NEW INFORMATION PROCESSING THEORY AND METHODS FOR EXPLOITING SPARSITY IN WIRELESS SYSTEMS

by

Waheed Uz Zaman Bajwa

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Electrical Engineering)

at the

UNIVERSITY OF WISCONSIN – MADISON

2009

To those two beautiful souls who are not here today to cherish this moment with me:

My father, Muhammad Zia Ullah Bajwa, and my brother Badi Uz Zaman Bajwa.

— May the two of you rest in peace —

ACKNOWLEDGEMENTS

This section has unexpectedly turned out to be the most challenging aspect of writing this dissertation. This is because far too many people have helped me—directly or indirectly and academically or nonacademically—over too many years to name them all here without accidentally omitting someone. Still, I have to try and express my gratitude to all those who have enabled me to be where I am today. It goes without saying that any omissions in the following are purely accidental and truly regrettable.

First and foremost, I would like to thank my advisors, Prof. Robert Nowak and Prof. Akbar Sayeed, for their encouragement, guidance, and support during the past six years. To state that Rob and Akbar have contrasting styles of advising would be an understatement for anyone who knows them really well. Nevertheless, their unique styles are precisely what made my experience as a graduate student all the more enriching and I am greatly indebted to them for providing me with this wonderful opportunity. In addition, I would also like to thank Rob separately for financially supporting my graduate studies and for instilling in me the love of "all things signal processing."

I would also like to thank the other three members of my dissertation committee, Prof. Stark Draper, Prof. Barry Van Veen, and Prof. Stephen Wright, for their careful reading of the dissertation draft and numerous helpful suggestions. I would in particular like to thank Steve (and Prof. Mario Figueiredo) for always being there with the latest and greatest software code that I needed for my simulations, and Stark for his advice, support, and help on a number of academic and nonacademic matters. In addition, I would also like to express my sincerest gratitude to the professors at the University of Wisconsin-Madison, including Professors Nigel Boston, James Bucklew, John Gubner, Robert Nowak, Akbar Sayeed, Barry Van Veen, and Stephen Wright, who taught me the very basics of scientific thinking and reasoning.

Next, I would like to gratefully acknowledge a few selfless individuals who were highly instrumental in ensuring that a person with humble beginnings can boast today to have a Ph.D. degree from one of the best research universities in the world. These include John and Tashia Morgridge (for generously supporting my first year of graduate studies), Prof. Farrukh Kamran and Prof. Muhammad Maud (for their unwavering trust in my abilities and for arranging financial support for my bachelor's degree), and Dr. Sohail Naqvi (for his support and mentorship).

I would also like to thank all my current and former colleagues in the Department of Electrical and Computer Engineering for their rather enviable companionship. I would in particular like to express my gratitude to Laura Balzano, Rui Castro, Brian Eriksson, Jarvis Haupt, Vasanthan Raghavan, Aarti Singh, and Thiagarajan Sivanadyan for making my stay at the department so enjoyable. I would also like to thank Jarvis separately for being an excellent collaborator and an enthusiastic poker teacher. In addition, I would especially like to thank the staff of International Reach at the University of Wisconsin-Madison and the Muslim community and Pakistani community in Madison area for making my life outside campus so fun and colorful.

Last but not least, I would like to thank all my friends and family members—both within and outside the United States—for their constant support and encouragement. I would in particular like to thank my brother Qasim Uz Zaman Bajwa, my sister, Amina Jamil, and their spouses for their incessant love and for always being there for the "youngling of the house." In addition, I would also like to thank my in-laws, who—barely three weeks into my marriage—allowed me to take their daughter so far away from them for such a long time.

Finally, it is often said that behind every successful man is a woman. In my case, however, this success of mine is attributable to two wonderful women and two charming girls: My mother, Nasim Zia, who has always placed the desires of her children before her own desires; the love of my life, Khadija Waheed, who is a remarkably strong woman, a wonderful human being, and an extremely loving wife; and my daughters, Sheza Bajwa and Shanzay Bajwa, who can bright up the darkest of days with their smiles. *Ammi*, Khadija, Sheza, and Shanzay: Thank you for all your sacrifices during the course of this degree and for bringing so much joy and happiness to my life! In the end, I conclude by reminding myself that "*All praise be to God, the Lord of the Worlds*."

TABLE OF CONTENTS

Page

Page

LIST OF TABLES

LIST OF FIGURES

Figure Page

Figure Page

ABSTRACT

The work presented in this dissertation revolves around three major research thrusts: (i) efficient acquisition of data from physical sources, (ii) reliable transmission of data from one point to another, and (iii) optimal extraction of meaningful information from given data. The common theme underlying these (often intertwined) research thrusts is what can be termed as the "blessing of sparsity": *while real-world data might live in a very high-dimensional space, the critical information conveyed by that data is often embedded in a much lower-dimensional (often non-linear) manifold of the observation space*.

The thesis of this dissertation is that "*Joint exploitation of the sparsity of real-world data by the acquisition, transmission, and information extraction (processing) operations allows design of new computationally efficient and nearly optimal information processing algorithms that—despite being agnostic to the underlying information embeddings—can reduce the amount of data collected without incurring any reduction in the information content as measured by some fidelity criterion*." In order to support our thesis, we have developed new theory and methods in the dissertation for some of the fundamental problems arising in wireless systems that involve sparse (or approximately sparse) data. In the process, we have also made a number of significant scholarly contributions in the diverse areas of compressed sensing, wireless communications, and wireless sensor networks.

First, as part of our contribution in the area of compressed sensing, we have abstractly studied in the dissertation three classes of "structured sensing vectors" that are given by the rows of either Toeplitz matrices, Gabor matrices, or "low-rank projections" of unitary matrices. Collectively, these three sensing-vector classes arise naturally in many application areas and we have rigorously proved using various tools from linear algebra, statistics, and probability theory in Banach spaces

that collections of sensing vectors belonging to these classes can also successfully encode and decode high-dimensional sparse data.

Second, as part of our contribution in the area of wireless communications, we have formalized the notion of sparse multipath channels and developed a new framework in the dissertation for estimating sparse channels in time, frequency, and space. In particular, we have established that the proposed channel estimation framework—which is based on our work on structured sensing vectors and is accordingly termed as "compressed channel sensing"—achieves a target reconstruction error using far less energy and, in many instances, latency and bandwidth than that dictated by the traditional training-based channel estimation methods.

Finally, as part of our contribution in the area of wireless sensor networks, we have proposed and analyzed new distributed algorithms in the dissertation that are capable of efficiently accomplishing the task of information extraction in resource-constrained wireless sensor networks using minimal energy and bandwidth. The basic idea behind our proposed approach is to combine processing and communication into a single operation designed to maximize the potential gain in information *per operation*. Using this procedure, we have shown that critical information in sensor network data can be reliably obtained at a distant fusion center as long as the total number of "information processing operations" carried out in the network is proportional to the "intrinsic" dimension of the information embedding.

Chapter 1

Introduction

> "*Something marvelous has been happening to humankind. Information is moving faster and becoming more plentiful, and people everywhere are benefiting from this change. But there's a surprising postscript to this story. When it comes to information, it turns out that one can have too much of a good thing.*"
>
> — *David Shenk*, Author of *Data Smog* (1997)

1.1 Motivation

While the expression "we live in an information age" has undoubtedly become one of the most worn out clichés of the 21st century, there is no denying the fact that the information revolution has had a profound impact on the lives of ordinary people and scientific researchers alike. The dominant trend underscoring the evolution of this phenomenon can be summed up in one phrase: *Data, data, and more data*. Thanks to spectacular technological advances of the last two decades, scientists and engineers have been able to build devices and study systems that are capable of generating massive quantities of data, on scales considered unimaginable until recently. Paradoxically, however, this sheer abundance of (raw) data is also threatening to become the Achilles' heel of the information revolution: Computational and analytical tools developed in the 20th century for the extraction of information from data are fast becoming irrelevant in the face of large problem sizes necessitated by today's applications. Therefore, the challenge facing us today is to devise a new computationally efficient set of *information processing* tools that can effectively cope with this relentless barrage of data.

It is generally recognized in this regard that the basic operations of acquisition, transmission, and processing are interdependent and—in order to attain optimal performance—they must be jointly optimized in problems involving large collections of data. Despite the need for joint optimization, however, our fundamental understanding of this complex problem is very limited, owing in part to the absence of a well-developed mathematical theory. As a result, information processing tools over the last few decades have been largely developed by separating them from the data acquisition and transmission. Despite the past success of this modular approach, however, there is now an imminent need for a reconnection between acquisition, transmission, and processing if we are to successfully manage the 21st century data deluge. Only by looking at these three operations through a unified lens can we characterize the relationships among them, reveal fundamental trade-offs between the sizes of data compilations and the quality of retained information, and devise radically new, but highly efficient, approaches to information extraction by *simultaneously* exploiting data redundancy at *all* stages of the problem. This dissertation is one such small, but nonetheless significant, undertaking in this direction.

1.2 Thesis Statement

The work presented in this dissertation revolves around three major research thrusts: (i) efficient acquisition of data from physical sources, (ii) reliable transmission of data from one point to another, and (iii) optimal extraction of meaningful information from given data. The common theme underlying these (often intertwined) research thrusts is what can be termed as the *blessing of sparsity*. The task of gleaning information from data, aptly termed information processing, hinges on our ability to gather a large collection of observations that adequately capture the underlying phenomenon of interest. The more observations we gather, however, the harder it becomes to make sense out of the collected data. The phrase “curse of dimensionality” is often used in scientific and engineering circles to describe this paradox. Nevertheless, it has long been observed that data in the real-world are often approximately sparse: While the data might live in a very high-dimensional space, the critical information conveyed by the data is often embedded in a much lower-dimensional (often non-linear) manifold of the observation space. Intuitively, this

means that one needs to focus resources only on this lower-dimensional manifold of the data, provided the information embedding can be learned in a computationally efficient manner. This has been the key idea behind the success of many information processing tools used for compression, estimation, data mining, pattern recognition, etc.

Classical information processing methods, however, were not designed to cope with the kind of explosive data growth that we are seeing today. In particular, the effectiveness of these methods is getting constrained by their inability to learn the lower-dimensional information embeddings (in large data sets) in a computationally tractable manner. This necessitates a fundamental rethinking of the data gathering and information processing problem, which brings us to the thesis of this dissertation, stated as follows:

> *Joint exploitation of the sparsity of real-world data by the acquisition, transmission, and processing operations allows design of new computationally efficient and nearly optimal information processing algorithms that—despite being agnostic to the underlying information embeddings—can reduce the amount of data collected without incurring any reduction in the information content as measured by some fidelity criterion.*

1.3 Major Contributions

In order to support our thesis, we have developed new theory and methods in the dissertation for some of the fundamental problems arising in wireless systems that involve sparse (or approximately sparse) data. In the process, we have also made a number of significant scholarly contributions in the diverse areas of compressed sensing, wireless communications, and wireless sensor networks. Below, we highlight some of the primary aspects of these contributions.

Compressed Sensing

Compressed sensing is a relatively new area of theoretical research that lies at the intersection of a number of other research areas such as signal processing, statistics, and computational harmonic analysis, and describes a new acquisition paradigm in which sparse (or approximately sparse),

high-dimensional data can be well-approximated by a small number of its (nonadaptive, linear) projections onto a collection of sensing vectors [1–5]. At an abstract level, there are two main ingredients to a compressed sensing problem:

[1] Designing a collection of sensing vectors that can adequately capture critical information in the high-dimensional data.

[2] Designing computationally efficient reconstruction methods that can faithfully reproduce data from the resulting projections.

In particular, with regard to **[1]**, some of the earliest work in the compressed sensing literature has established the sufficiency of using either independent realizations of certain zero-mean random variables or rows of certain unitary matrices as sensing vectors [6–12]. From an implementation viewpoint, however, it is not always possible to use these *unstructured* sensing vectors for acquisition purposes in many application areas due to the physics of the underlying problems.

It is in this context that we abstractly study in the dissertation three specific classes of *structured sensing vectors* that are given by the rows of either Toeplitz matrices, Gabor matrices, or "low-rank projections" of unitary matrices.[1] Collectively, these three sensing-vector classes arise naturally in many application areas such as time-invariant and time-varying linear system identification [13–15], time-frequency analysis [16], coded aperture imaging [17], sampling theory [18], and radar and seismic imaging [19,20]. As part of our **first major contribution**, which appears in **Chapter 3** of the dissertation, we rigorously prove using various tools from linear algebra, statistics, and probability theory in Banach spaces that collections of sensing vectors belonging to these classes can also successfully encode and decode high-dimensional sparse data.

Wireless Communications

Wireless communication systems have emerged as the vital backbone of information revolution over the last two decades. In particular, coherent communication systems are generally far

[1] Note that the term *projection* is not being used here in the usual linear algebra sense; see Section 3.5 for further details on this.

more efficient than the non-coherent ones, but require that the channel response be known at the receiver [21,22]. In practice, however, the channel response is seldom—if ever—available to communication systems a priori and the channel needs to be (periodically) estimated at the receiver to reap the benefits of coherent communication. As such, training-based methods—which probe the channel in time, frequency, and space with known signals and reconstruct the channel response from the output signals—are most commonly used to accomplish this task [23].

Traditional training-based channel estimation methods, typically comprising of linear reconstruction techniques (such as the maximum likelihood or the minimum mean squared error estimators), are known to be optimal for rich multipath channels [24–32]. However, physical arguments and growing experimental evidence suggest that wireless channels encountered in practice exhibit a sparse structure that gets pronounced as the signal space dimension gets large (e.g., due to large bandwidth or large number of antennas) [33–37]. Such *sparse channels* can be characterized with significantly fewer parameters compared to the maximum number dictated by the angle-delay-Doppler spread of the channel. Abstractly, all the relevant information about a sparse channel is embedded in an unknown low-dimensional manifold of the high-dimensional channel space and the challenge is to learn this embedding without resorting to probing the entire channel space.

As part of our **second major contribution**, which appears in **Chapter 4** of the dissertation, we formalize the notion of sparse multipath channels and develop a new framework for estimating sparse channels in time, frequency, and space. In particular, we establish that the proposed channel estimation framework—which is based on our work on structured compressed sensing vectors (matrices) and is accordingly termed as *compressed channel sensing*—achieves a target reconstruction error using far less energy and, in many instances, latency and bandwidth than that dictated by the traditional training-based methods.

Wireless Sensor Networks

Sensor networking is an emerging technology that promises an unprecedented ability to monitor the physical world via a spatially distributed network of small and inexpensive wireless devices

that have the ability to self-organize into a well-connected network [38]. A wide range of applications of sensor networks are being envisioned in a number of areas, including geographical monitoring (e.g., habitat monitoring, precision agriculture), industrial control (e.g., in a power plant or a submarine), business management (e.g., inventory tracking with radio frequency identification tags), homeland security (e.g., tracking and classifying moving targets) and health care (e.g., patient monitoring, personalized drug delivery) [39]. The essential task in many such applications of sensor networks is to extract relevant information about the sensed data—which we call *networked data* to emphasize both the distributed nature of the data and the fact that the data may be shared over the underlying communications infrastructure of the network—and deliver it with a desired fidelity to a (usually) distant fusion center. The overall goal in the design of sensor networks is to execute this task with least consumption of network resources—energy and bandwidth being the most limited resources, typically.

As part of our **third major contribution**, which appears in **Chapter 5** of the dissertation, we develop new distributed algorithms that are capable of efficiently accomplishing the task of information extraction in resource-constrained wireless sensor networks. Our approach represents a departure from existing methodologies from an architectural and protocol viewpoint, and involves a novel combination of techniques from nonparametric statistics, compressed sensing, and wireless communications to effectively straddle the two extremes of: (i) in-network data processing followed by transmission of sufficient statistics to the fusion center, and (ii) communication of raw data to the fusion center followed by out-of-network information extraction. The basic idea behind the proposed approach—inspired by recent results in wireless communications [40–43]—is to combine processing and communication into a single operation designed to maximize the potential gain in information *per operation*. Using this procedure, we show that critical information in the networked data can be obtained at the fusion center as long as the total number of "*information processing operations*" carried out in the network is proportional to the *intrinsic* dimension of the information embedding.

A few other remarkable features of the proposed framework include: (i) it requires almost no explicit collaboration among sensing nodes, (ii) consistent estimates can be obtained at the fusion

center under mild assumptions on the structure of the networked data even if the total network power consumption tends to zero asymptotically, and (iii) consistent (though necessarily suboptimal) estimates can be obtained at the fusion center even if no prior knowledge is assumed about the structure of the networked data.

1.4 Notational Convention

Here, we present some general and basic notation that we have tried to use coherently throughout this dissertation. Any exceptions to this notational convention, while rare, are explicitly mentioned in the body of the dissertation.

- **Set and Function Notation:** We use $\mathbb{R}, \mathbb{C}$, and $\mathbb{N}$ to denote the sets of all real numbers, complex numbers, and positive integers (usually starting from 1), respectively. Given a collection of sets $\{\mathcal{X}_i\}_{i=1}^n$, we use $\mathcal{X}_1 \times \cdots \times \mathcal{X}_n$ to denote their cartesian product. Given two integers n and m, we use the shorthand notation $[n \ldots m]$ to denote the set of all consecutive integers between (and including) n and m: $[n \ldots m] \stackrel{def}{=} \{n, n+1, \ldots, m\}$, where the symbol $\stackrel{def}{=}$ means "*equality by-virtue-of definition*." Given any $x \in \mathbb{R}$, we use $\lfloor x \rfloor$ and $\lceil x \rceil$ to denote the largest integer less than or equal to x and the smallest integer greater than or equal to x, respectively. Given any $z \in \mathbb{C}$, we use z^* to denote the conjugate of z. We also use $|\cdot|$ to denote both the magnitude of a real- or complex-valued quantity x and the cardinality of a finite set $\mathcal{X}$. Given any constant $c > 1$, we sometimes use the shorthand notation $\mathrm{polylog}(x)$ to denote the function $\log^c(x)$. In addition, we define the indicator function $1_{\mathcal{X}}(x)$ to take the value 1 if $x \in \mathcal{X}$ and 0 otherwise, while we use δ_{ij} to denote the Kronecker delta, which takes value 1 if $i = j$ and 0 otherwise. Finally, we use $x \sim F(a, b)$ to denote a random variable x with the cumulative distribution function $F(a, b)$. In particular, we use $F(a, b) = \mathcal{N}(m_1, \sigma_1^2)$ to denote a Gaussian distribution with mean m_1 and variance σ_1^2, while we use $F(a, b) = \mathcal{CN}(m_2, \sigma_2^2)$ to denote a circularly symmetric, complex-Gaussian distribution with mean m_2 and variance σ_2^2.

- **Linear Algebra Notation:** We use bold-faced, upper-case letters, such as $\mathbf{A}$ and $\mathbf{B}$, to denote matrices. Similarly, we use bold-faced, lower-case letters, such as $\mathbf{x}$ and $\mathbf{y}$, to denote vectors. Further, unless explicitly stated, we take all the vectors to be column vectors. Given any $n \times m$ matrix $\mathbf{A}$, we use $\mathsf{rank}(\mathbf{A}), \mathsf{trace}(\mathbf{A})$, and $\mathsf{vec}(\mathbf{A})$ to denote the rank of $\mathbf{A}$, the trace of $\mathbf{A}$, and the $nm \times 1$ vectorized version of $\mathbf{A}$ (obtained by stacking all its columns), respectively. We also use $\|\mathbf{A}\|_2, \|\mathbf{A}\|_F$, and $\|\mathbf{A}\|_{\max}$ to denote the spectral norm of $\mathbf{A}$ (the largest singular value of $\mathbf{A}$), the Frobenius norm of $\mathbf{A}$, and the max norm of $\mathbf{A}$ (absolute value of the largest-magnitude entry of $\mathbf{A}$), respectively. Sometimes, we also use the shorthand notation $\mathbf{A} \in \mathbb{C}^{n \times m}$ to denote a complex-valued matrix $\mathbf{A}$ that has n rows and m columns. Given any $n \times 1$ vector $\mathbf{x}$, we use $\|\mathbf{x}\|_p, \|\mathbf{x}\|_0$, and $\mathsf{diag}(\mathbf{x})$ to denote the usual ℓ_p-norm of $\mathbf{x}$, the number of nonzero entries of $\mathbf{x}$, and the $n \times n$ diagonal-matrix version of $\mathbf{x}$ (obtained by placing its entries on the main diagonal of a square matrix), respectively. We further use $\mathbf{I}_n, \mathbf{O}_n, \mathbf{1}_n$, and $\mathbf{0}_n$ to denote $n \times n$ identity matrices, $n \times n$ all-zeros matrices, $n \times 1$ all-ones vectors, and $n \times 1$ all-zeros vectors, respectively. In addition, we use the superscripts $(\cdot)^{\mathsf{T}}, (\cdot)^{\mathsf{H}}$, and $(\cdot)^{\dagger}$ to denote the operations of transposition, conjugate transposition, and (Moore–Penrose) pseudoinverse, respectively. Finally, we use $\langle \cdot, \cdot \rangle$ to denote an inner product between two vectors that is linear in the first argument, while we use $\otimes$ and $\odot$ to denote the Kronecker product and the Hadamard product, respectively.

- **Scaling Notation:** We establish scaling relationships between different quantities using Landau's notation. Specifically, if $f(x)$ and $g(x)$ are positive-valued functions of $x \in \mathbb{R}$, then we write $f(x) = O\big(g(x)\big)$ and $g(x) = \Omega\big(f(x)\big)$ if there exist some $c_o > 0$ and some $x_o \in \mathbb{R}$ such that $f(x) \leq c_o\, g(x)\ \forall\ x \geq x_o$, while we write $f(x) = \Theta\big(g(x)\big)$ if $f(x) = O\big(g(x)\big)$ and $g(x) = O\big(f(x)\big)$. Occasionally, with a slight abuse of notation, we also write $f(x) = O\big(g(x)\big)$ even though we really mean $f(x) = \Theta\big(g(x)\big)$. In addition, we sometimes also express the scaling relationships using Hardy's notation for compactness: $f(x) \preceq g(x), f(x) \succeq g(x)$, and $f(x) \asymp g(x)$ in place of $f(x) = O\big(g(x)\big), f(x) = \Omega\big(g(x)\big)$, and $f(x) = \Theta\big(g(x)\big)$, respectively. Finally, we write $f(x) \sim g(x)$ if there exists some positive-valued function $h(x)$ such that $f(x) \preceq h(x)$ and $g(x) \preceq h(x)$.

1.5 Dissertation Outline

The rest of this dissertation is organized as follows. In Chapter 2 of the dissertation, we briefly review the key compressed sensing results that are the most relevant to our discussion in the rest of the dissertation.

In Chapter 3 of the dissertation, we prove using tools from linear algebra, statistics, and probability theory in Banach spaces that collections of structured sensing vectors given by the rows of certain Toeplitz matrices, Gabor matrices, and low-rank projections of unitary matrices can also successfully encode and decode high-dimensional sparse data.

In Chapter 4 of the dissertation, we motivate the idea of compressed channel sensing for estimating sparse single- and multiple-antenna channels and, using results from Chapter 2 and Chapter 3, rigorously establish that compressed channel sensing achieves a target reconstruction error using far less energy and, in many instances, latency and bandwidth than that dictated by the traditional training-based methods.

Finally, in Chapter 5 of the dissertation, we develop and analyze an energy efficient distributed architecture for estimation of both sparse and approximately sparse networked data in resource-constrained wireless sensor networks.

Together, Chapters 3–5 constitute the major original research contributions of the dissertation. As an organizational convention, we have tried to make each of these chapters as self-contained as possible and—instead of the more general practice of concluding the dissertation with a discussion chapter—we have opted to conclude each chapter with a discussion section of its own.

Chapter 2

Theory of Compressed Sensing: A Brief Overview

2.1 Introduction

In signal processing, the purpose of sampling (or sensing) is to accurately capture the salient information in a signal of interest using as few samples as possible. A question that often comes up then in designing sampling systems is: *what is the minimum number of samples needed to ensure perfect recovery of the original signal?* The Nyquist–Shannon sampling theorem, which forms the basis of modern-day signal processing, provides a satisfactory answer to this question for the class of bandlimited signals: *signals that are bandlimited to W Hz can be perfectly recovered from their samples as long as the (uniform) sampling rate exceeds W samples per second.* The theory of compressed sensing (CS) can be thought of as a generalization of this traditional sampling theory (applicable only to bandlimited signals) to a much broader class of signals.

In order to rigorously motivate and carefully review the theoretical underpinnings of CS, we begin with the following classical linear measurement model

$$\nu_i = \mathbf{a}_i^{\mathsf{H}} \boldsymbol{\beta}\,,\ i = 1, \ldots, n \tag{2.1}$$

where $\mathbf{a}_i^{\mathsf{H}} \in \mathbb{C}^p$ is a known row vector, termed as a *sensing vector*, and $\boldsymbol{\beta} \in \mathbb{C}^p$ is a nonzero, deterministic but unknown vector. The model (2.1) corresponds to a nonadaptive measurement process that senses a discrete signal $\boldsymbol{\beta} \in \mathbb{C}^p$ by taking n linear measurements of the signal. This measurement model can also be written compactly using the matrix-vector representation

$$\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta} \tag{2.2}$$

where $\boldsymbol{\nu} \in \mathbb{C}^n$ is termed as the *observation* or *measurement vector*, and the *sensing matrix* $\mathbf{A} \in \mathbb{C}^{n \times p}$ is comprised of the n sensing vectors as its rows. The goal then is to reliably recover $\boldsymbol{\beta}$ from the knowledge of $\boldsymbol{\nu}$ and $\mathbf{A}$.

Conventional wisdom in solving (2.2) for $\boldsymbol{\beta}$ follows the basic principle of elementary linear algebra [44]: one needs $n \geq p$ to ensure a successful (and unique) recovery of $\boldsymbol{\beta}$ from $\boldsymbol{\nu}$. This conventional wisdom is indeed true in general. However, CS—a relatively new area of theoretical research that lies at the intersection of a number of other research areas such as signal processing, statistics, and computational harmonic analysis—suggests that the condition $n \geq p$ can be relaxed under certain circumstances. Specifically, if one assumes that $\boldsymbol{\beta}$ is intrinsically low-dimensional—in the sense that only a few entries of $\boldsymbol{\beta}$ are nonzero—then one can seek *sparse* solutions to (2.2). The search for sparse solutions of (2.2) completely transforms the problem at hand and can lead to successful recovery of $\boldsymbol{\beta}$ even when n is much smaller than p.

At a fundamental level, the theory of CS—sometimes also referred to as the theory of sparse approximation or the theory of sparse signal representation—deals with the case of $n \ll p$ and attempts to answer the following questions:

[Q1] What conditions does $\mathbf{A}$ need to satisfy to ensure successful recovery of a sparse $\boldsymbol{\beta}$?

[Q2] Can the solution to (2.2) be reliably obtained in practice using polynomial-time solvers?

[Q3] What performance guarantees can be given for various practical solvers when $\boldsymbol{\nu}$ is corrupted by either stochastic noise or deterministic perturbation?

A number of researchers have successfully addressed these questions, and their extensions to less restrictive notions of sparsity, over the past few years. In particular, the celebrated success of CS theory—as evidenced by its applications in areas as diverse as coding and information theory [6,45], sampling theory [18,46], imaging [47,48], and sensor networks [49–53]—can primarily be attributed to the following research breakthroughs:

[1] A relatively small number—typically much smaller than p—of appropriately designed sensing vectors can capture most of the salient information in a signal $\boldsymbol{\beta}$ that is either sparse

(has only a few nonzero entries) or approximately sparse (when reordered by magnitude, its entries decay rapidly).

[2] The signal $\boldsymbol{\beta}$ in this case can be reliably reconstructed from (noiseless or noisy) $\boldsymbol{\nu}$ by making use of tractable convex optimization programs, efficient greedy algorithms, or fast iterative thresholding methods.

At this point, the CS literature is growing so rapidly that it is difficult to do any justice to its achievements and results in this chapter alone. Instead, we briefly review the key CS results in this chapter that are the most relevant to our discussion in the dissertation; we refer the reader to [1,2,4] for a tutorial overview of some of the foundational developments and to [5] for some of the recent advances in this field.

2.2 Necessary and Sufficient Conditions for Recovery of Sparse Signals

We begin by revisiting the problem of recovering $\boldsymbol{\beta}$ from $\boldsymbol{\nu}$ with the added constraint that $\boldsymbol{\beta}$ is S-sparse (i.e., no more than S entries of $\boldsymbol{\beta}$ are nonzero). Mathematically, this can be expressed using the so-called "ℓ_0-norm" notation

$$\|\boldsymbol{\beta}\|_0 \stackrel{def}{=} \#\{i : |\beta_i| \neq 0\} \leq S. \tag{2.3}$$

Note that an implicit assumption underlying this notion of signal sparsity is that $S \ll p$ (in particular, we have that $S < 0.5p$). Now suppose that either the null-space of $\mathbf{A}$ contains $\boldsymbol{\beta}$, i.e., $\mathbf{A}\boldsymbol{\beta} = \mathbf{0}$, or $\mathbf{A}$ maps another distinct S-sparse signal, say $\boldsymbol{\beta}'$, to the same observation vector $\boldsymbol{\nu}$, i.e., $\mathbf{A}\boldsymbol{\beta}' = \boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta}$. One could not possibly hope to recover $\boldsymbol{\beta}$ in this case since the measurement vector does not provide (i) any information about $\boldsymbol{\beta}$ in the former scenario, and (ii) enough information about $\boldsymbol{\beta}$ in the latter scenario. We therefore have the following theorem and corollary from linear algebra.

Theorem 2.1 Any arbitrary S-sparse signal $\boldsymbol{\beta}$ can be uniquely recovered from $\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta}$ only if every $n \times 2S$ submatrix of $\mathbf{A}$ has full column rank.

Corollary 2.2 Any arbitrary S-sparse signal $\boldsymbol{\beta}$ can be uniquely recovered from $\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta}$ only if the number of observations $n \geq 2S$.

The proof of Theorem 2.1 is rather elementary in nature and is given in Section 2.5.1. Also, recall that the rank of a matrix is upper bounded by the minimum of the number of rows and the number of columns of the matrix. Corollary 2.2 therefore follows trivially from Theorem 2.1.

The property that every $n \times 2S$ submatrix of $\mathbf{A}$ has full column rank was studied in [54, 55] for the uniqueness of sparse solutions of underdetermined systems of equations. We term this property as the *unique representation property* (URP) following the terminology in [54].

Definition 2.3 (Unique Representation Property) An $n \times p$ matrix $\mathbf{A}$ is said to have the URP of order $2S$ if every $n \times 2S$ submatrix of $\mathbf{A}$ has full column rank.

The importance of URP for the study of the uniqueness of sparse solutions was first unraveled in [54]. In particular, it was shown in [54, 55] that URP of order $2S$ is also a sufficient condition for unique recovery of S-sparse $\boldsymbol{\beta}$ from (2.2). Specifically, define the combinatorial optimization program (P_0) as

$$\boldsymbol{\beta}_0 = \underset{\tilde{\boldsymbol{\beta}} \in \mathbb{C}^p}{\arg\min} \|\tilde{\boldsymbol{\beta}}\|_0 \quad \text{subject to} \quad \boldsymbol{\nu} = \mathbf{A}\tilde{\boldsymbol{\beta}} \tag{P_0}$$

then we have the following theorem regarding the equivalence of $\boldsymbol{\beta}_0$ and the true $\boldsymbol{\beta}$.

Theorem 2.4 If the sensing matrix $\mathbf{A}$ satisfies URP of order $2S$ then any arbitrary S-sparse signal $\boldsymbol{\beta}$ can be uniquely recovered from $\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta}$ as a solution to the optimization program (P_0).

The proof of this theorem is given in Section 2.5.2 for the sake of completion; similar versions of the proof can also be found in [54, 55].

Unfortunately, a straightforward approach to solving (P_0) seems hopeless since it is an NP-hard problem [56]. The computational intractability of (P_0) has over the years led researchers to develop many heuristic (tractable) approximations of the problem, including convex relaxations of (P_0) [54, 57], greedy algorithms [58, 59], and iterative thresholding methods [60, 61]. The results achieved so far in the CS literature range from identifying conditions under which (P_0) has

the same solution as its heuristic approximations, to conditions under which the approximations yield a reliable sparse solution even when $\boldsymbol{\beta}$ is not truly sparse, to conditions under which the approximations yield a robust solution in a stochastic or an adversarial noise setting. Some of the strongest results in this regard have been obtained for the convex optimization based approach to solving (2.2) for a sparse $\boldsymbol{\beta}$. As such, we focus only on those recovery/reconstruction methods in the sequel that are based on (or inspired by) convex relaxation of (P_0)—see [5] for the references of approximate solutions based on greedy algorithms and iterative thresholding methods.

2.3 Sufficient Conditions for Practical Recovery of Sparse Signals

In the literature, a frequently discussed alternative to the computational intractability of (P_0) is to regularize the problem by replacing the (highly discontinuous) ℓ_0-norm with an ℓ_p- "norm" for some $p \in (0, 1]$ [54]. While this is a practical strategy, little can be guaranteed in terms of whether a local minimum of the resulting problem will actually be a good approximation to the global minimum of (P_0). Instead, a better strategy is to *convexify* the problem by replacing the ℓ_0-norm with the ℓ_1-norm

$$\boldsymbol{\beta}_1 = \underset{\tilde{\boldsymbol{\beta}} \in \mathbb{C}^p}{\arg\min} \|\tilde{\boldsymbol{\beta}}\|_1 \quad \text{subject to} \quad \boldsymbol{\nu} = \mathbf{A}\tilde{\boldsymbol{\beta}} \tag{BP}$$

which results in a global minimum because of the convex nature of the problem [62]. This optimization program, which goes by the name of *basis pursuit* in the signal processing literature, is computationally tractable because it can also be recast as a linear program [57].

We now discuss the performance guarantees of basis pursuit and specify the conditions under which solving (BP) is equivalent to solving (P_0). Clearly, this equivalence cannot be expected for all sensing matrices $\mathbf{A}$ that satisfy URP of order $2S$, since this would contradict the known NP-hardness of (P_0) in the general case. Nevertheless, the initial success of CS theory is largely in part due to the seminal works of Candès and Tao [6, 8], Candès, Romberg and Tao [7, 9], and Donoho [10] that established that (BP) can produce the globally optimal solution of (P_0) under mildly stronger conditions on $\mathbf{A}$. Proofs of these remarkable initial results all rely on the same property of the sensing matrix, namely that any collection of $2S$ columns of (appropriately normalized)

$\mathbf{A}$ should behave almost like an isometry. One concise way to state this condition is through the *restricted isometry property* (RIP), first introduced in [6]. The RIP, defined below, can be leveraged to establish a series of fundamental results in CS.

Definition 2.5 (Restricted Isometry Property) An $n \times p$ matrix $\mathbf{A}$ having unit ℓ_2-norm columns is said to have the RIP of order S with parameter δ_S if there exists some $\delta_S \in (0, 1)$ such that

$$(1 - \delta_S)\|\tilde{\boldsymbol{\beta}}\|_2^2 \leq \|\mathbf{A}\tilde{\boldsymbol{\beta}}\|_2^2 \leq (1 + \delta_S)\|\tilde{\boldsymbol{\beta}}\|_2^2 \tag{2.4}$$

holds for all S-sparse vectors $\tilde{\boldsymbol{\beta}}$. In this case, we sometimes make use of the shorthand notation $\mathbf{A} \in RIP(S, \delta_S)$ to state that $\mathbf{A}$ satisfies the RIP of order S with parameter δ_S.

The initial contributions to the theory of CS established, essentially, that (BP) and (P_0) have identical solutions for all S-sparse signals $\boldsymbol{\beta}$ if an appropriately normalized $\mathbf{A}$ satisfies the RIP of order $2S$ with a sufficiently small parameter δ_{2S}. The following theorem—a generalization of the earlier results—also describes the recovery of signals that are not exactly sparse.

Theorem 2.6 (Noiseless Recovery [63]) Let $\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta}$ be an $n \times 1$ vector of observations of any deterministic but unknown signal $\boldsymbol{\beta} \in \mathbb{C}^p$. Assume that the columns of $\mathbf{A}$ have unit ℓ_2-norms and further let $\mathbf{A} \in RIP(2S, 0.3)$. Then the vector $\boldsymbol{\beta}_1$ obtained as the solution of (BP) satisfies

$$\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}\|_2^2 \leq c_0 \frac{\|\boldsymbol{\beta} - \boldsymbol{\beta}_S\|_1^2}{S} \tag{2.5}$$

where $\boldsymbol{\beta}_S$ is the vector formed by setting all but the S largest (in magnitude) entries of $\boldsymbol{\beta}$ to zero, and $c_0 > 0$ is a constant given by

$$c_0 = 4\left(\frac{1 + \delta_{2S}}{1 - 3\delta_{2S}}\right)^2. \tag{2.6}$$

Remark 2.7 The statement of Theorem 2.6 is a slight variation on [63, Theorem 1.2], which arises due to the complex-valued setup here as opposed to the real-valued one in [63]. Specifically, in the case of a real-valued setup, one only requires that $\mathbf{A} \in RIP(2S, 0.41)$ and the constant c_0 in that case can be given by

$$c_0 = 4\left(\frac{1 - \delta_{2S} + \sqrt{2}\delta_{2S}}{1 - \delta_{2S} - \sqrt{2}\delta_{2S}}\right)^2. \tag{2.7}$$

Note that Theorem 2.6 guarantees that the recovery of $\boldsymbol{\beta}$ is exact in the case when $\boldsymbol{\beta}$ has no more than S nonzero entries (since $\boldsymbol{\beta}_S = \boldsymbol{\beta}$ in that case). It is worth mentioning at this point that the idea to use the ℓ_1-norm as a sparsity-inducing objective function existed as early as in 1973 in the geophysics literature [64]. In fact, Santosa and Symes developed this idea further in 1986 and proved that a variation of (BP) (termed basis pursuit denoising) succeeds in recovering sparse spike trains under moderate restrictions [19]. However, it is only recently that researchers have been able to get the most rigorous results concerning the equivalence between (BP) and (P_0).

2.3.1 Compressed Sensing Matrices

It is clear from the definition of RIP that the condition $\mathbf{A} \in RIP(2S, 0.3)$ is essentially a statement about the singular values of all $n \times 2S$ submatrices of $\mathbf{A}$. However, the definition of RIP and the statement of Theorem 2.6 make no mention of either (i) how to design sensing matrices that satisfy the RIP of order $2S$ or (ii) how to check if a given sensing matrix satisfies the RIP of order $2S$. Nevertheless, while no algorithms are known to date that can check the RIP for a given matrix in polynomial time, one of the reasons that has led to the widespread applicability of CS theory in various application areas is the revelation that certain probabilistic constructions of matrices satisfy the RIP with high probability. In this regard, the following theorems are representative of the relevant results that can be found in the CS literature.

Theorem 2.8 (Independent and Identically Distributed Matrices [11]) Let $\mathbf{A}$ be an $n \times p$ matrix whose entries are drawn in an independent and identically distributed (i.i.d.) fashion from one of the following zero-mean distributions, each having variance $1/n$:

- $a_{i,j} \overset{i.i.d.}{\sim} \mathcal{N}(0, 1/n)$,
- $a_{i,j} \overset{i.i.d.}{\sim} \begin{cases} 1/\sqrt{n} & \text{with probability } 1/2 \\ -1/\sqrt{n} & \text{with probability } 1/2 \end{cases}$,
- $a_{i,j} \overset{i.i.d.}{\sim} \begin{cases} \sqrt{3/n} & \text{with probability } 1/6 \\ 0 & \text{with probability } 2/3 \\ -\sqrt{3/n} & \text{with probability } 1/6 \end{cases}$.

For each integer $S \in \mathbb{N}$, and for any $\delta_S \in (0,1)$ and any $c_1 < \delta_S^2(3-\delta_S)/48$, set

$$c_2 = \frac{192 \log(12/\delta_S)}{3\delta_S^2 - \delta_S^3 - 48c_1}. \tag{2.8}$$

Then whenever $n \geq c_2 S \log p$, $\mathbf{A} \in RIP(S, \delta_S)$ with probability exceeding $1 - \exp(-c_1 n)$.

Theorem 2.9 (Subsampled Unitary Matrices [12]) Let $\mathbf{U}$ be any $p \times p$ unitary matrix. Choose a subset Ω of cardinality $n \stackrel{def}{=} |\Omega|$ uniformly at random from the set $[1 \ldots p]$. Further, let $\mathbf{A}$ be the $n \times p$ matrix obtained by sampling n rows of $\mathbf{U}$ corresponding to the indices in Ω and renormalizing the resulting columns so that they have unit ℓ_2-norms. For each integer $p, S > 2$, and for any $t > 1$ and any $\delta_S \in (0,1)$, let

$$n \geq (c_3 \, \mu_{\mathbf{U}}^2 \, tS \log p) \log(tS \log p) \log^2 S \tag{2.9}$$

then the subsampled matrix $\mathbf{A} \in RIP(S, \delta_S)$ with probability exceeding $1 - 10\exp(-c_4 \delta_S^2 t)$. Here, $\mu_{\mathbf{U}} \stackrel{def}{=} \sqrt{p} \max_{i,j} |u_{i,j}|$ is termed as the *coherence* of the unitary matrix $\mathbf{U}$, and $c_3, c_4 > 0$ are absolute constants that do not depend on n, p, or S.

Corollary 2.10 (Polynomial Probability of Success [12]) Let $\mathbf{U}$ be any $p \times p$ unitary matrix with entries of magnitude $O(1/\sqrt{p})$. Then for each integer $p, S > 2$, and for any $\delta_S \in (0,1)$, the (appropriately normalized) matrix $\mathbf{A}$ obtained by sampling $n = \Omega(S \log^5 p)$ rows of $\mathbf{U}$ uniformly at random satisfies $RIP(S, \delta_S)$ with probability exceeding $1 - p^{-O(\delta_S^2)}$.

Remark 2.11 The original specification of the results in [12] assumed that $\mu_{\mathbf{U}} = O(1/\sqrt{p})$ and $\delta_S = 0.5$, but the proofs actually provide more general results for arbitrary $\mu_{\mathbf{U}}$ and δ_S. In addition, the subset Ω in [12] corresponds to Bernoulli sampling of the set $[1 \ldots p]$. That is, let $\zeta_1, \ldots, \zeta_p$ be independent Bernoulli random variables taking the value 1 with probability n/p. Then,

$$\Omega = \{i : \zeta_i = 1\}. \tag{2.10}$$

Nevertheless, it has been shown in [7] that if the subsampled unitary matrix $\mathbf{A} \in RIP(S, \delta_S)$ with probability $1 - \eta$ for the Bernoulli sampling model, then $\mathbf{A} \in RIP(S, \delta_S)$ with probability $1 - 2\eta$ for the uniformly-at-random sampling model. Hence, the statement of Theorem 2.9 above.

Note that Corollary 2.10 trivially follows from Theorem 2.9 by taking $t = \Theta(\log p)$. The preceding discussion in this section and Theorem 2.6 essentially guarantee that *practical* recovery of sparse signals from (2.2) is possible (with high probability) using only $n = \Omega\left(S \times \mathrm{polylog}(p)\right)$ observations. In this sense, the near-optimality of noiseless CS is evident.

2.3.2 Remark on Minimum ℓ_2-Norm Reconstruction

Another classical approach to the computational intractability of (P_0) is to convexify the problem by replacing the ℓ_0-norm with the ℓ_2-norm

$$\boldsymbol{\beta}_2 = \underset{\tilde{\boldsymbol{\beta}} \in \mathbb{C}^p}{\arg\min} \|\tilde{\boldsymbol{\beta}}\|_2 \quad \text{subject to} \quad \boldsymbol{\nu} = \mathbf{A}\tilde{\boldsymbol{\beta}} \tag{P_2}$$

which also results in a global minimum because of the convex nature of the problem. Geometrically, the collection of all solutions to (2.2) is an affine subspace of $\mathbb{C}^p$ and (P_2) selects that element of this subspace which is the closest to the origin. As such, $\boldsymbol{\beta}_2$ is sometimes also called the *minimum-energy solution*.

The key advantage that (P_2) has over other convex approximations of (P_0) is that it has a nice closed-form solution given by the Moore–Penrose pseudoinverse of $\mathbf{A}$: $\boldsymbol{\beta}_2 = \mathbf{A}^\dagger \boldsymbol{\nu}$. However, (P_2) has two key problems that make it highly unsuitable for recovery of sparse signals [57]:

[1] Because of the geometry of the problem, $\boldsymbol{\beta}_2$ is not very likely to be sparse.

[2] Little can be guaranteed in terms of whether $\boldsymbol{\beta}_2$ will be a good approximation to $\boldsymbol{\beta}$.

2.4 Sufficient Conditions for Reliable Reconstruction of Sparse Signals

From an implementation viewpoint, one cannot expect to measure a real-world signal $\boldsymbol{\beta}$ without any errors. Instead, a more plausible scenario is to assume that the observation vector $\boldsymbol{\nu}$ is corrupted by some additive noise

$$\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta} + \boldsymbol{\eta} \tag{2.11}$$

where $\boldsymbol{\eta} \in \mathbb{C}^n$ is either a deterministic (but unknown) perturbation, or it is a vector whose entries are i.i.d. realizations of some zero-mean random variable. This problem has been studied by a

number of researchers in the recent past [9, 65–70] and it turns out that the CS theory can be used in either case to obtain results that are in some sense parallel to those in the noiseless case. The only difference here being that the notion of exact recovery no longer applies—it is replaced by the notion of reliable reconstruction. Below, we briefly discuss some of what is currently known in the context of reliable reconstruction of sparse signals.

2.4.1 Reconstruction in the Presence of Bounded Noise

We begin by considering that the observation vector $\boldsymbol{\nu}$ is corrupted with a bounded perturbation vector $\boldsymbol{\eta}$: $\|\boldsymbol{\eta}\|_2 \leq \epsilon$ and study conditions under which $\boldsymbol{\beta}$ can be reliably reconstructed from $\boldsymbol{\nu}$. In this case, one may reconsider (P_0) and define an error-tolerant version of it as follows

$$\boldsymbol{\beta}_0 = \underset{\tilde{\boldsymbol{\beta}} \in \mathbb{C}^p}{\arg\min} \|\tilde{\boldsymbol{\beta}}\|_0 \quad \text{subject to} \quad \|\boldsymbol{\nu} - \mathbf{A}\tilde{\boldsymbol{\beta}}\|_2 \leq \epsilon\,. \tag{P_0^ϵ}$$

Loosely speaking, (P_0^ϵ) aims to do roughly the same thing as (P_0) would do on noiseless observations $\mathbf{A}\boldsymbol{\beta}$. Results establishing the stability and near-optimality of (P_0^ϵ) can be found in [66, 71].

Similar to the case of (P_0), however, (P_0^ϵ) is impractical to solve in general. Following the rationale of the previous section, we can instead replace the ℓ_0-norm in (P_0^ϵ) with the ℓ_1-norm and get the following error-tolerant variant of (BP)

$$\boldsymbol{\beta}_1 = \underset{\tilde{\boldsymbol{\beta}} \in \mathbb{C}^p}{\arg\min} \|\tilde{\boldsymbol{\beta}}\|_1 \quad \text{subject to} \quad \|\boldsymbol{\nu} - \mathbf{A}\tilde{\boldsymbol{\beta}}\|_2 \leq \epsilon\,. \tag{BPIC}$$

This optimization program, which we term as the *basis pursuit with inequality constraint*, is convex in nature and can be solved in a computationally tractable manner by recasting it as a linear optimization problem under quadratic inequality constraints [72]. Finally, the following theorem establishes that (BPIC) guarantees stable reconstruction of $\boldsymbol{\beta}$ from (2.11) in a deterministic (or adversarial) noise setting.

Theorem 2.12 (Noisy Reconstruction [63]) Let $\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta} + \boldsymbol{\eta}$ be an $n \times 1$ vector of observations of any deterministic but unknown signal $\boldsymbol{\beta} \in \mathbb{C}^p$, where the noise vector satisfies $\|\boldsymbol{\eta}\|_2 \leq \epsilon$. Assume that the columns of $\mathbf{A}$ have unit ℓ_2-norms and further let $\mathbf{A} \in RIP(2S, 0.3)$. Then the

vector $\boldsymbol{\beta}_1$ obtained as the solution of (BPIC) satisfies

$$\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}\|_2^2 \leq c_0 \left(c_0' \epsilon + \frac{\|\boldsymbol{\beta} - \boldsymbol{\beta}_S\|_1}{\sqrt{S}} \right)^2 \tag{2.12}$$

where c_0 and $\boldsymbol{\beta}_S$ are as defined earlier in Theorem 2.6, and $c_0' > 0$ is a constant given by

$$c_0' = 2\left(1 + \delta_{2S}\right)^{-\frac{1}{2}}. \tag{2.13}$$

2.4.2 Reconstruction in the Presence of Stochastic Noise

In many applications of practical interest, it is typically assumed that the observation vector $\boldsymbol{\nu}$ is corrupted by a stochastic noise vector $\boldsymbol{\eta}$ whose entries are i.i.d. realizations of a zero-mean, circularly complex, Gaussian random variable with variance σ^2. One of the first theoretical results in the (real) stochastic noise setting was established in [67] using an unconstrained error-tolerant version of (P_0), given by

$$\boldsymbol{\beta}_0 = \underset{\tilde{\boldsymbol{\beta}} \in \mathbb{C}^p}{\arg\min} \left(\frac{1}{2} \|\boldsymbol{\nu} - \mathbf{A}\tilde{\boldsymbol{\beta}}\|_2 + \lambda \|\tilde{\boldsymbol{\beta}}\|_0 \right) \tag{P_0^λ}$$

where the parameter $\lambda > 0$ is a function of p and σ^2. It is worth mentioning at this point that (P_0^λ) is the Lagrangian of (P_0^ϵ) and the two are related in the sense that any solution of (P_0^λ) for a particular λ corresponds to a solution of (P_0^ϵ) with an appropriate choice of ϵ. Strictly speaking, however, (P_0^ϵ) and (P_0^λ) are two distinct optimization programs.

Since (P_0^λ) requires solving a combinatorial program much like (P_0), a practical solution is to use an unconstrained error-tolerant version of (BP) by replacing ℓ_0 with the ℓ_1-norm in (P_0^λ)

$$\boldsymbol{\beta}_1 = \underset{\tilde{\boldsymbol{\beta}} \in \mathbb{C}^p}{\arg\min} \left(\frac{1}{2} \|\boldsymbol{\nu} - \mathbf{A}\tilde{\boldsymbol{\beta}}\|_2 + \lambda \|\tilde{\boldsymbol{\beta}}\|_1 \right). \tag{BPDN}$$

This optimization program goes by the name of *basis pursuit denoising* in the signal processing community [57], while it is known as *lasso* in the statistics literature [73]. The solution to (BPDN) can be found in a computationally tractable manner using standard convex optimization techniques since its objective is an unconstrained convex function [72]. Convex programs of the form (BPDN) have been extensively studied by researchers in the past in many different application areas [19, 57, 60, 73, 74]. However, very little attention has been paid in these and similarly related works

to develop a rigorous correspondence between (P_0^λ) and (BPDN) in the stochastic setting. In particular, while (BPDN) has been known to perform well in practice in a number of situations, results suggesting that (BPDN) gives reconstruction error bounds similar to those of (P_0^λ) have been reported only very recently in the literature [69, 70].

We now present another constrained optimization based method, which is in some sense related to (BPDN), for reliable reconstruction of $\boldsymbol{\beta}$ from (2.11) in the stochastic noise setting

$$\boldsymbol{\beta}_{1,\infty} = \underset{\tilde{\boldsymbol{\beta}} \in \mathbb{C}^p}{\arg\min} \|\tilde{\boldsymbol{\beta}}\|_1 \quad \text{subject to} \quad \|\mathbf{A}^{\mathsf{H}}(\boldsymbol{\nu} - \mathbf{A}\tilde{\boldsymbol{\beta}})\|_\infty \leq \lambda \,. \tag{DS}$$

This convex optimization program—which goes by the name of *Dantzig selector*—guarantees near-optimal reconstruction of $\boldsymbol{\beta}$ based on the RIP characterization of the sensing matrix [68]. Before stating the theoretical performance of (DS), however, it is worth pointing out the main reasons that make the Dantzig selector an integral part of our discussion on reliable reconstruction of sparse signals in the presence of stochastic noise:

[1] It is one of the few reconstruction methods in the CS literature that are guaranteed to perform near-optimally vis-à-vis stochastic noise—the others being (P_0^λ) and (BPDN).

[2] Unlike the combinatorial optimization program (P_0^λ), it is highly computationally tractable since it can be recast as a linear program.

[3] It comes with the cleanest and most interpretable reconstruction error bounds that we know for both sparse and approximately sparse signals.

Finally, note that some of the recent results in the literature seem to suggest that (BPDN) also enjoys many of the useful properties of (DS), including the reconstruction error bounds that appear very similar to those of (DS) [69, 70]. As such, making use of (BPDN) in practical settings can sometimes be more computationally attractive because of the availability of a wide range of efficient software packages, such as GPSR [75] and SpaRSA [76], for solving it. However, since a RIP-based characterization of (BPDN) that parallels that of (DS) does not exist to date, we limit ourselves in this chapter to discussing the results for (DS) only. The original specification of the

following theorem in [68] in this regard assumed a specific signal class, but the proof actually provides a more general oracle result.

Theorem 2.13 (The Dantzig Selector [68]) Let $\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta}+\boldsymbol{\eta}$ be an $n\times 1$ vector of observations of any deterministic but unknown signal $\boldsymbol{\beta} \in \mathbb{C}^p$, where the entries of $\boldsymbol{\eta}$ are independently distributed as $\mathcal{CN}(0, \sigma^2)$. Assume that the columns of $\mathbf{A}$ have unit ℓ_2-norms and further let $\mathbf{A} \in RIP(2S, 0.3)$ for some integer $S \geq 1$. Choose $\lambda = \sqrt{2\sigma^2(1+a)\log p}$ for any $a \geq 0$. Then the vector $\boldsymbol{\beta}_{1,\infty}$ obtained as the solution of (DS) satisfies

$$\|\boldsymbol{\beta}_{1,\infty} - \boldsymbol{\beta}\|_2^2 \leq c_0'' \min_{1\leq m\leq S} \left(\lambda\sqrt{m} + \frac{\|\boldsymbol{\beta} - \boldsymbol{\beta}_m\|_1}{\sqrt{m}}\right)^2 \tag{2.14}$$

with probability exceeding $1 - 2\left(\sqrt{\pi(1+a)\log p}\cdot p^a\right)^{-1}$. The constant $c_0'' = 16/\left(1-3\delta_{2S}\right)^2$, and as in Theorem 2.6, $\boldsymbol{\beta}_m$ is the vector formed by setting all but the m largest (in magnitude) entries of the true signal $\boldsymbol{\beta}$ to zero.

Remark 2.14 Notice that the reconstruction error in (2.14) is essentially comprised of two factors. One factor is due to the "estimation error" (or variance) that arises from determining m unknown quantities from noisy data, while the other is due to the "approximation error" (or bias) arising from estimating the unknown signal $\boldsymbol{\beta}$ using only m components. For a given signal class, the best rate of error decay is obtained by balancing the two terms. That is, the best choice of m is the value m_* such that

$$\|\boldsymbol{\beta} - \boldsymbol{\beta}_{m_*}\|_1 \approx \lambda m_*. \tag{2.15}$$

Thus, to make the optimal rates achievable, the sensing matrix should be chosen to satisfy RIP of order $2S$ such that S is at least as large as the "effective sparsity" m_*.

Note that Theorem 2.13 differs in two key respects from the results stated in [68] for the Dantzig selector. First, the probability of failure in Theorem 2.13 is twice the probability of failure obtained in [68]. This difference stems from the fact that the results in [68] are established only for the real-valued setup. In particular, [68, Section 3] proves for the case of $\eta_i \overset{i.i.d.}{\sim} \mathcal{N}(0,\sigma^2)$ that

$$\Pr\left(\boldsymbol{\beta}_{1,\infty} \text{ does not satisfy (2.14)}\right) \leq \Pr\left(\|\mathbf{A}^{\mathsf{T}}\boldsymbol{\eta}\|_\infty > \lambda\right) < \left(\sqrt{\pi(1+a)\log p}\cdot p^a\right)^{-1} \tag{2.16}$$

for the choice of λ in Theorem 2.13. The arguments underlying Theorem 2.13 for the complex case are almost the same as those for the real case. The only difference in the arguments in the complex case is due to the fact that $\eta_i \overset{i.i.d.}{\sim} \mathcal{CN}(0, \sigma^2)$, which results in

$$\Pr\left(\boldsymbol{\beta}_{1,\infty} \text{ does not satisfy (2.14)}\right) \leq \Pr\left(\|\mathbf{A}^{\mathsf{H}}\boldsymbol{\eta}\|_\infty > \lambda\right) < 2\left(\sqrt{\pi(1+a)\log p}\cdot p^a\right)^{-1}. \quad (2.17)$$

The second inequality in (2.17) is a consequence of the following lemma, proved in Section 2.5.3.

Lemma 2.15 Let $\mathbf{A}$ be an $n \times p$ matrix having unit ℓ_2-norm columns. Further, let $\boldsymbol{\eta}$ be an $n \times 1$ vector having entries independently distributed as $\mathcal{CN}(0, \sigma^2)$. Then for any $u > 0$

$$\Pr\left(\|\mathbf{A}^{\mathsf{H}}\boldsymbol{\eta}\|_\infty > \sigma u\right) < \frac{4p}{\sqrt{2\pi}} \cdot \frac{\exp(-u^2/2)}{u}. \quad (2.18)$$

Second, the sufficient condition stated in the original result in [68] for (DS) to succeed in reconstructing $\boldsymbol{\beta}$ is that $\mathbf{A} \in RIP(2S, \delta_{2S})$ such that $\delta_{2S} + \theta_{S,2S} < 1$, where $\theta_{S,2S}$ is called the $S, 2S$*-restricted orthogonality constant* (ROC) of $\mathbf{A}$. In general, the S, S'-ROC of $\mathbf{A}$ is defined as the smallest quantity such that

$$\left|\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}'\rangle\right| \leq \theta_{S,S'}\|\boldsymbol{\alpha}\|_2\|\boldsymbol{\alpha}'\|_2 \quad (2.19)$$

holds for all vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ having no more than S and S' nonzero entries, respectively, such that the nonzero entries of $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}'$ occur at disjoint indices. Nevertheless, the modified condition $\mathbf{A} \in RIP(2S, 0.3)$ stated in Theorem 2.13 is a simple consequence of the following lemma, which can be used to bound the $S, 2S$-ROC using δ_{2S}. The proof of this lemma appears in Section 2.5.4.

Lemma 2.16 Let $\mathbf{A}$ be an $n \times p$ matrix having unit ℓ_2-norm columns and assume without loss of generality that $S' \geq S$. Then the S, S'-ROC of $\mathbf{A}$ can be upper bounded as

$$\theta_{S,S'} \leq C\,\delta_{S+\lceil \frac{S'}{2} \rceil} \quad (2.20)$$

where $C = \sqrt{2}$ in a real-valued setup, while $C = 2$ in a complex-valued setup.

Finally, we conclude our review of CS by pointing out that Theorem 2.13 differs significantly from Theorem 2.12. Indeed, applying the deterministic noise results of Theorem 2.12 directly to

the stochastic noise setting (in which case $\|\boldsymbol{\eta}\|_2 \sim \sqrt{n}\sigma$) only guarantees that the resulting error scales like the *number of observations* times the noise power: $\|\boldsymbol{\beta}_1 - \boldsymbol{\beta}\|_2^2 = O(n\sigma^2)$. On the other hand, Theorem 2.13 results in a much better reconstruction error bound, with the error scaling like the *sparsity level* times the noise power. In other words, the estimation error bound of (DS) is adaptive to the sparsity level, while the error bound of (BPIC) is not. The difference in the two reconstruction error bounds could be significant, especially when the number of observations is far greater than the sparsity (or effective sparsity) of the signal.

2.5 Appendix

2.5.1 Proof of Theorem 2.1

Let $T \subset [1 \dots p]$ be a subset of cardinality $2S$ and assume that there exists an $n \times 2S$ submatrix $\mathbf{A}_T$ of $\mathbf{A}$ that does not have full column rank. Here, T corresponds to the indices of the columns of $\mathbf{A}$ that make up the submatrix $\mathbf{A}_T$. Note that the assumption $\mathsf{rank}(\mathbf{A}_T) < 2S$ means that there exists a $2S$-sparse vector $\boldsymbol{\beta}'$ such that $\mathbf{A}\boldsymbol{\beta}' = \mathbf{0}$ and $\{i : |\beta'_i| \neq 0\} = T$.

Next, partition T into two disjoint sets T_1 and T_2 of cardinality S each. That is, $T_1 \cup T_2 = T$, $T_1 \cap T_2 = \phi$, and $|T_1| = |T_2| = S$. Further, define S-sparse vectors $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ using the sets T_1 and T_2, respectively, as follows

$$\beta_{1,i} \stackrel{def}{=} \begin{cases} \beta'_i\,, & \text{if } i \in T_1; \\ 0\,, & \text{otherwise;} \end{cases} \qquad \text{and} \qquad \beta_{2,i} \stackrel{def}{=} \begin{cases} \beta'_i\,, & \text{if } i \in T_2; \\ 0\,, & \text{otherwise.} \end{cases} \tag{2.21}$$

It then follows from the definitions of $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ that

$$\mathbf{A}\boldsymbol{\beta}' = \mathbf{A}(\boldsymbol{\beta}_1 + \boldsymbol{\beta}_2) = \mathbf{0} \quad \Longrightarrow \quad \mathbf{A}\boldsymbol{\beta}_1 = \mathbf{A}\boldsymbol{\beta}'_2 \tag{2.22}$$

where $\boldsymbol{\beta}'_2 \stackrel{def}{=} -\boldsymbol{\beta}_2$ is also an S-sparse vector. The relation (2.22) shows that if any $n \times 2S$ submatrix of $\mathbf{A}$ does not have full column rank then there exist more than one S-sparse vector in $\mathbb{C}^p$ that get mapped to the same vector in $\mathbb{C}^n$ by the matrix $\mathbf{A}$. Therefore, every $n \times 2S$ submatrix of $\mathbf{A}$ must have full column rank to ensure unique recovery of any arbitrary S-sparse signal $\boldsymbol{\beta}$ from $\boldsymbol{\nu} = \mathbf{A}\boldsymbol{\beta}$. This completes the proof of the theorem. ∎

2.5.2 Proof of Theorem 2.4

We prove this theorem by contradiction. Suppose that $\mathbf{A}$ satisfies URP of order $2S$ but $\boldsymbol{\beta}_0 \neq \boldsymbol{\beta}$. This means that $\|\boldsymbol{\beta}_0\| \leq S$ (otherwise $\boldsymbol{\beta}_0$ cannot be a solution to (P_0)) and

$$\mathbf{A}\boldsymbol{\beta} = \mathbf{A}\boldsymbol{\beta}_0 \quad \Longrightarrow \quad \mathbf{A}\boldsymbol{\beta}' = \mathbf{0} \tag{2.23}$$

where $\boldsymbol{\beta}' \stackrel{def}{=} \boldsymbol{\beta} - \boldsymbol{\beta}_0$ is at most a $2S$-sparse vector (since both $\boldsymbol{\beta}$ and $\boldsymbol{\beta}_0$ are S-sparse vectors). Next, let $T \subset [1 \dots p]$ be such that $i \in T$ if and only if $|\beta'_i| \neq 0$ and define $\mathbf{A}_T$ to be a submatrix obtained by collecting all the columns of $\mathbf{A}$ corresponding to the indices in T. Note that $|T| \leq 2S$ and it is clear from the definition of URP that if $\mathbf{A}$ satisfies URP of order $2S$ then it also satisfies URP of order $|T|$. But we have from (2.23) that $\mathbf{A}_T$ has a nontrivial null space (since $\boldsymbol{\beta}' \neq 0$), which is a contradiction of the assumption that $\mathbf{A}$ satisfies URP of order $|T|$. Hence, $\boldsymbol{\beta}_0 = \boldsymbol{\beta}$ and this completes the proof of the theorem. ∎

2.5.3 Proof of Lemma 2.15

Assume without loss of generality that $\sigma = 1$, since the general case follows from a simple rescaling argument. Let $\mathbf{a}_1, \dots, \mathbf{a}_p \in \mathbb{C}^n$ be the p columns of $\mathbf{A}$ and define

$$z_i \stackrel{def}{=} \mathbf{a}_i^{\mathsf{H}} \boldsymbol{\eta}, \; i = 1, \dots, p. \tag{2.24}$$

Note that the z_i's are identically (but not independently) distributed as $z_i \sim \mathcal{CN}(0,1)$, which follows from the fact that $\eta_i \stackrel{i.i.d.}{\sim} \mathcal{CN}(0,1)$ and the columns of $\mathbf{A}$ have unit ℓ_2-norms. The rest of the proof is pretty elementary and follows from the facts that

$$\begin{aligned}
\Pr\left(\|\mathbf{A}^{\mathsf{H}}\boldsymbol{\eta}\|_\infty > u\right) &\stackrel{def}{=} \Pr\left(\max_{i=1,\dots,p} |z_i| > u\right) \\
&\stackrel{(a)}{\leq} p \cdot \Pr\left(|\mathrm{Re}(z_1)|^2 + |\mathrm{Im}(z_1)|^2 > u^2\right) \\
&\stackrel{(b)}{\leq} 2p \cdot \Pr\left(|\mathrm{Re}(z_1)| > \frac{u}{\sqrt{2}}\right) = 2p \cdot 2Q(u) \\
&\stackrel{(c)}{<} \frac{4p}{\sqrt{2\pi}} \cdot \frac{\exp(-u^2/2)}{u}\,.
\end{aligned} \tag{2.25}$$

Here, (a) follows by taking a union bound over the event $\bigcup_i\{|z_i| > u\}$, (b) follows from taking a union bound over the event $\{|\text{Re}(z_1)| > u/\sqrt{2}\} \cup \{|\text{Im}(z_1)| > u/\sqrt{2}\}$ and noting that the real and imaginary parts of z_i's are identically distributed as $\mathcal{N}(0, \frac{1}{2})$, and (c) follows by upper bounding the *complementary cumulative distribution function* as $Q(u) < \frac{1}{\sqrt{2\pi}u}\exp(-\frac{1}{2}u^2)$ [77]. ∎

2.5.4 Proof of Lemma 2.16

The proof of this lemma relies on the *polarization identity* [78], which expresses the inner product $\langle\cdot,\cdot\rangle$ in a vector space over a field $\mathbb{K}$ in terms of its induced norm $\|\mathbf{x}\| \stackrel{def}{=} \sqrt{\langle\mathbf{x},\mathbf{x}\rangle}$ as follows

$$\langle\mathbf{x},\mathbf{y}\rangle = \begin{cases} \frac{1}{4}\left(\|\mathbf{x}+\mathbf{y}\|^2 - \|\mathbf{x}-\mathbf{y}\|^2\right), & \mathbb{K} = \mathbb{R}\,, \\ \frac{1}{4}\Big[\left(\|\mathbf{x}+\mathbf{y}\|^2 - \|\mathbf{x}-\mathbf{y}\|^2\right) + j\left(\|\mathbf{x}+j\mathbf{y}\|^2 - \|\mathbf{x}-j\mathbf{y}\|^2\right)\Big], & \mathbb{K} = \mathbb{C}\,. \end{cases} \tag{2.26}$$

We begin by focussing on the case of $\mathbb{K} = \mathbb{C}$, since the proof for $\mathbb{K} = \mathbb{R}$ follows from similar arguments. Let $T \subset [1 \ldots p]$ be a subset of cardinality S' corresponding to the indices of nonzero entries of $\boldsymbol{\alpha}'$. Next, partition T into disjoint sets T_1 and T_2 of cardinality $\lceil\frac{S'}{2}\rceil$ and $(S' - \lceil\frac{S'}{2}\rceil)$, respectively. That is, $T_1 \cup T_2 = T$, $T_1 \cap T_2 = \phi$, and $|T_1| = \lceil\frac{S'}{2}\rceil$ and $|T_2| = (S' - \lceil\frac{S'}{2}\rceil)$. Further, define a $|T_1|$-sparse vector $\boldsymbol{\alpha}_1$ and a $|T_2|$-sparse vector $\boldsymbol{\alpha}_2$ as follows

$$\alpha_{1,i} \stackrel{def}{=} \begin{cases} \alpha'_i\,, & \text{if } i \in T_1; \\ 0\,, & \text{otherwise;} \end{cases} \quad \text{and} \quad \alpha_{2,i} \stackrel{def}{=} \begin{cases} \alpha'_i\,, & \text{if } i \in T_2; \\ 0\,, & \text{otherwise.} \end{cases} \tag{2.27}$$

It then follows from the triangle inequality and the definitions of $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$ that

$$\left|\langle\mathbf{A}\boldsymbol{\alpha},\mathbf{A}\boldsymbol{\alpha}'\rangle\right| \leq \left|\langle\mathbf{A}\boldsymbol{\alpha},\mathbf{A}\boldsymbol{\alpha}_1\rangle\right| + \left|\langle\mathbf{A}\boldsymbol{\alpha},\mathbf{A}\boldsymbol{\alpha}_2\rangle\right|. \tag{2.28}$$

Next, focus initially on $\left|\langle\mathbf{A}\boldsymbol{\alpha},\mathbf{A}\boldsymbol{\alpha}_1\rangle\right|$ and observe that because of the disjoint supports of $\boldsymbol{\alpha}$ and $\boldsymbol{\alpha}_1$, we have similar to the case in [68, Lemma 2.1]

$$(1-\delta_{S+|T_1|})(\|\boldsymbol{\alpha}\|_2^2 + \|\boldsymbol{\alpha}_1\|_2^2) \leq \|\mathbf{A}\boldsymbol{\alpha} \pm \mathbf{A}\boldsymbol{\alpha}_1\|_2^2 \leq (1+\delta_{S+|T_1|})(\|\boldsymbol{\alpha}\|_2^2 + \|\boldsymbol{\alpha}_1\|_2^2). \tag{2.29}$$

Paralleling the proof of [68, Lemma 2.1], we have from the polarization identity (2.26) and the above expression that

$$\begin{aligned}\Big|\mathrm{Re}\big(\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}_1\rangle\big)\Big| &= \frac{1}{4}\Big|\|\mathbf{A}\boldsymbol{\alpha} + \mathbf{A}\boldsymbol{\alpha}_1\|_2^2 - \|\mathbf{A}\boldsymbol{\alpha} - \mathbf{A}\boldsymbol{\alpha}_1\|_2^2\Big| \\ &\leq \frac{\delta_{S+|T_1|}}{2}\cdot(\|\boldsymbol{\alpha}\|_2^2 + \|\boldsymbol{\alpha}_1\|_2^2). \end{aligned} \tag{2.30}$$

Further, we can also get an expression similar to (2.29) for $\|\mathbf{A}\boldsymbol{\alpha} \pm j\mathbf{A}\boldsymbol{\alpha}_1\|_2^2$, which leads to the following upper bound due to (2.26)

$$\Big|\mathrm{Im}\big(\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}_1\rangle\big)\Big| \leq \frac{\delta_{S+|T_1|}}{2}\cdot(\|\boldsymbol{\alpha}\|_2^2 + \|\boldsymbol{\alpha}_1\|_2^2). \tag{2.31}$$

The two upper bounds (2.30) and (2.31) can now be combined together to yield

$$\begin{aligned}\big|\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}_1\rangle\big| &= \sqrt{\Big|\mathrm{Re}\big(\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}_1\rangle\big)\Big|^2 + \Big|\mathrm{Im}\big(\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}_1\rangle\big)\Big|^2} \\ &\leq \frac{\delta_{S+|T_1|}}{\sqrt{2}}\cdot(\|\boldsymbol{\alpha}\|_2^2 + \|\boldsymbol{\alpha}_1\|_2^2) \overset{(a)}{=} \sqrt{2}\delta_{S+|T_1|}\|\boldsymbol{\alpha}\|_2\|\boldsymbol{\alpha}_1\|_2 \end{aligned} \tag{2.32}$$

where (a) follows by noticing the simple fact that $\big|\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}_1\rangle\big| = \frac{\|\boldsymbol{\alpha}\|_2}{\|\boldsymbol{\alpha}_1\|_2}\big|\langle \frac{\|\boldsymbol{\alpha}_1\|_2}{\|\boldsymbol{\alpha}\|_2}\mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}_1\rangle\big|$. Similarly, by following identical arguments, it can be shown that

$$\big|\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}_2\rangle\big| \leq \sqrt{2}\delta_{S+|T_2|}\|\boldsymbol{\alpha}\|_2\|\boldsymbol{\alpha}_2\|_2. \tag{2.33}$$

Finally, we can upper bound $\big|\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}'\rangle\big|$ using (2.28), (2.32), and (2.33) as follows

$$\begin{aligned}\big|\langle \mathbf{A}\boldsymbol{\alpha}, \mathbf{A}\boldsymbol{\alpha}'\rangle\big| &\leq \sqrt{2}\delta_{S+|T_1|}\|\boldsymbol{\alpha}\|_2\|\boldsymbol{\alpha}_1\|_2 + \sqrt{2}\delta_{S+|T_2|}\|\boldsymbol{\alpha}\|_2\|\boldsymbol{\alpha}_2\|_2 \\ &\overset{(b)}{\leq} \sqrt{2}\delta_{S+|T_1|}\|\boldsymbol{\alpha}\|_2\Big(\|\boldsymbol{\alpha}_1\|_2 + \|\boldsymbol{\alpha}_2\|_2\Big) \\ &\overset{(c)}{\leq} 2\delta_{S+|T_1|}\|\boldsymbol{\alpha}\|_2\|\boldsymbol{\alpha}'\|_2. \end{aligned} \tag{2.34}$$

Here, (b) is a consequence of the fact that $\delta_{S+|T_2|} \leq \delta_{S+|T_1|}$ (since $|T_2| \leq |T_1|$) and (c) follows from the fact that $\|\boldsymbol{\alpha}_1\|_2 + \|\boldsymbol{\alpha}_2\|_2 \leq \sqrt{2}\|\boldsymbol{\alpha}'\|_2$ for $\boldsymbol{\alpha}_1$ and $\boldsymbol{\alpha}_2$ have disjoint supports. The lemma can now be established from the fact that, by definition, the S, S'-ROC of $\mathbf{A}$ is the smallest quantity that satisfies the last inequality in (2.34). ∎

Chapter 3

Compressed Sensing with Structured Matrices

3.1 Introduction

The field of sparse approximation—or compressed sensing (CS), as it is commonly called today—was arguably born almost 30 years ago out of the desire to solve underdetermined inverse problems in some application areas, such as seismic imaging [64], NMR spectroscopy [79], and array signal processing [54]. The apparent success of the applied community in solving these seemingly ill-posed problems spurred the research community's surge of interest in understanding the fundamental theoretical limits of these problems. And—coming back full circle to where we began—recent mathematical contributions of the research community (some of which were outlined in Chapter 2) have now inspired dozens of applied publications on reconstruction of sparse signals in various application areas—far too many to be listed here (see, e.g., [5, *Applications of Compressive Sensing*]).

Despite these advancements, however, a number of key technical challenges still need to be overcome in order to fully bridge the gap between theory and practice in many application areas of interest. In particular, given a linear or a nonlinear inverse problem involving sparse signals, the following two questions need to be satisfactorily answered before the practitioner can confidently make use of CS-based signal reconstruction methods:

[Q1] Can the inverse problem be transformed into the canonical CS observation model (2.2)?

[Q2] Can the ensuing sensing matrix be guaranteed to have the requisite conditions set forth in the CS literature for reliable signal reconstruction?

We focus exclusively on **[Q2]** in this part of the dissertation; we return to **[Q1]** in Chapter 4 and Chapter 5 of the dissertation in the context of estimation of sparse multipath channels and estimation of sparse networked data, respectively.

3.2 Structured Compressed Sensing Matrices

In many application areas, such as coding theory [6], imaging [47], and sensor networks [52], it turns out that transforming a given problem into the canonical CS setting requires the most amount of work. Once this transformation is carried out successfully, guaranteeing that the resulting sensing matrix satisfies conditions such as the restricted isometry property (RIP) is an easy consequence of fundamental results such as Theorem 2.8 and Theorem 2.9 in the CS literature. However, in a number of other application areas, such as linear system identification [14, 15], coded aperture imaging [17], and sampling theory [18], the sensing matrices tend to have a lot more structure to them due to the physics of the underlying problems. We use the term *structured compressed sensing matrices* for such matrices so as to distinguish them from the canonical (i.i.d. and subsampled unitary) CS matrices studied in Chapter 2.

The peculiar nature of structured CS matrices implies that existing results pertaining to i.i.d. and subsampled unitary matrices are not applicable in their case. In the past, researchers have often resorted to numerical simulations to prove the efficacy of structured CS matrices arising in various practical settings [13, 80, 81]. Nevertheless, rigorously proving the theoretical limits of structured CS matrices seems important for the credibility of the proposed research. It is in this context that we abstractly study three specific classes of structured CS matrices, namely, *Toeplitz matrices*, *Gabor matrices*, and *structurally-subsampled unitary matrices*, in this chapter. Collectively, these three matrix classes arise naturally in many application areas such as time-invariant and time-varying linear system identification [13–15], time-frequency analysis [16], coded aperture imaging [17], sampling theory [18], and radar and seismic imaging [19, 20], and our goal is to prove that sensing matrices belonging to these classes satisfy $RIP(2S, 0.3)$ with high probability. As shown in Chapter 2, this will be sufficient to guarantee reliable reconstruction of sparse signals using structured CS matrices belonging to the aforementioned matrix classes.

Before proceeding further, however, let us introduce some notation (originally used by Rudelson and Vershynin in [12]) that will greatly facilitate the mathematical analysis in the sequel. Recall from Chapter 2 that an $n \times p$ matrix $\mathbf{A} \in RIP(S, \delta_S)$ when the following inequality holds for some constant $\delta_S \in (0, 1)$

$$\max_{\substack{\|\tilde{\boldsymbol{\beta}}\|_0 \leq S \\ \tilde{\boldsymbol{\beta}} \neq \mathbf{0}}} \frac{\left| \|\mathbf{A}\tilde{\boldsymbol{\beta}}\|_2^2 - \|\tilde{\boldsymbol{\beta}}\|_2^2 \right|}{\|\tilde{\boldsymbol{\beta}}\|_2^2} \leq \delta_S \quad \Longleftrightarrow \quad \max_{\substack{\|\tilde{\boldsymbol{\beta}}\|_0 \leq S \\ \tilde{\boldsymbol{\beta}} \neq \mathbf{0}}} \frac{\left| \tilde{\boldsymbol{\beta}}^{\mathsf{H}} (\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p) \tilde{\boldsymbol{\beta}} \right|}{\tilde{\boldsymbol{\beta}}^{\mathsf{H}} \tilde{\boldsymbol{\beta}}} \leq \delta_S \,. \tag{3.1}$$

The expression on the right-hand side of (3.1) looks intriguingly similar to a bound on the spectral norm of the $p \times p$ matrix $(\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p)$, except that the maximum is taken over only a restricted subset of $\mathbb{C}^p$ [44]. Nevertheless, it is easy to see that

$$\max_{\substack{\|\tilde{\boldsymbol{\beta}}\|_0 \leq S \\ \tilde{\boldsymbol{\beta}} \neq \mathbf{0}}} \frac{\left| \tilde{\boldsymbol{\beta}}^{\mathsf{H}} (\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p) \tilde{\boldsymbol{\beta}} \right|}{\tilde{\boldsymbol{\beta}}^{\mathsf{H}} \tilde{\boldsymbol{\beta}}} \leq \delta_S \quad \Longleftrightarrow \quad \max_{\substack{T \subset [1 \ldots p] \\ |T| \leq S}} \left\| \mathbf{A}_T^{\mathsf{H}} \mathbf{A}_T - \mathbf{I}_{|T|} \right\|_2 \leq \delta_S \tag{3.2}$$

where $\mathbf{A}_T$ denotes an $n \times |T|$ submatrix of $\mathbf{A}$ obtained by collecting all the columns of $\mathbf{A}$ corresponding to the indices in set T. We can write this expression in a compact form with the help of a non-negative function $\|\cdot\|_{T,S} : \mathbb{C}^{p \times p} \to [0, \infty)$ defined as follows

$$\|\mathbf{M}\|_{T,S} \stackrel{def}{=} \max_{\substack{T \subset [1 \ldots p] \\ |T| \leq S}} \left\| \mathbf{M}_{T \times T} \right\|_2 \tag{3.3}$$

where $\mathbf{M}_{T \times T}$ denotes a $|T| \times |T|$ submatrix of $\mathbf{M}$ obtained by collecting all the entries of $\mathbf{M}$ corresponding to the indices in set $T \times T$. Going back to (3.2), we can alternatively say that an $n \times p$ matrix $\mathbf{A} \in RIP(S, \delta_S)$ for some constant $\delta_S \in (0, 1)$ when

$$\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} \leq \delta_S \tag{3.4}$$

and we will strive to prove this inequality in the sequel for the three matrix classes. Finally, we conclude this section with a lemma that will be extremely helpful in proving the RIP for structurally-subsampled unitary matrices. The proof of the following lemma is a trivial consequence of the definition of $\|\cdot\|_{T,S}$ and is therefore omitted here.

Lemma 3.1 The function $\|\cdot\|_{T,S} : \mathbb{C}^{p \times p} \to [0, \infty)$ defines a norm—which we term as (T, S)-norm—on the vector space $\mathbb{C}^{p \times p}$. Hence, $\mathcal{B} \stackrel{def}{=} (\mathbb{C}^{p \times p}, \|\cdot\|_{T,S})$ is a Banach space.

3.3 On the RIP of Toeplitz Matrices

We begin our discussion of structured CS matrices by first studying Toeplitz matrices. Recall that a Toeplitz matrix is a matrix in which every (left to right) diagonal is constant. Therefore, Toeplitz matrices are completely specified by the entries in their first rows and first columns. Since convolution between two discrete-time sequences can be construed as a matrix-vector multiplication, with one of the sequences converted into a Toeplitz matrix, such matrices frequently arise in applications involving linear, time-invariant systems [82].

In this section, we primarily focus on two—somewhat related, but still distinct—forms of Toeplitz matrices. Specifically, let $\mathcal{A}_k \stackrel{def}{=} \{a_i \in \mathbb{C}\}_{i=1}^{k}$ denote a k-length *generating sequence* for some $k \in \mathbb{N}$. The first form of Toeplitz matrices considered in this section, which we term as "full" Toeplitz matrices, is generated from the sequence $\mathcal{A}_k$ for $k \geq 1$ as follows

$$\mathbf{A} = \begin{bmatrix} a_1 & & 0 \\ a_2 & \ddots & \\ \vdots & \ddots & a_1 \\ a_k & & a_2 \\ & \ddots & \vdots \\ 0 & & a_k \end{bmatrix}. \tag{3.5}$$

Here, the full Toeplitz matrix $\mathbf{A}$ has dimensions $n \times p$ such that $n \stackrel{def}{=} k + (p-1)$. The second form of Toeplitz matrices considered in this section, which we term as "partial" Toeplitz matrices, more closely resembles the canonical underdetermined setting in the CS literature and is generated from the sequence $\mathcal{A}_k$ for $k \geq p$ as follows

$$\mathbf{A} = \begin{bmatrix} a_p & a_{p-1} & \dots & a_2 & a_1 \\ a_{p+1} & a_p & \dots & a_3 & a_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ a_k & a_{k-1} & \dots & a_{n+1} & a_n \end{bmatrix}. \tag{3.6}$$

The partial Toeplitz matrix in the above expression has dimensions $n \times p$ such that $n \stackrel{def}{=} k - (p-1)$. Notice that when $k \geq p$, the partial Toeplitz matrix described in (3.6) above is a submatrix of the

full Toeplitz matrix in (3.5). In contrast, when $k < p$, every row of the full Toeplitz matrix has at least one zero entry, and it is just a scaled version of $\mathbf{I}_p$ in the limiting case of $k = 1$.

The main question we address in this section is whether full and partial Toeplitz matrices generated from a *Rademacher sequence* $\mathcal{A}_k$ satisfy RIP. Here, we term a sequence as a Rademacher sequence if its elements independently take the values $\pm c$ with probability $1/2$ for some quantity $c > 0$ (in other words, if the elements of the sequence are independent symmetric Bernoulli or Rademacher random variables). Note that initial results in the CS literature that considered random sensing matrices either required statistical independence among the entries of $\mathbf{A}$ [11] or—at the very least—required statistical independence among the rows of $\mathbf{A}$ [12, 83]. The problem considered in this section, however, is significantly more challenging since the Toeplitz structure in (3.5) and (3.6) introduces statistical dependence among the rows of $\mathbf{A}$ and hence, existing techniques can no longer be used to address this problem. Instead, we develop a novel technique in the sequel that facilitates analysis in the presence of such dependencies among the rows of a sensing matrix.

3.3.1 Main Results

Before establishing the main claims of this section—that (full or partial) Toeplitz matrices with entries drawn independently from a Rademacher distribution satisfy RIP—we recall the simplified definition of RIP from (3.4) of Section 3.2

$$\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} = \max_{\substack{T \subset [1 \ldots p] \\ |T| \leq S}} \left\|\mathbf{A}_T^{\mathsf{H}}\mathbf{A}_T - \mathbf{I}_{|T|}\right\|_2 \leq \delta_S. \tag{3.7}$$

In other words, to establish RIP for a given sensing matrix $\mathbf{A}$, one needs to bound the spectral norms of all square submatrices of $(\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p)$ having no more than S rows/columns. Trivially from the definition of the spectral norm, however, we have that

$$\left\|\mathbf{A}_{T'}^{\mathsf{H}}\mathbf{A}_{T'} - \mathbf{I}_{|T'|}\right\|_2 \leq \left\|\mathbf{A}_T^{\mathsf{H}}\mathbf{A}_T - \mathbf{I}_{|T|}\right\|_2 \tag{3.8}$$

for any $T' \subset T$. Therefore, we only need to bound the spectral norms of all $S \times S$ submatrices of $(\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p)$. We now state one of the most useful and easily applied tools—known as *Geršgorin's disc theorem*—that can give bounds for the spectral norm of a matrix. The following result—stated

as a lemma without proof—seems to have first appeared in 1931 in [84] and its proof can be found in any standard text on matrix analysis such as [44].

Lemma 3.2 (Geršgorin) Let $\mathbf{M} \in \mathbb{C}^{S\times S}$ and $m_{i,j}, i, j = 1, \ldots, S$, denote the entries of $\mathbf{M}$. Then every eigenvalue of $\mathbf{M}$ lies in at least one of the S discs defined below

$$\mathcal{D}_i(\mathbf{M}) \stackrel{def}{=} \Big\{ z \in \mathbb{C} : |z - m_{i,i}| \leq \sum_{\substack{j=1 \\ j\neq i}}^{S} |m_{i,j}| \Big\},\ i = 1, \ldots, S. \tag{3.9}$$

The Geršgorin's disc theorem allows us to prove a very powerful auxiliary lemma that relates the (T, S)-norm of $(\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p)$ to $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$ as follows.

Lemma 3.3 Let $\mathbf{A}$ be an $n \times p$ matrix having unit ℓ_2-norm columns. Then

$$\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} \leq (S-1)\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}. \tag{3.10}$$

Proof: To prove this lemma, consider any arbitrary $T \subset [1 \ldots p]$ of cardinality $|T| = S$ and let $\mathbf{G}(T) \stackrel{def}{=} (\mathbf{A}_T^{\mathsf{H}}\mathbf{A}_T - \mathbf{I}_S)$. It is easy to see that each entry $g_{i,j}(T), i, j = 1, \ldots, S$, of $\mathbf{G}(T)$ corresponds to one of the entries in the *hollow* Gram matrix $\mathbf{G} \stackrel{def}{=} (\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p)$. Therefore, we have the trivial inequality

$$\|\mathbf{G}(T)\|_{\max} \leq \|\mathbf{G}\|_{\max}. \tag{3.11}$$

Further, we have that the main-diagonal entries $g_{i,i}(T), i = 1, \ldots, S$, of $\mathbf{G}(T)$ are zero since the main-diagonal entries $g_{i,i}, i = 1, \ldots, p$, of $\mathbf{G}$ are zero—a simple consequence of the fact that each column of $\mathbf{A}$ has unit ℓ_2-norm. Therefore, the centers of all Geršgorin's discs associated with $\mathbf{G}(T)$ are zero and, from (3.11), their radii are upperbounded by $(S-1)\|\mathbf{G}\|_{\max}$. Finally, note that this assertion is true regardless of the choice of T. That is,[1]

$$\mathcal{D}_i(\mathbf{G}(T)) \subset \{x \in \mathbb{R} : |x| \leq (S-1)\|\mathbf{G}\|_{\max}\},\ i = 1, \ldots, S,\ \forall\, T \tag{3.12}$$

which implies that $\|\mathbf{G}(T)\|_2 \leq (S-1)\|\mathbf{G}\|_{\max}\ \forall\, T$. This completes the proof of the lemma since this leads to $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} = \max_T \|\mathbf{G}(T)\|_2 \leq (S-1)\|\mathbf{G}\|_{\max}$. ∎

[1] Note that Geršgorin's discs in the complex plane are in fact *Geršgorin's intervals* on the real line in this case since we are dealing with Hermitian matrices.

An immediate consequence of Lemma 3.3 is that it allows us to prove that a given (appropriately normalized) sensing matrix $\mathbf{A} \in RIP(S, \delta_S)$ by showing that $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$ is close to δ_S/S (with high probability). Next, we provide a lemma that will be helpful in the proofs in describing the tail behavior of the distribution of $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$. The following result—known as *Hoeffding's inequality*—first appeared in this form in 1963 in [85].

Lemma 3.4 (Hoeffding's Inequality) Let $x_1, x_2, \ldots, x_N$ be (real-valued) independent bounded random variables satisfying $a_i \leq x_i \leq b_i, i = 1, \ldots, N$, almost surely. Define $S_N = \sum_{i=1}^{N} x_i$, then for any $t > 0$

$$\Pr\left(\left|S_N - \mathbb{E}[S_N]\right| \geq t\right) \leq 2\exp\left(-\frac{2t^2}{\sum_{i=1}^{N}(b_i - a_i)^2}\right). \tag{3.13}$$

We are now ready to establish RIP for the full Toeplitz matrices described in (3.5) when the generating sequence $\mathcal{A}_k$ is a Rademacher sequence.

Theorem 3.5 Let the elements of the generating sequence $\mathcal{A}_k = \{a_i\}_{i=1}^{k}$ be i.i.d. realizations of Rademacher random variables taking values $\pm 1/\sqrt{k}$ with probability $1/2$. Let $\mathbf{A}$—as defined in (3.5)—be the $n \times p$ full Toeplitz matrix generated by the sequence $\mathcal{A}_k$, where $n \stackrel{def}{=} k + (p-1)$. Then for any $\delta_S \in (0, 1)$, there exist constants c_1 and c_2 depending only on δ_S such that whenever $k \geq c_2 S^2 \log p$, $\mathbf{A} \in RIP(S, \delta_S)$ with probability exceeding $1 - \exp(-c_1 k/S^2)$.

Proof: Trivially, since the columns of $\mathbf{A}$ have—by construction—unit ℓ_2-norms, we have from Lemma 3.3 that $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} \leq (S-1)\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$. Therefore, we only need to study $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$ in order to prove that $\mathbf{A} \in RIP(S, \delta_S)$. Next, note that

$$\begin{aligned}\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} &= \max_{i \neq j} \left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right| \\ &\stackrel{(a)}{=} \max_{i < j} \left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right| \end{aligned} \tag{3.14}$$

where $\mathbf{a}_i \in \mathbb{C}^n$ denotes the i-th column of $\mathbf{A}$ and (a) follows since $\left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right| = \left|\langle \mathbf{a}_i, \mathbf{a}_j \rangle\right|$. We can explicitly write an expression for the inner product $\langle \mathbf{a}_j, \mathbf{a}_i \rangle$, assuming $1 \leq i < j \leq p$, as

$$\langle \mathbf{a}_j, \mathbf{a}_i \rangle = \sum_{q=1}^{k-\Delta} a_q a_{q+\Delta} \tag{3.15}$$

where $\Delta \stackrel{def}{=} (j-i)$. Observe that $\left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right|$ cannot be bounded through the use of Hoeffding's inequality since the terms in the above sum are not mutually independent. For example, consider the case of $i = 1, j = 2$, and $k = 5$. Then $\langle \mathbf{a}_2, \mathbf{a}_1 \rangle = a_1 a_2 + a_2 a_3 + a_3 a_4 + a_4 a_5$, and the first two terms are dependent (due to a_1), as are the second and third terms (due to a_2), etc. Notice, however, that the first and third terms and the second and fourth terms are independent, which suggests that the entire sum can be written as two sums of mutually independent terms.

We now prove that this is true in general. That is, $\langle \mathbf{a}_j, \mathbf{a}_i \rangle$ for any $i < j$ can always be written as two sums having mutually independent terms. To establish this claim, rearrange the summands in (3.15) as follows

$$\begin{aligned} \langle \mathbf{a}_j, \mathbf{a}_i \rangle &= \sum_{q=1}^{\Delta} \sum_{r=1}^{\lfloor \frac{k-q}{\Delta} \rfloor} a_{q+(r-1)\Delta}\, a_{q+r\Delta} \\ &= S_{1,\Delta} + S_{2,\Delta} + \cdots + S_{\Delta,\Delta} \end{aligned} \tag{3.16}$$

where $S_{q,\Delta} \stackrel{def}{=} \sum_{r=1}^{\lfloor \frac{k-q}{\Delta} \rfloor} a_{q+(r-1)\Delta}\, a_{q+r\Delta}$, $q = 1, \ldots, \Delta$. Notice that (i) all $S_{q,\Delta}$'s in the above expression are mutually independent since every element of the generating sequence $\mathcal{A}_k$ only appears in at most one of the $S_{q,\Delta}$'s, and (ii) every term within any $S_{q,\Delta}$ is only dependent with its adjacent terms in the same $S_{q,\Delta}$. Consequently, if we index the individual summands in (3.16) from 1 to $(k-\Delta)$ then it is easy to see that all the odd-indexed terms in (3.16) will be mutually independent, as will be the even-indexed ones. Finally, partitioning the indexed sum (3.16) into odd- and even-indexed terms, followed by a reindexing yields

$$\langle \mathbf{a}_j, \mathbf{a}_i \rangle = \underbrace{\sum_{q_1=1}^{\lceil \frac{k-\Delta}{2} \rceil} a'_{q_1}}_{S_{k_1,\Delta}} + \underbrace{\sum_{q_2=1}^{\lfloor \frac{k-\Delta}{2} \rfloor} a'_{q_2}}_{S_{k_2,\Delta}} \tag{3.17}$$

where $\{a'_{q_1}\}$ and $\{a'_{q_2}\}$ consist of $k_1 \stackrel{def}{=} \max\left\{0, \lceil \frac{k-\Delta}{2} \rceil\right\}$ and $k_2 \stackrel{def}{=} \max\left\{0, \lfloor \frac{k-\Delta}{2} \rfloor\right\}$ i.i.d. Rademacher random variables, respectively, that take values $\pm 1/k$ with probability $1/2$.

We are now ready to bound $\left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right|$ by applying Lemma 3.4 to its components sums $S_{k_1,\Delta}$ and $S_{k_2,\Delta}$ having k_1 and k_2 terms, respectively, as follows

$$\begin{aligned}\Pr\left(\left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right| > \frac{\delta_S}{S}\right) &\leq \Pr\left(\left\{|S_{k_1,\Delta}| > \frac{\delta_S}{2S}\right\} \cup \left\{|S_{k_2,\Delta}| > \frac{\delta_S}{2S}\right\}\right) \\ &\overset{(b)}{\leq} 2\max\left\{\Pr\left(|S_{k_1,\Delta}| > \frac{\delta_S}{2S}\right), \Pr\left(|S_{k_2,\Delta}| > \frac{\delta_S}{2S}\right)\right\} \\ &\overset{(c)}{\leq} 2\max\left\{2\exp\left(-\frac{k^2\delta_S^2}{8k_1S^2}\right), 2\exp\left(-\frac{k^2\delta_S^2}{8k_2S^2}\right)\right\}. \end{aligned} \tag{3.18}$$

Here, (b) follows from a simple union bounding argument and (c) follows from Lemma 3.4. Next, notice from the definitions of k_1 and k_2 that $k/2 \geq k_1 \geq k_2$ for any $\Delta \geq 1$. Therefore, we can further simplify (3.18) as follows

$$\Pr\left(\left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right| > \frac{\delta_S}{S}\right) \leq 4\exp\left(-\frac{k\delta_S^2}{4S^2}\right). \tag{3.19}$$

We have now established that $\left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right| \leq \delta_S/S$ with probability exceeding $1 - 4\exp(-\frac{k\delta_S^2}{4S^2})$ for any $i < j$. To prove that $\mathbf{A} \in RIP(S, \delta_S)$, however, we need to evaluate the probability that $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} \leq \delta_S/S$. To this end, we make use of (3.14) and obtain

$$\begin{aligned}\Pr\left(\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} > \frac{\delta_S}{S}\right) &= \Pr\left(\max_{i<j}\left|\langle \mathbf{a}_j, \mathbf{a}_i \rangle\right| > \frac{\delta_S}{S}\right) \\ &\overset{(d)}{\leq} 2p(p-1)\exp\left(-\frac{k\delta_S^2}{4S^2}\right) \end{aligned} \tag{3.20}$$

where (d) follows from taking union bound over a total of $p(p-1)/2$ events. This implies that for $c_1 < \delta_S^2/4$ and $k \geq (12/(\delta_S^2 - 4c_1))S^2 \log p$, we have $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} \leq (S-1)\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} < \delta_S$ with probability exceeding $1 - \exp(-c_1 k/S^2)$, which is what we needed to show. ∎

Given the structural similarities between full and partial Toeplitz matrices, the tools and techniques used in the proof of the previous theorem can also be applied in the case of partial Toeplitz matrices. This leads us to the second main result of this section.

Theorem 3.6 Let the elements of the generating sequence $\mathcal{A}_k = \{a_i\}_{i=1}^k$ be i.i.d. realizations of Rademacher random variables taking values $\pm 1/\sqrt{n}$ with probability $1/2$. Let $\mathbf{A}$—as defined in (3.6)—be the $n \times p$ partial Toeplitz matrix generated by the sequence $\mathcal{A}_k$, where $n \overset{def}{=} k - (p-1)$.

Then for any $\delta_S \in (0,1)$, there exist constants c_1' and c_2' depending only on δ_S such that whenever $n \geq c_2' S^2 \log p$, $\mathbf{A} \in RIP(S, \delta_S)$ with probability exceeding $1 - \exp(-c_1' n/S^2)$.

Proof: The proof of this theorem proceeds in a similar fashion to the proof of Theorem 3.5. In particular, the inner product $\langle \mathbf{a}_j, \mathbf{a}_i \rangle$, assuming $1 \leq i < j \leq p$, in this case can be expressed as

$$\langle \mathbf{a}_j, \mathbf{a}_i \rangle = \sum_{q=1}^{n} a_{q+\Delta_1} \, a_{q+\Delta_1+\Delta_2} \tag{3.21}$$

where $\Delta_1 \stackrel{def}{=} (p-j)$ and $\Delta_2 \stackrel{def}{=} (j-i)$. Further, (3.21) can also be shown to consist of two sums having mutually independent terms by rearranging the summands in the expression as follows

$$\begin{aligned} \langle \mathbf{a}_j, \mathbf{a}_i \rangle &= \sum_{q=1}^{\Delta_2} \sum_{r=1}^{\left\lfloor \frac{n+\Delta_2-q}{\Delta_2} \right\rfloor} a_{q+\Delta_1+(r-1)\Delta_2} \, a_{q+\Delta_1+r\Delta_2} \\ &= S_{1,i,j} + S_{2,i,j} + \cdots + S_{\Delta_2,i,j} \end{aligned} \tag{3.22}$$

where $S_{q,i,j} \stackrel{def}{=} \sum_{r=1}^{\left\lfloor \frac{n+\Delta_2-q}{\Delta_2} \right\rfloor} a_{q+\Delta_1+(r-1)\Delta_2} \, a_{q+\Delta_1+r\Delta_2}, q = 1, \ldots, \Delta_2$. It is easy to see in this case too that if we index the summands in (3.22) from 1 to n then all the odd-indexed terms in the resulting sum will be mutually independent, as will be the even-indexed ones. Finally, partitioning the indexed sum (3.22) into odd- and even-indexed terms yields

$$\langle \mathbf{a}_j, \mathbf{a}_i \rangle = \underbrace{\sum_{q_1=1}^{\lceil \frac{n}{2} \rceil} a'_{q_1}}_{S_{n_1,i,j}} + \underbrace{\sum_{q_2=1}^{\lfloor \frac{n}{2} \rfloor} a'_{q_2}}_{S_{n_2,i,j}} \tag{3.23}$$

where $\{a'_{q_1}\}$ and $\{a'_{q_2}\}$ in this case consist of $n_1 \stackrel{def}{=} \lceil \frac{n}{2} \rceil$ and $n_2 \stackrel{def}{=} \lfloor \frac{n}{2} \rfloor$ i.i.d. Rademacher random variables, respectively, that take values $\pm 1/n$ with probability $1/2$.

We can now bound $\left| \langle \mathbf{a}_j, \mathbf{a}_i \rangle \right|$ by once again applying Lemma 3.4 to its components sums $S_{n_1,i,j}$ and $S_{n_2,i,j}$ as follows

$$\begin{aligned} \Pr\left(\left| \langle \mathbf{a}_j, \mathbf{a}_i \rangle \right| > \frac{\delta_S}{S} \right) &\leq 2 \max\left\{ \Pr\left(|S_{n_1,i,j}| > \frac{\delta_S}{2S} \right), \Pr\left(|S_{n_1,i,j}| > \frac{\delta_S}{2S} \right) \right\} \\ &\leq 2 \max\left\{ 2\exp\left(-\frac{n^2 \delta_S^2}{8 n_1 S^2} \right), 2\exp\left(-\frac{n^2 \delta_S^2}{8 n_2 S^2} \right) \right\}. \end{aligned} \tag{3.24}$$

Finally, notice from the definitions of n_1 and n_2 that $n > n_1 \geq n_2$. Therefore, we can further simplify (3.24) and obtain the following probabilistic bound for $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$

$$\begin{aligned}\Pr\left(\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} > \frac{\delta_S}{S}\right) &= \Pr\left(\max_{i<j} \left|\langle \mathbf{a}_j, \mathbf{a}_i\rangle\right| > \frac{\delta_S}{S}\right) \\ &\leq 2p(p-1)\exp\left(-\frac{n\delta_S^2}{8S^2}\right). \end{aligned} \tag{3.25}$$

The above expression implies that for $c_1' < \delta_S^2/8$ and $n \geq (24/(\delta_S^2 - 8c_1'))S^2 \log p$, we have $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} \leq (S-1)\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} < \delta_S$ with probability exceeding $1 - \exp(-c_1' n/S^2)$, since $\mathbf{A}$ has unit ℓ_2-norm columns. This completes the proof of the theorem. ∎

3.4 On the RIP of Gabor Matrices

In this section, we continue our discussion of structured CS matrices by focusing on the class of Gabor matrices, which is a natural extension of the class of Toeplitz matrices. Specifically, let $\mathcal{A}_k \stackrel{def}{=} \{a_i \in \mathbb{C}\}_{i=1}^k$ denote a k-length generating sequence for some $k \in \mathbb{N}$. Further, let $L \geq 1$ and $M \geq 0$ be two integer parameters such that $M \leq (k-1)/2$, and define $\mathbf{T}$ to be an $n \times L$ Toeplitz matrix that is generated from the sequence $\mathcal{A}_k$ as follows

$$\mathbf{T} = \begin{bmatrix} a_1 & & 0 \\ a_2 & \ddots & \\ \vdots & \ddots & a_1 \\ a_k & & a_2 \\ & \ddots & \vdots \\ 0 & & a_k \end{bmatrix} \tag{3.26}$$

where $n \stackrel{def}{=} k + (L-1)$. Finally, let $\boldsymbol{\omega}_m \stackrel{def}{=} \begin{bmatrix} e^{j0\omega_{m,k}} & \dots & e^{j(n-1)\omega_{m,k}} \end{bmatrix}^{\mathsf{T}}$ denote the collection of n samples of a discrete sinusoid with frequency $\omega_{m,k} \stackrel{def}{=} 2\pi\frac{m}{k}, m \in [-M \dots M]$, and define corresponding $n \times n$ diagonal *modulation matrices* as $\mathbf{W}_m = \mathsf{diag}(\boldsymbol{\omega}_m)$. Then the "full" Gabor matrix generated from the sequence $\mathcal{A}_k$ is a block matrix of the form

$$\mathbf{A} = \begin{bmatrix} \mathbf{W}_{-M}\mathbf{T} & \dots & \mathbf{W}_0\mathbf{T} & \dots & \mathbf{W}_M\mathbf{T} \end{bmatrix}. \tag{3.27}$$

Here, the Gabor matrix $\mathbf{A}$—or Gabor system matrix, as it is sometimes called in the literature [86]—has dimensions $n \times p$ such that $p \stackrel{def}{=} L(2M+1)$. In words, whereas the columns of a Toeplitz matrix are given by downward linear shifts of $\mathcal{A}_k$, the columns of a Gabor matrix are given by downward linear shifts *and modulations* (frequency shifts) of the generating sequence. In particular, it is trivial to see from (3.5) and (3.27) that a Gabor matrix reduces to a Toeplitz matrix for the case of $M = 0$.

Gabor matrices of the form (3.27) naturally arise in applications involving linear, time-varying systems, and are frequently encountered by researchers working in the areas of communications, signal and image processing, and optics [87]. Similar to the case of Toeplitz matrices, the main question we address in this section is whether Gabor matrices generated from a Rademacher sequence $\mathcal{A}_k$ satisfy RIP. In particular, note that since a Gabor matrix generated from a random sequence has statistical dependence both among *and within* its rows (unless $M = 0$), it is quite reasonable to expect—when compared with Toeplitz matrices—a stricter requirement on k for Gabor matrices to satisfy RIP. Using some of the tools described in the last section, however, we prove in the sequel that Gabor matrices also only require $k = \Omega(S^2 \log p)$ in order to satisfy RIP of order S, despite the fact that they tend to have more structural dependencies.

Remark 3.7 It is instructive to note here that—unlike the case of full Toeplitz matrices—Gabor matrices can be either underdetermined or overdetermined depending on the choice of the parameters k (the length of the generating sequence), L (the number of shifts), and M (the number of one-sided modulations). In particular, for the special case of $M = (k-1)/2$ and $k = L > 1$, we have that the $n \times p$ Gabor matrices are highly underdetermined with $n \approx \sqrt{p}/2$.

3.4.1 Main Result

The main tool that we will be using to establish the claim that a Gabor matrix $\mathbf{A}$ with entries drawn independently from a Rademacher distribution satisfies RIP is again Lemma 3.3. Before proceeding with a formal proof of the claim, however, it is helpful to state a complex version of Hoeffding's inequality that will be used in the proof to describe the tail behavior of the distribution

of $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$. The following lemma is an easy consequence of the original Hoeffding's inequality and its proof is given in Section 3.7.1.

Lemma 3.8 (Complex Hoeffding's Inequality) Let $x_1, x_2, \ldots, x_N$ be complex-valued independent random variables satisfying $|x_i| \leq a_i, i = 1, \ldots, N$, almost surely. Define $S_N = \sum_{i=1}^{N} x_i$, then for any $t > 0$

$$\Pr\left(\left|S_N - \mathbb{E}[S_N]\right| \geq t\right) \leq 4\exp\left(-\frac{t^2}{4\sum_{i=1}^{N} a_i^2}\right). \tag{3.28}$$

We are now ready to establish RIP for Gabor matrices of the form (3.27) when the generating sequence $\mathcal{A}_k$ is a Rademacher sequence.

Theorem 3.9 Let the elements of the generating sequence $\mathcal{A}_k = \{a_i\}_{i=1}^{k}$ be i.i.d. realizations of Rademacher random variables taking values $\pm 1/\sqrt{k}$ with probability $1/2$. Let $\mathbf{A}$—as defined in (3.27)—be the $n \times p$ Gabor matrix generated by the sequence $\mathcal{A}_k$, where $n \stackrel{def}{=} k + (L-1)$ and $p \stackrel{def}{=} L(2M+1)$ for integers $L \geq 1$ and $M \geq 0$ such that $M \leq (k-1)/2$. Then for any $\delta_S \in (0,1)$, there exist constants c_3 and c_4 depending only on δ_S such that whenever $k \geq c_4 S^2 \log p$, $\mathbf{A} \in RIP(S, \delta_S)$ with probability exceeding $1 - \exp(-c_3 k/S^2)$.

Proof: In the following, we use $\mathbf{a}_{m,\ell} \in \mathbb{C}^n$ to denote the ℓ-th column of the m-th block, $\mathbf{W}_m\mathbf{T}$, of the Gabor matrix $\mathbf{A}$, where $-M \leq m \leq M$ and $1 \leq \ell \leq L$. It can be easily seen from the definition of $\mathbf{A}$ in (3.27) that

$$\|\mathbf{a}_{m,\ell}\|_2^2 = \sum_{q=1}^{k} \left|a_q\, e^{j(\ell+q-2)\omega_{m,k}}\right|^2 = 1. \tag{3.29}$$

Therefore, we have from Lemma 3.3 that $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} \leq (S-1)\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$ and—similar to the case of Theorem 3.5—we only need to study

$$\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} = \max_{(m,\ell)\neq(m',\ell')} \left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| \tag{3.30}$$

in order to prove that $\mathbf{A} \in RIP(S, \delta_S)$. Further, note that since $\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| = \left|\langle \mathbf{a}_{m,\ell}, \mathbf{a}_{m',\ell'}\rangle\right|$, the following relationship also holds trivially

$$\left\{\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| : (m,\ell) \neq (m',\ell')\right\} = \left\{\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| : m \neq m', \ell = \ell'\right\} \bigcup \left\{\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| : \ell < \ell'\right\}. \tag{3.31}$$

The above expression suggests that there are two main cases that we need to evaluate here in order to calculate $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max}$. In the first case, corresponding to $m \neq m'$ and $\ell = \ell'$, we have the following equality from the orthogonality of discrete sinusoids

$$\begin{aligned}\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell} \rangle &= \sum_{q=1}^{k} a_q^2 \, e^{j(\ell+q-2)(\omega_{m',k}-\omega_{m,k})} \\ &= \frac{1}{k}\sum_{n=0}^{k-1} e^{jn\omega_{(m'-m),k}} = 0. \end{aligned} \tag{3.32}$$

In the second case, corresponding to $\ell < \ell'$, define $\Delta \stackrel{def}{=} (\ell' - \ell)$ and note that the inner product $\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell} \rangle$ in this case can be written as

$$\begin{aligned}\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell} \rangle &= \sum_{q=1}^{k-\Delta} a_q a_{q+\Delta} \, e^{j(\ell'+q-2)\omega_{m',k}} e^{-j(\ell+q+\Delta-2)\omega_{m,k}} \\ &\stackrel{(a)}{=} \sum_{q=1}^{k-\Delta} a_q a_{q+\Delta} \, e^{j(\ell'+q-2)\omega_{(m'-m),k}} \end{aligned} \tag{3.33}$$

where (a) simply follows from the observation that $\ell' = \ell + \Delta$. Notice that we are once against faced with the situation that the terms contributing to the inner product $\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell} \rangle$ are not mutually independent. It is instructive at this point to compare (3.33) with (3.15). Given the similarities between the two expressions, it is easy to see that—similar to the case of (3.15)—the sum in (3.33) can also be written as two sums having mutually independent terms as follows

$$\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell} \rangle = \underbrace{\sum_{q_1=1}^{\lceil \frac{k-\Delta}{2} \rceil} a'_{q_1} \, e^{j\phi_{q_1}}}_{S_{k_1}} + \underbrace{\sum_{q_2=1}^{\lfloor \frac{k-\Delta}{2} \rfloor} a'_{q_2} \, e^{j\phi_{q_2}}}_{S_{k_2}} \tag{3.34}$$

where $\{a'_{q_1}\}$ and $\{a'_{q_2}\}$ consist of $k_1 \stackrel{def}{=} \max\left\{0, \lceil \frac{k-\Delta}{2} \rceil\right\}$ and $k_2 \stackrel{def}{=} \max\left\{0, \lfloor \frac{k-\Delta}{2} \rfloor\right\}$ i.i.d. Rademacher random variables, respectively, that take values $\pm 1/k$ with probability $1/2$, and $\{\phi_{q_1}\}$ and $\{\phi_{q_2}\}$ are some arbitrary (but deterministic) phase factors. Note that while the component sums S_{k_1} and S_{k_2} in (3.34) depend upon the parameters m, m', ℓ, and ℓ', we have suppressed this explicit dependence so as not to clutter the notation.

We can now bound $\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right|$ for the case $\ell < \ell'$ by applying Lemma 3.8 to S_{k_1} and S_{k_2} having k_1 and k_2 terms, respectively, as follows

$$\begin{aligned}\Pr\left(\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| > \frac{\delta_S}{S}\right) &\leq \Pr\left(\left\{|S_{k_1}| > \frac{\delta_S}{2S}\right\} \cup \left\{|S_{k_2}| > \frac{\delta_S}{2S}\right\}\right) \\ &\overset{(b)}{\leq} 2\max\left\{\Pr\left(|S_{k_1}| > \frac{\delta_S}{2S}\right), \Pr\left(|S_{k_2}| > \frac{\delta_S}{2S}\right)\right\} \\ &\overset{(c)}{\leq} 2\max\left\{4\exp\left(-\frac{k^2\delta_S^2}{16k_1S^2}\right), 4\exp\left(-\frac{k^2\delta_S^2}{16k_2S^2}\right)\right\}. \end{aligned} \tag{3.35}$$

Here, (b) follows from a simple union bounding argument and (c) follows from the complex Hoeffding's inequality. Next, notice from the definitions of k_1 and k_2 that $k/2 \geq k_1 \geq k_2$ for any $\Delta \geq 1$. Therefore, we can further simplify (3.35) as follows

$$\Pr\left(\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| > \frac{\delta_S}{S}\right) \leq 8\exp\left(-\frac{k\delta_S^2}{8S^2}\right). \tag{3.36}$$

We now have from (3.32) and (3.36) that $\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| \leq \delta_S/S$ with probability exceeding $1 - 8\exp(-\frac{k\delta_S^2}{8S^2})$ for any $(m,\ell) \neq (m',\ell')$. Finally, to prove that $\mathbf{A} \in RIP(S,\delta_S)$, we can make use of (3.30) to evaluate the probability that $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} \leq \delta_S/S$ as follows

$$\Pr\left(\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} > \frac{\delta_S}{S}\right) \overset{(d)}{\leq} 4p(p-1)\exp\left(-\frac{k\delta_S^2}{8S^2}\right)$$

where (d) follows from taking union bound over a total of $p(p-1)/2$ events (since the cardinality of the set $\left\{\left|\langle \mathbf{a}_{m',\ell'}, \mathbf{a}_{m,\ell}\rangle\right| : (m,\ell) \neq (m',\ell')\right\}$ is $p(p-1)/2$). This implies that for $c_3 < \delta_S^2/8$ and $k \geq (32/(\delta_S^2 - 8c_3))S^2\log p$, we have $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S} \leq (S-1)\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} < \delta_S$ with probability exceeding $1 - \exp(-c_3k/S^2)$, which is what we needed to show. ∎

3.5 On the RIP of Structurally-Subsampled Unitary Matrices

In this section, we further our discussion of structured CS matrices by studying the class of structurally-subsampled unitary matrices, which—as we will shortly see—is a generalization of the class of subsampled unitary matrices. Specifically, let $\mathcal{A}_p \overset{def}{=} \{a_i \in \mathbb{C}\}_{i=1}^p$ denote a p-length generating sequence. Further, let $1 \leq k \leq p$ be an integer parameter that divides p and define $m \overset{def}{=} p/k$. Next, define $\mathbf{R}$ to be an $m \times p$ *row-mixing matrix* that is generated from the sequence

$\mathcal{A}_p$ as follows

$$\mathbf{R} = \begin{bmatrix} a_1 & \dots & a_k & & & & & 0 \\ & & & a_{k+1} & \dots & a_{2k} & & \\ & & & \ddots & \ddots & \ddots & \ddots & \\ 0 & & & & & & a_{p-k+1} & \dots & a_p \end{bmatrix}. \tag{3.37}$$

Finally, let $\mathbf{U}$ be any $p \times p$ unitary matrix and choose a subset Ω of cardinality $n \stackrel{def}{=} |\Omega|$ uniformly at random from the set $[1 \dots m]$. Then the structurally-subsampled unitary matrix $\mathbf{A}$ generated from $(\mathbf{A}_p, \mathbf{U})$ is a submatrix of $\mathbf{X} \stackrel{def}{=} \mathbf{RU}$ obtained by sampling n rows of $\mathbf{X}$ corresponding to the indices in Ω and renormalizing the resulting columns so that they have unit ℓ_2-norms. It is trivial to see from this description that a structurally-subsampled unitary matrix reduces to a subsampled unitary matrix for the case of $a_1 = \cdots = a_p = 1$ and $k = 1$. In order to better motivate the use of structurally subsampled unitary matrices in CS applications, however, let us revisit the context in which subsampled unitary matrices arise in the CS literature.

Subsampled unitary matrices were originally introduced—and partially analyzed—in the modern CS literature in [7]. Initially, the focus in [7] was on sensing matrices that correspond to subsampled Fourier matrices. Later, the analysis was advanced further by Candès and Tao and Rudelson and Vershynin in [8] and [12], respectively, where they established—among other things—that the idea that a small (random) collection of transform-domain samples suffice to encode a sparse signal extends well beyond the Fourier domain. Together, the results of [7, 8, 12] form the basis of the so-called *principle of incoherent measurements*, stated as follows.

> It is best to acquire samples of a sparse signal in a maximally incoherent transform domain $\mathbf{U}$, where the incoherence is measured by the coherence parameter $\mu_{\mathbf{U}}$—the smaller the coherence parameter, the greater the incoherence.[2]

The statement of Theorem 2.9 makes this principle mathematically precise by requiring a randomly subsampled unitary matrix to have $n = \Omega(\mu_{\mathbf{U}}^2 S \times \mathrm{polylog}(p))$ rows in order for it to satisfy RIP with

[2]The coherence parameter gets its name from the fact that we can write $\mu_{\mathbf{U}} \stackrel{def}{=} \sqrt{p}\max_{i,j} |\langle \mathbf{u}_i, \mathbf{e}_j \rangle|$, where $\mathbf{u}_i$ denotes the i-th column of $\mathbf{U}^{\mathsf{H}}$ and $\mathbf{e}_j$ denotes the j-th column of $\mathbf{I}_p$.

high probability. Note that since $\mu_{\mathbf{U}} \stackrel{def}{=} \sqrt{p}\max_{i,j}|u_{i,j}|$, we have that (i) the coherence of a unitary matrix cannot be smaller than 1, and (ii) unitary matrices with entries of magnitude $O(1/\sqrt{p})$ are maximally incoherent. In other words, transform domains such as Fourier, composition of Fourier and wavelet, and Hadamard are all maximally incoherent and are, therefore, particularly well-suited for acquisition of sparse signals.

It turns out that a number of real-world signal acquisition systems already adhere to the principle of incoherent measurements due to various physical and/or technological reasons. For example, data acquired by magnetic resonance imaging (MRI) scanners naturally correspond to Fourier-domain samples of the object being imaged [88]. Similarly, channel measurements collected by a communications receiver using multicarrier modulation inherently correspond to Fourier-domain samples of the channel being estimated [89]. As such, there is a natural fit between the theory of subsampled unitary matrices and these two applications, as noted in, e.g., [14, 48]. If the object being imaged happens to be sparse in the spatial domain then applying Theorem 2.9 to MRI can potentially speed up the scan time by reducing the number of samples required for successful image reconstruction [48]. Similarly, if the (single-antenna) channel being estimated happens to have a sparse impulse response then applying Theorem 2.9 to channel estimation can potentially increase the efficiency of the communications system by reducing the number of measurements required for successful channel estimation [14].

Contrary to these examples, however, our interest in this section is in acquisition systems that—despite sensing sparse signals in an incoherent domain—cannot sample *individual* coefficients in the transform domain. This indeed happens in a number of real-world systems because of a multitude of physical constraints and/or technological limitations. For example, the impulse response of a multiple-antenna channel generally lives in a three-dimensional (3-D) space but a communications receiver using multicarrier modulation can only acquire 2-D projections of its 3-D Fourier-domain samples (physical constraint). Similarly, it is generally desirable to project an ultrawideband signal with limited spectral content onto a smaller spectral band before sampling it since random nonuniform sampling to acquire the signal can be very sensitive to timing errors (technological constraint).

In the parlance of CS, the sensing matrices in both the aforementioned cases now correspond to subsampled versions of $\mathbf{X} = \mathbf{R}\mathbf{U}$ (instead of $\mathbf{U}$), where $\mathbf{R}$ is the row-mixing matrix and $\mathbf{U}$ is the transform domain matrix. In particular, it is easy to see that the theory of subsampled unitary matrices is not easily extendable to structurally-subsampled unitary matrices, except for the trivial case of $\mathbf{R}$ being a diagonal matrix (in other words, $k = 1$). As such, the main question we address in this section is whether structurally-subsampled unitary matrices generated from $(\mathbf{A}_p, \mathbf{U})$ satisfy RIP for the nontrivial case of $k > 1$ when $\mathbf{A}_p$ is a Rademacher sequence.

3.5.1 Main Result

The main tools that we will be using to establish the claim that a structurally-subsampled unitary matrix $\mathbf{A}$ with entries drawn independently from a Rademacher distribution satisfies RIP come from the classical theory of probability in Banach spaces [90]. The general roadmap for the proof is similar to [12, Theorem 3.3], which is now a well-established technique in the CS literature for establishing RIP of subsampled matrices [18, 91]. In particular, the proof relies heavily on an upper bound on expected (T, S)-norm of sum of independent rank-one matrices that was established in [12, Lemma 3.8]. In the following, we describe the basic steps taken to establish a formal proof of our stated claim.

First, we assume that the elements of the generating sequence $\mathcal{A}_p = \{a_i\}_{i=1}^p$ are i.i.d. realizations of Rademacher random variables taking values ± 1 with probability $1/2$, and the sensing matrix $\mathbf{A}$ is generated from $(\mathcal{A}_p, \mathbf{U})$ according to a Bernoulli sampling model. That is, let $\zeta_1, \ldots, \zeta_m$ be independent Bernoulli random variables taking the value 1 with probability n/m, where $m \stackrel{def}{=} p/k$ for some positive integer k that divides p. Then,

$$\Omega \stackrel{def}{=} \{i : \zeta_i = 1\} \tag{3.38}$$

and the structurally-subsampled unitary matrix $\mathbf{A}$ is a (normalized) $|\Omega| \times p$ submatrix of $\mathbf{X} \stackrel{def}{=} \mathbf{R}\mathbf{U}$ obtained by sampling $|\Omega|$ rows of $\mathbf{X}$ corresponding to the indices in Ω and renormalizing the resulting columns by $\sqrt{m/n}$. Here, $\mathbf{R}$ is the $m \times p$ row-mixing matrix of the form (3.37). Note that—unlike the uniformly-at-random sampling model—the cardinality of Ω in Bernoulli sampling

model is a random variable with $\mathbb{E}[|\Omega|] = n$. We then have the following lemma that shows that the Gram matrix $\mathbf{A}^{\mathsf{H}}\mathbf{A} = \mathbf{I}_p$ in expectation.

Lemma 3.10 Let $\mathbf{A}$ be an $n \times p$ structurally-subsampled unitary matrix generated from a Rademacher sequence $\mathcal{A}_p$ according to a Bernoulli sampling model (as described above). Then,

$$\mathbb{E}[\mathbf{A}^{\mathsf{H}}\mathbf{A}] = \mathbf{I}_p. \tag{3.39}$$

Proof: Let $\mathbf{a}_i^{\mathsf{H}} \in \mathbb{C}^p$ denote the i-th row of $\mathbf{A}$. Then it is easy to see that $\mathbf{A}^{\mathsf{H}}\mathbf{A}$ can be written as a sum of rank-one matrices as follows

$$\mathbf{A}^{\mathsf{H}}\mathbf{A} = \sum_{i=1}^{|\Omega|} \mathbf{a}_i\mathbf{a}_i^{\mathsf{H}} = \frac{m}{n}\sum_{i=1}^{p} \zeta_i \mathbf{x}_i\mathbf{x}_i^{\mathsf{H}} \quad \Rightarrow \quad \mathbb{E}[\mathbf{A}^{\mathsf{H}}\mathbf{A}] = \mathbb{E}[\mathbf{X}^{\mathsf{H}}\mathbf{X}] \tag{3.40}$$

where $\mathbf{x}_i^{\mathsf{H}}$ denotes the i-th row of $\mathbf{X}$. Next, from the definition of $\mathbf{X}$, we can write an expression for $\mathbf{x}_i^{\mathsf{H}}$ in terms of elements of the generating sequence $\mathcal{A}_p$ and the rows of $\mathbf{U}$ as follows

$$\mathbf{x}_i^{\mathsf{H}} = \sum_{\ell=1}^{k} a_{(i-1)k+\ell}\, \mathbf{u}_{(i-1)k+\ell}^{\mathsf{H}}\,,\ i = 1, \ldots, m \tag{3.41}$$

where $\mathbf{u}_i^{\mathsf{H}}$ denotes the i-th row of $\mathbf{U}$. With the help of the above expression, we can further write the (i,j)-th entry of $\mathbf{X}$ as

$$x_{i,j} = \sum_{\ell=1}^{k} a_{(i-1)k+\ell}\, u_{(i-1)k+\ell,j}\,,\ i = 1, \ldots, m,\ j = 1, \ldots, p \tag{3.42}$$

where $u_{i,j}$ denotes the (i,j)-th entry of $\mathbf{U}$. It is then easy to see from this expression that

$$\begin{aligned}
\mathbb{E}[x_{i,j}^* x_{i,j'}] &= \sum_{q=1}^{k}\sum_{r=1}^{k} \mathbb{E}\big[a_{(i-1)k+q}\, a_{(i-1)k+r}\big]\, u_{(i-1)k+q,j}^*\, u_{(i-1)k+r,j'} \\
&= \sum_{q=1}^{k} u_{(i-1)k+q,j}^*\, u_{(i-1)k+q,j'},\ j, j' = 1, \ldots, p.
\end{aligned} \tag{3.43}$$

Finally, define the Gram matrix $\mathbf{G} \overset{def}{=} \mathbf{X}^{\mathsf{H}}\mathbf{X}$ and note from (3.43) that the expected value of the (i,j)-th entry of $\mathbf{G}$, $g_{i,j} = \sum_{\ell=1}^{m} x_{\ell,i}^* x_{\ell,j}$, is given by

$$\begin{aligned}
\mathbb{E}[g_{i,j}] &= \sum_{\ell=1}^{m} \mathbb{E}[x_{\ell,i}^* x_{\ell,j}] = \sum_{\ell=1}^{m}\sum_{q=1}^{k} u_{(\ell-1)k+q,i}^*\, u_{(\ell-1)k+q,j} \\
&= \sum_{\ell=1}^{p} u_{\ell,i}^* u_{\ell,j} \overset{(a)}{=} \delta_{ij},\ i, j = 1, \ldots, p
\end{aligned} \tag{3.44}$$

where (a) follows from the fact that $\mathbf{U}$ is a unitary matrix. This completes the proof of the lemma since (3.44) implies that $\mathbb{E}[\mathbf{G}] = \mathbf{I}_p \Rightarrow \mathbb{E}[\mathbf{A}^{\mathsf{H}}\mathbf{A}] = \mathbf{I}_p$ from (3.40). ∎

Second, we use Lemma 3.10 to establish that $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S}$ cannot be too large in expectation for large-enough values of n. The proof of this result, however, is a little more involved and makes use of a number of auxiliary lemmas. The most important lemma that we will need in this regard is the following one, due to Rudelson and Vershynin [12, Lemma 3.8].

Lemma 3.11 (Rudelson–Vershynin) Let $\mathbf{z}_1, \ldots, \mathbf{z}_r$, $r \leq p$, be vectors in $\mathbb{C}^p$ with uniformly bounded entries, $\|\mathbf{z}_i\|_\infty \leq K$ for all i. Further, let $\{\varepsilon_i\}$ be i.i.d. Rademacher random variables taking values ± 1 with probability $1/2$. Then

$$\mathbb{E}\left[\Big\| \sum_{i=1}^{r} \varepsilon_i \mathbf{z}_i \mathbf{z}_i^{\mathsf{H}} \Big\|_{T,S}\right] \leq B(r) \cdot \Big\| \sum_{i=1}^{r} \mathbf{z}_i \mathbf{z}_i^{\mathsf{H}} \Big\|_{T,S}^{1/2} \tag{3.45}$$

where $B(r) \stackrel{def}{=} c_5\, K\sqrt{S}\log(S)\sqrt{\log p}\sqrt{\log r}$ for some absolute constant $c_5 > 0$.

In order to make use of Lemma 3.11, however, we require the entries of $\mathbf{A}$ to be uniformly bounded by some number K. To this end, we will make use of the classical Khintchine inequality for i.i.d. Rademacher random variables [90, Lemma 4.1].

Lemma 3.12 (Khintchine Inequality) Let $\{\varepsilon_i\}$ be i.i.d. Rademacher random variables taking values ± 1 with probability $1/2$. For any $s \in (0, \infty)$, there exist positive finite constants C_s and D_s depending on s only such that for any finite sequence $\{\alpha_i\}$ of complex numbers

$$C_s \Big(\sum_i |\alpha_i|^2 \Big)^{1/2} \leq \Big(\mathbb{E}\Big[\Big| \sum_i \varepsilon_i \alpha_i \Big|^s \Big] \Big)^{1/s} \leq D_s \Big(\sum_i |\alpha_i|^2 \Big)^{1/2}. \tag{3.46}$$

In the sequel, we will only be concerned with the upper bound in Khintchine inequality. In that regard, Haagerup proved in [92] that the best constant D_s in (3.46) for the case of real numbers is

$$D_s^* \stackrel{def}{=} \begin{cases} 1, & \text{if } 0 < s \leq 2, \\ 2^{1/2} \Big(\frac{\Gamma((s+1)/2)}{\sqrt{\pi}} \Big)^{1/s}, & \text{if } 2 < s < \infty, \end{cases} \tag{3.47}$$

where $\Gamma(z) \stackrel{def}{=} \int_0^\infty t^{z-1} e^{-t} dt$ is the Gamma function. Note that D_s^* is also a valid constant in the case of complex numbers, since if the upper bound in the Khintchine inequality holds for real

numbers with some constant then it also holds for complex numbers with the same constant. We are now ready to prove that the entries of the structurally-subsampled unitary matrix $\mathbf{A}$ cannot be too large. The technique used to establish the following lemma is very similar to that of [18, Lemma 5].

Lemma 3.13 Let $\mathbf{A}$ be an $n \times p$ structurally-subsampled unitary matrix generated from a Rademacher sequence $\mathcal{A}_p$ according to a Bernoulli sampling model (as described earlier). Then for any integer $p > 2$ and any $r \in [2, 2\log p]$, we have

$$\Big(\mathbb{E}\big[\|\mathbf{A}\|_{\max}^r\big]\Big)^{1/r} \leq \sqrt{\frac{m}{n}}\Big(\mathbb{E}\big[\|\mathbf{X}\|_{\max}^r\big]\Big)^{1/r} \leq \sqrt{\frac{16\mu_{\mathbf{U}}^2 \log p}{n}}. \tag{3.48}$$

Proof: Note that since $\mathbf{A}$ is just a normalized submatrix of $\mathbf{X}$, the first inequality in the lemma is a trivial consequence of the definition of $\mathbf{A}$. The second inequality in the lemma can be established by first focusing on the (i,j)-th entry of the matrix $\mathbf{X} = \mathbf{RU}$. It is easy to see from the definition of $x_{i,j}$ in (3.42) that for any $s \in [2, \infty)$

$$\begin{aligned}\Big(\mathbb{E}\big[|x_{i,j}|^s\big]\Big)^{1/s} &\overset{(a)}{\leq} D_s^*\bigg(\sum_{\ell=1}^{k}\big|u_{(i-1)k+\ell,j}\big|^2\bigg)^{1/2} \\ &\overset{(b)}{\leq} \frac{D_s^*\mu_{\mathbf{U}}}{\sqrt{m}}\,,\ i = 1,\ldots,m,\ j = 1,\ldots,p\end{aligned} \tag{3.49}$$

where (a) follows from the Khintchine inequality (Lemma 3.12) and (b) follows from the fact that $|u_{i,j}| \leq \mu_{\mathbf{U}}/\sqrt{p}\ \forall\ i,j = 1,\ldots,p$. Next, we make use of the above inequality and the Hölder's inequality to bound the r-th moment of $\|\mathbf{X}\|_{\max}$. Note that for any $2 \leq r \leq s$, the following holds

$$\begin{aligned}\Big(\mathbb{E}\big[\|\mathbf{X}\|_{\max}^r\big]\Big)^{1/r} &\overset{(c)}{\leq} \Big(\mathbb{E}\big[\|\mathbf{X}\|_{\max}^s\big]\Big)^{1/s} \\ &\overset{(d)}{\leq} (mp)^{1/s}\Big\|\mathbb{E}\big[|\mathbf{X}|^s\big]^{1/s}\Big\|_{\max} \\ &\overset{(e)}{\leq} \frac{(mp)^{1/s}D_s^*\mu_{\mathbf{U}}}{\sqrt{m}}.\end{aligned} \tag{3.50}$$

Here, the notation $\mathbb{E}\big[|\mathbf{X}|^s\big]^{1/s}$ is meant to signify an $m \times p$ matrix obtained by taking entry-wise s-th absolute moments of $\mathbf{X}$. Note that (c) in the above expression is a simple consequence of the Hölder's inequality, (d) follows from the fact that $\Big(\mathbb{E}\big[\|\mathbf{X}\|_{\max}^s\big]\Big)^{1/s} \leq \Big(\sum_{i,j}\mathbb{E}\big[|x_{i,j}|^s\big]\Big)^{1/s}$, and

(e) simply follows from (3.49). Now choose s to be the smallest even integer such that it satisfies $s \geq 2\log p \geq r$, then we have from Stirling's formula

$$D_s^* = \left(\frac{s!}{2^{s/2}\cdot(s/2)!}\right)^{1/s} \leq 2^{1/2s}\sqrt{s/e} \overset{(f)}{\leq} 2^{5/4}\sqrt{\log p/e} \tag{3.51}$$

where (f) mainly follows from the fact that, by definition, we have $s \leq 2\log p + 2 \leq 4\log p$ for any $p > 2$. Finally, note that since $2\log p \leq s \;\Rightarrow\; (mp)^{1/s} \leq p^{2/s} \leq e$. This trivial observation, combined with (3.50) and (3.51), gives us the second inequality in the lemma. ∎

All the pieces are now in place to bound $\mathbb{E}\big[\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S}\big]$ using Lemma 3.11 and techniques developed in probability in Banach spaces [90].

Lemma 3.14 Let $\mathbf{A}$ be an $n \times p$ structurally-subsampled unitary matrix generated from a Rademacher sequence $\mathcal{A}_p$ according to a Bernoulli sampling model (as described earlier). Then for any integer $p > 2$ and any $\epsilon \in (0,1)$, we have

$$\mathbb{E}\big[\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S}\big] \leq \epsilon \tag{3.52}$$

provided $n \geq c_6\,\epsilon^{-2}\mu_{\mathbf{U}}^2\,S\log^3 p\log^2 S$ for some absolute constant $c_6 > 0$.

Proof: To establish this lemma, first define $E \overset{def}{=} \mathbb{E}\big[\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S}\big]$. Next, we have from Lemma 3.10 that $\mathbb{E}[\mathbf{A}^{\mathsf{H}}\mathbf{A}] = \mathbf{I}_p$. Therefore, it follows from (3.40) that

$$\mathbf{I}_p = \mathbb{E}_{\zeta',\mathbf{x}'}\left[\frac{m}{n}\sum_{i=1}^{p}\zeta_i'\mathbf{x}_i'\mathbf{x}_i'^{\mathsf{H}}\right] \tag{3.53}$$

where $\{\zeta_i'\}$ and $\{\mathbf{x}_i'^{\mathsf{H}}\}$ are independent copies of the Bernoulli random variables $\{\zeta_i\}$ and the (random) row vectors $\{\mathbf{x}_i^{\mathsf{H}}\}$, respectively. Consequently, we have from (3.40) and (3.53) that

$$\begin{aligned}
E &= \mathbb{E}_{\zeta,\mathbf{x}}\left[\left\|\frac{m}{n}\sum_{i=1}^{p}\Big(\zeta_i\mathbf{x}_i\mathbf{x}_i^{\mathsf{H}} - \mathbb{E}_{\zeta',\mathbf{x}'}\big[\zeta_i'\mathbf{x}_i'\mathbf{x}_i'^{\mathsf{H}}\big]\Big)\right\|_{T,S}\right] \\
&\overset{(a)}{\leq} \mathbb{E}_{\zeta,\zeta',\mathbf{x},\mathbf{x}'}\left[\left\|\frac{m}{n}\sum_{i=1}^{p}\Big(\zeta_i\mathbf{x}_i\mathbf{x}_i^{\mathsf{H}} - \zeta_i'\mathbf{x}_i'\mathbf{x}_i'^{\mathsf{H}}\Big)\right\|_{T,S}\right] \\
&\overset{(b)}{=} \mathbb{E}_{\zeta,\zeta',\mathbf{x},\mathbf{x}',\varepsilon}\left[\left\|\frac{m}{n}\sum_{i=1}^{p}\varepsilon_i\Big(\zeta_i\mathbf{x}_i\mathbf{x}_i^{\mathsf{H}} - \zeta_i'\mathbf{x}_i'\mathbf{x}_i'^{\mathsf{H}}\Big)\right\|_{T,S}\right] \\
&\overset{(c)}{\leq} 2\,\mathbb{E}_{\zeta,\mathbf{x},\varepsilon}\left[\left\|\frac{m}{n}\sum_{i=1}^{p}\varepsilon_i\zeta_i\mathbf{x}_i\mathbf{x}_i^{\mathsf{H}}\right\|_{T,S}\right].
\end{aligned} \tag{3.54}$$

Here, $\{\varepsilon_i\}$ are i.i.d. Rademacher random variables taking values ± 1 with probability $1/2$ that are independent of all other random variables. Note that (a) in the above expression follows from Jensen's inequality since $\|\cdot\|_{T,S}$ is a norm (Lemma 3.1), (b) follows from the fact that $\Big\{\zeta_i \mathbf{x}_i \mathbf{x}_i^{\mathsf{H}} - \zeta_i' \mathbf{x}_i' \mathbf{x}_i'^{\mathsf{H}}\Big\}$ is a sequence of independent, symmetric (matrix-valued) random variables and therefore has the same distribution as $\Big\{\varepsilon_i\big(\zeta_i \mathbf{x}_i \mathbf{x}_i^{\mathsf{H}} - \zeta_i' \mathbf{x}_i' \mathbf{x}_i'^{\mathsf{H}}\big)\Big\}$ [90, pp. 150-151], and (c) follows from the triangle inequality (again, since $\|\cdot\|_{T,S}$ is a norm) and the fact that the distributions of $\big\{\zeta_i'\big\}$ and $\big\{\mathbf{x}_i'^{\mathsf{H}}\big\}$ coincide with the distributions of $\big\{\zeta_i\big\}$ and $\big\{\mathbf{x}_i^{\mathsf{H}}\big\}$, respectively.

We can now once again appeal to (3.40) and write $\frac{m}{n}\sum_{i=1}^{p}\varepsilon_i\zeta_i\mathbf{x}_i\mathbf{x}_i^{\mathsf{H}} = \sum_{i=1}^{|\Omega|}\varepsilon_i\mathbf{a}_i\mathbf{a}_i^{\mathsf{H}}$. Therefore, we have from (3.54) and Lemma 3.11 that

$$E \leq 2\,\mathbb{E}_{\varepsilon,\mathbf{a}}\left[\Big\|\sum_{i=1}^{|\Omega|}\varepsilon_i\mathbf{a}_i\mathbf{a}_i^{\mathsf{H}}\Big\|_{T,S}\right] \leq 2\,\mathbb{E}_{\mathbf{a}}\left[B(|\Omega|)\cdot\Big\|\sum_{i=1}^{|\Omega|}\mathbf{a}_i\mathbf{a}_i^{\mathsf{H}}\Big\|_{T,S}^{1/2}\right] \tag{3.55}$$

where $B(|\Omega|) = c_5\,\|\mathbf{A}\|_{\max}\sqrt{S}\log(S)\sqrt{\log p}\sqrt{\log|\Omega|}$. Next, we can use Cauchy-Schwarz inequality and Lemma 3.13 to obtain

$$\begin{aligned}\mathbb{E}\left[B(|\Omega|)\cdot\Big\|\sum_{i=1}^{|\Omega|}\mathbf{a}_i\mathbf{a}_i^{\mathsf{H}}\Big\|_{T,S}^{1/2}\right] &\leq c_5\sqrt{S}\log(S)\sqrt{\log p}\,\times \\ &\times\ \sqrt{\frac{m}{n}}\Big(\mathbb{E}\big[\|\mathbf{X}\|_{\max}^2\big]\Big)^{1/2}\Big(\mathbb{E}\big[\log|\Omega|\big]\Big)^{1/2}\left(\mathbb{E}\left[\Big\|\sum_{i=1}^{|\Omega|}\mathbf{a}_i\mathbf{a}_i^{\mathsf{H}}\Big\|_{T,S}\right]\right)^{1/2}. \end{aligned}\tag{3.56}$$

Further, we also have from Lemma 3.13 that $\sqrt{\frac{m}{n}}\Big(\mathbb{E}\big[\|\mathbf{X}\|_{\max}^2\big]\Big)^{1/2} \leq \sqrt{\frac{16\mu_{\mathbf{U}}^2\log p}{n}}$, from Jensen's inequality that $\mathbb{E}\big[\log|\Omega|\big] \leq \log n$, and from triangle inequality (by adding and subtracting an identity matrix inside the norm) that $\big\|\sum_{i=1}^{|\Omega|}\mathbf{a}_i\mathbf{a}_i^{\mathsf{H}}\big\|_{T,S} \leq \big\|\sum_{i=1}^{|\Omega|}\mathbf{a}_i\mathbf{a}_i^{\mathsf{H}} - \mathbf{I}_p\big\|_{T,S} + 1$. Collecting all these facts together, we can rewrite (3.55) as

$$E \leq \frac{8c_5\mu_{\mathbf{U}}\sqrt{S}\log(S)\log(p)\sqrt{\log n}}{\sqrt{n}}\cdot(E+1)^{1/2} \tag{3.57}$$

which implies that $E \leq \epsilon$ whenever $n \geq c_6\,\epsilon^{-2}\mu_{\mathbf{U}}^2\,S\log^3 p\log^2 S$, where the constant $c_6 \stackrel{def}{=} 256\,c_5^2$. Here, the last assertion follows from the simple observation that $z \leq b\sqrt{z+1} \;\Rightarrow\; z \leq 2b$, whenever $b \leq 1$. This completes the proof of the lemma. ∎

Finally, we show that $\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S}$ concentrates around its mean with high probability. To establish this fact, however, we need one additional classical result from the theory of probability in

Banach spaces. The following result is originally due to Ledoux and Talagrand [90, Theorem 6.17] and appears in the following form in [12, Theorem 3.10].

Theorem 3.15 (Ledoux–Talagrand) Let $\mathcal{B} \stackrel{def}{=} (X, \|\cdot\|_X)$ be a Banach space. Further, let $\{Y_i\}_{i=1}^N$ be independent, symmetric random variables in $\mathcal{B}$ such that $\|Y_i\|_X \leq B$ for every i almost surely. Finally, define $Y \stackrel{def}{=} \left\| \sum_{i=1}^N Y_i \right\|_X$. Then for any integers $r \geq q$, any $t > 0$, and some absolute constant $c_7 > 0$, Y satisfies

$$\Pr\left(Y \geq 8q\mathbb{E}[Y] + 2rB + t\right) \leq \left(\frac{c_7}{q}\right)^r + 2\exp\left(-\frac{t^2}{256q\mathbb{E}[Y]^2}\right). \tag{3.58}$$

We are now ready to establish RIP for structurally-subsampled unitary matrices generated from $(\mathbf{A}_p, \mathbf{U})$ when $\mathbf{A}_p$ is a Rademacher sequence.

Theorem 3.16 Let the elements of the generating sequence $\mathcal{A}_p = \{a_i\}_{i=1}^p$ be i.i.d. realizations of Rademacher random variables taking values ± 1 with probability $1/2$. Further, let $\mathbf{R}$—as defined in (3.37)—be the $m \times p$ row-mixing matrix generated by the sequence $\mathcal{A}_p$, where $m \stackrel{def}{=} p/k$ for an integer parameter $k \in [1 \ldots p]$ that divides p. Choose a subset Ω of cardinality $n \stackrel{def}{=} |\Omega|$ uniformly at random from the set $[1 \ldots m]$. Finally, let $\mathbf{U}$ be any $p \times p$ unitary matrix, and let $\mathbf{A}$ be the $n \times p$ matrix obtained by sampling n rows of $\mathbf{X} \stackrel{def}{=} \mathbf{RU}$ corresponding to the indices in Ω and renormalizing the resulting columns by $\sqrt{m/n}$. Then for each integer $p, S > 2$, and for any $z > 1$ and any $\delta_S \in (0, 1)$, there exist absolute constants c_8 and c_9 such that whenever

$$n \geq c_8 z \mu_{\mathbf{U}}^2 S \log^3 p \log^2 S \tag{3.59}$$

the matrix $\mathbf{A} \in RIP(S, \delta_S)$ with probability exceeding $1 - 20\max\left\{\exp\left(-c_9\delta_S^2 z\right), p^{-1}\right\}$.

Proof: We begin by recalling the result established in [7, Section 2.3], which states that if it can be shown that subsampled matrices in a particular class satisfy RIP with probability exceeding $1 - \eta$ for the Bernoulli sampling model, then it follows that subsampled matrices belonging to the same class satisfy RIP with probability exceeding $1 - 2\eta$ for the uniformly-at-random sampling model. As such, we initially assume that the structurally-subsampled unitary matrix $\mathbf{A}$ is generated from $(\mathbf{A}_p, \mathbf{U})$ according to a Bernoulli sampling model.

Next, consider the Banach space $\mathcal{B} \stackrel{def}{=} (\mathbb{C}^{p\times p}, \|\cdot\|_{T,S})$ and define random variables $\{\mathbf{Y}_i\}_{i=1}^p$ and $\{\widetilde{\mathbf{Y}}_i\}_{i=1}^p$ that take values in $\mathcal{B}$ as follows

$$\mathbf{Y}_i \stackrel{def}{=} \frac{m}{n}\zeta_i \mathbf{x}_i\mathbf{x}_i^{\mathsf{H}} - \frac{1}{p}\mathbf{I}_p, \quad \text{and} \quad \widetilde{\mathbf{Y}}_i \stackrel{def}{=} \frac{m}{n}\Big(\zeta_i \mathbf{x}_i\mathbf{x}_i^{\mathsf{H}} - \zeta_i'\mathbf{x}_i'\mathbf{x}_i'^{\mathsf{H}}\Big), \quad i = 1,\ldots,p \tag{3.60}$$

where, as before, $\{\zeta_i\}$ are the Bernoulli random variables arising in the Bernoulli sampling model, $\{\mathbf{x}_i^{\mathsf{H}}\}$ denote the rows of $\mathbf{X} = \mathbf{R}\mathbf{U}$, and $\{\zeta_i'\}$ and $\{\mathbf{x}_i'^{\mathsf{H}}\}$ are independent copies of $\{\zeta_i\}$ and $\{\mathbf{x}_i^{\mathsf{H}}\}$, respectively. In other words, each random variable $\widetilde{\mathbf{Y}}_i \stackrel{def}{=} \mathbf{Y}_i - \mathbf{Y}_i'$ is a symmetric version of the corresponding random variable $\mathbf{Y}_i$, where $\mathbf{Y}_i'$ denotes an independent copy of $\mathbf{Y}_i$. In particular, we have that the random variable $\sum_{i=1}^p \widetilde{\mathbf{Y}}_i$ in $\mathcal{B}$ is a symmetric version of $\sum_{i=1}^p \mathbf{Y}_i$ and, as a consequence, the following symmetrization inequalities hold for all $u > 0$ [90, Chapter 6]

$$\mathbb{E}\bigg[\Big\|\sum_{i=1}^p \widetilde{\mathbf{Y}}_i\Big\|_{T,S}\bigg] \le 2\mathbb{E}\bigg[\Big\|\sum_{i=1}^p \mathbf{Y}_i - \mathbb{E}\big[\sum_{i=1}^p \mathbf{Y}_i\big]\Big\|_{T,S}\bigg], \tag{3.61}$$

$$\Pr\left(\Big\|\sum_{i=1}^p \mathbf{Y}_i\Big\|_{T,S} > 2\mathbb{E}\bigg[\Big\|\sum_{i=1}^p \mathbf{Y}_i\Big\|_{T,S}\bigg] + u\right) \le 2\Pr\left(\Big\|\sum_{i=1}^p \widetilde{\mathbf{Y}}_i\Big\|_{T,S} > u\right). \tag{3.62}$$

There are two key observations that can be made here. First, we can bound the expected value of $\widetilde{\mathbf{Y}} \stackrel{def}{=} \|\sum_{i=1}^p \widetilde{\mathbf{Y}}_i\|_{T,S}$ using (3.61) and Lemma 3.14 since (i) $\mathbb{E}\big[\sum_{i=1}^p \mathbf{Y}_i\big] = 0$ (Lemma 3.10), and (ii) $\mathbf{Y} \stackrel{def}{=} \|\sum_{i=1}^p \mathbf{Y}_i\|_{T,S} = \|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{T,S}$. Second, we can obtain a large-deviation bound for $\mathbf{Y}$ using (3.62) and Theorem 3.15 since—by construction—$\{\widetilde{\mathbf{Y}}_i\}_{i=1}^p$ are independent, symmetric random variables in $\mathcal{B}$. Before can use Thereom 3.15 to characterize the tail behavior of $\widetilde{\mathbf{Y}}$, however, we need to establish that $\max_i \|\widetilde{\mathbf{Y}}_i\|_{T,S} \le B$ for some B.

To this end, we first establish that $\max_i \big\{\sqrt{\frac{m}{n}}\|\mathbf{x}_i^{\mathsf{H}}\|_\infty, \sqrt{\frac{m}{n}}\|\mathbf{x}_i'^{\mathsf{H}}\|_\infty\big\}$ cannot be too large with high probability. Specifically, note from Lemma 3.13 that we have for $r = 2\log p$

$$\begin{aligned}
\Pr\left(\sqrt{\frac{m}{n}}\big\|\mathbf{X}\big\|_{\max} > \sqrt{\frac{16\,e\mu_{\mathbf{U}}^2 \log p}{n}}\right) &\le \Pr\left(\big\|\mathbf{X}\big\|_{\max} > \sqrt{e}\Big(\mathbb{E}\big[\|\mathbf{X}\|_{\max}^r\big]\Big)^{1/r}\right) \\
&= \Pr\left(\big\|\mathbf{X}\big\|_{\max}^r > e^{r/2}\cdot\mathbb{E}\big[\|\mathbf{X}\|_{\max}^r\big]\right) \\
&\stackrel{(a)}{\le} \frac{\mathbb{E}\big[\|\mathbf{X}\|_{\max}^r\big]}{e^{r/2}\cdot\mathbb{E}\big[\|\mathbf{X}\|_{\max}^r\big]} = p^{-1}
\end{aligned} \tag{3.63}$$

where (a) follows from a simple application of Markov's inequality. Next, define $B_1 \stackrel{def}{=} \frac{16\, e\mu_{\mathbf{U}}^2 \log p}{n}$. Then we have from (3.63) that

$$\Pr\left(\left\{\sqrt{\frac{m}{n}}\|\mathbf{X}\|_{\max} > \sqrt{B_1}\right\} \bigcup \left\{\sqrt{\frac{m}{n}}\|\mathbf{X}'\|_{\max} > \sqrt{B_1}\right\}\right) \stackrel{(b)}{\leq} 2p^{-1} \tag{3.64}$$

where $\mathbf{X}'$ is comprised of $\left\{\mathbf{x}_i'^{\mathsf{H}}\right\}$ as its rows (in other words, $\mathbf{X}'$ is an independent copy of $\mathbf{X}$), and (b) follows from a simple union bounding argument. Further, we also have

$$\begin{aligned}
\max_i \left\|\widetilde{\mathbf{Y}}_i\right\|_{T,S} &= \max_i \left\|\frac{m}{n}\Big(\zeta_i \mathbf{x}_i \mathbf{x}_i^{\mathsf{H}} - \zeta_i' \mathbf{x}_i' \mathbf{x}_i'^{\mathsf{H}}\Big)\right\|_{T,S} \\
&\stackrel{(c)}{\leq} \max_i \left\{\left\|\frac{m}{n}\mathbf{x}_i \mathbf{x}_i^{\mathsf{H}}\right\|_{T,S} + \left\|\frac{m}{n}\mathbf{x}_i' \mathbf{x}_i'^{\mathsf{H}}\right\|_{T,S}\right\} \\
&\stackrel{(d)}{\leq} \max_i \left\{S\Big(\sqrt{\frac{m}{n}}\|\mathbf{x}_i^{\mathsf{H}}\|_\infty\Big)^2 + S\Big(\sqrt{\frac{m}{n}}\|\mathbf{x}_i'^{\mathsf{H}}\|_\infty\Big)^2\right\} \\
&\stackrel{(e)}{\leq} S\Big(\frac{m}{n}\|\mathbf{X}\|_{\max}^2 + \frac{m}{n}\|\mathbf{X}'\|_{\max}^2\Big)
\end{aligned} \tag{3.65}$$

where (c) mainly follows from triangle inequality, (d) is a simple consequence of the definition of (T,S)-norm, and (e) follows from the fact that $\|\mathbf{X}\|_{\max} \stackrel{def}{=} \max_i \|\mathbf{x}_i^{\mathsf{H}}\|_\infty$ (and in the same way, $\|\mathbf{X}'\|_{\max} \stackrel{def}{=} \max_i \|\mathbf{x}_i'^{\mathsf{H}}\|_\infty$). It is then easy to see from (3.63) and (3.65) that we have $\max_i \|\widetilde{\mathbf{Y}}_i\|_{T,S} \leq 2SB_1$ with probability exceeding $1 - 2p^{-1}$.

Finally, define the event $E \stackrel{def}{=} \left\{\max_i \|\widetilde{\mathbf{Y}}_i\|_{T,S} \leq 2SB_1\right\}$. Then, conditioned on this event, we have from (3.61), Lemma 3.14 and Theorem 3.15 that whenever $n \geq c_6\, \epsilon^{-2}\mu_{\mathbf{U}}^2\, S \log^3 p \log^2 S$

$$\Pr\left(\widetilde{\mathbf{Y}} \geq 16q\epsilon + 4rSB_1 + t \big| E\right) < \left(\frac{c_7}{q}\right)^r + 2\exp\left(-\frac{t^2}{1024q\epsilon^2}\right) \tag{3.66}$$

for any integer $r \geq q$, any $t > 0$, and any $\epsilon \in (0,1)$. Next, choose $q = \lceil ec_7 \rceil$, $t = 32\sqrt{q}\eta\epsilon$, and $r = \lceil \frac{t}{2SB_1} \rceil$ for some $\eta > 1$. Further, define a new constant $c_8 \stackrel{def}{=} \max\left\{e\sqrt{q}, c_6\right\}$ and let $n \geq c_8\, \epsilon^{-2}\mu_{\mathbf{U}}^2\, S \log^3 p \log^2 S$. Note that this choice of n ensures $r \geq q$, resulting in

$$\Pr\left(\widetilde{\mathbf{Y}} \geq (16q + 96\sqrt{q}\eta)\epsilon \big| E\right) < \exp\left(-\frac{\sqrt{q}\eta\epsilon n}{3\mu_{\mathbf{U}}^2 S \log p}\right) + 2\exp\left(-\eta^2\right). \tag{3.67}$$

We can now get rid of the conditioning in the above expression by noting that $\Pr(E^c) \leq 2p^{-1}$, which in turn implies

$$\Pr\left(\widetilde{\mathbf{Y}} \geq (16q + 96\sqrt{q}\eta)\epsilon\right) < \exp\left(-\frac{\sqrt{q}\eta\epsilon n}{3\mu_{\mathbf{U}}^2 S \log p}\right) + 2\exp\left(-\eta^2\right) + 2p^{-1}. \tag{3.68}$$

In the end, what remains to be shown is that $\mathbf{Y} = \| \sum_{i=1}^{p} \mathbf{Y}_i \|_{T,S} = \| \mathbf{A}^{\mathsf{H}} \mathbf{A} - \mathbf{I}_p \|_{T,S} \leq \delta_S$ with high probability. To this end, note that if $n \geq c_8 \, \epsilon^{-2} \mu_{\mathbf{U}}^2 \, S \log^3 p \log^2 S$ then $\mathbb{E}[\mathbf{Y}] \leq \epsilon$ from Lemma 3.14. Consequently, we get from (3.62) and (3.68) that

$$\Pr\left(\mathbf{Y} \geq (2 + 16q + 96\sqrt{q}\eta)\epsilon \right) < 2 \exp\left(-\frac{\sqrt{q}\eta\epsilon n}{3\mu_{\mathbf{U}}^2 S \log p} \right) + 4 \exp\left(-\eta^2 \right) + 4p^{-1}. \tag{3.69}$$

Finally, define $c_6' \stackrel{def}{=} (2 + 16q + 96\sqrt{q})$ and note that $c_6' \eta \, \epsilon > (2 + 16q + 96\sqrt{q}\eta)\epsilon$ since $\eta > 1$. If we now choose $\eta = \frac{\delta_S}{c_6' \epsilon}$ then $\frac{\sqrt{q}\eta\epsilon n}{3\mu_{\mathbf{U}}^2 S \log p} > \eta^2$ and, therefore, (3.69) can be simplified as

$$\Pr\left(\mathbf{Y} \geq \delta_S \right) < 10 \max\left\{ \exp\left(-c_9 \delta_S^2 z \right), p^{-1} \right\} \tag{3.70}$$

where $c_9 \stackrel{def}{=} 1/c_6'$ and $z \stackrel{def}{=} 1/\epsilon^2$. The claim of the theorem now trivially follows from the discussion at the start of the proof. ∎

3.6 Discussion

3.6.1 Toeplitz Matrices

We have established in Section 3.3 of this chapter that full and partial Toeplitz matrices generated from a Rademacher sequence $\mathcal{A}_k$ satisfy RIP. However, note that similar results also hold in settings where the entries of the generating sequence are independently drawn from either any bounded zero-mean distribution or certain unbounded zero-mean distributions, such as the Gaussian distribution [93]. Similarly, it is easy to see that the results of Theorem 3.5 and Theorem 3.6 also apply directly to Hankel matrices, which are Toeplitz-structured matrices whose entries are constant along anti-diagonals. In addition, the proof technique utilized to obtain the results of Theorem 3.6 can also be used to establish RIP for right-shifted (or left-shifted) $n \times p$ partial circulant matrices that are generated from a random sequence $\mathcal{A}_k$ for $k = p \geq n$ as follows

$$\mathbf{A} = \begin{bmatrix} a_p & a_{p-1} & \dots & \dots & \dots & a_3 & a_2 & a_1 \\ a_1 & a_p & \dots & \dots & \dots & a_4 & a_3 & a_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ a_{n-1} & a_{n-2} & \dots & a_1 & a_p & \dots & a_{n+1} & a_n \end{bmatrix}. \tag{3.71}$$

It is also instructive to note at this point that using Lemma 3.3 to establish the RIP of a sensing matrix is not without its limitations. Specifically, for a general $n \times p$ matrix $\mathbf{A}$ with unit ℓ_2-norm columns and $n < p$, it can be verified that [94, Theorem 2.3]

$$\|\mathbf{A}^{\mathsf{H}}\mathbf{A} - \mathbf{I}_p\|_{\max} \geq \sqrt{\frac{p-n}{n(p-1)}} \approx \frac{1}{\sqrt{n}} \tag{3.72}$$

which would essentially imply—at best—an $n = \Omega(S^2)$ requirement on the number of observations for $\mathbf{A}$ to satisfy RIP, similar to what we obtain in the proof of Theorem 3.6. In particular, Lemma 3.3 in conjunction with (3.72) leads to the requirement that $\Omega(S^2)$ observations are needed in order for an i.i.d. random matrix to satisfy RIP. On the other hand, we know from Theorem 2.8 that $\Omega(S \log p)$ observations suffice to guarantee RIP for i.i.d. random matrices.

Therefore, while one might be tempted to conclude that the highly structured nature of Toeplitz matrices results in an increase in the number of observations required for RIP to be guaranteed (compared to canonical CS matrices), such a conclusion does not follow from the results established in Section 3.3. In fact, it is quite possible that Toeplitz matrices do satisfy RIP when k (or n) $= \Omega\left(S \times \mathrm{polylog}(p)\right)$, but the mathematical tools currently at our disposal are insufficient to establish this stronger result. The takeaway message here is that Lemma 3.3 provides a relatively straightforward—but possibly suboptimal—approach to establishing RIP for structured sensing matrices.

In the end, we conclude the discussion of our results on Toeplitz matrices with a brief overview of the connections between the results of Section 3.3 and some related existing works. To the best of our knowledge, Toeplitz-structured matrices were introduced in the modern CS literature by Tropp et al. in [80] in the context of *random filters*. (It is worth noting though that Toeplitz matrices have been considered in the sparse approximation literature even before the advent of CS—see, e.g., [95].) Nevertheless, no theoretical guarantees were provided by the authors in [80] and the results reported therein relied exclusively on numerical simulations.

The first theoretical results for using Toeplitz-structured matrices in CS applications were established by us in [96]. In particular, using a novel combination of existing results on i.i.d. CS matrices and equitable coloring of graphs, we showed in [96, Theorem 1] that $n \times p$ partial Toeplitz matrices satisfy RIP of order S with high probability, provided $n = \Omega\left(S^3 \log p\right)$. This sufficient

condition is of course more restrictive than what we established in Section 3.3, where we reduced the exponent on S by one order of magnitude.

Finally, the work that is perhaps most closely related to the results presented in Section 3.3 is that of DeVore [97]. Specifically, it has been shown in [97] that certain deterministic constructions of $n \times p$ sensing matrices satisfy RIP of order S, provided $n = q^2$ and $p = q^{r+1}$ for some prime integer q and some integer $r \in (0, q)$, and $n = \Omega\left(S^2 \log p\right)$. Among these deterministic constructions is also a special type of a sensing matrix $\mathbf{A}$ that has the property that the elements of the matrix $\mathbf{A}$ satisfy

$$a_{i+1,j+\ell} = a_{i,j} \tag{3.73}$$

for $\ell \stackrel{def}{=} p/n$, where the arithmetic on the indices is done modulo p. However, note that such a construction leads to *block-circulant matrices* in which every row corresponds to ℓ right cyclic shifts of the row above it, as opposed to the generalization of Theorem 3.6 discussed in this section, which applies to the "traditional" circulant matrices in which every row corresponds to a *single* cyclic shift of the row above.

3.6.2 Gabor Matrices

In Section 3.4 of this chapter, we successfully established that (underdetermined or overdetermined) Gabor matrices generated from a Rademacher sequence satisfy RIP. In particular, we showed in Theorem 3.9 that Gabor matrices generated from a sequence $\mathcal{A}_k$ perform as well as Toeplitz matrices generated from the same sequence (modulo some constants), despite the fact that Gabor matrices tend to have a lot more structure to them. This is somewhat surprising, especially since the proof of Theorem 3.9 relies heavily on some of the analytical techniques developed in Section 3.3.1 in the context of Toeplitz matrices. Note that because of this very reason, some of the previous comments regarding fundamental limits of Toeplitz matrices also apply to Gabor matrices. Specifically, while Theorem 3.9 asserts that $k = \Omega(S^2 \log p)$ is sufficient to guarantee RIP for Gabor matrices, the analysis carried out in Section 3.4 is inconclusive as far as the necessity of this condition is concerned.

Finally, with regards to the connections with CS literature, the work in Section 3.4 is closely related to the recent work of Pfander et al. [98], and Herman and Strohmer [20]. Both [98] and [20] study the problem of identifying matrices that have a sparse representation in some *matrix dictionary (basis)*; [98] looks at this problem in an abstract setting, while [20] studies it from a radar perspective. In particular, the authors in [20, 98] examine—as part of their analysis—some properties of matrix dictionaries comprising of time-frequency shift matrices, which are nothing but Gabor matrices. Nevertheless, the work in [20, 98] differs from the work presented in Section 3.4 in three important respects.

First, the $n \times p$ Gabor matrices studied in [20, 98] are restricted to have the number of rows $n = \sqrt{p}$, which corresponds to $n = L = (2M + 1)$. Second, both [20, 98] assume that the columns of $\mathbf{A}$ correspond to downward *cyclic shifts* of $\mathcal{A}_k$, as opposed to downward *linear shifts* considered in this work. In other words, the $n \times L$ Toeplitz matrix $\mathbf{T}$ in (3.26) is actually an $n \times n$ circulant matrix in [20, 98]. Third, and perhaps this is the most important difference, the emphasis in [20, 98] is on finding the coherence [66] of Gabor matrices, while we focus on showing that Gabor matrices satisfy RIP. Consequently, while we have established in Section 3.4 that—given an appropriate generating sequence—Gabor matrices can successfully reconstruct *any* sparse or *approximately sparse* signal, the results in [20, 98] only guarantee successful reconstruction of *most* sparse signals supported on a *random set* $T \subset [1 \dots p]$.

3.6.3 Structurally-Subsampled Unitary Matrices

In Section 3.5 of this chapter, we introduced and analyzed a new class of structured CS matrices—termed as structurally-subsampled unitary matrices—that can be thought of as a generalization of the class of subsampled unitary matrices. In particular, we successfully established in Theorem 3.16 that structurally-subsampled unitary matrices generated from $(\mathbf{A}_p, \mathbf{U})$, with $\mathbf{A}_p$ being a Rademacher sequence, perform *nearly* as well as subsampled unitary matrices generated from the same unitary matrix $\mathbf{U}$. Specifically, Theorem 3.16 for structurally-subsampled unitary matrices differs from Theorem 2.9 for subsampled unitary matrices by only a factor of $\log p$. Note that this difference is primarily a consequence of the fact that the maximum magnitude of the entries in

a subsampled unitary matrix is trivially given by $\mu_{\mathbf{U}}/\sqrt{n}$, whereas we could only bound the maximum magnitude of the entries in structurally-subsampled unitary matrices by $\mu_{\mathbf{U}}\sqrt{\log p/n}$. However, it remains to be seen whether this is a fundamental characteristic of structurally-subsampled unitary matrices or just an artifact of the proof technique employed in Lemma 3.13. It is also instructive to note at this point that since the results for structurally-subsampled unitary matrices should coincide with those for subsampled unitary matrices for the case of $k = 1$, it is heuristically plausible to conjecture that the performance of structurally-subsampled unitary matrices should deviate from that of subsampled unitary matrices by a factor that is a function of k (instead of p). Such a conclusion, however, does not follow from the results established in Section 3.5.

Finally, we conclude this chapter with a brief discussion of the connections between the results of Section 3.5 and some existing works. As noted earlier, the work in Section 3.5 is closely related in terms of the general proof technique to the work of Romberg [91] and Tropp et al. [18] in general, and Rudelson and Vershynin [12] in particular. This is primarily a consequence of the fact that the arguments used by Rudelson and Vershynin in [12] are substantially simpler (and tighter) than, for instance, the ones used in [8] to establish RIP of subsampled matrices.

In terms of the actual problem, however, our work on structurally-subsampled unitary matrices is most closely related to the recent work of Tropp et al. [18], where they propose a sub-Nyquist sampling architecture—termed *random demodulator*—to acquire sparse bandlimited signals. In particular, it is shown in [18] that the overall action of the random demodulator on a sparse bandlimited signal can be accurately described in terms of a sensing matrix, which the authors term as a *random demodulator matrix*. However, it is easy to see from [18, Section IV-B] that a random demodulator matrix is just a structurally-subsampled unitary matrix with $\mathbf{U}$ being a Fourier matrix and $k = p/n$ (in other words, no subsampling). In this regard, our work in Section 3.5 can also be thought of a generalization of the RIP analysis of a random demodulator matrix carried out in [18]. Based on the preceding discussion, it is perhaps best to think of structurally-subsampled unitary matrices as filling the void between the two extremes of subsampled unitary matrices (maximum subsampling) and random demodulator matrices (no subsampling) through the choice of the design parameter k (with k ranging from 1 to p/n).

3.7 Appendix

3.7.1 Proof of Lemma 3.8

To establish this lemma, first define $S_{N,R} \stackrel{def}{=} \sum_{i=1}^{N} \text{Re}(x_i)$ and $S_{N,I} \stackrel{def}{=} \sum_{i=1}^{N} \text{Im}(x_i)$. Further, notice that since $|x_i| \leq a_i, i = 1, \ldots, N$, we equivalently have that $|\text{Re}(x_i)| \leq a_i, i = 1, \ldots, N$, and $|\text{Im}(x_i)| \leq a_i, i = 1, \ldots, N$. Therefore,

$$\begin{aligned}
\Pr\left(\left|S_N - \mathbb{E}[S_N]\right| \geq t\right) &= \Pr\left(\left(S_{N,R} - \mathbb{E}[S_{N,R}]\right)^2 + \left(S_{N,I} - \mathbb{E}[S_{N,I}]\right)^2 \geq t^2\right) \\
&\stackrel{(a)}{\leq} \Pr\left(\left|S_{N,R} - \mathbb{E}[S_{N,R}]\right| \geq \frac{t}{\sqrt{2}}\right) + \Pr\left(\left|S_{N,I} - \mathbb{E}[S_{N,I}]\right| \geq \frac{t}{\sqrt{2}}\right) \\
&\stackrel{(b)}{\leq} 2\exp\left(-\frac{t^2}{4\sum_{i=1}^{N} a_i^2}\right) + 2\exp\left(-\frac{t^2}{4\sum_{i=1}^{N} a_i^2}\right) \qquad (3.74)
\end{aligned}$$

Here, (a) follows from a simple union bounding argument and (b) follows from the original Hoeffding's inequality. This completes the proof of the lemma. ∎

Chapter 4

Estimation of Sparse Multipath Channels

4.1 Introduction

Wireless technology has had and continues to have a profound impact on our society. It is arguably one of the leading drivers of the Information Revolution in the 21st century and is expected to play an increasingly important role in the global economic growth as the world transitions from a manufacturing-based economy to an information-based economy. Despite having a history of more than a century of rapid technological advancements, the field of wireless communications remains far from mature. A number of key technical challenges still need to be overcome in order to realize our vision of a future with ubiquitous wireless connectivity. Foremost among these challenges is designing wireless systems that not only support data rates comparable to that of wired systems, but also enable increased mobility while maintaining constant, reliable connectivity under resource constraints—energy, latency, and bandwidth constraints being the strictest among them. Successfully addressing this and similar challenges requires significant technical advances on multiple fronts. One such front is the development of signal processing techniques for estimating multipath wireless channels using minimal resources, and this chapter summarizes the findings of some of our recent efforts in this direction.

4.1.1 Background

In a typical terrestrial environment, a radio signal emitted from a transmitter is reflected, diffracted, and scattered from the surrounding objects, and arrives at the receiver as a superposition of multiple attenuated, delayed, and phase- and/or frequency-shifted copies of the original

signal. This superposition of multiple copies of the transmitted signal, called multipath signal components, is the defining characteristic of terrestrial wireless systems, and is both a curse and a blessing from a communications viewpoint. On the one hand, this *multipath signal propagation* leads to fading—fluctuations in the received signal strength—that severely impacts the rate and reliability of communication [21, 99]. On the other hand, research in the last decade has shown that multipath propagation also results in an increase in the number of degrees of freedom (DoF) available for communication, which—if utilized effectively—can lead to significant gains in the rate (multiplexing gain) and/or reliability (diversity gain) of communication [22, 100]. The impact of fading versus diversity/multiplexing gain on performance critically depends on the amount of channel state information (CSI) available to the system. For example, knowledge of instantaneous CSI at the receiver (coherent reception) enables exploitation of delay, Doppler, and/or spatial diversity to combat fading, while further gains in rate and reliability are possible if (even partial) CSI is available at the transmitter as well [21].

In practice, CSI is seldom—if ever—available to communication systems a priori and the channel needs to be (periodically) estimated at the receiver in order to reap the benefits of additional DoF afforded by multipath propagation. As such, two classes of methods are commonly employed to estimate multipath channels at the receiver. In *training-based channel estimation* methods, the transmitter multiplexes signals that are known to the receiver (henceforth referred to as training signals) with data-carrying signals in time, frequency, and/or code domain, and CSI is obtained at the receiver from knowledge of the training and received signals. In *blind channel estimation* methods, CSI is acquired at the receiver by making use of the statistics of data-carrying signals only. Although theoretically feasible, blind estimation methods typically require complex signal processing at the receiver and often entail inversion of large data-dependent matrices, which also makes them highly prone to error propagation in rapidly-varying channels. Training-based methods, on the other hand, require relatively simple receiver processing and lead to decoupling of the data-detection module from the channel-estimation module at the receiver, which reduces receiver complexity even further. As such, training-based methods are widely prevalent in modern wireless

systems [23] and we therefore focus exclusively on them in the sequel; see [101] for an overview of blind approaches to channel estimation.

One of the first analytical studies of training-based estimation methods for multipath channels was authored by Cavers in 1991 [24]. Since then, there has been a growing body of literature devoted to the design and analysis of training-based methods for various classes of channels. These works often highlight two salient aspects of training-based methods, namely, *sensing* and *reconstruction*. Sensing corresponds to the design of training signals used by the transmitter to probe the channel, while reconstruction is the problem of processing the corresponding channel output at the receiver to recover the CSI. The ability of a training-based method to accurately estimate the channel depends critically on both the design of training signals and the application of effective reconstruction strategies. Much of the work in the channel estimation literature is based on the implicit assumption of a *rich* underlying multipath environment in the sense that the number of DoF in the channel are expected to scale linearly with the signal space dimension (product of signaling bandwidth, symbol duration, and minimum of the number of transmit and receive antennas). As a result, training-based methods proposed in such works are mainly comprised of linear reconstruction techniques, which are known to be optimal for rich multipath channels, thereby more or less reducing the problem of channel estimation to that of designing optimal training signals for various channel classes [24–32].

Numerous experimental studies undertaken by various researchers in the recent past have shown though that physical wireless channels encountered in practice tend to exhibit *sparse* structures at high signal space dimension in the sense that majority of the channel DoF end up being either zero or nearly zero when operating at large bandwidths and symbol durations and/or with large plurality of antennas [33–37]. However, traditional training-based methods—relying on linear reconstruction schemes at the receiver—seem incapable of exploiting the inherent low-dimensionality of such sparse channels, thereby leading to overutilization of the key communication resources of energy, latency, and bandwidth. Recently, a number of researchers have tried to address this problem and proposed training signals and reconstruction strategies that are tailored to the anticipated characteristics of sparse multipath channels [13, 95, 102–106]. But much of the

emphasis in these studies has been directed towards establishing the feasibility of the proposed sparse-channel estimation methods numerically rather than analytically. A major drawback of this approach is that the methods detailed in the previous investigations lack a quantitative theoretical analysis of their performance in terms of the reconstruction error.

4.1.2 Chapter Outline

By leveraging key ideas from the theory of compressed sensing (some of which were discussed in Chapter 2 and Chapter 3), we recently proposed new training-based estimation methods for various classes of sparse single- and multiple-antenna channels that are provably more effective than their traditional counterparts [14,15,107]. In particular, we analytically showed in [14,15,107] that the proposed training-based methods achieve a target reconstruction error using far less energy and, in many instances, latency and bandwidth than that dictated by the traditional methods. As in the case of previous research, the exact nature of training signals employed by our proposed methods varies with the type of signaling waveforms used for sensing (e.g., single- or multi-carrier signaling waveforms) and the class to which the underlying multipath channel belongs (e.g., frequency- or doubly-selective channel). However, a common theme underlying all our training-based methods is the use of sparsity-inducing mixed-norm optimization criteria such as the Dantzig selector and the lasso (or basis pursuit denoising) for reconstruction at the receiver. These criteria, which have arisen out of recent advances in the theory of compressed sensing, have been discussed in extensive detail in Chapter 2 of the dissertation. In the spirit of compressed sensing, we term this particular approach to estimating sparse multipath channels as *compressed channel sensing* (CCS); the analogy here being that CCS requires far fewer communication resources to estimate sparse channels than do the traditional training-based methods.

The goal of this chapter is to complement our existing work on sparse-channel estimation by providing an expanded view of the key ideas underlying the theory of CCS. In order to accomplish this goal, we focus on five specific classes of multipath channels within the chapter, namely, frequency- and doubly-selective single-antenna channels, and nonselective, frequency-selective,

and doubly-selective multiple-antenna channels. For each of these five channel classes, the discussion in the chapter focusses on the nature of the training signals used for probing a sparse channel, the reconstruction method used at the receiver for recovering the CSI, and quantification of the reconstruction error in the resulting estimate. In terms of modeling of the sparse channels within each channel class, we use a virtual representation of physical multipath channels that represents the expansion of the time-frequency response of a channel in terms of multi-dimensional Fourier basis functions. It is worth noting though that the main ideas presented in the chapter can be generalized to channel models that make use of a basis other than the Fourier one, provided the expansion basis effectively exposes the sparse nature of the underlying multipath environment and can be made available to the receiver a priori.

4.2 Multipath Wireless Channel Modeling

Signal propagation in a wireless channel over multiple spatially distributed paths gives rise to a large number of propagation parameters. However, exact knowledge of these parameters is not critical for reliable communication of data over the channel. Rather, from a purely communication-theoretic perspective, we are only interested in characterizing the *interaction* between the physical propagation environment and the spatio-temporal signal space associated with the transmitter and the receiver. In this section, we review a virtual modeling framework for multipath wireless channels that captures this interaction through Nyquist sampling of the physical propagation environment in the angle-delay-Doppler space. As we will later see, this framework plays a key role in subsequent developments in the chapter since it not only exposes the relationship between the distribution of physical paths within the angle-delay-Doppler space and the sparsity of channel DoF, but also sets the stage for the application of compressed sensing theory and methods to sparse-channel estimation.

4.2.1 Physical Characterization of Multipath Wireless Channels

Consider, without loss of generality, a multiple-antenna wireless channel with half-wavelength spaced uniform linear arrays (ULAs) at the transmitter and receiver. Let N_T and N_R denote the

number of transmit and receive antennas, respectively. It is customary to model a multipath wireless channel $\boldsymbol{\mathcal{H}}$ as a linear, time-varying system [21, 22, 99, 108]. The corresponding (complex) baseband transmitted signal and (noiseless) channel output are related as

$$\boldsymbol{\mathcal{H}}(\mathbf{x}(t)) = \int_{\mathbb{R}} \mathbf{H}(t,f)\mathbf{X}(f)e^{j2\pi ft}df \tag{4.1}$$

where $\boldsymbol{\mathcal{H}}(\mathbf{x}(t))$ is the N_R-dimensional channel output, $\mathbf{X}(f)$ is the (element-wise) Fourier transform of the N_T-dimensional transmitted signal $\mathbf{x}(t)$, and $\mathbf{H}(t,f)$ is the $N_R \times N_T$ *time-varying frequency response* matrix of the channel. Note that temporal variations of the channel frequency response arise due to the relative motion between the transmitter, receiver, and the multipath propagation environment.

Channel and Signaling Parameters

Bello did some seminal work on the characterization of linear, time-varying systems [108] and introduced an equivalent representation of $\mathbf{H}(t,f)$, termed as the *delay-Doppler spreading function* $\mathbf{C}(\nu,\tau)$, by exploiting time-frequency duality theory. Specifically, it was established in [108] that the time-varying frequency response $\mathbf{H}(t,f)$ and the delay-Doppler spreading function $\mathbf{C}(\nu,\tau)$ of a linear, time-varying system constitute a two-dimensional Fourier transform pair

$$\mathbf{H}(t,f) = \iint_{\mathbb{R}^2} \mathbf{C}(\nu,\tau)e^{j2\pi\nu t}e^{-j2\pi\tau f}d\nu d\tau \iff \mathbf{C}(\nu,\tau) = \iint_{\mathbb{R}^2} \mathbf{H}(t,f)e^{-j2\pi\nu t}e^{j2\pi\tau f}dtdf. \tag{4.2}$$

The delay-Doppler spreading function of a wireless channel can be used to define two key channel parameters:

[1] The *delay spread* of a channel, τ_{max}, is defined as the range of values of τ over which $\mathbf{C}(\nu,\tau)$ is essentially nonzero: $\mathbf{C}(\nu,\tau) = 0$ for all $\tau \notin [0,\tau_{max}]$.

[2] The (two-sided) *Doppler spread* of a channel, ν_{max}, is defined as the range of values of ν over which $\mathbf{C}(\nu,\tau)$ is essentially nonzero: $\mathbf{C}(\nu,\tau) = 0$ for all $\nu \notin [-\nu_{max}/2, \nu_{max}/2]$.

Note that since $\mathbf{H}(t,f)$ is a time-frequency dual of $\mathbf{C}(\nu,\tau)$, the Doppler spread and the delay spread of a channel are a measure of the variations of $\mathbf{H}(t,f)$ in time and frequency. Specifically,

Table 4.1 Classification of wireless channels on the basis of channel and signaling parameters

Channel Classification	$W\tau_{max}$	$T\nu_{max}$
Nonselective Channels	$\ll 1$	$\ll 1$
Frequency-Selective Channels	≥ 1	$\ll 1$
Time-Selective Channels	$\ll 1$	≥ 1
Doubly-Selective Channels	≥ 1	≥ 1

the larger the values of ν_{max} and τ_{max}, the faster the variations of $\mathbf{H}(t, f)$ in time and frequency, respectively, and vice versa. The product of the delay spread and the Doppler spread, $\tau_{max}\nu_{max}$, is termed as the *spread factor* of the channel. The channel spread factor is a measure of the area of channel's (rectangular) support in the delay-Doppler space and estimating a channel having $\tau_{max}\nu_{max} > 1$ can often be an ill-posed problem, even in the absence of noise [109, 110]. Instead, we limit the discussion in this chapter to *underspread channels*, characterized by $\tau_{max}\nu_{max} \ll 1$, which is fortunately true of most wireless channels [99].[1]

Finally, throughout this chapter we implicitly consider signaling over wireless channels using symbols of duration T and (two-sided) bandwidth W, $\mathbf{x}(t) = \mathbf{0}_{N_T}$ for all $t \notin [0, T]$ and $\mathbf{X}(f) = \mathbf{0}_{N_T}$ for all $f \notin [-W/2, W/2]$, thereby giving rise to a *temporal signal space* of dimension $N_o \stackrel{def}{=} TW$ [112]. In addition, we assume that $T \gg \tau_{max}$ and $W \gg \nu_{max}$ so that intersymbol interference in time and frequency is negligible. Note that the signaling parameters T and W, together with the delay spread and the Doppler spread of a channel, can be used to broadly classify wireless channels as nonselective, frequency selective, time selective, or doubly selective; see Table 4.1 for a definition of each of these channel classes. As noted earlier, we limit ourselves in the sequel to primarily discussing frequency-selective and doubly-selective channels in the single-antenna setting ($N_T = N_R = 1$) and to nonselective, frequency-selective, and doubly-selective channels in the multiple-antenna setting.

[1] It is worth mentioning here though that part of the discussion in this chapter is also applicable to underwater acoustic communication channels, even though they may not be underspread [111].

Discrete-Path Physical Model

Due to scattering from objects in the surrounding environment, multipath wireless channels give rise to multiple attenuated, delayed, Doppler- and phase-shifted copies of the transmitted signal at the receiver. This physical reality can be accurately reflected through the use of a discrete-path channel model in which $\mathbf{H}(t,f)$ is expressed in terms of the underlying physical paths as

$$\mathbf{H}(t,f) = \sum_{n=1}^{N_p} \beta_n \mathbf{a}_R(\theta_{R,n}) \mathbf{a}_T^{\mathsf{H}}(\theta_{T,n}) e^{j2\pi\nu_n t} e^{-j2\pi\tau_n f} \tag{4.3}$$

$$\Rightarrow \quad \boldsymbol{\mathcal{H}}(\mathbf{x}(t)) = \sum_{n=1}^{N_p} \beta_n \mathbf{a}_R(\theta_{R,n}) \mathbf{a}_T^{\mathsf{H}}(\theta_{T,n}) e^{j2\pi\nu_n t} \mathbf{x}(t-\tau_n) \tag{4.4}$$

corresponding to signal propagation from the transmitter to the receiver over N_p physical paths; here, β_n is the complex path gain, $\theta_{R,n}$ the normalized angle of arrival (AoA) at the receiver, $\theta_{T,n}$ the normalized angle of departure (AoD) at the transmitter, $\tau_n \in [0, \tau_{max}]$ the (relative) delay, and $\nu_n \in [-\nu_{max}/2, \nu_{max}/2]$ the Doppler shift associated with the n-th path.

The $N_T \times 1$ vector $\mathbf{a}_T(\theta_T)$ and the $N_R \times 1$ vector $\mathbf{a}_R(\theta_R)$ in (4.3) denote the array steering and response vectors, respectively, for transmitting/receiving a signal in the direction θ_T/θ_R. These vectors are periodic in θ with unit period and are given by [113–115]

$$\mathbf{a}_T(\theta_T) \stackrel{def}{=} \frac{1}{\sqrt{N_T}} \begin{bmatrix} 1 & e^{-j2\pi\theta_T} & \dots & e^{-j2\pi(N_T-1)\theta_T} \end{bmatrix}^{\mathsf{T}} \tag{4.5}$$

$$\mathbf{a}_R(\theta_R) \stackrel{def}{=} \frac{1}{\sqrt{N_R}} \begin{bmatrix} 1 & e^{-j2\pi\theta_R} & \dots & e^{-j2\pi(N_R-1)\theta_R} \end{bmatrix}^{\mathsf{T}} \tag{4.6}$$

while the normalized AoD and AoA, θ_T and θ_R, are related to the physical (measured with respect to array broadside) AoD and AoA, ϕ_T and ϕ_R, as

$$\theta_T = \frac{\sin(\phi_T)}{2} \quad \text{and} \quad \theta_R = \frac{\sin(\phi_R)}{2}\,, \tag{4.7}$$

respectively, under the assumption of half-wavelength spaced ULAs at the transmitter and receiver. We further assume maximum *angular spread* for physical AoDs and AoAs, $(\phi_{T,n}, \phi_{R,n}) \in [-\pi/2, \pi/2] \times [-\pi/2, \pi/2]$, which means that $(\theta_{T,n}, \theta_{R,n}) \in [-1/2, 1/2] \times [-1/2, 1/2]$ and that there is a one-to-one correspondence between the physical and the normalized angles.

4.2.2 Virtual Representation of Multipath Wireless Channels

While the discrete-path channel model (4.3) is an accurate reflection of the physical reality, it is difficult to analyze and estimate owing to its *nonlinear* dependence on a potentially large number of parameters $\{(\beta_n, \theta_{R,n}, \theta_{T,n}, \tau_n, \nu_n)\}$. However, while accurate (nonlinear) estimation of AoAs, AoDs, delays, and Doppler shifts is critical in radar applications, it is not crucial in a communications context. Instead, because of the finite (transmit and receive) array apertures, signaling bandwidth, and symbol duration, $\mathbf{H}(t, f)$ can be well-approximated by a linear (in parameters) counterpart, known as a *virtual channel model*, with the aid of a four-dimensional Fourier series expansion [108, 115–118].

Sampling in Angle-Delay-Doppler

On an abstract level, virtual representation of a multipath wireless channel $\boldsymbol{\mathcal{H}}$ provides a discretized approximation of its time-varying frequency response by uniformly sampling the angle-delay-Doppler space at the Nyquist rate:

$$\Delta\theta_R \stackrel{def}{=} 1/N_R \, , \; \Delta\theta_T \stackrel{def}{=} 1/N_T \, , \; \Delta\tau \stackrel{def}{=} 1/W \, , \; \Delta\nu \stackrel{def}{=} 1/T \, . \tag{4.8}$$

Specifically, note that even though the physical time-varying frequency response $\mathbf{H}(t, f)$ may exhibit arbitrarily high spectral and temporal variations, the receiver *effectively* sees only its restriction to $[0, T]$ and $[-W/2, W/2]$ in time and frequency, respectively, $\mathbf{H}(t, f)1_{[-W/2,W/2]}(f)1_{[0,T]}(t)$, because of the finite signaling bandwidth and symbol duration. Further, the restricted channel response $\mathbf{H}(t, f)1_{[-W/2,W/2]}(f)1_{[0,T]}(t)$ also admits a four-dimensional Fourier series representation, since it has a finite support in the spatio-spectral-temporal space.

Mathematically speaking, the virtual representation of $\boldsymbol{\mathcal{H}}$, given by [108, 115–118]

$$\widetilde{\mathbf{H}}(t, f) = \sum_{i=1}^{N_R} \sum_{k=1}^{N_T} \sum_{\ell=0}^{L-1} \sum_{m=-M}^{M} H_v(i, k, \ell, m) \mathbf{a}_R \left(\frac{i}{N_R} \right) \mathbf{a}_T^{\mathsf{H}} \left(\frac{k}{N_T} \right) e^{j2\pi \frac{m}{T} t} e^{-j2\pi \frac{\ell}{W} f} \tag{4.9}$$

is simply a truncation of the Fourier series expansion of $\mathbf{H}(t, f)1_{[-W/2,W/2]}(f)1_{[0,T]}(t)$, where the truncation is justified on the basis of the fact that the Fourier series coefficients $\{H_v(i, k, \ell, m)\}$ of $\mathbf{H}(t, f)1_{[-W/2,W/2]}(f)1_{[0,T]}(t)$ are significantly nonvanishing only for $(i, k, \ell, m) \in [1 \ldots N_R] \times$

$[1 \dots N_T] \times [0 \dots L-1] \times [-M \dots M]$. As such, the virtual representation $\widetilde{\mathbf{H}}(t,f)$ approximates the channel time-varying frequency response $\mathbf{H}(t,f)$ in the sense that

$$\boldsymbol{\mathcal{H}}(\mathbf{x}(t)) = \int_{\mathbb{R}} \mathbf{H}(t,f)\mathbf{X}(f)e^{j2\pi ft}df \approx \int_{\mathbb{R}} \widetilde{\mathbf{H}}(t,f)\mathbf{X}(f)e^{j2\pi ft}df \; . \tag{4.10}$$

It is interesting to note that due to the fixed angle-delay-Doppler sampling of $\mathbf{H}(t,f)$, which defines the spatio-spectral-temporal Fourier basis functions in (4.9), $\widetilde{\mathbf{H}}(t,f)$ is a *linear* channel representation that is completely characterized by the *virtual channel coefficients* $\{H_v(i,k,\ell,m)\}$. These virtual channel coefficients can be computed from $\mathbf{H}(t,f)$ as

$$H_v(i,k,\ell,m) = \frac{1}{TW}\int_0^T \int_{-\frac{W}{2}}^{\frac{W}{2}} \mathbf{a}_R^{\mathsf{H}}\left(\frac{i}{N_R}\right)\mathbf{H}(t,f)\,\mathbf{a}_T\left(\frac{k}{N_T}\right)e^{-j2\pi\frac{m}{T}t}e^{j2\pi\frac{\ell}{W}f}dtdf \tag{4.11}$$

and the total number of these channel coefficients from (4.9) is given by $D \stackrel{def}{=} N_R N_T L(2M+1)$, where $N_R, N_T, L \stackrel{def}{=} \lceil W\tau_{max} \rceil + 1$, and $M \stackrel{def}{=} \lceil T\nu_{max}/2 \rceil$ represent the maximum number of *resolvable* AoAs, AoDs, delays, and (one-sided) Doppler shifts within the angle-delay-Doppler spread of the channel, respectively.[2] In the literature, L and M are sometimes also referred to as the *discrete* delay spread and (one-sided) *discrete* Doppler spread of the channel, respectively.

Path Partitioning in Angle-Delay-Doppler

An important and insightful property of the virtual representation $\widetilde{\mathbf{H}}(t,f)$ is that its coefficients $\{H_v(i,k,\ell,m)\}$ partition the N_p physical propagation paths into approximately disjoint subsets. Specifically, define the following subsets of paths, associated with each coefficient $H_v(i,k,\ell,m)$, based on the resolution in angle, delay, and Doppler:

$$\begin{aligned}
S_{R,i} &\stackrel{def}{=} \{n : \theta_{R,n} \in (i/N_R - 1/2N_R, i/N_R + 1/2N_R]\}\,, \\
S_{T,k} &\stackrel{def}{=} \{n : \theta_{T,n} \in (k/N_T - 1/2N_T, k/N_T + 1/2N_T]\}\,, \\
S_{\tau,\ell} &\stackrel{def}{=} \{n : \tau_n \in (\ell/W - 1/2W, \ell/W + 1/2W]\}\,, \\
S_{\nu,m} &\stackrel{def}{=} \{n : \nu_n \in (m/T - 1/2T, m/T + 1/2T]\}\,.
\end{aligned} \tag{4.12}$$

[2] With a slight abuse of *ceiling* notation, we use the convention here that $\lceil W\tau_{max} \rceil = 0$ and $\lceil T\nu_{max}/2 \rceil = 0$ for $W\tau_{max} \ll 1$ and $T\nu_{max} \ll 1$, respectively.

For example, $S_{\tau,\ell}$ denotes the set of all paths whose *physical* delays lie within the resolution bin $(\ell/W - 1/2W, \ell/W + 1/2W]$ of size $\Delta\tau$ centered around the ℓ-th *virtual* delay $\hat{\tau}_\ell \stackrel{def}{=} \ell/W$ in (4.9). Then, by substituting the discrete-path physical model (4.3) in (4.11), it can be shown that the virtual channel coefficients in (4.11) are related to the physical propagation paths as [115, 118]

$$H_v(i,k,\ell,m) = \sum_{n=1}^{N_p} \beta_n \, f_{N_R}(i/N_R - \theta_{R,n}) f^*_{N_T}(k/N_T - \theta_{T,n}) e^{-j\pi(m - T\nu_n)} \times \\ \times \; \mathrm{sinc}(m - T\nu_n)\,\mathrm{sinc}(\ell - W\tau_n) \approx \sum_{n \in S_{R,i} \cap S_{T,k} \cap S_{\tau,\ell} \cap S_{\nu,m}} \beta_n \quad (4.13)$$

where a phase and attenuation factor has been absorbed in the β_n's in (4.13). The smoothing kernels $f_{N_R}(\theta_R)$ and $f_{N_T}(\theta_T)$ in (4.13) are the Dirichlet kernels, defined as

$$f_N(\theta) \stackrel{def}{=} \frac{1}{N} \sum_{i=0}^{N-1} e^{-j2\pi i\theta} \quad (4.14)$$

while the sinc kernel is defined as $\mathrm{sinc}(x) \stackrel{def}{=} \sin(\pi x)/(\pi x)$. Note that the approximation in (4.13) is due to the sidelobes of the Dirichlet and sinc kernels induced by the finite signaling parameters, and the approximation gets more accurate with increasing T, W, N_R, and N_T.

The relation (4.13) signifies that each virtual channel coefficient $H_v(i,k,\ell,m)$ is approximately equal to the sum of the complex gains of all physical paths whose angles, delays, and Doppler shifts lie within an *angle-delay-Doppler resolution bin* of size $\Delta\theta_R \times \Delta\theta_T \times \Delta\tau \times \Delta\nu$ centered around the sampling point $(\hat{\theta}_{R,i}, \hat{\theta}_{T,k}, \hat{\tau}_\ell, \hat{\nu}_m) \stackrel{def}{=} (i/N_R, k/N_T, \ell/W, m/T)$ in the angle-delay-Doppler space. In other words, the virtual representation $\widetilde{\mathbf{H}}(t,f)$ effectively captures the underlying multipath environment comprising of N_p physical paths through D resolvable paths, thereby reducing the task of estimating a multipath wireless channel $\boldsymbol{\mathcal{H}}$ to that of reconstructing its virtual channel coefficients $\{H_v(i,k,\ell,m)\}$.

Special Case: Frequency-Selective Single-Antenna Channels

The development in this chapter has so far been carried out for the general case of a multipath wireless channel with arbitrary delay and Doppler spreads, and possibly having multiple antennas

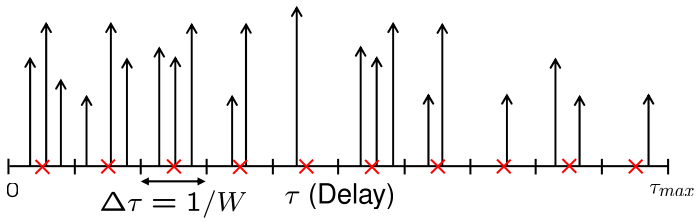

Figure 4.1 Schematic illustration of the virtual representation of a frequency-selective single-antenna channel. Each physical propagation path has associated with it a complex gain β_n (represented by an impulse of height $|\beta_n|$ in the delay space) and a delay $\tau_n \in [0, \tau_{max}]$. The virtual channel coefficients $\{H_v(\ell)\}$ correspond to samples of a smoothed version of the channel response taken at the virtual delays $\{\hat{\tau}_\ell = \ell/W\}$ in the delay space (represented by $\times$'s in the schematic).

at the transmitter and/or receiver. Nevertheless, the preceding discussion and analysis can be easily specialized to any class of multipath wireless channels.

As an illustration, first consider the special case of a frequency-selective single-antenna channel: $N_T = N_R = 1, W\tau_{max} > 1$, and $T\nu_{max} \ll 1$. Since $T\nu_{max} \ll 1$ in this case, there is negligible temporal variation of the channel frequency response over the symbol duration, $H(t, f) \approx H(f)$, and the physical channel model and its corresponding virtual representation reduce to [cf. (4.3), (4.4), (4.9), and (4.10)]

$$\mathcal{H}(x(t)) = \sum_{n=1}^{N_p} \beta_n x(t - \tau_n) \approx \sum_{\ell=0}^{L-1} H_v(\ell) x\left(t - \frac{\ell}{W}\right). \tag{4.15}$$

Further, as illustrated in Figure 4.1, the virtual channel coefficients $\{H_v(\ell)\}$ in this case are related to the physical propagation paths as [cf. (4.13)]

$$H_v(\ell) = \sum_{n=1}^{N_p} \beta_n \text{sinc}(\ell - W\tau_n) \approx \sum_{n \in S_{\tau,\ell}} \beta_n \,. \tag{4.16}$$

In other words, a frequency-selective single-antenna channel resolves physical paths in the delay space only and the total number of the resolvable paths (channel coefficients) in this case is given by $D = L = \lceil W\tau_{max} \rceil + 1$.

Special Case: Nonselective Multiple-Antenna Channels

Next, consider the special case of a nonselective multiple-antenna channel: N_T and/or $N_R > 1$, and $W\tau_{max}$ and $T\nu_{max} \ll 1$. In this case, there is negligible spectral and temporal variation of the

channel frequency response over the signaling bandwidth and symbol duration, $\mathbf{H}(t,f) \approx \mathbf{H}$, and the physical channel model and its corresponding virtual representation reduce to

$$\boldsymbol{\mathcal{H}}(\mathbf{x}(t)) = \sum_{n=1}^{N_p} \beta_n \mathbf{a}_R(\theta_{R,n}) \mathbf{a}_T^{\mathsf{H}}(\theta_{T,n}) \mathbf{x}(t) \approx \sum_{i=1}^{N_R} \sum_{k=1}^{N_T} H_v(i,k) \mathbf{a}_R\left(\frac{i}{N_R}\right) \mathbf{a}_T^{\mathsf{H}}\left(\frac{k}{N_T}\right) \mathbf{x}(t) \ . \quad (4.17)$$

The relation (4.17) signifies that a nonselective multiple-antenna channel resolves physical paths in the transmit- and receive-angle space only. The total number of resolvable paths (channel coefficients) in this case is given by $D = N_R N_T$, while the relationship between the virtual channel coefficients $\{H_v(i,k)\}$ and the physical paths is given by

$$H_v(i,k) = \sum_{n=1}^{N_p} \beta_n \, f_{N_R}(i/N_R - \theta_{R,n}) f_{N_T}^*(k/N_T - \theta_{T,n}) \approx \sum_{n \in S_{R,i} \cap S_{T,k}} \beta_n \ . \quad (4.18)$$

4.3 Sparse Multipath Wireless Channels

4.3.1 Modeling

The virtual representation (4.9) of a multipath wireless channel $\boldsymbol{\mathcal{H}}$ signifies that the maximum number of DoF in the channel is

$$D = N_R N_T L(2M+1) \approx \tau_{max} \nu_{max} N_R N_T T W \quad (4.19)$$

which corresponds to the maximum number of angle-delay-Doppler resolution bins in the virtual representation, and reflects the maximum number of resolvable paths within the four-dimensional channel spread. However, the level of diversity and/or multiplexing gain afforded by $\boldsymbol{\mathcal{H}}$ is governed by the number of nonvanishing $\{H_v(i,k,\ell,m)\}$. Therefore, the *actual* number of DoF, d, in the channel corresponds to the number of nonvanishing virtual channel coefficients.

Definition 4.1 (Channel Degrees of Freedom) Let $d \stackrel{def}{=} \big|\{(i,k,\ell,m) \, : \, |H_v(i,k,\ell,m)| > 0\}\big|$ denote the number of nonvanishing virtual channel coefficients that significantly contribute to the channel power. Then d reflects the actual (or effective) number of DoF in the channel.

Remark 4.2 In a communications context, "nonvanishing" is a relative term that can only be qualified on the basis of the operating received signal-to-noise ratio (SNR). For the sake of this exposition, and because it suffices to illustrate the principles, we make the simplified assumption that all but d of the virtual channel coefficients are identically zero.

Trivially, we have that the channel DoF $d \leq D$; also, by virtue of (4.13), $d = D$ if there are at least $N_p \geq D$ physical paths distributed in a way within the channel spread such that each angle-delay-Doppler resolution bin is populated by at least one path (see Figure 4.2).

Much of the work in the existing channel estimation literature is based on the implicit assumption of a rich scattering environment in which there are sufficiently many paths uniformly distributed within the angle-delay-Doppler spread of the channel so that $d \approx D$ for any choice of the signaling parameters. This assumption can be traced back to the seminal works of Bello [108] and Kennedy [119] on the wide-sense stationary uncorrelated scattering channel model, and more recently to the independent and identically distributed (i.i.d.) model for multiple-antenna channels proposed by Telatar [120], and Foschini and Gans [121].

However, numerous past and recent channel measurement campaigns have shown that propagation paths in many physical channels tend to be distributed as clusters within their respective channel spreads [33–37]. These clusters of paths physically correspond to large-scale objects in the scattering environment (e.g., buildings and hills in an outdoor propagation environment), while multipath components within a cluster arise as a result of scattering from small-scale structures of the corresponding large-scale reflector (e.g., windows of a building or trees on a hill). As we vary the spatio-temporal signaling parameters in such multipath-cluster channels by increasing the number of antennas, signaling bandwidth, and/or symbol duration, a point comes where $\Delta\theta_R, \Delta\theta_T, \Delta\tau$, and/or $\Delta\nu$ become smaller than the interspacings between the multipath clusters, thereby leading to the situation depicted in Figure 4.2 where not every resolution bin of size $\Delta\theta_R \times \Delta\theta_T \times \Delta\tau \times \Delta\nu$ contains a physical path. This implies that wireless channels with clustered multipath components tend to have far fewer than D nonzero virtual channel coefficients when operated at large bandwidths and symbol durations and/or with large plurality of antennas. We refer to such wireless

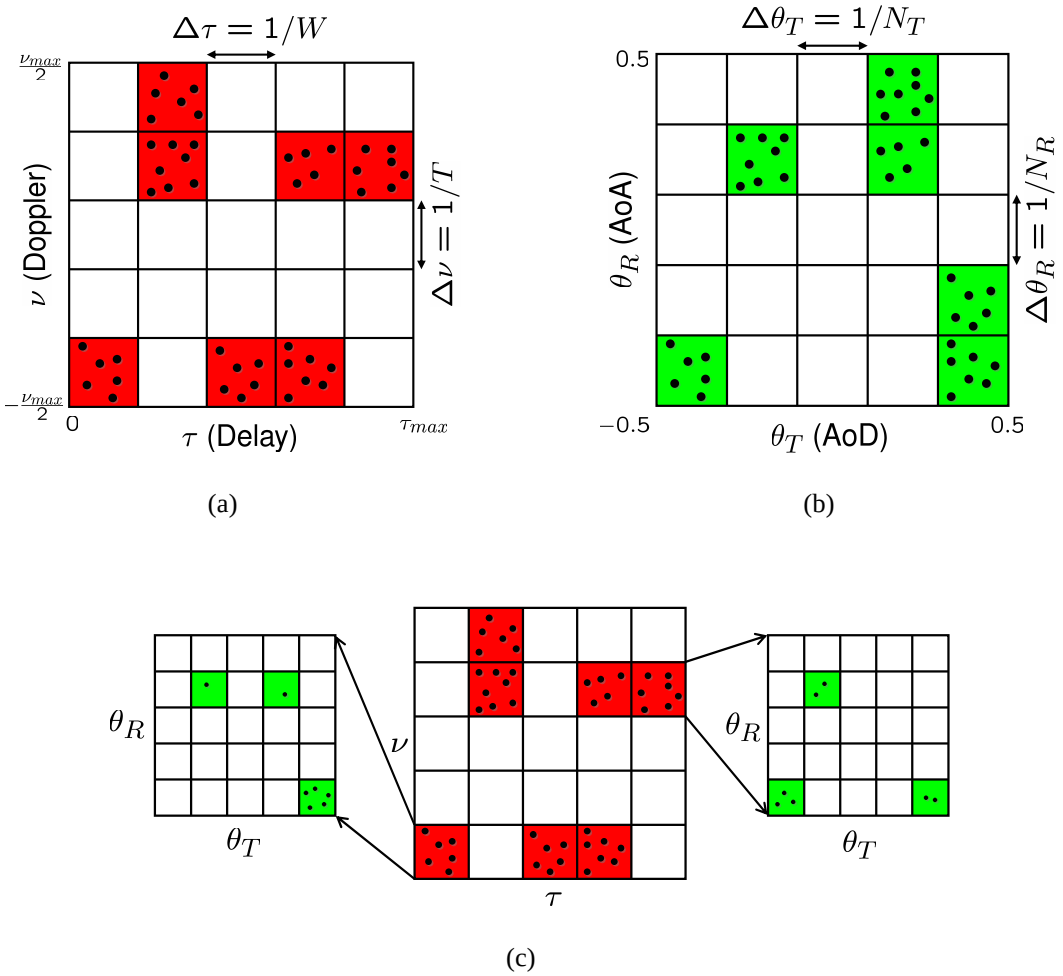

Figure 4.2 Schematic illustration of the virtual channel representation (VCR) and the channel sparsity pattern (SP). Each square represents a resolution bin associated with a distinct virtual channel coefficient. The total number of these squares equals D. The shaded squares represent the SP, $\mathcal{S}_d$, corresponding to the $d \ll D$ nonzero channel coefficients, and the dots represent the paths contributing to each nonzero coefficient. (a) VCR and SP in delay-Doppler: $\{H_v(\ell, m)\}_{\mathcal{S}_d}$. (b) VCR and SP in angle: $\{H_v(i, k)\}_{\mathcal{S}_d}$. (c) VCR and SP in angle-delay-Doppler: $\{H_v(i, k, \ell, m)\}_{\mathcal{S}_d}$. The paths contributing to a fixed nonzero delay-Doppler coefficient, $H_v(\ell_o, m_o)$, are further resolved in angle to yield the conditional SP in angle: $\{H_v(i, k, \ell_o, m_o)\}_{\mathcal{S}_d(\ell_o, m_o)}$.

channels as *sparse multipath channels* and formalize this notion of multipath sparsity in the following definition.

Definition 4.3 (Sparse Multipath Wireless Channels) Let d denote the number of DoF of a multipath wireless channel $\boldsymbol{\mathcal{H}}$. We say that the channel is d-sparse if it satisfies $d \ll D$.

Sparse multipath channels correspond to a sparse distribution of resolvable paths in the angle-delay-Doppler space. Sparsity in angle-delay-Doppler leads to *correlation* or *coherence* in space-frequency-time due to the duality between the angle-delay-Doppler and the space-frequency-time domains. The nature of this channel correlation in space, time, and frequency is influenced by the *locations* of the d nonzero channel coefficients within the D resolution bins in the angle-delay-Doppler space. This information about a multipath wireless channel can be captured through the notion of the *channel sparsity pattern*.

Definition 4.4 (Channel Sparsity Pattern) The set $\mathcal{S}_d \stackrel{def}{=} \big\{(i,k,\ell,m) : |H_v(i,k,\ell,m)| > 0\big\}$ is termed as the channel sparsity pattern of a d-sparse channel. That is, $\mathcal{S}_d$ is the set of indices of the $d = |\mathcal{S}_d|$ nonzero virtual channel coefficients of $\boldsymbol{\mathcal{H}}$.

Essentially, the sparsity pattern $\mathcal{S}_d$ characterizes the d-dimensional subspace of the D-dimensional channel space that is excited by the d nonzero virtual channel coefficients $\{H_v(i,k,\ell,m)\}_{\mathcal{S}_d}$ representing the DoF in the channel (see Figure 4.2). This means that the instantaneous CSI of sparse channels is completely characterized by $\{H_v(i,k,\ell,m)\}_{\mathcal{S}_d}$.

Finally, while statistical characterization of a sparse multipath channel $\boldsymbol{\mathcal{H}}$ is critical from a communication-theoretic viewpoint, either Bayesian (random) or non-Bayesian formulation of $\boldsymbol{\mathcal{H}}$ suffice from the channel estimation perspective. In this chapter, we stick to the non-Bayesian paradigm and assume that both the channel sparsity pattern $\mathcal{S}_d$ and the corresponding virtual coefficients $\{H_v(i,k,\ell,m)\}_{\mathcal{S}_d}$ are deterministic but unknown.

Special Case: Delay Sparsity

In order to further motivate the idea of multipath sparsity, take the simple example of a single-antenna transmitter-receiver pair communicating at large bandwidth in a static environment. The

underlying (frequency-selective) multipath channel in this setting is best described as a linear, time-invariant system whose channel response consists of only a few dominant echoes (due to the reflections of the transmitted signal from the surrounding objects such as buildings and hills).

In this case, the channel frequency response is given by $H(f) = \sum_{n=1}^{N_p} \beta_n e^{-j2\pi\tau_n f}$ and its virtual representation can be expressed as $\widetilde{H}(f) = \sum_{\ell=0}^{L-1} H_v(\ell) e^{-j2\pi\frac{\ell}{W}f}$. Here, $H_v(\ell)$ is approximately equal to the sum of gains of all echoes (paths) whose delays lie within the resolution bin $B_{\tau,\ell} = (\frac{\ell}{W} - \frac{1}{2W}, \frac{\ell}{W} + \frac{1}{2W}]$ and there are a total of $D\ (= L) = \lceil W\tau_{max} \rceil + 1$ of these virtual channel coefficients. However, since the channel consists of only a few dominant echoes (or clusters of echoes), a large number of the resolution bins $\{B_{\tau,\ell}\}$ would contain no echoes under the assumption of large-enough signaling bandwidth. Therefore, majority of the channel coefficients in this case would be zero, $d = \big|\{\ell : |H_v(\ell)| > 0\}\big| \ll L$, and we accordingly term the underlying multipath channel as d-sparse.

Remark 4.5 It is worth mentioning at this point that sparsity in multipath wireless channels is inherently tied to the choice of signaling parameters: channels with small-enough values of N_R, N_T, T, and W are bound to have $d \approx D$. Consider, for example, the case of the so-called Brazil B channel [122]—a digital television channel with six dominant echoes and a delay spread of 12.7 μs—as an illustration of this fact. Figure 4.3(b) and Figure 4.3(c) compare two virtual representations of a particular realization of the Brazil B channel [see Figure 4.3(a)] under the signaling bandwidths of 25 MHz and 5 MHz, respectively. For the sake of this illustration, we have assumed an operating received SNR of 30 dB and shown only the channel coefficients in the two representations that are above the corresponding noise floor. It can be easily seen from the two figures that while only 14% of the virtual channel coefficients are nonvanishing at $W = 25$ MHz, this number increases to 78% when the signaling bandwidth is reduced to 5 MHz. Nevertheless, the trend in modern wireless systems is to operate at high spatio-temporal signal space dimension (defined as: $N_s \stackrel{def}{=} \min\{N_T, N_R\}\, TW$). As such, sparse channels are becoming more and more ubiquitous in today's communications landscape.

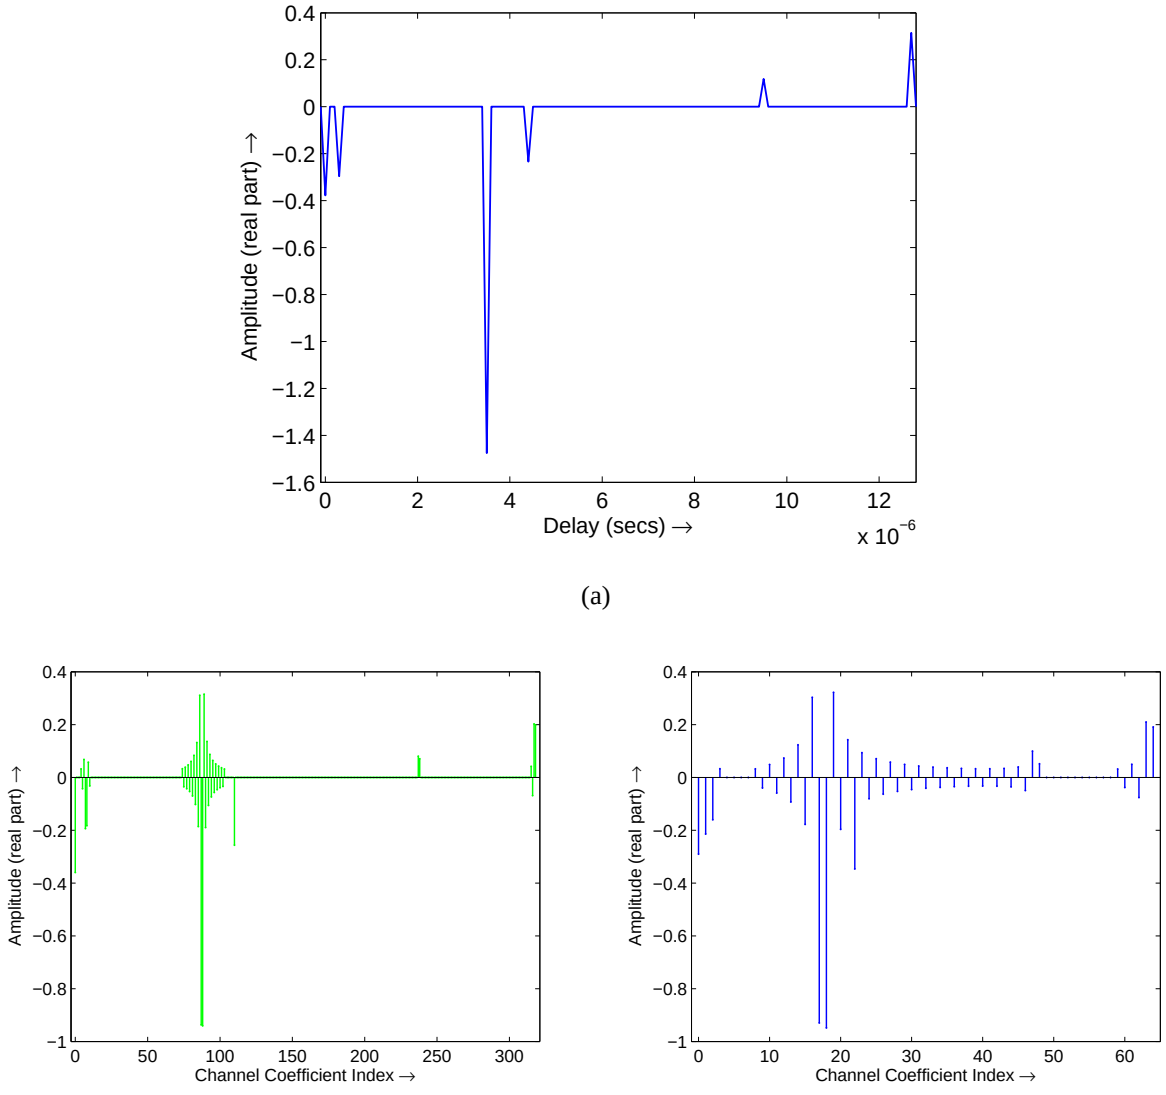

(a)

(b)

(c)

Figure 4.3 Sparsity comparisons of the virtual representations of Brazil B channel [122] under two different signaling bandwidths (only real parts of the complex gains are shown). (a) Realization of the physical response of Brazil B channel in the delay space ($\tau_1 = 0$ μs, $\tau_2 = 0.3$ μs, $\tau_3 = 3.5$ μs, $\tau_4 = 4.4$ μs, $\tau_5 = 9.5$ μs, and $\tau_6 = 12.7$ μs). (b) Virtual channel representation corresponding to $W = 25$ MHz. (c) Virtual channel representation corresponding to $W = 5$ MHz.

4.3.2 Sensing

In wireless systems that rely on training-based methods for channel estimation, the transmission of a symbol takes the form

$$\mathbf{x}(t) = \mathbf{x}_{tr}(t) + \mathbf{x}_{data}(t)\,,\ 0 \leq t \leq T \tag{4.20}$$

where $\mathbf{x}_{tr}(t)$ and $\mathbf{x}_{data}(t)$ represent the training signal and data-carrying signal, respectively. Because of the linearity of $\boldsymbol{\mathcal{H}}$, and under the assumption of $\mathbf{x}_{tr}(t)$ being orthogonally multiplexed with $\mathbf{x}_{data}(t)$ in time, frequency, and/or code domain, the resulting signal at the receiver can be partitioned into two noninterfering components: one corresponding to $\mathbf{x}_{tr}(t)$ and the other corresponding to $\mathbf{x}_{data}(t)$. In order to estimate $\boldsymbol{\mathcal{H}}$, training-based methods ignore the received data and focus only on the training component of the received signal, given by

$$\mathbf{y}_{tr}(t) = \boldsymbol{\mathcal{H}}(\mathbf{x}_{tr}(t)) + \mathbf{z}_{tr}(t)\,,\ 0 \leq t \leq T + \tau_{max} \tag{4.21}$$

where $\mathbf{z}_{tr}(t)$ is an N_R-dimensional complex additive white Gaussian noise (AWGN) signal introduced by the receiver circuitry.

As a first step towards estimating $\boldsymbol{\mathcal{H}}$, the (noisy) received training signal $\mathbf{y}_{tr}(t)$ is matched filtered with the transmitted waveforms at the receiver to obtain an equivalent discrete-time representation of (4.21). The exact form of this representation depends on a multitude of factors such as selectivity of the channel (nonselective, frequency selective, etc.), type of the signaling waveform used for sensing (single- or multi-carrier), and number of transmit and receive antennas. While this gives rise to a large number of possible scenarios to be examined, each one corresponding to a different combination of these factors, it is shown in Section 4.5 and Section 4.6 that algebraic manipulations of the matched-filtered output in each case result in the following general linear form at the receiver

$$\underbrace{\begin{bmatrix}\mathbf{y}_1 & \dots & \mathbf{y}_{N_R}\end{bmatrix}}_{\mathbf{Y}} = \sqrt{\frac{\mathcal{E}}{N_T}}\,\mathbf{X}\,\underbrace{\begin{bmatrix}\mathbf{h}_{v,1} & \dots & \mathbf{h}_{v,N_R}\end{bmatrix}}_{\mathbf{H}_v} + \underbrace{\begin{bmatrix}\mathbf{z}_1 & \dots & \mathbf{z}_{N_R}\end{bmatrix}}_{\mathbf{Z}}\,. \tag{4.22}$$

The quantity $\mathcal{E}$ here denotes the total transmit energy budget for training purposes, defined as

$$\mathcal{E} \stackrel{def}{=} \int_0^T \left\|\mathbf{x}_{tr}(t)\right\|_2^2 dt \tag{4.23}$$

while $\mathcal{E}/N_T$ is defined as the *average* training energy budget *per transmit antenna*. Further, the vectors $\mathbf{h}_{v,i}, i = 1, \ldots, N_R$, in (4.22) are $N_T L(2M+1)$-dimensional complex vectors comprising of the channel coefficients $\{H_v(i,k,\ell,m)\}$, and we let the AWGN matrix $\mathbf{Z}$ have zero-mean, unit-variance, independent complex-Gaussian entries. In other words, $\mathcal{E}$ is taken to be a measure of the training SNR at each receive antenna. Finally, the *sensing matrix* $\mathbf{X}$ in (4.22) is a complex-valued matrix having $D/N_R = N_T L(2M+1)$ columns that are normalized in a way such that $\|\mathbf{X}\|_F^2 = D/N_R$. The exact form and dimensions of $\mathbf{X}$ (and hence the dimensions of $\mathbf{Y}$ and $\mathbf{Z}$) are completely determined by $\mathbf{x}_{tr}(t)$ and the class to which $\boldsymbol{\mathcal{H}}$ belongs; concrete representations of $\mathbf{X}$ corresponding to the various training signals and channel configurations studied in the chapter can be found in Section 4.5 and Section 4.6.

Example 4.6 (Sensing of Frequency-Selective Single-Antenna Channels) In order to further illustrate the ideas presented in this section, consider again the case of a frequency-selective single-antenna channel. The general linear form (4.22) in this case can be expressed as

$$\mathbf{y} = \sqrt{\mathcal{E}}\,\mathbf{X}\mathbf{h}_v + \mathbf{z} \tag{4.24}$$

where $\mathbf{h}_v$ is an L-dimensional complex vector consisting of the channel coefficients $\{H_v(\ell)\}$. In addition, due to the linear, time-invariant nature of the channel, it is easy to see that if single-carrier waveforms are used for signaling in this setting then $\mathbf{X}$ has a Toeplitz form. On the other hand, if multi-carrier waveforms are used for signaling then $\mathbf{X}$ corresponds to a (scaled) submatrix of an $N_o \times N_o$ unitary discrete Fourier transform (DFT) matrix. The two signaling scenarios are discussed further in extensive details in Section 4.5.1.

4.3.3 Reconstruction

As noted in Section 4.1.1, training-based channel estimation methods are characterized by the two distinct—but highly intertwined—operations of sensing and reconstruction. The reconstruction aspect of a training-based method involves designing either a linear or a nonlinear procedure that produces an estimate of $\mathbf{H}_v$ at the receiver from the knowledge of $\mathcal{E}, \mathbf{X}$, and $\mathbf{Y}$:

$$\mathbf{H}_v^{\mathsf{est}} \stackrel{def}{=} \mathbf{H}_v^{\mathsf{est}}(\mathcal{E}, \mathbf{X}, \mathbf{Y}) \tag{4.25}$$

where the notation is meant to signify the dependence of $\mathbf{H}_v^{\text{est}}$ on $\mathcal{E}, \mathbf{X}$, and $\mathbf{Y}$. The resulting estimate also has associated with it a reconstruction error given by $\mathbb{E}\left[\|\mathbf{H}_v - \mathbf{H}_v^{\text{est}}\|_F^2\right]$, where $\mathbb{E}$ denotes the expectation with respect to the distribution of $\mathbf{Z}$. The corresponding sensing component at the transmitter involves probing the channel with a training signal that minimizes this reconstruction error[3]

$$\mathbf{x}_{tr}^{\text{opt}}(t) \stackrel{def}{=} \arg\min_{\mathbf{x}_{tr}(t)} \mathbb{E}\left[\Delta\left(\mathbf{H}_v^{\text{est}}\right)\right] \tag{4.26}$$

where we have used the shorthand notation $\Delta(\mathbf{H}) \stackrel{def}{=} \|\mathbf{H}_v - \mathbf{H}\|_F^2$ in the above equation. As a measure of its spectral efficiency, the resulting training signal also has associated with it the concept of *temporal training dimensions*, defined as

$$N_{tr} \stackrel{def}{=} \#\left\{\text{temporal signal space dimensions occupied by } \mathbf{x}_{tr}^{\text{opt}}(t)\right\}.$$

Since each one of the $N_o = TW$ temporal signal space dimensions utilized for training means one less dimension available for communication, the effectiveness of a particular training-based method for a fixed training SNR $\mathcal{E}$ is measured in terms of both the (temporal) training dimensions, N_{tr}, dedicated to $\mathbf{x}_{tr}^{\text{opt}}(t)$ and the ensuing reconstruction error $\mathbb{E}[\Delta(\mathbf{H}_v^{\text{est}}(\mathbf{x}_{tr}^{\text{opt}}))]$.

Traditional training-based methods, such as those in [24–32], have been developed under the implicit assumption that the number of DoF, d, in $\boldsymbol{\mathcal{H}}$ is roughly the same as the maximum *possible* number of its DoF: $d \approx D$. One direct consequence of this assumption has been that linear procedures have become the de-facto standard for reconstruction in much of the existing channel estimation literature. In particular, with a few exceptions such as [13, 95, 102–106], nearly all training-based methods proposed in the past make use of the minimum least squares (LS) error criterion—or its Bayesian counterpart, the minimum mean squared error criterion, for a Bayesian formulation of $\boldsymbol{\mathcal{H}}$—to obtain an estimate of $\mathbf{H}_v$ from $\mathbf{Y}$

$$\mathbf{H}_v^{\text{LS}} = \arg\min_{\mathbf{H}} \left\| \mathbf{Y} - \sqrt{\frac{\mathcal{E}}{N_T}}\,\mathbf{X}\mathbf{H} \right\|_F^2. \tag{4.27}$$

[3]Recall that $\mathbf{H}_v^{\text{est}}$ depends on the training signal $\mathbf{x}_{tr}(t)$ through $\mathbf{X}$.

This is a well-known problem in the statistical estimation literature (see, e.g., [123, Chapter 8]) and its closed-form solution is given by

$$\mathbf{H}_v^{\mathsf{LS}} = \sqrt{\frac{N_T}{\mathcal{E}}}\,\mathbf{X}^{\dagger}\mathbf{Y} \tag{4.28}$$

where $\mathbf{X}^{\dagger}$ is the Moore–Penrose pseudoinverse of $\mathbf{X}$. In order to ensure that (4.27) returns a physically meaningful estimate—in the sense that $\mathbf{H}_v^{\mathsf{LS}}$ returns $\mathbf{H}_v$ in the noiseless case—reconstruction methods based on the LS error criterion further require that the sensing matrix $\mathbf{X}$ has at least as many rows as D/N_R, resulting in the following form for $\mathbf{H}_v^{\mathsf{LS}}$

$$\mathbf{H}_v^{\mathsf{LS}} = \sqrt{\frac{N_T}{\mathcal{E}}}\,(\mathbf{X}^{\mathsf{H}}\mathbf{X})^{-1}\mathbf{X}^{\mathsf{H}}\mathbf{Y} \tag{4.29}$$

where it is assumed that the training signal $\mathbf{x}_{tr}(t)$ is such that $\mathbf{X}$ has full column rank. The reconstruction error of a LS-based channel estimation method in this case can be lower bounded in terms of the following theorem.

Theorem 4.7 Given the general linear form (4.22), and under the assumption that the sensing matrix $\mathbf{X}$ has full column rank, the reconstruction error of $\mathbf{H}_v^{\mathsf{LS}}$ is lower bounded as

$$\mathbb{E}\left[\Delta\left(\mathbf{H}_v^{\mathsf{LS}}\right)\right] \geq \frac{D \cdot N_T}{\mathcal{E}} \tag{4.30}$$

with equality if and only if $\mathbf{X}$ has orthonormal columns.

Proof: To establish this theorem, substitute (4.22) in (4.29) and note that

$$\begin{aligned}
\mathbb{E}\left[\Delta\left(\mathbf{H}_v^{\mathsf{LS}}\right)\right] &= \mathbb{E}\left[\left\|\sqrt{\frac{N_T}{\mathcal{E}}}\,(\mathbf{X}^{\mathsf{H}}\mathbf{X})^{-1}\mathbf{X}^{\mathsf{H}}\mathbf{Z}\right\|_F^2\right] \\
&\overset{(a)}{=} \frac{N_T}{\mathcal{E}} \cdot \mathbb{E}\left[\mathsf{trace}\left((\mathbf{X}^{\mathsf{H}}\mathbf{X})^{-1}\mathbf{X}^{\mathsf{H}}\mathbf{Z}\mathbf{Z}^{\mathsf{H}}\mathbf{X}(\mathbf{X}^{\mathsf{H}}\mathbf{X})^{-1}\right)\right] \\
&\overset{(b)}{=} \frac{N_T N_R}{\mathcal{E}} \cdot \mathsf{trace}\left((\mathbf{X}^{\mathsf{H}}\mathbf{X})^{-1}\right).
\end{aligned} \tag{4.31}$$

Here, (a) is simply the definition of the Frobenius norm and (b) follows from (a) since the inner product of any two rows $\mathbf{z}_i^{\mathsf{H}}$ and $\mathbf{z}_j^{\mathsf{H}}$ of the AWGN matrix $\mathbf{Z}$ is zero in expectation: $\mathbb{E}[\mathbf{z}_i^{\mathsf{H}}\mathbf{z}_j] = N_R\delta_{ij}$.

Next, note that the Gram matrix $\mathbf{X}^{\mathsf{H}}\mathbf{X}$ is positive definite under the assumption of $\mathbf{X}$ being full column rank and let $\{\lambda_i\}$ denote the collection of D/N_R (strictly positive) eigenvalues of $\mathbf{X}^{\mathsf{H}}\mathbf{X}$. Since it is assumed that $\|\mathbf{X}\|_F^2 = D/N_R$, we have from the definition of Frobenius norm that

$$\sum_{i=1}^{D/N_R} \lambda_i = \frac{D}{N_R}. \tag{4.32}$$

Then, from elementary linear algebra, we have

$$\begin{aligned} \mathsf{trace}\left((\mathbf{X}^{\mathsf{H}}\mathbf{X})^{-1}\right) = \sum_{i=1}^{D/N_R} \frac{1}{\lambda_i} &= \frac{D}{N_R}\left(\frac{\sum_{i=1}^{D/N_R}\frac{1}{\lambda_i}}{\frac{D}{N_R}}\right) \\ &\overset{(c)}{\geq} \frac{D}{N_R}\left(\frac{\frac{D}{N_R}}{\sum_{i=1}^{D/N_R}\lambda_i}\right) \overset{(d)}{=} \frac{D}{N_R} \end{aligned} \tag{4.33}$$

where (c) follows from the arithmetic-harmonic means inequality [124], while (d) is a consequence of (4.32). Substituting (4.33) in (4.31) yields the desired inequality in (4.30). Finally, from the arithmetic-harmonic means inequality, the equality in (c) in (4.33) holds if and only if $\lambda_1 = \lambda_2 = \cdots = \lambda_{D/N_R}$ [124], resulting in the condition that $\mathbf{X}$ must have orthonormal columns for the equality to hold in (4.30). This completes the proof of the theorem. ∎

Note that an immediate consequence of Theorem 4.7 is that an optimal training signal $\mathbf{x}_{tr}^{\mathsf{opt}}(t)$ for LS-based estimation methods is the one that leads to $\mathbf{X}^{\mathsf{H}}\mathbf{X} = \mathbf{I}_{N_T L(2M+1)}$. As such, much of the emphasis in the previously proposed channel estimation methods has been on designing training signals that are spectrally efficient and that result in sensing matrices having either orthonormal or nearly-orthonormal columns [24–32].

4.4 Compressed Channel Sensing: Main Results

The preceding discussion brings forth several salient characteristics of traditional training-based methods. Below, we recount three specific ones that we perceive to be the most relevant to our discussion in the sequel:

[1] Traditional training-based methods more or less rely on LS-based linear reconstruction strategies, such as the one in (4.29), at the receiver to obtain an estimate of $\mathbf{H}_v$.

[2] Because of their reliance on linear reconstruction procedures, the training signals used in traditional methods must be such that the resulting sensing matrix $\mathbf{X}$ has at least D/N_R rows. As noted in Table 4.2, depending upon the type of signaling waveforms used for training and the channel class to which $\boldsymbol{\mathcal{H}}$ belongs, this requirement often translates into the condition that the number of temporal training dimensions dedicated to $\mathbf{x}_{tr}(t)$ must be at least as large as the maximum number of DoF per receive antenna in $\boldsymbol{\mathcal{H}}$: $N_{tr} = \Omega\left(\frac{D}{N_R}\right)$; see Section 4.5 and Section 4.6 for further details on this condition.

[3] Finally, regardless of the eventual choice of training signals, the reconstruction error in traditional methods is given by $\mathbb{E}[\Delta(\mathbf{H}_v^{\mathsf{LS}})] = \Omega\left(\frac{D \cdot N_T}{\mathcal{E}}\right)$.

In the light of the above observations, a natural question to ask is: *how good is the performance of traditional training-based methods?* In fact, if one assumes that $\boldsymbol{\mathcal{H}}$ is not sparse (in other words, $d = D$) then it is easy to argue the optimality of these methods:

[1] In the case of nonsparse channels, the LS estimate $\mathbf{H}_v^{\mathsf{LS}}$ is also the maximum-likelihood estimate of $\mathbf{H}_v$ (see, e.g., [123, Chapter 8]).

[2] The reconstruction error lower bound (4.30) in this case is also the Cramer–Rao lower bound [123, Chapter 3], which—as noted in Theorem 4.7—can be achieved in most instances through an appropriate choice of the training signal.

However, it is arguable whether LS-based channel estimation methods are also optimal for the case when $\boldsymbol{\mathcal{H}}$ is d-sparse. In particular, note that sparse channels are completely characterized by $2d$ parameters—corresponding to the locations and values of nonzero virtual channel coefficients. Our estimation theory intuition therefore suggests that perhaps $\mathbb{E}[\Delta(\mathbf{H}_v^{\mathsf{est}})] = \Omega\left(\frac{d \cdot N_T}{\mathcal{E}}\right)$ and, for signaling and channel configurations that require $N_{tr} = \Omega\left(\frac{D}{N_R}\right)$ in the case of LS-based estimation methods, $N_{tr} = \Omega\left(\frac{d}{N_R}\right)$ are the actual fundamental limits in sparse-channel estimation.

Table 4.2 Summary and comparison of CCS results for the signaling and channel configurations studied in Chapter 4

Channel Classification	Signaling	Traditional Methods		Compressed Channel Sensing	
		Recon. Error	Condition	Recon. Error	Condition [a]
Frequency-Selective Single-Antenna ($D = L$)	Single-Carrier	$\succeq \frac{D}{\mathcal{E}}$	—	$\preceq \frac{d}{\mathcal{E}} \cdot \log D$	$N_o \succeq d^2 \cdot \log D$
	Multi-Carrier	$\succeq \frac{D}{\mathcal{E}}$	$N_{tr} \succeq D$	$\preceq \frac{d}{\mathcal{E}} \cdot \log D$	$N_{tr} \succeq d \cdot \log^5 N_o$
Doubly-Selective Single-Antenna ($D = L(2M+1)$)	Single-Carrier	$\succeq \frac{D}{\mathcal{E}}$	—	$\preceq \frac{d}{\mathcal{E}} \cdot \log D$	$N_o \succeq d^2 \cdot \log D$
	Multi-Carrier	$\succeq \frac{D}{\mathcal{E}}$	$N_{tr} \succeq D$	$\preceq \frac{d}{\mathcal{E}} \cdot \log D$	$N_{tr} \succeq d \cdot \log^5 N_o$
Nonselective Multiple-Antenna ($D = N_R N_T$)	—	$\succeq \frac{D \cdot N_T}{\mathcal{E}}$	$N_{tr} \succeq \frac{D}{N_R}$	$\preceq \frac{d \cdot N_T}{\mathcal{E}} \cdot \log D$	$N_{tr} \succeq \frac{d}{N_R} \cdot \log N_T$
Frequency-Selective Multiple-Antenna ($D = N_R N_T L$)	Multi-Carrier	$\succeq \frac{D \cdot N_T}{\mathcal{E}}$	$N_{tr} \succeq \frac{D}{N_R}$	$\preceq \frac{d \cdot N_T}{\mathcal{E}} \cdot \log D$	$N_{tr} \succeq \frac{d}{N_R} \cdot \log^6 N_o$
Doubly-Selective Multiple-Antenna ($D = N_R N_T L(2M+1)$)	Multi-Carrier	$\succeq \frac{D \cdot N_T}{\mathcal{E}}$	$N_{tr} \succeq \frac{D}{N_R}$	$\preceq \frac{d \cdot N_T}{\mathcal{E}} \cdot \log D$	$N_{tr} \succeq \frac{d}{N_R} \cdot \log^6 N_o$

[a] The last three conditions are for the case when the conditional sparsity of each AoA equals the average AoA sparsity.

In Section 4.5 and Section 4.6, we present new training-based estimation methods for seven particular signaling and channel configurations (see Table 4.2) and show that our intuition is indeed correct (modulo polylogarithmic factors). In particular, a key feature of our approach to estimating sparse multipath channels—originally proposed in [14] for frequency-selective single-antenna channels and later generalized in [15, 107] to some of the other channel classes—is the use of a sparsity-inducing mixed-norm optimization criterion for reconstruction at the receiver that is based on recent advances in the theory of compressed sensing (cf. Chapter 2). This makes our approach—termed as compressed channel sensing (CCS)—fundamentally different from the traditional training-based methods: the former relies on a nonlinear reconstruction procedure while the latter utilizes linear reconstruction techniques. Note that a number of researchers in the recent past have also proposed various training-based methods for sparse multipath channels that are based on nonlinear reconstruction techniques [13, 95, 102–106]. The thing that distinguishes CCS from the prior work is that the CCS framework is highly amenable to analysis.

Specifically, in order to give a summary of the results to come, define the *conditional* sparsity pattern associated with the i-th resolvable AoA to be $\mathcal{S}_d(i) \stackrel{def}{=} \{(i,k,\ell,m) : (i,k,\ell,m) \in \mathcal{S}_d\}$. Then it is shown in the sequel that:

[R1] The performance of CCS in terms of the reconstruction error is provably better than the traditional methods. The training signals and reconstruction procedures specified by CCS for the signaling and channel configurations studied in this chapter ensure that the CCS reconstruction error $\Delta(\mathbf{H}_v^{\mathsf{ccs}}) = O\left(\frac{d \cdot N_T}{\mathcal{E}} \cdot \log D\right)$ with high probability.

[R2] CCS is provably more spectrally efficient than the traditional methods. Assume that the conditional sparsity of each resolvable AoA is equal to the average AoA sparsity; in other words, $|\mathcal{S}_d(i)| = \frac{d}{N_R}, i = 1, \ldots, N_R$. Then while traditional methods require that the number of training dimensions $N_{tr} = \Omega\left(\frac{D}{N_R}\right)$ for certain signaling and channel configurations, CCS only requires that $N_{tr} = \Omega\left(\frac{d}{N_R} \times \text{polylog factor}\right)$ for the same configurations.

Conversely, **[R1]** and **[R2]** together imply that CCS achieves a target reconstruction error using far less energy and, in many instances, latency and bandwidth than that dictated by the traditional training-based methods.

Table 4.2 provides a compact summary of the CCS results as they pertain to the seven signaling and channel configurations studied in this chapter and compares them to the corresponding results for traditional training-based methods. One thing to point out in this table is the CCS condition $N_o = \Omega(d^2 \cdot \log D)$ when using single-carrier signaling waveforms for estimating single-antenna channels. This condition *seems* to be nonexistent for traditional methods. Note, however, that in order to make the columns of $\mathbf{X}$ as close to orthonormal as possible—a necessary condition for the LS-based reconstruction to achieve the lower bound of (4.30)—traditional methods implicitly require that the temporal signal space dimensions be as large as possible: $N_o \nearrow \infty$. As such, the CCS condition is in fact a relaxation of this implicit requirement for traditional methods.

As is evident from the preceding discussion and analysis, the performance of CCS is a significant improvement over that of traditional training-based methods when it comes to sparse-channel estimation. And while we have purposely avoided providing concrete details of the CCS framework up to this point so as not to clutter the presentation, the rest of the chapter is primarily devoted to discussing the exact form of training signals and reconstruction procedures used by CCS for the configurations listed in Table 4.2. However, since CCS builds on top of the theoretical framework provided by compressed sensing, the reader may want to revisit at this point Chapter 2 in general and Chapter 3 in particular before proceeding further.

4.5 Compressed Channel Sensing: Single-Antenna Channels

4.5.1 Estimating Sparse Frequency-Selective Channels

For a single-antenna channel that is frequency-selective (also, see "*Special Case: Frequency-Selective Single-Antenna Channels*"), the virtual representation (4.9) of the channel reduces to

$$\widetilde{H}(f) = \sum_{\ell=0}^{L-1} H_v(\ell) e^{-j2\pi \frac{\ell}{W} f} \tag{4.34}$$

and the corresponding received training signal is given by [cf. (4.21)]

$$y_{tr}(t) \approx \sum_{\ell=0}^{L-1} H_v(\ell) x_{tr}(t - \ell/W) + z_{tr}(t)\,,\ 0 \le t \le T + \tau_{max}. \tag{4.35}$$

In general, two types of signaling waveforms are commonly employed to communicate over a frequency-selective channel, namely, (single-carrier) *spread spectrum* (SS) waveforms and (multi-carrier) *orthogonal frequency division multiplexing* (OFDM) waveforms. We begin our discussion of the CCS framework for sparse frequency-selective channels by focusing first on SS signaling and then on OFDM signaling.

Spread Spectrum Signaling

In the case of SS signaling, the training signal $x_{tr}(t)$ can be represented as

$$x_{tr}(t) = \sqrt{\mathcal{E}} \sum_{n=0}^{N_o-1} x_n g(t - nT_c)\,,\ 0 \le t \le T \tag{4.36}$$

where $g(t)$ is the *chip waveform* having unit energy ($\int |g(t)|^2 dt = 1$), $T_c \approx 1/W$ is the chip duration, and $\{x_n \in \mathbb{C}\}$ is the N_o-dimensional spreading code associated with the training signal, also having unit energy ($\sum_n |x_n|^2 = 1$). In this case, chip-matched filtering the received training signal (4.35) yields an equivalent discrete-time representation [22]

$$y_n = \sqrt{\mathcal{E}} \sum_{\ell=0}^{L-1} H_v(\ell) x_{n-\ell} + z_n,\ n = 0, 1, \ldots, \tilde{N}_o - 1 \tag{4.37}$$

where $\{z_n\}$ is a zero-mean, unit-variance, complex AWGN sequence and $\tilde{N}_o \stackrel{def}{=} (N_o + L - 1)$. This input-output relation can also be expressed as a matrix-vector product

$$\underbrace{\begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ \vdots \\ y_{\tilde{N}_o-2} \\ y_{\tilde{N}_o-1} \end{bmatrix} = \sqrt{\mathcal{E}} \begin{bmatrix} x_0 & & 0 \\ x_1 & \ddots & \\ \vdots & \ddots & x_0 \\ x_{N_o-1} & & x_1 \\ & \ddots & \vdots \\ 0 & & x_{N_o-1} \end{bmatrix} \begin{bmatrix} H_v(0) \\ H_v(1) \\ \vdots \\ H_v(L-1) \end{bmatrix} + \begin{bmatrix} z_0 \\ z_1 \\ \vdots \\ \vdots \\ z_{\tilde{N}_o-2} \\ z_{\tilde{N}_o-1} \end{bmatrix}}_{\mathbf{y} = \sqrt{\mathcal{E}}\,\mathbf{X}\mathbf{h}_v + \mathbf{z}} \tag{4.38}$$

where $\mathbf{h}_v \in \mathbb{C}^L$ represents the vector of channel coefficients $\{H_v(\ell)\}$, $\mathbf{X}$ denotes an $\tilde{N}_o \times L$ full Toeplitz (convolutional) matrix generated by the spreading code $\{x_n\}$, and $\mathbf{z}$ is an AWGN vector that is distributed as $\mathcal{CN}(\mathbf{0}_{\tilde{N}_o}, \mathbf{I}_{\tilde{N}_o})$.

Note that (4.38) is the single-antenna version of the standard form (4.22). Therefore, from (4.30), the reconstruction error of LS-based training methods in this case is simply given by $\mathbb{E}[\Delta(\mathbf{h}_v^{\mathsf{LS}})] = \Omega(L/\mathcal{E})$. We now describe the CCS approach to estimating frequency-selective channels using SS signaling, and show that for d-sparse channels it leads to an improvement of a factor of about L/d in the reconstruction error (modulo a logarithmic factor).

$\big[$**CCS-1**$\big]$ – SS Training and Reconstruction

Training: Pick the spreading code $\{x_n\}_{n=0}^{N_o-1}$ associated with $x_{tr}(t)$ to be a sequence of i.i.d. Rademacher random variables taking values $+1/\sqrt{N_o}$ or $-1/\sqrt{N_o}$ with probability $1/2$ each.

Reconstruction: Fix any $a \geq 0$ and pick $\lambda = \sqrt{2\mathcal{E}(1+a)\log L}$. The CCS estimate of $\mathbf{h}_v$ is then given as a solution to the Dantzig selector

$$\mathbf{h}_v^{\mathsf{ccs}} = \underset{\mathbf{h}\in\mathbb{C}^L}{\arg\min} \left\|\mathbf{h}\right\|_1 \quad \text{subject to} \quad \left\|\sqrt{\mathcal{E}}\,\mathbf{X}^{\mathsf{H}}\big(\mathbf{y} - \sqrt{\mathcal{E}}\,\mathbf{X}\mathbf{h}\big)\right\|_\infty \leq \lambda\,. \tag{4.39}$$

We now summarize the performance of $\big[$**CCS-1**$\big]$ in terms of the reconstruction error.

Theorem 4.8 Suppose that the number of temporal signal space dimensions $N_o \geq 4c_2\, d^2 \log L$. Then for any $\delta_{2d} \in (0, 0.3]$, the CCS estimate of $\mathbf{h}_v$ satisfies

$$\Delta\left(\mathbf{h}_v^{\mathsf{ccs}}\right) \leq c_0^2 \cdot \frac{d}{\mathcal{E}} \cdot \log L \tag{4.40}$$

with probability exceeding $1 - 2\max\left\{2\big(\pi(1+a)\log L \cdot L^{2a}\big)^{-1/2}, \exp(-\frac{c_1 N_o}{4d^2})\right\}$. Here, the constants $c_1 > 0$ and $c_2 > 0$ are the same as in Theorem 3.5 (with δ_{2d} in place of δ_S), while the constant $c_0 = 4\sqrt{2(1+a)}/(1-3\delta_{2d})$.

Proof: The proof of this theorem is a trivial consequence of Theorem 2.13 and Theorem 3.5. Specifically, fix some $\delta_{2d} \in (0, 0.3]$ and note that the frequency-selective channel being d-sparse

means that $\|\mathbf{h}_v\|_0 \leq d \ll L$. Therefore, we have from Theorem 2.13 that the reconstruction error satisfies (4.40) with probability exceeding $1 - 2\left(\sqrt{\pi(1+a)\log L}\cdot L^a\right)^{-1}$, provided the full Toeplitz matrix $\mathbf{X} \in RIP(2d, \delta_{2d})$. On the other hand, Theorem 3.5 guarantees that the matrix $\mathbf{X} \in RIP(2d, \delta_{2d})$ with probability exceeding $1 - \exp(-\frac{c_1 N_o}{4d^2})$, provided $N_o \geq 4c_2\, d^2 \log L$. The claim of the theorem therefore follows by taking a union bound over the two *failure* events. ∎

OFDM Signaling

If OFDM signaling is used instead of SS signaling for training and communication purposes then the training signal takes the form

$$x_{tr}(t) = \sqrt{\frac{\mathcal{E}}{N_{tr}}} \sum_{n \in \mathcal{S}_{tr}} g(t) e^{j2\pi \frac{n}{T} t},\ 0 \leq t \leq T \tag{4.41}$$

where $g(t)$ is a prototype pulse having unit energy, $\mathcal{S}_{tr} \subset \mathcal{S} \stackrel{def}{=} [0 \ldots N_o - 1]$ is the set of indices of *pilot tones* used for training, and N_{tr}—the number of temporal training dimensions—denotes the total number of pilot tones, $N_{tr} \stackrel{def}{=} |\mathcal{S}_{tr}|$, and is a measure of the spectral efficiency of $x_{tr}(t)$. In this case, matched filtering the received training signal (4.35) with the OFDM basis waveforms $\left\{g(t)e^{j2\pi \frac{n}{T} t}\right\}_{\mathcal{S}_{tr}}$ yields [22]

$$y_n = \sqrt{\mathcal{E}} H_n + z_n,\ n \in \mathcal{S}_{tr} \tag{4.42}$$

where $\left\{z_n\right\}$ again is a zero-mean, unit-variance, complex AWGN sequence, while the *OFDM channel coefficients* $\left\{H_n\right\}$ are given by [cf. (4.34)]

$$H_n \approx \frac{1}{\sqrt{N_{tr}}} \widetilde{H}(f)\big|_{f=\frac{n}{T}} = \frac{1}{\sqrt{N_{tr}}}\, \mathbf{u}_n^{\mathsf{T}} \mathbf{h}_v,\ n \in \mathcal{S}_{tr}. \tag{4.43}$$

Here, similar to the case of SS signaling, $\mathbf{h}_v \in \mathbb{C}^L$ is the vector of channel coefficients $\left\{H_v(\ell)\right\}$, whereas $\mathbf{u}_n^{\mathsf{T}} \stackrel{def}{=} \begin{bmatrix} e^{-j0\omega_{n,N_o}} & \ldots & e^{-j(L-1)\omega_{n,N_o}} \end{bmatrix}$ denotes the collection of L samples of a discrete sinusoid with frequency $\omega_{n,N_o} \stackrel{def}{=} 2\pi \frac{n}{N_o}, n \in \mathcal{S}_{tr}$. It is then easy to see from (4.42) and (4.43) that stacking the received training data $\left\{y_n\right\}_{S_{tr}}$ into an N_{tr}-dimensional vector $\mathbf{y}$ again yields the standard linear form

$$\mathbf{y} = \sqrt{\mathcal{E}}\, \mathbf{X}\mathbf{h}_v + \mathbf{z}. \tag{4.44}$$

The main difference here being that the sensing matrix $\mathbf{X}$ is now an $N_{tr} \times L$ matrix that is comprised of $\left\{ \frac{1}{\sqrt{N_{tr}}} \mathbf{u}_n^{\mathsf{T}} : n \in \mathcal{S}_{tr} \right\}$ as its rows, while the AWGN vector $\mathbf{z}$ is distributed as $\mathcal{CN}(\mathbf{0}_{N_{tr}}, \mathbf{I}_{N_{tr}})$.

The important thing to note from (4.44) is that the form of $\mathbf{X}$ in the case of OFDM signaling imposes the condition that $N_{tr} \geq L$ for $\mathbf{X}$ to have full column rank. In order to estimate a frequency-selective channel using OFDM signaling, traditional methods—such as [25]—therefore require that $N_{tr} = \Omega(L)$ and, from (4.30), at best yield $\mathbb{E}[\Delta(\mathbf{h}_v^{\mathsf{LS}})] = \Omega(L/\mathcal{E})$. In contrast, we now outline the CCS approach to this problem and quantify its advantage over traditional methods for sparse channels. Note that an implicit assumption we will be making in the sequel is that $N_o \gg L$, which follows from the fact that a basic premise in OFDM systems is that $T \gg \tau_{max}$ [21].

$\big[$**CCS -2**$\big]$ – OFDM Training and Reconstruction

Training: Pick $\mathcal{S}_{tr}$—the set of indices of pilot tones—to be a set of N_{tr} indices sampled uniformly at random (without replacement) from the set $\mathcal{S} = [0 \ldots N_o - 1]$.

Reconstruction: Fix any $a \geq 0$ and pick $\lambda = \sqrt{2\mathcal{E}(1+a)\log L}$. The CCS estimate of $\mathbf{h}_v$ is then given as a solution to the Dantzig selector

$$\mathbf{h}_v^{\mathsf{CCS}} = \underset{\mathbf{h} \in \mathbb{C}^L}{\arg\min} \left\| \mathbf{h} \right\|_1 \quad \text{subject to} \quad \left\| \sqrt{\mathcal{E}}\, \mathbf{X}^{\mathsf{H}} \big(\mathbf{y} - \sqrt{\mathcal{E}}\, \mathbf{X}\mathbf{h} \big) \right\|_\infty \leq \lambda \,. \tag{4.45}$$

Below, we summarize the performance of $\big[$**CCS -2**$\big]$ in terms of both the spectral efficiency and the reconstruction error.

Theorem 4.9 Suppose that $N_o, d > 2$, and let the number of pilot tones $N_{tr} \geq 2c_3 d \log^5 N_o$. Then for any $\delta_{2d} \in (0, 0.3]$, the CCS estimate of $\mathbf{h}_v$ satisfies

$$\Delta\left(\mathbf{h}_v^{\mathsf{CCS}}\right) \leq c_0^2 \cdot \frac{d}{\mathcal{E}} \cdot \log L \tag{4.46}$$

with probability exceeding $1 - 2\max\left\{ 2\big(\pi(1+a)\log L \cdot L^{2a}\big)^{-1/2}, 10 N_o^{-c_4 \delta_{2d}^2} \right\}$. Here, the constants $c_3 > 0$ and $c_4 > 0$ are the same as in Theorem 2.9, while the constant $c_0 = 4\sqrt{2(1+a)}/(1 - 3\delta_{2d})$.

Proof: The proof of this theorem is a consequence of Theorem 2.9 and Theorem 2.13. Specifically, fix some $\delta_{2d} \in (0, 0.3]$ and note that since we are assuming $\|\mathbf{h}_v\|_0 \leq d \ll L$, Theorem 2.13 asserts that the reconstruction error trivially satisfies (4.46) with probability exceeding $1 - 2\left(\sqrt{\pi(1+a)\log L}\cdot L^a\right)^{-1}$, provided the sensing matrix $\mathbf{X} \in RIP(2d, \delta_{2d})$. On the other hand, observe that $\mathbf{X}$ in this case simply corresponds to an $N_{tr} \times L$ *column submatrix* of an $N_{tr} \times N_o$ (randomly) subsampled unitary matrix $\mathbf{A}$ generated from an $N_o \times N_o$ unitary (Fourier) matrix $\mathbf{U}$ that consists of

$$\left\{ \frac{1}{\sqrt{N_o}} \left[e^{-j0\omega_{n,N_o}} \quad \ldots \quad e^{-j(N_o-1)\omega_{n,N_o}} \right] : n = 0, 1, \ldots, N_o - 1 \right\}$$

as its rows. Therefore, by taking $t \stackrel{def}{=} \log N_o$ (and since $\mu_{\mathbf{U}} = 1$ in this case), Theorem 2.9 guarantees that $\mathbf{A} \in RIP(2d, \delta_{2d})$ with probability exceeding $1 - 10N_o^{-c_4\delta_{2d}^2}$, provided we have $N_{tr} \geq 2c_3 d \log^5 N_o$. Finally, note from the definition of RIP that since $\mathbf{X}$ is a column submatrix of $\mathbf{A}$, $\mathbf{A} \in RIP(2d, \delta_{2d})$ implies that $\mathbf{X} \in RIP(2d, \delta_{2d})$. The claimed probability bound therefore trivially follows by taking a union bound over the two failure events. ∎

4.5.2 Estimating Sparse Doubly-Selective Channels

In the case of a single-antenna channel that is doubly selective, the virtual representation of the channel reduces to [cf. (4.9)]

$$\widetilde{H}(t, f) = \sum_{\ell=0}^{L-1} \sum_{m=-M}^{M} H_v(\ell, m) e^{-j2\pi \frac{\ell}{W} f} e^{j2\pi \frac{m}{T} t} \tag{4.47}$$

and the corresponding received training signal can be written as

$$y_{tr}(t) \approx \sum_{\ell=0}^{L-1} \sum_{m=-M}^{M} H_v(\ell, m) e^{j2\pi \frac{m}{T} t} x_{tr}(t - \ell/W) + z_{tr}(t)\,,\ 0 \leq t \leq T + \tau_{max}\,. \tag{4.48}$$

Signaling waveforms that are often used to communicate over a doubly-selective channel can be broadly categorized as (single-carrier) SS waveforms and (multi-carrier) *short-time Fourier* (STF) waveforms—a generalization of OFDM waveforms for doubly-selective channels [125, 126]. Below, we discuss the specifics of the CCS framework for sparse doubly-selective channels as it pertains to both SS and STF signaling waveforms.

Spread Spectrum Signaling

The SS training signal $x_{tr}(t)$ in the case of a doubly-selective channel has exactly the same form as given in (4.36). The difference here is that the chip-matched-filtered output in this case looks different from the one in (4.37). Specifically, define again $\tilde{N}_o \stackrel{def}{=} N_o + L - 1$. Then chip-matched filtering the received training signal (4.48) yields [13]

$$y_n = \sqrt{\mathcal{E}} \sum_{\ell=0}^{L-1} \sum_{m=-M}^{M} H_v(\ell, m) e^{j2\pi \frac{m}{N_o} n} x_{n-\ell} + z_n \,,\ n = 0, 1, \ldots, \tilde{N}_o - 1. \tag{4.49}$$

Now, in order to represent this received training data into the standard form (4.22), first define an $\tilde{N}_o$-length sequence of vectors $\left\{\mathbf{x}_n \in \mathbb{C}^L\right\}$ comprising of the spreading code $\left\{x_n\right\}$ as follows

$$\mathbf{x}_n^{\mathsf{T}} \stackrel{def}{=} \begin{bmatrix} x_n & x_{n-1} & \ldots & x_{n-(L-1)} \end{bmatrix},\ n = 0, 1, \ldots, \tilde{N}_o - 1 \tag{4.50}$$

where the notational understanding is that $x_i \stackrel{def}{=} 0$ for any $i \neq [0 \ldots N_o - 1]$. Next, define

$$\widetilde{\mathbf{H}} \stackrel{def}{=} \begin{bmatrix} H_v(0,-M) & H_v(0,-M+1) & \ldots & H_v(0,M) \\ H_v(1,-M) & H_v(1,-M+1) & \ldots & H_v(1,M) \\ \vdots & \vdots & & \vdots \\ H_v(L-1,-M) & H_v(L-1,-M+1) & \ldots & H_v(L-1,M) \end{bmatrix} \tag{4.51}$$

to be the $L \times (2M+1)$ matrix of channel coefficients. Note that each column of this channel matrix represents the impulse response of the doubly-selective channel corresponding to a *fixed* Doppler shift. Finally, let $\left\{\mathbf{u}_n \in \mathbb{C}^{2M+1}\right\}$ denote an $\tilde{N}_o$-length sequence of *modulation vectors* given by

$$\mathbf{u}_n^{\mathsf{T}} = \begin{bmatrix} e^{-jM\omega_{n,N_o}} & e^{-j(M-1)\omega_{n,N_o}} & \ldots & e^{jM\omega_{n,N_o}} \end{bmatrix},\ n = 0, 1, \ldots, \tilde{N}_o - 1 \tag{4.52}$$

where we again have $\omega_{n,N_o} \stackrel{def}{=} 2\pi \frac{n}{N_o}$. Then it is easy to see that the sequence $\left\{y_n\right\}$ in (4.49) can also be expressed as

$$\begin{aligned} y_n &= \sqrt{\mathcal{E}}\, \mathbf{x}_n^{\mathsf{T}} \widetilde{\mathbf{H}} \mathbf{u}_n + z_n \stackrel{(a)}{=} \sqrt{\mathcal{E}} \left(\mathbf{u}_n^{\mathsf{T}} \otimes \mathbf{x}_n^{\mathsf{T}}\right) \mathsf{vec}(\widetilde{\mathbf{H}}) + z_n \\ &= \sqrt{\mathcal{E}} \left(\mathbf{u}_n^{\mathsf{T}} \otimes \mathbf{x}_n^{\mathsf{T}}\right) \mathbf{h}_v + z_n,\ n = 0, 1, \ldots, \tilde{N}_o - 1 \end{aligned} \tag{4.53}$$

where $\mathbf{h}_v \stackrel{def}{=} \mathsf{vec}(\widetilde{\mathbf{H}}) \in \mathbb{C}^{L(2M+1)}$ is the vector of virtual channel coefficients, and (a) simply follows from the basic identity that relates the vec-function with the Kronecker product [44]: $\mathsf{vec}(\mathbf{ABC}) = (\mathbf{C}^{\mathsf{T}} \otimes \mathbf{A})\mathsf{vec}(\mathbf{B})$ for arbitrary (complex- or real-valued) matrices $\mathbf{A}, \mathbf{B}$, and $\mathbf{C}$.

The expression (4.53) suggests that the input-output relation (4.49) can also be expressed as a matrix-vector product by stacking the received training data $\{y_n\}$ as follows

$$\underbrace{\begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_{\tilde{N}_o-2} \\ y_{\tilde{N}_o-1} \end{bmatrix} = \sqrt{\mathcal{E}} \begin{bmatrix} \mathbf{u}_0^{\mathsf{T}} \otimes \mathbf{x}_0^{\mathsf{T}} \\ \mathbf{u}_1^{\mathsf{T}} \otimes \mathbf{x}_1^{\mathsf{T}} \\ \vdots \\ \mathbf{u}_{\tilde{N}_o-2}^{\mathsf{T}} \otimes \mathbf{x}_{\tilde{N}_o-2}^{\mathsf{T}} \\ \mathbf{u}_{\tilde{N}_o-1}^{\mathsf{T}} \otimes \mathbf{x}_{\tilde{N}_o-1}^{\mathsf{T}} \end{bmatrix} \begin{bmatrix} H_v(0,-M) \\ H_v(1,-M) \\ \vdots \\ H_v(L-1,M) \end{bmatrix} + \begin{bmatrix} z_0 \\ z_1 \\ \vdots \\ z_{\tilde{N}_o-2} \\ z_{\tilde{N}_o-1} \end{bmatrix}}_{\mathbf{y}=\sqrt{\mathcal{E}}\,\mathbf{X}\mathbf{h}_v+\mathbf{z}} \tag{4.54}$$

where the sensing matrix $\mathbf{X}$ has dimensions $\tilde{N}_o \times L(2M+1)$, while the AWGN vector $\mathbf{z}$ is distributed as $\mathcal{CN}(\mathbf{0}_{\tilde{N}_o}, \mathbf{I}_{\tilde{N}_o})$. Note that under the assumption that the doubly-selective channel is underspread ($\tau_{max}\nu_{max} \ll 1$), we have that $TW \gg \tau_{max}\nu_{max}TW \;\Rightarrow\; \tilde{N}_o > L(2M+1)$. This fact—combined with the form of $\mathbf{X}$—ensures that the sensing matrix in this case has full column rank and training-based methods can use the LS criterion (4.27) without further conditions, resulting in $\mathbb{E}[\Delta(\mathbf{h}_v^{\mathsf{LS}})] = \Omega(L(2M+1)/\mathcal{E})$. Below, we describe the CCS approach to estimating doubly-selective channels using SS signaling, which is markedly similar to $\big[$**CCS-1**$\big]$, and provide an upper bound on the corresponding reconstruction error for d-sparse channels that is significantly better than $\Omega(L(2M+1)/\mathcal{E})$.

$\big[$**CCS-3**$\big]$ – SS Training and Reconstruction

Training: Same as in the case of $\big[$**CCS-1**$\big]$.

Reconstruction: Fix any $a \geq 0$ and pick $\lambda = \sqrt{2\mathcal{E}(1+a)\log L(2M+1)}$. The CCS estimate of $\mathbf{h}_v$ is then given as a solution to the Dantzig selector

$$\mathbf{h}_v^{\mathsf{CCS}} = \underset{\mathbf{h}\in\mathbb{C}^{L(2M+1)}}{\arg\min} \big\|\mathbf{h}\big\|_1 \quad \text{subject to} \quad \Big\|\sqrt{\mathcal{E}}\,\mathbf{X}^{\mathsf{H}}\big(\mathbf{y} - \sqrt{\mathcal{E}}\,\mathbf{X}\mathbf{h}\big)\Big\|_\infty \leq \lambda\,. \tag{4.55}$$

Theorem 4.10 Suppose that the number of temporal signal space dimensions satisfies the constraint $N_o \geq 4c_4\, d^2 \log L(2M+1)$. Then for any $\delta_{2d} \in (0, 0.3]$, the CCS estimate of $\mathbf{h}_v$ satisfies

$$\Delta\left(\mathbf{h}_v^{\mathsf{ccs}}\right) \leq c_0^2 \cdot \frac{d}{\mathcal{E}} \cdot \log L(2M+1) \tag{4.56}$$

with probability $\geq 1 - 2\max\left\{2\Big(\pi(1+a)\log L(2M+1)\cdot\big(L(2M+1)\big)^{2a}\Big)^{-1/2}, \exp(-\frac{c_3 N_o}{4d^2})\right\}$. Here, the constants $c_3 > 0$ and $c_4 > 0$ are the same as in Theorem 3.9 (with δ_{2d} in place of δ_S), while the constant $c_0 = 4\sqrt{2(1+a)}/(1-3\delta_{2d})$.

Proof: The key ingredient in the proof of this theorem is the observation that the sensing matrix $\mathbf{X}$ in (4.54) is in fact an overdetermined Gabor matrix (with parameters L and M) that is generated from the spreading code $\{x_n\}$ (cf. Section 3.4). The statement of the theorem then follows trivially from Theorem 2.13 and Theorem 3.9 using arguments similar to those made in the proof of Theorem 4.8. ∎

STF Signaling

In the case of STF signaling, which is a generalization of OFDM signaling to counteract the time selectivity of doubly-selective channels [125, 126], the training signal $x_{tr}(t)$ is of the form

$$x_{tr}(t) = \sqrt{\frac{\mathcal{E}}{N_{tr}}} \sum_{(r,s)\in\mathcal{S}_{tr}} g(t - rT_o)e^{j2\pi s W_o t},\ 0 \leq t \leq T \tag{4.57}$$

where $g(t)$ is again a prototype pulse having unit energy, $\mathcal{S}_{tr} \subset \mathcal{S} \stackrel{def}{=} [0 \ldots N_t - 1] \times [0 \ldots N_f - 1]$ is the set of indices of STF pilot tones used for training, and N_{tr}—a measure of the spectral efficiency of $x_{tr}(t)$—denotes the total number of pilot tones: $N_{tr} \stackrel{def}{=} |\mathcal{S}_{tr}|$. Here, the parameters $T_o \in [\tau_{max}, 1/\nu_{max}]$ and $W_o \in [\nu_{max}, 1/\tau_{max}]$ correspond to the time and frequency separation of the STF basis waveforms $\left\{g(t - rT_o)e^{j2\pi s W_o t}\right\}_{\mathcal{S}}$ in the N_o-dimensional time-frequency plane, respectively, and are chosen so that $T_o W_o = 1$ (which gives rise to an orthogonal STF basis [126]). Note that the total number of STF basis waveforms available for communication and training purposes in this case is $N_t N_f = N_o$, where $N_t \stackrel{def}{=} T/T_o$ and $N_f \stackrel{def}{=} W/W_o$.[4]

[4]Note that signaling over an orthogonal STF basis can be thought of as block OFDM signaling with OFDM symbol duration T_o and block length $N_t = T/T_o$.

One of the main advantages of signaling over doubly-selective channels using STF waveforms is that the basis functions $\left\{g(t-rT_o)e^{j2\pi sW_o t}\right\}$ serve as approximate eigenfunctions for sufficiently underspread channels [126]. Specifically, matched filtering the received training signal (4.48) in this case with the STF basis waveforms $\left\{g(t-rT_o)e^{j2\pi sW_o t}\right\}_{\mathcal{S}_{tr}}$ yields [126]

$$y_{r,s} = \sqrt{\mathcal{E}}\, H_{r,s} + z_{r,s}\,,\ (r,s) \in \mathcal{S}_{tr} \tag{4.58}$$

where $\left\{z_{r,s}\right\}$ is a zero-mean, unit-variance, complex AWGN sequence, while the *STF channel coefficients* $\left\{H_{r,s}\right\}$ are given by [cf. (4.47)]

$$\begin{aligned} H_{r,s} &\approx \frac{1}{\sqrt{N_{tr}}}\,\widetilde{H}(t,f)\big|_{(t,f)=(rT_o,sW_o)} = \frac{1}{\sqrt{N_{tr}}}\,\mathbf{u}_{f,s}^{\mathsf{T}}\widetilde{\mathbf{H}}\mathbf{u}_{t,r} \\ &= \frac{1}{\sqrt{N_{tr}}}\left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}}\right)\mathsf{vec}(\widetilde{\mathbf{H}}) = \frac{1}{\sqrt{N_{tr}}}\left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}}\right)\mathbf{h}_v\,,\ (r,s) \in \mathcal{S}_{tr} \end{aligned} \tag{4.59}$$

where $\widetilde{\mathbf{H}}$ is the $L\times(2M+1)$ *channel matrix* defined earlier in (4.51), $\mathbf{h}_v \stackrel{def}{=} \mathsf{vec}(\widetilde{\mathbf{H}}) \in \mathbb{C}^{L(2M+1)}$ is the vector of virtual channel coefficients, and the vectors $\mathbf{u}_{t,r}^{\mathsf{T}} \in \mathbb{C}^{2M+1}$ and $\mathbf{u}_{f,s}^{\mathsf{T}} \in \mathbb{C}^{L}$ are defined as $\mathbf{u}_{t,r}^{\mathsf{T}} \stackrel{def}{=} \left[e^{-jM\omega_{r,N_t}} \quad \ldots \quad e^{jM\omega_{r,N_t}}\right]$ and $\mathbf{u}_{f,s}^{\mathsf{T}} \stackrel{def}{=} \left[e^{-j0\omega_{s,N_f}} \quad \ldots \quad e^{-j(L-1)\omega_{s,N_f}}\right]$, respectively.[5] It can now be seen from (4.58) and (4.59) that stacking the received training data $\left\{y_{r,s}\right\}$ in this case into an N_{tr}-dimensional vector $\mathbf{y}$ yields the standard linear form

$$\mathbf{y} = \sqrt{\mathcal{E}}\,\mathbf{X}\mathbf{h}_v + \mathbf{z} \tag{4.60}$$

where the AWGN vector $\mathbf{z}$ is distributed as $\mathcal{CN}(\mathbf{0}_{N_{tr}}, \mathbf{I}_{N_{tr}})$, while the $N_{tr} \times L(2M+1)$ sensing matrix $\mathbf{X}$ is comprised of $\left\{\frac{1}{\sqrt{N_{tr}}}\left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}}\right) : (r,s) \in \mathcal{S}_{tr}\right\}$ as its rows. Consequently, traditional methods impose the condition $N_{tr} = \Omega(L(2M+1))$ in order to satisfy the requirement that $\mathbf{X}$ has full column rank in this setting and yield—at best—$\mathbb{E}[\Delta(\mathbf{h}_v^{\mathsf{LS}})] = \Omega(L(2M+1)/\mathcal{E})$. We now describe the CCS approach to estimating d-sparse doubly-selective channels using STF signaling, which not only has a lower reconstruction error than the LS-based approach but is also spectrally more efficient.

[5] It is instructive to remind the reader at this point that $T_o \in [\tau_{max}, 1/\nu_{max}]$ and $W_o \in [\nu_{max}, 1/\tau_{max}]$ imply here that $N_t \geq 2M+1$ and $N_f \geq L$, respectively.

$\big[$**CCS - 4**$\big]$ – STF Training and Reconstruction

Training: Pick $\mathcal{S}_{tr}$—the set of indices of pilot tones—to be a set of N_{tr} ordered pairs sampled uniformly at random (without replacement) from the set $\mathcal{S} = [0 \ldots N_t - 1] \times [0 \ldots N_f - 1]$.

Reconstruction: Fix any $a \geq 0$ and pick $\lambda = \sqrt{2\mathcal{E}(1+a)\log L(2M+1)}$. The CCS estimate of $\mathbf{h}_v$ is then given as a solution to the Dantzig selector

$$\mathbf{h}_v^{\mathsf{ccs}} = \underset{\mathbf{h} \in \mathbb{C}^{L(2M+1)}}{\arg\min} \big\|\mathbf{h}\big\|_1 \quad \text{subject to} \quad \Big\|\sqrt{\mathcal{E}}\,\mathbf{X}^{\mathsf{H}}\big(\mathbf{y} - \sqrt{\mathcal{E}}\,\mathbf{X}\mathbf{h}\big)\Big\|_{\infty} \leq \lambda\,. \tag{4.61}$$

Theorem 4.11 Suppose that $N_o, d > 2$, and let the number of pilot tones $N_{tr} \geq 2c_3 d \log^5 N_o$. Then for any $\delta_{2d} \in (0, 0.3]$, the CCS estimate of $\mathbf{h}_v$ satisfies

$$\Delta\left(\mathbf{h}_v^{\mathsf{ccs}}\right) \leq c_0^2 \cdot \frac{d}{\mathcal{E}} \cdot \log L(2M+1) \tag{4.62}$$

with probability $\geq 1 - 2\max\left\{2\Big(\pi(1+a)\log L(2M+1)\cdot\big(L(2M+1)\big)^{2a}\Big)^{-1/2}, 10N_o^{-c_4\delta_{2d}^2}\right\}$. Here, the constant $c_0 = 4\sqrt{2(1+a)}/(1-3\delta_{2d})$, while the constants $c_3 > 0$ and $c_4 > 0$ are the same as in Theorem 2.9.

Proof: To begin with, consider an $N_t \times N_t$ unitary (Fourier) matrix $\mathbf{U}_t$ that consists of

$$\left\{\frac{1}{\sqrt{N_t}}\left[e^{-jM\omega_{r,N_t}} \quad \ldots \quad e^{j(N_t-M-1)\omega_{r,N_t}}\right] : r = 0, 1, \ldots, N_t - 1\right\}$$

as its rows. In a similar fashion, consider another $N_f \times N_f$ unitary matrix $\mathbf{U}_f$ that consists of

$$\left\{\frac{1}{\sqrt{N_f}}\left[e^{-j0\omega_{s,N_f}} \quad \ldots \quad e^{-j(N_f-1)\omega_{s,N_f}}\right] : s = 0, 1, \ldots, N_f - 1\right\}$$

as its rows. Next, let $\mathbf{U} \stackrel{def}{=} \mathbf{U}_t \otimes \mathbf{U}_f$ be the Kronecker product of these two unitary matrices. The key thing to note here is that since the Kronecker product of two unitary matrices is a unitary matrix [44], the $N_o \times N_o$ matrix $\mathbf{U}$ is also a unitary matrix (recall that $N_t N_f = N_o$). It is then easy to see from (4.59) and (4.60) that $\mathbf{X}$ in this case simply corresponds to an $N_{tr} \times L(2M+1)$ column submatrix of an $N_{tr} \times N_o$ (randomly) subsampled unitary matrix $\mathbf{A}$ generated from the

$N_o \times N_o$ unitary matrix $\mathbf{U}$. The statement of the theorem then follows trivially from Theorem 2.9 and Theorem 2.13 using arguments similar to those made in the proof of Theorem 4.9. ∎

This concludes our discussion of the CCS framework for single-antenna channels; see Table 4.2 for a summary of the results presented in this section.

4.6 Compressed Channel Sensing: Multiple-Antenna Channels

4.6.1 Estimating Sparse Nonselective Channels

For a multiple-antenna (MIMO) channel that is nonselective (also, see "*Special Case: Nonselective Multiple-Antenna Channels*"), the virtual representation (4.9) of the channel reduces to

$$\widetilde{\mathbf{H}} = \sum_{i=1}^{N_R} \sum_{k=1}^{N_T} H_v(i,k) \mathbf{a}_R\left(\frac{i}{N_R}\right) \mathbf{a}_T^{\mathsf{H}}\left(\frac{k}{N_T}\right) = \mathbf{A}_R \mathbf{H}_v^{\mathsf{T}} \mathbf{A}_T^{\mathsf{H}}. \tag{4.63}$$

Here, $\mathbf{A}_R$ and $\mathbf{A}_T$ are $N_R \times N_R$ and $N_T \times N_T$ unitary (Fourier) matrices that are defined as

$$\mathbf{A}_R \overset{def}{=} \left[\mathbf{a}_R\big(\tfrac{1}{N_R}\big) \quad \mathbf{a}_R\big(\tfrac{2}{N_R}\big) \quad \ldots \quad \mathbf{a}_R\big(1\big)\right], \quad \mathbf{A}_T \overset{def}{=} \left[\mathbf{a}_T\big(\tfrac{1}{N_T}\big) \quad \mathbf{a}_T\big(\tfrac{2}{N_T}\big) \quad \ldots \quad \mathbf{a}_T\big(1\big)\right] \tag{4.64}$$

whereas $\mathbf{H}_v = \left[\mathbf{h}_{v,1} \quad \ldots \quad \mathbf{h}_{v,N_R}\right]$ is an $N_T \times N_R$ matrix of virtual channel coefficients in which the i-th column $\mathbf{h}_{v,i} \in \mathbb{C}^{N_T}$ consists of all channel coefficients $\big\{H_v(i,k)\big\}$ that are associated with the i-th resolvable AoA.

Generally speaking, the vector-valued training signal used to probe a nonselective MIMO channel can be written as

$$\mathbf{x}_{tr}(t) = \sqrt{\frac{\mathcal{E}}{N_T}} \sum_{n=0}^{N_{tr}-1} \tilde{\mathbf{x}}_n \, g\left(t - \frac{n}{W}\right), \; 0 \le t \le \frac{N_{tr}}{W} \tag{4.65}$$

where $g(t)$ is a prototype pulse having unit energy, $\big\{\tilde{\mathbf{x}}_n \in \mathbb{C}^{N_T}\big\}$ is the (vector-valued) training sequence having energy $\sum_n \|\tilde{\mathbf{x}}_n\|_2^2 = N_T$, and N_{tr}—the number of temporal training dimensions—denotes the total number of time slots dedicated to training in this setting. Trivially, matched filtering the received training signal $\mathbf{y}_{tr}(t) = \widetilde{\mathbf{H}}\mathbf{x}_{tr}(t) + \mathbf{z}_{tr}(t)$ in this case with time-shifted versions of the prototype pulse yields

$$\tilde{\mathbf{y}}_n = \sqrt{\frac{\mathcal{E}}{N_T}}\, \widetilde{\mathbf{H}} \tilde{\mathbf{x}}_n + \tilde{\mathbf{z}}_n, \; n = 0, \ldots, N_{tr} - 1 \tag{4.66}$$

where $\left\{\tilde{\mathbf{y}}_n \in \mathbb{C}^{N_R}\right\}$ is the (vector-valued) received training sequence and the AWGN vectors $\left\{\tilde{\mathbf{z}}_n\right\}$ are independently distributed as $\mathcal{CN}(\mathbf{0}_{N_R}, \mathbf{I}_{N_R})$. Next, the received training vectors $\left\{\tilde{\mathbf{y}}_n\right\}$ can be row-wise stacked into an $N_{tr} \times N_R$ matrix to obtain the following system of equations

$$\underbrace{\begin{bmatrix} \tilde{\mathbf{y}}_0^{\mathsf{T}} \\ \tilde{\mathbf{y}}_1^{\mathsf{T}} \\ \vdots \\ \tilde{\mathbf{y}}_{N_{tr}-1}^{\mathsf{T}} \end{bmatrix} = \sqrt{\frac{\mathcal{E}}{N_T}} \begin{bmatrix} \tilde{\mathbf{x}}_0^{\mathsf{T}} \\ \tilde{\mathbf{x}}_1^{\mathsf{T}} \\ \vdots \\ \tilde{\mathbf{x}}_{N_{tr}-1}^{\mathsf{T}} \end{bmatrix} \mathbf{A}_T^* \mathbf{H}_v \mathbf{A}_R^{\mathsf{T}} + \begin{bmatrix} \tilde{\mathbf{z}}_0^{\mathsf{T}} \\ \tilde{\mathbf{z}}_1^{\mathsf{T}} \\ \vdots \\ \tilde{\mathbf{z}}_{N_{tr}-1}^{\mathsf{T}} \end{bmatrix}}_{\widetilde{\mathbf{Y}} = \sqrt{\frac{\mathcal{E}}{N_T}} \widetilde{\mathbf{X}} \mathbf{A}_T^* \mathbf{H}_v \mathbf{A}_R^{\mathsf{T}} + \widetilde{\mathbf{Z}}}. \tag{4.67}$$

Finally, post-multiplying the $N_{tr} \times N_R$ matrix $\widetilde{\mathbf{Y}}$ with $\mathbf{A}_R^*$ yields the standard linear form (4.22)

$$\mathbf{Y} = \sqrt{\frac{\mathcal{E}}{N_T}} \mathbf{X} \mathbf{H}_v + \mathbf{Z}. \tag{4.68}$$

Here, $\mathbf{Y} \stackrel{def}{=} \widetilde{\mathbf{Y}} \mathbf{A}_R^*$, $\mathbf{X} \stackrel{def}{=} \widetilde{\mathbf{X}} \mathbf{A}_T^*$, and $\mathbf{Z} \stackrel{def}{=} \widetilde{\mathbf{Z}} \mathbf{A}_R^*$. Further, note that since $\mathbf{A}_R^*$ is a unitary matrix, the entries of the matrix $\mathbf{Z}$ are still independently distributed as $\mathcal{CN}(0, 1)$.

It is easy to see from the structure of the sensing matrix $\mathbf{X}$ in the above expression that the number of rows of $\mathbf{X}$ is exactly equal to the number of temporal training dimensions, N_{tr}. In order to estimate nonselective MIMO channels, traditional training-based methods such as those in [29, 30] therefore require that $N_{tr} = \Omega(N_T)$ so as to ensure that $\mathbf{X}$ has full column rank and produce an estimate that satisfies $\mathbb{E}[\Delta(\mathbf{H}_v^{\mathsf{LS}})] = \Omega(N_R N_T^2/\mathcal{E})$. In contrast, we now describe the CCS approach to this problem for d-sparse channels and quantify its performance in terms of the reconstruction error and spectral efficiency. Before proceeding further, however, it is instructive to recall that the conditional sparsity pattern associated with the i-th resolvable AoA is defined as $\mathcal{S}_d(i) \stackrel{def}{=} \left\{(i,k) : (i,k) \in \mathcal{S}_d\right\}$, while the *maximum* conditional AoA sparsity is accordingly defined as $\bar{d} \stackrel{def}{=} \max_i \left|\mathcal{S}_d(i)\right|$.

[CCS - 5] – Training and Reconstruction

Training: Pick the training sequence $\left\{\tilde{\mathbf{x}}_n, n = 0, \ldots, N_{tr} - 1\right\}$ associated with $\mathbf{x}_{tr}(t)$ to be a sequence of i.i.d. Rademacher random vectors in which each entry independently takes the value $+1/\sqrt{N_{tr}}$ or $-1/\sqrt{N_{tr}}$ with probability $1/2$ each.

Reconstruction: Fix any $a \geq 0$ and pick $\lambda = \sqrt{2\mathcal{E}(1+a)(\log N_R N_T)/N_T}$. Next, define

$$\mathbf{h}_{v,i}^{\mathsf{CCS}} = \underset{\mathbf{h}\in\mathbb{C}^{N_T}}{\arg\min} \left\|\mathbf{h}\right\|_1 \quad \text{subject to} \quad \left\| \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{X}^{\mathsf{H}}\big(\mathbf{y}_i - \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{X}\mathbf{h}\big)\right\|_\infty \leq \lambda,\ i = 1, \ldots, N_R$$

where $\mathbf{y}_i \in \mathbb{C}^{N_{tr}}$ denotes the i-th column of the matrix $\mathbf{Y}$. The CCS estimate of $\mathbf{H}_v$ is then simply given as follows: $\mathbf{H}_v^{\mathsf{CCS}} = \Big[\mathbf{h}_{v,1}^{\mathsf{CCS}} \quad \ldots \quad \mathbf{h}_{v,N_R}^{\mathsf{CCS}}\Big]$.

Theorem 4.12 Suppose that the number of training time slots (temporal training dimensions) satisfies $N_{tr} \geq 2c_2\bar{d}\log N_T$. Then for any $\delta_{2\bar{d}} \in (0, 0.3]$, the CCS estimate of $\mathbf{H}_v$ satisfies

$$\Delta\left(\mathbf{H}_v^{\mathsf{CCS}}\right) \leq c_0^2 \cdot \frac{d \cdot N_T}{\mathcal{E}} \cdot \log N_R N_T \tag{4.69}$$

with probability exceeding $1 - 2\max\left\{2\Big(\pi(1+a)\log N_R N_T \cdot (N_R N_T)^{2a}\Big)^{-1/2}, \exp(-c_1 N_{tr})\right\}$. Here, the constants $c_1 > 0$ and $c_2 > 0$ are the same as in Theorem 2.8 (with $\delta_{2\bar{d}}$ in place of δ_S), while the constant $c_0 = 4\sqrt{2(1+a)}/(1 - 3\delta_{2\bar{d}})$.

Proof: The claims of this theorem can be established using Theorem 2.8 and a slight modification of the proof of Theorem 2.13 in [68]. To begin with, fix some $\delta_{2\bar{d}} \in (0, 0.3]$ and note that since the $N_{tr} \times N_R$ matrix $\widetilde{\mathbf{X}}$ in (4.67) consists of the training vectors $\left\{\tilde{\mathbf{x}}_n\right\}$ as its row, the condition $N_{tr} \geq 2c_2\bar{d}\log N_T$ implies that the i.i.d. (binary) matrix $\widetilde{\mathbf{X}} \in RIP(2\bar{d}, \delta_{2\bar{d}})$ with probability exceeding $1 - \exp(-c_1 N_{tr})$ (cf. Theorem 2.8). This, in turn, implies that the matrix $\mathbf{X} = \widetilde{\mathbf{X}}\mathbf{A}_R^*$ also satisfies $RIP(2\bar{d}, \delta_{2\bar{d}})$ with the same probability, since $\mathbf{A}_R^*$ is a unitary matrix [11].

Next, notice that the i-th column of the matrix $\mathbf{Y}$ can be expressed as [cf. (4.67), (4.68)]

$$\mathbf{y}_i = \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{X}\mathbf{h}_{v,i} + \mathbf{z}_i,\ i = 1, \ldots, N_R \tag{4.70}$$

where $\mathbf{h}_{v,i}$ is the i-th column of $\mathbf{H}_v$, while the vector $\mathbf{z}_i \sim \mathcal{CN}(\mathbf{0}_{N_{tr}}, \mathbf{I}_{N_{tr}})$ denotes the i-th column of the AWGN matrix $\mathbf{Z}$. It is also instructive to note here that—by definition—we have the constraints $\|\mathbf{h}_{v,i}\|_0 \leq \bar{d}, i = 1, \ldots, N_R$. We now revisit a technical detail in the proof of Theorem 2.13 in [68, Section 3]. Specifically, define $c_i(\mathbf{h}) \stackrel{def}{=} \left\| \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{X}^{\mathsf{H}}\big(\mathbf{y}_i - \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{X}\mathbf{h}\big)\right\|_\infty,\ i = 1, \ldots, N_R$. Then it follows from the preceding discussion and [68, Section 3] that

$$\|\mathbf{h}_{v,i}^{\mathsf{CCS}} - \mathbf{h}_{v,i}\|_2^2 \leq c_0^2 \cdot \frac{\|\mathbf{h}_{v,i}\|_0 \cdot N_T}{\mathcal{E}} \cdot \log N_R N_T \tag{4.71}$$

provided $c_i(\mathbf{h}_{v,i}) \leq \lambda, i = 1, \ldots, N_R$. Therefore, since $\Delta\left(\mathbf{H}_v^{\mathsf{ccs}}\right) = \sum_{i=1}^{N_R} \|\mathbf{h}_{v,i}^{\mathsf{ccs}} - \mathbf{h}_{v,i}\|_2^2$ and also since $\sum_{i=1}^{N_R} \|\mathbf{h}_{v,i}\|_0 = d$, it is easy to conclude from this analysis that $\mathbf{H}_v^{\mathsf{ccs}}$ satisfies the reconstruction error bound (4.69) with probability exceeding

$$1 - \Pr\left(\Big\{\bigcup_{i=1}^{N_R} \big\{c_i(\mathbf{h}_{v,i}) > \lambda\big\}\Big\} \bigcup \big\{\mathbf{X} \notin RIP(2\bar{d}, \delta_{2\bar{d}})\big\}\right). \tag{4.72}$$

Finally, since we already have that $\Pr(\mathbf{X} \notin RIP(2\bar{d}, \delta_{2\bar{d}})) \leq \exp(-c_1 N_{tr})$, the only remaining thing is to upperbound the probability of the event $\bigcup_{i=1}^{N_R} \big\{c_i(\mathbf{h}_{v,i}) > \lambda\big\}$. To this end, first define

$$w_{i,j} \stackrel{def}{=} \langle \mathbf{z}_i, \mathbf{x}_j \rangle,\ i = 1, \ldots, N_R,\ j = 1, \ldots, N_T \tag{4.73}$$

where $\mathbf{x}_j \in \mathbb{C}^{N_{tr}}$ denote the j-th column of $\mathbf{X}$. Note that the $w_{i,j}$'s are identically (but not independently) distributed as $\mathcal{CN}(0,1)$, which follows from the fact that $\mathbf{z}_i \sim \mathcal{CN}(\mathbf{0}_{N_{tr}}, \mathbf{I}_{N_{tr}})$ and the columns of $\mathbf{X}$ have—by construction—unit ℓ_2-norms. Next, observe that $c_i(\mathbf{h}_{v,i}) > \lambda$ if and only if $\|\mathbf{X}^{\mathsf{H}}\mathbf{z}_i\|_\infty > \sqrt{2(1+a)\log N_R N_T}$ and, therefore, we have

$$\begin{aligned}
\Pr\left(\bigcup_{i=1}^{N_R} \big\{c_i(\mathbf{h}_{v,i}) > \lambda\big\}\right) &= \Pr\left(\max_i \|\mathbf{X}^{\mathsf{H}}\mathbf{z}_i\|_\infty > \sqrt{2(1+a)\log N_R N_T}\right) \\
&= \Pr\left(\max_{i,j} |w_{i,j}| > \sqrt{2(1+a)\log N_R N_T}\right) \\
&\stackrel{(a)}{\leq} 2N_R N_T \cdot \Pr\left(\big|\mathrm{Re}(w_{1,1})\big| > \sqrt{(1+a)\log N_R N_T}\right) \\
&\stackrel{(b)}{<} 2\Big(\pi(1+a)\log N_R N_T \cdot (N_R N_T)^{2a}\Big)^{-1/2}
\end{aligned} \tag{4.74}$$

where (a) follows from a simple union bounding argument (applied twice), while (b) mainly follows from the fact that $\Pr(|x| > u) < \frac{2}{\sqrt{2\pi}u} \exp\left(-\frac{1}{2}u^2\right)$ for $x \sim \mathcal{N}(0,1)$ and $u > 0$ [77]. The claimed probability bound in the theorem now follows by taking a final union bound over the events $\bigcup_{i=1}^{N_R} \big\{c_i(\mathbf{h}_{v,i}) > \lambda\big\}$ and $\big\{\mathbf{X} \notin RIP(2\bar{d}, \delta_{2\bar{d}})\big\}$. ∎

Remark 4.13 It is worth mentioning here that the statement of Theorem 4.12 remains unchanged if the training sequence $\big\{\tilde{\mathbf{x}}_n, n = 0, \ldots, N_{tr}-1\big\}$ described in $\big[\mathbf{CCS-5}\big]$ probes the MIMO channel in the so-called *beamspace* (instead of in the *antenna space*): $\big\{\tilde{\mathbf{x}}_n = \mathbf{A}_T\mathbf{x}_n, n = 0, \ldots, N_{tr}-1\big\}$, where now the beamspace vectors $\mathbf{x}_n \in \mathbb{C}^{N_T}$ are i.i.d. Rademacher random vectors in which

each entry independently takes the value $+1/\sqrt{N_{tr}}$ or $-1/\sqrt{N_{tr}}$ with probability $1/2$ each. In addition, we can also describe a variant of $\big[$**CCS - 5**$\big]$ in which a completely different (non-binary) training sequence is used without significantly altering the statement of Theorem 4.12. Specifically, let $\mathcal{S}_{tr}$ be a (sorted) set of $N_{tr} = 2c_3\bar{d}\log^5 N_T$ indices sampled uniformly at random (without replacement) from the set $[1 \dots N_T]$, and define the N_{tr} vectors in the training sequence to be $\big\{\tilde{\mathbf{x}}_n = \mathbf{e}_i : i \in \mathcal{S}_{tr}\big\}$. Here, the constant $c_3 > 0$ is the same as in Theorem 2.9, while $\mathbf{e}_i$ is used to denote the i-th column of $\mathbf{I}_{N_T}$. Then it can be shown that the reconstruction error bound (4.69) still holds in this case—the only difference here being that the second term in the $\max$ expression in Theorem 4.12 changes to $10N_T^{-c_4\delta^2_{2\bar{d}}}$, where the constant $c_4 > 0$ is given in Theorem 2.9.

Finally, before concluding this discussion, it is worth evaluating the minimum number of temporal training dimensions required for the CCS approach to succeed in the case of sparse nonselective MIMO channels. From the structure of the training sequence in $\big[$**CCS - 5**$\big]$ and the statement of Theorem 4.12, we have that CCS requires the number of training dimensions to satisfy $N_{tr} = \Omega(\bar{d}\log N_T)$, which—modulo the logarithmic factor—is never worse than $N_{tr} = \Omega(N_T)$ for traditional methods. In fact, for the case when the scattering geometry is such that the conditional AoA sparsity is equal to the average AoA sparsity ($\bar{d} = d/N_R$), we have that CCS requires the number of training dimensions to satisfy $N_{tr} = \Omega\left(\frac{d}{N_R} \cdot \log N_T\right)$.

4.6.2 Estimating Sparse Frequency-Selective Channels

In the case of a MIMO channel that is frequency selective, the virtual representation of the channel can be written as [cf. (4.9)]

$$\widetilde{\mathbf{H}}(f) = \sum_{\ell=0}^{L-1} \mathbf{A}_R \mathbf{H}_v^{\mathsf{T}}(\ell) \mathbf{A}_T^{\mathsf{H}} e^{-j2\pi \frac{\ell}{W} f} \tag{4.75}$$

where $\mathbf{H}_v(\ell) \stackrel{def}{=} \Big[\mathbf{h}_{v,1}(\ell) \quad \dots \quad \mathbf{h}_{v,N_R}(\ell)\Big]$, $\ell = 0, \dots, L-1$, are $N_T \times N_R$ matrices in which the i-th column $\mathbf{h}_{v,i}(\ell) \in \mathbb{C}^{N_T}$ consists of the coefficients $\big\{H_v(i,k,\ell)\big\}$, while $\mathbf{A}_R$ and $\mathbf{A}_T$ are the unitary (Fourier) matrices defined in (4.64).

As in the case of single-antenna channels, both SS and OFDM waveforms can be used to communicate over a frequency-selective MIMO channel. For the sake of this exposition, however,

we limit ourselves to OFDM signaling only. Paralleling the OFDM formulation for single-antenna channels, the corresponding training signal in the case of MIMO channels can be expressed as

$$\mathbf{x}_{tr}(t) = \sqrt{\frac{\mathcal{E}}{N_T}} \sum_{n \in \mathcal{S}_{tr}} \tilde{\mathbf{x}}_n \, g(t) e^{j2\pi \frac{n}{T} t}, \; 0 \le t \le T \tag{4.76}$$

where N_{tr} here again denotes the total number of pilot tones (equivalently, the number of temporal training dimensions), $\mathcal{S}_{tr} \subset \mathcal{S} \stackrel{def}{=} [0 \ldots N_o - 1]$ is the corresponding set of indices of pilot tones used for training ($|\mathcal{S}_{tr}| = N_{tr}$), and $\left\{\tilde{\mathbf{x}}_n \in \mathbb{C}^{N_T}\right\}$ is the (vector-valued) training sequence having energy $\sum_{\mathcal{S}_{tr}} \|\tilde{\mathbf{x}}_n\|_2^2 = N_T$. Matched filtering the received training signal $\mathbf{y}_{tr}(t) = \boldsymbol{\mathcal{H}}(\mathbf{x}_{tr}(t)) + \mathbf{z}_{tr}(t)$ with the OFDM basis waveforms $\left\{g(t)e^{j2\pi \frac{n}{T} t}\right\}_{\mathcal{S}_{tr}}$ yields in this case [21]

$$\tilde{\mathbf{y}}_n = \sqrt{\frac{\mathcal{E}}{N_T}} \mathbf{H}_n \tilde{\mathbf{x}}_n + \tilde{\mathbf{z}}_n, \; n \in \mathcal{S}_{tr} \tag{4.77}$$

where the AWGN vectors $\left\{\tilde{\mathbf{z}}_n\right\}$ are independently distributed as $\mathcal{CN}(\mathbf{0}_{N_R}, \mathbf{I}_{N_R})$, while the (matrix-valued) OFDM channel coefficients $\left\{\mathbf{H}_n\right\}$ are given by [cf. (4.75)]

$$\mathbf{H}_n \approx \widetilde{\mathbf{H}}(f)\big|_{f=\frac{n}{T}} = \sum_{\ell=0}^{L-1} \mathbf{A}_R \mathbf{H}_v^{\mathsf{T}}(\ell) \mathbf{A}_T^{\mathsf{H}} e^{-j2\pi \frac{\ell}{N_o} n}, \; n \in \mathcal{S}_{tr}. \tag{4.78}$$

Next, in order to represent the received training data (4.77) into the standard form (4.22), define row vectors $\left\{\mathbf{y}_n^{\mathsf{T}} \stackrel{def}{=} \tilde{\mathbf{y}}_n^{\mathsf{T}} \mathbf{A}_R^*\right\}_{\mathcal{S}_{tr}}$ and note from (4.77) and (4.78) that

$$\mathbf{y}_n^{\mathsf{T}} = \sqrt{\frac{\mathcal{E}}{N_T}} \tilde{\mathbf{x}}_n^{\mathsf{T}} \mathbf{A}_T^* \sum_{\ell=0}^{L-1} \mathbf{H}_v(\ell) e^{-j2\pi \frac{\ell}{N_o} n} + \mathbf{z}_n^{\mathsf{T}}, \; n \in \mathcal{S}_{tr} \tag{4.79}$$

where entries of the transformed noise vectors $\left\{\mathbf{z}_n^{\mathsf{T}} \stackrel{def}{=} \tilde{\mathbf{z}}_n^{\mathsf{T}} \mathbf{A}_R^*\right\}_{\mathcal{S}_{tr}}$ are still (mutually) independently distributed as $\mathcal{CN}(0,1)$ due to the unitary nature of $\mathbf{A}_R^*$. Now let $y_n(i), i = 1, \ldots, N_R$, denote the i-th entry of $\mathbf{y}_n^{\mathsf{T}}$, then it can be seen from (4.79) that

$$\begin{aligned} y_n(i) &= \sqrt{\frac{\mathcal{E}}{N_T}} \tilde{\mathbf{x}}_n^{\mathsf{T}} \mathbf{A}_T^* \sum_{\ell=0}^{L-1} \mathbf{h}_{v,i}(\ell) e^{-j2\pi \frac{\ell}{N_o} n} + z_n(i) \\ &= \sqrt{\frac{\mathcal{E}}{N_T}} \tilde{\mathbf{x}}_n^{\mathsf{T}} \mathbf{A}_T^* \mathbf{H}_{v,i} \mathbf{u}_n + z_n(i) \stackrel{(a)}{=} \sqrt{\frac{\mathcal{E}}{N_T}} (\mathbf{u}_n^{\mathsf{T}} \otimes \tilde{\mathbf{x}}_n^{\mathsf{T}} \mathbf{A}_T^*) \mathsf{vec}(\mathbf{H}_{v,i}) + z_n(i) \\ &\stackrel{(b)}{=} \sqrt{\frac{\mathcal{E}}{N_T}} \tilde{\mathbf{x}}_n^{\mathsf{T}} (\mathbf{u}_n^{\mathsf{T}} \otimes \mathbf{A}_T^*) \mathbf{h}_{v,i} + z_n(i), \quad i = 1, \ldots, N_R, \; n \in \mathcal{S}_{tr}. \end{aligned} \tag{4.80}$$

Here, $z_n(i)$ denotes the i-th entry of $\mathbf{z}_n^{\mathsf{T}}$, the matrix $\mathbf{H}_{v,i} \stackrel{def}{=} \begin{bmatrix} \mathbf{h}_{v,i}(0) & \mathbf{h}_{v,i}(1) & \dots & \mathbf{h}_{v,i}(L-1) \end{bmatrix}$ is an $N_T \times L$ matrix that consists of all channel coefficients $\{H_v(i,k,\ell)\}$ that are associated with the i-th resolvable AoA, $\mathbf{u}_n^{\mathsf{T}} \stackrel{def}{=} \begin{bmatrix} e^{-j0\omega_{n,N_o}} & \dots & e^{-j(L-1)\omega_{n,N_o}} \end{bmatrix}$ denotes the collection of L samples of a discrete sinusoid with frequency $\omega_{n,N_o} \stackrel{def}{=} 2\pi\frac{n}{N_o}$, and $\mathbf{h}_{v,i} \stackrel{def}{=} \mathsf{vec}(\mathbf{H}_{v,i})$ is just a vectorized version of $\mathbf{H}_{v,i}$. Finally, note that (a) in (4.80) follows from the identity $\mathsf{vec}(\mathbf{ABC}) = (\mathbf{C}^{\mathsf{T}} \otimes \mathbf{A})\mathsf{vec}(\mathbf{B})$, while (b) follows from the identity $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}$ [44].

It can now be easily seen from (4.79) and (4.80) that stacking the rows vectors $\{\mathbf{y}_n^{\mathsf{T}}\}_{\mathcal{S}_{tr}}$ into an $N_{tr} \times N_R$ matrix $\mathbf{Y}$ yields the standard linear form

$$\mathbf{Y} = \sqrt{\frac{\mathcal{E}}{N_T}}\,\mathbf{X}\mathbf{H}_v + \mathbf{Z} \tag{4.81}$$

where $\mathbf{H}_v = \begin{bmatrix} \mathbf{h}_{v,1} & \dots & \mathbf{h}_{v,N_R} \end{bmatrix}$ is the $N_T L \times N_R$ channel matrix whose i-th column $\mathbf{h}_{v,i} \in \mathbb{C}^{N_T L}$ consists of all channel coefficients $\{H_v(i,k,\ell)\}$ that are associated with the i-th resolvable AoA, while the matrix $\mathbf{X}$ is an $N_{tr} \times N_T L$ matrix comprising of $\{\tilde{\mathbf{x}}_n^{\mathsf{T}}(\mathbf{u}_n^{\mathsf{T}} \otimes \mathbf{A}_T^*) : n \in \mathcal{S}_{tr}\}$ as its rows. The key things to note here are that (i) the form of the sensing matrix $\mathbf{X}$ here once again dictates that $N_{tr} = \Omega(N_T L)$ for the traditional methods such as those in [31, 32] to obtain a meaningful estimate of $\mathbf{H}_v$, and (ii) we have from (4.30) that $\mathbb{E}[\Delta(\mathbf{H}_v^{\mathsf{LS}})] = \Omega(N_R N_T^2 L/\mathcal{E})$ in that case. In contrast, we now provide the CCS approach to estimating d-sparse frequency-selective MIMO channels using OFDM signaling and quantify its performance advantage over traditional methods. The following discussion once again makes use of the definition of maximum conditional sparsity within the AoA spread of the channel: $\bar{d} \stackrel{def}{=} \max_i \big|\{(i,k,\ell) : (i,k,\ell) \in \mathcal{S}_d\}\big|$.

$\big[$**CCS - 6**$\big]$ – OFDM Training and Reconstruction

Training: Pick $\mathcal{S}_{tr}$—the set of indices of pilot tones—to be a set of N_{tr} indices sampled uniformly at random (without replacement) from the set $\mathcal{S} = [0 \dots N_o - 1]$. Further, define the corresponding sequence of training vectors $\{\tilde{\mathbf{x}}_n, n \in \mathcal{S}_{tr}\}$ associated with $\mathbf{x}_{tr}(t)$ to be a sequence of i.i.d. Rademacher random vectors in which each entry independently takes the value $+1/\sqrt{N_{tr}}$ or $-1/\sqrt{N_{tr}}$ with probability $1/2$ each.

Reconstruction: Fix any $a \geq 0$ and pick $\lambda = \sqrt{2\mathcal{E}(1+a)(\log N_R N_T L)/N_T}$. Next, define

$$\mathbf{h}_{v,i}^{\mathsf{ccs}} = \underset{\mathbf{h}\in\mathbb{C}^{N_T L}}{\arg\min} \left\|\mathbf{h}\right\|_1 \quad \text{subject to} \quad \left\| \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{X}^{\mathsf{H}} \big(\mathbf{y}_i - \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{X}\mathbf{h}\big) \right\|_{\infty} \leq \lambda,\ i = 1, \ldots, N_R$$

where $\mathbf{y}_i \in \mathbb{C}^{N_{tr}}$ denotes the i-th column of the matrix $\mathbf{Y}$. The CCS estimate of $\mathbf{H}_v$ is then simply given as follows: $\mathbf{H}_v^{\mathsf{ccs}} = \Big[\mathbf{h}_{v,1}^{\mathsf{ccs}} \quad \ldots \quad \mathbf{h}_{v,N_R}^{\mathsf{ccs}}\Big]$.

Theorem 4.14 Suppose that $N_o, \bar{d} > 2$, and let the number of pilot tones $N_{tr} \geq (2c_8/c_9)\bar{d}\log^6 N_o$. Then for any $\delta_{2\bar{d}} \in (0, 0.3]$, the CCS estimate of $\mathbf{H}_v$ satisfies

$$\Delta\left(\mathbf{H}_v^{\mathsf{ccs}}\right) \leq c_0^2 \cdot \frac{d \cdot N_T}{\mathcal{E}} \cdot \log N_R N_T L \tag{4.82}$$

with probability exceeding $1 - 4\max\left\{\Big(\pi(1+a)\log N_R N_T L \cdot (N_R N_T L)^{2a}\Big)^{-1/2}, 10N_o^{-\delta_{2\bar{d}}^2}\right\}$. Here, the constants $c_8, c_9 > 0$ are the same as in Theorem 3.16, while $c_0 = 4\sqrt{2(1+a)}/(1 - 3\delta_{2\bar{d}})$.

Proof: To begin with, consider an $N_o \times N_o$ unitary (Fourier) matrix $\mathbf{U}_f$ that consists of

$$\left\{ \frac{1}{\sqrt{N_o}} \Big[e^{-j0\omega_{n,N_o}} \quad \ldots \quad e^{-j(N_o-1)\omega_{n,N_o}} \Big] : n = 0, 1, \ldots, N_o - 1 \right\}$$

as its rows. Next, let $\mathbf{U} \stackrel{def}{=} \mathbf{U}_f \otimes \mathbf{A}_T^*$ be the Kronecker product of the two unitary matrices $\mathbf{U}_f$ and $\mathbf{A}_T^*$. Note that $\mathbf{U}$ itself is a unitary matrix since it is the Kronecker product of two unitary matrices. Now the key thing to observe here is that the sensing matrix $\mathbf{X}$ comprising of $\left\{\tilde{\mathbf{x}}_n^{\mathsf{T}}(\mathbf{u}_n^{\mathsf{T}} \otimes \mathbf{A}_T^*) : n \in \mathcal{S}_{tr}\right\}$ as its rows is just an $N_{tr} \times N_T L$ column submatrix of an $N_{tr} \times N_T N_o$ structurally-subsampled unitary matrix $\mathbf{A}$ (with parameters $k = N_T$ and $n = N_{tr}$) that is generated from the $N_T N_o \times N_T N_o$ unitary matrix $\mathbf{U}$ and a Rademacher sequence (cf. Section 3.5). The statement of the theorem then follows from Theorem 3.16 and a slight modification of the proof of Theorem 2.13 in [68] using arguments similar to those made in the proof of Theorem 4.12. ∎

One key observation from the description of the training signal in $\big[$**CCS - 6**$\big]$ is that CCS requires $N_{tr} = \Omega(\bar{d} \cdot \log^6 N_o)$ for d-sparse frequency-selective MIMO channels. In particular, for the case of conditional AoA sparsity being equal to the average AoA sparsity, this implies that CCS requires $N_{tr} = \Omega(\frac{d}{N_R} \cdot \log^6 N_o)$ in this setting as opposed to $N_{tr} = \Omega(N_T L)$ for traditional methods—a significant improvement in terms of the training spectral efficiency when operating at large bandwidths and with large plurality of antennas since, by definition, $d \ll N_T L$.

4.6.3 Estimating Sparse Doubly-Selective Channels

In the most general case of a MIMO channel that is doubly selective, the virtual representation of the channel can be written as [cf. (4.9) and Figure 4.2(c)]

$$\widetilde{\mathbf{H}}(t,f) = \sum_{\ell=0}^{L-1} \sum_{m=-M}^{M} \mathbf{A}_R \mathbf{H}_v^{\mathsf{T}}(\ell,m) \mathbf{A}_T^{\mathsf{H}} e^{j2\pi \frac{m}{T} t} e^{-j2\pi \frac{\ell}{W} f} \tag{4.83}$$

where $\mathbf{H}_v(\ell,m) \stackrel{def}{=} \Big[\mathbf{h}_{v,1}(\ell,m) \quad \ldots \quad \mathbf{h}_{v,N_R}(\ell,m)\Big]$, $\ell = 0,\ldots,L-1,\ m = -M,\ldots,M$, are $N_T \times N_R$ matrices in which the i-th column $\mathbf{h}_{v,i}(\ell,m) \in \mathbb{C}^{N_T}$ consists of the channel coefficients $\big\{H_v(i,k,\ell,m)\big\}$, while $\mathbf{A}_R$ and $\mathbf{A}_T$ are the unitary (Fourier) matrices defined in (4.64).

As in the case of frequency-selective MIMO channels, we limit our discussion in this section to multi-carrier signaling only, even though both single- and multi-carrier waveforms can be used to communicate over doubly-selective MIMO channels. In particular, paralleling the multi-carrier formulation for single-antenna channels, the N_T-dimensional STF training signal in the case of MIMO channels can be expressed as

$$\mathbf{x}_{tr}(t) = \sqrt{\frac{\mathcal{E}}{N_{tr}}} \sum_{(r,s)\in\mathcal{S}_{tr}} \tilde{\mathbf{x}}_{r,s}\, g(t-rT_o) e^{j2\pi s W_o t},\ 0 \le t \le T \tag{4.84}$$

where N_{tr} here again denotes the total number of STF pilot tones used for training, the corresponding set $\mathcal{S}_{tr} \subset \mathcal{S} \stackrel{def}{=} [0\ldots N_t-1]\times[0\ldots N_f-1]$ is the set of indices of the pilot tones ($|\mathcal{S}_{tr}| = N_{tr}$), and $\big\{\tilde{\mathbf{x}}_{r,s} \in \mathbb{C}^{N_T}\big\}$ is the (vector-valued) training sequence having energy $\sum_{\mathcal{S}_{tr}} \|\tilde{\mathbf{x}}_{r,s}\|_2^2 = N_T$. Finally, as in Section 4.5.2, the STF basis parameters $T_o \in [\tau_{max}, 1/\nu_{max}]$ and $W_o \in [\nu_{max}, 1/\tau_{max}]$ are chosen here so that $T_o W_o = 1$, resulting in a total of $N_t N_f = N_o$ STF basis functions, where $N_t \stackrel{def}{=} T/T_o$ and $N_f \stackrel{def}{=} W/W_o$.

Next, under the assumption that the doubly-selective MIMO channel is sufficiently underspread so that the STF basis functions serve as approximate eigenfunctions of the channel, the received training signal $\mathbf{y}_{tr}(t) = \boldsymbol{\mathcal{H}}(\mathbf{x}_{tr}(t)) + \mathbf{z}_{tr}(t)$ can be matched filtered with $\Big\{g(t-rT_o)e^{j2\pi s W_o t}\Big\}$ to yield [126]

$$\tilde{\mathbf{y}}_{r,s} = \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{H}_{r,s}\tilde{\mathbf{x}}_{r,s} + \tilde{\mathbf{z}}_{r,s},\ (r,s) \in \mathcal{S}_{tr} \tag{4.85}$$

where the AWGN vectors $\{\tilde{\mathbf{z}}_{r,s}\}$ are independently distributed as $\mathcal{CN}(\mathbf{0}_{N_R}, \mathbf{I}_{N_R})$, while the STF channel coefficients $\{\mathbf{H}_{r,s}\}$ are given by [cf. (4.83)]

$$\mathbf{H}_{r,s} \approx \widetilde{\mathbf{H}}(t,f)\big|_{(t,f)=(rT_o, sW_o)} = \sum_{\ell=0}^{L-1} \sum_{m=-M}^{M} \mathbf{A}_R \mathbf{H}_v^{\mathsf{T}}(\ell, m) \mathbf{A}_T^{\mathsf{H}} e^{j2\pi \frac{m}{N_t} r} e^{-j2\pi \frac{\ell}{N_f} s}, \ (r,s) \in \mathcal{S}_{tr}. \quad (4.86)$$

Now define row vectors $\{\mathbf{y}_{r,s}^{\mathsf{T}} \stackrel{def}{=} \tilde{\mathbf{y}}_{r,s}^{\mathsf{T}} \mathbf{A}_R^*\}$, and note from (4.85) and (4.86) that

$$\mathbf{y}_{r,s}^{\mathsf{T}} = \sqrt{\frac{\mathcal{E}}{N_T}}\, \tilde{\mathbf{x}}_{r,s}^{\mathsf{T}} \mathbf{A}_T^* \underbrace{\sum_{\ell=0}^{L-1} \sum_{m=-M}^{M} \mathbf{H}_v(\ell, m) e^{j2\pi \frac{m}{N_t} r} e^{-j2\pi \frac{\ell}{N_f} s}}_{\mathbf{H}_v^{r,s}} + \mathbf{z}_{r,s}^{\mathsf{T}}, \ (r,s) \in \mathcal{S}_{tr} \quad (4.87)$$

where the entries of the transformed noise vectors $\{\mathbf{z}_{r,s}^{\mathsf{T}} \stackrel{def}{=} \tilde{\mathbf{z}}_{r,s}^{\mathsf{T}} \mathbf{A}_R^*\}_{\mathcal{S}_{tr}}$ are again independently distributed as $\mathcal{CN}(0,1)$ due to the unitary nature of $\mathbf{A}_R^*$.

In order to represent the received training data (4.87) into the standard form (4.22), we first focus on the $N_T \times N_R$ matrix $\mathbf{H}_v^{r,s}$ defined in (4.87). Specifically, if we let $h_v^{r,s}(k,i)$ denote the k-th entry of the i-th column of $\mathbf{H}_v^{r,s} \stackrel{def}{=} \sum\limits_{\ell,m} \mathbf{H}_v(\ell,m) e^{j2\pi \frac{m}{N_t} r} e^{-j2\pi \frac{\ell}{N_f} s}$ then it is easy to see that

$$\begin{aligned} h_v^{r,s}(k,i) &= \sum_{\ell=0}^{L-1} \sum_{m=-M}^{M} H_v(i,k,\ell,m) e^{j2\pi \frac{m}{N_t} r} e^{-j2\pi \frac{\ell}{N_f} s} \\ &= \mathbf{u}_{f,s}^{\mathsf{T}} \mathbf{H}_{v,i,k} \mathbf{u}_{t,r} = \left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}}\right) \mathsf{vec}(\mathbf{H}_{v,i,k}) \\ &= \left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}}\right) \mathbf{h}_{v,i,k}, \ k = 1, \ldots, N_T, \ i = 1, \ldots, N_R, \ (r,s) \in \mathcal{S}_{tr} \end{aligned} \quad (4.88)$$

where the vectors $\mathbf{u}_{t,r}^{\mathsf{T}} \in \mathbb{C}^{2M+1}$ and $\mathbf{u}_{f,s}^{\mathsf{T}} \in \mathbb{C}^{L}$ are defined as $\mathbf{u}_{t,r}^{\mathsf{T}} \stackrel{def}{=} \begin{bmatrix} e^{-jM\omega_{r,N_t}} & \ldots & e^{jM\omega_{r,N_t}} \end{bmatrix}$ and $\mathbf{u}_{f,s}^{\mathsf{T}} \stackrel{def}{=} \begin{bmatrix} e^{-j0\omega_{s,N_f}} & \ldots & e^{-j(L-1)\omega_{s,N_f}} \end{bmatrix}$, respectively, the matrix $\mathbf{H}_{v,i,k}$ is an $L \times (2M+1)$ matrix that consists of all channel coefficients $\{H_v(i,k,\ell,m)\}$ that are associated with the i-th resolvable AoA and k-th resolvable AoD

$$\mathbf{H}_{v,i,k} \stackrel{def}{=} \begin{bmatrix} H_v(i,k,0,-M) & H_v(i,k,0,-M+1) & \ldots & H_v(i,k,0,M) \\ H_v(i,k,1,-M) & H_v(i,k,1,-M+1) & \ldots & H_v(i,k,1,M) \\ \vdots & \vdots & & \vdots \\ H_v(i,k,L-1,-M) & H_v(i,k,L-1,-M+1) & \ldots & H_v(i,k,L-1,M) \end{bmatrix} \quad (4.89)$$

and $\mathbf{h}_{v,i,k} \stackrel{def}{=} \mathsf{vec}(\mathbf{H}_{v,i,k}) \in \mathbb{C}^{L(2M+1)}$ is just a vectorized version of $\mathbf{H}_{v,i,k}$. Therefore, if we use $\mathbf{h}_{v,i}^{r,s} \in \mathbb{C}^{N_T}$ to denote the i-th column of the matrix $\mathbf{H}_v^{r,s}$, then it follows from (4.88) that

$$\mathbf{h}_{v,i}^{r,s} = \underbrace{\begin{bmatrix} \mathbf{h}_{v,i,1}^{\mathsf{T}} \\ \mathbf{h}_{v,i,2}^{\mathsf{T}} \\ \vdots \\ \mathbf{h}_{v,i,N_T}^{\mathsf{T}} \end{bmatrix}}_{\mathbf{H}_{v,i}} (\mathbf{u}_{t,r} \otimes \mathbf{u}_{f,s}) \tag{4.90}$$

where the $N_T \times L(2M+1)$ matrix $\mathbf{H}_{v,i}$ defined in the above expression consists of all channel coefficients $\{H_v(i,k,\ell,m)\}$ that are associated with the i-th resolvable AoA.[6]

Finally, if we let $y_{r,s}(i), i = 1, \ldots, N_R$, denote the i-th entry of the row vector $\mathbf{y}_{r,s}^{\mathsf{T}}$, then it can be see from (4.87) and (4.90) that

$$\begin{aligned} y_{r,s}(i) &= \sqrt{\frac{\mathcal{E}}{N_T}}\, \tilde{\mathbf{x}}_{r,s}^{\mathsf{T}} \mathbf{A}_T^* \mathbf{H}_{v,i} \left(\mathbf{u}_{t,r} \otimes \mathbf{u}_{f,s}\right) + z_{r,s}(i) \\ &= \sqrt{\frac{\mathcal{E}}{N_T}}\, \left(\left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}}\right) \otimes \tilde{\mathbf{x}}_{r,s}^{\mathsf{T}} \mathbf{A}_T^*\right) \mathsf{vec}(\mathbf{H}_{v,i}) + z_{r,s}(i) \\ &\stackrel{(a)}{=} \sqrt{\frac{\mathcal{E}}{N_T}}\, \tilde{\mathbf{x}}_{r,s}^{\mathsf{T}} \left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}} \otimes \mathbf{A}_T^*\right) \mathbf{h}_{v,i} + z_{r,s}(i), \quad (r,s) \in \mathcal{S}_{tr}. \end{aligned} \tag{4.91}$$

Here, $z_{r,s}(i)$ denotes the i-th entry of $\mathbf{z}_{r,s}^{\mathsf{T}}$, while $\mathbf{h}_{v,i} \stackrel{def}{=} \mathsf{vec}(\mathbf{H}_{v,i}) \in \mathbb{C}^{N_T L(2M+1)}$ is just a vectorized version of $\mathbf{H}_{v,i}$. Note that (a) in (4.91) mainly follows from the identity $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}$ and the fact that the Kronecker product is associative [44]. It can now be easily seen from (4.87) and (4.91) that stacking the row vectors $\{\mathbf{y}_{r,s}^{\mathsf{T}}\}_{\mathcal{S}_{tr}}$ in this case into an $N_{tr} \times N_R$ matrix $\mathbf{Y}$ again yields the standard linear form

$$\mathbf{Y} = \sqrt{\frac{\mathcal{E}}{N_T}}\, \mathbf{X}\mathbf{H}_v + \mathbf{Z} \tag{4.92}$$

where $\mathbf{H}_v = \begin{bmatrix} \mathbf{h}_{v,1} & \ldots & \mathbf{h}_{v,N_R} \end{bmatrix}$ is the $N_T L(2M+1) \times N_R$ channel matrix whose i-th column $\mathbf{h}_{v,i} \in \mathbb{C}^{N_T L(2M+1)}$ consists of all channel coefficients $\{H_v(i,k,\ell,m)\}$ that are associated with the

[6]The reader can easily verify from the definition of $\otimes$ that the operation of transposition is distributive over the Kronecker product [44]: $(\mathbf{A} \otimes \mathbf{B})^{\mathsf{T}} = \mathbf{A}^{\mathsf{T}} \otimes \mathbf{B}^{\mathsf{T}}$.

i-th resolvable AoA, while the sensing matrix $\mathbf{X}$ is an $N_{tr} \times N_T L(2M+1)$ matrix comprising of $\left\{\tilde{\mathbf{x}}_{r,s}^{\mathsf{T}} \left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}} \otimes \mathbf{A}_T^*\right) : (r,s) \in \mathcal{S}_{tr}\right\}$ as its rows.

Note that, similar to the case of nonselective and frequency-selective MIMO channels, the expression (4.92) requires that $N_{tr} = \Omega(N_T L(2M+1))$ for the traditional methods to reliably use the LS criterion (4.27) in the case of doubly-selective MIMO channels. Further, even when this condition is satisfied, we have from (4.30) that $\mathbb{E}[\Delta(\mathbf{H}_v^{\mathsf{LS}})] = \Omega(N_R N_T^2 L(2M+1)/\mathcal{E})$ at the very best. Instead, we now provide the CCS approach to estimating d-sparse doubly-selective MIMO channels using STF signaling and quantify its performance advantage over traditional methods in terms of both the spectral efficiency and the reconstruction error. Before proceeding further, however, recall that the maximum conditional sparsity within the AoA spread of the channel is defined as $\bar{d} \stackrel{def}{=} \max_i \left|\left\{(i,k,\ell,m) : (i,k,\ell,m) \in \mathcal{S}_d\right\}\right|$.

$\left[\textbf{CCS - 7}\right]$ – STF Training and Reconstruction

---

Training: Pick $\mathcal{S}_{tr}$—the set of indices of pilot tones—to be a set of N_{tr} ordered pairs sampled uniformly at random (without replacement) from $\mathcal{S} = [0 \ldots N_t - 1] \times [0 \ldots N_f - 1]$. Further, define the corresponding sequence of training vectors $\left\{\tilde{\mathbf{x}}_{r,s}, (r,s) \in \mathcal{S}_{tr}\right\}$ associated with $\mathbf{x}_{tr}(t)$ to be a sequence of i.i.d. Rademacher random vectors in which each entry independently takes the value $+1/\sqrt{N_{tr}}$ or $-1/\sqrt{N_{tr}}$ with probability $1/2$ each.

Reconstruction: Fix some $a \geq 0$ and pick $\lambda = \sqrt{2\mathcal{E}(1+a)(\log N_R N_T L(2M+1))/N_T}$. Further, define $N_T L(2M+1)$-dimensional vectors

$$\mathbf{h}_{v,i}^{\mathsf{ccs}} = \underset{\mathbf{h} \in \mathbb{C}^{N_T L(2M+1)}}{\arg\min} \left\|\mathbf{h}\right\|_1 \quad \text{subject to} \quad \left\| \sqrt{\frac{\mathcal{E}}{N_T}} \mathbf{X}^{\mathsf{H}} \left(\mathbf{y}_i - \sqrt{\frac{\mathcal{E}}{N_T}} \mathbf{X}\mathbf{h}\right) \right\|_\infty \leq \lambda, \; i = 1, \ldots, N_R$$

where $\mathbf{y}_i \in \mathbb{C}^{N_{tr}}$ denotes the i-th column of the matrix $\mathbf{Y}$. The CCS estimate of the MIMO channel matrix $\mathbf{H}_v$ is then simply given as follows: $\mathbf{H}_v^{\mathsf{ccs}} = \left[\mathbf{h}_{v,1}^{\mathsf{ccs}} \quad \ldots \quad \mathbf{h}_{v,N_R}^{\mathsf{ccs}}\right]$.

Theorem 4.15 Suppose that $N_o, \bar{d} > 2$, and let the number of pilot tones $N_{tr} \geq (2c_8/c_9)\bar{d} \log^6 N_o$. Then for any $\delta_{2\bar{d}} \in (0, 0.3]$, the CCS estimate of $\mathbf{H}_v$ satisfies

$$\Delta\left(\mathbf{H}_v^{\mathsf{ccs}}\right) \leq c_0^2 \cdot \frac{d \cdot N_T}{\mathcal{E}} \cdot \log N_R N_T L(2M+1) \tag{4.93}$$

with probability $\geq 1 - 4\max\left\{\left(\pi(1+a)\log N_R N_T L(2M+1)\cdot\left(N_R N_T L(2M+1)\right)^{2a}\right)^{-1/2}, 10N_o^{-\delta_{2\bar{d}}^2}\right\}$. Here, the constants $c_8 > 0$ and $c_9 > 0$ are the same as in Theorem 3.16, while the constant $c_0 = 4\sqrt{2(1+a)}/(1-3\delta_{2\bar{d}})$.

Proof: To begin with, consider an $N_t \times N_t$ unitary (Fourier) matrix $\mathbf{U}_t$ that consists of

$$\left\{\frac{1}{\sqrt{N_t}}\left[e^{-jM\omega_{r,N_t}} \quad \ldots \quad e^{j(N_t-M-1)\omega_{r,N_t}}\right] : r = 0, 1, \ldots, N_t - 1\right\}$$

as its rows. In a similar fashion, consider another $N_f \times N_f$ unitary matrix $\mathbf{U}_f$ that consists of

$$\left\{\frac{1}{\sqrt{N_f}}\left[e^{-j0\omega_{s,N_f}} \quad \ldots \quad e^{-j(N_f-1)\omega_{s,N_f}}\right] : s = 0, 1, \ldots, N_f - 1\right\}$$

as its rows. Next, let $\mathbf{U} \stackrel{def}{=} \mathbf{U}_t \otimes \mathbf{U}_f \otimes \mathbf{A}_T^*$ be the Kronecker product of the unitary matrices $\mathbf{U}_t$, $\mathbf{U}_f$, and $\mathbf{A}_T^*$. Note that $\mathbf{U}$ is also a unitary matrix since it is the Kronecker product of three unitary matrices. Now, similar to the proof of Theorem 4.14, the key thing to observe here is that the sensing matrix $\mathbf{X}$ comprising of $\left\{\tilde{\mathbf{x}}_{r,s}^{\mathsf{T}}\left(\mathbf{u}_{t,r}^{\mathsf{T}} \otimes \mathbf{u}_{f,s}^{\mathsf{T}} \otimes \mathbf{A}_T^*\right) : (r,s) \in \mathcal{S}_{tr}\right\}$ as its rows is just an $N_{tr} \times N_T L(2M+1)$ column submatrix of an $N_{tr} \times N_T N_o$ structurally-subsampled unitary matrix $\mathbf{A}$ (with parameters $k = N_T$ and $n = N_{tr}$) that is generated from the $N_T N_o \times N_T N_o$ unitary matrix $\mathbf{U}$ and a Rademacher sequence (cf. Section 3.5). The statement of the theorem then follows from Theorem 3.16 and a slight modification of the proof of Theorem 2.13 in [68] using arguments similar to those made in the proof of Theorem 4.12. ∎

This concludes our discussion of the CCS framework for MIMO channels; see Table 4.2 for a summary of the results presented in this section.

4.7 Discussion

There is a large body of physical evidence that suggests that multipath signal components in many wireless channels tend to be distributed as clusters within their respective channel spreads. Consequently, as the world transitions from single-antenna communication systems operating at small bandwidths (typically in the megahertz range) to multiple-antenna ones operating at large bandwidths (possibly in the gigahertz range), the representation of such channels in appropriate

bases starts to look sparse. This has obvious implications for the design and implementation of training-based channel estimation methods. Since—by definition—the intrinsic dimension, d, of sparse multipath channels tends to be much smaller than their extrinsic dimension, D, one expects to estimate them using far fewer communication resources than that dictated by traditional methods based on the LS criterion. Equally importantly, however, sparsity of multipath channels also has implications for the design and implementation of the communication aspects of a wireless system that is equipped with a limited-rate feedback channel. First, if the channel-estimation module at the receiver yields a sparse estimate of the channel (something which LS-based reconstruction fails to accomplish) then—even at a low rate—that estimate can also be reliably fed back to the transmitter. Second, this reliable knowledge of the channel sparsity structure at both the transmitter and the receiver can be exploited by agile transceivers, such as the ones in [127], for improved communication performance.

In this chapter, we have described a new approach to estimating multipath channels that have a sparse representation in the Fourier basis. Our approach is based on some of the recent advances in the theory of compressed sensing and is accordingly termed as compressed channel sensing (CCS). Ignoring polylogarithmic factors, two distinct features of CCS are: (i) it has a reconstruction error that scales like $O(d)$ as opposed to $\Omega(D)$ for traditional methods, and (ii) it requires the number of temporal training dimensions, N_{tr}, to scale like $N_{tr} = \Omega(d/N_R)$ for certain signaling and channel configurations as opposed to $N_{tr} = \Omega(D/N_R)$ for traditional methods.

Before concluding our discussion, it is also worth commenting on some theoretical and practical aspects of CCS that have not been addressed earlier in this chapter. First, while there is no discussion of the optimality of CCS in here, we have established in [14, 15] that its performance for single-antenna sparse channels comes within a (poly)logarithmic factor of an (unrealizable) training-based method that clairvoyantly knows the channel sparsity pattern (also, see the accompanying numerical simulations in [14, 15]). Somewhat similar arguments can be made to argue the near-optimal nature of CCS for multiple-antenna sparse channels also. Second, the main ideas underlying the theory of CCS can be easily generalized to channel representations that make use of a basis other than the Fourier one and to other application areas such as high-resolution radar

imaging. Third, one expects the representation of real-world multipath channels in certain bases to be often approximately sparse because of the so-called *leakage* effect. While our primary focus in this chapter has been on characterizing the performance of CCS for exactly sparse channels, it works equally well for approximately sparse channels thanks to the near-optimal nature of the Dantzig selector; see, e.g., Theorem 2.13. Finally, and perhaps most importantly for the success of the envisioned wireless systems, we believe that CCS can be leveraged to design efficient training-based methods for estimating sparse *network* channels—a critical component of the emerging area of cognitive radio in which wireless transceivers sense and adapt to the wireless environment for enhanced spectral efficiency and interference management [128].

Chapter 5

Estimation of Sparse Networked Data

5.1 Introduction

Sensor networking is an emerging technology that promises an unprecedented ability to monitor the physical world via a spatially distributed network of small and inexpensive wireless devices that have the ability to self-organize into a well-connected network. A typical wireless sensor network (WSN), as shown in Figure 5.1, consists of a large number of wireless sensor nodes, spatially distributed over a region of interest, that can sense (and potentially actuate) the physical environment in a variety of modalities, including acoustic, seismic, thermal, and infrared [38]. A wide range of applications of sensor networks are being envisioned in a number of areas, including geographical monitoring (e.g., habitat monitoring, precision agriculture), industrial control (e.g., in a power plant or a submarine), business management (e.g., inventory tracking with radio frequency identification tags), homeland security (e.g., tracking and classifying moving targets) and health care (e.g., patient monitoring, personalized drug delivery) [39].

The essential task in many such applications of sensor networks is to extract relevant information about the sensed data—which we call *networked data* to emphasize both the distributed nature of the data and the fact that the data may be shared over the underlying communications infrastructure of the network—and deliver it with a desired fidelity to a (usually) distant destination, termed as the fusion center (FC). The overall goal in the design of sensor networks is to execute this task with least consumption of network resources—energy and bandwidth being the most limited resources, typically. In this regard, the relevant metrics of interest are: (i) the average total network power consumption P_{tot} for estimating the networked data, (ii) the distortion D in the

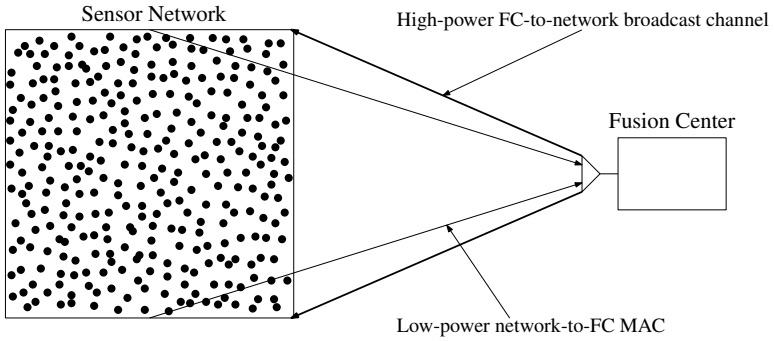

Figure 5.1 Sensor network with a fusion center (FC). Black dots denote sensor nodes. FC can communicate to the network over a high-power broadcast channel but the multiple-access channel (MAC) from the network to the FC is power constrained.

estimate, and (iii) the latency L incurred in obtaining the estimate (which is defined as the number of network-to-FC channel uses per estimate). It is also generally recognized that jointly optimizing the operations of sensing, processing, and communication can often lead to very energy efficient operation of sensor networks.

In this chapter, we propose a distributed joint source-channel communication architecture for energy efficient estimation of sensor field data at the FC. Under mild assumptions on the *spatial smoothness* of the signal field (cf. Section 5.2), we analyze the corresponding relationships between power, distortion, and latency as well as their scaling behavior with the number of sensor nodes. Our approach is inspired by recent results in wireless communications [40, 42, 43] and represents a new, non-traditional attack on the problem of sensing, processing, and communication in distributed wireless sensing systems. Rather than digitally encoding and transmitting samples from individual sensors, we consider an alternate encoding paradigm based on the projections of samples from many sensors onto appropriate spatial basis functions (e.g., local polynomials, wavelets). The joint source-channel communication architecture at the heart of our approach is an energy efficient method for communicating such projections to the FC—the projections are communicated in a phase-coherent fashion over the network-to-FC multiple-access channel (MAC). This architecture was first proposed and analyzed by us in [43] in the context of spatially homogeneous signal fields. In this chapter, we generalize that approach to a broader class of signals classified as either *compressible* or *sparse* (see Section 5.2).

The power of our proposed approach is that, in principle, one can choose to acquire samples in the domain of any basis that is particularly well-suited to the spatial structure of the signal field being sensed (e.g., smooth signals tend to be well-approximated in the Fourier basis and wavelet bases tend to be well-suited for the approximation of piecewise smooth signals [129]). Thus, if one has reasonable prior knowledge about the signal (e.g., spatial statistics or smoothness characteristics of the sensed field), then each sensing operation maximizes the potential gain in information per sample. More generally, however, we may have little prior knowledge about the sensed field. And, in some applications, the physical phenomenon of interest may contain time-varying spatial edges or boundaries that separate very different physical behaviors in the measured signal field (e.g., an oceanic oil spill, limited spatial distributions of hazardous biochemical agents). To handle such scenarios, we introduce the concept of *compressive wireless sensing* (CWS) in the later part of the chapter that is inspired by results in the theory of compressed sensing and fits perfectly into our proposed source-channel communication architecture.

The key idea in CWS is that neither the sensor nodes nor the FC need to know/specify the optimal basis elements in advance, and rests on the fact that a relatively small number of random projections of a compressible or sparse signal contain most of its salient information. Thus, in essence, CWS is a universal scheme based on delivering random projections of the networked data to the FC in an efficient manner. Under the right conditions, the FC can recover a good approximation of the data from these random projections. Nevertheless, this *universality* comes at the cost of a less favorable power-distortion-latency relationship that is a direct consequence of not exploiting prior knowledge of the signal field in the choice of projections that are communicated to the FC. This trade-off between universality and prior knowledge in CWS is quantified in Section 5.6.

5.1.1 Chapter Outline

The rest of this chapter is organized as follows. In Section 5.2, we describe the system model and associated assumptions on the networked data and the communication channel. In particular, in Section 5.2.1, we formalize the notions of compressible and sparse signals. In Section 5.3, we review the optimal distortion scaling benchmarks for compressible and sparse signals under the

assumption that the sensor measurements are available to the FC without any added cost or noise due to communications. In Section 5.4, we develop the basic building block in our source-channel communication architecture for computing and communicating projections of the networked data to the FC. Using this basic building block, we describe and analyze an energy efficient distributed estimation scheme in Section 5.5 that achieves the distortion scaling benchmarks of Section 5.3 for both compressible and sparse signals under the assumption of sufficient prior knowledge about the compressing (and sparse) basis. In Section 5.6, we introduce the concept of CWS for the case when sufficient prior knowledge about the compressing/sparse basis is not available and analyze the associated power-distortion-latency scaling laws. Up to this point, we operate under the assumptions that the network is fully synchronized and transmissions from the sensor nodes do not undergo fading. We relax these assumptions in Section 5.7 and study the impact of fading and imperfect phase synchronization on the scaling laws obtained in Sections 5.4, 5.5, and 5.6. Finally, we present some simulation results in Section 5.8 to illustrate the proposed methodologies and concluding remarks are provided in Section 5.9.

5.2 System Model and Assumptions

We begin by considering a WSN with n nodes observing some physical phenomenon in space and discrete-time[1], where each node takes a noisy sample at time index k of the form

$$x_j^k \;=\; s_j^k + w_j^k,\; j = 1, \ldots, n,\; k \in \mathbb{N} \tag{5.1}$$

and the noiseless samples $\left\{s_j^k,\; k \in \mathbb{N}\right\}$ at each sensor correspond to a deterministic *but unknown* sequence in $\mathbb{R}$. We further assume that $|s_j^k| \leq B$ $(\forall\, j = 1, \ldots, n,\; k \in \mathbb{N})$ for some known constant $B > 0$ that is determined by the sensing range of the sensors, and the measurement errors $\left\{w_j^k\right\}$ are zero-mean (real-valued) Gaussian random variables with variance σ_w^2 that are independent and identically distributed (i.i.d.) across space and time.

Note that the observed data $\left\{x_j^k = s_j^k + w_j^k\right\}_{j=1}^n$ at time k can be considered as a vector $\mathbf{x}^k \in \mathbb{R}^n$ such that $\mathbf{x}^k = \mathbf{s}^k + \mathbf{w}^k$, where $\mathbf{s}^k \in \mathbb{R}^n$ is the noiseless networked data and $\mathbf{w}^k \sim \mathcal{N}(\mathbf{0}_n, \sigma_w^2 \mathbf{I}_n)$

[1]The discrete-time model is an abstraction of the fact that the field is being temporally sampled at some rate of T_s seconds that depends upon the physics of the observed phenomenon.

is the measurement noise vector. Therefore, the physical phenomenon under observation can be characterized by the deterministic but unknown sequence of n-dimensional vectors

$$\mathbf{S} \stackrel{def}{=} \left\{\mathbf{s}^k\right\}_{k\in N} = \left\{\mathbf{s}^1, \mathbf{s}^2, \dots\right\}. \tag{5.2}$$

Furthermore, we assume no dependence between different time snapshots of the physical phenomenon. Note that if we were to model $\mathbf{S}$ as a stochastic signal, this would be equivalent to saying that $\mathbf{S}$ is a discrete (vector-valued) memoryless source.

5.2.1 Networked Data Model

It is a well-known fact in the field of transform coding that real-world signals can often be efficiently approximated and encoded in terms of Fourier, wavelet or other related transform representations [130–134]. For example, smooth signals can be accurately approximated using a truncated Fourier or wavelet series, and signals and images of bounded variation can be represented very well in terms of a relatively small number of wavelet coefficients [10, 129, 135]. Indeed, features such as smoothness and bounded variation are found in images, video, audio, and various other types of data, as evident from the success of familiar compression standards such as JPEG, MPEG and MP3 that are based on Fourier and wavelet transforms.

We take the transform coding point of view in modeling the signal observed by the sensor nodes. Specifically, we assume that the physical phenomenon described by $\mathbf{S}$ is (deterministic and) spatially compressible in the sense that each noiseless snapshot $\mathbf{s}^k$ is well-approximated by a linear combination of m vectors taken from an orthonormal basis of $\mathbb{R}^n$. We formalize this notion in the following definition.

Definition 5.1 (Compressible Signals) Let $\boldsymbol{\Psi} \stackrel{def}{=} \{\boldsymbol{\psi}_i\}_{i=1}^n$ be an orthonormal basis of $\mathbb{R}^n$. Denote the coefficients of $\mathbf{s}^k$ in this basis (inner products between $\mathbf{s}^k$ and the basis vectors $\boldsymbol{\psi}_i$) by writing $\theta_i^k \stackrel{def}{=} \boldsymbol{\psi}_i{}^{\mathsf{T}}\mathbf{s}^k = \sum_{j=1}^n \psi_{ij} s_j^k$. Re-index these coefficients of $\mathbf{s}^k$ and the corresponding basis vectors such that

$$|\theta_1^k| \geq |\theta_2^k| \geq \cdots \geq |\theta_n^k|. \tag{5.3}$$

The *best m-term approximation* of $\mathbf{s}^k$ in terms of $\boldsymbol{\Psi}$ is given by

$$\mathbf{s}^{k,(m)} \stackrel{def}{=} \sum_{i=1}^{m} \theta_i^k \boldsymbol{\psi}_i \tag{5.4}$$

and we say that $\mathbf{S}$ is α-compressible in $\boldsymbol{\Psi}$ (or that $\boldsymbol{\Psi}$ is the α-compressing basis of $\mathbf{S}$) if the *average squared-error* behaves like

$$\begin{aligned} \frac{1}{n} \left\| \mathbf{s}^k - \mathbf{s}^{k,(m)} \right\|_2^2 &\stackrel{def}{=} \frac{1}{n} \sum_{j=1}^{n} \left(s_j^k - s_j^{k,(m)} \right)^2 \\ &\leq C_o \, m^{-2\,\alpha}, \quad k \in \mathbb{N} \end{aligned} \tag{5.5}$$

for some constants $C_o > 0$ and $\alpha \geq 1/2$, where the parameter α governs the degree to which $\mathbf{S}$ is compressible with respect to $\boldsymbol{\Psi}$.

In addition, we will also consider the special case where, instead of being merely compressible, $\mathbf{S}$ is spatially sparse in the sense that each noiseless temporal sample $\mathbf{s}^k$ can be fully described by a few $\boldsymbol{\Psi}$-coefficients. We formalize this notion as follows.

Definition 5.2 (Sparse Signals) We say that $\mathbf{S}$ is M-sparse in $\boldsymbol{\Psi}$ (or that $\boldsymbol{\Psi}$ is the M-sparse basis of $\mathbf{S}$) if the following holds

$$\mathbf{s}^k = \sum_{i \in \mathcal{I}^k} \theta_i^k \boldsymbol{\psi}_i \,, \quad k \in \mathbb{N} \tag{5.6}$$

where $\mathcal{I}^k \subset [1 \ldots n]$, $k \in \mathbb{N}$, and $\max_k |\mathcal{I}^k| \leq M < n$, i.e., each networked data vector $\mathbf{s}^k$ has at most $M < n$ nonzero coefficients corresponding to some basis $\boldsymbol{\Psi}$ of $\mathbb{R}^n$.

Remark 5.3 An equivalent definition of compressibility or sparsity may be defined by assuming that, for some $0 < p \leq 1$ and some $R = R(n) > 0$, the $\boldsymbol{\Psi}$-coefficients of $\mathbf{s}^k$ belong to an ℓ_p ball of radius R [8, 10, 136], i.e.,

$$\left(\sum_{j=1}^{n} |\theta_j^k|^p \right)^{1/p} \leq R, \quad k \in \mathbb{N} \,. \tag{5.7}$$

To see that this is indeed an equivalent definition, first note that (5.7) can hold only if the cardinality of the set $\{\theta_j^k : |\theta_j^k| > 1/N, \, N \in \mathbb{N}, \, j = 1, 2, \ldots, n\}$ is upper bounded by $R\, N^{1/p}$ [136, 137].

Hence, the ℓ_p constraint of (5.7) in turn requires that the j-th largest (and re-indexed according to magnitude) coefficient θ_j^k is smaller than or equal to $R\, j^{-1/p}$, resulting in

$$\begin{aligned} \left\|\mathbf{s}^k - \mathbf{s}^{k,(m)}\right\|_2^2 &= \sum_{j=m+1}^{n} |\theta_j^k|^2 \\ &\leq C_p\, R^2\, m^{1-2/p}, \quad k \in \mathbb{N} \end{aligned} \tag{5.8}$$

for some constant C_p that depends only on p [8, 10, 136]. Thus, our definition of compressible signals is equivalent to assuming that the ordered $\boldsymbol{\Psi}$-coefficients of each networked data vector $\mathbf{s}^k$ exhibit a power law decay

$$|\theta_j^k| \leq R\, j^{-1/p}, \quad j = 1, \ldots, n,\ k \in \mathbb{N} \tag{5.9}$$

where $1/p = \alpha + 1/2$ and $R = \sqrt{\frac{n\, C_o}{C_p}}$ in our case [cf. (5.5), (5.8)]. Indeed, power law decays like this arise quite commonly in nature and we refer the readers to [8, 10, 131, 137] for some of those instances. Finally, with regard to the notion of sparsity, note that the ℓ_p constraint of (5.7) simply reduces to measuring the number of nonzero $\boldsymbol{\Psi}$-coefficients as $p \to 0$ and thus, corresponds to our definition of sparse signals with $R = M$.[2]

Remark 5.4 The above networked data model can be relaxed to allow temporal dependence between time snapshots of the physical phenomenon by assuming *spatio-temporal* compressibility (or sparsity) of the source signal in an appropriate space-time basis. While a detailed analysis of this setup is beyond the scope of this chapter, some of the techniques presented in this chapter can be extended to incorporate this scenario.

Remark 5.5 Note that while we are not concerned in this chapter with the issue of sensor placement (sampling) in the signal field, the choice of a good compressing basis is inherently coupled with the sensors' locations within the WSN. For example, while Fourier basis would suffice as a compressing basis for a sensor network observing a smooth signal field in which sensors are placed on a uniform grid, random (irregular) placement of sensors within the same field may warrant the use of an irregular wavelet transform as the appropriate compressing basis [138].

[2] For an M-sparse signal, no particular decay structure is assumed for the M nonzero coefficients of $\mathbf{s}^k$ in $\boldsymbol{\Psi}$.

5.2.2 Communication Setup

Given the observation vector $\mathbf{x}^k$ at time k, the aim of the sensor nodes (and the network as a whole) is to communicate a reliable-enough estimate $\widehat{\mathbf{s}}^k$ of the networked data vector $\mathbf{s}^k$ to a distant FC, where the reliability is measured in terms of the mean-squared error (MSE). Before proceeding further, however, we shall make the following assumptions concerning communications between the sensor nodes and the FC:

[1] Each sensor and the FC are equipped with a single omni-directional antenna and sensors communicate to the FC over a narrowband additive white Gaussian noise (AWGN) multiple-access channel (MAC), where each channel use is characterized by transmission over a period of T_c seconds. Furthermore, the FC can communicate to the sensor nodes over an essentially noise-free broadcast channel.

[2] Transmissions from the sensor nodes to the FC do not suffer any fading [139–141], which would indeed be the case in many remote sensing applications, such as desert border monitoring, with little or no scatterers in the surrounding environment and static sensor nodes having a strong line-of-sight connection to the FC [21].

[3] Each sensor knows its distance from the FC and thus, can calculate the channel path gain $\sqrt{h_j}$ given by [139–141]

$$\sqrt{h_j} \stackrel{def}{=} \frac{1}{{d_j}^{\zeta/2}}, \; j = 1, 2, \ldots, n \tag{5.10}$$

where $1 \leq d_j \leq d_u < \infty$ is the distance between the sensor at location j and the FC, and $\zeta \geq 2$ is the path-loss exponent [21, 142]. In principle, even when the distances and/or path loss exponent are unknown, these channel gains could be estimated at the FC using received signal strength and communicated back to the sensors during network initialization.

[4] The network is fully synchronized with the FC in the following sense [140, 141]: (i) *Carrier Synchronization*: All sensors have a local oscillator synchronized to the receiver carrier frequency; (ii) *Time Synchronization*: For each channel use, the relative timing error between sensors' transmissions is much smaller than the channel symbol duration T_c; and (iii)

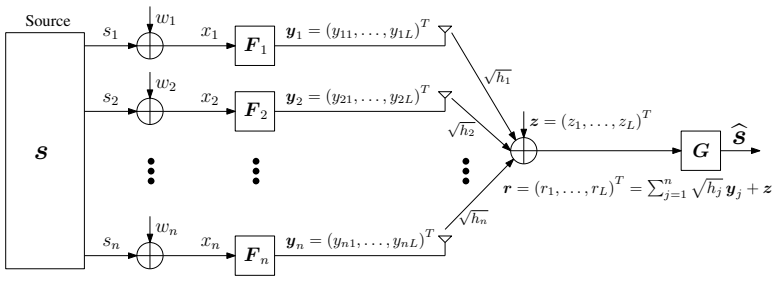

Figure 5.2 L-channel use snapshot of the sensor network per source observation. The superscript corresponding to the time index has been dropped in the figure to simplify notation.

Phase Synchronization: Sensors' transmissions arrive at the FC in a phase coherent fashion, which can be achieved by employing the distributed phase synchronization schemes described in [143, 144].

[5] Sensor transmissions are constrained to a sum transmit power of P per channel use. Specifically, let y_j be the transmission of sensor j in any channel use. Then, it is required that

$$\sum_{j=1}^{n} \mathbb{E}\left[|y_j|^2\right] \;\leq\; P. \tag{5.11}$$

[6] The network is allowed L network-to-FC channel uses per source observation, which we term as the latency of the system. If, for example, these L channel uses were to be employed using time division multiple access (TDMA) then this would require that the temporal sampling time $T_s \;\geq\; L\,T_c$; hence, the term latency. In a system with no bandwidth constraints, this could also be interpreted as the effective bandwidth of the network-to-FC MAC.

Given this communication setup, an estimation scheme corresponds to designing n source-channel encoders $(\boldsymbol{F}_1, \ldots, \boldsymbol{F}_n)$—one for each sensor node, and the decoder $\boldsymbol{G}$ for the FC such that at each time instant k, given the observations $\left\{x_j^\kappa\right\}_{\kappa=1}^k$ up to time k at node j, the encoders generate an L-tuple $\mathbf{y}_j^k \stackrel{def}{=} \boldsymbol{F}_j\left(\left\{x_j^\kappa\right\}_{\kappa=1}^k\right) = \begin{bmatrix} y_{j1}^k & \ldots & y_{jL}^k \end{bmatrix}^{\mathsf{T}}$ corresponding to L-channel uses per source observation (that also satisfy the power constraint of (5.11)). And at the end of the L-th channel use, the decoder $\boldsymbol{G}$ produces an estimate $\widehat{\mathbf{s}}^k$ of the networked data vector $\mathbf{s}^k$ given by $\widehat{\mathbf{s}}^k \stackrel{def}{=} \boldsymbol{G}\left(\left\{\mathbf{r}^\kappa\right\}_{\kappa=1}^k\right)$, where $\mathbf{r}^\kappa = \sum_{j=1}^n \sqrt{h_j}\,\mathbf{y}_j^\kappa + \mathbf{z}^\kappa$ and $\mathbf{z}^\kappa \sim \mathcal{N}(\mathbf{0}_L, \sigma_z^2 \mathbf{I}_L)$ is the MAC

AWGN vector corresponding to the L-channel uses at time instant κ (see Figure 5.2), and the goal of the sensor network is to minimize (i) the average total network power consumption per source observation

$$P_{\text{tot}} \stackrel{def}{=} \lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \sum_{\ell=1}^{L} \sum_{j=1}^{n} \mathbb{E}\left[\left|y_{j\ell}^{k}\right|^2\right]; \tag{5.12}$$

(ii) the mean-squared error distortion measure

$$D \stackrel{def}{=} \lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k - \widehat{\mathbf{s}}^k\right\|_2^2\right]; \tag{5.13}$$

and (iii) the latency L (# of channel uses per source observation) of the system.[3] Thus, for a fixed number of sensor nodes n, the performance of any estimation scheme is characterized by the triplet $(P_{\text{tot}}(n), D(n), L(n))$ and rather than obtaining an exact expression for this triplet, our goal would be to analyze how do these three quantities scale with n for a given scheme. Moreover, minimization of all three quantities in the triplet is sometimes a conflicting requirement and there is often a trade-off involved between minimizing P_{tot}, D and L, and we shall also be analyzing this power-distortion-latency trade-off as a function of n.

Remark 5.6 Notice that implicit in this formulation is the fact that no collaboration among the sensor nodes is allowed for the purposes of signal estimation, i.e., encoder $\boldsymbol{F}_j$ does not have access to the inputs of any sensor other than sensor j.

Remark 5.7 Note that while stating the performance metrics of power and latency, we have ignored the cost of initializing the sensor network (primarily corresponding to the cost of channel gain estimation/phase synchronization algorithms under the current communication setup and the cost of initial route/topology discovery algorithms under the more traditional multi-hop communication setups). This is because the average cost of this initialization (over time) tends to zero as k—the time scale of the network operation—tends to infinity. Of course, in practice, a one-time initialization may not suffice and these procedures may have to be repeated from time to time, but

[3] Notice that with the distortion metric as defined in (5.13), the MSE of any arbitrary length signal can at worst be a constant since $\lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k\right\|_2^2\right] \leq B^2$.

we will assume that the corresponding costs are negligible compared to the routine sensing and communication operations.

5.3 Optimal Distortion Scaling in a Centralized System

In this section, we consider a system in which the sensor measurements $\left\{x_j^k\right\}_{j=1}^n$ at each time instant k are assumed to be available at the FC with no added cost or noise due to communications, and we review the corresponding classical estimation theory results (see, e.g., [123,131,145]). Note that such a system corresponds to a sensor network with a noise-free network-to-FC MAC and thus, the optimal distortion scaling achievable under this *centralized* setting serves as a benchmark for assessing the distortion related performance of any scheme under the original setup.

5.3.1 Compressible Signals

Given the observation vector $\mathbf{x}^k$ at the FC, an optimal centralized estimator for an α-compressible signal can be easily constructed by projecting $\mathbf{x}^k$ onto the m basis vectors of $\boldsymbol{\Psi}$ corresponding to m largest (in the absolute sense) $\boldsymbol{\Psi}$-coefficients of $\mathbf{s}^k$ (see, e.g., [131]), i.e., if $\boldsymbol{\Psi}_m^k$ is the $n \times m$ matrix of those basis vectors, where the superscript k indicates that the re-indexing in (5.3) may be a function of the time index k, then $\mathbf{s}^k$ can be estimated as

$$\widehat{\mathbf{s}}_{\text{cen}}^k \stackrel{def}{=} \boldsymbol{\Psi}_m^k \left(\boldsymbol{\Psi}_m^{k\,\mathsf{T}} \mathbf{x}^k\right) = \mathbf{s}^{k,(m)} + \boldsymbol{\Psi}_m^k \left(\boldsymbol{\Psi}_m^{k\,\mathsf{T}} \mathbf{w}^k\right) \tag{5.14}$$

which results in

$$\mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k - \widehat{\mathbf{s}}_{\text{cen}}^k\right\|_2^2\right] = \frac{1}{n}\left\|\mathbf{s}^k - \mathbf{s}^{k,(m)}\right\|_2^2 + \frac{1}{n}\mathbb{E}\left[\left\|\boldsymbol{\Psi}_m^k \left(\boldsymbol{\Psi}_m^{k\,\mathsf{T}} \mathbf{w}^k\right)\right\|_2^2\right] \tag{5.15}$$

$$\leq C_o\, m^{-2\alpha} + \left(\frac{m}{n}\right)\sigma_w^2\,. \tag{5.16}$$

Furthermore, from (5.15), we also have the trivial lower bound of

$$\mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k - \widehat{\mathbf{s}}_{\text{cen}}^k\right\|_2^2\right] \geq \frac{1}{n}\mathbb{E}\left[\left\|\boldsymbol{\Psi}_m^k \left(\boldsymbol{\Psi}_m^{k\,\mathsf{T}} \mathbf{w}^k\right)\right\|_2^2\right] = \left(\frac{m}{n}\right)\sigma_w^2 \tag{5.17}$$

and combining the upper and lower bounds of (5.16) and (5.17), we obtain

$$\left(\frac{m}{n}\right)\sigma_w^2 \leq \mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k - \widehat{\mathbf{s}}_{\text{cen}}^k\right\|_2^2\right] \leq C_o\, m^{-2\alpha} + \left(\frac{m}{n}\right)\sigma_w^2\,. \tag{5.18}$$

From this expression, we see that the choice of m affects the classic bias-variance trade-off [145]: increasing m causes the bound $C_o\, m^{-2\alpha}$ on the approximation error $\frac{1}{n}\left\|\mathbf{s}^k - \mathbf{s}^{k,(m)}\right\|_2^2$ (the squared "bias") to decrease, but causes the stochastic component of the error due to the measurement noise $\frac{1}{n}\,\mathbb{E}\left[\left\|\boldsymbol{\Psi}_m^k\left(\boldsymbol{\Psi}_m^{k\,\mathsf{T}}\mathbf{w}^k\right)\right\|_2^2\right] = \left(\frac{m}{n}\right)\sigma_w^2$ (the "variance") to increase. The upper bound is tight, in the sense that there exist signals for which the upper bound is achieved, and in such cases the upper bound is minimized (by choice of m) by making the approximation error and the stochastic component of the error scale at the same rate, i.e.,

$$m^{-2\alpha} \asymp \frac{m}{n} \quad \Longleftrightarrow \quad m \asymp n^{1/(2\alpha+1)} \tag{5.19}$$

resulting in the following expression for optimal distortion scaling of an α-compressible signal in a centralized system[4]

$$D_{\text{cen}}^{*} \;=\; \lim_{K\to\infty}\frac{1}{K}\sum_{k=1}^{K}\mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k - \widehat{\mathbf{s}}_{\text{cen}}^k\right\|_2^2\right] \;\asymp\; n^{-2\alpha/(2\alpha+1)}. \tag{5.20}$$

5.3.2 Sparse Signals

Similar to a compressible signal, an optimal centralized estimator for an M-sparse signal corresponds to projecting the observation vector onto the M basis vectors of $\boldsymbol{\Psi}$ corresponding to M nonzero $\boldsymbol{\Psi}$-coefficients of $\mathbf{s}^k$ (see, e.g., [123]), i.e., if $\boldsymbol{\Psi}_M^k$ is the $n \times M$ matrix of those basis vectors, then $\mathbf{s}^k$ can be estimated as

$$\widehat{\mathbf{s}}_{\text{cen}}^k \;\stackrel{def}{=}\; \boldsymbol{\Psi}_M^k\left(\boldsymbol{\Psi}_M^{k\,\mathsf{T}}\mathbf{x}^k\right) \;=\; \mathbf{s}^k \;+\; \boldsymbol{\Psi}_M^k\left(\boldsymbol{\Psi}_M^{k\,\mathsf{T}}\mathbf{w}^k\right) \tag{5.21}$$

which results in the usual parametric rate

$$\mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k - \widehat{\mathbf{s}}_{\text{cen}}^k\right\|_2^2\right] \;=\; \frac{1}{n}\,\mathbb{E}\left[\left\|\boldsymbol{\Psi}_M^k\left(\boldsymbol{\Psi}_M^{k\,\mathsf{T}}\mathbf{w}^k\right)\right\|_2^2\right] \;=\; \left(\frac{M}{n}\right)\sigma_w^2 \tag{5.22}$$

resulting in the following expression for optimal distortion scaling of an M-sparse signal in a centralized system

$$D_{\text{cen}}^{*} \;=\; \left(\frac{M}{n}\right)\sigma_w^2 \;\asymp\; \frac{M}{n}\,. \tag{5.23}$$

[4] $*$ in D_{cen}^{*} refers to the fact that this is the *optimal* centralized distortion scaling.

Note that it might very well be that M—the number of degrees of freedom (DoF) of an M-sparse signal—scales with the number of nodes n in the network. For example, two-dimensional piecewise constant fields with one-dimensional boundaries separating constant regions can be compressed using the discrete wavelet transform and have $M \asymp n^{1/2}\log(n)$ nonzero wavelet coefficients [146]. Therefore, we model M as $M \asymp n^{\mu}$, where $0 \leq \mu < 1$ and hence, the inclusion of M in the scaling relation in (5.23).

Remark 5.8 Note that the optimal distortion scaling relations of (5.20) and (5.23) for compressible and sparse signals have been obtained under the assumption that the FC has precise knowledge of the ordering of coefficients of $\mathbf{s}^k$ in the compressing basis (indices of nonzero coefficients of $\mathbf{s}^k$ in the sparse basis). This is not necessarily a problem in a centralized setting and in cases where this information is not available, coefficient thresholding methods can be used to automatically select the appropriate basis elements from the noisy data, and these methods obey error bounds that are within a constant or logarithmic factor of the ones given above (see, e.g., [147, 148]).

5.4 Distributed Projections in Wireless Sensor Networks

In this section, we develop the basic communication architecture that acts as a building block of our proposed estimation scheme. As evident from the previous section, each DoF of a compressible or sparse signal corresponds to projection of networked data onto an n-dimensional vector in $\mathbb{R}^n$ and at the heart of our approach is a distributed method of communicating such projections to the FC in a power efficient manner by exploiting the spatial averaging inherent in an AWGN MAC.

To begin, assume that the goal of the sensor network is to obtain an estimate of the projection of networked data, corresponding to each observation of the physical phenomenon, onto a vector in $\mathbb{R}^n$ at the FC. That is, let us suppose that at each time instant k, we are interested in obtaining an estimate $\widehat{v}^k$ of

$$v^k \stackrel{def}{=} \boldsymbol{\varphi}^{\mathsf{T}}\mathbf{s}^k = \sum_{j=1}^{n} \varphi_j s_j^k \tag{5.24}$$

where $\boldsymbol{\varphi} \in \mathbb{R}^n$. One possibility for realizing this goal is to nominate a clusterhead in the network and then, assuming all the sensor nodes know their respective φ_j's and have constructed

routes which form a spanning tree through the network to the clusterhead, each sensor node can locally compute $\varphi_j x_j^k = \varphi_j(s_j^k + w_j^k)$ and these values can be aggregated up the tree to obtain $\widehat{v}^k = \sum_{j=1}^n \varphi_j x_j^k$ at the clusterhead, which can then encode and transmit this estimate to the FC. However, even if we ignore the communication cost of delivering $\widehat{v}^k$ from the clusterhead to the FC, it is easy to check that such a scheme requires at least n transmissions. For a similar reason, gossip algorithms such as the ones described in [149, 150], while known for their robustness in the face of changing network topology, might not be the schemes of first choice for these types of applications.

Another, more promising, alternative is to exploit recent results concerning uncoded (analog) coherent transmission schemes in WSNs [40–43]. The proposed distributed joint source-channel communication architecture requires only one channel use per source observation ($L_v = 1$) and is based on the notion of so-called "matched source-channel communication" [42, 43]: the structure of the network communication architecture should "match" the structure of the optimal estimator. Under the current setup, this essentially involves phase-coherent, low-power, analog transmission of appropriately weighted sample values directly from the nodes in the network to the FC via the AWGN network-to-FC MAC and the required projection is implicitly computed at the FC as a result of the spatial averaging in the MAC. In light of the communication setup of Section 5.2, full characterization of this architecture essentially entails characterization of the corresponding *scalar-output* source-channel encoders $(F_1, \ldots, F_n)$ at the sensor nodes and the *scalar-input* decoder G at the FC, where scalar nature of the encoders and the decoder is owing to the fact that (by construction) $L_v = 1$ in this scenario.

To begin with, each sensor encoder F_j in this architecture corresponds to simply multiplying the sensor measurement x_j^k with $\left(\sqrt{\frac{\rho}{h_j}}\,\varphi_j\right)$ to obtain[5]

$$y_j^k \stackrel{def}{=} F_j\left(\left\{x_j^\kappa\right\}_{\kappa=1}^k\right) = \sqrt{\frac{\rho}{h_j}}\,\varphi_j x_j^k, \quad j = 1, \ldots, n \tag{5.25}$$

where $\rho > 0$ is a scaling factor used to satisfy sensors' sum transmit power constraint P, and all the nodes coherently transmit their respective y_j^k's in an analog fashion over the network-to-FC

[5]Practical schemes of how each sensor encoder might get access to its respective φ_j is discussed in Section 5.5.3.

MAC. Under the synchronization assumption of Section 5.2 and the additive nature of an AWGN MAC, the corresponding received signal at the FC is given by

$$\begin{aligned} r^k \;&=\; \sum_{j=1}^{n} \sqrt{h_j}\, y_j^k + z^k \;=\; \sqrt{\rho} \sum_{j=1}^{n} \varphi_j x_j^k + z^k \\ &=\; \sqrt{\rho}\, \upsilon^k + \left(\sqrt{\rho}\, \boldsymbol{\varphi}^{\mathsf{T}} \mathbf{w}^k + z^k \right) \end{aligned} \tag{5.26}$$

where $z^k \sim \mathcal{N}(0, \sigma_z^2)$ is the MAC AWGN at time k (independent of $\mathbf{w}^k$). In essence, the encoders $(F_1, \ldots, F_n)$ correspond to delivering to the FC a noisy projection of $\mathbf{s}^k$ onto $\boldsymbol{\varphi}$ that is scaled by $\sqrt{\rho}$ [cf. (5.26)]. Given r^k, the decoder G corresponds to a simple re-scaling of the received signal

$$\begin{aligned} \widehat{\upsilon}^k \;&\stackrel{def}{=}\; G\left(\{r^\kappa\}_{\kappa=1}^k \right) \;=\; \frac{r^k}{\sqrt{\rho}} \\ &=\; \upsilon^k + \boldsymbol{\varphi}^{\mathsf{T}} \mathbf{w}^k + \frac{z^k}{\sqrt{\rho}}\,. \end{aligned} \tag{5.27}$$

We are now ready to characterize the power-distortion-latency triplet $(P_{\text{tot},\upsilon}\,, D_\upsilon\,, L_\upsilon)$ of the proposed joint source-channel communication architecture for computing distributed projections of networked data.[6]

Theorem 5.9 Let $\boldsymbol{\varphi} \in \mathbb{R}^n$ and let $\upsilon^k = \boldsymbol{\varphi}^{\mathsf{T}} \mathbf{s}^k$. Given the sensor network model of Section 5.2, the joint source-channel communication scheme described by the encoders in (5.25) and the decoder in (5.27) can achieve the following end-to-end distortion by employing only one channel use per source observation

$$\begin{aligned} D_\upsilon \;&\stackrel{def}{=}\; \lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \mathbb{E}\left[\left| \upsilon^k - \widehat{\upsilon}^k \right|^2 \right] \\ &=\; \sigma_w^2 \left\| \boldsymbol{\varphi} \right\|_2^2 \;+\; \left(\frac{\sigma_z^2\, {d_u}^\zeta\, (B^2 + \sigma_w^2)}{\lambda P} \right) \left\| \boldsymbol{\varphi} \right\|_2^2 \end{aligned} \tag{5.28}$$

where $\widehat{\upsilon}^k$ is the estimate of υ^k at the FC, σ_w^2 is the measurement noise variance, σ_z^2 is the channel noise variance, B is the bound on $|s_j^k|$, d_u is the bound on the maximum distance between the sensor nodes and the FC, ζ is the path-loss exponent, P is the sum transmit power constraint per

[6] $(P_{\text{tot},\upsilon}\,, D_\upsilon\,, L_\upsilon)$ triplet here corresponds to power, distortion, and latency of the projection coefficient as opposed to (P_{tot}, D, L) in Section 5.2 that corresponds to power, distortion, and latency required to estimate the entire signal.

channel use, and $\lambda = \lambda(n) \in (0,1]$ is a design parameter used to control total network power consumption. Moreover, the total network power consumption per source observation associated with achieving this distortion is given by

$$\lambda P \left(\frac{\sigma_w^2}{{d_u}^{\zeta} \left(B^2 + \sigma_w^2\right)} \right) \;\leq\; P_{\text{tot},\upsilon} \stackrel{def}{=} \lim_{K\to\infty} \frac{1}{K} \sum_{k=1}^{K} \sum_{j=1}^{n} \mathbb{E}\left[\left| y_j^k \right|^2 \right] \;\leq\; \lambda P. \tag{5.29}$$

Proof: To establish this theorem, first observe that (5.27) implies that $\forall\, k \in \mathbb{N}$

$$\begin{aligned} \mathbb{E}\left[\left| \upsilon^k - \widehat{\upsilon}^k \right|^2 \right] \;&=\; \mathbb{E}\left[\left| \boldsymbol{\varphi}^{\mathsf{T}} \mathbf{w}^k + \frac{z^k}{\sqrt{\rho}} \right|^2 \right] \\ &=\; \sigma_w^2 \left\| \boldsymbol{\varphi} \right\|_2^2 \;+\; \frac{\sigma_z^2}{\rho} \end{aligned} \tag{5.30}$$

resulting in the following expression for the projection coefficient MSE

$$D_{\upsilon} \;=\; \sigma_w^2 \left\| \boldsymbol{\varphi} \right\|_2^2 \;+\; \frac{\sigma_z^2}{\rho}\,. \tag{5.31}$$

As for obtaining an expression for $P_{\text{tot},\upsilon}$, note that (5.25) implies that $\forall\, k \in \mathbb{N}$,

$$\begin{aligned} \rho\; \sigma_w^2 \sum_{j=1}^{n} |\varphi_j|^2 \;&\leq\; \sum_{j=1}^{n} \mathbb{E}\left[\left| y_j^k \right|^2 \right] \\ &=\; \sum_{j=1}^{n} \mathbb{E}\left[\frac{\rho}{h_j} \left(s_j^k + w_j^k \right)^2 |\varphi_j|^2 \right] \\ &\leq\; \rho\, {d_u}^{\zeta} \left(B^2 + \sigma_w^2\right) \sum_{j=1}^{n} |\varphi_j|^2 \end{aligned} \tag{5.32}$$

and, therefore,

$$\rho \;=\; \lambda\, P \left(\frac{1}{{d_u}^{\zeta} \left(B^2 + \sigma_w^2\right) \left\| \boldsymbol{\varphi} \right\|_2^2} \right) \tag{5.33}$$

would suffice to satisfy the sum transmit power constraint of (5.11), where $\lambda = \lambda(n) \in (0,1]$ is a power scaling factor to be used by the designer of a WSN to control total network power consumption. This in turn results in the following expression for total network power consumption per source observation

$$\lambda P \left(\frac{\sigma_w^2}{{d_u}^{\zeta} \left(B^2 + \sigma_w^2\right)} \right) \;\leq\; P_{\text{tot},\upsilon} \;\leq\; \lambda P. \tag{5.34}$$

Finally, substituting the value of ρ from (5.33) in (5.31) yields (5.28), thereby completing the proof of the theorem. ∎

Notice that the projection coefficient distortion D_υ achieved by the proposed joint source-channel communication architecture has been expressed in terms of two separate contributions [cf. (5.28), (5.31)], the first of which is independent of the proposed communication scheme. This term is solely due to the noisy observation process ($\sigma_w^2 \neq 0$) and scales like $\|\boldsymbol{\varphi}\|_2^2$. The second contribution is primarily due to the noisy communication channel and scales like $\|\boldsymbol{\varphi}\|_2^2/\lambda$. Moreover, given the observation model of Section 5.2, it is easy to check that $D_\upsilon^* \asymp \|\boldsymbol{\varphi}\|_2^2$ is the best that any (centralized or distributed) scheme can hope to achieve in terms of an order relation for distortion scaling [123]. Therefore, for optimal distortion scaling, it is sufficient that the second term in (5.31) also scales like $\|\boldsymbol{\varphi}\|_2^2$ and hence, $\lambda = O(1)$ would suffice to ensure that

$$D_\upsilon \asymp \|\boldsymbol{\varphi}\|_2^2 \asymp D_\upsilon^*. \tag{5.35}$$

Consequently, the total network power consumption associated with achieving this optimal distortion scaling would be given by $P_{\text{tot},\upsilon} = O(1)$ [cf. (5.34)]. We summarize this insight as follows.

Corollary 5.10 Let $\boldsymbol{\varphi} \in \mathbb{R}^n$ and let $\upsilon^k = \boldsymbol{\varphi}^{\mathsf{T}}\mathbf{s}^k$. Given the sensor network model of Section 5.2 and assuming that the system parameters $(B, \sigma_w^2, \sigma_z^2, d_u, \zeta, P)$ do not vary with the number of nodes n in the network, the joint source-channel communication scheme described by the encoders in (5.25) and the decoder in (5.27) can obtain an estimate $\widehat{\upsilon}^k$ of υ^k at the FC, such that $D_\upsilon \asymp \|\boldsymbol{\varphi}\|_2^2 \asymp D_\upsilon^*$, by employing only one channel use per source observation, $L_\upsilon = 1$, and using a fixed amount of total network power, $P_{\text{tot},\upsilon} = O(1)$.

Observation 1 While the original problem has been setup under a *fixed* sum transmit power constraint P, one of the significant implications of the preceding analysis is that even if one allows P to grow with the number of nodes in the network—say, e.g., $P = O(n)$—one cannot improve on the distortion scaling law of $O(\|\boldsymbol{\varphi}\|_2^2)$. In other words, when it comes to estimating a single projection coefficient in the presence of noise, using more than a fixed amount of total power per channel use is wasteful as the distortion due to the measurement noise (first term in (5.31)) is the limiting factor in the overall distortion scaling.

Observation 2 Even though the joint source-channel communication architecture described in this section is meant to be a building block for the signal estimation scheme, the architecture is important in its own right too. Often times, for example, rather than obtaining an estimate of the networked data at the FC, the designer of a WSN is merely interested in obtaining the estimates of a few of its linear summary statistics. And, given that any linear summary statistic is nothing but the projection of networked data onto a vector in $\mathbb{R}^n$, preceding analysis implies that one can obtain such linear summary statistics at the FC with minimal distortion (and latency) and consumption of only a small amount of total network power.

Example 5.11 (Networked Data Average) To illustrate this idea further, consider a specific case where the designer of a WSN is interested in obtaining an estimate of the average $\bar{s}^k = \frac{1}{n}\sum_{j=1}^{n} s_j^k$ of networked data at each time instant k. This would correspond to the projection vector being given by $\boldsymbol{\varphi} = \begin{bmatrix} 1/n & \dots & 1/n \end{bmatrix}^{\mathsf{T}}$ and thus, using the communication architecture described in this section, an estimate of $\bar{s}^k$ can be obtained at the FC such that $D_{\bar{s}} \asymp 1/n \asymp D_{\bar{s}}^{*}$ (the parametric rate), $L_{\bar{s}} = 1$, and $P_{\text{tot},\bar{s}} = O(1)$.

5.5 Distributed Estimation from Noisy Projections: Known Subspace

In this section, we build upon the joint source-channel communication architecture of Section 5.4 and using it as a basic building block, present a completely decentralized scheme for efficient estimation of networked data at the FC. The analysis in this section is carried out under the assumption that the designer of the WSN has complete knowledge of the basis in which $\mathbf{S}$ is compressible (or sparse) as well as precise knowledge of the ordering of its coefficients in the compressing basis (indices of nonzero coefficients in the sparse basis) at each time instant k. We refer to this scenario as the *known subspace* case and, under this assumption, analyze the corresponding power-distortion-latency scaling laws of the proposed scheme as a function of number of sensor nodes in the network. As to the question of whether the known subspace assumption is a reasonable one, the answer depends entirely on the underlying physical phenomenon. For example, if the signal is smooth or bandlimited, then the Fourier or wavelet coefficients can be ordered (or partially ordered) from low frequency/resolution to high frequency/resolution. Alternatively, if

the physical phenomenon under observation happened to be spatially Hölder smooth at each time instant k, then it would be quite reasonable to treat the resulting networked data under the known subspace category (see, e.g., [43, 151]).

5.5.1 Estimation of Compressible Signals

To begin with, let $\boldsymbol{\Psi} = \{\boldsymbol{\psi}_i\}_{i=1}^{n}$ be the compressing basis of $\mathbf{S}$ such that $\frac{1}{n}\left\|\mathbf{s}^k - \mathbf{s}^{k,(m)}\right\|_2^2 = O(m^{-2\alpha})\ \forall\ k \in \mathbb{N}$. In Section 5.4, we showed that using the communication scheme described by the encoders in (5.25) and the decoder in (5.27), one projection per snapshot can be efficiently communicated to the FC by employing only one channel use ($L_v = 1$). By a simple extension of the encoders/decoder structure of Section 5.4, however, the network can equally well communicate $L\ (>1)$ projections per snapshot in L consecutive channel uses (one channel use *per projection* per snapshot). Essentially, at each time instant k, the L-tuples generated by the encoders $\boldsymbol{F}_j$ are given by (cf. Section 5.2, Figure 5.2)

$$\begin{aligned}\mathbf{y}_j^k \;&=\; \boldsymbol{F}_j\left(\left\{x_j^\kappa\right\}_{\kappa=1}^{k}\right) \;=\; \begin{bmatrix} y_{j1}^k & \dots & y_{jL}^k \end{bmatrix}^{\mathsf{T}} \\ &=\; \sqrt{\frac{\rho}{h_j}}\begin{bmatrix} \psi_{1j}\,x_j^k & \dots & \psi_{Lj}\,x_j^k \end{bmatrix}^{\mathsf{T}}, \quad j = 1,\dots,n \end{aligned} \tag{5.36}$$

where $\rho = (\lambda P)/(d_u{}^{\zeta}(B^2+\sigma_w^2))$, and at the end of the L-th channel use, the received signal at the input of the decoder $\boldsymbol{G}$ is given by

$$\begin{aligned}\mathbf{r}^k \;&=\; \sum_{j=1}^{n}\sqrt{h_j}\,\mathbf{y}_j^k \;+\; \mathbf{z}^k \\ &=\; \sqrt{\rho}\,\begin{bmatrix} \sum_{j=1}^{n}\psi_{1j}\,x_j^k & \dots & \sum_{j=1}^{n}\psi_{Lj}\,x_j^k \end{bmatrix}^{\mathsf{T}} \;+\; \mathbf{z}^k \\ &=\; \sqrt{\rho}\,\boldsymbol{\theta}_L^k \;+\; \left(\sqrt{\rho}\,\boldsymbol{\Psi}_L^{k\,\mathsf{T}}\mathbf{w}^k \;+\; \mathbf{z}^k\right) \end{aligned} \tag{5.37}$$

where $\boldsymbol{\Psi}_L^k$ is the $n \times L$ matrix of the basis vectors corresponding to L largest (in magnitude) $\boldsymbol{\Psi}$-coefficients of $\mathbf{s}^k$, $\boldsymbol{\theta}_L^k \overset{def}{=} \begin{bmatrix}\theta_1^k & \dots & \theta_L^k\end{bmatrix}^{\mathsf{T}} = \boldsymbol{\Psi}_L^{k\,\mathsf{T}}\mathbf{s}^k$ and $\mathbf{z}^k \sim \mathcal{N}(\mathbf{0}_L, \sigma_z^2\mathbf{I}_L)$ is the MAC AWGN vector (independent of $\mathbf{w}^k$). Thus, at the end of the L-th channel use, the decoder has access to L scaled, noisy projections of $\mathbf{s}^k$ onto L distinct elements of $\boldsymbol{\Psi}$ and, using these noisy projections, it

produces an estimate of the networked data vector $\mathbf{s}^k$ given by

$$\begin{aligned}\widehat{\mathbf{s}}^k \;&=\; \boldsymbol{G}\left(\{\mathbf{r}^\kappa\}_{\kappa=1}^k\right) \;=\; \boldsymbol{\Psi}_L^k\left(\frac{\mathbf{r}^k}{\sqrt{\rho}}\right)\\ &=\; \mathbf{s}^{k,(L)} \;+\; \boldsymbol{\Psi}_L^k\left(\boldsymbol{\Psi}_L^{k\,\mathsf{T}}\mathbf{w}^k\right) \;+\; \frac{\boldsymbol{\Psi}_L^k\,\mathbf{z}^k}{\sqrt{\rho}}\,. \end{aligned} \tag{5.38}$$

Notice the intuitively pleasing similarity between $\widehat{\mathbf{s}}^k$ and $\widehat{\mathbf{s}}^k_{\text{cen}}$ [cf. (5.14), (5.38)]: the first two terms in the above expression correspond identically to the centralized estimate of a compressible signal (with m replaced by L) and the last term is introduced due to the noisy MAC communication. In particular, this results in the following expression for distortion of a compressible signal at the FC

$$\begin{aligned}\left(\frac{L}{n}\right)\sigma_w^2 \;+\; \left(\frac{L}{n}\right)\left(\frac{\sigma_z^2\,{d_u}^{\zeta}\,(B^2+\sigma_w^2)}{\lambda P}\right) \;\leq D \;\leq\; C_o\,L^{-2\alpha} \;+\\ +\;\left(\frac{L}{n}\right)\sigma_w^2 \;+\left(\frac{L}{n}\right)\left(\frac{\sigma_z^2\,{d_u}^{\zeta}\,(B^2+\sigma_w^2)}{\lambda P}\right).\end{aligned} \tag{5.39}$$

Finally, simple manipulations along the lines of the ones in Section 5.4 result in the following expression for total network power consumption

$$\lambda\,L\,P\left(\frac{\sigma_w^2}{{d_u}^{\zeta}\,(B^2+\sigma_w^2)}\right) \;\leq\; P_{\text{tot}} \;\leq\; \lambda\,L\,P. \tag{5.40}$$

The above two expressions essentially govern the interplay between P_{tot}, D, and L of the proposed distributed estimation scheme and we shall analyze this interplay in further details in the sequel.

Minimum Power and Latency for Optimal Distortion Scaling

Similar to the case of distortion scaling in the centralized setting, (5.39) shows that the choice of number of projections per snapshot in the distributed setting also results in a bias-variance trade-off: increasing L causes the bound $C_o\,L^{-2\alpha}$ on the approximation error $\frac{1}{n}\left\|\mathbf{s}^k-\mathbf{s}^{k,(L)}\right\|_2^2$ to decrease, but causes the stochastic components of the error due to the measurement noise $\frac{1}{n}\,\mathbb{E}\left[\left\|\boldsymbol{\Psi}_L^k\left(\boldsymbol{\Psi}_L^{k\,\mathsf{T}}\mathbf{w}^k\right)\right\|_2^2\right] \;=\; \left(\frac{L}{n}\right)\sigma_w^2$ and the communication noise $\frac{1}{n}\,\mathbb{E}\left[\left\|\boldsymbol{\Psi}_L^k\,\mathbf{z}^k/\sqrt{\rho}\right\|_2^2\right] \;=\; \left(\frac{L}{n}\right)\cdot\left(\frac{\sigma_z^2\,{d_u}^{\zeta}\left(B^2+\sigma_w^2\right)}{\lambda P}\right)$ to increase. Consequently, the tightest upper bound scaling in (5.39) is attained by making the approximation error, the measurement noise error, and the communication noise

error scale (as a function of n) at the same rate. That is, assuming that the system parameters $(C_o, B, \sigma_w^2, \sigma_z^2, d_u, \zeta, P)$ do not depend on n,

$$L^{-2\alpha} \asymp \frac{L}{n} \asymp \frac{L}{\lambda n} \tag{5.41}$$

implying that L must be chosen, independently of λ, as

$$L \asymp n^{1/(2\alpha+1)} \tag{5.42}$$

which in turn requires that $\lambda = O(1)$, resulting in the following expression for optimal distortion scaling

$$D^* \asymp L^{-2\alpha} \asymp n^{-2\alpha/(2\alpha+1)} \tag{5.43}$$

that has the same scaling behavior as that of D^*_{cen} [cf. (5.20)]. Moreover, the total network power consumption associated with achieving this optimal distortion scaling is given by [see (5.40)]

$$P_{\text{tot}} \asymp L \asymp n^{1/(2\alpha+1)}. \tag{5.44}$$

Combining (5.42), (5.43), and (5.44), we can also compactly characterize the relationship between optimal distortion scaling and the associated power and latency requirements in terms of the following expression

$$D^* \sim {P_{\text{tot}}}^{-2\alpha} \sim L^{-2\alpha}. \tag{5.45}$$

Note that this expression does not mean that a WSN with fixed number of sensor nodes using more power and/or latency can provide better accuracy. Rather, power, distortion, and latency are functions of the number of nodes and the above relation indicates how the three performance metrics behave with respect to each other as the density of nodes increases.

Remark 5.12 Equation (5.44) shows that the total network power requirement of our proposed scheme for optimal distortion scaling $(D^* \asymp n^{-2\alpha/(2\alpha+1)})$ is given by $P_{\text{tot}} \asymp n^{1/(2\alpha+1)}$. A natural question is: *How good is this scheme in terms of power scaling?* While a comparison with all conceivable schemes does not seem possible, in order to give an idea of the performance of our

proposed scheme we compare it to a setup where all the nodes in the network noiselessly communicate their measurements to a designated cluster of $1 \leq \tilde{n} \leq n$ nodes. Each node in the cluster computes the required L projections of the measurement data for each snapshot and then all the $\tilde{n}$ nodes coherently transmit these (identical) projections to the FC over the MAC; in this case, the $\tilde{n} \times 1$ MAC is effectively transformed into a point-to-point AWGN channel with an $\tilde{n}$-fold power-pooling (beamforming) gain. One extreme, $\tilde{n} = 1$, corresponds to a single clusterhead (no beamforming gain), whereas the other extreme, $\tilde{n} = n$, corresponds to maximum beamforming gain. Note that in our proposed scheme, nodes transmit coherently (and hence benefit from power-pooling) but there is no data exchange between them. An exact comparison of our scheme with the above setup involving in-network data exchange is beyond the scope of this exposition since quantifying the cost of required in-network communication is challenging and requires making additional assumptions. Thus, we ignore the cost of in-network communication and provide a comparison just based on the cost of communicating the projections to the FC—though, in general, we expect the in-network cost to increase with the size $\tilde{n}$ of the cluster. Under this assumption, the analysis in Section 5.10.1 shows that our scheme requires less communication power compared to the $\tilde{n} = 1$ case, whereas it requires more power compared to the $\tilde{n} = n$ case. In particular, the power scaling achieved by our proposed scheme (for optimal distortion scaling) is identical to that in the case when there are $\tilde{n} = \frac{n}{L} \asymp n^{2\alpha/(2\alpha+1)}$ nodes in the designated cluster to coherently communicate the required $L \asymp n^{1/(2\alpha+1)}$ projection coefficients to the FC. Note that since $n^{2\alpha/(2\alpha+1)} \nearrow n$ for highly compressible signals ($\alpha \gg 1$), the performance of our proposed estimation scheme in this case approaches that of the $\tilde{n} = n$ extreme, without incurring any overhead of in-network communication.

Power-Distortion-Latency Scaling Laws for Consistent Estimation

Preceding analysis shows that in order to achieve the optimal centralized distortion scaling $n^{-2\alpha/(2\alpha+1)}$, the network must expend power P_{tot} and incur latency L that scale (with n) at a sublinear rate of $n^{1/(2\alpha+1)}$. This may pose a bottleneck in deploying dense WSNs for certain types of applications that might require extended battery life or faster temporal sampling of the physical

phenomenon. Cursory analysis of (5.39) and (5.40), however, shows that it is possible to lower these power and latency requirements at the expense of *sub-optimal* distortion scaling, and in the sequel, we shall be analyzing these power-distortion-latency scaling regimes.

Notice that under the assumption of system parameters $(C_o, B, \sigma_w^2, \sigma_z^2, d_u, \zeta, P)$ not varying with n, L and λ are the only two quantities that bear upon the required network power and achievable distortion of the estimation scheme [see (5.39), (5.40)]. Therefore, we begin by treating L (effective number of projections per snapshot) as an independent variable and model its scaling behavior as $L \asymp n^{\beta}$ for $\beta \in (0,1)$, while we model the scaling behavior of λ as $\lambda \asymp n^{-\delta}$ for $\delta \in [0,\infty)$ (recall, $0 < \lambda \leq 1$).[7] Note that $\beta = \beta^* \stackrel{def}{=} 1/(2\alpha+1)$ has already been solved previously (resulting in $\delta = \delta^* \stackrel{def}{=} 0$) and corresponds to the optimal distortion scaling of (5.43).

Bias-Limited Regime. Recall that $L \asymp n^{1/(2\alpha+1)}$ is the critical scaling of the number of projections at which point the distortion component due to the approximation error scales at the same rate as the distortion component due to the measurement noise [cf. (5.41), (5.42)]. If, however, we let L ($\asymp n^{\beta}$) scale at a rate such that $\beta < \beta^*$, then the first term in the upper bound in (5.39) that is due to the approximation error (bias term) starts to dominate the second term that is due to the measurement noise and, ignoring constants, the resulting distortion at the FC scales as

$$n^{-1+\beta+\delta} \preceq D \preceq n^{-2\alpha\beta} + n^{-1+\beta+\delta} \tag{5.46}$$

and the corresponding choice of *optimal* δ is given by

$$\delta = 1-(2\alpha+1)\beta \tag{5.47}$$

where optimal here refers to the fact that (i) $\delta < 1-(2\alpha+1)\beta$ is wasteful of power since distortion component due to the approximation error (first term in the upper bound in (5.46)) in that case decays slower than the distortion component to the communication noise (second term in the upper bound in (5.46)), and (ii) $\delta > 1-(2\alpha+1)\beta$ is wasteful of projections (i.e., latency) since distortion component due to the approximation error in that case decays faster than the distortion component

[7] There is nothing particular about choosing L as the independent variable except that it makes the analysis slightly easier. Nevertheless, we might as well start off by treating λ as the independent variable and reach exactly the same conclusions.

due to the communication noise. With this balancing of L and λ, distortion goes to zero at the rate

$$D \asymp n^{-2\alpha\beta} \tag{5.48}$$

as long as the chosen $\beta \in \big(0, 1/(2\alpha+1)\big)$, and the corresponding total network power consumption is given by [cf. (5.40)]

$$P_{\text{tot}} \asymp n^{(2\alpha+2)\beta-1} \tag{5.49}$$

resulting in the following expression for power-distortion-latency scaling relationship in the bias-limited regime

$$D \sim P_{\text{tot}}{}^{\frac{-2\alpha\beta}{(2\alpha+2)\beta-1}} \sim L^{-2\alpha}. \tag{5.50}$$

Variance-Limited Regime. On the other hand, if we let L scale at a rate such that $\beta > \beta^*$, then the second term in the upper bound in (5.39) that is due to the measurement noise (variance term) starts to dominate the bias term and the resulting distortion at the FC scales as

$$D \asymp n^{-1+\beta} + n^{-1+\beta+\delta} \tag{5.51}$$

and the corresponding choice of optimal δ is given by $\delta = 0$ ($= \delta^*$). This implies that as long as the chosen $\beta \in \big(1/(2\alpha+1), 1\big)$, distortion in the variance-limited regime goes to zero at the rate

$$D \asymp n^{-1+\beta} \tag{5.52}$$

and the corresponding total network power consumption is given by

$$P_{\text{tot}} \asymp n^{\beta} \tag{5.53}$$

resulting in the following expression for the power-distortion-latency scaling relationship in the variance-limited regime

$$D \sim P_{\text{tot}}{}^{-\frac{1}{\beta}+1} \sim L^{-\frac{1}{\beta}+1}. \tag{5.54}$$

Notice that as $\beta \to \beta^*$, both (5.50) and (5.54) collapse to the power-distortion-latency scaling relationship of (5.45), indicating that the optimal distortion scaling D^* corresponds to the transition

point between the bias-limited and variance-limited regimes. Thus, (5.50) and (5.54) completely characterize the power-distortion-latency scaling relationship of the proposed distributed estimation scheme for a compressible signal in the known subspace case. This scaling relationship is also illustrated in Figure 5.3, where the scaling exponents of P_{tot} and D are plotted against $\beta \in (0,1)$ (the chosen scaling exponent of L) for different values of α.

Observation 3 Analysis of (5.50), (5.54), and Figure 5.3 shows that (i) any distortion scaling that is achievable in the variance-limited regime is also achievable in the bias-limited regime, and (ii) scaling of P_{tot} in the variance-limited regime is uniformly worse than in the bias-limited regime. This implies that any WSN observing an α-compressible signal in the known subspace case should be operated only either in the bias-limited regime or at the optimal distortion scaling point, i.e., $\beta \in (0, 1/(2\alpha+1)]$. Thus, given α and a target distortion scaling of $D \asymp n^{-\gamma}$, $0 < \gamma \leq 2\alpha/(2\alpha+1)$, the number of projections computed by the WSN per snapshot needs to be scaled as $L \asymp n^{\beta}$, where $\beta = \gamma/2\alpha$ [cf. (5.50)], and the corresponding total network power consumption would be given by (5.49). Alternatively, the designer of a WSN could also reverse the roles of D and L by specifying a target latency scaling and obtaining the corresponding distortion (and power) scaling expression.

Observation 4 Another implication of the analysis carried out in this section is that the more compressible a signal is in a particular basis (i.e., the higher the value of α), the easier it is to estimate that signal in the bias-limited regime/at the optimal distortion scaling point (easier in terms of an improved power-distortion-latency relationship). On the other hand, note that the power-distortion-latency scaling in the variance-limited regime is completely independent of the parameter α [cf. (5.54)].

Observation 5 One of the most significant implication of the preceding analysis is that, while operating in the bias-limited regime, if β is chosen to be such that $\beta < 1/(2\alpha+2)$ then the scaling exponent of P_{tot} would be negative (cf. (5.49), Figure 5.3). This is remarkable since it shows that, in principle, consistent signal estimation is possible ($D \searrow 0$ as $n \to \infty$) even if the total network power consumption P_{tot} goes to zero!

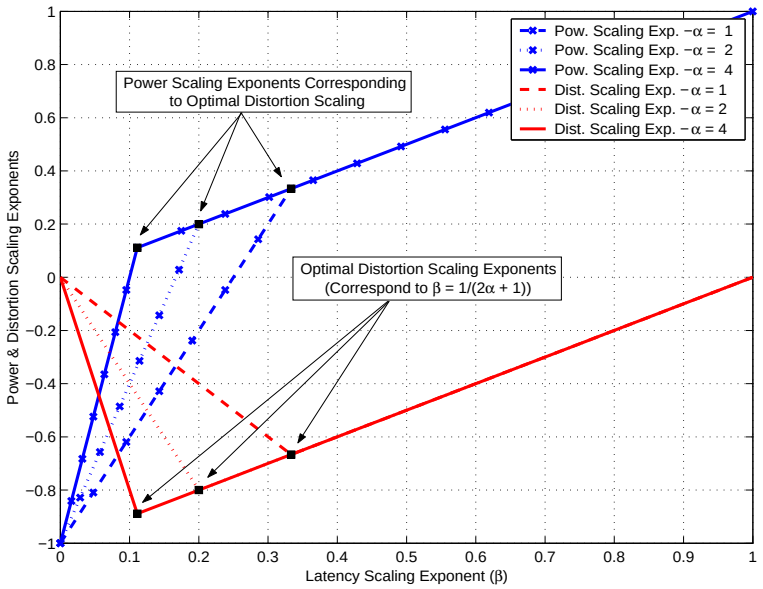

Figure 5.3 Power-distortion-latency scaling relationship of compressible signals in the known subspace case. The scaling exponents of P_{tot} and D are plotted against $\beta \in (0, 1)$ for different values of α. The filled black square on each curve corresponds to the operating point for optimal distortion scaling ($\beta = \beta^*$), with bias-limited and variance-limited regimes corresponding to the curve on its left and right side, respectively.

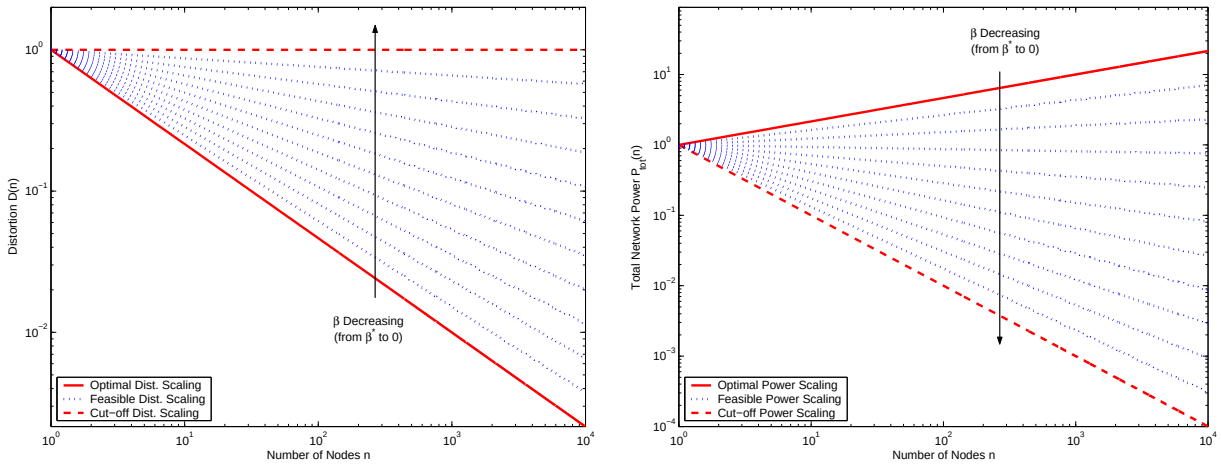

(a) Different scaling curves for D

(b) Different scaling curves for P_{tot}

Figure 5.4 Power-Density trade-off for compressible signals in the known subspace case. Various power and distortion scaling curves, each one corresponding to a different value of β, are plotted on a log-log scale against the number of nodes for $\alpha = 1$. The dashed curves are the cut-off scalings for consistent signal estimation ($\beta \searrow 0$).

Power-Density Trade-off

Viewed in a different way, Observation 5 also reveals a remarkable power-density trade-off inherent in our approach: *increasing the sensor density, while keeping the latency requirements the same, reduces the total network power consumption required to achieve a target distortion level.* This essentially follows from the fact that the power-distortion scaling law in the bias-limited regime (including the optimal distortion scaling point) follows a conservation relation given by [cf. (5.48), (5.49)]

$$P_{\text{tot}}\, D \asymp n^{2\beta-1}. \tag{5.55}$$

Specifically, let $\beta_2 < \beta_1$ denote two latency scalings in the bias-limited regime and let $n_2 > n_1$ denote the corresponding number of nodes needed to achieve a target distortion level $D(n_1, \beta_1) = D(n_2, \beta_2) = D_o$. Then, we have from (5.48) that

$$D(n_1, \beta_1) = D(n_2, \beta_2) \quad \Rightarrow \quad n_1^{-2\alpha\beta_1} = n_2^{-2\alpha\beta_2} \tag{5.56}$$

and therefore, the corresponding latency requirements are rather trivially related to each other as

$$\frac{L_2(n_2, \beta_2)}{L_1(n_1, \beta_1)} = \frac{n_2^{\beta_2}}{n_1^{\beta_1}} = 1. \tag{5.57}$$

Moreover, it follows from (5.55) that the total network power consumptions in the two cases are related by

$$\frac{P_{\text{tot}}(n_2, \beta_2)}{P_{\text{tot}}(n_1, \beta_1)} = \frac{n_2^{2\beta_2-1}}{n_1^{2\beta_1-1}} = \frac{n_1}{n_2} \tag{5.58}$$

where we have used the fact that (5.56) implies that $n_1^{2\beta_1} = n_2^{2\beta_2}$. Relations (5.57) and (5.58) show that increasing the sensor density by a factor of N, while keeping the number of projections (per snapshot) communicated by the network to the FC the same, reduces the total network power required to attain a given target distortion by a factor of N.

This power-density trade-off is also illustrated in Figure 5.4, where various power and distortion scaling curves (corresponding to different values of β) are plotted on a log-log scale against the number of nodes for $\alpha = 1$. For the sake of illustration, these plots assume that the constants of

proportionality in the scaling relations are unity. In order to illustrate the power-density trade-off, suppose that we want to attain a target distortion of $D_o = 0.02$. With optimal distortion scaling (solid curve in Figure 5.4(a) corresponding to $\beta_1 = \beta^* = 1/3$), the desired distortion can be attained with $n_1 \approx 353$ nodes, consuming a total network power $P_{\text{tot}}(n_1, \beta_1) \approx 7.07$, as calculated from the solid curve in Figure 5.4(b). On the other hand, however, if we operate on a sub-optimal distortion scaling curve (say, e.g., the third dotted feasible curve from the bottom in Figure 5.4(a) corresponding to $\beta_2 = 8/33$), we would attain the desired distortion of $D_o = 0.02$ with $n_2 \approx 3192$ nodes—roughly a factor of 9 increase in sensor density—but would only consume a total network power of $P_{\text{tot}}(n_2, , \beta_2) \approx 0.78$, as calculated from the third dotted feasible curve from the top in Figure 5.4(b). Thus, as predicted, increasing the sensor density roughly by a factor of 9 reduces the total network power consumption by a factor of 9, while the latency requirements stay exactly the same ($n_1^{\beta_1} = n_2^{\beta_2}$).

5.5.2 Estimation of Sparse Signals

The analysis for the estimation of an M-sparse signal in the known subspace case using the joint source-channel communication architecture of Section 5.4 can be carried out along almost the same lines as for compressible signals in Section 5.5.1, with the only obvious difference being that $\mathbf{S}$ now lies exactly in an M-dimensional subspace of $\mathbb{R}^n$ and therefore, L has to be taken exactly equal to M, i.e., is no longer a variable parameter in the hands of the designer of a WSN. This results in the following expressions for the end-to-end distortion at the FC and the corresponding total network power consumption and system latency

$$D \;=\; \left(\frac{M}{n}\right)\sigma_w^2 \;+\; \left(\frac{M}{n}\right)\left(\frac{\sigma_z^2\, {d_u}^{\zeta}\,(B^2+\sigma_w^2)}{\lambda P}\right), \tag{5.59}$$

$$\lambda\, MP\left(\frac{\sigma_w^2}{{d_u}^{\zeta}\,(B^2+\sigma_w^2)}\right) \;\leq\; P_{\text{tot}} \;\leq\; \lambda\, MP, \tag{5.60}$$

$$L \;=\; M. \tag{5.61}$$

Next, recall that while the scaling behavior of M is modeled as $M \asymp n^{\mu}$, $0 \leq \mu < 1$, the choice of μ is not in our hands in this case and, instead, depends upon the underlying physical phenomenon. Essentially, μ here plays the role analogous to that of α in the compressible case.

Therefore, unlike the compressible signal case, the only controllable parameter in this case is the power scaling factor λ, modeled as $\lambda \asymp n^{-\delta}$ for $\delta \in [0, \infty)$, and in the sequel, we analyze the effect of various scaling behaviors of λ on the power-distortion-latency scaling relationship of the proposed estimation scheme.

Power-Distortion-Latency Scaling Laws for Consistent Estimation

We start out by first analyzing the optimal distortion scaling that is achievable for an M-sparse signal. Notice that for fastest distortion reduction, the first term (due to the measurement noise) in (5.59) should scale at the same rate as the second term (due to the communication noise). This in turn requires that $\lambda = O(1)$ (or $\delta = \delta^* \stackrel{def}{=} 0$), resulting in the following expression for optimal distortion scaling

$$D^* \asymp \frac{M}{n} \asymp n^{-1+\mu} \tag{5.62}$$

which has the same scaling behavior as that of D^*_{cen} [cf. (5.23)]. Moreover, the total network power consumption associated with achieving this optimal distortion scaling is given by

$$P_{\text{tot}} \asymp L = M \asymp n^{\mu}. \tag{5.63}$$

Equations (5.61), (5.62) and (5.63) can also be combined together for $\mu \in (0, 1)$ to express the relationship between optimal distortion scaling and the corresponding power and latency requirements in terms of the following expression

$$D^* \sim {P_{\text{tot}}}^{-\frac{1}{\mu}+1} \sim L^{-\frac{1}{\mu}+1}. \tag{5.64}$$

Notice that the above relationship has the same form as that of (5.54) (the power-distortion-latency relationship of a compressible signal in the variance-limited regime) which is precisely what one would expect since there is no bias related distortion component for a sparse signal [cf. (5.59)].

Remark 5.13 Equations (5.61), (5.62), and (5.63) show that for the case of $\mu = 0$, i.e., $M = O(1)$, optimal distortion scaling $D^* \asymp n^{-1}$ can be obtained by consuming only a fixed amount of total network power and incurring a fixed latency, i.e., $P_{\text{tot}} = L = O(1)$. This result is similar in spirit

to the one obtained in [42] that primarily studies the case analogous to that of a sparse signal with non-scaling DoF ($M = O(1)$), albeit assuming Gaussian sources and multiple FCs (see Theorem 1 and Theorem 3 therein).

Power-limited Regime. On the other hand, if we take $\delta > 0$ then the distortion component due to the communication noise (second term in (5.59)) starts to dominate the distortion component due to the measurement noise (first term in (5.59)) and, ignoring the constant parameters, the resulting distortion at the FC scales as

$$D \asymp n^{-1+\mu+\delta} \tag{5.65}$$

in the power-limited regime. This implies that as long as $\delta \in (0, 1-\mu)$, distortion can still be driven to zero, albeit at a slower, *sub-optimal* rate of $n^{-1+\mu+\delta}$ ($\succ D^*$). In particular, this means that D can be asymptotically driven to zero even if the total network power P_{tot} ($\asymp n^{\mu-\delta}$) scales just a little faster than $n^{2\mu-1}$ [cf. (5.60)]. This observation is similar in spirit to the one made for compressible signals since it shows that, in principle, consistent signal estimation is possible in the limit of a large number of nodes for $\mu \in [0, 1/2]$ (i.e., the number of DoF M scaling at most as fast as $\sqrt{n}$) even if the total network power P_{tot} goes to zero. Finally, this power-distortion-latency scaling relationship in the power-limited regime can be expressed as

$$D \sim P_{\text{tot}}^{\frac{-1+\mu+\delta}{\mu-\delta}} \sim L^{\frac{-1+\delta}{\mu}+1}. \tag{5.66}$$

Discussion

While quite similar in spirit, there are still some key differences between the power-distortion-latency scaling laws of the proposed estimation scheme for compressible and sparse signals. To begin with, unlike for compressible signals, the latency scaling requirements for sparse signals are dictated by the underlying physical phenomenon ($L = M \asymp n^{\mu}$) and cannot be traded-off for power and/or distortion without making further assumptions on the decay characteristics of the M nonzero coefficients of $\mathbf{S}$. Secondly, the scenario of consistent signal estimation of sparse signals with decaying total network power consumption exists if and only if the number of DoF M scales

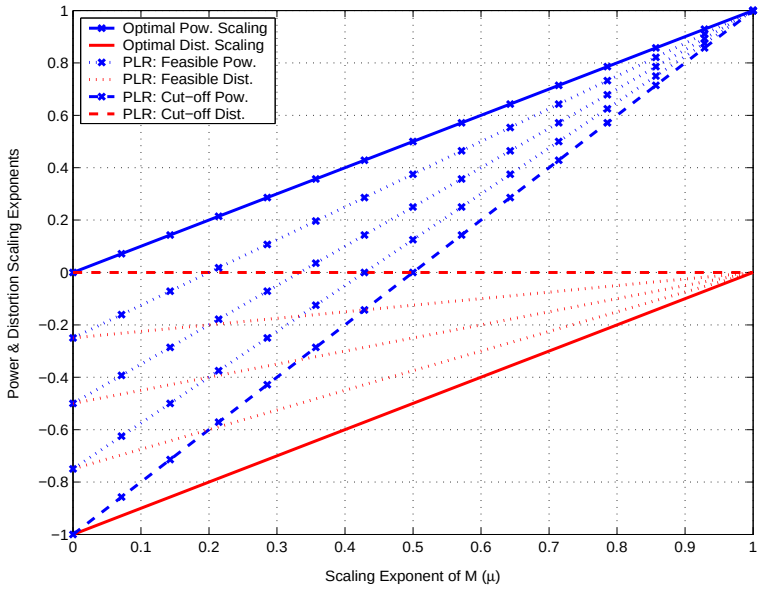

Figure 5.5 Power-distortion-latency scaling relationship of sparse signals in the known subspace case. The scaling exponents of P_{tot} and D are plotted against the scaling exponent of M ($\asymp n^{\mu}$) for $0 \leq \mu < 1$, while the scaling exponent of L is the same as that of M. The solid curves correspond to the optimal distortion scaling exponents and the corresponding total network power scaling exponents, while the dotted curves correspond to various power-limited regime (PLR) scalings that result in consistent signal estimation. The dashed curves are the cut-off scalings for consistent signal estimation.

at a rate less than or equal to $\sqrt{n}$, i.e., $\mu \leq 1/2$ (see Figure 5.5).[8] And finally, as a flip side to this observation, the power-density trade-off for sparse signals exists only when $0 \leq \mu < 1/2$, happens to be a function of μ, and is not as pronounced for $0 < \mu < 1/2$. Specifically, for an increase in the sensor density by a factor of N, the total network power consumption requirements can only be reduced by a factor of $N^{1-2\mu}$, $0 \leq \mu < 1/2$, in order to attain the same target distortion for a sparse signal [cf. (5.65)].

5.5.3 Communicating the Projection Vectors to the Network

Recall that an implicit requirement for employing the proposed distributed estimation scheme in the known subspace case is that the sensor encoders have access to the respective projection vectors' elements at each time instant k [cf. (5.36)]. In this subsection, we address the issue of how one might communicate this information to the sensor nodes. One viable option in this regard could be the pre-storage of relevant information in each sensor node. However, pre-storage of the entire compressing (sparse) basis $\boldsymbol{\Psi}$ or a subset of it, $\{\boldsymbol{\psi}_i\}_{i=1}^{L}$, where $1 \leq L \leq n$, in each sensor node is not feasible in large-scale WSNs since this would require at least n storage elements per sensor node. Instead, a better alternative is to store only the corresponding *nonzero* elements of the L projection vectors, $\{\psi_{ij} : \psi_{ij} \neq 0\}_{i=1}^{L}$, in the j-th sensor node. In the context of [43], for example, this would mean having only $O(1)$ storage elements per sensor node, since the structure of the proposed projection vectors in [43] is such that the cardinality of the set $\{\psi_{ij} : \psi_{ij} \neq 0\}_{i=1}^{L}$ is identically equal to one $\forall\ j = 1, \ldots, n$. Other instances when pre-storage might be a feasible option could be, for example, when the projection vectors' elements come from an analytical expression. Pre-storage, however, suffers from the drawback that sensor nodes pre-stored with one compressing (sparse) basis vectors might not be readily deployable in signal fields compressible (sparse) in some other basis.

Another more feasible, but not always practical, approach to the communication of projection vectors to the network could be that the FC transmits this information over the FC-to-network

[8] Recall that for compressible signals, this observation holds true for all α.

broadcast channel at either the start of the estimation process or at the start of each network-to-FC channel use. For basis $\boldsymbol{\Psi}$ whose vectors have some sort of spatial regularity in their structure such that they do not require addressing each sensor node individually (e.g., vectors describable by a few parameters such as in [43]), this could be readily accomplished by broadcasting a few command signals from the FC to the network. One could also increase the addressing resolution of the FC by equipping it with multiple transmit antennas and using some of the techniques described in [152]. However, depending upon the structure of the compressing (sparse) basis, this approach may require the FC to be able to address each sensor node individually which may or may not be practical in large-scale dense WSNs. We will show in Section 5.6, however, that one benefit of *compressive wireless sensing* is a straightforward treatment of this issue.

5.6 Distributed Estimation from Noisy Projections: Unknown Subspace

In Section 5.5, we proposed an efficient distributed estimation scheme that achieves the optimal centralized distortion scaling D^*_{cen} for both compressible and sparse signals under the assumption that the WSN has complete knowledge of the basis in which $\mathbf{S}$ is compressible or sparse. Generally speaking, however, even if the basis in which $\mathbf{S}$ is compressible (sparse) is known, it is quite likely that the precise ordering of its coefficients (indices of its nonzero coefficients) in that basis at each time instant k might not be known ahead of time—a scenario that we refer to as the *unknown subspace* or *adaptive subspace* case. As an example, consider the following simple case. Suppose $\mathbf{S}$ is very sparse in some basis $\boldsymbol{\Psi} = \{\boldsymbol{\psi}_i\}_{i=1}^n$ such that each temporal sample $\mathbf{s}^k$ has only one nonzero coefficient of amplitude $\sqrt{n}\,B$ corresponding to *some* element $\boldsymbol{\psi}_i$ of $\boldsymbol{\Psi}$ and i is drawn at random from the set $[1 \ldots n]$. This is an example of the case where we know the basis in which $\mathbf{S}$ is sparse but do not know the indices of its nonzero coefficients in that basis.

One naïve approach to this problem would be to use the distributed estimation scheme described in Section 5.5. However, since the network does not have a precise knowledge of the index of the true basis vector, it would need to be determined by trial and error (e.g., deterministically or randomly selecting basis vectors in some fashion). As an illustration, consider a randomized selection process: the network computes the projection of the sensor data onto $\boldsymbol{\psi}_i$ and i is selected

L times uniformly at random (without replacement) from the set $[1 \dots n]$. Ignoring the distortion due to the measurement and communication noise, the squared distortion error would be 0 at the FC if the spike in the $\boldsymbol{\Psi}$ domain corresponds to one of the uniformly picked $\boldsymbol{\psi}_i$'s and B^2 otherwise, and the probability of not finding the spike in L trials is $\prod_{i=0}^{L-1} \left(1 - \frac{1}{n-i}\right)$. If n is large enough and $L \ll n$, we can approximate the resulting distortion by $D \approx \left(1 - \frac{1}{n}\right)^L B^2 \approx e^{-L/n} B^2 \to B^2$ as $n \to \infty$, i.e., equivalent to the MSE that is achievable even without any information.

Another more general, and perhaps relevant, example is a situation in which the signal field is spatially piecewise smooth. Signals of this type do lie in a low-dimensional subspace of the wavelet domain, but precisely which subspace depends on the locations of the change points in the signal, which of course are unlikely to be known a priori. Broadly speaking, any signal that is generally spatially smooth apart from some localized sharp changes or edges will essentially lie in a low-dimensional subspace of a multiresolution basis such as wavelets or curvelets, but the subspace will be a function of the time index k and thus, will preclude the use of methods like the one in Section 5.5 that require prior specification of the basis vectors to be used in the projection process. This is where the *universality* of compressive wireless sensing (CWS) scheme, presented in this section, comes into play. As we shall see, CWS provides us with a consistent estimation scheme ($D \searrow 0$ as node density increases), even if little or no prior knowledge about the networked data is assumed, while P_{tot} and L grow at most sub-linearly with the number of nodes in the network.

5.6.1 Compressive Wireless Sensing

Recall that if $\upsilon^k = \boldsymbol{\varphi}^{\mathsf{T}} \mathbf{s}^k = \sum_{j=1}^n \varphi_j s_j^k$ is the projection of $\mathbf{s}^k$ onto a vector $\boldsymbol{\varphi} \in \mathbb{R}^n$ then, using the communication architecture described in Section 5.4 and consuming only $O(1)$ amount of power, the FC can obtain an estimate of υ^k in one channel use that is given by

$$\widehat{\upsilon}^k = \upsilon^k + \boldsymbol{\varphi}^{\mathsf{T}} \mathbf{w}^k + \widetilde{z}^k \tag{5.67}$$

where $\widetilde{z}^k \sim \mathcal{N}(0, \sigma_z^2/\rho)$ is the scaled MAC AWGN [cf. (5.27)]. The basic idea behind CWS is that instead of projecting the sensor network data onto a subset of a deterministic basis of $\mathbb{R}^n$, the FC tries to reconstruct $\mathbf{s}^k$ from *random projections* of the sensor network data. Specifically, let

$\{\boldsymbol{\phi}_i \in \mathbb{R}^n\}_{i=1}^n$ be an i.i.d. sequence of (normalized) Rademacher random vectors, i.e., $\{\phi_{ij}\}_{j=1}^n = \pm 1/\sqrt{n}$, each with probability $1/2$, and the FC tries to reconstruct $\mathbf{s}^k$ by projecting $\mathbf{x}^k$ onto L of these random vectors.[9] Because the entries of each projection vector $\boldsymbol{\phi}_i$ are generated at random, observations of this form are called random projections of the signal.

Remark 5.14 An important consequence of using Rademacher random vectors for projection purposes is that each sensor can locally draw the elements of the projection vectors $\{\boldsymbol{\phi}_i\}_{i=1}^L$ in an efficient manner by simply using its network address as the seed of a pseudo-random number generator (see, e.g., [153]). Moreover, given these seed values and the number of nodes in the network, the FC can also easily reconstruct the vectors $\{\boldsymbol{\phi}_i\}_{i=1}^L$. Therefore, in addition to being a universal estimation scheme, CWS has an added advantage that no extended information concerning the projection vectors needs to be communicated to (or stored inside) the sensor nodes (cf. Section 5.5.3).

After employing L random (Rademacher) projections, the corresponding projection estimates at the FC are simply given by

$$\widehat{v}_i^k \;=\; \boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{s}^k + \boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k + \widetilde{z}_i^k,\; i = 1, \ldots, L \tag{5.68}$$

where $\mathbf{w}^k \;=\; \begin{bmatrix} w_1^k & \ldots & w_n^k \end{bmatrix}^{\mathsf{T}}$, and $\{w_j^k\}_{j=1}^n$ and $\{\widetilde{z}_i^k\}_{i=1}^L$ are i.i.d. zero-mean Gaussian random variables, independent of each other and $\{\phi_{ij}\}$, with variances σ_w^2 and σ_z^2/ρ, respectively. The reconstruction process can be described as follows—let $\mathcal{S}_q$ denote a countable collection of candidate reconstruction vectors such that

$$\mathcal{S}_q \;\subset\; \{\mathbf{y} \in \mathbb{R}^n : |y_j| \leq B,\; j = 1, \ldots, n\} \tag{5.69}$$

and define a CWS estimate $\widehat{\mathbf{s}}^k$ as

$$\widehat{\mathbf{s}}^k \;=\; \underset{\mathbf{s}\in\mathcal{S}_q}{\arg\min} \left\{ \widehat{R}(\mathbf{s}) + \frac{c(\mathbf{s})\log(2)}{L\epsilon} \right\}. \tag{5.70}$$

[9]The L Rademacher vectors are to be generated independently at each time instant k. However, we omit the superscript corresponding to the time index to simplify the notation.

The first term in the objective function is the empirical risk, defined as

$$\widehat{R}(\mathbf{s}) \stackrel{def}{=} \frac{1}{L}\sum_{i=1}^{L}\left(\widehat{v}_i^k - \boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{s}\right)^2 \tag{5.71}$$

which measures the average (Euclidean) distance between the observations $\{\widehat{v}_i^k\}_{i=1}^L$ and the projections of a given candidate vector $\mathbf{s}$ onto the corresponding Rademacher vectors $\{\boldsymbol{\phi}_i\}_{i=1}^L$. The quantity $c(\mathbf{s})$ in the second term is a non-negative number assigned to each candidate vector in $\mathcal{S}_q$ such that $\sum_{\mathbf{s}\in\mathcal{S}_q} 2^{-c(\mathbf{s})} \leq 1$, and is designed to penalize candidate vectors proportional to their complexity [see (5.76)]. Finally, $\epsilon > 0$ is a constant (independent of L and n) that controls the relative contribution of the complexity term to the objective function. In the context of [67, Theorem 1], $\sigma_w^2 = 0$ and so ϵ depends only the sample bound B and the noise variance σ_z^2/ρ.

In order to apply the results of [67] to the observation model (5.68), the effect of the *projected noise* terms $\{\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k\}_{i=1}^L$ needs to be determined. First, suppose that the projection vectors $\{\boldsymbol{\phi}_i\}_{i=1}^L$ were mutually orthogonal. In that case, it is easy to see that the projected noises are equivalent (in distribution) to i.i.d. zero-mean Gaussian noises with variance σ_w^2. In addition, note that $\{\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k\}$ and $\{\phi_{ij}\}$ are independent. To see this, notice that for any fixed vector $\mathbf{g} \in \mathbb{R}^n$, the joint characteristic function of $\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k$ and $\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{g}$ can be factored into the product of the individual characteristic functions, i.e.,

$$\mathbb{E}\left[e^{j\nu_1\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k + j\nu_2\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{g}}\right] \;=\; \mathbb{E}\left[e^{j\nu_1\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k}\right]\cdot\mathbb{E}\left[e^{j\nu_2\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{g}}\right], \tag{5.72}$$

and taking $\mathbf{g}$ to be a vector that has one at location 'j' and zero at all other locations establishes the independence of $\{\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k\}$ and $\{\phi_{ij}\}$. In this case, if we pick

$$\rho \;=\; \frac{P}{{d_u}^{\zeta}\left(B^2+\sigma_w^2\right)} \tag{5.73}$$

then the observations in (5.68) would be equivalent (in distribution) to observations of the form

$$\widehat{v}_i^k \;=\; \boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{s}^k + \eta_i^k, \quad i = 1,\ldots,L, \tag{5.74}$$

where $\{\eta_i^k\}_{i=1}^L$ are i.i.d. zero-mean Gaussian random variables—*independent* of $\{\phi_{ij}\}$—with variance $\sigma^2 = \sigma_w^2 + \sigma_z^2\,{d_u}^{\zeta}\left(B^2+\sigma_w^2\right)/P$, and the results of [67, Theorem 1] can be applied directly.

On the other hand, our model only assumes that the vectors $\{\boldsymbol{\phi}_i\}_{i=1}^{L}$ are mutually orthogonal in expectation; hence, the projected noise is colored—if $\boldsymbol{\Phi}$ is the $L \times n$ matrix whose rows are $\{\boldsymbol{\phi}_i^{\mathsf{T}}\}_{i=1}^{L}$ then, given $\boldsymbol{\Phi}$, the projected noise vector $\boldsymbol{\Phi}\mathbf{w}^k$ is a zero-mean Gaussian vector with covariance matrix $\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathsf{T}}\sigma_w^2$. Without loss of generality, however, we can assume that the projected noise $\{\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k\}_{i=1}^{L}$ behaves approximately like white Gaussian noise and consequently, use the observation model of (5.74) for further analysis. This approximation is motivated by the asymptotic results presented in [8, Section IV-B] which show that the extreme eigenvalues of $\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathsf{T}}$ are almost surely (a.s.) contained in the interval $\left[(1-\sqrt{c})^2, (1+\sqrt{c})^2\right]$ in the limit as $L, n \to \infty$ with $L/n \to c$. Since we assume L grows sublinearly with n so $L/n \to 0$ in our case and consequently, all the eigenvalues of $\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathsf{T}}$ tend to 1 a.s. In other words, $\{\boldsymbol{\phi}_i\}_{i=1}^{L}$ become mutually orthogonal asymptotically, and the degree of coloring becomes negligible for large values of n (this approximation is also shown to work well in practice—see Section 5.8).

Explicit bounds on the reconstruction error using the CWS estimate of (5.70) can then be obtained by first assuming that we can find a basis $\boldsymbol{\Psi}$ at the FC in which the signal $\mathbf{S}$ is α-compressible and then, using this compressing basis in the reconstruction process by defining $\mathcal{S}_q$ and $c(\mathbf{s})$ in terms of $\boldsymbol{\Psi}$. Specifically, let

$$\boldsymbol{\Theta}_q \stackrel{def}{=} \left\{\boldsymbol{\theta} \in \mathbb{R}^n : |(\boldsymbol{\Psi}\boldsymbol{\theta})_j| \leq B,\ \theta_j \text{ uniformly quantized to } n^q \text{ levels},\ j = 1, \ldots, n\right\} \tag{5.75}$$

be a set of quantized candidate solutions in the transform domain $\boldsymbol{\Psi}$, so that $\mathcal{S}_q = \{\mathbf{s} \in \mathbb{R}^n : \mathbf{s} = \boldsymbol{\Psi}\boldsymbol{\theta},\ \boldsymbol{\theta} \in \boldsymbol{\Theta}_q\}$. Furthermore, let the penalty term $c(\mathbf{s}) = c(\boldsymbol{\theta})$ be

$$c(\boldsymbol{\theta}) \stackrel{def}{=} (1+q)\log(n)\|\boldsymbol{\theta}\|_0. \tag{5.76}$$

Then, the optimization problem (5.70) essentially reduces to solving the problem

$$\widehat{\boldsymbol{\theta}}^k = \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}_q}{\arg\min}\left\{ \left\| \widehat{\boldsymbol{\upsilon}}_L^k - \boldsymbol{\Phi}_L^{\mathsf{T}}\boldsymbol{\Psi}\boldsymbol{\theta} \right\|_2^2 + \frac{(1+q)\log(2)\log(n)}{\epsilon}\|\boldsymbol{\theta}\|_0 \right\} \tag{5.77}$$

where $\widehat{\boldsymbol{\upsilon}}_L^k \stackrel{def}{=} \begin{bmatrix}\widehat{\upsilon}_1^k & \ldots & \widehat{\upsilon}_L^k\end{bmatrix}^{\mathsf{T}}$, $\boldsymbol{\Phi}_L$ is the $n \times L$ matrix of Rademacher projection vectors $\{\boldsymbol{\phi}_i\}_{i=1}^{L}$, and $\widehat{\boldsymbol{\theta}}^k$ is the estimate of the representation of $\mathbf{s}^k$ in the compressing basis $\boldsymbol{\Psi}$, i.e., $\widehat{\mathbf{s}}^k \stackrel{def}{=} \boldsymbol{\Psi}\widehat{\boldsymbol{\theta}}^k$. As

shown in [67], for an α-compressible $\mathbf{S}$, such an estimate would satisfy[10]

$$D \preceq \left(\frac{L}{\log(n)}\right)^{-2\alpha/(2\alpha+1)} \tag{5.78}$$

while, for an M-sparse signal, this would result in

$$D \preceq \left(\frac{L}{M\log(n)}\right)^{-1}. \tag{5.79}$$

5.6.2 Power-Distortion-Latency Scaling Laws

Recall that in order to achieve the distortion scaling of (5.78) and (5.79), the network had to employ L network-to-FC MAC uses per source observation, each one corresponding to a projection of $\mathbf{s}^k$ onto a random Rademacher vector. And while the projection vectors in this case happened to be random as opposed to the analysis carried out in Section 5.4, it is a simple exercise to show that with the scaling factor ρ as given in (5.73), each projection of the noisy networked data onto a (random) Rademacher vector still consumes only $O(1)$ amount of power. Therefore, power-distortion-latency scaling relationship of the CWS scheme for the case when $\mathbf{S}$ is α-compressible can be given by

$$D \sim \left(\frac{P_{\text{tot}}}{\log(n)}\right)^{-2\alpha/(2\alpha+1)} \sim \left(\frac{L}{\log(n)}\right)^{-2\alpha/(2\alpha+1)} \tag{5.80}$$

while for an M-sparse signal with M and L scaling as $M \asymp n^{\mu}$, $0 \leq \mu < 1$, and $L \asymp n^{\beta}$, $0 < \beta < 1$, it can be given by

$$D \sim \log(n)\, {P_{\text{tot}}}^{-1+\frac{\mu}{\beta}} \sim \log(n)\, L^{-1+\frac{\mu}{\beta}}. \tag{5.81}$$

Comparison of these power-distortion-latency relationships with the ones achievable in Section 5.5 yields an interesting insight: regardless of the compressibility (sparsity) of $\mathbf{S}$, if there is enough prior knowledge about the underlying physical phenomenon, the distortion achievable under CWS would always be greater than the one achievable in the known subspace case, when using

[10] The stated results hold for all $q \geq 1$; the explicit dependencies of the leading constants on the quantization parameter q are derived in [67].

the same amount of power and latency *and* identical reconstruction basis. As an example, whereas one can obtain a distortion scaling of $D \asymp n^{-2\alpha/(2\alpha+1)}$ by employing $(P_{\text{tot}} \asymp) L \asymp n^{1/(2\alpha+1)}$ projections for the estimation of an α-compressible signal in the known subspace case, the distortion scaling in the unknown subspace case, when using the same number of projections can only be given by $D \preceq n^{-2\alpha/(2\alpha+1)^2}$—a significantly slower decay [cf. (5.45), (5.80)].

On the other hand, by virtue of a toy example, we have already seen at the start of this section what could happen to the distortion scaling in the known subspace case if the known subspace assumption is false and that is where the universality of CWS comes into play: given sufficient prior knowledge about the underlying signal field, CWS can be far from optimal but under circumstances where there is *little* or *no* knowledge available about the signal field, CWS should be the scheme of choice for estimating compressible or sparse networked data.

5.7 Impact of Fading and Imperfect Phase Synchronization

The joint source-channel communication architecture presented in Section 5.4 for computing distributed projections in WSNs (and extended to estimation of compressible/sparse signals in Sections 5.5 and 5.6) is analyzed under the assumptions that the network is fully synchronized and transmissions from the sensor nodes do not undergo fading. This assumption may not hold in practice for sensor network deployments in scattering environments and due to drifts in phases of sensor oscillators. Therefore, we relax these assumptions in this section and study the impact of fading and imperfect phase synchronization on the previously obtained scaling laws. In particular, we establish that (i) the power-distortion-latency laws of Sections 5.4 and 5.5 continue to hold as long as the random channel gains of received signals at the FC (due to fading and phase synchronization errors) have a nonzero mean, and (ii) the CWS scaling laws continue to hold as long as the mean of these random channel gains is not too small.

5.7.1 Distributed Projections in Wireless Sensor Networks

We begin by analyzing the impact of fading and imperfect phase synchronization on the power-distortion-latency scaling law of the proposed communication scheme for computing distributed

projections (cf. Theorem 5.9). This is accomplished by assuming that the communication scheme is still described by the encoders in (5.25) but, as a result of narrowband fading and phase synchronization errors, each sensor's transmitted signal is received at the FC after multiplication by a random channel gain $\gamma_j^k \stackrel{def}{=} g_j^k \cos(\Delta_j^k)$, $j = 1, 2, \ldots, n$, where the random variables $\{g_j^k\}$ and $\{\Delta_j^k\}$ are i.i.d. (across sensors), which are also assumed to be independent of each other [21,142].[11] Note that $\{g_j^k\}$ are non-negative valued random variables—typically modeled as Rayleigh, Rician or log-normal distributed—and correspond to random fading envelopes of received signals at the FC, whereas $\{\Delta_j^k\}$ model the combined effect of random phase-shifts due to multipath scattering and imperfect phase synchronization between the sensors and the FC. We assume that the precise values and distributions of these random variables are not available to the sensors or the FC, but their means are known at the FC.

Consequently, as a result of fading and imperfect phase synchronization, the FC receives

$$\begin{aligned} r^k &= \sum_{j=1}^{n} \sqrt{h_j}\, \gamma_j^k y_j^k + z^k \;=\; \sqrt{\rho} \sum_{j=1}^{n} \varphi_j \gamma_j^k x_j^k + z^k \\ &= \sqrt{\rho}\, \boldsymbol{\varphi}^{\mathsf{T}}\big(\boldsymbol{\gamma}^k \odot \mathbf{s}^k\big) + \Big(\sqrt{\rho}\, \boldsymbol{\varphi}^{\mathsf{T}}\big(\boldsymbol{\gamma}^k \odot \mathbf{w}^k\big) + z^k\Big) \end{aligned} \tag{5.82}$$

where $\boldsymbol{\gamma}^k \stackrel{def}{=} \begin{bmatrix} \gamma_1^k & \ldots & \gamma_n^k \end{bmatrix}^{\mathsf{T}}$, $\odot$ represents a Hadamard product (element-wise multiplication), and ρ is still given by the expression in (5.33). Clearly, this coincides with the received signal in (5.26) if and only if $\boldsymbol{\gamma}^k = \mathbf{1}_n$, where $\mathbf{1}_n$ denotes an n-dimensional vector of all ones. However, by a slight modification of the decoder in (5.27), it can be shown that the scaling law established in Theorem 5.10 is still achievable as long as the distribution of random channel gains is such that the network remains at least "barely synchronized" in the sense that $\mathbb{E}[\gamma_j^k] \stackrel{def}{=} \overline{\gamma} \neq 0$. This condition would be satisfied, e.g., if $\Delta_j^k \sim \mathfrak{unif}[-\pi + \epsilon,\, \pi - \epsilon]$ for any $\epsilon > 0$. The modified decoder G in this scenario is given by

$$\begin{aligned} \widehat{v}^k &= G\left(r^k\right) \;=\; \frac{r^k}{\overline{\gamma}\sqrt{\rho}} \\ &= \boldsymbol{\varphi}^{\mathsf{T}}\big(\widetilde{\boldsymbol{\gamma}}^k \odot \mathbf{s}^k\big) + \boldsymbol{\varphi}^{\mathsf{T}}\big(\widetilde{\boldsymbol{\gamma}}^k \odot \mathbf{w}^k\big) + \frac{z^k}{\overline{\gamma}\sqrt{\rho}} \end{aligned} \tag{5.83}$$

[11] Recall that we are doing real-signaling; the random channel gains are, therefore, given by $\gamma_j^k = g_j^k \cos(\Delta_j^k)$ instead of $\gamma_j^k = g_j^k e^{\boldsymbol{\jmath}\Delta_j^k}$.

where $\widetilde{\boldsymbol{\gamma}}^k \stackrel{def}{=} \left[\frac{\gamma_1^k}{\overline{\gamma}} \quad \ldots \quad \frac{\gamma_n^k}{\overline{\gamma}}\right]^{\mathsf{T}}$, and the achievable distortion using this modified decoder can be characterized by the following result.

Theorem 5.15 Let $\boldsymbol{\varphi} \in \mathbb{R}^n$ and let $\upsilon^k = \boldsymbol{\varphi}^{\mathsf{T}}\mathbf{s}^k$. Suppose that the random channel gains $\{\gamma_j^k\}$ due to fading and imperfect phase synchronization are i.i.d across sensors and have a nonzero mean $(\mathbb{E}[\gamma_j^k] = \overline{\gamma} \neq 0)$.[12] Then, given the sensor network model of Section 5.2, the joint source-channel communication scheme described by the encoders in (5.25) and the modified decoder in (5.83) can achieve the following end-to-end distortion by employing only one channel use per source observation

$$D_\upsilon \;\leq\; \left(\frac{\sigma_w^2\,\overline{\overline{\gamma}} + B^2\big(\overline{\overline{\gamma}} - \overline{\gamma}^2\big)}{\overline{\gamma}^2}\right)\|\boldsymbol{\varphi}\|_2^2 \;+\; \left(\frac{\sigma_z^2\, {d_u}^{\zeta}\,(B^2+\sigma_w^2)}{\lambda\overline{\gamma}^2 P}\right)\|\boldsymbol{\varphi}\|_2^2 \tag{5.84}$$

where $\widehat{\upsilon}^k$ is the estimate of υ^k at the FC and $\overline{\overline{\gamma}} \stackrel{def}{=} \mathbb{E}[|\gamma_j^k|^2] \leq 1$ (mainly because of the law of conservation of energy).

Proof: To establish this theorem, note that (5.83) implies that $\forall\, k \in \mathbb{N}$

$$\begin{aligned}
\mathbb{E}\left[\left|\upsilon^k - \widehat{\upsilon}^k\right|^2\right] \;&=\; \mathbb{E}\left[\left|\Big(\boldsymbol{\varphi}^{\mathsf{T}}\big(\widetilde{\boldsymbol{\gamma}}^k \odot \mathbf{s}^k\big) - \boldsymbol{\varphi}^{\mathsf{T}}\mathbf{s}^k\Big) + \boldsymbol{\varphi}^{\mathsf{T}}\big(\widetilde{\boldsymbol{\gamma}}^k \odot \mathbf{w}^k\big) + \frac{z^k}{\overline{\gamma}\sqrt{\rho}}\right|^2\right] \\
&\stackrel{(a)}{=}\; \mathbb{E}\left[\left|\boldsymbol{\varphi}^{\mathsf{T}}\Big(\big(\widetilde{\boldsymbol{\gamma}}^k - \mathbf{1}_n\big) \odot \mathbf{s}^k\Big)\right|^2\right] \;+\; \frac{\sigma_w^2\,\overline{\overline{\gamma}}}{\overline{\gamma}^2}\|\boldsymbol{\varphi}\|_2^2 \;+\; \frac{\sigma_z^2}{\overline{\gamma}^2\rho} \\
&=\; \sum_{j=1}^{n}\big(\varphi_j s_j^k\big)^2\mathbb{E}\left[\Big(\frac{\gamma_j^k}{\overline{\gamma}} - 1\Big)^2\right] \;+\; \frac{\sigma_w^2\,\overline{\overline{\gamma}}}{\overline{\gamma}^2}\|\boldsymbol{\varphi}\|_2^2 \;+\; \frac{\sigma_z^2}{\overline{\gamma}^2\rho} \\
&\stackrel{(b)}{\leq}\; \frac{B^2\big(\overline{\overline{\gamma}} - \overline{\gamma}^2\big)}{\overline{\gamma}^2}\|\boldsymbol{\varphi}\|_2^2 \;+\; \frac{\sigma_w^2\,\overline{\overline{\gamma}}}{\overline{\gamma}^2}\|\boldsymbol{\varphi}\|_2^2 \;+\; \frac{\sigma_z^2}{\overline{\gamma}^2\rho}
\end{aligned} \tag{5.85}$$

where (a) essentially follows from the fact that the random channel gain vector $\boldsymbol{\gamma}^k$ is independent of the zero-mean measurement noise vector $\mathbf{w}^k$ and zero-mean communication noise z^k, and (b) primarily follows from the fact that $|s_j^k| \leq B$. Finally, to complete the proof of the theorem, we substitute in (5.85) the value of ρ from (5.33) and take the limit in k to obtain (5.84). ∎

[12] At the expense of some extra notation, the scaling laws stated in this section can be obtained even when $\{\gamma_j^k\}$ are not identically distributed (as long as they are independent across sensors and have nonzero means). For the sake of this exposition, however, and because it suffices to illustrate the principles, we focus only on the i.i.d. case.

Remark 5.16 A corresponding lower bound on the projection coefficient distortion D_υ under the modified decoder of (5.83) is given by (5.28). This follows trivially from (a) in (5.85) and the fact that $\overline{\gamma}^2 \leq \overline{\overline{\gamma}} \leq 1$.

Remark 5.17 Since the structure of source-channel encoders $(F_1, \ldots, F_n)$ remains unchanged under fading and imperfect phase synchronization, the total network power consumption associated with achieving the distortion in (5.84) is still given by (5.34).

Notice that even under the effects of fading and imperfect synchronization, the projection coefficient distortion D_υ achieved by the proposed joint source-channel communication architecture (using the modified decoder of (5.83)) is given by a sum of two separate terms, the first of which scales like $\|\boldsymbol{\varphi}\|_2^2$, while the second term that is primarily due to the noisy communication channel scales like $\|\boldsymbol{\varphi}\|_2^2/\lambda$ (cf. (5.84), Remark 5.16). Comparing this observation with the scaling law established in Section 5.4 shows that Theorem 5.15 describes the same distortion scaling behavior as Theorem 5.9, with the only difference being that the scaling constants are now different (they depend upon the second-order statistics of channel gains). In particular, $L_\upsilon = 1$ and $P_{\text{tot},\upsilon} = O(1)$ is still sufficient to ensure that $D_\upsilon \asymp \|\boldsymbol{\varphi}\|_2^2 \asymp D_\upsilon^*$, as long as $\overline{\overline{\gamma}} \neq 0$ (cf. Corollary 5.10).

5.7.2 Distributed Estimation from Noisy Projections: Known Subspace

Similar to the case of estimation of a single projection coefficient under the effects of fading and imperfect phase synchronization, it is a simple exercise to show that by using the joint source-channel communication scheme described by the encoders in (5.36) and under a slightly modified decoder $\boldsymbol{G}$ given by

$$\widehat{\mathbf{s}}^k \;=\; \boldsymbol{G}\left(\mathbf{r}^k\right) \;\stackrel{def}{=}\; \boldsymbol{\Psi}_L^k\left(\frac{\mathbf{r}^k}{\overline{\gamma}\sqrt{\rho}}\right) \tag{5.86}$$

the end-to-end distortion of an α-compressible signal at the FC in the presence of fading and phase synchronization errors can be upper bounded by

$$D \;\leq\; C_o\, L^{-2\alpha} \;+\; \left(\frac{L}{n}\right)\left(\frac{\sigma_w^2\,\overline{\overline{\gamma}} + B^2\left(\overline{\overline{\gamma}} - \overline{\gamma}^2\right)}{\overline{\gamma}^2}\right) \;+\; \left(\frac{L}{n}\right)\left(\frac{\sigma_z^2\, d_u{}^{\zeta}\left(B^2 + \sigma_w^2\right)}{\lambda\overline{\gamma}^2 P}\right) \tag{5.87}$$

and lower bounded by the expression in the lower bound of (5.39), as long as $\{\gamma_j^k\}$ are i.i.d. across sensors and $\overline{\gamma} \neq 0$. Ignoring constants, this implies that the resulting distortion of an α-compressible signal in this scenario still scales as

$$\left(\frac{L}{n}\right) + \left(\frac{L}{\lambda n}\right) \preceq D \preceq L^{-2\alpha} + \left(\frac{L}{n}\right) + \left(\frac{L}{\lambda n}\right) \tag{5.88}$$

i.e., has the same scaling behavior as that of D in (5.39). Similarly, it can be shown that using the modified decoder of (5.86) (with L replaced by M), the end-to-end distortion of an M-sparse signal in this scenario would scale as [cf. (5.59)]

$$D \asymp \left(\frac{M}{n}\right) + \left(\frac{M}{\lambda n}\right). \tag{5.89}$$

Moreover, given that the structure of source-channel encoders $(\boldsymbol{F}_1, \ldots, \boldsymbol{F}_n)$ remains unchanged under fading and imperfect phase synchronization, the total network power consumption per source observation associated with achieving these distortion scalings for compressible and sparse signals would still be given by (5.40) and (5.60), respectively. Comparison of these results with the ones obtained in Section 5.5 shows that the previously established power-distortion-latency scaling laws for estimation of compressible and sparse signals in the known subspace case continue to hold under the effects of fading and imperfect phase synchronization, provided $\{\gamma_j^k\}$ have a nonzero mean and the FC uses the modified decoder of (5.86).

Remark 5.18 Note that these results are similar in spirit to some of the earlier results obtained in the context of joint source-channel communication for distributed estimation of sources—see, e.g., [42, 154–156]. In particular, those results also indicate that fading (and/or imperfect synchronization) tends to have no effect on the distortion scaling as long as the random channel gains have nonzero means.

5.7.3 Compressive Wireless Sensing

In the presence of phase synchronization errors only (no fading, i.e., $\gamma_j^k = \cos(\Delta_j^k)$ only), CWS observations are given by $(\forall\, i = 1, \ldots, L)$

$$\widehat{v}_i^k \;=\; \boldsymbol{\phi}_i^{\mathsf{T}}\big(\widetilde{\boldsymbol{\gamma}}^k \odot \mathbf{s}^k\big) + \boldsymbol{\phi}_i^{\mathsf{T}}\big(\widetilde{\boldsymbol{\gamma}}^k \odot \mathbf{w}^k\big) + \frac{z_i^k}{\overline{\gamma}\sqrt{\rho}} \tag{5.90}$$

where $\widetilde{\boldsymbol{\gamma}}^k = \left[\frac{\gamma_1^k}{\overline{\gamma}} \quad \ldots \quad \frac{\gamma_n^k}{\overline{\gamma}}\right]^{\mathsf{T}}$ [see (5.68)]. Defining the vector $\widetilde{\boldsymbol{\gamma}}^k$ as $\widetilde{\boldsymbol{\gamma}}^k \stackrel{def}{=} \mathbf{1}_n + \boldsymbol{\delta}^k$ and substituting into the above gives

$$\widehat{v}_i^k \;=\; \boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{s}^k + \boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k + \boldsymbol{\phi}_i^{\mathsf{T}}\big(\boldsymbol{\delta}^k \odot (\mathbf{s}^k + \mathbf{w}^k)\big) + \frac{z_i^k}{\overline{\gamma}\sqrt{\rho}} \tag{5.91}$$

where $\boldsymbol{\delta}^k$ is a zero-mean random vector with i.i.d. entries given by $\delta_j^k = \widetilde{\gamma}_j^k - 1$, $j = 1, \ldots, n$. Comparing this with (5.68), we see that the net effect of phase synchronization errors is the introduction of a new noise-like term of the form $\boldsymbol{\phi}_i^{\mathsf{T}}\big(\boldsymbol{\delta}^k \odot (\mathbf{s}^k + \mathbf{w}^k)\big)$. Foregoing a rigorous theoretical analysis of the effects of this contribution, we instead assume (by the Central Limit Theorem) that it is approximately Gaussian distributed, in which case it can be treated like the projected noise $\{\boldsymbol{\phi}_i^{\mathsf{T}}\mathbf{w}^k\}$, as in Section 5.6.1. Further, assuming that $\Delta_j^k \stackrel{\text{i.i.d.}}{\sim} \mathfrak{unif}[-b,\, b]$ and that b is small, we can use a one-term Taylor series approximation of the variance of the new zero-mean noise contribution. The result is that each CWS observation is again given by (5.74) but the equivalent noise variance is given by $\sigma^2 = \sigma_w^2 + \sigma_z^2\, d_u{}^{\zeta}\, (B^2 + \sigma_w^2)\,/(\overline{\gamma}^2 P) + (B^2 + \sigma_w^2) b^4/45$, the last term in the expression being the contribution of the new phase synchronization error term.

More generally, if we also define the fading envelope of each sensor's transmissions as $g_j^k \stackrel{def}{=} 1 + \epsilon_j^k$, then the overall random channel gain of each received signal becomes $\gamma_j^k = g_j^k \cos(\Delta_j^k) = (1 + \epsilon_j^k)(1 + \delta_j^k) = 1 + \epsilon_j^k + \delta_j^k + \epsilon_j^k \delta_j^k$. The net result of this is a new noise-like term of the form $\boldsymbol{\phi}_i^{\mathsf{T}}\big((\boldsymbol{\epsilon}^k + \boldsymbol{\delta}^k + \boldsymbol{\epsilon}^k \boldsymbol{\delta}^k) \odot (\mathbf{s}^k + \mathbf{w}^k)\big)$. With appropriate modeling of the ϵ_j^k terms, the additional variance due to this contribution can also be computed and the optimization problem in (5.77) can be updated accordingly. This approach was used in the simulations and appears to work well in practice for a range of phase synchronization errors, with or without mild fading, as can be seen from Figure 5.10.

5.8 Simulation Results

In this section, we present a few simulation results to numerically demonstrate some of the power-distortion-latency relationships of our scheme under both known and unknown/adaptive subspace assumptions. All signals discussed in this section are contaminated with zero-mean additive white Gaussian measurement noise of variance $\sigma_w^2 = 1$, i.e., the baseline MSE of all signals

is taken to be 1. Moreover, the measurement SNR of all signals, defined to be $\left(\|\mathbf{s}^k\|_2^2/n\right)/\sigma_w^2$, is given by $\text{SNR}_{\text{meas}} = 20$ dB, and the received communication SNR for each projection, defined to be ρ/σ_z^2, is given by $\text{SNR}_{\text{comm}} = 0$ dB (unless otherwise stated).

The first simulation result, corresponding to Figure 5.6, illustrates the distortion scaling D of a spatially piecewise smooth signal field with the number of projections L using both CWS and known subspace case reconstructions, where the signal field is sampled by $n = 8192$ sensor nodes in a noisy manner. Such signals tend to be compressible in the Haar domain with $\alpha = 1$ and this value of α was also verified numerically. For the purposes of known subspace reconstruction, the observation vector is projected onto L Haar basis elements corresponding to L largest coefficients of the noiseless vector using the scheme described in Section 5.5, while for the case of CWS reconstruction, the observation vector is projected onto L random Rademacher vectors. The resultant reconstruction MSEs are shown in the figure using solid curves (on a log-log scale), while the dotted curve and dashed curve in the figure correspond to linear fit of CWS distortion curve and reconstruction MSE in a centralized setting (i.e., $\sigma_z^2 = 0$), respectively. Finally, the total network power consumption for both CWS and known subspace case reconstructions is given by $P_{\text{tot}} \asymp L$, owing to the fact that we have chosen $\lambda = O(1)$ in this simulation.

As predicted by the theory, distortion curve for known subspace reconstruction in Figure 5.6 hits its minimum at a point where the distortion due to the approximation error is balanced by the distortion due to the observation and communication noise, and starts to rise after $L \approx 70$ projections since each subsequent projection contributes only a small amount of signal but a larger amount of noise. Note that minimum distortion in the centralized setting is attained for $L \approx 90$ projections. This is because distortion scaling constants in the known subspace case depend upon σ_w^2 and σ_z^2/ρ [see (5.39)], while $\sigma_z^2 = 0$ in the centralized case. For the case of CWS, distortion scaling follows a slope of -1.48 that turns out to be better than the expected value of $-2\alpha/(2\alpha+1) = -2/3$ [see (5.78)]. This, however, does not contradict the results reported in Section 5.6 since we only have upper bounds for distortion scaling in the CWS case. Finally, Figure 5.7 illustrates the fact that varying the received communication SNR per projection has no

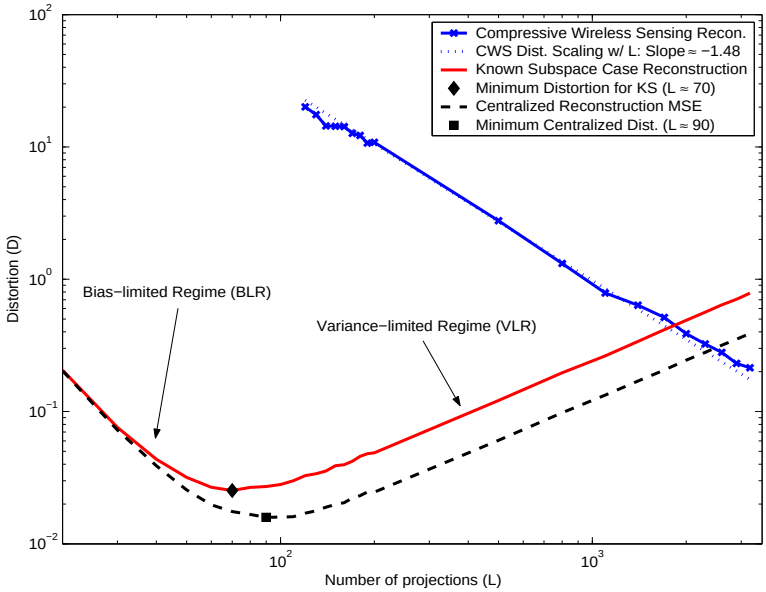

Figure 5.6 Distortion scaling of a fixed length α-compressible signal as a function of number of projections L under both known and unknown subspace assumptions (log-log scale): number of sensor nodes n = 8192; α = 1 (in Haar basis); baseline MSE (σ_w^2) = 1; measurement SNR = 20 dB; received communication SNR per projection = 0 dB.

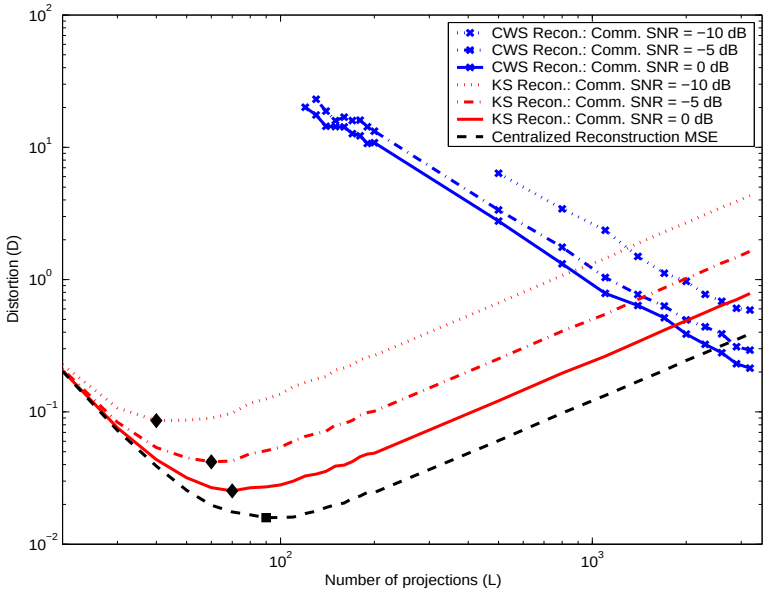

Figure 5.7 Distortion scaling of a fixed length α-compressible signal as a function of number of projections L for various values of received communication SNR per projection under both known and unknown subspace assumptions (log-log scale): number of sensor nodes n = 8192; α = 1 (in Haar basis); baseline MSE (σ_w^2) = 1; measurement SNR = 20 dB.

effect on the scaling behavior of known subspace and CWS reconstruction MSEs (except a change in the scaling constants).

The second simulation result, corresponding to Figure 5.8, illustrates the distortion scaling D of an M-sparse signal field with the number of sensor nodes using both CWS and known subspace case reconstructions, where we also scale the number of DoF in the signal as $M \asymp n^{\mu} = n^{1/3}$ in the Haar basis. For the purposes of known subspace case reconstruction, the observation vector is projected onto $L = M$ Haar basis elements corresponding to the M nonzero coefficients of the noiseless vector using the scheme described in Section 5.5, while for the case of CWS reconstruction, the observation vector is projected onto $L \asymp \log(n)\, n^{1/2} M \asymp \log(n)\, n^{5/6}$ random Rademacher vectors. The resultant reconstruction MSEs are shown in the figure using solid curves (on a log-log scale), while the dotted and dashed curves in the figure correspond to linear fit of known subspace/CWS distortion curves and reconstruction MSE in a centralized setting, respectively. Finally, the total network power consumption for CWS and known subspace case reconstructions is given by $P_{\text{tot}} \asymp \log(n)\, n^{5/6}$ and $P_{\text{tot}} \asymp n^{1/3}$, respectively, owing to the fact that we have chosen $\lambda = O(1)$ in this simulation.

As predicted by the theory, the distortion scaling curve for known subspace reconstruction in this case tends to follow a slope of $-1 + \mu \approx -0.64$ [see (5.62)]. Similarly, the distortion scaling curve for CWS reconstruction in this case can be expressed as $D \asymp M \log(n)/L \approx n^{-0.52}$—again in accordance with the theory [see (5.79)]. Finally, Figure 5.9 and Figure 5.10 illustrate the robustness of our proposed scheme to a range of phase synchronization errors, with or without fading, under both known and unknown/adaptive subspace assumptions.

5.9 Discussion

In this chapter, we have presented a distributed joint source-channel communication architecture for estimation of networked data at the FC and analyzed the corresponding power-distortion-latency relationships as a function of the number of sensor nodes. Our approach is built on distributed computation of appropriately chosen projections of the sensor data at the fusion center.

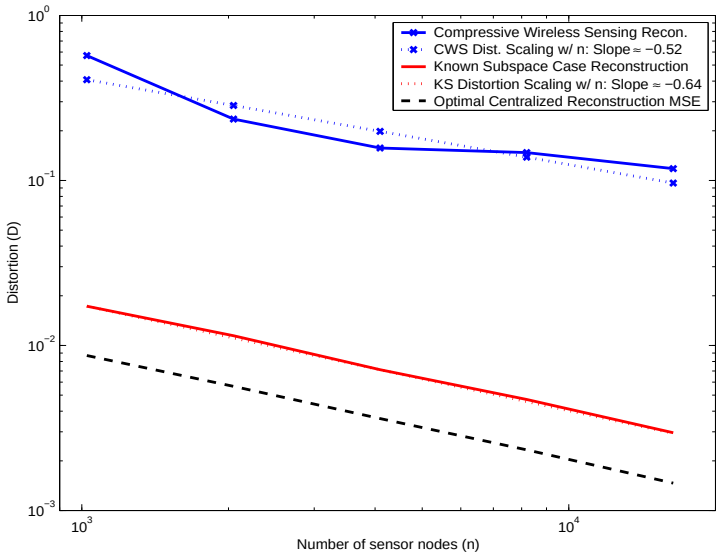

Figure 5.8 Distortion scaling of an M-sparse signal as a function of number of sensor nodes n under both known and unknown subspace assumptions (log-log scale): number of nonzero coefficients $M \asymp n^{1/3}$ (in Haar basis); baseline MSE (σ_w^2) = 1; measurement SNR = 20 dB; received communication SNR per projection = 0 dB; number of projections—Known subspace case reconstruction: $L = M \asymp n^{1/3}$, CWS reconstruction: $L \asymp \log(n)\, n^{1/2} M \asymp \log(n)\, n^{5/6}$.

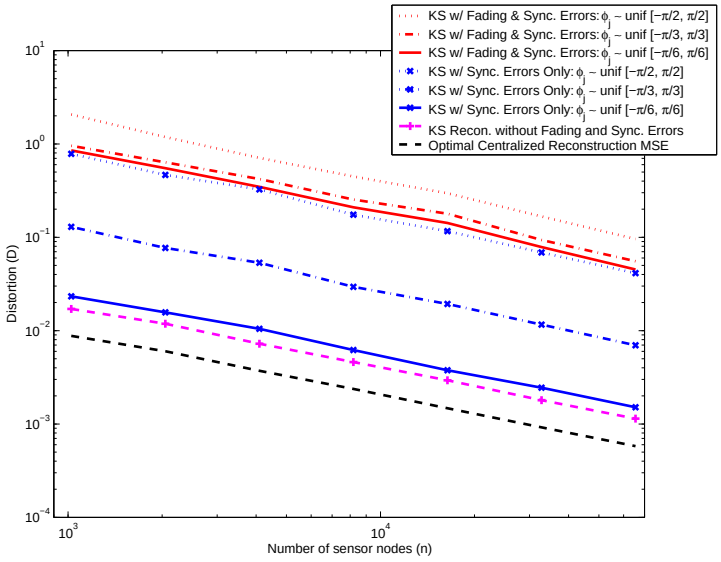

Figure 5.9 Distortion scaling of an M-sparse signal as a function of number of sensor nodes n under the effects of fading and phase synchronization errors (Known subspace case reconstruction only): number of nonzero coefficients $M \asymp n^{1/3}$ (in Haar basis); baseline MSE (σ_w^2) = 1; measurement SNR = 20 dB; received communication SNR per projection = 0 dB; fading envelope: Rayleigh distributed; number of projections $L = M \asymp n^{1/3}$.

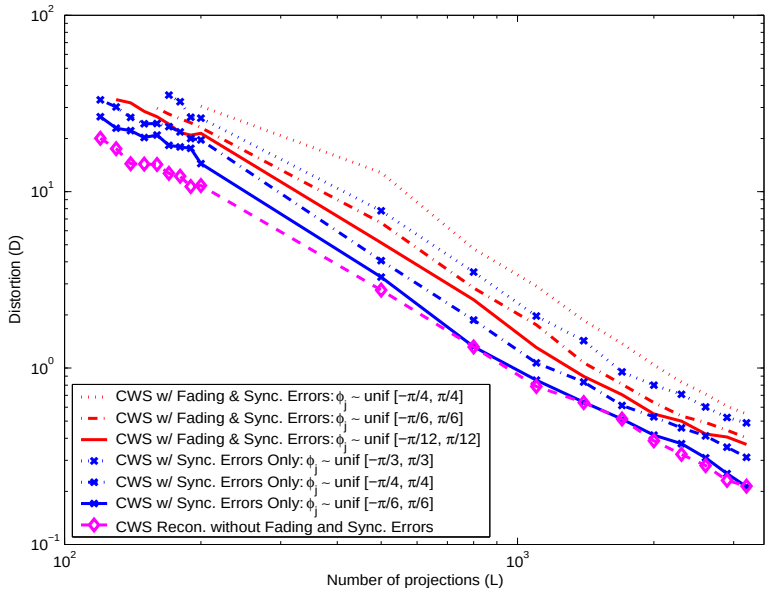

Figure 5.10 Distortion scaling of a fixed length α-compressible signal as a function of number of projections L under the effects of fading and phase synchronization errors (CWS reconstruction only): number of sensor nodes n = 8192, α = 1 (in Haar basis), baseline MSE (σ_w^2) = 1, measurement SNR = 20 dB, received communication SNR per projection = 0 dB; fading envelope: Rician distributed (K-factor of 7.5).

Phase-coherent transmissions from the sensors enable exploitation of the distributed beamforming gain for dramatic reductions in power consumption. A few distinct features of our approach are: 1) processing and communication are combined into one distributed projection operation, 2) it requires almost no in-network processing and communication, and 3) given sufficient prior knowledge about the networked data, asymptotically consistent signal estimation is possible even if the total network power consumption goes to zero.

In addition, we have also introduced and analyzed a universal estimation scheme—compressive wireless sensing (CWS)—that provides asymptotically consistent signal estimates, even if little or no prior knowledge about the networked data is assumed. Furthermore, power and latency requirements in CWS grow at most sub-linearly with the number of nodes in the network. This universality, however, comes at the cost of less favorable power-distortion-latency relationship: the absence of sufficient prior knowledge about the signal field leads to probing the entire n-dimensional space using random projections instead of focusing on the subspace of interest. However, for precisely the same reason, CWS has the ability to capture part of signal under all circumstances and does not require reprogramming of the network for different sensing scenarios—different hypotheses on the signal field structure can be tested at the fusion center via the reconstruction algorithms. Furthermore, projecting the sensor network data onto a fixed subspace may result in a distortion much greater than the one achievable by CWS if prior information about the signal field is inaccurate. Therefore, we contend that CWS should be the estimation scheme of choice in cases when either little prior knowledge about the sensed field is available or confidence level about the accuracy of the available knowledge is low.

Finally, we conclude this chapter with a brief overview of the connections between the results of this chapter and some related existing works. First, let us comment on the signal model used in this chapter. We assume that the physical phenomenon under observation is characterized by an unknown but deterministic sequence of vectors in $\mathbb{R}^n$, where each vector in the sequence is α-compressible or M-sparse in some orthonormal basis of $\mathbb{R}^n$ (see Section 5.2). Alternative assumptions that are commonly used in previous work are that the signal field is either a realization of a stationary (often bandlimited) random field with some known correlation function [157–160],

or it is fully described by a certain number of degrees of freedom (often less than n) that are random in nature [42, 161]. All of these signal models, however, express a notion of smoothness or complexity in the signal field, and the decay characteristics of the correlation function (e.g., the rate of decay) or the number of DoF in the field play a role analogous to that of α and M in this work. Essentially, the choice between a deterministic or a stochastic model is mostly a matter of taste and mathematical convenience, the latter being more prevalent when it comes to information-theoretic analysis of the problem (also, see [131] and the discussion therein). However, the deterministic formulation can be more readily generalized to include inhomogeneities, such as boundaries, in the signal field [146].

Second, it is generally recognized that the basic operations of sensing (acquisition), processing (computation), and communication in sensor networks are interdependent and, in general, they must be jointly optimized to attain optimal trade-offs between power, distortion, and latency. This joint optimization may be viewed as a form of distributed joint source-channel communication (or coding), involving both estimation (compression) and communication. Despite the need for optimized joint source-channel communication, our fundamental understanding of this complex problem is very limited, owing in part to the absence of a well-developed network information theory [162]. As a result, a majority of research efforts have tried to address either the compression or the communication aspects of the problem. Recent results on joint source-channel communication for distributed estimation or detection of sources in sensor networks [40–43, 155, 161, 163], although relatively few, are rather promising and indicate that limited node cooperation can sometimes greatly facilitate optimized source-channel communication and result in significant energy savings that more than offset the cost of cooperation. Essentially, for a given signal field, the structure of the optimal estimator dictates the structure of the corresponding communication architecture. To the best of our knowledge, the most comprehensive treatment of this problem to date (in the context of WSNs) has been carried out by Gastpar and Vetterli in [42] (see also [161]). While some of our work is inspired by and similar in spirit to [42], Gastpar and Vetterli have primarily studied the case of finite number of independent sources that is analogous to that of an M-sparse signal, albeit assuming Gaussian DoF and multiple FCs. Moreover, the number of DoF in [42] is

assumed to be fixed and does not scale with the number of nodes in the network. Our work, in contrast, not only extends the results of [42] to the case when the number of DoF of an M-sparse signal scales with n, but also applies to a broader class of signal fields and gives new insights into the power-distortion-latency relationships for both compressible and sparse signals (cf. Section 5.5). Furthermore, we also present extensions of our methodology to situations in which very limited prior information about the signal field is available.

Third, in the context of compressed sensing theory, while the idea of using random projections for the estimation of networked data has recently received some attention in the research community, the focus has primarily been on the compression or estimation aspects of the problem (see, e.g., [51, 67, 164, 165]), and our work is the first to carefully investigate the potential of using random projections from a source-channel communication perspective (cf. Section 5.6).

Finally, from an architectural and protocol viewpoint, most existing works in the area of sensor data estimation emphasize the networking aspects by focusing on multi-hop communication schemes and in-network data processing and compression (see, e.g., [146, 157, 159, 160]). This typically requires a significant level of networking infrastructure (e.g., routing algorithms), and existing works generally assume this infrastructure as given. Our approach, in contrast to these methods, eliminates the need for in-network communications and processing, and instead requires phase synchronization among nodes that imposes a relatively small burden on network resources and can be achieved, in principle, by employing distributed synchronization/beamforming schemes, such as those described in [143, 144]. Although we use the common term "sensor network" to refer to such systems, the systems we envision often act less like networks and more like coherent ensembles of sensors and thus, our proposed wireless sensing system is perhaps more accurately termed a "*sensor ensemble*" that is appropriately queried by an "*information retriever*" (FC) to acquire the desired information about the networked data.

5.10 Appendix

5.10.1 In-Network Collaboration: Power-Distortion Trade-off Revisited

Analysis in Sections 5.5.1 and 5.5.2 shows that $\lambda = O(1)$ is necessary for optimal distortion scaling in estimation of compressible and sparse signals, resulting in $P_{\text{tot}} \asymp n^{1/(2\alpha+1)}$ and $P_{\text{tot}} \asymp M$, respectively. In this appendix we partially address the question: *How good is the power-distortion scaling of our proposed scheme?* While a comparison with all conceivable distributed estimation schemes does not seem possible, we compare the performance of the proposed scheme (which does not require data exchange between nodes) to a more favorable and idealized setup where the nodes in the network can communicate their observations in an error-free manner to a designated cluster of $1 \leq \tilde{n} \leq n$ nodes. We do not make any assumptions on the nature of in-network communication and also ignore the incurred cost on energy consumption (since quantifying this cost requires making additional system-specific assumptions). Thus, our performance comparison is solely based on the power required for network-to-FC communication to achieve optimal distortion scaling. Note that $\tilde{n} = 1$ corresponds to all nodes routing their measurements to a single clusterhead in the network (using perhaps multi-hop communications), while $\tilde{n} = n$ corresponds to all the nodes in the network noiselessly sharing their data with each other (using perhaps gossip algorithms). Once $\tilde{n}$ nodes in the network have access to the entire observation vector $\mathbf{x}^k$ following each snapshot, they compute the required L projection coefficients (with respect to a given basis) and then coherently transmit the resulting projection coefficients to the FC using a sum transmit power of P per channel use. This effectively transforms the cluster-to-FC MAC into a point-to-point AWGN channel with $\tilde{n}$-fold power-pooling gain due to coherent beamforming of identical data.

We focus on estimation of α-compressible signals. Specifically, we assume that $\tilde{n}$ nodes in the designated cluster have access to identical estimates of the required L $\mathbf{\Psi}$-coefficients at the end of the data-exchange stage, i.e.,

$$\hat{\theta}_\ell^k \;=\; \boldsymbol{\psi}_\ell^{\mathsf{T}} \mathbf{x}^k \;=\; \theta_\ell^k + \boldsymbol{\psi}_\ell^{\mathsf{T}} \mathbf{w}^k, \quad \ell = 1, \ldots, L. \tag{5.92}$$

By a simple extension of the encoder structure of Section 5.5, the transmitting cluster of $\tilde{n}$ nodes coherently beamforms these L projection coefficients per snapshot in L consecutive channel uses as follows

$$\begin{aligned}\mathbf{y}_j^k \;&=\; \boldsymbol{F}_j\left(\{\hat{\theta}_\ell^k\}_{\ell=1}^L\right) \;=\; \left[y_{j1}^k \quad \ldots \quad y_{jL}^k\right]^{\mathsf{T}} \\ &=\; \frac{1}{\sqrt{h_j}}\left[\sqrt{\rho_1}\hat{\theta}_1^k \quad \ldots \quad \sqrt{\rho_L}\,\hat{\theta}_L^k\right]^{\mathsf{T}},\; j=1,\ldots,\tilde{n} \end{aligned} \tag{5.93}$$

where $\{\rho_\ell\}_{\ell=1}^L$ are scaling factors used to satisfy the sum power constraint P in each of the L channel uses. At the end of the L-th channel use, the received signal at the input of the decoder $\boldsymbol{G}$ is simply given by

$$\begin{aligned}\mathbf{r}^k \;&=\; \sum_{j=1}^{\tilde{n}}\sqrt{h_j}\,\mathbf{y}_j^k \;+\; \mathbf{z}^k \\ &=\; \tilde{n}\left[\sqrt{\rho_1}\,\hat{\theta}_1^k \quad \ldots \quad \sqrt{\rho_L}\,\hat{\theta}_L^k\right]^{\mathsf{T}} \;+\; \mathbf{z}^k \\ &=\; \tilde{n}\,\boldsymbol{\Gamma}\,\boldsymbol{\theta}_L^k \;+\; \tilde{n}\,\boldsymbol{\Gamma}\left(\boldsymbol{\Psi}_L^{k\,\mathsf{T}}\mathbf{w}^k\right) \;+\; \mathbf{z}^k \end{aligned} \tag{5.94}$$

where $\boldsymbol{\Gamma} \stackrel{def}{=} \mathbf{diag}(\sqrt{\rho_1},\ldots,\sqrt{\rho_L})$, $\boldsymbol{\theta}_L^k = (\theta_1^k,\ldots,\theta_L^k)^{\mathsf{T}}$, $\mathbf{z}^k \sim \mathcal{N}(\mathbf{0}_L, \sigma_z^2\mathbf{I}_L)$ is an AWGN vector, and $\tilde{n}$ is the power-pooling gain of identical coherent transmissions from $\tilde{n}$ nodes. An estimate of the noiseless networked data can be formed at the FC as

$$\begin{aligned}\widehat{\mathbf{s}}^k \;&=\; \boldsymbol{G}\left(\mathbf{r}^k\right) \;=\; \boldsymbol{\Psi}_L^k\left(\frac{\boldsymbol{\Gamma}^{-1}\mathbf{r}^k}{\tilde{n}}\right) \\ &=\; \mathbf{s}^{k,(L)} \;+\; \boldsymbol{\Psi}_L^k\left(\boldsymbol{\Psi}_L^{k\,\mathsf{T}}\mathbf{w}^k\right) \;+\; \boldsymbol{\Psi}_L^k\left(\frac{\boldsymbol{\Gamma}^{-1}\mathbf{z}^k}{\tilde{n}}\right). \end{aligned} \tag{5.95}$$

As for fixing the values of $\{\rho_\ell\}$, note that (5.93) implies that $\forall\ \ell = 1,\ldots,L$,

$$\begin{aligned}\sum_{j=1}^{\tilde{n}}\mathbb{E}\left[\left|y_{j\ell}^k\right|^2\right] \;&=\; \sum_{j=1}^{\tilde{n}}\mathbb{E}\left[\frac{\rho_\ell}{h_j}\left|\hat{\theta}_\ell^k\right|^2\right] \\ &=\; \mathbb{E}\left[\left|\theta_\ell^k + \boldsymbol{\psi}_\ell^{\mathsf{T}}\mathbf{w}^k\right|^2\right]\sum_{j=1}^{\tilde{n}}\frac{\rho_\ell}{h_j} \\ &\leq\; \tilde{n}\,\rho_\ell\,{d_u}^{\zeta}\left(\frac{n\,C_o}{C_p}\ell^{-2\alpha-1} + \sigma_w^2\right) \end{aligned} \tag{5.96}$$

where the upper bound essentially follows from the fact that the squared magnitudes of the ordered coefficients $\{\theta_\ell^k\}$ in the case of compressible signals are bounded as $|\theta_\ell^k|^2 \leq \frac{n\,C_o}{C_p}\ell^{-2\alpha-1}$ [see (5.9)]. This implies that

$$\rho_\ell \;=\; \frac{\lambda\,P}{\tilde{n}\,{d_u}^{\zeta}\left(n\,\widetilde{C}_o\,\ell^{-2\alpha-1}+\sigma_w^2\right)},\;\ell=1,\ldots,L \tag{5.97}$$

would suffice to satisfy the sum power constraint of P for each of the L channel uses, where $\widetilde{C}_o \stackrel{def}{=} C_o/C_p$ and $\lambda \in (0,1]$ is again the power scaling factor for controlling total network power consumption. We are now ready to state the distortion achievable for an α-compressible signal under the assumption of in-network collaboration.

Theorem 5.19 Given the sensor network model of Section 5.2 for an α-compressible signal and under the assumption of in-network collaboration enabling $\tilde{n}$ nodes in the network to have access to the entire observation vector $\mathbf{x}^k$ at each time instant k, the beamforming strategy described by the encoders in (5.93) and the decoder in (5.95) can achieve the following end-to-end distortion by employing L channel uses per source observation

$$\left(\frac{L}{n}\right)\sigma_w^2 \;+\; \left(\frac{1}{\tilde{n}}\right)\left(\frac{\sigma_z^2\,\widetilde{C}_o}{\lambda\,P}\right) \;\leq\; D \;\leq\; C_o\,L^{-2\alpha} + \\ +\;\left(\frac{L}{n}\right)\sigma_w^2\Big(1+\sigma_z^2\,{d_u}^{\zeta}\Big) \;+\; \left(\frac{1}{\tilde{n}}\right)\left(\frac{2\,\sigma_z^2\,{d_u}^{\zeta}\,\widetilde{C}_o}{\lambda\,P}\right). \tag{5.98}$$

Proof: To establish this theorem, first observe that (5.95) implies that $\forall\; k \in \mathbb{N}$

$$\begin{aligned}\mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k-\widehat{\mathbf{s}}^k\right\|_2^2\right] \;&\leq\; C_o\,L^{-2\alpha} \;+\; \left(\frac{L}{n}\right)\sigma_w^2 \;+\; \left(\frac{1}{n}\right)\sum_{\ell=1}^{L}\frac{\sigma_z^2}{\tilde{n}^2\,\rho_\ell} \\ &=\; C_o\,L^{-2\alpha} \;+\; \left(\frac{L}{n}\right)\sigma_w^2 \;+\; \left(\frac{\sigma_z^2\,{d_u}^{\zeta}\,\widetilde{C}_o}{\tilde{n}\,\lambda\,P}\right)\sum_{\ell=1}^{L}\ell^{-2\alpha-1} \;+\; \left(\frac{L}{n\,\tilde{n}}\right)\sigma_z^2\,{d_u}^{\zeta}\,\sigma_w^2 \\ &\leq\; C_o\,L^{-2\alpha} \;+\; \left(\frac{L}{n}\right)\sigma_w^2\Big(1+\sigma_z^2\,{d_u}^{\zeta}\Big) \;+\; \left(\frac{1}{\tilde{n}}\right)\left(\frac{2\,\sigma_z^2\,{d_u}^{\zeta}\,\widetilde{C}_o}{\lambda\,P}\right)\end{aligned} \tag{5.99}$$

where the last inequality follows from the fact that

$$\sum_{\ell=1}^{L}\ell^{-2\alpha-1} \;\leq\; 1+\int_1^L x^{-2\alpha-1}dx \;\leq\; 2 \tag{5.100}$$

and $\frac{L}{n\tilde{n}} \leq \frac{L}{n}$. Furthermore, from (5.95), we also have a lower bound of

$$\mathbb{E}\left[\frac{1}{n}\left\|\mathbf{s}^k - \widehat{\mathbf{s}}^k\right\|_2^2\right] \geq \left(\frac{L}{n}\right)\sigma_w^2 + \left(\frac{1}{n}\right)\sum_{\ell=1}^{L}\frac{\sigma_z^2}{\tilde{n}^2\,\rho_\ell}$$
$$\geq \left(\frac{L}{n}\right)\sigma_w^2 + \left(\frac{1}{\tilde{n}}\right)\left(\frac{\sigma_z^2\,\widetilde{C}_o}{\lambda\,P}\right). \tag{5.101}$$

Finally, combining the upper and lower bounds of (5.99) and (5.101), and taking the limit in k yields (5.98), thus completing the proof of the theorem. ∎

Remark 5.20 Under the assumption of $\tilde{n}$ nodes coherently transmitting the identical data, the cluster-to-FC MAC is effectively transformed into a point-to-point multiple-input single-output AWGN channel. Consequently, while the distortion expression in (5.98) has been obtained corresponding to an analog beamforming strategy of (5.93), a similar expression for distortion can be obtained by appropriately transforming the compressible source model into a stochastic one, and employing standard rate-distortion and capacity-cost analysis (in other words, by employing "digital" beamforming).

Remark 5.21 Note that the last term in the upper and lower bounds in (5.98) corresponds to the distortion component due to the noisy communication channel. The factor of $\tilde{n}$ in that term corresponds to the power-pooling gain due to coherent transmission of identical data: the greater the number of nodes coherently beamforming the identical data, the greater the power-pooling gain. Comparing this communication noise term to the last term in the upper and lower bounds in (5.39) shows that, in terms of scaling, the performance of the proposed estimation scheme of Section 5.5 is equivalent to that of an in-network collaboration based system that has a beamforming cluster consisting of $\tilde{n} = \frac{n}{L}$ nodes.

Analysis of (5.98) reveals that for optimal distortion scaling under the in-network collaboration assumption, $L \asymp n^{1/(2\alpha+1)}$ and the distortion component due to the communication noise should also scale at least as $\frac{L}{n} \asymp n^{-2\alpha/(2\alpha+1)}$. Consequently, this implies that as long as the extent of in-network collaboration is such that $\tilde{n} < n^{2\alpha/(2\alpha+1)}$, one cannot achieve the optimal distortion scaling under a fixed transmit power constraint of P: the power constraint itself needs to be

scaled up as $P \asymp n^{2\alpha/(2\alpha+1)}/\tilde{n}$ to achieve optimal distortion scaling. On the other hand, if the extent of in-network collaboration is such that $\tilde{n} > n^{2\alpha/(2\alpha+1)}$ then, in fact, λ need not be given by $\lambda = O(1)$. Rather, in that situation, it can be scaled down as $\lambda \asymp n^{2\alpha/(2\alpha+1)}/\tilde{n}$. Going back to the two extremes of $\tilde{n} = 1$ and $\tilde{n} = n$, this means that for the case of a single clusterhead in the network, we have $P_{\text{tot}} = O(\lambda\, L\, P) = O(n)$ and for the case where all the nodes in the network act as a big clusterhead, we have $P_{\text{tot}} = O(\lambda\, L\, P) = O(1)$. Essentially, as the cardinality of the beamforming cluster $\tilde{n}$ scales up as $1 \nearrow n$, the total network power scales down from $O(n)$ to $O(1)$. Remarkably, the proposed estimation scheme of Section 5.5 achieves the performance equivalent to that of a cluster with $\tilde{n} = n^{2\alpha/(2\alpha+1)}$ nodes, without requiring any in-network collaboration. Furthermore, while we have ignored the cost of in-network communication, we expect that it will increase monotonically with increase in the size of the beamforming cluster $\tilde{n}$.

REFERENCES

[1] J. A. Tropp, "Topics in sparse approximation," Ph.D. dissertation, The University of Texas at Austin, Austin, TX, 2004.

[2] E. J. Candès, "Compressive sampling," in *Proc. Int. Congr. of Mathematicians*, vol. III, Madrid, Spain, Aug. 2006, pp. 1433–1452.

[3] *IEEE Signal Processing Mag., Special Issue on Compressive Sampling*, vol. 25, no. 2, Mar. 2008.

[4] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Review*, vol. 51, no. 1, pp. 34–81, Feb. 2009.

[5] (2009) The compressive sensing resources webpage. Rice University. [Online]. Available: http://www.dsp.ece.rice.edu/cs

[6] E. J. Candès and T. Tao, "Decoding by linear programming," *IEEE Trans. Inform. Theory*, vol. 51, no. 12, pp. 4203–4215, Dec. 2005.

[7] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.

[8] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, Dec. 2006.

[9] E. J. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, Mar. 2006.

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

[11] R. Baraniuk, M. Davenport, R. A. DeVore, and M. B. Wakin, "A simple proof of the restricted isometry property for random matrices," in *Constructive Approximation*. New York, NY: Springer, 2008.

[12] M. Rudelson and R. Vershynin, "On sparse reconstruction from Fourier and Gaussian measurements," *Commun. Pure Appl. Math.*, no. 8, pp. 1025–1045, Aug. 2008.

[13] W. Li and J. C. Preisig, "Estimation of rapidly time-varying sparse channels," *IEEE J. Oceanic Eng.*, vol. 32, no. 4, pp. 927–939, Oct. 2007.

[14] W. U. Bajwa, J. Haupt, G. Raz, and R. Nowak, "Compressed channel sensing," in *Proc. 42nd Annu. Conf. Information Sciences and Systems (CISS '08)*, Princeton, NJ, Mar. 2008, pp. 5–10.

[15] W. U. Bajwa, A. M. Sayeed, and R. Nowak, "Learning sparse doubly-selective channels," in *Proc. 45th Annu. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sep. 2008, pp. 575–582.

[16] L. Cohen, *Time-Frequency Analysis*. Upper Saddle River, NJ: Prentice-Hall, 1994.

[17] R. F. Marcia and R. M. Willett, "Compressive coded aperture superresolution image reconstruction," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP '08)*, Las Vegas, NV, Apr. 2008, pp. 833–836.

[18] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk, "Beyond Nyquist: Efficient sampling of sparse bandlimited signals," submitted. [Online]. Available: arXiv:0902.0026v1

[19] F. Santosa and W. W. Symes, "Linear inversion of band-limited reflection seismograms," *SIAM J. Sci. Statist. Comput.*, vol. 7, no. 4, pp. 1307–1330, Oct. 1986.

[20] M. A. Herman and T. Strohmer, "High-resolution radar via compressed sensing," *IEEE Trans. Signal Processing*, vol. 57, no. 6, pp. 2275–2284, Jun. 2009.

[21] A. Goldsmith, *Wireless Communications*. New York, NY: Cambridge University Press, 2005.

[22] D. Tse and P. Viswanath, *Fundamentals of Wireless Communication*. Cambridge, U.K.: Cambridge University Press, 2005.

[23] L. Tong, B. M. Sadler, and M. Dong, "Pilot-assisted wireless transmissions," *IEEE Signal Processing Mag.*, vol. 21, no. 6, pp. 12–25, Nov. 2004.

[24] J. K. Cavers, "An analysis of pilot symbol assisted modulation for Rayleigh fading channels," *IEEE Trans. Veh. Technol.*, vol. 40, no. 4, pp. 686–693, Nov. 1991.

[25] R. Negi and J. Cioffi, "Pilot tone selection for channel estimation in a mobile OFDM system," *IEEE Trans. Consumer Electron.*, vol. 44, no. 3, pp. 1122–1128, Aug. 1998.

[26] M.-A. R. Baissas and A. M. Sayeed, "Pilot-based estimation of time-varying multipath channels for coherent CDMA receivers," *IEEE Trans. Signal Processing*, vol. 50, no. 8, pp. 2037–2049, Aug. 2002.

[27] X. Ma, G. B. Giannakis, and S. Ohno, "Optimal training for block transmissions over doubly selective wireless fading channels," *IEEE Trans. Signal Processing*, vol. 51, no. 5, pp. 1351–1366, May 2003.

[28] A. P. Kannu and P. Schniter, "Design and analysis of MMSE pilot-aided cyclic-prefixed block transmissions for doubly selective channels," *IEEE Trans. Signal Processing*, vol. 56, no. 3, pp. 1148–1160, Mar. 2008.

[29] T. L. Marzetta, "BLAST training: Estimating channel characteristics for high capacity space-time wireless," in *Proc. 37th Annu. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sep. 1999.

[30] B. Hassibi and B. M. Hochwald, "How much training is needed in multiple-antenna wireless links?" *IEEE Trans. Inform. Theory*, vol. 49, no. 4, pp. 951–963, Apr. 2003.

[31] I. Barhumi, G. Leus, and M. Moonen, "Optimal training design for MIMO OFDM systems in mobile wireless channels," *IEEE Trans. Signal Processing*, vol. 51, no. 6, pp. 1615–1624, Jun. 2003.

[32] H. Minn and N. Al-Dhahir, "Optimal training signals for MIMO OFDM channel estimation," *IEEE Trans. Wireless Commun.*, vol. 5, no. 5, pp. 1158–1168, May 2006.

[33] A. Saleh and R. Valenzuela, "A statistical model for indoor multipath propagation," *IEEE J. Select. Areas Commun.*, vol. 5, no. 2, pp. 128–137, Feb. 1987.

[34] A. F. Molisch, "Ultrawideband propagation channels-Theory, measurement, and modeling," *IEEE Trans. Veh. Technol.*, vol. 54, no. 5, pp. 1528–1545, Sep. 2005.

[35] Z. Yan, M. Herdin, A. M. Sayeed, and E. Bonek, "Experimental study of MIMO channel statistics and capacity via the virtual channel representation," University of Wisconsin-Madison, Tech. Rep., Feb. 2007.

[36] N. Czink, X. Yin, H. Ozcelik, M. Herdin, E. Bonek, and B. H. Fleury, "Cluster characteristics in a MIMO indoor propagation environment," *IEEE Trans. Wireless Commun.*, vol. 6, no. 4, pp. 1465–1475, Apr. 2007.

[37] L. Vuokko, V.-M. Kolmonen, J. Salo, and P. Vainikainen, "Measurement of large-scale cluster power characteristics for geometric channel models," *IEEE Trans. Antennas Propagat.*, vol. 55, no. 11, pp. 3361–3365, Nov. 2007.

[38] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," in *Proc. 5th Annu. ACM/IEEE Int. Conf. Mobile Computing and Networking (MobiCom '99)*, Seattle, WA, Aug. 1999, pp. 263–270.

[39] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: A survey," *Computer Networks*, vol. 38, no. 4, pp. 393–422, Mar. 2002.

[40] M. Gastpar and M. Vetterli, "Source-channel communication in sensor networks," in *Proc. 2nd Int. Workshop Information Processing in Sensor Networks (IPSN '03)*, Palo Alto, CA, Apr. 2003, pp. 162–177.

[41] K. Liu and A. M. Sayeed, "Optimal distributed detection strategies for wireless sensor networks," in *Proc. 42nd Annu. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Oct. 2004.

[42] M. Gastpar and M. Vetterli, "Power, spatio-temporal bandwidth, and distortion in large sensor networks," *IEEE J. Select. Areas Commun.*, vol. 23, no. 4, pp. 745–754, Apr. 2005.

[43] W. U. Bajwa, A. M. Sayeed, and R. Nowak, "Matched source-channel communication for field estimation in wireless sensor networks," in *Proc. 4th Int. Symp. Information Processing in Sensor Networks (IPSN '05)*, Los Angeles, CA, Apr. 2005, pp. 332–339.

[44] P. Lancaster and M. Tismenetsky, *The Theory of Matrices*, 2nd ed. Orlando, FL: Academic Press, 1985.

[45] M. Akcakaya and V. Tarokh, "A frame construction and a universal distortion bound for sparse representations," *IEEE Trans. Signal Processing*, vol. 56, no. 6, pp. 2443–2450, Jun. 2008.

[46] M. Mishali and Y. C. Eldar, "Blind multiband signal reconstruction: Compressed sensing for analog signals," *IEEE Trans. Signal Processing*, vol. 57, no. 3, pp. 993–1009, Mar. 2009.

[47] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk, "Single-pixel imaging via compressive sampling," *IEEE Signal Processing Mag.*, vol. 25, no. 2, pp. 83–91, Mar. 2008.

[48] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Processing Mag.*, vol. 25, no. 2, pp. 72–82, Mar. 2008.

[49] W. U. Bajwa, J. Haupt, A. M. Sayeed, and R. Nowak, "A universal matched source-channel communication scheme for wireless sensor ensembles," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing (ICASSP '06)*, vol. 5, Toulouse, France, May 2006, pp. 1153–1156.

[50] ——, "Compressive wireless sensing," in *Proc. 5th Int. Conf. Information Processing in Sensor Networks (IPSN '06)*, Nashville, TN, Apr. 2006, pp. 134–142.

[51] M. Rabbat, J. Haupt, A. Singh, and R. Nowak, "Decentralized compression and predistribution via randomized gossiping," in *Proc. 5th Int. Conf. Information Processing in Sensor Networks (IPSN '06)*, Nashville, TN, Apr. 2006, pp. 51–59.

[52] W. U. Bajwa, J. Haupt, A. M. Sayeed, and R. Nowak, "Joint source–channel communication for distributed estimation in sensor networks," *IEEE Trans. Inform. Theory*, vol. 53, no. 10, pp. 3629–3653, Oct. 2007.

[53] J. Haupt, W. U. Bajwa, M. Rabbat, and R. Nowak, "Compressed sensing for networked data," *IEEE Signal Processing Mag.*, vol. 25, no. 2, pp. 92–101, Mar. 2008.

[54] I. F. Gorodnitsky and B. D. Rao, "Sparse signal reconstruction from limited data using FOCUSS: A re-weighted minimum norm algorithm," *IEEE Trans. Signal Processing*, vol. 45, no. 3, pp. 600–616, Mar. 1997.

[55] D. L. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization," *Proc. Natl. Acad. Sci.*, vol. 100, no. 5, pp. 2197–2202, Mar. 2003.

[56] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, no. 2, pp. 227–234, Apr. 1995.

[57] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scientific Comput.*, vol. 20, no. 1, pp. 33–61, Jan. 1998.

[58] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Processing*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.

[59] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Asilomar Conf. Signals, Systems and Computers*, Pacific Grove, CA, Nov. 1993, pp. 40–44.

[60] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.*, vol. 57, no. 11, pp. 1413–1457, Aug. 2004.

[61] M. Fornasier and H. Rauhut, "Iterative thresholding algorithms," *Appl. Comput. Harmon. Anal.*, vol. 25, no. 2, pp. 187–208, 2008.

[62] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed. Springer, 2006.

[63] E. J. Candès, "The restricted isometry property and its implications for compressed sensing," in *C. R. Acad. Sci., Ser. I*, Paris, 2008, vol. 346, pp. 589–592.

[64] J. F. Claerbout and F. Muir, "Robust modeling with erratic data," *Geophysics*, vol. 38, no. 5, pp. 826–844, Oct. 1973.

[65] J.-J. Fuchs, "Recovery of exact sparse representations in the presence of bounded noise," *IEEE Trans. Inform. Theory*, vol. 51, no. 10, pp. 3601–3608, Oct. 2005.

[66] D. L. Donoho, M. Elad, and V. N. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inform. Theory*, vol. 52, no. 1, pp. 6–18, Jan. 2006.

[67] J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," *IEEE Trans. Inform. Theory*, vol. 52, no. 9, pp. 4036–4048, Sep. 2006.

[68] E. J. Candès and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n," *Ann. Statist.*, vol. 35, no. 6, pp. 2313–2351, Dec. 2007.

[69] P. J. Bickel, Y. Ritov, and A. B. Tsybakov, "Simultaneous analysis of lasso and Dantzig selector," *Ann. Statist.*, vol. 37, no. 4, pp. 1705–1732, Aug. 2009.

[70] Z. Ben-Haim, Y. C. Eldar, and M. Elad, "Near-oracle performance of basis pursuit under random noise," submitted. [Online]. Available: arXiv:0903.4579v1

[71] R. Gribonval, R. M. Figueras i Ventura, and P. Vandergheynst, "A simple test to check the optimality of a sparse signal approximation," *Signal Processing*, vol. 86, no. 3, pp. 496–510, 2006.

[72] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge University Press, 2004.

[73] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Roy. Statist. Soc. Ser. B*, vol. 58, no. 1, pp. 267–288, 1996.

[74] J.-J. Fuchs, "Extension of the Pisarenko method to sparse linear arrays," *IEEE Trans. Signal Processing*, vol. 45, no. 10, pp. 2413–2421, Oct. 1997.

[75] M. A. T. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Select. Topics Signal Processing*, vol. 1, no. 4, pp. 586–597, Dec. 2007.

[76] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. Signal Processing*, vol. 57, no. 7, pp. 2479–2493, Jul. 2009.

[77] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*. Upper Saddle River, NJ: Prentice Hall, 1998.

[78] E. Kreyszig, *Introductory Functional Analysis with Applications*. New York, NY: John Wiley and Sons, Inc., 1978.

[79] G. Wagner, P. Schmieder, A. S. Stern, and J. C. Hoch, "Application of nonlinear sampling schemes to COSY-type spectra," *J. Biomolecular NMR*, vol. 3, no. 5, pp. 569–576, Sep. 1993.

[80] J. A. Tropp, M. B. Wakin, M. F. Duarte, D. Baron, and R. G. Baraniuk, "Random filters for compressive sampling and reconstruction," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP '06)*, Toulouse, France, May 2006, pp. 872–875.

[81] M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz, "Single-shot compressive spectral imaging with a dual-disperser architecture," *Opt. Express*, vol. 15, no. 21, pp. 14 013–14 027, 2007.

[82] R. M. Gray, *Toeplitz and Circulant Matrices: A Review*. Hanover, MA: Now Publishers, 2006.

[83] M. J. Wainwright, "Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1-constrained quadratic programming (Lasso)," *IEEE Trans. Inform. Theory*, vol. 55, no. 5, pp. 2183–2202, May 2009.

[84] S. A. Geršgorin, "Über die Abgrenzung der Eigenwerte einer Matrix," *Izv. Akad. Nauk SSSR Ser. Fiz.-Mat.*, vol. 6, pp. 749–754, 1931.

[85] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J. Amer. Statist. Assoc.*, vol. 58, no. 301, pp. 13–30, Mar. 1963.

[86] J. Lawrence, G. E. Pfander, and D. Walnut, "Linear independence of Gabor systems in finite dimensional vector spaces," *J. Fourier Anal. Appl.*, vol. 11, no. 6, pp. 715–726, Dec. 2005.

[87] H. G. Feichtinger and T. Strohmer, Eds., *Gabor Analysis and Algorithms: Theory and Applications*. Cambridge, MA: Birkhäuser Boston, 1998.

[88] G. A. Wright, "Magnetic resonance imaging," *IEEE Signal Processing Mag.*, vol. 14, no. 1, pp. 56–66, Jan. 1997.

[89] J. A. C. Bingham, "Multicarrier modulation for data transmission: An idea whose time has come," *IEEE Commun. Mag.*, vol. 28, no. 5, pp. 5–14, May 1990.

[90] M. Ledoux and M. Talagrand, *Probability in Banach Spaces*. New York, NY: Springer-Verlag, 1991.

[91] J. Romberg, "Compressive sensing by random convolution," submitted. [Online]. Available: http://users.ece.gatech.edu/~justin/Publications.html

[92] U. Haagerup, "The best constants in the Khintchine inequality," *Studia Math.*, vol. 70, no. 3, pp. 231–283, 1981.

[93] J. Haupt, W. U. Bajwa, G. Raz, and R. Nowak, "Toeplitz compressed sensing matrices with applications to sparse channel estimation," submitted. [Online]. Available: http://www.ece.wisc.edu/~nowak/sub08_toep.pdf

[94] T. Strohmer and R. Heath, "Grassmanian frames with applications to coding and communication," *Appl. Comput. Harmon. Anal.*, vol. 14, no. 3, pp. 257–275, May 2003.

[95] S. F. Cotter and B. D. Rao, "Sparse channel estimation via matching pursuit with application to equalization," *IEEE Trans. Commun.*, vol. 50, no. 3, pp. 374–377, Mar. 2002.

[96] W. U. Bajwa, J. Haupt, G. Raz, S. J. Wright, and R. Nowak, "Toeplitz-structured compressed sensing matrices," in *Proc. 14th IEEE/SP Workshop Statistical Signal Processing (SSP '07)*, Madison, WI, Aug. 2007, pp. 294–298.

[97] R. A. DeVore, "Deterministic constructions of compressed sensing matrices," *J. Complexity*, vol. 23, pp. 918–925, Aug. 2007.

[98] G. E. Pfander, H. Rauhut, and J. Tanner, "Identification of matrices having a sparse representation," *IEEE Trans. Signal Processing*, vol. 56, no. 11, pp. 5376–5388, Nov. 2008.

[99] J. G. Proakis, *Digital Communications*, 4th ed. New York, NY: McGraw-Hill, 2001.

[100] L. Zheng and D. N. C. Tse, "Diversity and multiplexing: A fundamental tradeoff in multiple-antenna channels," *IEEE Trans. Inform. Theory*, vol. 49, no. 5, pp. 1073–1096, May 2003.

[101] L. Tong and S. Perreau, "Multichannel blind identification: From subspace to maximum likelihood methods," *Proc. IEEE*, vol. 86, no. 10, pp. 1951–1968, Oct. 1998.

[102] S. F. Cotter and B. D. Rao, "The adaptive matching pursuit algorithm for estimation and equalization of sparse time-varying channels," in *Proc. 34th Asilomar Conf. Signals, Systems and Computers*, vol. 2, Pacific Grove, CA, Oct./Nov. 2000, pp. 1772–1776.

[103] W. Dongming, H. Bing, Z. Junhui, G. Xiqi, and Y. Xiaohu, "Channel estimation algorithms for broadband MIMO-OFDM sparse channel," in *Proc. 14th IEEE Int. Symp. Personal, Indoor and Mobile Radio Communications (PIMRC '02)*, vol. 2, Beijing, China, Sep. 2003, pp. 1929–1933.

[104] C. Carbonelli, S. Vedantam, and U. Mitra, "Sparse channel estimation with zero tap detection," *IEEE Trans. Wireless Commun.*, vol. 6, no. 5, pp. 1743–1753, May 2007.

[105] G. Tauböck and F. Hlawatsch, "A compressed sensing technique for OFDM channel estimation in mobile environments: Exploiting channel sparsity for reducing pilots," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP '08)*, Las Vegas, NV, Apr. 2008, pp. 2885–2888.

[106] M. Sharp and A. Scaglione, "Application of sparse signal recovery to pilot-assisted channel estimation," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP '08)*, Las Vegas, NV, Apr. 2008, pp. 3469–3472.

[107] W. U. Bajwa, A. M. Sayeed, and R. Nowak, "Compressed sensing of wireless channels in time, frequency, and space," in *Proc. 42nd Asilomar Conf. Signals, Systems and Computers*, Pacific Grove, CA, Oct. 2008, pp. 2048–2052.

[108] P. Bello, "Characterization of randomly time-variant linear channels," *IEEE Trans. Commun.*, vol. 11, no. 4, pp. 360–393, Dec. 1963.

[109] T. Kailath, "Measurements on time-variant communication channels," *IRE Trans. Inform. Theory*, vol. 8, no. 5, pp. 229–236, Sep. 1962.

[110] P. Bello, "Measurement of random time-variant linear channels," *IEEE Trans. Inform. Theory*, vol. 15, no. 4, pp. 469–475, Jul. 1969.

[111] D. B. Kilfoyle and A. B. Baggeroer, "The state of the art in underwater acoustic telemetry," *IEEE J. Oceanic Eng.*, vol. 25, no. 1, pp. 4–27, Jan. 2000.

[112] D. Slepian, "On bandwidth," *Proc. IEEE*, vol. 64, no. 3, pp. 292–300, Mar. 1976.

[113] B. D. Van Veen and K. M. Buckley, "Beamforming: A versatile approach to spatial filtering," *IEEE ASSP Mag.*, vol. 5, no. 2, pp. 4–24, Apr. 1988.

[114] H. Krim and M. Viberg, "Two decades of array signal processing research: The parametric approach," *IEEE Signal Processing Mag.*, vol. 13, no. 4, pp. 67–94, Jul. 1996.

[115] A. M. Sayeed, "Deconstructing multiantenna fading channels," *IEEE Trans. Signal Processing*, vol. 50, no. 10, pp. 2563–2579, Oct. 2002.

[116] A. M. Sayeed and B. Aazhang, "Joint multipath-Doppler diversity in mobile wireless communications," *IEEE Trans. Commun.*, vol. 47, no. 1, pp. 123–132, Jan. 1999.

[117] A. M. Sayeed and V. Veeravalli, "The essential degrees of freedom in space-time fading channels," in *Proc. 13th IEEE Int. Symp. Personal, Indoor and Mobile Radio Communications (PIMRC '02)*, vol. 4, Lisbon, Portugal, Sep. 2002, pp. 1512–1516.

[118] A. M. Sayeed, "A virtual representation for time- and frequency-selective correlated MIMO channels," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP '03)*, vol. 4, Hong Kong, Apr. 2003, pp. 648–651.

[119] R. S. Kennedy, *Fading Dispersive Communication Channels*. Wiley-Interscience, 1969.

[120] I. E. Telatar, "Capacity of multi-antenna Gaussian channels," *Eur. Trans. Telecommun.*, vol. 10, no. 6, pp. 585–595, Nov. 1999.

[121] G. J. Foschini and M. J. Gans, "On limits of wireless communications in a fading environment when using multiple antennas," *Wireless Pers. Commun.*, vol. 6, no. 3, pp. 311–335, Mar. 1998.

[122] "Final report on the comparative trials of the digital television systems," SET/ABERT Study, Brazilian Telecommunications Agency (Anatel), Jun. 2000.

[123] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ: Prentice Hall, 1993.

[124] G. H. Hardy, J. E. Littlewood, and G. Pólya, *Inequalities*, 2nd ed. Cambridge, U.K.: Cambridge University Press, 1988.

[125] W. Kozek and A. F. Molisch, "Nonorthogonal pulseshapes for multicarrier communications in doubly dispersive channels," *IEEE J. Select. Areas Commun.*, vol. 16, no. 8, pp. 1579–1589, Oct. 1998.

[126] K. Liu, T. Kadous, and A. M. Sayeed, "Orthogonal time-frequency signaling over doubly dispersive channels," *IEEE Trans. Inform. Theory*, vol. 50, no. 11, pp. 2583–2603, Nov. 2004.

[127] A. M. Sayeed and V. Raghavan, "Maximizing MIMO capacity in sparse multipath with reconfigurable antenna arrays," *IEEE J. Select. Topics Signal Processing*, vol. 1, no. 1, pp. 156–166, Jun. 2007.

[128] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE J. Select. Areas Commun.*, vol. 23, no. 2, pp. 201–220, Feb. 2005.

[129] M. Vetterli, "Wavelets, approximation, and compression," *IEEE Signal Processing Mag.*, vol. 18, no. 5, pp. 59–73, Sep. 2001.

[130] R. A. DeVore, B. Jawerth, and B. J. Lucier, "Image compression through wavelet transform coding," *IEEE Trans. Inform. Theory*, vol. 38, no. 2, pp. 719–746, Mar. 1992.

[131] D. L. Donoho, M. Vetterli, R. A. DeVore, and I. Daubechies, "Data compression and harmonic analysis," *IEEE Trans. Inform. Theory*, vol. 44, no. 6, pp. 2435–2476, Oct. 1998.

[132] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd ed. Academic Press, 1999.

[133] V. K. Goyal, *Single and Multiple Description Transform Coding with Bases and Frames*. Philadelphia, PA: SIAM, 2001.

[134] A. Cohen, I. Daubechies, O. G. Guleryuz, and M. T. Orchard, "On the importance of combining wavelet-based nonlinear approximation with coding strategies," *IEEE Trans. Inform. Theory*, vol. 48, no. 7, pp. 1895–1921, Jul. 2002.

[135] R. A. Devore, "Nonlinear approximation," in *Acta Numerica*, A. Iserles, Ed. Cambridge, U.K.: Cambridge University Press, 1998, vol. 7, pp. 51–150.

[136] A. Cohen, W. Dahmen, and R. A. Devore, "Compressed sensing and best k-term approximation," *J. Amer. Math. Soc.*, vol. 22, no. 1, p. 211231, Jan. 2009.

[137] D. L. Donoho, "Sparse components of images and optimal atomic decompositions," in *Constructive Approximation*. New York, NY: Springer, Jan. 2001, vol. 17, no. 3, pp. 353–382.

[138] R. Wagner, H. Choi, R. Baraniuk, and V. Delouille, "Distributed wavelet transform for irregular sensor network grids," in *Proc. 13th IEEE/SP Workshop Statistical Signal Processing (SSP '05)*, Bordeaux, France, Jul. 2005, pp. 1196–1201.

[139] L.-L. Xie and P. R. Kumar, "A network information theory for wireless communication: Scaling laws and optimal operation," *IEEE Trans. Inform. Theory*, vol. 50, no. 5, pp. 748–767, May 2004.

[140] M. Gastpar and M. Vetterli, "On the capacity of large Gaussian relay networks," *IEEE Trans. Inform. Theory*, vol. 51, no. 3, pp. 765–779, Mar. 2005.

[141] H. El-Gamal, "On the scaling laws of dense wireless sensor networks: The data gathering channel," *IEEE Trans. Inform. Theory*, vol. 51, no. 3, pp. 1229–1234, Mar. 2005.

[142] T. S. Rappaport, *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ: Prentice-Hall, 2002.

[143] R. Mudumbai, J. Hespanha, U. Madhow, and G. Barriac, "Scalable feedback control for distributed beamforming in sensor networks," in *Proc. Int. Symp. Information Theory (ISIT '05)*, Sep. 2005, pp. 137–141.

[144] ——, "Distributed transmit beamforming using feedback control," submitted. [Online]. Available: arXiv:cs/0603072v1

[145] L. Wasserman, *All of Nonparametric Statistics*, ser. Springer Texts in Statistics. New York, NY: Springer-Verlag, 2006.

[146] R. Nowak, U. Mitra, and R. Willett, "Estimating inhomogeneous fields using wireless sensor networks," *IEEE J. Select. Areas Commun.*, vol. 22, no. 6, pp. 999–1006, Jun. 2004.

[147] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation by wavelet shrinkage," *Biometrika*, vol. 81, no. 3, pp. 425–455, 1994.

[148] D. L. Donoho and J. M. Johnstone, "Ideal denoising in an orthonormal basis chosen from a library of bases," in *C. R. Acad. Sci., Ser. I*, Paris, 1994, vol. 319, pp. 1317–1322.

[149] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah, "Gossip algorithms: Design, analysis and applications," in *Proc. IEEE Conf. Computer Communications (INFOCOM '05)*, vol. 3, Miami, FL, Mar. 2005, pp. 1653–1664.

[150] A. G. Dimakis, A. D. Sarwate, and M. Wainwright, "Geographic gossip : Efficient aggregation for sensor networks," in *Proc. 5th Int. Conf. Information Processing in Sensor Networks (IPSN '06)*, Nashville, TN, Apr. 2006, pp. 69–76.

[151] A. P. Korostelev and A. B. Tsybakov, *Minimax Theory of Image Reconstruction*, ser. Lecture Notes in Statistics. New York, NY: Springer-Verlag, 1993, vol. 82.

[152] T. Sivanadyan and A. M. Sayeed, "Active wireless sensing for rapid information retrieval in sensor networks," in *Proc. 5th Int. Conf. Information Processing in Sensor Networks (IPSN '06)*, Nashville, TN, Apr. 2006, pp. 85–92.

[153] P. L'Ecuyer, "Random number generation," in *Handbook of Computational Statistics: Concepts and Methods*, J. E. Gentle, W. Härdle, and Y. Mori, Eds. Springer-Verlag, 2004, ch. II.2, pp. 35–70.

[154] K. Liu, H. El-Gamal, and A. M. Sayeed, "On optimal parametric field estimation in sensor networks," in *Proc. 13th IEEE/SP Workshop Statistical Signal Processing (SSP '05)*, Bordeaux, France, Jul. 2005, pp. 1170–1175.

[155] G. Mergen and L. Tong, "Type-based estimation over multiaccess channels," *IEEE Trans. Signal Processing*, vol. 54, no. 2, pp. 613–626, Feb. 2006.

[156] P. Gao and C. Tepedelenlioglu, "Practical issues in parameter estimation over fading channels with TBMA wireless sensor networks," in *Proc. 49th IEEE GLOBECOM Technical Conference*, San Francisco, CA, Nov. 2006.

[157] S. D. Servetto, "On the feasibility of large-scale wireless sensor networks," in *Proc. 40th Annu. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, 2002.

[158] S. S. Pradhan, J. Kusuma, and K. Ramchandran, "Distributed compression in a dense microsensor network," *IEEE Signal Processing Mag.*, vol. 19, no. 2, pp. 51–60, Mar. 2002.

[159] D. Marco, E. Duarte-Melo, M. Liu, and D. Neuhoff, "On the many-to-one transport capacity of a dense wireless sensor network and the compressibility of its data," in *Proc. 2nd Int. Workshop Information Processing in Sensor Networks (IPSN '03)*, Apr. 2003, pp. 1–16.

[160] P. Ishwar, A. Kumar, and K. Ramchandran, "Distributed sampling for dense sensor networks: A bit-conservation principle," in *Proc. 2nd Int. Workshop Information Processing in Sensor Networks (IPSN '03)*, Apr. 2003, pp. 17–31.

[161] M. Gastpar, M. Vetterli, and P. L. Dragotti, "Sensing reality and communicating bits: A dangerous liaison," *IEEE Signal Processing Mag.*, vol. 23, no. 4, pp. 70–83, Jul. 2006.

[162] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York, NY: John Wiley and Sons, Inc., 1991.

[163] S. Cui, J. Xiao, A. Goldsmith, Z. Q. Luo, , and H. V. Poor, "Energy-efficient joint estimation in sensor networks: Analog vs. digital," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP '05)*, Philadelphia, PA, Mar. 2005, pp. 745–748.

[164] D. Baron, M. B. Wakin, M. F. Duarte, S. Sarvotham, and R. G. Baraniuk, "Distributed compressive sensing," submitted. [Online]. Available: arXiv:0901.3403v1

[165] M. F. Duarte, M. B. Wakin, D. Baron, and R. G. Baraniuk, "Universal distributed sensing via random projections," in *Proc. 5th Int. Conf. Information Processing in Sensor Networks (IPSN '06)*, Nashville, TN, Apr. 2006, pp. 177–185.