding the ℓ_2-norm $\mathbf{b}^T\mathbf{b}$ constant for every vector in the alphabet. This extra term allows one to obtain an equivalent full-rank system on which the standard SD algorithm can be applied. This efficient method can be readily extended to handle non-constant modulus constellations.

Interestingly, for our S-MAP detectors in (5.7) with lattice search of binary transmitted symbols, the norm term needed for regularization comes naturally from the Bernoulli prior. Specifically, with $p=2$ the reformulated S-MAP detectors in (5.7) can be equivalently written as

$$\begin{aligned}\hat{\mathbf{b}}^{\mathrm{MAP}} &= \arg\min_{\mathbf{b}\in\mathcal{A}_a^K} \frac{1}{2}[\mathbf{b}^T(\mathbf{H}^T\mathbf{H}+2\lambda\mathbf{I})\mathbf{b} - 2\mathbf{y}^T\mathbf{H}\mathbf{b}] \\ &= \arg\min_{\mathbf{b}\in\mathcal{A}_a^K} \frac{1}{2}\|\tilde{\mathbf{y}}' - \tilde{\mathbf{R}}\mathbf{b}\|_2^2 \end{aligned} \tag{5.30}$$

where $\tilde{\mathbf{R}}$ is the full rank $K\times K$ upper triangular matrix such that $\tilde{\mathbf{R}}^T\tilde{\mathbf{R}} = \mathbf{H}^T\mathbf{H}+2\lambda\mathbf{I}$, and $\tilde{\mathbf{y}}' := \tilde{\mathbf{R}}^{-T}\mathbf{H}^T\mathbf{y}$. Utilizing the metric of (5.30), the back substitution of DDD and the lattice point search of SSD can be implemented easily. Hence, these S-MAP detectors can be readily extended to under-determined systems. In this way, all the (sub-)optimal S-MAP detectors are applicable with less observations than unknowns in a CDMA system with low activity factor, but their SER performance will certainly be affected. In Sec. 5.6, we will provide simulated performance comparisons of the different optimal and sub-optimal S-MAP algorithms proposed, for a variable number of observations.

5.5.3 Group Lassoing block activity

The last generalization considered pertains to user (in)activity in a (quasi-)synchronous block fashion, where during a block of N_s symbol slots, user k remains (in)active independently from other users and across blocks. Concatenate the K user symbols across N_s time slots in the $K \times N_s$ matrix $\mathbf{B} := [\mathbf{b}(1) \dots \mathbf{b}(N_s)]$, where $\mathbf{b}(t)$ collects the symbols of all K users at slot t, and likewise for the receive-data matrix $\mathbf{Y}$ as well as the noise matrix $\mathbf{W}$, both of size $N \times N_s$. With these definitions, the counterpart of (5.1) for this block model is $\mathbf{Y} = \mathbf{HB} + \mathbf{W}$. Letting the $N_s \times 1$ vector $\breve{\mathbf{b}}_k := [b_k(1), \dots, b_k(N_s)]^T$ collect the N_s symbols of user k, it is useful to consider it drawn from an augmented (due to possible inactivity) block alphabet $\mathcal{A}_{a,N_s} := \mathcal{A}^{N_s} \bigcup \{\mathbf{0}_{N_s}\}$. Assuming again binary transmissions, the S-MAP *block* detector will now yield

$$\begin{aligned} \hat{\mathbf{B}}^{\text{MAP}} &= \arg \min_{\breve{\mathbf{b}}_k \in \mathcal{A}_{a,N_s}} \frac{1}{2}\|\mathbf{Y} - \mathbf{HB}\|_F^2 + \sum_{k=1}^{K} \frac{\lambda_b}{\sqrt{N_s}} \|\breve{\mathbf{b}}_k\|_0 \\ &= \arg \min_{\breve{\mathbf{b}}_k \in \mathcal{A}_{a,N_s}} \frac{1}{2}\|\mathbf{Y} - \mathbf{HB}\|_F^2 + \sum_{k=1}^{K} \lambda_b \|\breve{\mathbf{b}}_k\|_2 \end{aligned} \tag{5.31}$$

where $\lambda_b := \frac{1}{\sqrt{N_s}} \ln\left(\frac{1-p_a}{p_a/2^{N_s}}\right)$.

Similar to (5.7), the convex reformulation of the cost in (5.31) will lead to what is referred to in statistics as *Group Lasso* [99], which effects group sparsity on a block of symbols ($\breve{\mathbf{b}}_k$ in our case). This Group-Lasso based formulation is particularly handy for under-determined CDMA systems. In fact, the unconstrained version of (5.31) can be solved first to unveil the nonzero rows (i.e., the support) of $\mathbf{B}$, with improved reliability as N_s increases. Subsequently, standard sparsity-agnostic MUD schemes can be run on the estimated set of active users. Note that such a two-step approach works in the under-determined case, and also reduces number of symbols to be detected per time slot.

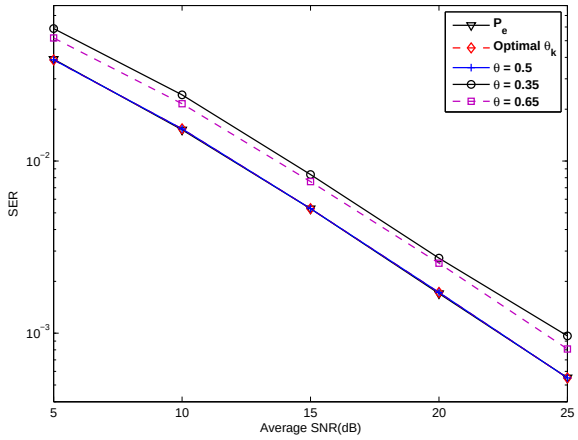

Figure 5.2: SER vs. average SNR (in dB) for RD-MUD with $N = 32$ and $K = 20$ and different quantization threshold θ's.

5.6 Simulations

We simulate a CDMA system with $K = 20$ users, each with activity factor $p_a = 0.3$. Random sequences with length $N = 32$ are used for spreading. We consider non-dispersive independent Rayleigh fading channels between AP and users, where the channel gain g_k of the k-th user is Rayleigh distributed with variance $E[g_k^2] = \sigma^2$. Thus, the average system SNR is set to be σ^2 since the AWGN $\mathbf{w}$ has unit variance.

Test Case 5.1 (Quantization thresholds for RD): First, we test the RD scheme with different quantization thresholds θ in (5.9). The optimum threshold for the k-th symbol is obtained as in (5.15) per channel realization $\mathbf{H}$. The resulting SER is compared for four choices of θ: 0.5, 0.35, 0.65, and $\hat{\theta}_k$. The theoretical minimum SER $\hat{P}_{e,k}^{\mathrm{RD}}$ in (5.16) using the optimum θ is also added for comparison. Fig. 5.2 shows that the SER curve with $\theta = 0.5$ comes very close to the one of the optimal $\hat{\theta}_k$, and thus constitutes a near-optimal choice in practice. Moreover, those two curves also coincide with the analytical SER formulation corresponding to $\hat{P}_{e,k}^{\mathrm{RD}}$, thus corroborating the closed-form expression in (5.16).

Test Case 5.2 (S-MAP MUD algorithms): Next, the RD, LD, DDD, and SSD MUD

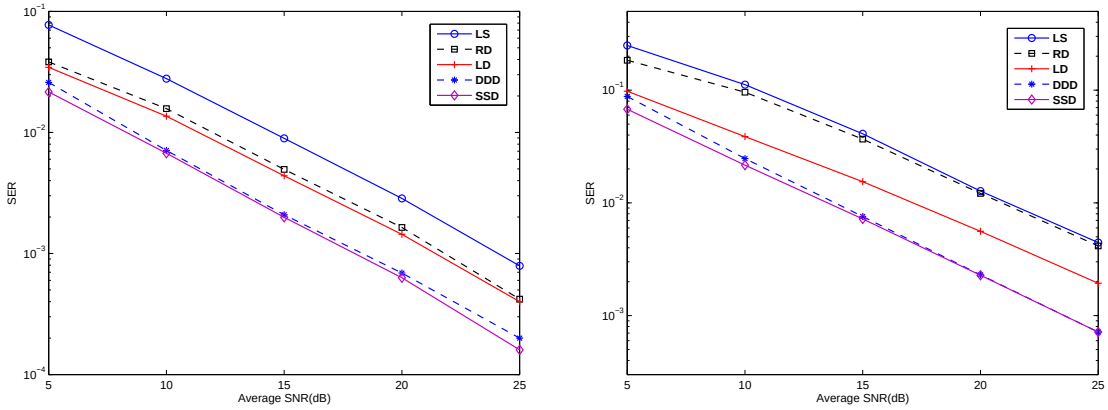

Figure 5.3: SER vs. average SNR (in dB) of sparsity-exploiting MUD algorithms with $N = 32$ and $K = 20$ for (a) BPSK, and (b) 4-PAM alphabets.

algorithms are all tested for both BPSK and 4-PAM constellations, and their SER performance is compared. For LD, the quadratic program in (5.17) is solved using the general-purpose SeDuMi toolbox [79]. The quantization rule chooses the nearest point in $\mathcal{A}_a$ for both RD and LD. For comparison, we also include the ordinary LS detector, which corresponds to the RD solution in (5.10) with $\lambda = 0$.

Fig. 5.3(a) shows that the LS detector exhibits the worst performance. This is intuitive since it neither exploits the finite alphabet nor the sparsity present in $\mathbf{b}$. The SSD exhibits the best performance at the price of highest complexity. The LD outperforms the RD algorithm, as predicted. It is interesting to observe that even at low SNR region the DDD algorithm is surprisingly competitive, especially in view of its low complexity that grows only linearly in the number of symbols K. The diversity orders for those detectors are basically the same. This is reasonable since independent Rayleigh fading channels between AP and users were simulated here. The corresponding curves for 4-PAM depicted in Fig. 5.3(b) follow the same trend. However, compared to Fig. 5.3(a), the RD algorithm degrades noticeably as its SER approaches the LS one. This is because choices of quantization thresholds become more influential as the constellation

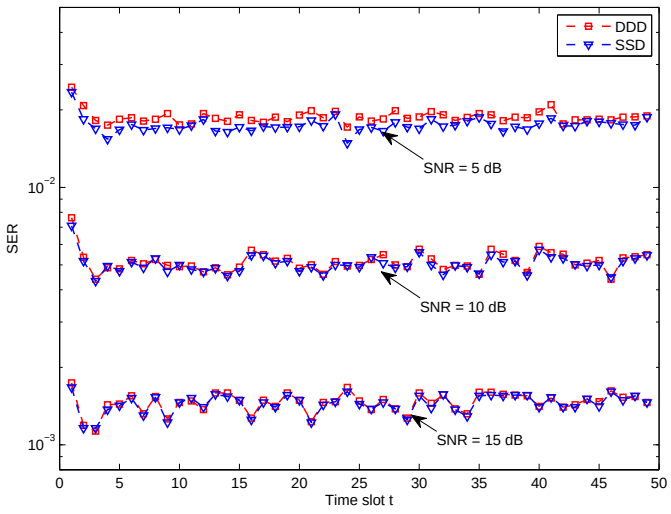

Figure 5.4: SER vs. time t of sparsity-exploiting MUD algorithms with RLS estimation of the activity factors.

size increases. As expected, the LD exhibits resilience to this influence. The DDD and SSD algorithms have almost identical performance in high-SNR region.

Test Case 5.3 ((In)activity across symbols): In this case, the user (in)activity is correlated across time slots. We model this random (in)activity process as a two-state (active-inactive) stationary Markov chain. The state transition matrix is $\mathbf{P} = [a \;\; (1-a); \;\; b \;\; (1-b)]$, where a is uniformly distributed over [0.8 0.85], and b over [0.05 0.1], for each user. For this model, the expected number of successive active slots is $1/(1-a)$, and $1/b$ for the inactive ones. Also, the limiting probability for the "active" state becomes $b/(1-a+b)$, taking values from the interval [0.2 0.4]. Note that the activity factor over time is still quite low. We use the RLS approach to estimate $\hat{p}_{a,k}(t)$ as in (5.29) using $\beta = 0.5$, and test both the DDD and SSD algorithms in solving the resultant S-MAP detection problem. The empirical SER is plotted in Fig. 5.4 across time for different SNR values. Clearly, the proposed scheme is effective in tracking the evolving time-correlated activity. For the same SNR value, it yields SER performance similar to the independent case in Fig. 5.3(a).

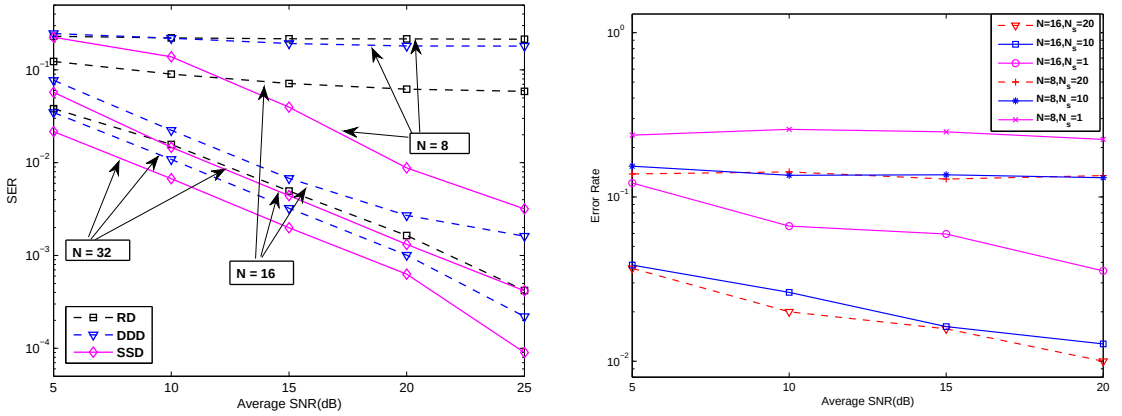

(a) (b)

Figure 5.5: (a) SER vs. average SNR (in dB) of sparsity-exploiting MUD algorithms with $N = 32, 16$, or 8 and $K = 20$; (b) Error rate for block activity vs. average SNR (in dB) of Group Lasso algorithm with $N = 16, 8$ and $N_s = 20, 10, 1$.

Test Case 5.4 (Under-determined CDMA systems): We also test the S-MAP MUD algorithms for under-determined systems, by varying N from 32 to 16 and 8. The results are depicted in Fig. 5.5(a). Since the RD is a simple linear detector, it is expected that once $N < K$, it will lose identifiability, and exhibits a considerably flat SER curve. At the same time, DDD still enjoys almost full diversity with a moderate choice of $N = 16$. Being the optimum detector, the SSD collects the full diversity even if $N = 8$; however, the other two kinds of detectors exhibit flat SER curves, as expected.

The Group Lasso scheme for recovering block activity is also included for the under-determined case. Fig. 5.5(b) illustrates the activity recovery error rate for different values of N and N_s. The number of observations N affects the diversity order, while the block size N_s influences the recovery accuracy.

Chapter 6

Conclusions and Directions for Future Research

In this Thesis we dealt with sparsity-cognizant algorithms for applications in communications, signal processing, and the smart power grid. In this final chapter, we provide a summary of the main results in the Thesis, and point out possible directions for future research.

6.1 Thesis Summary

In Chapter 2, a novel R-SE formulation is developed for power system monitoring, which adopts an overcomplete model for the *outlier* vector and exploits its underlying sparsity. The R-SE problem is then reformulated by leveraging the nonlinear relationship between legacy meter measurements and complex bus voltages. Hence, the nonconvex R-SE problem can be relaxed to a convex SDP one, which renders it efficiently solvable via existing interior-point methods. The SDR-based approach to R-SE jointly estimates the state of voltage phasors, and identifies the outlier using statistical prior information regarding the system model. In addition, the convex dual R-SE problem is analyzed to provide insights on the dual equivalence and guidelines to reduce complexity. Corroborating numerical tests on the 30-bus benchmark system demonstrate near-optimal performance of the novel approaches.

Corroborated by numerical examples, the proposed (R-)SE algorithm well approximates the global optimal solution, regardless of the initial guess. The SDR-based SE framework has also been extended to include linear synchrophasor measurements.

Another aspect of power system monitoring relates to comprehensive situational awareness, which is more stringent due to the local administration of the current power grid. To improve this aspect, Chapter 3 develops novel line-change identification algorithms, through the *overcomplete* representations capturing the sparse innovation manifested by line changes on phasor angle measurements. The proposed approach allows for identification of multiple (even cascaded) power lines in outage using two schemes with complementary strengths. The first one entails successive (greedy) inclusion of the lines in outage, and it is attractive for its computational efficiency. The second scheme regularizes an LS error criterion with the ℓ_1-norm of the sparse vector comprising the overcomplete representation coefficients, and thus converts the problem of line outage identification to a computationally tractable convex QP, allowing for a CD-based solver to compute the entire line-outage identification path for any prescribed level of sparsity.

Practical statistical tests are also adapted to identify the actual number of lines in outage. The framework developed for linear DC power flow models is shown possible to generalize to nonlinear AC models, as well as capable to accommodate prior information available on the internal system noise, and further flexible to cope with generic line changes. Numerical tests demonstrated the merits of the proposed schemes in unveiling multiple lines in outage, even when those reside at the external system for which phasor angle data are not available in real time.

While Chapters 2 and 3 mainly focused on leveraging sparsity for challenging and urgent issues closely related to smartly operating the power grid, this Thesis also aspires to touch upon general challenges in the framework of CS and sparse signal recovery, with intriguing applications in contemporary wireless communications and array signal processing. Specifically, Chapter 4 proposes a *sparse (S-) TLS* approach to fitting sparse, fully-perturbed, linear models, which complements both CS algorithms and the (regularized) TLS methods; while in Chapter 5 a sparse detection approach is pursued to deal with sparse signal vectors whose entries adhere to *finite-alphabet (FA)* constraints.

Specifically, for the S-TLS framework in Chapter 4, an innovative approach is developed to account for sparsity in estimating coefficient vectors of fully-perturbed linear

regression models. This approach enriches TLS criteria that have been traditionally used to fit such models with the ability to handle under-determined linear systems. The novel S-TLS framework also enables sparsity-exploiting approaches (CS, BP, and Lasso) to cope with perturbations present not only in the data but also in the regression matrix.

Near-optimum and reduced-complexity suboptimum solvers with global and local convergence guarantees were (or have been) also developed to optimize the generally nonconvex S-TLS criteria. They rely on bisection, branch-and-bound, or coordinate descent iterations, and have universal applicability regardless of whether perturbations are modeled as deterministic or random. Valuable generalizations have also been provided when prior information is available on the deterministic structure or statistics of the associated (augmented) data matrix. Under specific statistical models with errors-in-variables, the resultant (generally weighted and structured) S-TLS estimators are provably optimal in the MAP sense. Simulated tests corroborate the analytical claims, compare competing alternatives, and demonstrate the practical impact of the novel S-TLS framework to grid-based sparsity-exploiting approaches for cognitive radio sensing, and direction-of-arrival estimation with possibly uncalibrated antenna arrays.

In addition to S-TLS, the MUD problem of sparse symbol vectors emerging within CDMA systems having low-activity factor is considered in Chapter 5. Viewing user inactivity as augmenting the underlying alphabet, the *a priori* available sparsity information has been exploited in designing optimal S-MAP detectors, whose exponential complexity can be reduced by resorting to (sub-) optimal algorithms. Relaxed S-MAP detectors (RD and LD) come with low complexity but sacrifice optimality, because they ignore the finite-alphabet constraint. The second kind of detectors (DDD and SSD), searches over a subset of the alphabet and exhibits improved performance at increased complexity. Performance analysis has been provided for the RD and DDD algorithms, and closed-form expressions derived for the corresponding symbol error rates.

S-MAP detectors have been further generalized to deal with correlated user (in)activity across symbols by recursively estimating each user's activity factor online; and also with under-determined (a.k.a. overloaded CDMA) settings emerging when the spreading gain is smaller than the potential number of users. Coping with the latter becomes possible through regularization with the ℓ_0-norm prior, or, with the Group Lasso-based recovery

of the active user set. Numerical tests corroborate our analytical findings and quantify the relative performance of the novel sparsity-exploiting MUD algorithms.

6.2 Future research

The results in this Thesis open up interesting directions for a number of future research topics. Next, we outline two of them that we pursue currently. The first one is a direct extension of the SDR-based R-SE algorithm in Chapter 2, using distributed and parallel implementations for computational savings. The second one pertains to recursive S-TLS for online estimation of power system electromechanical modes.

6.2.1 Multi-area (R-)SE using distributed SDP

Power system SE has been traditionally performed at regional control centers with limited interaction. However, due to the deregulation of energy markets, large amounts of power are transferred over high-rate, long-distance lines spanning several control areas [42]. These so-termed tie lines, originally constructed for emergency situations, are now fully operational and must be accurately monitored. Since any control area can be strongly affected by events and decisions elsewhere, regional operators can no longer operate in a truly independent fashion. At the same time, central processing of the current energy market faces two serious limitations: i) vulnerability to unreliable telemetry; and ii) high computational complexity at the control center. These considerations clearly suggest that distributed SE among multiple control areas is essential for realizing the smart grid vision.

To complement the SDR-based (R-)SE framework in Chapter 2, it is of great interest to devise efficient distributed computation modules for the SDP problem (2.13). However, this task is challenged by the positive semidefiniteness of the matrix $\mathbf{V}$, involving the voltage states across all areas. To this end, suppose that all buses in the network are partitioned so that $\mathcal{N} = \bigcup_k \mathcal{N}_k$. Let also vector $\mathbf{v}_k$ collect the state variables of area $\mathcal{N}_k$. Notice that the state vectors could be overlapping, as the power meter measurements along tie-lines involve voltage phasors of a subset of neighboring buses. The outer-product $\mathbf{V}_k = \mathbf{v}_k \mathbf{v}_k^{\mathcal{H}}$ of area $\mathcal{N}_k$ for instance, is also overlapping with $\mathbf{V}_l$, for each neighboring area $l \in \mathcal{A}_k$, which can be denoted by the submatrix equality $\mathbf{V}_k^l = \mathbf{V}_l^k$.

Hence, if the error criterion per area k can be captured by the function $f_k(\cdot)$, which could be either the original WLS one in (2.3) or the robust one in (2.5), the centralized SE problem is given by

$$\begin{aligned} \min_{\mathbf{V}} \quad & \sum_k f_k(\mathbf{V}_k) \\ \text{s.to} \quad & \mathbf{V} \succeq \mathbf{0}, \mathbf{V}_k^l = \mathbf{V}_l^k, \ \forall l \in \mathcal{A}_k, \ \forall k. \end{aligned} \tag{6.1}$$

The difficulty in implementing (6.1) in a distributed fashion lies in the constraint $\mathbf{V} \succeq \mathbf{0}$, which couples all local state matrices. Our idea here is to exploit the chordal property of a graph, which has been also used for distributed SDP-based optimal power flow problem [60]. Interestingly, for the multi-area SE problem considered here, it is possible to argue that under the tree-graph connection of all local areas, the centralized SE (6.1) will be equivalent to

$$\begin{aligned} \min_{\mathbf{V}} \quad & \sum_k f_k(\mathbf{V}_k) \\ \text{s.to} \quad & \mathbf{V}_k \succeq \mathbf{0}, \mathbf{V}_k^l = \mathbf{V}_l^k, \ \forall l \in \mathcal{A}_k, \ \forall k. \end{aligned} \tag{6.2}$$

This equivalent formulation overcomes the SDP couplings, and allows for popular distributed implementation modules such as the alternating-direction method of multipliers. This approach turns out to be resilient to imperfect area-to-area communication links, possibly affected by additive noise and random link failures [113]. Leveraging the distributed solver of (6.2) will be an interesting direction to realize multi-area SE using SDR for estimating nonlinear power system states.

6.2.2 Recursive S-TLS for power system mode estimation

Low-frequency electromechanical modes play a vital role in small-signal stability analysis of power systems [59]. Recently, with the availability of synchrophasor measurements across the power grid, it becomes feasible to develop real-time electromechanical mode estimation algorithms that are capable of processing newly arrived measurement data online. Mode estimation from both ringdown and ambient measurement data can be formulated as a linear prediction problem [101]. Specifically, the scalar measurements $y(t)$ per time index t can be modeled as a sum of damped exponentials, whose parameters

are unknown and determine the system electromechanical modes for stability analysis. In order to estimate the parameters of the exponentials involved, consider the following linear prediction problem

$$y(t) = x_1 y(t-1) + x_2 y(t-2) + \cdots + x_n y(t-n) \tag{6.3}$$

where the linear coefficients stacked in the vector $\mathbf{x}$ can be used to compute the sought parameters. By collecting (6.3) across time t, one ends up with a linear regression problem involving a Hankel-structured data matrix. Recursive LS (RLS) has been applied to estimate the vector $\mathbf{x}$ online, offering reduced computational costs relative to batch processing [101].

However, two challenges have not been addressed by prior electromechanical mode estimation works. First, the noise in $y(t)$ actually affects both the regression matrix as well as the data vector. However, the RLS (and any LS-based approach for that matter) takes into account only perturbations in the data vector. Second, the linear model in (6.3) is formulated with a given order n, but in practice it is not clear how n should be selected. Clearly, the model order needs to be large enough to allow sufficient degrees of freedom in the model, but no too large as to cause data over-fitting. Usually, various model orders are tested and their results are compared to the true value.

To tackle the two aforementioned issues, it will be interesting to investigate a recursive S-TLS approach. The (S-)TLS allows model perturbations in both the data matrix and vector, while the sparsity-induced regularization in $\mathbf{x}$ can handle the possible data over-fitting when using a larger n. Our idea here is to solve the recursive S-TLS (RS-TLS) problem as follows

$$\begin{aligned} &\min_{e(t)} \ \textstyle\sum_{t=1}^{T} (1/T)[e(t)]^2 + \lambda \|\mathbf{x}\|_1 \\ &\text{s. to} \quad y(t) + e(t) = x_1[y(t-1) + e(t-1)] + \cdots + x_n[y(t-n) + e(t-n)]. \end{aligned} \tag{6.4}$$

Due to the Hankel structure of the data matrix, it is not enough to only update the data covariance, as in the RLS case. This challenge will be addressed by wisely weighting the data covariance at each recursion step. After recursively weighting the covariance matrix, we will update per time t the entries of $\mathbf{x}$ one by one. This subproblem can be solved using the scalar Lasso which becomes available in closed form. Further, incorporating outliers in the measurement model, along the lines of Chapter 2, will lead

to a robust RS-TLS mode estimator. Finally, we will generalize the model to include also probing signals as an input to the system. This will change the linear prediction model in (6.3) to a system identification setup with perturbations present in both input and output.

References

[1] A. Abur and A. Gomez-Exposito, *Power System State Estimation: Theory and Implementation*, New York, NY: Marcel Dekker, 2004.

[2] I. G. Akrotirianakis and C. A. Floudas, "Computational experience with a new class of convex underestimators: box-constrained NLP problems," *Journal of Global Optimization*, vol. 29, no. 3, pp. 249–264, Jul. 2004.

[3] D. Angelosante, E. Grossi, G. B. Giannakis, and M. Lops, "Sparsity-aware estimation of CDMA system parameters," *EURASIP Journal on Advances in Signal Processing*, vol. 2010, Jun. 2010, Article ID 417981.

[4] X. Bai, H. Wei, K. Fujisawa, and Y. Wang, "Semidefinite programming for optimal power flow problems," *Intl. Journal of Electrical Power & Energy Systems,* vol. 30, no. 6–7, pp. 383–392, July 2008.

[5] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Processing Magazine*, vol. 24, no. 4, pp. 118–121, Jul. 2007.

[6] J. A. Bazerque and G. B. Giannakis, "Distributed spectrum sensing for cognitive radio networks by exploiting sparsity," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, pp. 1847–1862, Mar. 2010.

[7] A. Beck, A. Ben-Tal, and M. Teboulle, "Finding a global optimal solution for a quadratically constrained fractional quadratic problem with applications to the regularized total least squares," *SIAM J. on Matrix Anal. and Appl.*, vol. 28, no. 2, pp. 425–445, 2006.

[8] A. R. Bergen and V. Vittal, *Power System Analysis*, 2nd ed., Upper Saddle River, N.J. : Prentice Hall, 2000.

[9] C. R. Berger, S. Zhou, J. C. Preisig, and P. Willett, "Sparse channel estimation for multicarrier underwater acoustic communication: From subspace methods to compressed sensing," *IEEE Trans. on Signal Process.*, vol. 58, no. 3, pp. 1708–1721, Mar. 2010.

[10] L. Berman and A. Feuer, "Robust patterns in recurrent sampling of multiband signals," *IEEE Trans. on Signal Processing*, vol. 56, no. 6, pp. 2326–2333, June 2008.

[11] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Belmont, MA: Athena Scientific, 1999.

[12] N. D. Blakeley, P. J. Bones, R. P. Millane, and P. Renaud, "Efficient frequency-domain sample selection for recovering limited-support images," *J. Opt. Soc. Am. A* vol. 20, no. 1, pp. 67–77, Jan. 2003.

[13] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge, UK: Cambridge University Press, 2004.

[14] J. Brodie, I. Daubechies, C. De Mol, D. Giannone, and I. Loris, "Sparse and stable Markowitz portfolios," *Proc. of National Academy of Sciences,* vol. 106, no. 30, pp. 12267–12272, 2009.

[15] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," *Comptes Rendus Mathematique*, vol. 346, no. 9–10, pp. 589–592, 2008.

[16] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: universal encoding strategies?" *IEEE Trans on Info. Theory*, vol. 52, pp. 5406–5425, Dec. 2006.

[17] E. Candès and M. B. Wakin, "An Introduction to Compressive Sampling," *IEEE Sign. Proc. Mag.*, vol. 25, pp. 21–30, Mar. 2008.

[18] M. Cetin and B. M. Sadler, "Semi-Blind Sparse Channel Estimation with Constant Modulus Symbols," *Proc. of Intl. Conf. on Acoust., Speech, and Signal Proc.*, Philadelphia, PA, March 2005.

[19] V. Cevher, M. F. Duarte, and R. G. Baraniuk, "Distributed target localization via spatial sparsity," in *Proc. of 16th European Signal Processing Conf.*, Lausanne, Switzerland, Aug. 25–29, 2008.

[20] D. H. Chae, P. Sadeghi, and R. A. Kennedy, "Effects of basis-mismatch in compressive sampling of continuous sinusoidal signals," in *Proc. of 2nd Intl. Conf. on Future Computer and Communication*, Wuhan, China, May 21–24, 2010.

[21] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. on Scientific Computing*, vol. 20, pp. 33–61, Jan. 1998.

[22] Y. Chi, A. Pezeshki, L. Scharf, and R. Calderbank, "Sensitivity to basis mismatch in compressed sensing," in *Proc. of International Conf. on Acoustics, Speech, and Signal Processing*, Dallas, TX, Mar. 14–19, 2010.

[23] S. F. Cotter and B. D. Rao, "Application of total least squares (TLS) to the design of sparse signal representation dictionaries," in *Proc. of Asilomar Conf. on Signals, Systems and Computers*, Pacific Grove, CA, Nov. 3–6, 2002.

[24] T. Cui and C. Tellambura, "An Efficient Generalized Sphere Decoder for Rank-Deficient MIMO Systems," *IEEE Communications Letters*, vol. 9, pp. 423–425, May 2005.

[25] E. Dall'Anese, J. A. Bazerque, H. Zhu, and G. B. Giannakis, "Group sparse total least-squares for cognitive spectrum sensing," *Proc. Wrkshp. Signal Proc. Advances in Wireless Comms. (SPAWC)*, pp. 96-100, San Francisco, CA, Jun. 26-29, 2011.

[26] M. O. Damen, H. E. Gamal, and G. Caire, "On Maximum-Likelihood Detection and the Search for the Closest Lattice Point," *IEEE Trans. on Information Theory*, vol. 49, pp. 2389–2402, Oct. 2003.

[27] A. Das and D. Kempe, "Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection," in *Proc. 28th Intl. Conference on Machine Learning*, Bellevue, WA, Jun. 29–Jul. 2, 2011.

[28] E. de Klerk, "Exploiting special structure in semidefinite programming: A survey of theory and applications," *European Journal of Operational Research*, vol. 201, no. 1, pp. 1–10, 2010.

[29] W. Dinkelbach, "On nonlinear fractional programming," *Management Science*, vol. 13, no. 7, pp. 492–498, Mar. 1967.

[30] D. L. Donoho, "Compressed sensing," *IEEE Trans on Info. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.

[31] R. Emami and A. Abur, "Tracking changes in the external network model," in *Proc. 42nd North American Power Symp.*, Sep. 26–28, 2010.

[32] J. Fan and R. Li, "Variable selection via nonconcave penalized likelihood and its oracle properties," *Journal of the American Statistical Association*, vol. 96, no. 456, pp. 1348–1360, Dec. 2001.

[33] A. K. Fletcher, S. Rangan, and V. K. Goyal, "A Sparsity Detection Framework for On-Off Random Access Channels," *IEEE Proc. of Intl. Symp. on Info. Theory*, pp. 169–173, Seoul, South Korea, June 28 - July 3, 2009.

[34] J. Friedman, T. Hastie, and R. Tibshirani, "Regularization paths for generalized linear models via coordinate descent," *Journal of Statistical Software*, vol. 33, no. 1, pp. 1–22, Feb. 2010.

[35] J.-J. Fuchs, "Multipath time-delay detection and estimation," *IEEE Trans. on Signal Process.*, vol. 47, no. 1, pp. 237–243, Jan. 1999.

[36] ——, "On the application of the global matched filter to DOA estimation with uniform circular arrays," *IEEE Trans. on Signal Process.*, vol. 49, no. 4, pp. 702–709, Apr. 2001.

[37] J. Gertler, *Fault Detection and Diagnosis in Engineering Systems*, New York: Marcel Dekker, 1998.

[38] G. B. Giannakis, N. Gatsis, V. Kekatos, S.-J. Kim, H. Zhu, and B. Wollenberg, "Signal processing for smart monitoring and optimization of power systems," *IEEE Signal Processing Magazine (Feature Article),* 2012 (accepted).

[39] G. B. Giannakis, Z. Liu, X. Ma, and S. Zhou, *Space-Time Coding for Broadband Wireless Communications,* John Wiley & Sons, Inc., Jan. 2007.

[40] G. B. Giannakis, G. Mateos, S. Farahmand, and H. Zhu, "USPACOR: Universal sparsity-controlling outlier rejection," *Proc. of Intl. Conf. on Acoustics, Speech and Signal Processing,* Prague, Czech Republic, May 22–27, 2011.

[41] M. X. Goemans and D. P. Williamson, "Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming," *Journal of the ACM*, vol. 42, no. 6, pp. 1115–1145, 1995.

[42] A. Gomez-Exposito, A. Abur, A. de la Villa Jaen, and C. Gomez-Quiles, "A multilevel state estimation paradigm for smart grids," *Proceedings of the IEEE*, vol. 99, no. 6, pp. 952–976, June 2011.

[43] A. Gomez-Exposito, A. Abur, P. Rousseaux, A. de la Villa Jaen, and C. Gomez-Quiles, "On the use of PMUs in power system state estimation," in *Proc. 17th Power Systems Computation Conference*, Stockholm, Sweden, Aug. 22-26, 2011.

[44] G. Goodwin, M. Seron, J.D. Doná, *Constrained Control and Estimation: An Optimisation Approach*, Berlin: Springer, 2004.

[45] D. Gorinevsky, S. Boyd, and S. Poll, "Estimation of faults in DC electrical power system," in *Proc. American Control Conf.*, Jun. 10–12, 2009.

[46] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming," version 1.21, Apr. 2011. [Online]. Available: `http://cvxr.com/cvx/`

[47] T. Güler, G. Gross, and M. Liu, "Generalized line outage distribution factors," *IEEE Trans. on Power Systems*, vol. 22, no. 2, pp. 879–881, May 2007.

[48] J. Guo, Y. Fu, Z. Li, and M. Shahidehpour, "Direct calculation of line outage distribution factors," *IEEE Trans. on Power Systems*, vol. 24, no. 3, pp. 1633–1634, Aug. 2009.

[49] M. H. Hansen and B. Yu, "Model selection and the principle of minimum description length," *Journal of the American Statistical Association*, vol. 96, no. 454, pp. 746-774, Jun. 2001.

[50] P. C. Hansen, "Regularization tools: A Matlab package for analysis and solution of discrete ill-posed problems," *Numerical Algorithms*, vol. 6, no. 1, pp. 1–35, Mar. 1994.

[51] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference, and prediction*, 2nd ed. New York: Springer, 2009.

[52] M. He and J. Zhang, "A dependency graph approach for fault detection and localization towards secure smart grid," *IEEE Trans. on Smart Grid*, vol. 2, no. 2, pp. 342–351, June 2011.

[53] M. A. Herman and T. Strohmer, "General deviants: an analysis of perturbations in compressive sensing," *IEEE J. of Selected Topics in Signal Processing*, vol. 4, pp. 342–349, Apr. 2010.

[54] R. A. Jabr, "Exploiting sparsity in SDP relaxations of the OPF problem," *IEEE Trans. on Power Systems*, 2012 (to appear).

[55] M. Jansson, A. L. Swindlehurst, and B. Ottersten, "Weighted subspace fitting for general array error models," *IEEE Transactions on Signal Processing*, vol. 46, no. 9, pp. 2484–2498, Sep. 1998.

[56] R. J. Kaye and F. F. Wu, "Analysis of linearized decoupled power flow approximations for steady-state security assessment," *IEEE Trans. on Circuits and Systems*, vol. 31, no. 7, pp. 623–636, Jul. 1984.

[57] V. Kekatos and G. B. Giannakis, "Distributed robust power system state estimation," *IEEE Transactions on Power Systems*, 2012 (submitted).

[58] O. Kosut, L. Jia, J. Thomas, and L. Tong, "Malicious data attacks on the smart grid," *IEEE Trans. Smart Grid*, vol. 2, no. 4, pp. 645–658, Dec. 2011.

[59] P. Kundur, *Power System Stability and Control*, New York: McGraw-Hill, 1994.

[60] A. Y.S. Lam, B. Zhang, and D. Tse, "Distributed algorithms for optimal power flow problem," Sep. 2011 (submitted). [Online]. Available: `http://arxiv.org/abs/1109.5229`

[61] J. Lavaei and S. H. Low, "Zero duality gap in optimal power flow problem," *IEEE Trans. on Power Systems*, vol. 27, no. 1, pp. 92–107, Feb. 2012.

[62] Y. Liu, M. K. Reiter, and P. Ning, "False data injection attacks against state estimation in electric power grids," in *Proc. ACM Conf. on Computer and Comm. Security*, Chicago, IL, Nov. 2009, pp. 21 –32.

[63] J. Löfberg, "Dualize it: Software for automatic primal and dual conversions of conic programs," *Optimization Methods & Software*, vol. 24, no. 3, pp. 313–325, Jun. 2009.

[64] Z.-Q. Luo, W.-K. Ma, A. M.-C. So, Y. Ye, and S. Zhang, "Semidefinite relaxation of quadratic optimization problems," *IEEE Signal Processing Magazine*, vol. 27, no. 3, pp. 20–34, May 2010.

[65] D. M. Malioutov, M. Çetin, and A. S. Willsky, "A sparse signal reconstruction perspective for source localization with sensor arrays," *IEEE Trans. on Signal Process.*, vol. 53, pp. 3010–3022, Aug. 2005.

[66] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Proc.*, vol. 41, pp. 3397–3415, Dec. 1993.

[67] I. Markovsky and S. Van Huffel, "Overview of total least-squares methods," *Signal Processing*, vol. 87, no. 10, pp. 2283–2302, Oct. 2007.

[68] A. Monticelli, "Electric power system state estimation," *P*roceedings of the IEEE, vol. 88, no. 2, pp. 262–282, Feb. 2000.

[69] O. Nestares, D. J. Fleet, and D. J. Heeger, "Likelihood functions and confidence bounds for total-least-squares problems," in *Proc. of Conf. on Computer Vision and Pattern Recognition*, Hilton Head Island, SC, Jun. 13–15, 2000.

[70] N. Ö. Önhon and M. Çetin, "A nonquadratic regularization-based technique for joint SAR imaging and model error correction," in *Proc. of SPIE*, vol. 7337, 2009.

[71] A. Owrang, M. Viberg, M. Nosratinia, and M. Rashidi, " A new method to compute optimal periodic sampling patterns," *Proc. of IEEE 14th Digital Signal Processing (DSP) Workshop*, Sedona, AZ, Jan. 4–7, 2011.

[72] P. Pad, F. Marvasti, K. Alishahi, and S. Akbari, "A class of errorless codes for overloaded synchronous wireless and optical CDMA systems," *IEEE Transactions on Information Theory*, vol. 55, no. 6, pp. 2705-2715, June 2009.

[73] R. R. Picard and R. D. Cook, "Cross-validation of regression models," *Journal of the American Statistical Association*, vol. 79, no. 387, pp. 575–583, Sep. 1984.

[74] Power Systems Test Case Archive, *University of Washington*. [Online.] Available: `http://www.ee.washington.edu/research/pstca/`

[75] A. H. Sayed, *Fundamentals of Adaptive Filtering,* John Wiley & Sons, 2003.

[76] A. Schellenberg, W. Rosehart, and J. Aguado, "Cumulant-based probabilistic optimal power flow (P-OPF) with Gaussian and Gamma distributions," *IEEE Trans. on Power Systems*, vol. 20, pp. 773–781, May 2005.

[77] D. M. Sima, S. Van Huffel, and G. H. Golub, "Regularized total least squares based on quadratic eigenvalue problem solvers," *BIT Numerical Mathematics*, vol. 44, no. 4, pp. 793–812, Dec. 2004.

[78] B. Stott, J. Jardim, and O. Alsac, "DC power flow revisited," *IEEE Trans. on Power Systems*, vol. 24, no. 3, pp. 1290–1300, Aug. 2009.

[79] J. F. Sturm, "Using SeDuMi 1.02, a matlab toolbox for optimization over symmetric cones," *Optimization Methods Software*, vol. 11–12, pp. 625–653, Aug. 1999. [Online]. Available: `http://sedumi.mcmaster.ca`

[80] J. E. Tate and T. J. Overbye, "Line outage detection using phasor angle measurements," *IEEE Trans. on Power Systems*, vol. 23, no. 4, pp. 1644–1652, Nov. 2008.

[81] J. E. Tate and T. J. Overbye, "Double line outage detection using phasor angle measurements," in *Proc. IEEE PES General Meeting*, July 26–30, 2009.

[82] Z. Tian, G. Leus, and V. Lottici, "Detection of Sparse Signals under Finite-Alphabet Constraints," *Proc. of Intl. Conf. on Acoust., Speech and Signal Proc.*, pp. 2349–2352, Taipei, April 2009.

[83] R. Tibshirani, "Regression shrinkage and selection via the Lasso," *Journal of Royal. Statist. Soc. B*, vol. 58, pp. 267–288, 1996.

[84] M. E. Tipping, "Sparse Bayesian Learning and the Relevance Vector Machine," *Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.

[85] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. on Info. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.

[86] P. Tseng, "Convergence of a block coordinate descent method for nondifferentiable minimization," *Journal of Optimization Theory and Applications*, vol. 109, no. 3, pp. 475–494, Jun. 2001.

[87] U.S.-Canada Power System Outage Task Force, "Final report on the August 14th blackout in the United States and Canada," Apr. 2004. [Online.] Available: `https://reports.energy.gov/`

[88] U.S.-Canada Power System Outage Task Force, "Final Report on the implementation of the task force recommendation," Sep. 2006. [Online.] Available: `http://nrcan.gc.ca/eneene/pdf/outpan-eng.pdf`

[89] U.S. Department of Energy, *The Smart Grid: An introduction*, 2008. [Online.] Available: `http://www.oe.energy.gov/SmartGridIntroduction.htm`

[90] S. Van Huffel and J. Vandewalle, *The Total Least Squares Problem: Computational Aspects and Analysis*, ser. Frontier in Applied Mathematics. Philadelphia: SIAM, 1991, vol. 9.

[91] R. Venkataramani and Y. Bresler, "Perfect reconstruction formulas and bounds on aliasing error in sub-Nyquist nonuniform sampling of multiband signals," *IEEE Trans. on Information Theory*, vol. 46, no. 6, pp. 2173–2183, Sep. 2000.

[92] R. Venkataramani and Y. Bresler, "Optimal sub-Nyquist nonuniform sampling and reconstruction for multiband signals," *IEEE Trans. on Signal Proc.*, vol. 49, no. 10, pp. 2301–2313, Oct. 2001.

[93] S. Verdu, *Multiuser Detection,* Cambridge, U.K., Cambridge Univ. Press, 1998.

[94] E. Viterbo and J. Boutros, "A Universal Lattice Code Decoder for Fading Channels," *IEEE Trans. on Inform. Theory* , vol. 45, pp. 1639–1642, July 1999.

[95] D. B. West, *Introduction to Graph Theory*, 2nd ed., Prentice Hall, 2001.

[96] A. J. Wood and B. F. Wollenberg, *Power Generation, Operation, and Control*, 2nd ed., New York, NY: John Wiley and Sons, 1996.

[97] W. A. Wulf, "Great achievements and grand challenges," *The Bridge*, vol. 30, no. 3/4, pp. 5–10, Fall 2010. [Online.] Available: `http://www.greatachievements.org/`

[98] W. Xu, M. Wang, and A. Tang, "Sparse recovery from nonlinear measurements with applications in bad data detection for power networks," Dec. 2011 (submitted). [Online]. Available: `http://arxiv.org/abs/1112.6234`

[99] M. Yuan and Y. Lin, "Model Selection and Estimation in Regression with Grouped Variables," *Journal of the Royal Statistical Soceity: Series B*, vol. 68, pp. 49–67, 2006.

[100] M. Zhou, V. A. Centeno, J. S. Thorp, and A. G. Phadke, "An alternative for including phasor measurements in state estimators," *IEEE Trans. on Power Systems,* vol. 21, no. 4, pp. 1930–1937, Nov. 2006.

[101] N. Zhou, J. W. Pierre, D. J. Trudnowski, R. T. Guttromson, , "Robust RLS methods for online estimation of power system electromechanical modes," *IEEE Trans. on Power Systems*, vol. 22, no. 3, pp. 1240–1249, Aug. 2007.

[102] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas, "Matpower: Steady-state operations, planning and analysis tools for power systems research and education," *IEEE Trans. on Power Systems*, vol. 26, no. 1, pp. 12–19, Feb. 2011.

[103] H. Zhu, "Signal processing algorithms for smart power system monitoring," *First Women's Workshop on Comms. and Signal Processing*, BIRS, Banff, Canada, Jul. 13-15, 2012.

[104] H. Zhu, A. Cano, and G. B. Giannakis, "Consensus-based distributed MIMO decoding using semidefinite relaxation," in *Proc. of 2nd Intl. Workshop on Comp. Advances in Multi-Sensor Adapt. Proc.*, St. Thomas, U.S. Virgin Islands, Dec. 12-14, 2007.

[105] H. Zhu and G. B. Giannakis, "Multi-area state estimation using distributed semidefinite programming for nonlinear power systems," *Proc. 3rd IEEE Intl Conf. on Smart Grid Communications (SmartGridComm)*, Tainan, Taiwan, Nov. 5-8, 2012 (invited).

[106] H. Zhu and G. B. Giannakis, "Robust state estimation for nonlinear power flow models," *Proc. 43rd North America Power Symposium (NAPS)*, UIUC, IL, Sep. 9-11, 2012.

[107] H. Zhu and G. B. Giannakis, "Nonlinear AC power system state estimation using semidefinite relaxation," *IEEE Transactions on Power Systems,* February 2012 (submitted).

[108] H. Zhu and G. B. Giannakis, "Sparse overcomplete representations for efficient identification of power line outages," *IEEE Transactions on Power Systems,* 2012 (accepted).

[109] H. Zhu and G. B. Giannakis, "Lassoing line outages in the smart power grid," in *Proc. of 2nd Intl. Conf. on Smart Grid Communications*, Brussels, Belgium, Oct. 17-20, 2011.

[110] H. Zhu and G. B. Giannakis, "Estimating the state of AC power systems using semidefinite programming," *Proc. of 43rd NAPS*, pp. 1-7, Boston, MA, Aug. 4-6, 2011.

[111] H. Zhu and G. B. Giannakis, "Exploiting sparse user activity in multiuser detection," *IEEE Transactions on Communications*, vol. 59, no. 2, pp. 454-465, February 2011.

[112] H. Zhu and G. B. Giannakis, "Sparsity-embracing multiuser detection for CDMA systems with low activity factor," *Proc. of IEEE Intl. Symp. on Info. Theory*, pp. 164–168, Seoul, South Korea, June 28 - July 3, 2009.

[113] H. Zhu, G. B. Giannakis, and A. Cano, "Distributed in-network channel decoding," *IEEE Trans. on Signal Proc.*, vol. 57, no. 10, pp. 3970-3983, October 2009.

[114] H. Zhu, G. B. Giannakis, and G. Leus, "Weighted and structured sparse total least-squares for perturbed compressive sampling," in *Proc. of Intl. Conf. on Acoust., Speech, and Sig. Proc.*, Prague, Czech Republic, May 22–27, 2011.

[115] H. Zhu, G. Leus, and G. B. Giannakis, "Sparsity-cognizant total least-squares for perturbed compressive sampling," *IEEE Transactions on Signal Processing*, vol. 59, no. 5, pp. 2002-2016, May 2011.

[116] H. Zhu, G. Leus, and G. B. Giannakis, "Sparse regularized total least-squares for sensing applications," in *Proc. of Wrkshp. on Signal Processing Advances in Wireless Communications*, Marrakech, Marocco, June 20 – 23, 2010.

[117] H. Zhu, G. Mateos, G. B. Giannakis, N. D. Sidiropoulos, and A. Banerjee, "Sparsity-aware overlapping co-clustering: Identifiability and fast algorithms," *IEEE Transactions on Signal Processing,* August 2012 (submitted).

[118] H. Zhu, G. Mateos, G. B. Giannakis, N. D. Sidiropoulos, and A. Banerjee, "Sparsity-cognizant overlapping co-clustering for behavior inference in social networks," in *Proc. of Intl. Conf. on Acoust., Speech, and Signal Proc.*, Dallas, Texas, March 14–19, 2010.

[119] H. Zou, "The adaptive Lasso and its oracle properties," *J. of the American Stat. Assoc.*, vol. 101, no. 476, pp. 1418–1429, Dec. 2006.

[120] A. Zymnis, S. Boyd, and D. Gorinevsky, "Relaxed maximum a posteriori fault identification," *Signal Processing*, vol. 89, no. 6, pp. 989–999, June 2009.

Appendix A

Acronyms

Care has been taken in this Thesis to minimize the use of acronyms, but this cannot always be achieved. This appendix contains a table of acronyms and their meanings, which are ordered alphabetically.

Table A.1: Acronyms

Acronym	Meaning
AC	alternating current
AP	access point
AWGN	additive white Gaussian noise
BB	branch-and-bound
BP	basis pursuit
CD	coordinate descent
CDMA	code division multiple access
CR	cognitive radio
CS	compressive sampling
CV	cross validation
DC	direct current
DDD	decision-directed detector
DoA	directions-of-arrival

Continued on next page

Table A.1 – continued from previous page

Acronym	Meaning
EIV	errors-in-variables
EMS	energy management system
ES	exhaustive search
FA	finite alphabet
FFT	fast Fourier transform
FVP	flat-voltage profile
GMRF	Gauss-Markov random field
i.i.d.	independent identically distributed
ISI	inter-symbol interference
LARS	least-angle regression
Lasso	least-absolute shrinkage and selection operator
LAV	least-absolute value
LD	Lasso detector
LNR	largest normalized residual
LODF	line outage distribution factor
LS	least-squares
MAP	maximum *a posteriori*
MDL	minimum description length
MIMO	multiple-input multiple-output
ML	maximum likelihood
MMSE	minimum mean-square error
MP	matching pursuit
MUD	multiuser detection
NERC	North-American Electric Reliability Corporation
OFDM	orthogonal frequency-division multiplexing
OMP	orthogonal matching pursuit
OPF	optimal power flow
PAM	pulse amplitude modulation
pdf	probability density function

Continued on next page

Table A.1 – continued from previous page

Acronym	Meaning
PMU	phasor measurement unit
PSD	power spectral density
PSK	phase-shift keying
QAM	quadrature amplitude modulation
QP	quadratic program
RD	Ridge detector
RF	radio frequency
RLS	recursive least-squares
ROC	receiver operating characteristic
R-SE	robust state estimation
RS-TLS	recursive sparse total least-squares
r.v.	random variable
SD	sphere decoding
SDP	semidefinite program
SDR	semidefinite relaxation
SDX	system data exchange
SE	state estimation
SE	Schnorr-Euchner
SER	symbol error rate
S-MAP	sparsity-exploiting maximum *a posteriori*
SNR	signal-to-noise ratio
SSD	sparse sphere decoder
S-TLS	sparse total least-squares
SVD	singular value decomposition
TLS	total least-squares
UAV	unmanned aerial vehicle
ULA	uniformly-spaced linear array
w.l.o.g.	without loss of generality
WLS	weighted least-squares

Continued on next page

Table A.1 – continued from previous page

Acronym	Meaning
w.p.	with probability
wrt	with respect to
WSN	wireless sensor network
WSS-TLS	weighted and structured sparse total least-squares

Appendix B

Technical Details in Chapter 2

B.1 Proof of Lemma 2.1

To establish (2.8a), use the injected power flow equation, and successively Kirchoff's law and (2.6a) to obtain $P_n + jQ_n = V_n I_n^{\mathcal{H}} = (V_n^{\mathcal{H}} I_n)^{\mathcal{H}} = (\mathbf{v}^{\mathcal{H}}\mathbf{e}_n\mathbf{e}_n^T\mathbf{i})^{\mathcal{H}} = (\mathbf{v}^{\mathcal{H}}\mathbf{Y}_n\mathbf{v})^{\mathcal{H}} = \mathbf{v}^{\mathcal{H}}\mathbf{Y}_n^{\mathcal{H}}\mathbf{v} = \mathrm{Tr}(\mathbf{Y}_n^{\mathcal{H}}\mathbf{V})$. Hence, P_n and Q_n are related to $\mathbf{V}$ using the real and imaginary parts of $\mathbf{Y}_n^{\mathcal{H}}$, respectively, as asserted in (2.8a) [cf. (2.7a)]. Likewise, (2.8b) follows after replacing with the line power flow equation and (2.7b), while $|V_n|^2$ is naturally the (n,n)-th entry of $\mathbf{V}$ as in (2.8c).

B.2 Proof of Proposition 2.2

First, notice that the outer-product $\hat{\mathbf{v}}\hat{\mathbf{v}}^{\mathcal{H}}$ is feasible for (2.10). Expressing the rank-1 matrix as $\hat{\mathbf{V}}_1 = \hat{\mathbf{v}}_1\hat{\mathbf{v}}_1^{\mathcal{H}}$ renders the complex $\hat{\mathbf{v}}_1$ also feasible for (2.5). As Lemma 2.1 establishes the equivalence between the costs in (2.5) and (2.10a), their solutions can be related henceforth.

For the equivalence between (2.10) and (2.11), Schur's complement lemma for (2.11c) ensures that $\chi_\ell \geq [z_\ell - \mathrm{Tr}(\mathbf{H}_\ell\mathbf{V})]^2\ \forall \ell$. Since (2.11a) minimizes a positively weighted sum of $\{\chi_\ell\}$, the equality is further guaranteed at the optimum; i.e., $\hat{\chi}_{2,\ell} = \left[z_\ell - \mathrm{Tr}(\mathbf{H}_\ell\hat{\mathbf{V}}_2)\right]^2$, $\forall \ell$. Substituting this back to (2.11a) shows that $\hat{\mathbf{V}}_1$ and $\hat{\mathbf{V}}_2$ achieve the same minimum cost, which completes the proof.

B.3 Proof of Proposition 2.3

Given how $\mathrm{Tr}(\mathbf{AV})$ appears in the Lagrangian $\mathcal{L}$, it follows that $\mathbf{V}$ is the Lagrange multiplier for the dual constraint $\mathbf{A} \preceq \mathbf{0}$ in (2.17c). Strong duality between these two convex problems holds because the primal R-SE problem admits a strictly feasible $\mathbf{V} = \mathbf{I}$ and $\chi_\ell = (z_\ell - \mathrm{Tr}(\mathbf{H}_\ell))^2 + 1 \ \forall \ell$. To show that (2.17) is also the dual of the original problem in (2.11), recall that (2.11b) is equivalent to having $\mathbf{V} = \mathbf{v}\mathbf{v}^{\mathcal{H}}$ for some $\mathbf{v}$. Therefore, the corresponding Lagrangian of the R-SE problem (2.11) can be obtained by substituting $\mathbf{V}$'s outer-product form into (2.16). Interestingly, the minimum of this new Lagrangian over $\boldsymbol{\chi}$ and $\mathbf{v}$ is no different from the one in (2.16), and thus the same dual problem in (2.17) follows.

B.4 PMU-aided SDR-based SE

Chapter 2 mostly focuses on the robust formulation for SE by controlling the sparsity of outliers, but with all results based on SDR readily applicable to the non-robust SE problem. This section will show how to generalize this novel SDR based SE to include linear PMU measurements.

Recent deployment of PMUs suggests complementing the nonlinear measurements collected by legacy meters to perform SE. Compared to legacy measurements, PMUs provide synchronous data that are *linear* functions of the state $\mathbf{v}$. If bus m is equipped with a PMU, then its voltage phasor V_m and related current phasors $\{I_{mn}\}_{n\in\mathcal{N}_m}$ are available to the control center with high accuracy. Hence, with adequate number of PMUs and wisely chosen placement buses, SE using only PMU data boils down to estimating a linear regression coefficient vector for which a batch WLS solution is available in closed form. However, installation and networking costs involved allow only for limited penetration of PMUs in the near future. This means that SE must be performed using *jointly* legacy meters and PMU measurements.

To this end, let $\boldsymbol{\zeta}_m = \boldsymbol{\Phi}_m\mathbf{v} + \boldsymbol{\varepsilon}_m$ collect the noisy PMU data at bus m, where the linear regression matrix $\boldsymbol{\Phi}_m$ is constructed in accordance with the bus index m and line admittances, while $\boldsymbol{\varepsilon}_m$ denotes the PMU measurement noise, assumed to be complex zero-mean Gaussian distributed with covariance $2\check{\sigma}_m^2\mathbf{I}$, independent across buses and from the legacy meter noise terms $\{\epsilon_\ell\}$. The SE task now amounts to estimating $\mathbf{v}$

given both $\mathbf{z}$ and $\{\boldsymbol{\zeta}_m\}_{m\in\mathcal{P}}$, where $\mathcal{P}\subseteq\mathcal{N}$ denotes the PMU-instrumented set of buses. Hence, the ML-optimal WLS cost in (2.3) must be augmented with the log-likelihood induced by PMU data, as

$$\hat{\mathbf{v}} := \arg\min_{\mathbf{v}} \sum_{\ell=1}^{L} w_\ell[z_\ell - h_\ell(\mathbf{v})]^2 + \sum_{m\in\mathcal{P}} \omega_m \|\boldsymbol{\zeta}_m - \boldsymbol{\Phi}_m\mathbf{v}\|_2^2 \tag{B.1}$$

where $\omega_m := 1/\check{\sigma}_m^2\ \forall m\in\mathcal{P}$. The augmented SE problem (B.1) is still nonconvex due to the quadratic dependence of legacy measurements in the wanted state $\mathbf{v}$.

Existing SE methods that account for PMU measurements can be categorized in two groups. The first one includes the so-termed hybrid SE approaches which utilize both PMU and legacy measurements in a WLS solver via iterative linearization; see e.g., [43]. Depending on the number of PMUs, the state can be either expressed using polar coordinates (similar to traditional WLS-based SE), or by rectangular coordinates (as the notation $\mathbf{v}$ here). The polar representation is preferred when legacy measurements are abundant, because it requires minor adaptations of the existing WLS-based SE. On the other hand, the polar representation is less powerful when it comes to exploiting the linearity of PMU measurements. With full penetration of PMUs in the future, the rectangular representation is expected to grow in popularity, especially if full observability can be ensured by the sole use of PMU measurements.

An alternative approach to including PMU data is through sequential SE [100], which entails two steps. The WLS-based SE is performed first based only on legacy measurements. This state estimates serve the role of linear "pseudo-measurements" for the subsequent step together with PMU measurements. The post-processing involves linear models only, and is efficiently computable. Clearly, sequential SE requires no modifications of existing SE modules, but loses the optimality offered by joint estimation. More severely, if the traditional SE based only on legacy measurements fails to converge to a global optimum, the post-processing including PMU data is unlikely to improve estimation accuracy.

Since both of these options for including PMU data suffer from the nonconvexity present with legacy measurements, the SDR technique is again well motivated to convexify the augmented SE to

$$\left\{\hat{\mathbf{X}},\hat{\boldsymbol{\chi}}\right\} := \arg\min_{\mathbf{X},\boldsymbol{\chi}} \mathbf{w}^T\boldsymbol{\chi} + \sum_{m\in\mathcal{P}} \omega_m \left[\mathrm{Tr}(\boldsymbol{\Phi}_m^{\mathcal{H}}\boldsymbol{\Phi}_m\mathbf{V}) - 2\mathrm{Re}(\boldsymbol{\zeta}_m^{\mathcal{H}}\boldsymbol{\Phi}_m\mathbf{v})\right] \tag{B.2a}$$

$$\text{s.to } \mathbf{X} = \begin{bmatrix} \mathbf{V} & \mathbf{v} \\ \mathbf{v}^{\mathcal{H}} & 1 \end{bmatrix} \succeq \mathbf{0}, \tag{B.2b}$$

$$\begin{bmatrix} -\chi_\ell & z_\ell - \mathrm{Tr}(\mathbf{H}_\ell \mathbf{V}) \\ z_\ell - \mathrm{Tr}(\mathbf{H}_\ell \mathbf{V}) & -1 \end{bmatrix} \preceq \mathbf{0} \quad \forall \ell. \tag{B.2c}$$

Similar to the SDR-based R-SE problem (2.13) with the additional constraint $\mathrm{rank}(\mathbf{X}) = 1$, the positive semi-definiteness of $\mathbf{X}$ can ensure $\mathbf{V} = \mathbf{v}\mathbf{v}^{\mathcal{H}}$. Substituting the latter into (B.2), and following the proof of Proposition 2.2 leads to the equivalence of rank-constrained (B.2) with the augmented WLS in (B.1). The SDP problem here also offers the advantages of (2.13), in terms of polynomial complexity and dual equivalence. From the solution $\hat{\mathbf{X}}$, either eigenvector approximation or randomization can be employed to generate vectors of length $N+1$. Using the first N entries of any such vector, a feasible state estimate can be formed by proper rescaling. With linear PMU measurements, the voltage angle ambiguity is no longer present, and the rescaling factor $\hat{c}^\star$ can be found by solving

$$\hat{c}^\star = \arg\min_{c \in \mathbb{C}} \sum_{\ell=1}^{L} w_\ell \left[z_\ell - |c|^2 \boldsymbol{\nu}^{\mathcal{H}} \mathbf{H}_\ell \boldsymbol{\nu} \right]^2 + \sum_{m \in \mathcal{P}} \omega_m \|\boldsymbol{\zeta}_m - c \boldsymbol{\Phi}_m \boldsymbol{\nu}\|_2^2. \tag{B.3}$$

This is a fourth-order polynomial minimization problem and numerically solvable. It is often the case that the PMUs have much higher accuracy compared to legacy meters, and their corresponding weights are much larger, which leads to $\omega_\ell \gg w_\ell$. Hence, it suffices to minimize only the dominant second summand in (B.3), and efficiently approximate the solution $\hat{c}^\star \approx (\sum_{m \in \mathcal{P}} \omega_m \boldsymbol{\zeta}_m^{\mathcal{H}} \boldsymbol{\Phi}_m \boldsymbol{\nu}) / (\sum_{m \in \mathcal{P}} \omega_m \|\boldsymbol{\Phi}_m \boldsymbol{\nu}\|_2^2)$. This approximation will be used in the numerical tests of the ensuing section.

Appendix C

Technical Details in Chapter 4

C.1 Proof of Lemma 4.1

To establish the equivalence of (4.5) with (4.3), simply eliminate $\mathbf{e}$ by substituting the constraint (4.3b) into the cost function of (4.3a). For (4.6), let $\mathbf{v} := \text{vec}([\mathbf{E}\ \mathbf{e}])$, and re-write the cost in (4.3a) as $\|[\mathbf{E}\ \mathbf{e}]\|_F^2 = \|\mathbf{v}\|_2^2$; and the constraint (4.3b) as $\mathbf{y} - \mathbf{A}\mathbf{x} = \mathbf{G}(\mathbf{x})\mathbf{v}$, where $\mathbf{G}(\mathbf{x}) := \mathbf{I} \otimes [\mathbf{x}^T, -1]$. With $\mathbf{x}$ fixed, the ℓ_1-norm can be dropped from (4.3a), and the reformulated optimization becomes: $\min_{\mathbf{v}}\ \|\mathbf{v}\|_2^2$ s. to $\mathbf{y} - \mathbf{A}\mathbf{x} = \mathbf{G}(\mathbf{x})\mathbf{v}$. But the latter is a minimum-norm LS problem, admitting the closed-form solution

$$\mathbf{v}(\mathbf{x}) = \mathbf{G}^T(\mathbf{x})[\mathbf{G}(\mathbf{x})\mathbf{G}^T(\mathbf{x})]^{-1}(\mathbf{y} - \mathbf{A}\mathbf{x}) = (1 + \|\mathbf{x}\|_2^2)^{-1}\mathbf{G}^T(\mathbf{x})(\mathbf{y} - \mathbf{A}\mathbf{x}) \tag{C.1}$$

where the second equality holds because $\mathbf{G}(\mathbf{x})\mathbf{G}^T(\mathbf{x}) = \|[\mathbf{x}^T, -1]\|_2^2\mathbf{I} = (1 + \|\mathbf{x}\|_2^2)\mathbf{I}$. Substituting (C.1) back into the cost $\|\mathbf{v}\|_2^2$, yields readily the fractional form in (4.6), which depends solely on $\mathbf{x}$.

C.2 Proof of Proposition 4.2

Given $\mathbf{y}$ and $\mathbf{A}$, the MAP approach to estimating both $\mathbf{x}_o$ and $\mathbf{A}_o$ in (4.4) amounts to maximizing with respect to (wrt) $\mathbf{x}$ and $\mathbf{E}$ the logarithm of the posterior pdf denoted as $\ln p[\mathbf{x}_o = \mathbf{x}, \mathbf{A}_o = \mathbf{A} + \mathbf{E}|\mathbf{y}, \mathbf{A}]$. Recalling that $\mathbf{x}_o$ and $\mathbf{A}_o$ are independent under (as4.1a), Bayes' rule implies that this is equivalent to: $\min_{\mathbf{x},\mathbf{E}} -\{\ln p[\mathbf{y}, \mathbf{A}|\mathbf{x}_o = \mathbf{x}, \mathbf{A}_o = \mathbf{A} + \mathbf{E}] + \ln p[\mathbf{x}_o = \mathbf{x}] + \ln p[\mathbf{A}_o = \mathbf{A} + \mathbf{E}]\}$, where the summands correspond to the

(conditional) log-likelihood and the log-prior pdfs, respectively. The log-prior associated with the Laplacian pdf of $\mathbf{x}_o$ is given by

$$\ln p[\mathbf{x}_o = \mathbf{x}] = \ln \prod_{\nu=1}^{n} [(\lambda/4) \exp(-\lambda |x_\nu|/2)] = -(\lambda/2) \sum_{\nu=1}^{n} |x_\nu| + n \ln(\lambda/4) \tag{C.2}$$

while the log-prior associated with the uniform pdf of $\mathbf{A}_o$ is constant under (as4.1a), and thus does not affect the MAP criterion. Conditioning the log-likelihood on $\mathbf{x}_o$ and $\mathbf{A}_o$, implies that the only sources of randomness in the data $[\mathbf{y}\ \mathbf{A}]$ are the EIV model perturbations, which under (as4.1) are independent, standardized Gaussian; thus, the conditional log-likelihood is $\ln p[\mathbf{y}, \mathbf{A} | \mathbf{x}_o = \mathbf{x}, \mathbf{A}_o = \mathbf{A} + \mathbf{E}] = \ln[\mathbf{e}_y = \mathbf{y} - (\mathbf{A} + \mathbf{E})\mathbf{x}] + \ln p[\mathbf{E}_A = \mathbf{E}]$. After omitting terms not dependent on the variables $\mathbf{x}$ and $\mathbf{E}$, the latter shows that the log-likelihood contributes to the MAP criterion two quadratic terms (sum of two Gaussian exponents): $(1/2)\{\|\mathbf{y} - (\mathbf{A} + \mathbf{E})\mathbf{x}\|_2^2 + \|\mathbf{E}\|\}_F^2\}$. Upon combining these quadratic terms with the ℓ_1-norm coming from the sum in (C.2), the log-posterior pdf boils down to the form minimized in (4.5), which per Lemma 4.1 is equivalent to (4.3), and thus establishes MAP optimality of S-TLS under (as4.1a).

Proceeding to prove optimality under (as4.1b), given again the data $\mathbf{y}$ and $\mathbf{A}$, consider the MAP approach now to estimate only $\mathbf{x}_o$ in (4.4), treating $\mathbf{A}_o$ as a nuisance parameter matrix that satisfies (as4.1b). MAP here amounts to maximizing (wrt $\mathbf{x}$ only) the criterion $\ln p[\mathbf{x}_o = \mathbf{x} | \mathbf{y}, \mathbf{A}]$; and Bayes' rule leads to the equivalent problem $\min_{\mathbf{x}} -\{\ln p[\mathbf{y}, \mathbf{A} | \mathbf{x}_o = \mathbf{x}] + \ln p[\mathbf{x}_o = \mathbf{x}]\}$. But conditioned on $\mathbf{x}_o$, (as4.1b) dictates that $\mathbf{A}_o$ and $[\mathbf{E}_A\ \mathbf{e}_y]$ are zero-mean Gaussian and independent. Thus, linearity of the EIV model (4.4) implies that $\mathbf{y}$ and $\mathbf{A}$ are zero-mean jointly Gaussian in the conditional log-likelihood. Since rows of $\mathbf{A}_o$ and $[\mathbf{E}_A\ \mathbf{e}_y]$ are (conditionally) i.i.d. under (as4.1b), the rows of matrix $[\mathbf{A}\ \mathbf{y}]$ are independent. In addition, the ρth-row of $[\mathbf{A}\ \mathbf{y}]$ denoted as $[\mathbf{a}_\rho^T\ y_\rho]$, has inverse (conditional) covariance matrix

$$\begin{aligned}
&\mathbb{E}\left[\begin{bmatrix} \mathbf{a}_\rho \\ y_\rho \end{bmatrix} [\mathbf{a}_\rho^T\ y_\rho] \,\middle|\, \mathbf{x}_o = \mathbf{x} \right]^{-1} \\
=&\begin{bmatrix} (\sigma^2+1)\mathbf{I} - \sigma^2 \mathbf{x}\mathbf{x}^T / \left(1 + \|\mathbf{x}\|_2^2\right) & \sigma^2 \mathbf{x} / \left(1 + \|\mathbf{x}\|_2^2\right) \\ \sigma^2 \mathbf{x}^T / \left(1 + \|\mathbf{x}\|_2^2\right) & 1 + \sigma^2 \|\mathbf{x}\|_2^2 / \left(1 + \|\mathbf{x}\|_2^2\right) \end{bmatrix}^{-1} \\
=&\frac{1}{\sigma^2+1} \left\{ \mathbf{I} + \frac{\sigma^2}{1 + \|\mathbf{x}\|_2^2} \begin{bmatrix} \mathbf{x} \\ -1 \end{bmatrix} [\mathbf{x}^T\ -1] \right\}
\end{aligned} \tag{C.3}$$

with determinant $1/(\sigma^2+1)^n$ not a function of $\mathbf{x}$. After omitting such terms not dependent on $\mathbf{x}$, and using the independence among rows and their inverse covariance in (C.3), the conditional log-likelihood boils down to the fractional form $\frac{\sigma^2}{2(\sigma^2+1)}\|\mathbf{y}-\mathbf{A}\mathbf{x}\|_2^2/\left(1+\|\mathbf{x}\|_2^2\right)$. Since the Laplacian parameter under (as4.1b) equals $2(\sigma^2+1)/(\lambda\sigma^2)$, the log-prior in (C.2) changes accordingly; and together with the fractional form of the log-likelihood reduces the negative log-posterior to the cost in (4.6). This establishes MAP optimality of the equivalent S-TLS in (4.3) for estimating only $\mathbf{x}_o$ in (4.4), under (as4.1b).

C.3 Proof of Proposition 4.3

Upon updating the lower and upper bounds, it holds per outer iteration $i \geq 1$ that $u_i - l_i \leq \frac{1}{2}(u_{i-1} - l_{i-1}) + \delta$; and by induction, $u_i - l_i \leq \left(\frac{1}{2}\right)^i u_0 + 2\delta$, when $l_0 = 0$. The latter implies that if the number of iterations $i \geq \left\lceil \ln\left(\frac{u_0}{\varepsilon - 2\delta}\right) / \ln(2) \right\rceil$, the distance $u_i - l_i \leq \varepsilon$ is satisfied.

Since per outer iteration Algorithm 4.1-a outputs $\mathbf{x}_g^\star \in \mathcal{X}_1(\mu)$, it holds that the updated $\mathbf{x}_\varepsilon^\star$ is also feasible. Further, the bisection process guarantees that $l_i \leq a^\star \leq f(\mathbf{x}_\varepsilon^\star) \leq u_i$ per iteration i. Since Algorithm 4.1-b ends with $u_i - l_i \leq \varepsilon$, the inequality in (4.13) follows readily.

C.4 Proof of Proposition 4.4

The argument relies on the basic convergence result in [86]. The alternating descent algorithm specified by (4.14) and (4.16) is a special case of the block coordinate descent method using the cyclic rule for minimizing the cost in (4.5). The first two summands of this cost are differentiable wrt the optimization variables, while the non-differential third term (ℓ_1-norm regularization) is separable in the entries of $\mathbf{x}$. Hence, the three summands satisfy the assumptions (B1)–(B3) and (C2) in [86]. Convergence of the iterates $\{\mathbf{E}(i), \mathbf{x}(i)\}$ to a coordinate minimum point of the cost thus follows by appealing to [86, Thm. 5.1]. Moreover, the first summand is Gâteaux-differentiable over its domain which is open. Hence, the cost in (4.5) is regular at each coordinate's minimum point, and every coordinate's minimum point becomes a stationary point; see [86, Lemma

3.1]. Monotonicity of the convergence follows simply because the cost per iteration may either reduce or maintain its value.

C.5 Proof of Lemma 4.5

Constraint (4.23b) can be solved uniquely for $\boldsymbol{\epsilon}^y$ to obtain $\boldsymbol{\epsilon}^y = (\boldsymbol{S}^y)^\dagger[\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x})\boldsymbol{\epsilon}^A - (\mathbf{y}-\mathbf{A}\mathbf{x})]$. Plug the latter with the definition of $\mathbf{r}(\mathbf{x})$ from (4.24) into the quadratic form in (4.23a) to recognize that (4.23) is equivalent to the unconstrained form in (4.25) with the $\boldsymbol{\epsilon}^y$ variable eliminated.

To arrive at (4.26), suppose that $\mathbf{x}$ is given and view the compact form of (4.23) (after ignoring $\lambda\|\mathbf{x}\|_1$) as the following weighted minimum-norm LS problem: $\min_{\mathbf{x},\boldsymbol{\epsilon}}\{\boldsymbol{\epsilon}^T\mathbf{W}\boldsymbol{\epsilon}+\lambda\|\mathbf{x}\|_1\}$ s.to $\boldsymbol{G}(\mathbf{x})\boldsymbol{\epsilon}=\mathbf{r}(\mathbf{x})$. Solving the latter in closed form expresses $\boldsymbol{\epsilon}$ in terms of $\mathbf{x}$ as: $\boldsymbol{\epsilon}(\mathbf{x}) = \mathbf{W}^{-1}\boldsymbol{G}^T(\mathbf{x})\left[\boldsymbol{G}(\mathbf{x})\mathbf{W}^{-1}\boldsymbol{G}^T(\mathbf{x})\right]^\dagger\mathbf{r}(\mathbf{x})$. Substitute now $\boldsymbol{\epsilon}(\mathbf{x})$ back into the cost, and reinstate $\lambda\|\mathbf{x}\|_1$, to obtain (4.26).

C.6 Proof of Proposition 4.7

The proof follows the lines used in proving the MAP optimality of (4.5) under (as4.1a) in Proposition 4.2. The log-prior pdf of $\mathbf{x}_o$ contains an ℓ_1-norm term as in (C.2), while the uniform prior on $\mathbf{p}_o^A$ is constant under (as4.1$'$). Furthermore, given the structure mapping $[\mathbf{A}\ \mathbf{y}]=\mathbf{S}(\mathbf{p})$, the conditional log-likelihood here can be expressed in terms of $\mathbf{x}$ and $\boldsymbol{\epsilon}^A$, as $\ln p[\mathbf{y},\mathbf{A}|\mathbf{x}_o=\mathbf{x},\mathbf{p}_o^A=\mathbf{p}^A+\boldsymbol{\epsilon}^A] = \ln p\left[\boldsymbol{\epsilon}_A=\boldsymbol{\epsilon}^A,\ \boldsymbol{\epsilon}_y=-(\boldsymbol{S}^y)^\dagger[\mathbf{y}-(\mathbf{A}+\sum_{k=1}^{n_A}\epsilon_k^A\mathbf{S}_k^A)\mathbf{x}] = (\boldsymbol{S}^y)^\dagger[\boldsymbol{S}^A(\mathbf{I}\otimes\mathbf{x})\boldsymbol{\epsilon}^A-\mathbf{r}(\mathbf{x})]\right]$. After omitting terms not dependent on $\mathbf{x}$ and $\boldsymbol{\epsilon}^A$, the conditional log-likelihood under the joint Gaussian distribution in (as4.1$'$) boils down to half of the quadratic cost in (4.25). Combining the latter with $\|\mathbf{x}\|_1$ from the log-prior pdf, it follows that maximizing the log-posterior pdf amounts to minimizing the unconstrained sum of the two, which establishes MAP optimality of the WSS-TLS estimator in (4.25).

Appendix D

Technical Details in Chapter 5

D.1 Proof of (5.21)

To solve the optimization problem (5.20), consider first the unconstrained solution to the LS part of the cost, which can be written as

$$b_k^{\mathrm{LS}} := \left(y_k' - \textstyle\sum_{\ell=k+1}^{K} R_{k\ell}\hat{b}_\ell^{\mathrm{DDD}}\right)/R_{kk}. \tag{D.1}$$

The detected symbol in (5.20) can be equivalently expressed as

$$\begin{aligned}\hat{b}_k^{\mathrm{DDD}} &= \arg\min_{b_k\in\mathcal{A}_a} f(b_k),\\ f(b_k) &:= \left(b_k^{\mathrm{LS}} - b_k\right)^2 + (2\lambda/R_{kk}^2)|b_k|_0.\end{aligned} \tag{D.2}$$

The solution $\hat{b}_k^{\mathrm{DDD}}$ can be obtained by comparing $f(0)$ with $\min_{b_k\in\mathcal{A}} f(b_k)$. Specifically, as the cost $f(b_k)$ is quadratic for $b_k \in \mathcal{A}$, the minimum is achieved at $f(\lfloor b_k^{\mathrm{LS}}\rceil)$, by quantizing b_k^{LS} to the nearest point in $\mathcal{A}$. Thus, $\hat{b}_k^{\mathrm{DDD}} = 0$ only if $f(0) \le f(\lfloor b_k^{\mathrm{LS}}\rceil)$, or equivalently, after using the definition of $f(\cdot)$, if $2b_k^{\mathrm{LS}}\lfloor b_k^{\mathrm{LS}}\rceil - \lfloor b_k^{\mathrm{LS}}\rceil^2 - 2\lambda/R_{kk}^2 \le 0$. This completes the proof of (5.21).

D.2 Proof of (5.22)

When $M = 2$, the DDD solution (5.21) reduces to

$$\hat{b}_k^{\mathrm{DDD}} = \mathrm{sign}(b_k^{\mathrm{LS}})\mathbb{1}(2|b_k^{\mathrm{LS}}| - 1 - 2\lambda/R_{kk}^2 > 0) \tag{D.3}$$

due to the fact that $\lfloor b_k^{\rm LS} \rceil = {\rm sign}(b_k^{\rm LS}) \in \{\pm 1\}$.

Recalling that $\check{\mathbf{b}}$ denotes the transmitted vector, and substituting $\mathbf{y} = \mathbf{H}\check{\mathbf{b}} + \mathbf{w}$ yields

$$\mathbf{y}' := \mathbf{Q}^T\mathbf{y} = \mathbf{R}\check{\mathbf{b}} + \mathbf{u} \tag{D.4}$$

where $\mathbf{u} := \mathbf{Q}^T\mathbf{w}$ is zero-mean Gaussian with identity covariance matrix. Supposing that there is no error propagation, the $b_k^{\rm LS}$ term in (D.3) becomes [cf. (D.1)]

$$\begin{aligned} b_k^{\rm LS} &= \left(\sum_{\ell=k}^{K} R_{k\ell}\check{b}_\ell + u_k - \sum_{\ell=k+1}^{K} R_{k\ell}\hat{b}_k^{\rm DDD} \right) \Big/ R_{kk} \\ &= \check{b}_\ell + u_k/R_{kk} \end{aligned} \tag{D.5}$$

where u_k denotes the k-th entry of $\mathbf{u}$.

To analyze the error probability for the detector in (D.3) consider the following three cases.

a) $\check{b}_k = 0$ is sent: under this case, $b_k^{\rm LS} = u_k/R_{kk}$ and a detection error emerges when $\hat{b}_k^{\rm DDD} = 1$ or -1. With the closed-form DDD detector (D.3) in mind, such an error occurs only if both $b_k^{\rm LS} \neq 0$ and $2|b_k^{\rm LS}| - 2\lambda/R_{kk}^2 - 1 > 0$ hold. The first case corresponds to $u_k \neq 0$ and the second one is equivalent to $|u_k| > |R_{kk}|/2 + \lambda/|R_{kk}|$, which is included in the event $u_k \neq 0$. Hence, to evaluate the error probability it suffices to consider only the case $|u_k| > |R_{kk}|/2 + \lambda/|R_{kk}|$.

b) $\check{b}_k = 1$ is sent: following the analysis in a), an error occurs if ${\rm sign}(R_{kk})u_k \leq -|R_{kk}|/2 + \lambda/|R_{kk}|$.

c) $\check{b}_k = -1$ is sent: following the analysis in a), an error occurs if ${\rm sign}(R_{kk})u_k \geq |R_{kk}|/2 - \lambda/|R_{kk}|$.

Given that the Gaussian distributed u_k has variance 1, the overall SER for the k-th entry b_k becomes

$$\begin{aligned} P_{e,k}^{\rm DDD} &= \textstyle\sum_{i=0,\pm 1} P({\rm error}|\check{b}_k = i)P(\check{b}_k = i) \\ &= 2(1-p_a)Q\left(\frac{|R_{kk}|}{2} + \frac{\lambda}{|R_{kk}|} \right) + p_a Q\left(\frac{|R_{kk}|}{2} - \frac{\lambda}{|R_{kk}|} \right) . \end{aligned} \tag{D.6}$$