UNIVERSITY OF CALIFORNIA,
IRVINE

On Convex Relaxations for Joint Transmit-Receive Design in Radar Systems

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Electrical Engineering

by

Sean Myles O'Rourke

Dissertation Committee:
Professor A. Lee Swindlehurst, Chair
Professor Glenn Healey
Doctor Muralidhar Rangaswamy (AFRL/RYAR)

2018

DEDICATION

To my father, who lit the way.
To my mother, who made sure I followed the light.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would like to thank, first and foremost, my father for giving me the inspiration 28 years ago to pursue this journey. Without the "swirlie-whirlies" and "bumbly-bumblies", statistical signal processing never enters my mind as a possibility. His guidance, love, and support drew me to and kept me places where I could intellectually explore in the way that I have. Next, but no less importantly, I thank my mother. In the face of turmoil and uncertainty, she made sure that I pushed forward, that I felt comfort, that I could take the path set in front of me. This dissertation is as much your project as it is mine.

I owe my advisor, A. Lee Swindlehurst, some of the deepest gratitude that I can ever muster. His willingness to take and keep me on as a student and to provide me with the space to explore this field is more than I could have ever asked from any person. His style of plain-spoken, but deeply illuminating technical understanding is a constant inspiration in my work and my writing. Thank you, Lee.

I am lucky to have Muralidhar Rangaswamy of the Air Force Research Laboratory, Sensors Directorate as a member of this committee, an early backer of my research, a mentor, and a friend. I am forever thankful that he saw my potential for the SMART Scholarship and as a government researcher. Without him, this dissertation does not exist. Period. Thank you, Murali.

I would also like to thank Glenn Healey for being a positive contributor to my career in many (sometimes unexpected!) ways. Advising me for my senior undergraduate project, serving on my defense commitee, hooding me at graduation, and somehow finding time to be a strong voice for engineering in baseball – thank you for it all.

My deep gratitude is extended to my primary co-author, Pawan Setlur, who originally inspired the idea for this project in 2013 by asking "Is this cost function convex?" and continues to be a great sounding board for and champion of my ideas. Thanks also to the rest of the Fully Adaptive Radar STAR team, past and present. This also extends to my leadership, coworkers, and other researchers at the Air Force Research Laboratory and Wright State Research Institute: Dr. Greg Cazzell, Dr. Linda Moore, Dr. Jeffrey Sanders, Jackie Barker-Touissant, Bill Baldygo, and Doug Hager; Dr. Clark Taylor; and the rest of AFRL/RYAP.

Numerous friends and colleagues helped me stay (somewhat) sane through this process. Not only are there too many to comprehensively list, it has become clear that Department of Defense regulations prohibit me from naming them. I would like to sincerely apologize to them – your contributions may be forced to be anonymous here, but they are not anonymous to me.

This work has been supported by the following funding sources: the Sensors Directorate of the Air Force Research Laboratory (AFRL/RYRT) under contract number FA8650-08-D-1303 to Dynetics, Inc.; the American Society of Engineering Education (ASEE) and the Department of Defense (DoD) through the 2013 Science, Mathematics, and Research for Transformation

(SMART) Scholarship Program; and the Air Force Office of Scientific Research through Project 13RY10COR under the Sensing, Surveillance, and Navigation Project and Project 17RYCOR481 under the Dynamic Data Driven Application Systems (DDDAS) Program.

I thank the IET for permission to include portions of Chapter 2 of this dissertation, which first appeared as a conference paper at RADAR 2017 – International Conference on Radar Systems that I co-authored with Drs. Setlur and Rangaswamy, who have given their permission to include this work in my dissertation. I also thank the IEEE for permission to include portions of Chapters 3 and 4 of this dissertation. The portions of Chapter 3 originally appeared in IEEE Transactions on Signal Processing. The portions of Chapter 4 are currently under review for the same journal. Both works were co-authored with Prof. Swindlehurst and Drs. Setlur and Rangaswamy, who have given their permission to include this work in my dissertation.

Finally, many thanks and much love to my partner in crime and the only person who will reliably call me "Dr. O'Rourke" (whether trolling or not) from now on. You have sacrificed constantly to be with me, from moving to new places, to enduring many weekends where I went to the office to write these words. The universe cannot contain my love and appreciation for you, so this paragraph will serve as the least sufficent proxy I can muster.

CURRICULUM VITAE

Sean Myles O'Rourke

EDUCATION

Doctor of Philosophy in Electrical Engineering **2018**
University of California, Irvine *Irvine, CA*

Master of Science in Electrical Engineering (Systems) **2011**
University of California, Irvine *Irvine, CA*

Bachelor of Science in Electrical Engineering *(cum laude)* **2008**
University of California, Irvine *Irvine, CA*

RESEARCH EXPERIENCE

Research Electronics Engineer **2015–2018**
Air Force Res. Lab., Sensors Directorate (AFRL/RYAP) *Wright-Patterson AFB, OH*

Graduate Research Assistant **2008–2014**
University of California, Irvine *Irvine, CA*

TEACHING EXPERIENCE

Teaching Assistant, ENG 198CW: Engineering Writing **2015**
University of California, Irvine *Irvine, CA*

Teaching Assistant, EECS 152B/CSE 135B: DSP Lab **2010, 2011, 2013, 2015**
University of California, Irvine *Irvine, CA*

Teaching Assistant, CSE 181CW: Engineering Writing **2009**
University of California, Irvine *Irvine, CA*

REFEREED JOURNAL PUBLICATIONS

Relaxed Biquadratic Optimization for Joint Filter-Signal Design in Signal-Dependent STAP **2018**
IEEE Transactions on Signal Processing

Quadratic Semidefinite Programming for Waveform-Constrained Joint Filter-Signal Design in STAP **2018**
submitted to IEEE Transactions on Signal Processing

CONFERENCE PUBLICATIONS

Constrained Least Squares, SDP, & QCQP Perspectives on Joint Biconvex Radar Receiver & Waveform Design **Oct. 2017**
Radar 2017 – International Conference on Radar Systems, Belfast, UK

Limited Field-of-View Multimodal Sensor Adaptation for Data Association **May 2012**
2012 IEEE Sensor Array and Multichannel Workshop, Hoboken, NJ

Closed-Loop Tracking Using Multimodal RF/EO Sensors **Nov. 2010**
2010 Asilomar Conference on Signals, Systems, and Computers, Pacific Grove, CA

ABSTRACT OF THE DISSERTATION

On Convex Relaxations for Joint Transmit-Receive Design in Radar Systems

By

Sean Myles O'Rourke

Doctor of Philosophy in Electrical Engineering

University of California, Irvine, 2018

Professor A. Lee Swindlehurst, Chair

This dissertation investigates joint design of transmit and receive resources in radar systems, given prior knowledge of a signal-dependent clutter environment, based on simultaneous convex relaxation. Motivated by waveform-adaptive space-time adaptive processing (WA-STAP), we first analyze traditional approaches to joint trans-receive design, which alternate between optimal signal designs for fixed filters and vice versa to maximize signal-to-interference-plus-noise ratio (SINR). First, we demonstrate that SINR & its variants are non-convex bi-quadratic programs (BQP). Next, we provide perspectives on current alternating methods used to solve BQPs in terms of linear semidefinite (SDP) & quadratically-constrained quadratic (QCQP) programs. We then develop a novel method of performing this joint design based on quadratic semidefinite programming (QSDP). Our proposed relaxation scheme first analyzes the power-constrained problem, and is accompanied by exploration of operator structures for signal-dependent clutter. This technique is then extended to more challenging waveform constraints – in particular, constant modulus and similarity constrained designs. A refinement scheme is devised to improve recovery of vector solutions from the relaxed matrix solution for the constant modulus problem. Finally, we apply the technique to a system model other than WA-STAP; namely, the reverberant sonar/radar channel. Numerical simulations are provided to show the superior performance of our method for SINR, detection, & resolution across multiple environmental scenarios.

Chapter 1

Introduction

As visions of autonomy dance in the minds of technologists across multiple industries, the matter of how to implement and enable such system behavior is still very much an open question. This is especially true in sensing systems like radar that can actively interrogate the environment for information relevant to itself or another decision maker. This active role is mostly unique among sensors and provides the system designer a variety of degrees of freedom to vest extract relevant information about the observed scene. Historically, these designs were done assuming minimal foreknowledge of what awaited the radar and, as such, much attention was focused on the optimal way to observe scattering returns. Even signal design has been traditionally focused on minimizing uncertainty and ambiguity upon reception, assuming that the world in front of a transmitter was a generally uncontrollable wilderness of electromagnetic scattering.

As such, the prospect of joint trans-receive design has historically been approached as a specifically iterative process. A transmit signal is chosen, the optimal receive filter for that signal is designed, the optimal transmit signal for that filter is designed, and so on and so forth. Sometimes this design is based not on pulse shapes, but on transmit spectra from

which the transmitter resource is synthesized. This makes the design problem explicitly multi-valued, since many transmit resources map to the same spectrum. This dissertation operates from the assumption that the *entire* trans-receive process can be synthesized from a given optimal spectrum or basis under given constraints in one simultaneous, rather than relying on alternating synthesis.

1.1 An introduction to Radar and Space-Time Adaptive Processing

In this work, we generally consider the following scenario: assume we have a side-looking airborne radar that transmits a given discrete pulse that repeats a number of times and is transmitted from a single antenna. We assume this array is calibrated through some prior method (see, *e.g.* [1, 2]). The scattered returns of this signal from the environment are received by a known calibrated antenna array. Typically, the goals of this system are twofold. One, maximize the probability of detecting a target at a given location and two, minimize the impact of noise, undesired interference, and non-target signal dependent scattering known as clutter.

As with many other information bearing electromagnetic systems, most notably communications, we can characterize the scattering response from all sources of energy in a channel model, one that can hopefully be linearized. The work of Guerci [3] provides a systematic way to accomplish this modeling for generic MIMO radar systems, for example. With a channel model in mind we can progress and consider means of processing the observed data. The most popular of these in the radar community in the last 30 years is space-time adaptive processing, or STAP. As the name implies, the data is processed by a space-time filter learned from the received non-target data. Said filter is applied to each fast-time range

gate to determine the presence or absence of a target and potentially course estimation of the location of said target, depending on which sample bin it lies within. This is what is known as a post-matched filter process, since it obtains the returns in each range gate by directly matching the temporal response of the environment with the known transmitted signal. Hence, in this case, as can be found in common radar references like [4], the statistics of the clutter are independent of the pulse shape used to observe the environment. Our initial motivating concept, and the primary application discussed in this dissertation, is so-called "waveform-adaptive STAP" (WA-STAP) as proposed by Setlur and Rangaswamy [5, 6, 7, 8]. This process considers the fast-time samples within a range gate. Therefore, the observed data must be processed by a space-*fast time*-slow time adaptive filter instead of the typical space-(slow) time filter, because the clutter response is now obviously dependent on the signals pulse shape. We remark, however, that this signal dependence of the clutter is not unique to WA-STAP. Other system models can and do exhibit signal dependent clutter responses, which we will address later in the dissertation.

1.2 The general problem space

Here, we discuss the historical perspectives and methods used for optimal and adaptive transmission and reception of radar signals. Our major guideposts will be what we term Van Trees' four design paradigms, adapted from [9, Chapter 10.5]. These design issues are as follows: First, we have conventional reception. In this design paradigm, the system designer assumes no knowledge of the clutter statistics and merely attempts to design a signal that will match a known target spectrum the best. This could involve, for example, operating frequency, pulse shape, or transmitted direction. An excellent example of this is a traditional air traffic control tracking radar, which transmits pulses in bands meant to

excite large aircraft and does so in their anticipated direction as the transmitting antenna is mechanically rotated.

The second paradigm is "optimum" reception. This paradigm assumes a competent signal designer has designed a good signal well-matched to the intended target and that the anticipated clutter response to that signal is fully known and well-characterized. This characterization is typically obtained by performing measurement campaigns using the candidate signals, and results in statistical clutter models fitted from the observed data. Given this prior knowledge of the clutter statistics, the optimal receive chain (or filter) is designed to maximize the signal-to-noise ratio, if detection is the primary concern, or some other figure of merit, like range or Doppler estimation error.

Matched transmission, the third paradigm, uses the conventional reception techniques of Paradigm 1. However, prior knowledge of clutter statistics (as in Paradigm 2) is used to design a signal that minimizes the clutter response, potentially at the cost of target gain. This is historically where most signal design efforts have been focused, regardless of the particular parameterization of the signal design. For example, minimizing ambiguity in either delay or Doppler can be regarded as a clutter rejection design where the clutter generating process is the uncertainty in the target process itself.

The final paradigm is that of "optimum" transmission. The most complicated of the four paradigms, this method proposes that foreknowledge of the clutter and other interfering processes is known or otherwise estimated. Then, an "optimum receiver is designed to mitigate these processes as well as having a signal design that minimizes non-target returns. This is the paradigm closest to our contribution. However, we note that Van Trees, and hence nearly every author who followed, consider these two design processes as coupled but separate. That is, a receive filter can only be designed for a given signal and vice versa, but not simultaneously. As mentioned previously, our goal is to join their design together from the outset.

Along with the design issues, Van Trees also articulated three signal design difficulties that, in our view, informed the development of future practical design algorithms. We list these difficulties here, and then describe some modern responses to them.

The first difficulty described is that waveforms that are generated by any design process can be exceedingly complicated. They may have large phase and amplitude changes, varying frequency, and in general not follow canonically studied pulse shapes. Second, if the environment changes, then the signal must also change. Third, all designs are inherently sensitive to the detailed assumptions made about the environment in which the transmitter is used. This can include clutter absorption spectra, nearby large scatterers, present adversarial jammers, etc.

While the intent behind these difficulties remains as relevant as they were when first articulated, advances in signal processing and computation have, in some sense, limited their scope. In terms of complicated waveforms, we note significant advances have been made in digital waveform synthesis, even for radar. For example, the explosion of software-defined radios in the last decade for communications had the knock-on effect of making software-defined *radars* possible [10]. Additionally, even arbitrary analog synthesis is possible within constraints and therefore complicated waveforms inaccessible for a previous age are now viable. See, for example, the extensive work on fully-adaptive radar for tracking and target identification (see, *e.g.* [11, 12]), which can adaptively change baseline parameters (bandwidth, pulse length, and carrier frequency) and transmit pulse shapes through the use of arbitrary function generators.

Second, while a change in environment may indeed necessitate a change in transmission or reception processes, that is the entire purpose of this area of research! Massive computing power permits us to harness decades of prior mapping and collection to either anticipate or identify fluctuations in the scattering environment, and then adaptively control both transmission and reception in response. Therefore, changing a signal in response to a change

in environment may be considered a feature, rather than a bug. Finally, we concede that, like any other system or modeling process, the designs we obtain are only "optimal" as far as our modeling assumptions remain valid. However, this again ties into the fully adaptive radar paradigm, and even earlier concepts such as knowledge-aided processing [13]. That is, while techniques like the one advanced in this dissertation rely on prior channel state knowledge, this can be updated as we sense the environment, learn the environment, and adapt to the environment.

1.3 Early History

As is typical in engineering and the sciences, the concepts we investigate in this document have a significant prior history of other researchers who have come before. Optimal signal design in the presence of dispersive channels historically has been the purview of research in communications. An early example of this is the work of Chalk [14], who derived optimum pulse shapes to mitigate co-channel interference. However, communication systems have a distinct advantage over radar for both transmit and receive design because the monster at the end of the book, so to speak, can talk back. That is, there is a distinct user to and from whom a distinct message is conveyed. Alas, in radar, with the notable exception of air traffic control or identification friend-or-foe (IFF) radars, we do not have the luxury of the targets talking back beyond how they scatter the transmitted signals. Despite this, there exists a wealth of literature on the application of information theory in radar. The earliest of these are summarized in a monograph by P.M. Woodward [15], where he presciently noted that even if an optimal scattering profile is identified, the best means of transmission and reception to obtain that profile "remains substantially unanswered." Other works on information theory followed, in particular the work of Mark Bell [16, 17], that we shall discuss later.

Early applications of signal design and clutter, owing to the technology of the time, focused heavily on analog spectra. For example, in a 1961 technical report, Manasse [18] designed a coding technique for analog pulses in the presence of a clutter power spectral density that is a scaled replica of the transmitted signal, which led to a simple solution: spectrally flat pulses with large time-bandwidth products, like the common linear frequency modulated (LFM) pulses. However, Manasse noted that "[t]his conclusion depends critically on the clutter model" they assumed – a rather simple one, indeed – and in the process, echoed the third design difficulty above. This concept carries forward through other seminal works such as Van Trees' reverberant signal design [19] and the work of Kooij [20] which designs an optimal signal spectrum in the presence of white ambient noise. We will discuss [20] in a further context below. Additional early research in clutter-mitigating signal designs can be found in [21, 22, 23, 24].

In addition to pure signal design, early receiver design considered many of the same paradigms. In a canonical example of adaptive receiver design, Applebaum [25] developed his (later eponymous) array receive filter that maximizes SINR in the presence of disturbances with a known covariance. This paradigm was also applied for detection and clutter by Rihaczek [26], as well as in the work of Bernie Widrow.

The first authors to truly consider sampled signal design (a precursor to digital signal design) to mitigate clutter were the pair of DeLong and Hofstetter [27, 28], and W.D. Rummler [29]. These researchers independently developed a technique that maximizes the SIR under either the target matched filter or the optimum clutter filter. We note, however, that DeLong and Hofstetter did not necessarily restrict themselves to coding a pulse train as Rummler did, though as we will see below the mathematical representation of both is remarkably similar. We will follow and describe DeLong and Hofstetter here since it is, in our opinion, a clearer description of the design concepts we use today, despite being initially designed for classical single input-single output (SISO) radar. Their derivation starts in the analog

domain and focuses heavily on the ambiguity function of the transmitted signal which covers both range and Doppler dependence. They also make the (now seemingly obvious) insight that maximization of SINR, subject only to energy constraints, is operationally suboptimal. Despite this, they recognize that such simplifications are often necessary in the design of the transmit signal and/or receive filter for computational purposes. This is where they arrive at the idea of designing a sampled signal, as the vectorized form permits the use of linear algebra techniques. Their iterative technique for either matched or clutter filter relies on the temporal symmetry property of a signals cross ambiguity function. That is, they exploit the unique SISO property that interchanging the transmit signal and the receive filter to their opposite roles does not change the SINR. In fact, these authors were also the first to identify the quadratic dependence of the clutter covariance on the signal.

1.4 Previous Methods

The process of signal design or joint signal-filter design has historically followed one of two main tracks. The first we term iteration. Here, given a figure of merit (usually SINR) and a preferred detection scheme, these algorithms iterate on a sequence of signals alone. They may employ either a target matched filter or the so-called optimum filter of Van Trees' Second Paradigm. What follows are some notable examples of this process. First, the aforementioned works of DeLong and Hofstetter iterate on a sequence of minimum whitened eigenvectors of the clutter cross-ambiguity function and consider both target-matched and optimum filters. Much later, the works of Pillai and others [30, 31, 32] considered an alternating minimization of the Fredholm integral eigenfunction and the optimum filter for both SISO and MIMO configurations. However, their signal sequence, as one might expect from eigenfunction optimization, is in a purely functional continuous-time form. In a series of papers from earlier this century, Friedlander [33, 34] considered alternating maximization of

the target to clutter ratio, enabled by gradient descent of what he termed a non-linear SINR cost function. He also made explicit comparisons with target matched signals, clutter minimizing signals, and target maximizing signals. Chen and Vaidyanathan [35] considered the sequential analysis based on minimizing eigenvectors in the context of MIMO radar detection of extended and temporally-correlated spatial channels. Unlike the work of Friedlander [34], these methods guarantee non-decreasing SINR at each subsequent iteration. We will revisit some conclusions of this work in Chapter 2 in our discussion of biconvex optimization and guarantee convergence of alternating minimization. Finally, the work of Stoica, He, and Li [36] return to gradient descent as applied to reverberant channel environments (that is, mere temporal shifts). This work also considers waveform constraints most notably, constant modulus and peak-to-average power ratio (PAPR) which, in their view, requires the use of gradient descent in their CREW(gra) algorithm. While they show monotonic convergence, due to their technique, it is not guaranteed. These authors also investigate a relaxation-then-synthesis technique, the design philosophy of which we will describe subsequently. A more recent example of this design philosophy can be found in [37] and [38]. Both papers investigate constant-modulus and similarity-constrained colocated MIMO filter-signal design. In [37], assuming an optimum filter, Aldayel, et al. solve, at each iteration, a quadratically-constrained quadratic program (QCQP) that maximizes the target response. Within this procedure, they employ a clever bisection refinement technique to successively enforce the phase-only similarity constraint through a linear operator. While following the iterative signal-only design philosophy, Cui, et al. [38] take a far different approach. These authors employ fractional programming and a Dinklebach-type algorithm to minimize the clutter response, generating a new signal at each step. They subsequently form the filter from the maximum generalized eigenvector of the target and noise-clutter covariance matrices.

The second major design process can best be described as relax, then iterate. By this we mean these works determine an optimal signal spectrum, either the power spectral density or auto-correlation/ambiguity function, then synthesize a signal from that spectrum then

designs a filter. The first notable example of this we have already seen in the work of Kooij [20], which designs the optimum frequency spectrum in the presence of white noise and scaled clutter for a classic monostatic radar. A more recent canonical example comes from the work of Kay [39, 40], who generalizes the work of Pillai, Kooij, and Manasse to general noise and interference spectra. In order to synthesize waveforms from the optimally designed spectra, Kay uses Durbins method to interpolate samples onto a given train of pulses. Furthermore, Kay was one of the first to demonstrate that the well known water filling solution (transmission should occur preferentially in bands with low noise and/or interference subject to their alignment with the desired target response) applies to radar as much as communications. This portion of his result has been confirmed from an information theoretical standpoint by Romero and Goodman [41] as well as Setlur, et al. [42]. As a preview for Chapter 3, we note that our perspective confirms the water filling nature of optimal trans receive design. The work of Aubry et al. [43, 44] is another contribution in the relax, then iterate family of algorithms. These works design a SIMO transmit pulse code through iterative relaxation of the auto correlation matrix, followed by a matched eigenfilter, which is repeated until convergence. Further analysis of these works will follow in Chapter 2, where we will demonstrate that these techniques cannot be generalized to pulse shape design, and have a critical flaw in their design assumptions. A more recent deployment of this design strategy is that of Cui, Li, and Rangaswamy [45], which was employed for colocated MIMO radars. The primary purpose of relaxation here is, as with [36], meant to provide a simpler means of incorporating modulus and waveform similarity constraints. They considered two related sequential algorithms which maximized different forms of the target response as a function of the relaxed signal spectrum – one directly based on concentration of the cost function given the optimum filter and one that depends on a sequentially-optimized filter-signal pair. In both cases, the algorithms in [45] rely on randomization to generate an optimal signal from the resulting spectrum at each step. This technique is also applied in recent works that consider uncertainty in the prior knowledge of the target or clutter

channel, like that in [46]. In these cases, the relaxation occurs not just in the signal design step, but in the filter design step as well; hence, at both stages, a semidefinite relaxation (and a potential randomization step) form the optimization and each signal and filter must be synthesized therefrom.

1.5 Research Motivation and Inspiration

This dissertations research is an outgrowth of a third design process that considers the waveform-filter design problem jointly; that is, the receive filter is intentionally designed to be jointly optimal with the transmitted signal rather than a free design parameter or merely a vector for further signal design. This process is embodied by two recent papers which take their cues from the previous two design processes above. The first of these is the work of Tang and Tang [47] who consider both alternating signal filter designs and relaxed signal filter designs for MIMO STAP. Under the relaxation paradigm, their techniques are very reminiscent of those of Aubry et al. [43, 44], and under the sequential concept, they employ power-like iterations, including the SWORD algorithm of Li, Guerci, and Xu [48]. The other is the work of Setlur and Rangaswamy [6, 8] which considers the waveform adaptive STAP above. Primarily, their technique alternates between an optimal design of the signal then beamformer and so on in the simplest case. This is also applied to a constant modulus design where the signal is synthesized exactly without randomization and relaxation. Additionally, they consider a more robust proximal method that can be used when the signal independent noise and interference covariance matrix is rank-deficient. This is unique in the literature, as many of the aforementioned authors do not consider the impact of poor estimation of channel parameters. Finally, this work directly demonstrated the non-convexity of the general problem computationally.

In light of this work and the other previous literature we are left with the following questions: First, is an alternating signal-then-filter algorithm the only way to obtain jointly optimal trans-receive pairs? Second, is this problem provably non-convex, and if so, is there any useful structure within the non-convexity that can be exploited? Third, does this problem, either in its original or relaxed forms, resemble any known problem types in the optimization literature? Finally, if this fits any known problem types, can convex optimization, possibly enabled by relaxation, be useful beyond signal-only design? The rest of this work will answer the final three questions in the affirmative, while demonstrating that the first question is decidedly false.

1.6 The Joint Design Problem & Related Forms

Since our motivating example is STAP, we consider variations on the following optimization problem throughout the rest of this dissertation: Minimize the overall variance of the signals received undesired sources such as noise, jamming, and other interferers both signal dependent and independent, such that the gain from a given direction within a particular arranged bin is fixed, and the transmitted pulse has an upper bound on its power.

Owing to the "variance" aspect, this type of problem is typically associated with convex quadratic optimization when either signal or filter is fixed and the other is chosen as the design parameter. However, when the joint problem is considered, this becomes, naturally, a biquadratic problem.

The biquadratic problem has been mildly investigated in the literature by, among others, [49, 50, 51, 52, 53, 54, 55]. Naturally, this is a highly non-convex problem and so guarantees about solutions are nearly nonexistent. However, as the problem is (convex) quadratic in each variable, there is some hope for sensible bounds and approximations. This biconvex

property has significant implications for any alternating minimization techniques, whether or not a relaxation is used, the details of which will be explored in Chapter 2.

Another potential means of attack is considering a joint problem – that is, the signal and filter are each considered projections of a single design parameter. Hence, the problem becomes a quartic optimization, similar to the form seen in [36]. Again, this form is obviously non-convex, but as we will show in Chapter 3 and beyond, it admits a convex matrix relaxation, which can be solved with conventional interior-point methods. This relaxation is itself a special type of optimization problem known as a quadratic semidefinite program (QSDP), which has some recent theory that we can employ to understand the joint problem in further detail; see, *e.g.* [56] and references therein. We will also show that due to the separability of the constraints, we can solve some of the Karush-Kuhn-Tucker (KKT) conditions to gain insight into the solutions obtained from the relaxation. It is here we note that the use of QSDPs for radar systems design is not entirely novel – Stoica, Li, and Xie [57] used a far simpler version of it for designing probing signals for MIMO radar over a decade ago – but our application to joint trans-receive design is, to our knowledge, the first such instance anywhere.

1.7 Organization of the Dissertation

With our motivations and the state of the prior art outlined, we now describe the organization of the rest of the dissertation, as well as the contributions described within.

Chapter 2 investigates the most generic forms of the considered problem and the most common design practice used to tackle them: alternating minimization. Our contribution is the first systematic comparison of the multiple design philosophies that fall under this

paradigm, as well as introduce the notion that monotonic convergence is guaranteed so long as the iterates contribute to the next step's constraint set.

Chapter 3 introduces the use of quadratic semidefinite programming as a technique to avoid explicitly iterative processes (like the ones explored in Chapter 2) and jointly optimize the trans-receive chain as a single entity, given foreknowledge of the channel. We consider a simplified power-inequality constraint and directly analyze the KKT conditions for further insight. This is also supplemented by more general QSDP theorems that can be applied to any other set of waveform constraints. Additionally, we can show that, in this case, there is a direct connection to the "waterfilling" analysis by [40].

Chapter 4 extends the work of Chapter 3 by incorporating two common waveform constraints – constant modulus and "similarity." We find that these constraints are easily incorporated into our relaxation framework and in many cases, we can obtain respectable approximate solutions in the same way as in the previous chapter for both constraint sets. In the cases where there is difficulty, we examine the need for further improvement in rank-one approximation, enabled by a simple refinement strategy.

While most of the work in this dissertation is focused on application to the waveform-adaptive STAP introduced by Setlur and Rangaswamy, Chapter 5 considers the extension of the QSDP method to other active sensing models that consider reverberant channels or signal-dependent interference and clutter. Thus, this chapter describes other relevant models and relation to the work of Chapters 3 and 4, and demonstrates that techniques described in the previous two chapters can be directly applied with minimal modification to great effect.

Chapter 6 will conclude the dissertation, providing summarizing thoughts and caveats, as well as a prospectus for future expansion of this work.

Chapter 2

Preliminaries and Iterative Designs

In the previous chapter, we provided an overview of the most common modern strategies and algorithms for joint trans-receive adaptivity – namely, explicitly iterative techniques that alternately optimize the signal and the beamformer/filter. These techniques have, for the most part, limited convergence analyses and, in any case, often neglect to make significant connections to the optimization literature. This chapter will provide that foundation, examining the connection of the nominal problem to the so-called biquadratic program, as well as additional mathematical rigor to show that any appropriately constructed alternating algorithm does indeed converge.

We will also propose a useful equivalent reduction of the power-bounded problem, one relatively generalizable to the waveform-constrained extensions we will pursue in Chapter 4. This reduction will provide critical insight necessary for the relaxation approaches in Chapters 3 and 4, and also demonstrate why some competing algorithms can have significant performance issues.

The rest of the chapter is organized as follows: Chapter 2.1 will discuss the general radar model that we will consider throughout the rest of the dissertation, as well as the overall design

problem. Chapter 2.2 demonstrates the connection of the design problem to biconvex and biquadratic programming, and discusses some of the geometrical aspects of the problem. In Chapter 2.3, we will provide a more complete overview of the existing alternating paradigms in light of this theory, including a brief discussion of the basics of linear semidefinite relaxation. Chapter 2.4 provides the aforementioned reduction and a general proof of convergence for any alternating minimization technique where the sets vary at each iteration, as is the case for the joint design problems existing in the literature. Finally, Chapter 2.5 provides some simulation results to demonstrate the convergence properties of the previous section, as well as demonstrate that the nominal alternating minimization of Setlur & Rangaswamy [8], and our extensions thereof, is a good baseline.

2.1 Active Sensing Models and Joint Trans-Receive Design

So far in this dissertation, we have discussed the differences between radar and other active sensing systems *qualitatively*. As engineers, however, mere qualitative description and analysis is inherently unsatisfying, and thus we require a more mathematical treatment of the systems under consideration. In this section, we shall outline first outline the radar model for our primary motivating application, waveform-adaptive STAP (WA-STAP). We will also briefly contrast the operation of WA-STAP with traditional STAP. Next, we will comment on the connection to a general discrete-time channel model for radar signal processing and how they may be generated by stochastic Green's functions, after the work of Guerci [3]. Finally, we will motivate the primary problem in this work, that of joint signal-filter design, by discussing the mechanics of detection in STAP and thus the necessary cost functions we will examine throughout.

2.1.1 Radar Fundamentals

We being by discussing the general contours of active RF sensing systems, and radar in particular. All radar systems are predicated on the simple concept of transmitting an electromagnetic signal through a number of antennas into a given scattering environment, then processing the scattered returns (or echoes) from said environment in response to the initial excitation, as received by another number of antennas. Depending on the strength and locality of these scatterers, the time scales examined, and the available degrees of freedom used to transmit and/or receive these echoes, we are then able to detect the presence or absence of the scatterers and extract the familiar localizing parameters that characterize their existence in a scene – range, Doppler shift/velocity, azimuth and elevation, etc.

In this work, we consider pulse-Doppler radar systems, which transmit a number of pulses over a finite window of time to excite the environment. This pulse-and-wait process forms the backbone of these systems' ability to simultaneously measure both range-to-target and the Doppler shift/target velocity. This is in contrast to continuous wave (CW) radars, which, as the name implies, continually transmit some waveform and use modulation of that waveform in some fashion to return both range and Doppler information. While CW systems are once again in vogue for numerous short-range applications (see, e.g., the radars used in Major League Baseball's Statcast$^{\text{TM}}$system [58]), we shall not examine them further here.

For now, let us describe the general radar model we will used throughout the rest of this dissertation; for the most part, we follow the notation and practice of [8]. Consider a radar that consists of a calibrated linear array of M identical sensor elements, where the first element is the phase center and also acts as the transmitter. Hence, we notionally consider what is known as a single-input multiple-output (SIMO) monostatic radar. We will show later, however, that our proposed process is amenable to other transmit/receive configurations, like colocated multiple-input multiple-output (MIMO) radars. As seen in

Figure 2.1 [8], we assume that the array's phase center is located at the physical location $\mathbf{x_r} = [x_\mathrm{r}\ y_\mathrm{r}\ z_\mathrm{r}]^T \in \mathbb{R}^3$ in the given Cartesian coordinate space at a particular time instant, and moves with a velocity vector $\dot{\mathbf{x}}_\mathbf{r} = [\dot{x}_\mathrm{r}\ \dot{y}_\mathrm{r}\ \dot{z}_\mathrm{r}]^T \in \mathbb{R}^3$ (with velocity magnitude v_r). We leave the description of the target depicted for later.

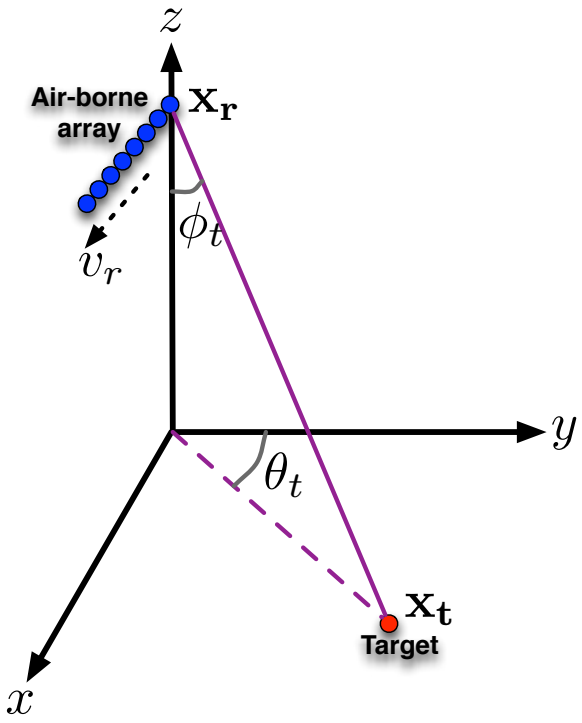

Figure 2.1: Simplified illustration of airborne radar scene considering ground target located at azimuth θ_t and elevation ϕ_t relative to the platform. ©2016 IEEE. Reprinted, with permission, from Setlur and Rangaswamy, IEEE Trans. Sig. Process., vol. 64, pp. 19–34 (2016) [8]

During the transmission period, the radar probes the environment with a number of pulses $s(t)$ of width T seconds and bandwidth B Hz at a carrier frequency f_o. Within each burst, we transmit L pulses at a rate of f_p (i.e. the pulses are transmitted every $T_p = 1/f_p$ seconds) and collect them in a coherent processing interval (CPI). Hence, the overall transmitted signal at a given point in time $t \in (0, T]$ is given by

$$u(t) = \sum_{l=0}^{L-1} s(t - lT_p) e^{\jmath 2\pi f_o (t - lT_p)}. \tag{2.1}$$

Observe that we do not directly impose any additional coding or shaping of the pulse; if it does occur, i.e. for spectral control purposes, then it is encoded into $s(t)$ directly.

As seen in Figure 2.1, we assume that the probed environment contains a non-fluctuating target located at the coordinate vector $\mathbf{x_t} \in \mathbb{R}^3$, moving with a velocity vector $\dot{\mathbf{x}}_\mathbf{t} = [\dot{x}_\mathrm{t}\ \dot{y}_\mathrm{t}\ \dot{z}_\mathrm{t}]^T \in \mathbb{R}^3$ (with velocity magnitude ν_t). Let us define the relative position and velocity vectors $\boldsymbol{\delta}\mathbf{x} = \mathbf{x_r} - \mathbf{x_t}$ and $\boldsymbol{\delta}\dot{\mathbf{x}} = \dot{\mathbf{x}}_\mathbf{r} - \dot{\mathbf{x}}_\mathbf{t}$, respectively. These positions correspond to the platform-relative azimuth θ_t and elevation ϕ_t, which are related to the relative position vector by

$$\boldsymbol{\delta}\mathbf{x} = R_t[\sin(\phi_t)\sin(\theta_t)\ \ \sin(\phi_t)\cos(\theta_t)\ \ \cos(\phi_t)]^T, \tag{2.2}$$

where $\|\cdot\|$ is the standard Euclidean norm, and $R_t = \|\mathbf{x_r} - \mathbf{x_t}\|$ is the slant range to the centroid of the target.

Considering only the target, the received signal at the mth array element for the lth pulse is

$$s_{ml}(t) = \rho_t s(t - lT_p - \tau_m)e^{j2\pi(f_o + f_{d,m})(t - lT_p - \tau_m)} \tag{2.3}$$

where τ_m is the round-trip time delay, $f_{d,m}$ is the observed Doppler shift induced by the relative motion, and ρ_t is a complex scattering coefficient that is a function of physical target characteristics. If, as in Figure 2.1, the array is located along the platform's local x-axis, the reflected signal returns to a location $\mathbf{x_r} + m\mathbf{d}$, where $\mathbf{d} = [d\ 0\ 0]$ when it impinges on the mth element ($m \in \{0, 1, \ldots, M-1\}$)[1] Hence, the position differential we must consider for the return half of the target echo is $\mathbf{x_r} + m\mathbf{d} - \mathbf{x_t} = \boldsymbol{\delta}\mathbf{x} + m\mathbf{d}$. Then, the time delay τ_m is

$$\tau_m = \frac{R_t}{c} + \frac{\|\boldsymbol{\delta}\mathbf{x} + m\mathbf{d}\|}{c}. \tag{2.4}$$

[1]We define the phase center to be the so-called zeroth element. The fifth element we define to be Leeloo.

In most practical radar cases [4], especially airborne radars, the interelement spacing (usually fractions of a wavelength) is far smaller than the slant range, that is $d \ll R_t$. Under this assumption, we can approximate the delay as

$$\tau_m \approx \tau_t + \frac{m\mathbf{d}^T\boldsymbol{\delta}\mathbf{x}}{c} = \tau_t + \frac{m\,d\,\sin(\phi_t)\sin(\theta_t)}{c}. \tag{2.5}$$

where we have defined $\tau_t = 2\frac{R_t}{c}$ as the bulk target delay. While the term dependent on m in (2.5) can be significant in certain cases, it is rendered moot if the common narrowband assumption is invoked – that is, if we assume the time it takes for the signal wavefront to completely impinge upon the array is far smaller than the pulse length itself [4]. Hence, in the narrowband case, $\tau_m \approx \tau_t$ for all m.

Using these assumptions, we can also approximate the Doppler shift $f_{d,m}$ as being independent of the element index; that is, $f_{d,m} \approx f_d$ for all m where

$$f_d = 2f_o\frac{(\boldsymbol{\delta}\dot{\mathbf{x}})^T\boldsymbol{\delta}\mathbf{x}}{cR_t} = 2f_o\frac{[\sin(\phi_t)\sin(\theta_t)\ \ \sin(\phi_t)\cos(\theta_t)\ \ \cos(\phi_t)](\boldsymbol{\delta}\dot{\mathbf{x}})}{c}. \tag{2.6}$$

Given the approximations in (2.5) and (2.6), once the target return signal is downmixed to baseband, we have

$$s_{ml}(t) = \rho_t s(t - lT_p - \tau_t)e^{\jmath 2\pi(f_o+f_{d,m})(t-lT_p-\tau_m)} \tag{2.7}$$

where we have absorbed constant phase terms into ρ_t and have further assumed that the Doppler phase is insignificant on a fast-time scale.

Let us now assume that we discretely sample the pulse $s(t)$ into N samples, resulting in the sample vector $\mathbf{s} = [s_1, s_2, \cdots, s_N]^T \in \mathbb{C}^N$, and the data is aligned to a common reference. At the target range gate τ_t, the combined target response over the entire CPI can be represented

by a vector $\mathbf{y}_t \in \mathbb{C}^{NML}$ given by

$$\mathbf{y_t} = \rho_t \mathbf{v}_t(f_d) \otimes \mathbf{s} \otimes \mathbf{a}_t(\theta_t, \phi_t) \tag{2.8}$$

where ρ_t is defined above, the vector $\mathbf{v}_t(f_d) = [1 \; e^{-\jmath 2\pi f_d T_p} \; \cdots \; e^{-\jmath 2\pi f_d (L-1) T_p}]^T \in \mathbb{C}^L$ is the Doppler steering vector, and the vector $\mathbf{a}_t(\theta_t, \phi_t) = [1 \; e^{-\jmath 2\pi \vartheta_t} \; \cdots \; e^{-\jmath 2\pi (M-1)\vartheta_t}]^T \in \mathbb{C}^M$ is the spatial steering vector where $\vartheta_t = d \sin(\theta_t) \sin(\phi_t) f_o / c$ is the spatial frequency. We note that this particular arrangement of the Kronecker products in (2.8) arises from stacking first the array response at each time sample n, then the combined fast time-spatial response for each pulse. One could consider other arrangements of the response, but they are functionally identical up to a permutation matrix.

It is here we point out that the presence of fast-time samples $\mathbf{s}$ in $\mathbf{y_t}$ is the heart of the WA-STAP model, and differentiates it from the standard STAP scenario. Traditional STAP operates on spatio-Doppler responses from the radar after matched filtering of the array response with the pulse shape [59, 60]; that is, the modeled returns are constant within a given range bin and the response vector depends only on the angle and Doppler. This gives rise to the traditional response vector $\mathbf{y}_{\mathbf{t},\text{STAP}} = \rho_t \mathbf{v}_t(f_d) \otimes \mathbf{a}_t(\theta_t, \phi_t)$. However, recent advances in the literature have considered incorporating fast-time data for more accurate clutter modeling, both transmitter- and jammer-induced [61, 62]. Furthermore, more recent MIMO STAP papers (see, e.g. [47]) have also adopted the WA-STAP model to account for pulse design choices. The net impact on the traditional STAP datacube is reflected in Figure 2.2 [8]. Essentially, each range bin is now further subdivided into N sub-"bins" – however, as we will soon see, we do not process these sub-bins individually, but as a collection.

The inclusion of fast-time information is mandatory for performing any realistic signal design because, as mentioned in Chapter 1, the response of any scatterer is highly dependent on

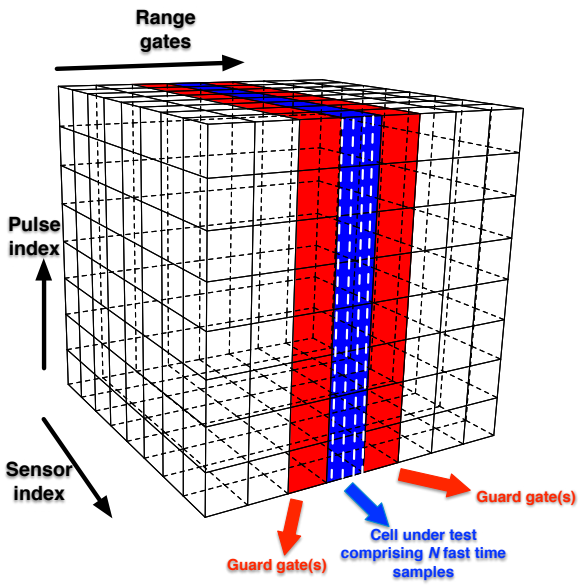

Figure 2.2: Waveform Adaptive STAP Datacube, prior to matched filtering/range compression. The dashed lines indicate the N subdivisions of the range gate under consideration. ©2016 IEEE. Reprinted, with permission, from Setlur and Rangaswamy, IEEE Trans. Sig. Process., vol. 64, pp. 19–34 (2016) [8]

the pulse shape and thus the pulse's spectral content. This is especially true for monostatic SIMO applications, whose sole degrees of transmit design freedom is that pulse shape.

While we would like the response from any target to return to the array as cleanly as possible, in reality, $\mathbf{y}_{\mathtt{t}}$ is corrupted by a variety of undesired returns from the environment – noise, signal-independent interference, and clutter:

$$\tilde{\mathbf{y}} = \mathbf{y}_{\mathtt{t}} + \mathbf{y}_{\mathtt{n}} + \mathbf{y}_{\mathtt{i}} + \mathbf{y}_{\mathtt{c}} = \mathbf{y}_{\mathtt{t}} + \mathbf{y}_{\mathtt{u}}, \tag{2.9}$$

where the subscripts n, i, c, and u stand for noise, interference, clutter and undesired, respectively. We assume that these undesired energy sources are statistically uncorrelated from each other. We shall subsequently describe each of these sources, starting with the noise.

We assume the noise is zero mean and identically distributed across all sensors, pulses, and fast time samples. The covariance matrix of $\mathbf{y}_n$ is denoted $\mathbf{R}_{\mathtt{n}} \in \mathbb{C}^{NML \times NML}$. The simplest model is the traditional white noise model, given by $\mathbf{R}_{\mathtt{n}} = \sigma_{\mathtt{n}}^2 \mathbf{I}_{NML}$, where $\sigma_{\mathtt{n}}^2$ is the assumed noise variance. However, due to the pulse-shape dependence of the model, $\mathbf{R}_{\mathtt{n}}$ may in fact exhibit a more predictable structure; in particular, a Toeplitz or circulant form may be appropriate to model the correlation of the noise on a pulse subsample-by-pulse subsample basis.

The interference term consists of K known interference sources that correspond to a zero-mean random process spread over all pulses and fast time samples. Assume that the kth interferer is located at the azimuth-elevation pair (θ_k, ϕ_k). In the lth PRI, we assume that the waveform is a complex continuous-time signal $\alpha_{kl}(t)$ which, when sampled, yields the vector $\boldsymbol{\alpha}_{kl} \in \mathbb{C}^N$, similar in form to $\mathbf{s}$. Stacked across all PRIs, we obtain the random vector $\boldsymbol{\alpha}_k = [\boldsymbol{\alpha}_{k0}^T \, \boldsymbol{\alpha}_{k1}^T \, \cdots \, \boldsymbol{\alpha}_{k(L-1)}^T]^T \in \mathbb{C}^{NL}$, whose covariance matrix we define as $\mathbf{E}\{\boldsymbol{\alpha}_k \boldsymbol{\alpha}_k^H\} = \mathbf{R}_{\boldsymbol{\alpha},k}$. Then, the response from the kth interferer can be modeled as

$$\mathbf{y}_{\mathtt{i},k} = \boldsymbol{\alpha}_k \otimes \mathbf{a}_{\mathtt{i}}(\theta_k, \phi_k) \tag{2.10}$$

where, as with the target, $\mathbf{a}_{\mathtt{i}}(\theta_k, \phi_k)$ is the array response to the interferer. The covariance of $\mathbf{y}_{\mathtt{i},k}$ is then $\mathbf{R}_{\boldsymbol{\alpha},k} \otimes \mathbf{a}_{\mathtt{i}}(\theta_k, \phi_k)\mathbf{a}_{\mathtt{i}}(\theta_k, \phi_k)^H$. As above, we may consider this covariance to be white, i.e., $\mathbf{R}_{\boldsymbol{\alpha},k} = \sigma_{\mathtt{i}}^2 \mathbf{I}_{NL}$, but due to the sampling structure employed for WA-STAP, a temporal correlation model may be more appropriate. If we assume the interferers are uncorrelated with each other, then the overall covariance of the combined signal-independent interference $\mathbf{y}_{\mathtt{i}}$ is

$$\mathbf{R}_{\mathtt{i}} = \sum_{k=1}^{K} \mathbf{R}_{\boldsymbol{\alpha},k} \otimes \mathbf{a}_{\mathtt{i}}(\theta_k, \phi_k)\mathbf{a}_i(\theta_k, \phi_k)^H. \tag{2.11}$$

In future sections, we will use the notation $\mathbf{R}_{\mathtt{ni}} = \mathbf{R}_{\mathtt{n}} + \mathbf{R}_{\mathtt{i}}$ to denote the combined noise and interference covariance matrix.

In airborne radar applications, the most significant clutter source is the ground, which produces returns persistent throughout all range gates up to the horizon. Though other clutter sources exist, like large discrete objects, vegetation, and targets not currently being surveilled, the specific stochastic model we apply here only concerns ground clutter. However, we note that a later formulation *may* be amenable to considering those other sources in a manner that recalls efforts in the literature on channel estimation (see, for example, [3]). Regardless of the source, a major assumption in all works on signal-dependent interference mitigation is that there is some prior knowledge of the clutter's response to a given input signal. While this can lead to a "chicken-or-egg" scenario, as mentioned by [8], such information can be obtained in a variety of ways: physics-based scattering models, estimation from topographical information and/or knowledge-aided databases (as in DARPA's KASSPER program [13]), or occasional direct probing and estimation of the environment with "known-good" pilot signals (see, e.g., [3]). In this work, we will consider an idealized physics-based approach, but note this does not limit our method's applicability to more realistic scenarios.

For now, assume the presence of Q clutter patches, each comprising P distinct scatterers. As with the target, the return from the pth scatterer in the qth patch, located spatially at the azimuth-elevation pair (θ_{pq}, ϕ_{pq}), maintains a Kronecker structure given by

$$\gamma_{pq}\mathbf{v}(f_{\mathsf{c},pq}) \otimes \mathbf{s} \otimes \mathbf{a}(\theta_{pq}, \phi_{pq})$$

where the returned complex reflectivity is γ_{pq} and the Doppler shift observed by the platform is $f_{\mathsf{c},pq}$. This Doppler shift is solely induced by the platform motion, characterized by the aforementioned velocity vector $\dot{\mathbf{x}}_{\mathrm{r}}$, and is given by

$$f_{\mathsf{c},pq} = 2f_o \frac{\dot{\mathbf{x}}_{\mathrm{r}}^T[\sin(\phi_{pq})\sin(\theta_{pq})\ \sin(\phi_{pq})\cos(\theta_{pq})\ \cos(\phi_{pq})]}{c}. \tag{2.12}$$

Thus, the overall response from the qth clutter patch is

$$\mathbf{y}_{\mathsf{c},q} = \sum_{p=1}^{P} \gamma_{pq} \mathbf{v}(f_{\mathsf{c},pq}) \otimes \mathbf{s} \otimes \mathbf{a}(\theta_{pq}, \phi_{pq}). \tag{2.13}$$

In order to define the covariance matrix of the clutter, we define the qth combining matrix $\mathbf{B}_q \in \mathbb{C}^{NML \times P}$ as the matrix whose pth column is $\mathbf{v}(f_{\mathsf{c},pq}) \otimes \mathbf{s} \otimes \mathbf{a}(\theta_{pq}, \phi_{pq})$, $i \in \{1, \ldots, P\}$ and the covariance of the reflectivity vector $[\gamma_{1q}\, \gamma_{2q} \cdots \gamma_{Pq}]^T$ given by $\mathbf{R}_{\boldsymbol{\gamma}}^{pq} \in \mathbb{C}^{P \times P}$. Then, the overall covariance matrix for the patch is

$$\mathbf{R}_{\boldsymbol{\gamma}}^{q} = \mathbf{B}_q \mathbf{R}_{\boldsymbol{\gamma}}^{pq} \mathbf{B}_q^H \tag{2.14}$$

If we assume that the scatterers in one patch are uncorrelated with the scatterers in any other patch, then the total clutter response is $\mathbf{y}_{\mathsf{c}} = \sum_{q=1}^{Q} \mathbf{y}_{\mathsf{c},q}$ and its covariance is given by

$$\mathbf{R}_{\mathsf{c}} = \sum_{q=1}^{Q} \mathbf{R}_{\boldsymbol{\gamma}}^{q} \tag{2.15}$$

We will denote the overall undesired response covariance matrix as $\mathbf{R}_{\mathsf{u}}(\mathbf{s}) = \mathbf{R}_{\mathbf{ni}} + \mathbf{R}_{\mathsf{c}}(\mathbf{s})$.

Since this is a rather cumbersome model, we can simplify our description of the clutter as follows. Let us assume that our range resolution is large enough that we cannot resolve individual scatterers in each patch – as mentioned in [8], this is typical in STAP applications. Thus, we can regard each scatterer in the patch as lying within the same range gate and having approximately equal Doppler shifts, hence $f_{\mathsf{c},pq} \approx f_{\mathsf{c},q}$. Similarly, if we assume far-field operation, the scatterers will lie in approximately the same angular resolution cell centered at (θ_q, ϕ_q), which means $\theta_{pq} \approx \theta_q, \phi_{pq} \approx \phi_q$. Given this simplification, we can modify our representation of the patch response and its covariance.

Under this assumption, the per-patch clutter response is

$$\mathbf{y}_{\mathsf{c},q} = \gamma_q \mathbf{v}(f_{\mathsf{c},q}) \otimes \mathbf{s} \otimes \mathbf{a}(\theta_q, \phi_q)$$

where $\gamma_q = \sum_{p=1}^{P} \gamma_{pq}$ is the combined reflectivity of all scatterers within the patch. For the covariance, the combining matrix for the q-th clutter patch is given by $\mathbf{B_q} = [\mathbf{v}(f_{\mathsf{c},q}) \otimes \mathbf{s} \otimes \mathbf{a}(\theta_q, \phi_q), \ldots, \mathbf{v}(f_{\mathsf{c},q}) \otimes \mathbf{s} \otimes \mathbf{a}(\theta_q, \phi_q)]$. Since the deterministic patch response is repeated P times, via standard Kronecker product properties, this is equivalent to

$$\mathbf{B_q} = \mathbf{1}_P^T \otimes \mathbf{v}(f_{\mathsf{c},q}) \otimes \mathbf{s} \otimes \mathbf{a}(\theta_q, \phi_q). \tag{2.16}$$

More importantly, since $\mathbf{R}_{\boldsymbol{\gamma}}^q = \mathbf{B}_q \mathbf{R}_{\gamma}^{pq} \mathbf{B}_q^H$, we have

$$\mathbf{R}_{\gamma}^q = \overline{R}_{\gamma}^q (\mathbf{v}_q \otimes \mathbf{s} \otimes \mathbf{a}_q)(\mathbf{v}_q \otimes \mathbf{s} \otimes \mathbf{a}_q)^H \tag{2.17}$$

where $\overline{R}_{\gamma}^q = \mathbf{1}_P^T \mathbf{R}_{\gamma}^{pq} \mathbf{1}_P$ can be viewed as the total variance of the clutter reflectivity within the patch.

2.1.2 General Channel Models

In [63, 3], the authors consider a transfer function/matrix approach for simultaneous transmit and receive resource design in MIMO radar, similar to the typical literature on control theory and digital communications. These are predicated on identifying so-called stochastic Green's functions, which are signal-independent characterizations of the scattering properties of objects in the environment, and can be reasonably precalculated. In general, let us assume we have an active sensing system with M_T transmit resources arranged in the vector $\mathbf{s}$ that are transduced through the environment to M_R receive resources into a received vector $\mathbf{y}$, where $M_R \geq M_T$. If we assume that the transmit and receive resources are sufficiently

"independent" degrees of freedom then, for a given scatterer, we can identify a matrix $\mathbf{H}_{\mathrm{S}} \in \mathbb{C}^{M_R \times M_T}$ such that the response $\mathbf{y}_{\mathrm{S}} = \mathbf{H}_{\mathrm{S}}\mathbf{s}$. We can think of this as a generalization of traditional linear time-invariant (LTI) system theory, as each element of $\mathbf{H}_{\mathrm{S}}$ encodes a linear mapping between each transmit resource and each receive resource. If there are a number of scatterers, say S_{total}, each with their own channel response $\mathbf{H}_s$, the overall response to the input to the environment is given by

$$\tilde{\mathbf{y}} = \sum_{s=1}^{S_{\text{total}}} \mathbf{H}_s\mathbf{s} + \mathbf{y}_{\mathtt{ni}}, \tag{2.18}$$

where $\mathbf{y}_{\mathtt{ni}} \in \mathbb{C}^{M_R}$ is a signal-independent noise and interference term that operates similar to the above. A key to identifiability in these models is having an effective characterization or template for distinguishing "target"-like channel matrices from "clutter"-like channel matrices, a distinction that can vary depending on the application.

Using general properties of Kronecker products, we can reframe our model above to fit the channel model, as in [3]. For example, we can write the ideal response of the qth clutter patch as

$$\begin{aligned}\mathbf{v}(f_{\mathtt{c},q}) \otimes \mathbf{s} \otimes \mathbf{a}(\theta_q, \phi_q) &= (\mathbf{v}(f_{\mathtt{c},q}) \otimes \mathbf{I}_N \otimes \mathbf{a}(\theta_q, \phi_q))\mathbf{s} \\ &= \mathbf{\Gamma}_q\mathbf{s}\end{aligned}$$

where the clutter channel matrix $\mathbf{\Gamma}_q$ is implicitly defined. By the same token, the target response is equivalent to $\mathbf{Ts}$ where $\mathbf{T} = \mathbf{v}_t(f_d) \otimes \mathbf{I}_N \otimes \mathbf{a}_t(\theta_t, \phi_t)$ is the target channel matrix. With these forms, the overall received vector (2.9) in the range gate of interest is

$$\tilde{\mathbf{y}} = \rho_t\mathbf{Ts} + \sum_{q=1}^{Q} \gamma_q\mathbf{\Gamma}_q\mathbf{s} + \mathbf{y}_{\mathtt{n}} + \mathbf{y}_{\mathtt{i}}, \tag{2.19}$$

where the random reflectivities ρ_t, γ_q are as above. This channel model description will, in fact, form the basis of our understanding of the signal-independent clutter covariance in Chapter 3.

2.1.3 Detection and Joint Filter-Signal Design

With our models now defined, we now come to the question of what to do with these returns. Traditionally, STAP is focused on optimizing the detection of targets; that is, it would like to discriminate between two hypotheses: is a target present in a given range-Doppler-angle bin or not? At that point, a reasonable estimate of the target parameters may be obtained, or the detections may be output to another processing level for kinematic parameter estimation and tracking. We can formulate this process as a decision problem on the return sampled at a given range bin:

$$\begin{aligned} \mathscr{H}_0 : \quad & \tilde{\mathbf{y}} = \mathbf{y}_{\mathtt{u}} \\ \mathscr{H}_1 : \quad & \tilde{\mathbf{y}} = \mathbf{y}_{\mathtt{t}} + \mathbf{y}_{\mathtt{u}} \end{aligned}, \tag{2.20}$$

where the null hypothesis $\mathscr{H}_0$ indicates "no target" and the alternative hypothesis $\mathscr{H}_1$ indicates "target present." In order to make a decision, we require a decision statistic that we can threshold as appropriate. In general, this is facilitated by processing the return $\tilde{\mathbf{y}}$ with some form of filter[2]– say, $\mathbf{w} \in \mathbb{C}^{NML}$ – the output of which is given by $\mathbf{w}^H \tilde{\mathbf{y}}$. (If the general channel model is used, $\mathbf{w} \in \mathbb{C}^{M_R}$.) It has been shown that, for a Gauss-Gauss problem, maximizing the probability that $\mathscr{H}_1$ is chosen when $\mathscr{H}_1$ occurs (aka the probability of detection P_D) is equivalent to maximizing what we term the signal-to-interference-plus-noise ratio (SINR)[3].

[2]Also known as a weight vector or beamformer.

[3]More explicitly, we consider the signal-to-*clutter*-plus-interference-plus-noise ratio (SCINR), but for simplicity, we will collapse the clutter into "interference" for this purpose.

In our case, the SINR is given by the ratio of the expected energy from the target after processing to the expected energy from the undesired emitters (both signal-dependent and independent):

$$\text{SINR} = \frac{\mathbf{E}\{|\mathbf{w}^H \mathbf{y}_\mathtt{t}|^2\}}{\mathbf{E}\{|\mathbf{w}^H \mathbf{y}_\mathtt{u}|^2\}}, \tag{2.21}$$

where the expectations are taken over whatever relevant random parameters are in $\mathbf{y}_\mathtt{t}, \mathbf{y}_\mathtt{u}$. If we have prior knowledge of the channel characteristics as above, then this is

$$\text{SINR} = \frac{|\mathbf{w}^H \mathbf{T}\mathbf{s}|^2}{\mathbf{w}^H \mathbf{R}_\mathbf{u}(\mathbf{s})\mathbf{w}}. \tag{2.22}$$

Let there be a (potentially convex) set $\Omega_\mathbf{s}$ that describes properties we wish our signal $\mathbf{s}$ to have, like a power bound, similarity to another waveform, etc. The joint design problem most modern research considers is

$$\begin{aligned} \max_{\mathbf{w},\mathbf{s}} \quad & \text{SINR}(\mathbf{w}, \mathbf{s}) \\ \text{s.t.} \quad & \mathbf{s} \in \Omega_\mathbf{s}. \end{aligned} \tag{2.23}$$

Alternatively, we can consider what is known as the minimum-variance distortionless response (MVDR) strategy, first devised by Capon [64]. As the name implies, the technique designs a filter that aims to minimize the denominator in (2.23) while keeping the response to the target fixed at some level, say, $\kappa \in \mathbb{C}$. This results in the near-equivalent optimization problem

$$\begin{aligned} \min_{\mathbf{w},\mathbf{s}} \quad & \mathbf{w}^H \mathbf{R}_\mathbf{u}(\mathbf{s})\mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{T}\mathbf{s} = \kappa,\ \mathbf{s} \in \Omega_\mathbf{s}. \end{aligned} \tag{2.24}$$

Under this design strategy, $\text{SINR} = \frac{|\kappa|^2}{\mathbf{w}^H \mathbf{R}_\mathbf{u}(\mathbf{s})\mathbf{w}}$, always. We will soon see that, as κ varies, the denominator $\mathbf{w}^H \mathbf{R}_\mathbf{u}(\mathbf{s})\mathbf{w}$ scales to maintain the same "optimal" SINR.

Both (2.23) and (2.24) are non-convex problems in general, and thus are hard to solve on their own. We provided a proof of their non-convexity in [65], and [66, Lemma 1] states that "convexity" is only attained if every stationary point of the objectives is orthogonal to every clutter channel matrix $\mathbf{\Gamma}_q$ *and* the signal-independent covariance matrix $\mathbf{R_{ni}}$. We will soon see a variety of ways to handle this non-convexity, as well as a better characterization of it, in the sections below.

2.2 Biquadratic Programming and Nonconvexity

In this section, we discuss biconvex programming – that is, joint optimization of two multi-dimensional variables with objectives and constraint sets that are convex in each variable – and its application to our joint signal-filter design problem. In particular, we take a special interest in a subclass of biconvex programs, biquadratic programs (BQP/BiQP), because our primary objective function concerns optimization of a variance, which naturally leads to a quadratic structure.

2.2.1 Biconvex Optimization

We begin by considering the generalities of biconvex optimization [67]. Let $\mathbf{x} \in \mathbb{F}^n, \mathbf{y} \in \mathbb{F}^m$ be two vector variables of the given dimensions, drawn from a field $\mathbb{F}$. Nominally, this field may be the real numbers ($\mathbb{F} = \mathbb{R}$) or the complex numbers ($\mathbb{F} = \mathbb{C}$), but throughout this work, we will generally assume $\mathbb{F} = \mathbb{C}$. Assume we also have an objective function $f(\mathbf{x}, \mathbf{y}) : \mathbb{F}^n \times \mathbb{F}^m \to \mathbb{R}$ and a constraint set $\Upsilon \subseteq \mathbb{F}^n \times \mathbb{F}^m$. Taken together, this describes

the following optimization problem:

$$\begin{aligned} \min_{\mathbf{x},\mathbf{y}} \quad & f(\mathbf{x},\mathbf{y}) \\ \text{subject to} \quad & (\mathbf{x},\mathbf{y}) \in \Upsilon. \end{aligned} \tag{2.25}$$

Let us further define the following two restrictions of Υ, which are defined for fixed values $\hat{\mathbf{x}} \in \mathbb{F}^n, \hat{\mathbf{y}} \in \mathbb{F}^m$:

$$\Upsilon_X(\hat{\mathbf{x}}) = \{\mathbf{y} \in \mathbb{F}^m : (\hat{\mathbf{x}},\mathbf{y}) \in \Upsilon\} \tag{2.26}$$

$$\Upsilon_Y(\hat{\mathbf{y}}) = \{\mathbf{x} \in \mathbb{F}^m : (\mathbf{x},\hat{\mathbf{y}}) \in \Upsilon\}. \tag{2.27}$$

Notice that the subscript for each set denotes the variable that is *fixed*. With all of these in mind, we provide the following definition of a biconvex optimization problem:

Definition 2.1. *The optimization problem* (2.25) *is* biconvex *if, for every fixed* $\hat{\mathbf{x}} \in \mathbb{F}^n$ *and* $\hat{\mathbf{y}} \in \mathbb{F}^m$*,*

1. *the sets* $\Upsilon_X(\hat{\mathbf{x}})$ *and* $\Upsilon_Y(\hat{\mathbf{y}})$ *are each convex, and,*
2. *the functions* $f(\mathbf{x},\hat{\mathbf{y}})$ *and* $f(\hat{\mathbf{x}},\mathbf{y})$ *are each convex functions on* $\Upsilon_Y(\hat{\mathbf{y}})$ *and* $\Upsilon_X(\hat{\mathbf{x}})$*, respectively.*

It can be shown that, in general, biconvex problems are non-convex, which can make them seem difficult to solve. However, because they are convex in each variable, a simple solution strategy naturally arises – namely, alternating minimization.

The alternating minimization (AM) algorithm (also known as alternate convex search, or ACS), as the name implies, alternates between optimization of each variable while holding the other fixed, usually at the value of the previous iterate. More specifically, AM operates by first choosing an initial point[4] $\mathbf{x}_0 \in \mathcal{X}_0 \subseteq \mathbb{F}^n$. Then, for each iteration $k \in \mathbb{Z}^+$, AM solves

the reduced dimension optimizations

$$\mathbf{y}_k = \underset{\mathbf{y}\in\Upsilon_X(\mathbf{x}_k)}{\arg\min} \quad f(\mathbf{x}_k, \mathbf{y}) \tag{2.28a}$$

$$\mathbf{x}_{k+1} = \underset{\mathbf{x}\in\Upsilon_Y(\mathbf{y}_k)}{\arg\min} \quad f(\mathbf{x}, \mathbf{y}_k) \tag{2.28b}$$

until some convergence criterion is met, be it in the objective function $f(\cdot,\cdot)$ or the iterate pair $(\mathbf{x}_k, \mathbf{y}_k)$. Note that the order of this process could be swapped; that is, one could start with an initializer $\mathbf{y}_0 \in \mathcal{Y}_0 \subseteq \mathbb{F}^m$, solve for an optimal $\mathbf{x}_0$, and so on.

Before continuing, we make a few observations. First, *how* one chooses to optimize each variable is left unstated in this definition. Indeed, the choice of optimization technique at each step is the main differentiator between the classes of joint signal-filter design algorithms we will explore in Chapter 2.3 below. Second, the concept of convergence is also left quite open to interpretation. It is well known that, by construction, AM algorithms provide a monotonic cost function decrease at each iteration [67, Lemma 4.6], *i.e.*

$$f(\mathbf{x}_k, \mathbf{y}_{k-1}) \geq f(\mathbf{x}_k, \mathbf{y}_k) \geq f(\mathbf{x}_{k+1}, \mathbf{y}_k).$$

Furthermore, it has been shown [68] that if the constraint sets, which change at each iteration, "converge" in some sense to stationary sets and the objective function f retains certain properties, AM is guaranteed to converge to a given optimal value on those stationary sets. This is true even if the sets come from different metric spaces [69]. Later, we will show a specialized proof of this convergence in the context of our optimization problem.

[4]Initializers are often chosen to be jointly feasible, *i.e.* $\mathbf{x}_0 \in \Upsilon_Y(\hat{\mathbf{y}})$ for any $\hat{\mathbf{y}}$, or feasible only for the parts of $\Upsilon_Y(\hat{\mathbf{y}})$ independent of $\hat{\mathbf{y}}$, but this is not mandatory.

2.2.2 Biquadratic Optimization

Next, we consider a specific case of biconvex optimization that will prove useful for our characterization of SINR-focused joint designs and, later, identifying a simpler way to solve them. Biquadratic optimization, or biquadratic programs (BQP), arise when the dominant form of the objective function f is individually quadratic in each variable $\mathbf{x}, \mathbf{y}$.

As a standalone problem, specific coverage of the BQP in the literature, from the optimization community or otherwise, can be charitably described as sparse. In fact, there are only seven papers directly addressing the BQP, all of which have been released within the last decade [49, 50, 51, 52, 53, 54, 55]. This is somewhat puzzling given the depth and breadth of related problems both cited within these works and conceivable by an ordinary engineer – applications range from economics to quantum physics to the fully-adaptive radar concept we wish to pursue here. Nevertheless, this limitation allows a reasonable overview of the state-of-the-art without too much difficulty.

The generic BQP, in terms of real variables, is first defined by Ling, et al. [49]. Consider the vectors $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{y} \in \mathbb{R}^m$. The standard biquadratic optimization problem is

$$\begin{aligned} \min_{\mathbf{x},\mathbf{y}} \quad & b(\mathbf{x},\mathbf{y}) = \sum_{1\leq i,k\leq n, 1\leq j,l\leq m} b_{ijkl} x_i y_j x_k y_l \\ \text{subject to} \quad & \|\mathbf{x}\|_2 = 1, \|\mathbf{y}\|_2 = 1, \end{aligned} \tag{2.29}$$

where b_{ijkl} is the (i,j,k,l)th element of the 4th-order tensor $\mathcal{B} \in \mathbb{R}^{n\times m\times n\times m}$, the subscripted x and y are the appropriate element of the corresponding vector, and $\|\cdot\|_2$ is the standard Euclidean norm. Before we continue, it is worth noting that the objective function $b(\mathbf{x},\mathbf{y})$ can be rewritten using a tensor operator and the matrix inner product, viz

$$b(\mathbf{x},\mathbf{y}) = \mathcal{B}(\mathbf{x}\mathbf{x}^T) \bullet (\mathbf{y}\mathbf{y}^T) = \mathcal{B}^*(\mathbf{y}\mathbf{y}^T) \bullet (\mathbf{x}\mathbf{x}^T)$$

where $\mathbf{X} \bullet \mathbf{Y} = \mathrm{Tr}(\mathbf{Y}^T\mathbf{X})$ is the matrix inner product, $\mathcal{B}^*$ is the adjoint operator of $\mathcal{B}$, and the tensor-operator matrices are $\mathcal{B}(\mathbf{x}\mathbf{x}^T) = \sum_{i,k=1}^{n} b_{ijkl}x_ix_k$ and $\mathcal{B}^*(\mathbf{y}\mathbf{y}^T) = \sum_{j,l=1}^{m} b_{ijkl}y_jy_l$. If the operator $\mathcal{B}$ has a matricization (that is, a rearrangement of its elements into a matrix compatible with the products $\mathbf{x}\mathbf{x}^T$ or $\mathbf{y}\mathbf{y}^T$) that is at least positive semidefinite, then (2.29) is also biconvex.

In addition to defining the general framework of the BQP, [49] also provides some critical insight into the solvability of such problems, even under a simple constraint like the unit sphere. Specifically, they prove that both the BQP and its naive semidefinite relaxation are NP-hard problems, and that the initial general problem does not even admit a polynomial time *approximation* algorithm. This would seem to bode ill for any hope of reasonable solutions, but the authors demonstrate that polynomial-time approximation algorithms exist (with bounds approximately inversely proportional to the dimension of the variables) when the objective is square-free or has squared terms in only one of the variables. Finally, they demonstrate the creation of such algorithms, showing a tradeoff between speed and accuracy for sum-of-squares based solvers (better accuracy) versus convex semidefinite relaxation solvers (faster).

The pair of papers [50, 52] extend concepts from [49], focusing on quadratically constrained versions of the BQP, similar to most of the constraints our radar problem encounters. Both papers develop probabilistic bounds for these BQP types, as solved through a similar (but more complex and restricted) form of the semidefinite program considered previously. The key difference between the two is that [50] considers a more general non-homogeneous objective function which includes a function that depends on only one of the variables, which is homogenized by the addition of an arbitrary scaling. In either case, the solutions of the relaxed SDP provide inputs to a randomization scheme that results in feasible approximate solutions to the original BQP. Additionally, [52] makes the first claims about *complex-valued*

BQPs (BQPs that are functions of complex variables); namely, that they are equivalent to a relaxed linear semidefinite program so long as the number of constraints does not exceed 4.

Finally, [51] presents a familiar alternative optimization method to solve the standard BQP: an alternating direction method (ADM), which is a close cousin of the AM described above. This solver uses a modified form of the SDP relaxation to drive what is essentially an iterative solver of the primal-dual problem. Analysis within the work demonstrates that the ADM algorithm converges to a global solution (should one exist) almost surely, in a computationally more efficient fashion than any of the methods considered in [49]. We will see a general application of this concept applied later in Chapter 2.3. This work is related to similar investigations on non-negative quartic forms; a good review of these topics, especially in an optimization context, may be found in [70][5].

2.2.3 BQP for Fully Adaptive Radar

With the general mathematical concepts explained, we now turn to how they can be applied to the joint design problem. Consider an expanded version of (2.24) in the pair $(\mathbf{w}, \mathbf{s})$:

$$\begin{aligned} \min_{\mathbf{w},\mathbf{s}} \quad & \mathbf{w}^H \mathbf{R_{ni}} \mathbf{w} + \mathbf{w}^H \left(\sum_{q=1}^{Q} \overline{R}_{\gamma}^{q} \mathbf{\Gamma}_q \mathbf{s}\mathbf{s}^H \mathbf{\Gamma}_q^H \right) \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{T}\mathbf{s} = \kappa, \ \mathbf{s} \in \Omega_{\mathbf{s}} \end{aligned} . \tag{2.30}$$

where, again, $\Omega_{\mathbf{s}}$ is a signal-only constraint set. Clearly, if the signal $\mathbf{s}$ is fixed, then the objective function in (2.30) is a convex quadratic function in $\mathbf{w}$, because the matrix $\mathbf{R_{ni}} + \sum_{q=1}^{Q} \overline{R}_{\gamma}^{q} \mathbf{\Gamma}_q \mathbf{s}\mathbf{s}^H \mathbf{\Gamma}_q^H$ is positive semidefinite. If the filter $\mathbf{w}$ is fixed, we can rearrange the clutter

[5] Interestingly, the authors in [70] correctly identify [45] as a canonical example of applied quartic optimization.

covariance term and obtain the following additional form for the objective:

$$f_u(\mathbf{s}) = \mathbf{s}^H \mathbf{Z}_{\mathbf{c}}(\mathbf{w})\mathbf{s} + \mathbf{w}^H \mathbf{R}_{\mathbf{ni}}\mathbf{w} \tag{2.31}$$

where the "filter covariance" $\mathbf{Z}_{\mathbf{c}}(\mathbf{w})$ given by $\mathbf{Z}_{\mathbf{c}}(\mathbf{w}) = \sum_{q=1}^{Q} \overline{R}_{\gamma}^{q} \mathbf{\Gamma}_q^H \mathbf{w}\mathbf{w}^H \mathbf{\Gamma}_q$ is also a positive semidefinite matrix, by construction. This is also, obviously, a convex quadratic function in $\mathbf{s}$. We can also clearly define the total and restricted constraint sets for the variables $\mathbf{w}, \mathbf{s}$ and fixed values $\hat{\mathbf{w}}, \hat{\mathbf{s}}$ as

$$\Upsilon(\mathbf{w}, \mathbf{s}) = \{(\mathbf{w}, \mathbf{s}) : \mathbf{w}^H \mathbf{T}\mathbf{s} = \kappa,\ \mathbf{s} \in \Omega_{\mathbf{s}}\} \tag{2.32}$$

$$\Upsilon_{\mathbf{s}}(\hat{\mathbf{s}}) = \{\mathbf{w} : \mathbf{w}^H \mathbf{T}\hat{\mathbf{s}} = \kappa\} \tag{2.33}$$

$$\Upsilon_{\mathbf{w}}(\hat{\mathbf{w}}) = \{\mathbf{s} : \hat{\mathbf{w}}^H \mathbf{T}\mathbf{s} = \kappa,\ \mathbf{s} \in \Omega_{\mathbf{s}}\} \tag{2.34}$$

Since the Capon constraint is affine for a fixed parameter, $\Upsilon_{\mathbf{s}}(\hat{\mathbf{s}})$ is convex. So long as $\Omega_{\mathbf{s}}$ is convex, $\Upsilon_{\mathbf{s}}(\hat{\mathbf{s}})$ is convex, and consequently, $\Upsilon(\mathbf{w}, \mathbf{s})$ is biconvex. Hence, (2.30) is a biconvex program, by definition.

With these forms determined, we also can identify the tensor-operator that relates the bi-quadratic form (2.29) to the general joint design problem (2.24). Let $\mathcal{C} \in \mathbb{C}^{NML \times N \times NML \times N}$ denote the clutter tensor operator that is defined by the relationships $\mathcal{C}(\mathbf{s}\mathbf{s}^H) = \mathbf{R}_{\mathbf{c}}(\mathbf{s})$ and $\mathcal{C}^*(\mathbf{w}\mathbf{w}^H) = \mathbf{Z}_{\mathbf{c}}(\mathbf{w})$. For the baseline WA-STAP clutter model, this tensor is given element-by-element as

$$c_{ijkl} = \begin{cases} \sum_{q=1}^{Q} \bar{R}_q^{\gamma} v_{dq} v_{eq}^* a_{mq} a_{nq}^* & j = m + M(i-1) + NM(d-1), \\ & l = n + M(k-1) + NM(e-1) \\ 0 & \text{otherwise} \end{cases}$$

where m, n are integers in $\{1, 2, \dots M\}$ and d, e are integers in $\{1, 2, \dots, L\}$ that satisfy the conditions above, and the subscripted $v_{\cdot q}, a_{\cdot q}$ are the respective elements of the Doppler response $\mathbf{v}_{c,q}$ and spatial response $\mathbf{a}_{c,q}$ of the qth clutter patch.

We can also identify a joint form which will be useful later, and also demonstrates that the problem can be considered as a quartic vector optimization. Let $J = N(ML + 1)$ be the total number of joint parameters. If we treat the signal and beamformer as a single stacked variable, say $\mathbf{b} = [\mathbf{w}^T \, \mathbf{s}^T]^T \in \mathbb{C}^J$, we can find an equivalent form of the above optimization problem in the new variable. Define the selection matrices $\mathbf{\Psi}_W = [\mathbf{I}_{NML} \;\; \mathbf{0}_{NML \times N}]$ and $\mathbf{\Psi}_S = [\mathbf{0}_{N \times NML} \; \mathbf{I}_{NML}]$ such that $\mathbf{w} = \mathbf{\Psi}_W \mathbf{b}, \mathbf{s} = \mathbf{\Psi}_S \mathbf{b}$. Then, the joint quartic problem is given by

$$\begin{aligned} \min_{\mathbf{b}} \quad & \mathbf{b}^H \widetilde{\mathbf{R}}_{\mathbf{u}}(\mathbf{b}) \mathbf{b} \\ \text{s.t.} \quad & \mathbf{b}^H \widetilde{\mathbf{T}} \mathbf{b} = \kappa, \mathbf{\Psi}_S \mathbf{b} \in \Omega_{\mathbf{s}}, \end{aligned} \tag{2.35}$$

where $\widetilde{\mathbf{R}}_{\mathbf{u}}(\mathbf{b}) = \widetilde{\mathbf{R}}_{\mathbf{ni}} + \sum_{q=1}^{Q} \overline{R}_q^{\gamma} \widetilde{\mathbf{\Gamma}}_q \mathbf{b}\mathbf{b}^H \widetilde{\mathbf{\Gamma}}_q^H$, $\widetilde{\mathbf{R}}_{\mathbf{ni}} = \mathbf{\Psi}_W^T \mathbf{R}_{\mathbf{ni}} \mathbf{\Psi}_W$, $\widetilde{\mathbf{\Gamma}}_q = \mathbf{\Psi}_W^T \mathbf{\Gamma}_q \mathbf{\Psi}_S$, and $\widetilde{\mathbf{T}} = \mathbf{\Psi}_W^T \mathbf{T} \mathbf{\Psi}_S$ are "expanded" versions of the matrices seen in the previous problem. In this case, the clutter tensor expands to a tensor-operator $\widetilde{\mathcal{C}} \in \mathbb{C}^{J \times J \times J \times J}$, whose matrix form is given by $\widetilde{\mathbf{R}}_{\mathbf{c}}(\mathbf{b}) = \mathbf{\Psi}_W^T \mathbf{R}_{\mathbf{c}}(\mathbf{\Psi}_S \mathbf{b}) \mathbf{\Psi}_W = \widetilde{\mathcal{C}}(\mathbf{b}\mathbf{b}^H)$. In the WA-STAP case, its element-by-element description is given by

$$\widetilde{c}_{ijkl} = \begin{cases} c_{ijkl} & i, k \leq NML \text{ and } j, l \leq N \\ 0 & \text{otherwise} \end{cases}$$

where c_{ijkl} is the element of the "smaller" tensor-operator defined above.

The tensors $\mathcal{C}$ and $\widetilde{\mathcal{C}}$ have properties that will be useful later, which mostly relate to the fact that $\mathcal{C}$ preserves semidefiniteness. That is, given a PSD matrix $\mathbf{X}$, $\mathcal{C}(\mathbf{X})$ is also positive-semidefinite.

2.3 Sequential Algorithms for Joint Designs

In this section, we analyze the two primary classes of optimization techniques used for joint filter-signal design in the literature for signal-dependent clutter environments: signal-filter AM and spectral AM. The two classes stand on either end of a speed-accuracy tradeoff, mostly because of how they are employed. Both classes, however, will share a partial nomenclature for operation that will be used throughout the rest of this paper. As mentioned previously, AM alternates between the optimizations (2.28b) and (2.28a). In the filter-signal design context, we will refer to these as the $\mathbf{w}$-step and $\mathbf{s}$-step, where the terminology indicates which variable is being optimized. This will be true even when we discuss the optimization of signal or filter *spectra*, which are technically different variables from the vector inputs.

2.3.1 Signal-Filter Alternating Minimization

We begin our analysis with signal-filter AM. Signal-filter AM operates on the vectors $\mathbf{w}, \mathbf{s}$ individually, as vectors, and uses particular concentrations of the given constraint sets to inform the optimization technique used in each.

As mentioned above, there are multiple equivalent formulations of the same cost function. Let us define the target gain response matrix due to a given signal $\hat{\mathbf{s}}$ as $\mathbf{G}_S(\hat{\mathbf{s}}) = \mathbf{T}\hat{\mathbf{s}}\hat{\mathbf{s}}^H\mathbf{T}^H$ and the observable response for a given filter as $\mathbf{G}_W(\hat{\mathbf{w}}) = \mathbf{T}^H\hat{\mathbf{w}}\hat{\mathbf{w}}^H\mathbf{T}$. The most common form of signal-filter AM is predicated on the fact that, at each step, the optimal $\mathbf{w}$ ($\mathbf{s}$) are a maximum (minimum) generalized eigenvector[6]of the target response $\mathbf{G}_S(\hat{\mathbf{s}})$ ($\mathbf{G}_W(\hat{\mathbf{w}})$) and the undesired source covariance $\mathbf{R}_{\mathsf{u}}(\hat{\mathbf{s}})$ ($\mathbf{Z}_{\mathsf{c}}(\hat{\mathbf{w}})$). Due to certain waveform constraints, some algorithms execute a different signal optimization procedure in the $\mathbf{s}$-step; for example, in [38], the authors propose solving the fractional form (2.23) via a closed-form phase calculation under combined constant-modulus and infinity-norm similarity constraints. Another

recent example is Algorithm 3 of [47], who use the power-method-like SWORD technique of [48] to solve the fractional problem for separate constant modulus and two-norm similarity constraints. Regardless, these algorithms can be considered generally to be of lower complexity because, at each step, they operate on an M_R- (M_T-) dimensional problem. However, as noted in [8], these techniques can often be ill-conditioned and can require many iterations to converge, whatever the chosen criterion. Furthermore, the number of iterations needed for convergence is either not investigated or otherwise treated as a problem-dependent curiosity. Additionally, these require the prior knowledge of the components of $\mathbf{R}_{\mathtt{u}}(\hat{\mathbf{s}})$ and/or $\mathbf{Z}_{\mathtt{c}}(\hat{\mathbf{w}})$ to ensure that both matrices are positive definite. In a practical sense, this requires $\mathbf{R}_{\mathbf{ni}}$ (whether estimated or exactly known) to be positive definite for the former case, and $Q \geq M_T$ ($Q \geq N$ for WA-STAP) in the latter case. If such a guarantee cannot be made, then one must resort to more complicated methods, like the proximal alternating minimization of [7, 8].

For now, let us consider an illustrative example of signal-filter optimization, as described by [6, 8] for WA-STAP. Let the signal-only constraint set be a simple power constraint; i.e., $\Omega_{\mathbf{s}} = \{\mathbf{s} : \|\mathbf{s}\|_2^2 \leq P_o\}$. When signal-filter AM is used to solve (2.30), at the k-th iteration, we obtain the iterate pair

$$\mathbf{w}_k = \frac{\kappa^* \mathbf{R}_{\mathtt{u}}^{-1}(\mathbf{s}_{k-1})\mathbf{T}\mathbf{s}_{k-1}}{\mathbf{s}_{k-1}^H \mathbf{T}^H \mathbf{R}_{\mathtt{u}}^{-1}(\mathbf{s}_{k-1})\mathbf{T}\mathbf{s}_{k-1}} \tag{2.36a}$$

$$\mathbf{s}_k = \frac{\kappa \mathbf{F}^{-1}(\mathbf{w}_k, \lambda_k)\mathbf{T}^H \mathbf{w}_k}{\mathbf{w}_k^H \mathbf{T}\mathbf{F}^{-1}(\mathbf{w}_k, \lambda_k)\mathbf{T}^H \mathbf{w}_k} \tag{2.36b}$$

where $\mathbf{F}(\mathbf{w}_k, \lambda_k) = \mathbf{Z}_{\mathtt{c}}(\mathbf{w}_k) + \lambda_k \mathbf{I}_N$ and λ_k is a Lagrange parameter that relates to the power constraint at the given iteration. The general solution of λ_k for a fixed $\mathbf{w}_k$ is given by

[6] In the case of continuous-time optimization, it is the eigen*function* of a Fredholm integral, as noted by countless aforementioned authors.

$\lambda = \max[0, \lambda^\star]$, where $\lambda^\star$ is a real root of the equation

$$\lambda^\star \left(|\kappa|^2 \mathbf{w}_k^H \mathbf{T} \mathbf{F}_\star^{-2} \mathbf{T}^H \mathbf{w}_k - P_o (\mathbf{w}_k^H \mathbf{T} \mathbf{F}_\star^{-1} \mathbf{T}^H \mathbf{w}_k)^2 \right) = 0. \tag{2.37}$$

and $\mathbf{F}_\star = \mathbf{F}(\mathbf{w}_k, \lambda^\star)$. This expression is obtained through simple algebraic manipulation of the Lagrange function, and can be solved by any numerical rooting or unconstrained optimization technique. We shall subsequently refer to this algorithm, when it may be confused with a general AM procedure, as EigenAM.

As a practical matter, P_o is often large since the radar experiences a power loss factor of $1/R^4$, requiring large power budgets to cover large range swaths. With this effect in mind, the authors in [6, 8] demonstrated that $\lambda^\star \to 0$ when $P_o \gg \kappa$, regardless of the filter iterate $\mathbf{w}_k$. If the power *must* equal P_o, then the rooting solution (2.37) should be used.

Power Scaling: A Computational Compromise As a design compromise, the authors of [8] note that it is possible to save computational cycles by always setting $\lambda = 0$, avoiding the need to implement a difficult root finding routine to solve (2.37). This effectively makes the power constraint inactive, however, and all subsequent signal iterates $\mathbf{s}_k$ will be such that $\|\mathbf{s}_k\|_2 < P_o$. In cases where an exact solution is required – or, for example, when the matrix $\mathbf{Z}_\mathbf{c}(\mathbf{w}_k)$ in the $\mathbf{s}$-step is rank deficient – (2.37) may then be used if the resources are available.

The consequence of this choice is that if an exactly power-bounded solution is required and the computational resources for root-finding are not available, then the final iterate (say, $\mathbf{s}^\star$) must then be scaled back to meet the power constraint. The result of this scaling is as follows: Let $(\mathbf{w}^\star, \mathbf{s}^\star)$ be a solution of EigenAM where λ is forced to zero; that is, it is either a limit point or a solution which satisfies the desired objective tolerance. Then consider the following scaled solution, $(\frac{\|\mathbf{s}^\star\|}{\sqrt{P_o}}\mathbf{w}^\star, \frac{\sqrt{P_o}}{\|\mathbf{s}^\star\|}\mathbf{s}^\star)$. It is clear that this scaled solution now satisfies

the power equality, and still satisfies the Capon constraint. Furthermore, this scaling does not change the clutter objective, due to the structure; that is,

$$\left(\frac{||\mathbf{s}^\star||}{\sqrt{P_o}}\mathbf{w}^\star\right)^H \mathbf{R}_\mathbf{c}\left(\frac{\sqrt{P_o}}{||\mathbf{s}^\star||}\mathbf{s}^\star\right)\left(\frac{||\mathbf{s}^\star||}{\sqrt{P_o}}\mathbf{w}^\star\right) = \mathbf{w}^{\star H}\mathbf{R}_\mathbf{c}\left(\mathbf{s}^\star\right)\mathbf{w}^\star \tag{2.38}$$

However, the scaling does increase the signal-independent disturbance response by a factor of $||\mathbf{s}^*||^2/P_o$ when compared to the original solution. Furthermore, while the scaled iterates may produce a usable solution, there is no guarantee that they, in fact, satisfy the KKTs of the original problem.

2.3.2 Spectral Alternating Minimization

We continue by discussing the next class of algorithms typically used for joint design that we call spectral alternating minimization, or spectral AM. These techniques rely on relaxation of one or both of the vector quadratic terms in the SINR cost function, producing a transmission or observation spectrum/correlation matrix from which an optimal signal or filter must then be synthesized, either at each iteration or after convergence. In some cases, relaxation is considered only at one stage of the procedure; typically, this is the $\mathbf{s}$-step, where constraints are more complex and amenable to relaxation to linear constraints. In others, where a constraint set might itself be biquadratic (say, if the Capon constraint is transformed into a "minimum target gain for an uncertain target response" constraint), relaxation can be performed at both steps.

Before we discuss the design concept, we provide a quick overview of relaxing quadratic programs to semidefinite programs, and their concurrence. We follow some of the tutorial results provided in [71]. Consider, for example, a general quadratically constrained quadratic

problem in the variable $\mathbf{x} \in \mathbb{F}^n$:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbf{x}^H \mathbf{A}_0 \mathbf{x} \\ \text{s.t.} \quad & \mathbf{x}^H \mathbf{A}_i \mathbf{x} \trianglelefteq b_i, i \in 1, 2, \ldots, C \end{aligned} \tag{2.39}$$

where the matrix $\mathbf{A}_0 \in \mathbb{H}_n^+$, the constraint matrices $\mathbf{A}_i \in \mathbb{H}_n$ (though they may merely be complex valued), and the relationship $\trianglelefteq$ indicates one of $=, \geq, \leq$. Let $\mathbf{X} = \mathbf{x}\mathbf{x}^H \in \mathbb{H}_n^+$. Through the rules of the trace operator, it is clear that (2.39) can be rewritten as

$$\begin{aligned} \min_{\mathbf{X} \in \mathbb{H}_n^+} \quad & \text{tr}(\mathbf{A}_0 \mathbf{X}) \\ \text{s.t.} \quad & \text{tr}(\mathbf{A}_i \mathbf{X}) \trianglelefteq b_i, i \in 1, 2, \ldots, C \\ & \text{rank}(\mathbf{X}) = 1. \end{aligned} \tag{2.40}$$

Now, (2.40) is a non-convex problem, due to the rank constraint. If this is relaxed, we now have the convex matrix problem

$$\begin{aligned} \min_{\mathbf{X} \in \mathbb{H}_n^+} \quad & \text{tr}(\mathbf{A}_0 \mathbf{X}) \\ \text{s.t.} \quad & \text{tr}(\mathbf{A}_i \mathbf{X}) \trianglelefteq b_i, i \in 1, 2, \ldots, C, \end{aligned} \tag{2.41}$$

which can be efficiently solved. The above process is often called semidefinite relaxation, or SDR, regardless of the particular cost function or constraints being relaxed. The same relaxation process can be applied to fractional programs (like the one considered in (2.23)), and then transformed into a form like (2.41) through the use of the Charnes-Cooper transformation. This transformation functionally sets the denominator equal to one, thus adding one more trace constraint. Regardless of the relaxation, the purpose is to transform an NP-hard problem into a more tractable one. However, this requires some guarantee on whether or not the relaxation is tight – that is, if the solution $\mathbf{X}^\star$ to (2.41) is rank one and therefore

exactly solves (2.39). For a complex-valued QCQP/SDP relaxation, we have the following theorem:

Theorem 2.1. *Let the complex-valued linear SDP with affine constraints* (2.41) *and its dual both be feasible and bounded (i.e., solvable). Then,* (2.41) *always has an optimal solution* $\mathbf{X}^\star$ *such that*

$$\text{rank}(\mathbf{X}^\star) \leq \sqrt{C}.$$

This naturally leads to the following corollary about the tightness of the relaxed solution:

Corollary 2.1. *The relaxation of the complex-valued QCQP* (2.39) *to the complex-valued linear SDP* (2.41) *is guaranteed to be tight if the number of constraints* $C \leq 3$.

When the relaxation is not tight, the major question then becomes how to recover an acceptable (*i.e.*, close-to-optimal) rank-one solution $\mathbf{x}^{appx}$ from $\mathbf{X}^\star$. A common technique which we will explore in more depth in Chapter 3 is merely choosing the eigenvector of $\mathbf{X}^\star$ associated with the largest eigenvalue, as this is the best rank-one approximation in the Frobenius norm sense. Depending on the number and type of affine constraints C, a geometrically-based rank-reduction or equivalence process may be possible. For example, a rank reduction process is examined for separable (*i.e.*, multiple sets of variables) SDPs in [72]. The authors in [72] also show (see Corollary 4.6) that if $C \leq 5$ and two of $\mathbf{A}_i$ are positive semidefinite, then a rank-one solution is also possible. In a similar vein, when $C = 4$, Ai, Huang and Zhang [73] demonstrated a procedure to generate a decomposition of $\mathbf{X}^\star$ such that each rank-one component is feasible (*i.e.*, satisfies the constraints). This latter technique is employed across a variety of joint signal-filter design processes, one instance of which we will describe below.

Finally, a popular, if computationally expensive, means of finding approximate rank-one solutions relies on the equivalence of (2.41) to a *stochastic* version of (2.39). This process, known as randomization, assumes the following: let $\boldsymbol{\xi} \in \mathbb{F}^n$ be a random vector drawn from a zero-mean (complex) normal distribution with covariance $\mathbf{X}$; mathematically, $\boldsymbol{\xi} \sim \mathcal{CN}(\mathbf{0}_{n\times 1}, \mathbf{X})$. Then, consider the expectation of a quadratic form in $\boldsymbol{\xi}$; say, $\mathbf{E}_{\boldsymbol{\xi}}\{\boldsymbol{\xi}^H \mathbf{A}_0 \boldsymbol{\xi}\}$. After rearranging the expression by the rules of the trace operator and noting that $\mathbf{E}_{\boldsymbol{\xi}}\{\boldsymbol{\xi}\boldsymbol{\xi}^H\} = \mathbf{X}$, we have $\mathbf{E}_{\boldsymbol{\xi}}\{\boldsymbol{\xi}^H \mathbf{A}_0 \boldsymbol{\xi}\} = \mathrm{tr}(\mathbf{A}_0\mathbf{X})$. Therefore, the objective and the constraints are all satisfied in expectation by the relaxed solution. With this in mind, the randomization procedure proceeds as follows:

1. A number of random draws D are made from the distribution $\mathcal{CN}(\mathbf{0}_{n\times 1}, \mathbf{X}^\star)$, producing the vector collection $\{\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_D\}$.

2. If necessary, a problem dependent transformation is applied to $\{\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_D\}$ to create feasible points of the original problem, resulting in the transformed set $\{\boldsymbol{\xi}_1^f, \ldots, \boldsymbol{\xi}_D^f\}$.

3. The approximate solution $\mathbf{x}^{appx}$ is given by the random feasible vector that minimizes the original objective function; that is, $\mathbf{x}^{appx} = \arg\ \min_{d\in\{1,\ldots,D\}}\ \boldsymbol{\xi}_d^{f\,H} \mathbf{A}_0 \boldsymbol{\xi}_d^f$.

Depending on the problem, D may need to be quite large to get an acceptable solution. Further insights into the application of this technique for approximation bounds can be found in [71] and references therein.

With the preliminaries in place, we now discuss the application of linear SDR to the joint design problem. Depending on the constraints, there are multiple ways this can be realized. First, we discuss methods that relax in the **s**-step only. In these cases, for each iteration, the **w**-step generates a new filter $\mathbf{w}_k$ in a manner similar to (2.36a) or from an eigendecomposition, as in the signal-filter AM procedure. We note that this process must *correctly* enforce the Capon constraint (or its proxy) by constraining both the real and imaginary parts of the

term $\mathbf{w}^H\mathbf{T}\mathbf{s}$ (implicitly or explicitly), or else the returned iterate $\mathbf{w}_k$ is not optimal for that step and performance (convergence or otherwise) may be degraded. Then, for the $\mathbf{s}$-step, they consider either the clutter-minimizing QCQP

$$\begin{aligned} \min_{\mathbf{s}} \quad & \mathbf{s}^H\mathbf{Z}_{\mathbf{c}}(\mathbf{w}_k)\mathbf{s} \\ \text{s.t.} \quad & \mathbf{s}^H\mathbf{G}_W(\mathbf{w}_k)\mathbf{s} = |\kappa|^2 \qquad \mathbf{s} \in \Omega_{\mathbf{s}} \end{aligned} \tag{2.42}$$

or the target-response-maximizing QCQP

$$\begin{aligned} \max_{\mathbf{s}} \quad & \mathbf{s}^H\mathbf{G}_W(\mathbf{w}_k)\mathbf{s} \\ \text{s.t.} \quad & \mathbf{s}^H\mathbf{F}\left((\mathbf{w}_k^H\mathbf{R}_{\mathbf{ni}}\mathbf{w}_k), \mathbf{w}_k\right)\mathbf{s} = 1 \qquad \mathbf{s} \in \Omega_{\mathbf{s}}. \end{aligned} \tag{2.43}$$

These forms can then be relaxed to either

$$\begin{aligned} \min_{\mathbf{S}\in\mathbb{H}_+^{M_T}} \quad & \text{tr}(\mathbf{Z}_{\mathbf{c}}(\mathbf{w}_k)\mathbf{S}) \\ \text{s.t.} \quad & \text{tr}(\mathbf{G}_W(\mathbf{w}_k)\mathbf{S}) = |\kappa|^2 \qquad \mathbf{S} \in \mathcal{T}_R(\Omega_{\mathbf{s}}) \end{aligned} \tag{2.44}$$

or

$$\begin{aligned} \max_{\mathbf{S}\in\mathbb{H}_+^{M_T}} \quad & \text{tr}(\mathbf{G}_W(\mathbf{w}_k)\mathbf{S}) \\ \text{s.t.} \quad & \text{tr}(\mathbf{F}\left((\mathbf{w}_k^H\mathbf{R}_{\mathbf{ni}}\mathbf{w}_k), \mathbf{w}_k\right)\mathbf{S}) = 1 \qquad \mathbf{S} \in \mathcal{T}_R(\Omega_{\mathbf{s}}). \end{aligned} \tag{2.45}$$

where $\mathbf{S}$ is the positive semidefinite arbitrary-rank relaxation of $\mathbf{s}\mathbf{s}^H$. Here, $\mathcal{T}_R(\Omega_{\mathbf{s}}) \subseteq \mathbb{H}^n$ is a transformed version of $\Omega_{\mathbf{s}}$ where all dependences on $\mathbf{s}\mathbf{s}^H$ are now (typically affine) dependences on $\mathbf{S}$. For example, in a power-constrained case where $\Omega_{\mathbf{s}} = \{\mathbf{s} : \|\mathbf{s}\|_2^2 \leq P_o\}$, $\mathcal{T}_R(\Omega_{\mathbf{s}}) = \{\mathbf{S} : \text{tr}(\mathbf{S}) \leq P_o\}$. In the latter form (2.45), the transformation includes the need for another variable to appropriately account for scaling. For example, if a constant modulus constraint is enforced, i.e. $\Omega_{\mathbf{s}} = \{\mathbf{s} : |s_i| = \rho,\ i = 1, \ldots, M_T\}$, then the transformed set is

defined by a variable t such that

$$\mathcal{T}_R(\Omega_{\mathbf{s}}) = \{\mathbf{S} : \mathbf{S} = \mathbf{M}/t,\ \mathrm{diag}(\mathbf{M}) = t\rho\mathbf{1}_{M_T}, \mathbf{M} \in \mathbb{H}_+^{M_T}, t \geq 0\}$$

In these cases, the optimization is technically carried out over the unscaled matrix $\mathbf{M}$ and the scaling parameter t. For the most part, any technique that uses the Charnes-Cooper transformation like [45]'s SOA-2 algorithm, Algorithms 2 and 4 of [47] (henceforth TT-A2 and TT-A4), and Algorithm 2 of [44] (henceforth AA2) will optimize a form like (2.45). Then, finally, a signal $\mathbf{s}_k$ is synthesized from its relaxed spectrum $\mathbf{S}$. This can be achieved through rank-based decomposition [44] or randomization [45, 47]. Randomization is far more popular when a constant modulus constraint is applied, since there are then at least M_T constraints and, typically, $M_T \geq 4$. If randomization is used, another design choice can be *when* to synthesize a new signal, as this may lead to better convergence and lower computational load. For example, SOA-2 performs a randomization step at every iteration to generate a new signal $\mathbf{s}_k$, while TT-A2 & TT-A4 perform the randomization after convergence, forming a new *spectrum* $\mathbf{S}_k$ at each iteration.

Relaxation for the $\mathbf{w}$-step, as seen in works like [46], is carried out similarly to the above. In these cases, $\mathbf{w}\mathbf{w}^H$ is replaced with an arbitrary-rank matrix $\mathbf{W} \in \mathbb{H}_+^{M_R}$ in the relevant filter-equivalents to (2.42) and (2.43). Filter synthesis from the relaxed observing spectrum $\mathbf{W}$ may or may not happen at each iteration, and the same caveats from before still apply.

2.4 Perspectives on Solvability and Convergence

In this section, we provide two further examples of both signal-filter AM and spectral AM that are functionally equivalent to the EigenAM method, but illustrate the behavior of the algorithm as well as criteria for optimality and uniqueness of solutions. We also provide

a general proof of convergence of any AM scheme for bi-convex optimization and discuss the need to enforce constraints appropriately at every step, even if they necessarily seem redundant. This analysis was originally presented in [74]. Throughout this section, we assume the dimensions of the WA-STAP problem, though N and NML may be replaced by M_T and M_R respectively if a more generic analysis is desired.

2.4.1 QCQP Formulation

We begin by demonstrating how one step can be transformed to a QCQP which reveals a per-iterate constraint on the norm of $\mathbf{w}_k$. At the kth iteration, the $\mathbf{s}$-step is

$$\begin{aligned} \min_{\mathbf{s}} \quad & \mathbf{s}^H \mathbf{F}(\mathbf{w}_k, \lambda_k)\mathbf{s} \\ \text{s.t.} \quad & \mathbf{s}^H \mathbf{y}_{\mathbf{w}} = \kappa^* \qquad \|\mathbf{s}\|_2^2 \leq P_o \end{aligned} \tag{2.46}$$

where $\mathbf{y}_{\mathbf{w}} = \mathbf{T}^H \mathbf{w}_k$ is the "response" of the given target channel matrix to observation through filter $\mathbf{w}_k$ and $\mathbf{F}(\lambda_k, \mathbf{w}_k)$ is as defined above. For simplicity of notation, let us define $\mathbf{F}_k = \mathbf{F}(\mathbf{w}_k, \lambda_k)$ for this section. To be optimal, the weight vector must span some of the noise, interference, and/or clutter subspaces, and hence so must $\mathbf{y}_{\mathbf{w}}$. Let us then consider the orthogonal projections $\mathbf{P}_{\mathbf{w}} = \frac{1}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \mathbf{y}_{\mathbf{w}} \mathbf{y}_{\mathbf{w}}^H$ and $\mathbf{P}_{\mathbf{w}}^{\perp} = \mathbf{I}_N - \mathbf{P}_{\mathbf{w}}$, which are onto and away from the span of the observed target response, respectively. This corresponds to a decomposition of $\mathbb{C}^N$ that we can then apply to our signal variable; that is, any $\mathbf{s}$ can be decomposed as $\mathbf{s} = \mathbf{P}_{\mathbf{w}}\mathbf{s} + \mathbf{P}_{\mathbf{w}}^{\perp}\mathbf{s}$. Due to the Capon constraint and the form of $\mathbf{P}_{\mathbf{w}}$, it is clear that for any feasible $\mathbf{s}$

$$\mathbf{P}_{\mathbf{w}}\mathbf{s} = \frac{1}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \mathbf{y}_{\mathbf{w}} \mathbf{y}_{\mathbf{w}}^H \mathbf{s} = \frac{\kappa}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \mathbf{y}_{\mathbf{w}}.$$

Hence, any feasible $\mathbf{s}$ can be represented as

$$\mathbf{s} = \frac{\kappa}{\|\mathbf{y_w}\|_2^2}\mathbf{y_w} + \mathbf{P_w^\perp q} \tag{2.47}$$

for some $\mathbf{q} \in \mathbb{C}^N$.

We can then substitute this form of a feasible solution into (2.46) and find an equivalent problem to solve. After this substitution, and ignoring the effect of a constant term, we have the new equivalent problem

$$\begin{aligned} \min_{\mathbf{q}} \quad & \mathbf{q}^H \mathbf{P_w^\perp F}_k \mathbf{P_w^\perp q} + \frac{2}{\|\mathbf{y_w}\|_2^2}\mathrm{Re}\{\kappa\, \mathbf{q}^H \mathbf{P_w^\perp F}_k \mathbf{y_w}\} \\ \text{s.t.} \quad & \left\|\mathbf{P_w^\perp q}\right\|_2^2 \leq P_o - \frac{|\kappa|^2}{\|\mathbf{y_w}\|_2^2}. \end{aligned} \tag{2.48}$$

The optimal solution to (2.48), if it exists, is readily shown to be

$$\mathbf{q}(\gamma) = \frac{\kappa}{||\mathbf{y_w}||^2}\mathbf{A}(\gamma)\mathbf{y_w} \tag{2.49}$$

where $\mathbf{A}(\gamma) = \mathbf{P_w^\perp}(\mathbf{F}_k + \gamma \mathbf{P_w^\perp})^\dagger \mathbf{P_w^\perp F}_k$ and γ is another Lagrange parameter dictated by the power allocation to the orthogonal component $\mathbf{P_w^\perp q}$. The optimal value of γ is the solution to another polynomial, given by

$$\begin{aligned} & \gamma \left(\frac{|\kappa|^2}{\|\mathbf{y_w}\|_4^2} \mathbf{y_w}^H \mathbf{A}^H(\gamma) \mathbf{P_w^\perp A}(\gamma) \mathbf{y_w} - \left(P_o - \frac{|\kappa|^2}{\|\mathbf{y_w}\|_2^2} \right) \right) = 0 \\ & \gamma \geq 0. \end{aligned} \tag{2.50}$$

Finally, the waveform design solution at the kth iteration of this form of AM is given by

$$\mathbf{s}_k = \kappa \frac{\mathbf{y_w}}{||\mathbf{y_w}||^2} + \mathbf{P_w^\perp q}(\gamma^\star) \tag{2.51}$$

where $\gamma^\star$ is the solution of (2.50).

As with the Lagrange parameter λ_k, obtaining a solution for γ requires a root-finding process; hence, a similar tradeoff between accuracy and speed may be made by setting $\gamma = 0$. Note that this choice necessarily implies that, at each stage, the orthogonal power constraint is inactive.

One might also ask when this constraint has meaning to begin with. Indeed, since it is the only constraint in (2.48), if it is rendered impossible by the previous iterate $\mathbf{w}_k$, then there is no optimal solution whatsoever. Since the left hand side of the constraint is a squared norm, then the right hand side must be at least non-negative, or $P_o - \frac{|\kappa|^2}{\|\mathbf{y}_\mathbf{w}\|_2^2} \geq 0$. After rearranging and placing it in terms of the iterate $\mathbf{w}_k$, we have

$$\frac{|\kappa|^2}{P_o} \leq \left\|\mathbf{T}^H \mathbf{w}_k\right\|_2^2. \tag{2.52}$$

From this, we can identify when the constraint set has meaning – namely, when the power in the signal "matched" to the kth filter iterate $\mathbf{T}^H\mathbf{w}_k$ is at least as much as the desired normalized gain. When $P_o \gg |\kappa|^2$, then the feasible set is easily non-empty. This also gives us a good stopping criterion for the EigenAM problem, since if $\frac{|\kappa|^2}{P_o} = \left\|\mathbf{T}^H \mathbf{w}_k\right\|_2^2$, there is no need for $\mathbf{P}_\mathbf{w}^\perp \mathbf{q}$. That is, there are no more resources available to transmit orthogonally to what is already observed. One can think of this as a minor confirmation of the waterfilling concept shown by [40, 17]. Incidentally, (2.52) also leads to a lower bound on the norm of $\mathbf{w}_k$ Since $\left\|\mathbf{T}^H \mathbf{w}_k\right\|_2^2 \leq \|\mathbf{T}\|_F^2 \|\mathbf{w}_k\|_2^2$ by compatibility of norms, we have $\frac{|\kappa|^2}{P_o \|\mathbf{T}\|_F^2} \leq \|\mathbf{w}_k\|_2^2$. This condition will recur in some sense when we discuss the Slater condition for our relaxed problem in Chapter 3.

2.4.2 SDP Formulations

Next, we consider two approaches to obtain semidefinite programs that are equivalent to the projected QCQP in (2.48).

A Dual SDP

We can first use the dual formulation of (2.48) to obtain a simple SDP in two variables. Let $\alpha \geq 0$ be a Lagrange parameter that takes the place of γ above. The Lagrangian function of (2.48) is given by

$$\mathcal{L}(\mathbf{q}, \alpha) = \mathbf{q}^H \mathbf{P}_{\mathbf{w}}^{\perp}(\mathbf{F}_k + \alpha \mathbf{P}_{\mathbf{w}}^{\perp})\mathbf{P}_{\mathbf{w}}^{\perp}\mathbf{q} + \frac{2}{\|\mathbf{y_w}\|_2^2}\mathrm{Re}\{\kappa\, \mathbf{q}^H \mathbf{P}_{\mathbf{w}}^{\perp}\mathbf{F}_k\mathbf{y_w}\} - \alpha\left(P_o - \frac{|\kappa|^2}{\|\mathbf{y_w}\|_2^2}\right). \quad (2.53)$$

Hence, the dual problem of (2.48) is given by

$$\inf_{\mathbf{q}} \mathcal{L}(\mathbf{q}, \alpha) = \begin{cases} -\frac{|\kappa|^2}{\|\mathbf{y_w}\|_2^4}\mathbf{d}^H\mathbf{D}^{\dagger}(\alpha)\mathbf{d}^H - \alpha\left(P_o - \frac{|\kappa|^2}{\|\mathbf{y_w}\|_2^2}\right) & \mathbf{D}(\alpha) \succeq 0, \mathbf{d} \notin \mathcal{N}(\mathbf{D}(\alpha)) \\ -\infty & \text{otherwise} \end{cases} \quad (2.54)$$

where $\mathbf{D}(\alpha) = \mathbf{P}_{\mathbf{w}}^{\perp}(\mathbf{F}_k + \alpha\mathbf{P}_{\mathbf{w}}^{\perp})\mathbf{P}_{\mathbf{w}}^{\perp}$, $\mathbf{d} = \mathbf{P}_{\mathbf{w}}^{\perp}\mathbf{F}_k\mathbf{y_w}$. Defining another positive variable β and using the approach of Shor [75], we can reformulate the dual as the following SDP:

$$\begin{aligned} \max_{\alpha,\beta} \quad & \beta \\ \text{s.t.} \quad & \begin{bmatrix} \mathbf{D}(\alpha) & \frac{\kappa^*}{\|\mathbf{y_w}\|_2^2}\mathbf{d} \\ \frac{\kappa}{\|\mathbf{y_w}\|_2^2}\mathbf{d}^H & -\alpha\left(P_o - \frac{|\kappa|^2}{\|\mathbf{y_w}\|_2^2}\right) - \beta \end{bmatrix} \succeq 0, \qquad \alpha \geq 0 \end{aligned} \quad (2.55)$$

A Primal SDP

We can also form a semidefinite relaxation of (2.48) by considering the matrix $\mathbf{Q} = \begin{bmatrix} \mathbf{q}\mathbf{q}^H & \mathbf{q} \\ \mathbf{q}^H & 1 \end{bmatrix}$. Given this variable, we can rewrite (2.48) as

$$\begin{aligned} \min_{\mathbf{q}} \quad & \operatorname{tr}\left(\mathbf{Q} \begin{bmatrix} \mathbf{P}_{\mathbf{w}}^{\perp} \mathbf{F}_k \mathbf{P}_{\mathbf{w}}^{\perp} & \frac{\kappa^*}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \mathbf{d} \\ \frac{\kappa}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \mathbf{d}^H & 0 \end{bmatrix} \right) \\ \text{s.t.} \quad & \operatorname{tr}\left(\mathbf{Q} \begin{bmatrix} \mathbf{P}_{\mathbf{w}}^{\perp} & 0 \\ 0 & 0 \end{bmatrix} \right) \leq P_o - \frac{|\kappa|^2}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \\ & \mathbf{Q} = \begin{bmatrix} \mathbf{q}\mathbf{q}^H & \mathbf{q} \\ \mathbf{q}^H & 1 \end{bmatrix} \succeq 0 \end{aligned} \tag{2.56}$$

Dropping the implicit rank-one constraint and ensuring the final diagonal component of $\mathbf{Q}$ is fixed, we have the final relaxation

$$\begin{aligned} \min_{\mathbf{q}} \quad & \operatorname{tr}\left(\mathbf{Q} \begin{bmatrix} \mathbf{P}_{\mathbf{w}}^{\perp} \mathbf{F}_k \mathbf{P}_{\mathbf{w}}^{\perp} & \frac{\kappa^*}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \mathbf{d} \\ \frac{\kappa}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \mathbf{d}^H & 0 \end{bmatrix} \right) \\ \text{s.t.} \quad & \operatorname{tr}\left(\mathbf{Q} \begin{bmatrix} \mathbf{P}_{\mathbf{w}}^{\perp} & 0 \\ 0 & 0 \end{bmatrix} \right) \leq P_o - \frac{|\kappa|^2}{\|\mathbf{y}_{\mathbf{w}}\|_2^2} \\ & \mathbf{Q} \succeq 0, \; Q_{N,N} = 1 \end{aligned} \tag{2.57}$$

where $Q_{i,j}$ is the (i,j)th element of $\mathbf{Q}$.

With these in place, we can tie them together to demonstrate equivalence using strong duality. It can be readily shown that, given an appropriate change of variables, (2.55) is *also* the dual problem for (2.57). Thus, we can claim the following (see, e.g., [76, Appendix B] for the real variable case): Since the primal SDP (2.57) is a relaxation of the QCQP (2.48), then its optimal value is a lower bound on the QCQP's optimal value; that is, $\nu_{QCQP}^o \geq \nu_{SDP-P}^o$.

Furthermore, as a dual problem for (2.57), the dual SDP (2.55)'s value is a lower bound on the primal SDP, which thus implies $\nu^o_{QCQP} \geq \nu^o_{SDP-P} \geq \nu^o_{SDP-D}$.

Recall that a strictly feasible solution is a solution where all inequality constraints are inactive and the solution lies in the relative interior of any other convex constraint set. Furthermore, a strictly feasible solution for (2.48) exists only if (2.52) is a strict inequality. If there is a strictly feasible solution for (2.48), then strong duality between (2.48) and (2.55) holds, and $\nu^o_{SDP-D} = \nu^o_{QCQP}$. Clearly then, $\nu^o_{SDP-P} = \nu^o_{QCQP}$ as well, which implies strong duality holds between (2.48), (2.55), and (2.57). In fact, any rank-one optimal solution of (2.57) is therefore optimal for (2.48). This, in a roundabout way, also proves the equivalences mentioned above. Finally, due to this equivalence, an optimal solution $\mathbf{q}^\star$ from either (2.55) or synthesized from (2.57) may be substituted in (2.51) to obtain the optimal $\mathbf{s}^\star_k$.

2.4.3 Convergence of AM for Bi-convex Optimization

With these multiple perspectives on the EigenAM technique now defined, we now examine the conditions and assumptions under which any AM algorithm converges for a biconvex problem like (2.25).

Briefly, we will show that if limit points exist for a biconvex problem, then those limit points are also stationary. Furthermore, it will be clear that if two algorithms that ostensibly solve the same problem do not enforce the same alternating constraints, then the limit points attained by each algorithm are not guaranteed to coincide. However, if they are somehow equivalent when evaluated (for example, a matrix semidefinite constraint evaluated only on a rank-one matrix), then the limit points will also generally coincide.

Let $\mathbf{x}_k, \mathbf{y}_k$ be the individual iterates generated by an AM optimization of the biconvex problem (2.25). We start by stating two properties of (2.25) that might be considered implicit in Definition 2.1.

Property 2.1. *A functional relationship exists between the iterates at the k-th and $(k-1)$-th iteration.*

More mathematically, we require functions $f_1 : \mathbb{F}^M \to \mathbb{F}^N$ and $f_2 : \mathbb{F}^N \to \mathbb{F}^M$ such that

$$\begin{aligned} \mathbf{x}_k &= f_1(\mathbf{y}_{k-1}) = f_1(f_2(\mathbf{x}_{k-1})) = f_y(\mathbf{x}_{k-1}) \\ \mathbf{y}_k &= f_2(\mathbf{x}_{k-1}) = f_x(\mathbf{x}_{k-1}) \end{aligned} \tag{2.58}$$

where $f_y = f_1 \circ f_2, f_x = f_2$, and we assume the initializers are $\mathbf{x}_0 \in \mathcal{X}_0$, $\mathbf{y}_0 \in \mathcal{Y}_0$. Note that our dependence is predicated on the $\mathbf{x}$ iterate at iteration $k-1$; this is not unique, and we could redefine this to be in terms of $\mathbf{y}_{k-1}$.

Next, we consider the constraint sets involved:

Property 2.2. *The constraint sets for the k-th iteration have an explicit functional dependence on the $(k-1)$-th iterates.*

Mathematically, this is equivalent to saying that given the constraint sets $\mathcal{A}_k, \mathcal{B}_k$ for $\mathbf{x}_k, \mathbf{y}_k$, we can define them as

$$\begin{aligned} \mathcal{A}_k(\mathbf{x}_{k-1}) &= \{\mathbf{x} : (f_y(\mathbf{x}_{k-1}), \mathbf{x}) \in \Upsilon\} \\ \mathcal{B}_k(\mathbf{x}_{k-1}) &= \{\mathbf{y} : (\mathbf{y}, \mathbf{x}_{k-1}) \in \Upsilon\} \end{aligned} \tag{2.59}$$

From here on, we will omit the explicit dependence on the iterate $\mathbf{x}_{k-1}$ for the constraint sets and denote them merely as $\mathcal{A}_k, \mathcal{B}_k$. In a well-defined biconvex problem, we can directly identify the functions f_1, f_2 and the constraint sets $\mathcal{A}_k, \mathcal{B}_k$ – they are the relationships given by (2.28), so long as the sets are not dependent on the iterate index. As such, the problems

we wish to solve *should* have these functional relationships, and therefore we have

$$\mathbf{x}_k = \arg\min_{\mathbf{x}\in\mathcal{A}_k} f(\mathbf{x}, \mathbf{y}_{k-1}), \tag{2.60}$$

$$\mathbf{y}_k = \arg\min_{\mathbf{y}\in\mathcal{B}_k} f(\mathbf{x}_{k-1}, \mathbf{y}). \tag{2.61}$$

For the rest of this section, let us assume that the sequences $\{\mathbf{x}_k\}$ & $\{\mathbf{y}_k\}$ and each of their corresponding subsequences have a finite limit point. Furthermore, we assume that the objective f is uniformly continuous everywhere and its domain includes the Cartesian product of two metrizable supersets.

In order to continue, we require some definitions from topology & other general mathematical disciplines. Recall that a metric space is a set (say, $\mathcal{X}$) equipped with a function $d : \mathcal{X}\times\mathcal{X} \to \mathbb{R}$ known as a metric (or distance) defined between two points in the set which, loosely, outputs only non-negative arguments, is symmetric in its arguments, is subadditive, and, if zero, indicates the points are identical. If the points are in a vector space (as it is in our case), then a norm $\|\cdot\|$ is a typical choice of distance function; that is, for any two points $\mathbf{a}, \mathbf{b}$, the distance $d(\mathbf{a}, \mathbf{b}) = \|\mathbf{a} - \mathbf{b}\|$. Given this distance function, we can define the diameter of a set as follows:

Definition 2.2 (Diameter of a set)**.** *The diameter $D_{\mathcal{A}}$ of a set $\mathcal{A}$ is given by the maximum distance between any two points in the set; that is, $D_{\mathcal{A}} = \sup_{\mathbf{a},\mathbf{b}\in\mathcal{A}} d(\mathbf{a}, \mathbf{b})$.*

We can also define a metric that compares the "closeness" of two sets in the same metric space. This metric, known as the Hausdorff distance, is defined by the following:

Definition 2.3 (Hausdorff distance)**.** *Let $\mathcal{A}, \mathcal{B}$ be two sets in a metric space $(\mathcal{X}, d)$. The Hausdorff distance between $\mathcal{A}$ and $\mathcal{B}$ is defined as*

$$d_{\mathcal{H}}(\mathcal{A}, \mathcal{B}) = \max\left\{\sup_{\mathbf{a}\in\mathcal{A}}\inf_{\mathbf{b}\in\mathcal{B}} d(\mathbf{a}, \mathbf{b}),\ \sup_{\mathbf{b}\in\mathcal{B}}\inf_{\mathbf{a}\in\mathcal{A}} d(\mathbf{a}, \mathbf{b})\right\}. \tag{2.62}$$

Furthermore, let us denote the convex hull and the closure of a set $\mathcal{A}$ as $\mathrm{Conv}(\mathcal{A})$ and $\mathrm{Cl}(\mathcal{A})$, respectively. We now define the sets of all possible iterates originating from a particular set of initializers and operating versions thereof. First, given a set of initializers for $\mathbf{x}$ and $\mathbf{y}$, say, $\mathcal{X}_0$ and $\mathcal{Y}_0$, we define the iterate collections $\mathcal{C}_{\mathcal{X}_0}$ and $\mathcal{D}_{\mathcal{Y}_0}$ as the union of all possible iterate sequences that originate from every $\mathbf{x}_0 \in \mathcal{X}_0$, $\mathbf{y}_0 \in \mathcal{Y}_0$; mathematically, that is:

$$\mathcal{C}_{\mathcal{X}_0} := \bigcup_{\mathbf{x_0} \in \mathcal{X}_0} \mathcal{C}_{\mathbf{x_0}}, \ \mathcal{D}_{\mathcal{Y}_0} := \bigcup_{\mathbf{y_0} \in \mathcal{Y}_0} \mathcal{D}_{\mathbf{y_0}}$$

where

$$\begin{aligned} \mathcal{C}_{\mathbf{x}_0} &= \{\mathbf{x}_0\} \ \cup \ \left\{ \bigcup_{k \in \mathbb{Z}^+} \{\mathbf{x}_k : \mathbf{x}_k = \arg\min_{\mathbf{x} \in \mathcal{A}_k} f(\mathbf{x}, \mathbf{y}_{k-1})\} \right\} \\ \mathcal{D}_{\mathbf{y}_0} &= \{\mathbf{y}_0\} \ \cup \ \left\{ \bigcup_{k \in \mathbb{Z}^+} \{\mathbf{y}_k : \mathbf{y}_k = \arg\min_{\mathbf{y} \in \mathcal{B}_k} f(\mathbf{x}_{k-1}, \mathbf{y})\} \right\} \end{aligned} \tag{2.63}$$

are the iterate sets given a particular initializer $\mathbf{x}_0, \mathbf{y}_0$.

Since we are interested in convergence around these particular initializer sets, we define the two convex closures $\mathring{\mathcal{C}}_{\mathcal{X}_0}, \mathring{\mathcal{D}}_{\mathcal{Y}_0}$ as

$$\mathring{\mathcal{C}}_{\mathcal{X}_0} = \mathrm{Cl}(\mathrm{Conv}(\mathcal{C}_{\mathcal{X}_0})), \ \mathring{\mathcal{D}}_{\mathcal{Y}_0} = \mathrm{Cl}(\mathrm{Conv}(\mathcal{D}_{\mathcal{Y}_0})) \tag{2.64}$$

which are composed of every convex combination of these possible iterates. Furthermore, let us endow the topology $\mathscr{T}_x$ to $\mathring{\mathcal{C}}_{\mathcal{X}_0}$, and the topology $\mathscr{T}_y$ to $\mathring{\mathcal{D}}_{\mathcal{Y}_0}$, both of which are induced by the metric $d(\cdot, \cdot)$.

Let us now consider two arbitrary subsets of these closures, which may not be generated in a particular instantiation of AM: $\bar{\mathcal{C}}_x \subset \mathring{\mathcal{C}}_{\mathcal{X}_0}$ and $\bar{\mathcal{D}}_y \subset \mathring{\mathcal{D}}_{\mathcal{Y}_0}$. However, while the subsets are arbitrary, we are not interested in arbitrary sequences of these subsets. Rather, we are interested in sequences that are themselves capable of being generated by the AM procedure.

That is, we consider the sequences $\{\mathbf{x}_k^c\}$ and $\{\mathbf{y}_k^d\}$ such that

$$\mathbf{x}_k^c = \arg\min_{\mathbf{x}\in\bar{\mathcal{C}}_x} f(\mathbf{x}, \mathbf{y}_{k-1}^d),\ \mathbf{y}_k^d = \arg\min_{\mathbf{y}\in\bar{\mathcal{D}}_y} f(\mathbf{x}_{k-1}^c, \mathbf{y}). \tag{2.65}$$

Given these sequences, we make a critical assumption:

Lemma 2.1. *The sequences* $\{\mathbf{x}_k^c\}$ *and* $\{\mathbf{y}_k^d\}$ *have limit points identical to the sequences* $\{\mathbf{x}_k\}$ *and* $\{\mathbf{y}_k\}$*, denoted as* $\mathbf{x}^\star$ *and* $\mathbf{y}^\star$*, respectively*

We justify this assumption by considering convex closures of the grand iterate collections that are small; that is, their diameters are bounded as $D_{\text{Conv}(\mathcal{C}_{\mathcal{X}_0})} \leq \epsilon_1$ and $D_{\text{Conv}(\mathcal{D}_{\mathcal{Y}_0})} \leq \epsilon_2$, where ϵ_1, ϵ_2 are small, but positive, values. Hence, both sets of sequences will never be "far" from one another in this diameter sense. Furthermore, since both sets of sequences are generated from the same overall constraint sets, we will not run into the difficulty of a mismatch of limit points.

Given these preliminaries, we now prove, loosely, that if a limit point $(\mathbf{x}^\star, \mathbf{y}^\star)$ exists, then this limit point is also stationary.

Theorem 2.2 (Convergence of AM)**.** *Let the AM procedure, endowed with the above properties, run on* (2.25) *given initializers as mentioned above. Assume Lemma 2.1 is satisfied, the minimizers at each stage* k *are unique, and the constraint sets are stationary, i.e.*

$$\lim_{k\to\infty} d_{\mathcal{H}}(\mathcal{A}_k, \mathcal{A}) \to 0$$
$$\lim_{k\to\infty} d_{\mathcal{H}}(\mathcal{B}_k, \mathcal{B}) \to 0$$

for some fixed sets $\mathcal{A}, \mathcal{B}$*. If a limit point* $(\mathbf{x}^\star, \mathbf{y}^\star)$ *exists for any pair of iterate sequences* $\{\mathbf{x}_k\}$ *and* $\{\mathbf{y}_k\}$*, then it is also a stationary point.*

Proof. Since we assume Lemma 2.1 is satisfied, we consider only limit points that also characterize sequences considered in said Lemma.

We first consider the necessity of Property 2.1, which ensures simultaneous convergence in both the $\mathbf{x}$ and $\mathbf{y}$ iterates. If such a functional relationship did not exist, convergence in one would not influence convergence in the other, which might lead to cycling & a failure to converge entirely. Let us consider the following counterexample: Assume we have reached the k_oth iteration of an algorithm without such a dependence, and that the sequence beginning with the iterate $\mathbf{y}_{k_o}$ converges to the limit point $\mathbf{y}^\star$, i.e. $\mathbf{y}_{k_o+n} \to \mathbf{y}^*$ as $n \to \infty$. If the functional relationship did not exist, then the next $\mathbf{x}$ iterate, $\mathbf{x}_{k_o+1}$, is not guaranteed to be in a convergent sequence to the limit point $\mathbf{x}^*$. Furthermore, there will almost certainly then exist a $k > k_o$ where a sequence starting with $\mathbf{y}_k$ will *not* converge to $\mathbf{y}^\star$, due to the functional dependence, and we have reached our contradiction.

Next, we consider the interaction with Property 2.2. Together, these properties enforce the dependency of the constraint sets $\mathcal{A}_k$ and $\mathcal{B}_k$ on the previous iterates; hence, they are not arbitrary, and must generate iterates "cleanly". Furthermore, if the iterates converge, the stationarity of the constraint sets is enforced: $\lim_{k\to\infty} \mathcal{A}_k \to \mathcal{A}$, and $\lim_{k\to\infty} \mathcal{B}_k \to \mathcal{B}$[7]. Thus, if these sets do not converge, then there is no guarantee that the iterates themselves converge, and vice versa. This is also where the uniqueness of the solutions come into play – if the solutions at each iteration are not unique, then the sets generated by them are also not unique and may not converge to the stationary sets. However, we might relax this requirement if the process sequentially generated the same constraint sets and the same optimal value at different points.

[7]Indeed, these stable sets are likely functions of the limit points themselves!

Finally, let us consider the following optimization problem over the grand iterate collections:

$$\min_{\mathbf{x}\in\mathring{\mathcal{C}}_{\mathcal{X}_0},\mathbf{y}\in\mathring{\mathcal{D}}_{\mathcal{Y}_0}} f(\mathbf{x},\mathbf{y}) \tag{2.66}$$

We state that the solution to (2.66) is also the desired limit point $(\mathbf{x}^\star,\mathbf{y}^\star)$, and prove by contradiction. Assume that this is not true, and there is another limit point generated by this minimization; say, $(\mathbf{x}^\blacklozenge,\mathbf{y}^\blacklozenge)=\inf_{\mathbf{x}\in\mathring{\mathcal{C}}_{\mathcal{X}_0},\mathbf{y}\in\mathring{\mathcal{D}}_{\mathcal{Y}_0}} f(\mathbf{x},\mathbf{y})$.

By construction, this limit point must be a member of the convex hull, which means there exist weight sets α_k and β_k (where $\sum_k \alpha_k = 1, \sum_k \beta_k = 1$)

$$f(\mathbf{x}^\blacklozenge,\mathbf{y}^\blacklozenge) = f\left(\sum_k \alpha_k\mathbf{x}_k,\ \sum_k \beta_k\mathbf{y}_k\right).$$

for any given sequence of iterates $\{\mathbf{x}_k\},\{\mathbf{y}_k\}$. Since the problem is biconvex, we have the following inequality chain:

$$\begin{aligned} f(\mathbf{x}^\blacklozenge,\mathbf{y}^\blacklozenge) &\le \sum_k \alpha_k f\left(\mathbf{x}_k,\ \sum_k \beta_k\mathbf{y}_k\right) \\ &\le \alpha \inf_{\mathbf{x}_k} f\left(\mathbf{x_k},\ \sum_k \beta_k\mathbf{y_k}\right) + (1-\alpha)\sup_{\mathbf{x}_k} f\left(\mathbf{x}_k,\ \sum_k \beta_k\mathbf{y}_k\right) \\ &\le \inf_{\mathbf{x_k}} f\left(\mathbf{x}_k,\ \sum_k \beta_k\mathbf{y}_k\right) \end{aligned}$$

where the second inequality follows from Carathéodory's theorem on the real axis, and the final inequality is a special case when $\alpha = 1$. If we start with the final inequality and repeat the procedure for the $\mathbf{y}$ iterates, we finally have

$$f(\mathbf{x}^\blacklozenge,\mathbf{y}^\blacklozenge) \le \inf_{\mathbf{x}_k,\mathbf{y}_k} f(\mathbf{x}_k,\mathbf{y}_k) = f(\mathbf{x}^\star,\mathbf{y}^\star).$$

By Lemma 2.1, this is a contradiction, because we have considered a different subset from the same convex hull, but achieved a different ("better") limit point. Therefore, $(\mathbf{x}^{\blacklozenge}, \mathbf{y}^{\blacklozenge}) = (\mathbf{x}^{\star}, \mathbf{y}^{\star})$.

We can now apply the same technique as in [77, Prop. 2.7.1] to the problem in (2.66) by using the AM algorithm on this modified optimization problem and demonstrate that limit point is also a stationary point. □

Note that this proof indirectly demonstrates that if the stationary sets that two AM algorithms converge to are different, then not only are the limit points different, but since there may be a mismatch between actual problem and "solved" problem, convergence may suffer. We leave an illustration of this for the next section. Finally, we note that the final part of the proof shows that even if Lemma 2.1 were not true, we would still have $f(\mathbf{x}^{\blacklozenge}, \mathbf{y}^{\blacklozenge}) = f(\mathbf{x}^{\star}, \mathbf{y}^{\star})$ for (potentially non-equal) limit points $(\mathbf{x}^{\blacklozenge}, \mathbf{y}^{\blacklozenge})$ and $(\mathbf{x}^{\star}, \mathbf{y}^{\star})$, since they originate from the same initializer sequence.

2.5 Numerical Analysis

In this section, we validate the concurrence of AM and its SDP formulation given in Section 3 through a simulated example. Additionally, we will show that any of these algorithms, if appropriately rescaled to satisfy a power equality constraint, both can outperform a version of the technique in [44] and asymptotically approaches the optimal value of the relaxed problem that we will explain in Chapter 3, and was originally given in [65].

In both of these cases, we will let $\kappa = P_o = 1$ for two reasons: first, to demonstrate the inherent tradeoff in restricting $\lambda = 0$ when the aforementioned condition is not satisfied, and second, to facilitate comparison with the algorithm in [44] where these particular values are implicitly assumed.

The scenario presented to the solvers is as follows: We assume the radar transmits $L = 8$ pulses of $N = 8$ samples each, and receives them with a uniform linear array of $M = 5$ elements on a moving platform. We assume that the signal independent noise covariance matrix $\mathbf{R_n}$ is a Toeplitz matrix with the (i, j)th element equal to $\exp(-0.005|i-j|)$. A single interferer is present at the azimuth elevation pair (0.3941, $\pi/3$) radians with a fast time-slow time correlation given by $\exp(0.02n)$. The target is located at the azimuth-elevation pair (0, $\pi/3$) radians, moving at a relative normalized Doppler of -0.1443 to the platform. The clutter is described by 25 statistically independent patches whose nominal phase centers are at azimuths linearly spaced in the interval $[-\pi/2, \pi/2]$ radians and an elevation angle of 0.3 radians. Additional parameters are as in [8].

First, we compare the EigenAM procedure (henceforth, AM) with the theoretically equivalent QCQP in (2.48) and projected SDP in (2.57). Rank-one solutions, if they were not immediately available, were obtained by the rank-equivalence procedure of [73]. All algorithms were terminated after 20 iterations and the parameters λ, γ were directly determined by a line search at each iteration. Results of an example convergence run are demonstrated in Figure 2.3, where we compare them with the optimal value of the joint relaxed biquadratic program (RBQP) proposed in [65], which we will describe more fully in Chapter 3. Clearly, when fully implemented, AM, the QCQP method, and Projected SDP (the first set of curves) coincide exactly and converge rather quickly to a certain limit point in a manner reminiscent of Theorem 2.2. Additionally, there is an expected gap between the optimal values of the original problem and its relaxation. We note that this gap occurs partially because the $\mathbf{s}_k$ iterates converge to strictly feasible points, as we expected. The second set of curves shows the equivalent cost if the power constraint was always satisfied with equality, which we enforce by rescaling the optimal $(\mathbf{w}, \mathbf{s})$ pairs at each iteration such that $||\mathbf{s}_k||^2 = P_o$ and the Capon constraint is maintained. This is important because of the aforementioned strict feasibility. That is, the iterative solvers we propose prefer to sequentially lower the power until convergence, which makes comparison with techniques that enforce power limi-

tations with equality difficult. As noted above, however, these rescaled iterates may not be KKT points of the originally considered problem. Unsurprisingly, the three algorithms again coincide; however, unlike the inequality constrained iterates, the rescaled iterates clearly converge to an asymptote given by the optimal value of the relaxed problem, despite not being guaranteed to satisfy the original problem's KKTs. Hence, we have verified the equivalence and convergence properties of the aforementioned problems.

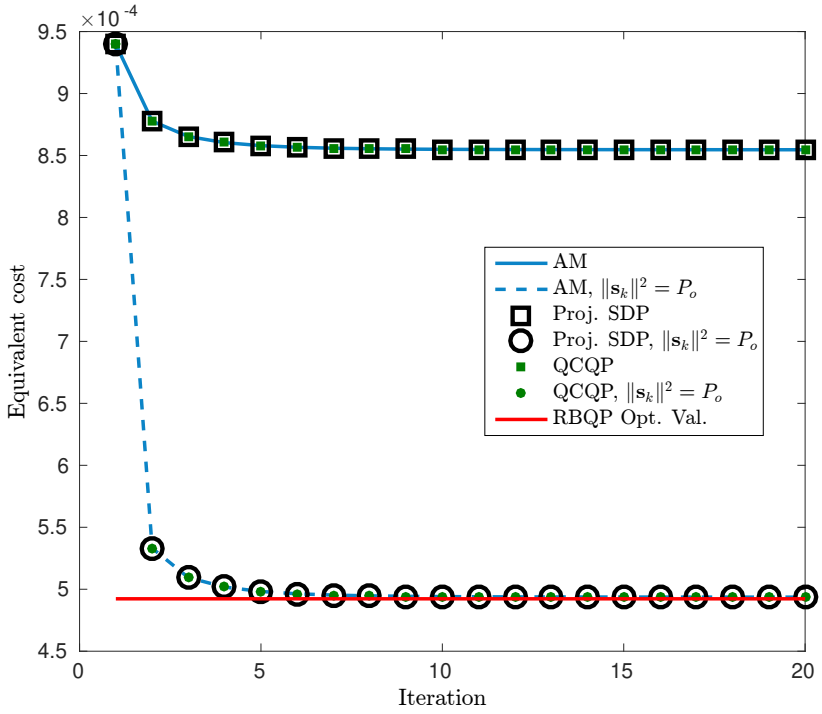

Figure 2.3: Convergence comparison of AM, QCQP, & projected SDP (unscaled and scaled).

Next, we compare the convergence properties of the equivalent methods above with two other algorithms. The first algorithm is a suboptimal form of projected SDP where the Lagrange multiplier λ is set to zero at each step. This solution is suboptimal since $\lambda = 0$ does not necessarily satisfy the KKTs when $P_o = \kappa$. Like the optimal version above, this version is terminated after 20 iterations. The second algorithm is a modified version of the algorithm in [44] (henceforth, AA2) where the similarity constraint is removed. While AA2 nominally

maximizes SINR, we can make a reasonable comparison of the equivalent objective for the other problems by examining the reciprocal.

In this case, we set the convergence parameter $\epsilon = 1$ – as convergence tolerances go, this may seem rather large, but it is our experience that even this can lead to extended runtimes for the algorithm. For comparison, we examine the mean objective value $\hat{\nu}$ attained by each algorithm (rescaled, if necessary) after 20 iterations or, in the case of AA2, when the algorithm was deemed to have converged if it occurred before 20 iterations. This was achieved by averaging the results of 50 Monte Carlo trials, with each algorithm initialized with a given randomly-generated signal for each trial. The results are presented in Table 2.1, where we express the resultant objective value in multiples of the optimal value of RBQP $\nu^{\star}_{RBQP}$ (here, 4.21×10^{-4}), which we have previously established as an empirical convergence asymptote.

Table 2.1: Mean convergence comparison of algorithms.

Algorithm	$\hat{\nu}/\nu^{\star}_{RBQP}$
AM/QCQP	2.12
SDP	2.06
SDP, $\lambda = 0$	1.33
AM/QCQP/SDP Rescaled	1.01
AA2	2.16

It is clear that AA2 converges to an effectively higher limit point than any of the generally equivalent forms above or their more comparable rescaled counterparts. This is also true for the suboptimal projected SDP; however, even this method outperforms AA2 in our experiments. This is critical, because AA2 always enforces the power-constraint with equality and hence does not have a per-iterate reduction in power to account for convergence issues. Furthermore, suboptimal projected SDP requires minimal rescaling (that is, $||\mathbf{s}_k||^2 \approx P_o$ for all iterations). Hence, it may be a more viable option than suboptimal AM for times when $P_o \approx \kappa$. This version of the algorithm frequently has iterates exceed the power constraint if λ is forced to zero and, when these iterates are rescaled, can produce solutions with higher

costs than any of the methods listed herein. See [78] for similar behavior after enforcing feasibility constraints in waveform design.

Finally, AA2 has the critical flaw that it incorrectly enforces the Capon constraint – namely, that it does not enforce it at all in the $\mathbf{s}$-step. This effective switch between objective functions and constraint sets destroys even the loose convergence properties described in Theorem 2.2. We will see that this has further consequences for performance in Chapter 4.

Chapter 3

Joint Relaxations for Power-Constrained Designs

While the iterative/sequential techniques examined in Chapter 2 are fairly popular and endlessly reinvestigated with new constraints, it is worthwhile to ask if the design process can be accomplished in a truly joint fashion, with minimal "guessing" of an initializer.

In this chapter, we propose a convex relaxation technique for simultaneous design of both transmit and receive resources. It is, at first blush, a very simple trick that can ostensibly lead to very large solver complexities. Indeed, a version of this relaxation method has been investigated in the biquadratic optimization literature; however, due to the special nature of the considered problem, it manifests as a linear semidefinite program. Our method considers the potentially more difficult quadratic semidefinite program (QSDP), which arises from the clutter tensor mentioned in Chapter 2 in the context of BQP. Due to the convex structure and unique separability properties of the objective and constraint functions, though, we can efficiently and analytically characterize the solutions (and thus their resulting quality) obtained through numerical solvers.

We will show this analytical process results in problem-dependent guarantees on rank, solution type, and other relevant quantities. Furthermore, it also results in multiple viable solution paths and bounds on optimality, which we will also examine in detail. This process can also provide simple conditions to ensure rank-one optimality in terms of the dual variables.

The organization of the rest of this chapter is as follows. Chapter 3.1 will discuss our relaxation technique, which combines the notion of completely-positive operators and the theory of QSDPs to transform the non-convex joint problem (2.35) into a convex problem. This will leverage the signal-independent nature of the channel matrices. Chapter 3.2 will discuss constraint qualifications and an initial analysis of the Karush-Kuhn-Tucker conditions, which are necessary and sufficient for solutions of the convex problem under consideration. The KKT analysis will continue in more specific terms for the power-constrained problem in Chapter 3.3 and include both power-bounded and non-power-bounded solution paths. Chapter 3.4 will demonstrate how the KKTs naturally coincide with a waterfilling-like interpretation of the design problem, providing a robust linkage to previous continuous-time design methodologies for signal-dependent interference sources. Finally, this chapter will conclude with Chapter 3.5, where the utility and superiority of our technique will be demonstrated through representative simulations of nominal waveform-adaptive STAP scenarios.

3.1 Relaxations & Quadratic Semidefinite Programs

As mentioned above, while the signal-filter problem (2.30) is a non-convex biquadratic program, it can at least be decoupled into quadratic programs that can then be relaxed to linear SDPs or solved through other reduced means. However, it seems, in some sense, misguided to completely rely on these alternating procedures because their convergence often depends on selecting "good" initializers in one of the variables. Furthermore, choosing initializers ba-

sed on pre-existing waveform or filter libraries can lead to the problem of producing merely small perturbations of the initial point, which is hardly a design problem at all.

Our path to avoid this quagmire focuses on using semidefinite relaxation on the joint vector problem (2.35). In this section, we describe the fundamentals of and the process behind our relaxation concept, which will carry through the rest of the dissertation.

3.1.1 Completely Positive Clutter Operators

We begin first by discussing a particular part of the theory of linear operators which relate to the characteristics of the signal-independent clutter tensor-operators $\mathcal{C}$ and $\widetilde{\mathcal{C}}$.

Let us consider a linear operator[1] $\mathcal{Q} : \mathbb{C}^{n\times n} \to \mathbb{C}^{m\times m}$; that is, given a matrix $\mathbf{M} \in \mathbb{C}^{n\times n}$, $\mathcal{Q}(\mathbf{M}) \in \mathbb{C}^{m\times m}$. Operators like $\mathcal{Q}$ are said to be *preserving* in some sense if, for every input $\mathbf{M}$ in some subset of $\mathbb{C}^{n\times n}$, the output matrix $\mathcal{Q}(\mathbf{M})$ is in the equivalent subset of $\mathbb{C}^{m\times m}$. We now define two important types of preserving operators: Hermitian-preserving and (positive) semidefinite-preserving.

Definition 3.1 (Hermitian-preserving operators)**.** *A linear operator $\mathcal{Q} : \mathbb{C}^{n\times n} \to \mathbb{C}^{m\times m}$ is* Hermitian-preserving *if, for every matrix $\mathbf{M} \in \mathbb{H}^n$, $\mathcal{Q}(\mathbf{M}) \in \mathbb{H}^m$.*

Example 3.1. *Let $\mathbf{A}$ be a 2×2 Hermitian matrix, given by*

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{12}^* & a_{22} \end{bmatrix}.$$

The operator $\mathcal{D} : \mathbb{C}^{2\times 2} \to \mathbb{C}^{2\times 2}$ that zeroes out the diagonal, producing

$$\mathcal{D}(\mathbf{A}) = \begin{bmatrix} 0 & a_{12} \\ a_{12}^* & 0 \end{bmatrix},$$

[1]The reason behind this notation choice will become relevant in the next section.

is a Hermitian-preserving operator.

The definition for semidefinite-preserving operators, which are naturally also Hermitian-preserving, is functionally identical:

Definition 3.2 (Semidefinite-preserving operators)**.** *A linear operator $\mathcal{Q} : \mathbb{C}^{n\times n} \to \mathbb{C}^{m\times m}$ is* semidefinite-preserving *if, for every matrix $\mathbf{M} \in \mathbb{H}_+^n$, $\mathcal{Q}(\mathbf{M}) \in \mathbb{H}_+^m$. That is, if $\mathbf{M} \succeq 0$, $\mathcal{Q}(\mathbf{M}) \succeq 0$ as well.*

Semidefinite-preserving operators are often also known as "positive operators" in the linear algebra literature.

One might think that semidefinite-preserving operators are the operator analogs of PSD matrices; however, Choi [79] proved that it is in fact a subset of these operators known as *completely positive* operators that fill this role. Completely positive operators are, essentially, operators that preserve semidefiniteness even when extended to block matrices whose blocks are elements of $\mathbb{H}^n$. Define the set $\mathbb{C}^{p\times p}(\mathbb{C}^{n\times n}) \subset \mathbb{C}^{pn\times pn}$ to be the set of $p \times p$ block matrices whose blocks are $n \times n$ matrices. A more technical definition of complete positivity then follows as:

Definition 3.3. *Let the block matrix sets be defined as above. A linear operator $\mathcal{Q} : \mathbb{C}^{n\times n} \to \mathbb{C}^{m\times m}$ is* p-positive *if the operator $\mathbf{1}_p \otimes \mathcal{Q} : \mathbb{C}^{p\times p}(\mathbb{C}^{n\times n}) \to \mathbb{C}^{p\times p}(\mathbb{C}^{m\times m})$ is semidefinite-preserving/positive. $\mathcal{Q}$ is* completely positive *if it is p-positive for all $p \in \mathbb{Z}^+$.*

While the definition above is compact, it can be hard to visualize how the repetition operator $\mathbf{1}_p \otimes \mathcal{Q}$ works. We provide the following example:

Example 3.2 (The operator $\mathbf{1}_p \otimes \mathcal{Q}$ & p-positivity)**.** *Assume the dimensions from Definition 3.3. Let $p = 2$, and assume we have the four matrices $\mathbf{M}_{11}, \mathbf{M}_{12}, \mathbf{M}_{21}, \mathbf{M}_{22} \in \mathbb{H}_+^n$. Then,*

the block matrix $\widetilde{\mathbf{M}} \in \mathbb{C}^{2\times 2}(\mathbb{C}^{n\times n})$ *given by*

$$\widetilde{\mathbf{M}} = \begin{bmatrix} \mathbf{M}_{11} & \mathbf{M}_{12} \\ \mathbf{M}_{21} & \mathbf{M}_{22} \end{bmatrix}$$

is also PSD. Let $\mathbf{P}$ *be the result of applying* $\mathbf{1}_2 \otimes \mathcal{Q}$ *to the above block matrix, or* $\mathbf{P} = [\mathbf{1}_2 \otimes \mathcal{Q}](\widetilde{\mathbf{M}})$*. Then,* $\mathbf{P}$ *is also representable as*

$$\mathbf{P} = \begin{bmatrix} \mathcal{Q}(\mathbf{M}_{11}) & \mathcal{Q}(\mathbf{M}_{12}) \\ \mathcal{Q}(\mathbf{M}_{21}) & \mathcal{Q}(\mathbf{M}_{22}) \end{bmatrix}.$$

If $\mathbf{P}$ *is also PSD, then* $\mathcal{Q}$ *is 2-positive, by definition. Should this positivity extend to any dimension of blocking, then* $\mathcal{Q}$ *is completely positive.*

The definition provided above is quite cumbersome and is incredibly difficult to check for any given operator. Thankfully, Choi also provided a theorem that allows for immediate characterization of completely positive operators in terms of their defined forms:

Theorem 3.1 (Theorem 1, [79])**.** *Let* $\mathcal{Q} : \mathbb{C}^{n\times n} \to \mathbb{C}^{m\times m}$ *be a linear operator. Then,* $\mathcal{Q}$ *is completely positive if and only if there exist* $m \times n$ *matrices* $\mathbf{V}_i$ *such that*

$$\mathcal{Q}(\mathbf{A}) = \sum_i \mathbf{V}_i \mathbf{A} \mathbf{V}_i^H \tag{3.1}$$

for all $\mathbf{A} \in \mathbb{C}^{n\times n}$*.*

Notice that the $\mathbf{V}_i$ do not necessarily need to form a linearly-independent set, though [79, Remarks 4 and 6] note that linear independence ensures the representation in (3.1) is "canonical," if not unique, and that no more than nm such matrices are needed. It is immediately clear by the form in Theorem 3.1 that any completely positive operator is also operator convex by definition. As such, any affine composition with or quadratic form resulting from

$\mathcal{Q}$ is also convex. There are multiple theorems in the literature that are essentially corollarys of Theorem 3.1. The most useful of these rely on the *matricization* $\langle\mathcal{Q}\rangle \in \mathbb{C}^{m^2\times n^2}$ of an operator $\mathcal{Q}$, which is implicitly defined as $\text{vec}(\mathcal{Q}(\mathbf{A})) = \langle\mathcal{Q}\rangle \text{vec}(\mathbf{A})$. Given this definition, Poluikis and Hill [80] provide the following corollary:

Corollary 3.1 (Theorem (1.2), [80])**.** *$\mathcal{Q}$ is completely positive iff there exist $\mathbf{V}_1, \ldots, \mathbf{V}_s \in \mathbb{C}^{m\times n}$ such that there is a matricization $\langle\mathcal{Q}\rangle = \sum_{i=1}^{s} \mathbf{V}_i^* \otimes \mathbf{V}_i$.*

The proof of this corollary can be directly obtained by applying the rules of vectorization to (3.1). There are also two further matrices one can define, given the reshaping operators Ψ and Γ [80], which map $\mathbb{C}^{m^2\times n^2}$ to $\mathbb{C}^{mn\times mn}$. For a completely positive operator $\mathcal{Q}$, we have the "vectorized" matrices (which are elements of $\mathbb{H}_+^{nm}$)

$$\Psi(\langle\mathcal{Q}\rangle) = \sum_{i=1}^{s} \text{vec}(\mathbf{V}_i)\,\text{vec}(\mathbf{V}_i)^H \tag{3.2}$$

$$\Gamma(\langle\mathcal{Q}\rangle) = \sum_{i=1}^{s} \text{vec}(\mathbf{V}_i^H)\,\text{vec}(\mathbf{V}_i^H)^H. \tag{3.3}$$

The positive semidefiniteness of these matrices is self-evident, as they are sums of positive semidefinite rank-one matrices. Recall that for any $m \times n$ matrix, $\text{vec}(\mathbf{V}^T) = \mathbf{K}_{m,n}\,\text{vec}(\mathbf{V})$ where $\mathbf{K}_{m,n}$ is the commutation matrix [81]. Then, the two matrix representations are related by $\Gamma(\langle\mathcal{Q}\rangle) = \mathbf{K}_{NML,N}\Psi(\langle\mathcal{Q}\rangle)^*\mathbf{K}_{NML,N}^T$.

The theory we have mentioned so far has significant connections to quantum physics and computing, particularly in terms of defining observability of a quantum state. In this realm, the matrices $\mathbf{V}_i$ are known as *Kraus operators*, and define a transfer function from the excitation of a quantum state to its observed result. The matrices $\Psi(\langle\mathcal{Q}\rangle), \Gamma(\langle\mathcal{Q}\rangle)$ are often known in these contexts, appropriately, as the *Choi matrices* of the operator. Henceforth, when we refer to the Choi matrix of a given operator $\mathcal{Q}$, we will mean the matrix $\Psi(\langle\mathcal{Q}\rangle)$.

In our case, however, we can immediately see the clear connection to the signal-independent channel matrices, the clutter tensor $\mathcal{C} : \mathbb{C}^{N\times N} \to \mathbb{C}^{NML\times NML}$ and its expansion $\widetilde{\mathcal{C}} : \mathbb{C}^{J\times J} \to \mathbb{C}^{J\times J}$. Recall that we originally defined these operators implicitly in terms of the clutter covariances $\mathbf{R}_{\mathsf{c}}(\mathbf{s}), \widetilde{\mathbf{R}}_{\mathsf{c}}(\mathbf{b})$, which is their evaluation on rank-one matrices:

$$\mathbf{R}_{\mathsf{c}}(\mathbf{s}) = \mathcal{C}(\mathbf{s}\mathbf{s}^H) \tag{3.4}$$

$$\widetilde{\mathbf{R}}_{\mathsf{c}}(\mathbf{b}) = \mathbf{\Psi}_W^T \mathbf{R}_{\mathsf{c}}(\mathbf{\Psi}_S \mathbf{b})\mathbf{\Psi}_W = \widetilde{\mathcal{C}}(\mathbf{b}\mathbf{b}^H) \tag{3.5}$$

Given arbitrary rank matrices $\mathbf{S} \in \mathbb{H}_+^N, \mathbf{B} \in \mathbb{H}_+^J$, though, we can explicitly define these operators in a way that aligns them with Theorem 3.1:

$$\mathcal{C}(\mathbf{S}) = \sum_{q=1}^{Q} \bar{R}_q^{\gamma} \mathbf{\Gamma}_q \mathbf{S} \mathbf{\Gamma}_q^H \tag{3.6}$$

$$\widetilde{\mathcal{C}}(\mathbf{B}) = \sum_{q=1}^{Q} \overline{R}_q^{\gamma} \widetilde{\mathbf{\Gamma}}_q \mathbf{B} \widetilde{\mathbf{\Gamma}}_q^H \tag{3.7}$$

This leads us to a theorem specific to the joint trans-receive design problem:

Theorem 3.2. *The clutter operators $\mathcal{C}$ and $\widetilde{\mathcal{C}}$ are both completely positive operators.*

Proof: If we set $n = N,\ m = NML,\ i = q$ and let $\mathbf{V}_i = \sqrt{\bar{R}_q^{\gamma}}\mathbf{\Gamma}_q$, then $\mathcal{C}$ is a completely positive operator by Theorem 3.1. Similarly, letting $n = m = J$ and $\mathbf{V}_i = \sqrt{\bar{R}_q^{\gamma}}\widetilde{\mathbf{\Gamma}}_q$ proves the complete positivity of $\widetilde{\mathcal{C}}$. ■

Note that $\widetilde{\mathcal{C}}(\mathbf{B}) = \mathbf{\Psi}_W^T \mathcal{C}(\mathbf{\Psi}_S \mathbf{B} \mathbf{\Psi}_S^T)\mathbf{\Psi}_W$, but one must be careful with naively applying this to a given optimization problem, because the correspondence to the covariances above notionally applies only when $\operatorname{rank}(\mathbf{B}) = 1$. We will deal with this issue later in the section.

Since $\mathcal{C}, \widetilde{\mathcal{C}}$ are both completely positive, we can also identify their Choi matrices:

$$\mathbf{C}_V = \sum_{q=1}^{Q} \bar{R}_q^{\gamma} \operatorname{vec}(\mathbf{\Gamma}_q) \operatorname{vec}(\mathbf{\Gamma}_q)^H \tag{3.8}$$

$$\widetilde{\mathbf{C}}_V = \sum_{q=1}^{Q} \bar{R}_q^{\gamma} \operatorname{vec}(\widetilde{\mathbf{\Gamma}}_q) \operatorname{vec}(\widetilde{\mathbf{\Gamma}}_q)^H \tag{3.9}$$

where we shall denote $\mathbf{C}_V = \Psi(\langle \mathcal{C} \rangle) \in \mathbb{H}_+^{N^2ML}, \widetilde{\mathbf{C}}_V = \Psi(\langle \widetilde{\mathcal{C}} \rangle) \in \mathbb{H}_+^{J^2}$, with the subscript V emphasizing that they arise from **v**ectorizing the clutter channel matrices. Recalling that $\widetilde{\mathbf{\Gamma}}_q = \mathbf{\Psi}_W^T \mathbf{\Gamma}_q \mathbf{\Psi}_S$, the rules of vec lead to a simpler equivalence between $\mathbf{C}_V$ and $\widetilde{\mathbf{C}}_V$:

$$\widetilde{\mathbf{C}}_V = (\mathbf{\Psi}_S \otimes \mathbf{\Psi}_W)^T \mathbf{C}_V (\mathbf{\Psi}_S \otimes \mathbf{\Psi}_W) \tag{3.10}$$

Since $\operatorname{rank}(\mathbf{\Psi}_S \otimes \mathbf{\Psi}_W) = \operatorname{rank}(\mathbf{\Psi}_S) \operatorname{rank}(\mathbf{\Psi}_W) = N^2ML$, $\mathbf{\Psi}_S \otimes \mathbf{\Psi}_W$ is of full column rank, and therefore $\operatorname{rank}(\widetilde{\mathbf{C}}_V) = \operatorname{rank}(\mathbf{C}_V)$. Furthermore, we can link $\mathbf{R}_{\mathbf{c}}(\mathbf{s})$ to $\mathbf{C}_V$ by observing that $\mathbf{R}_{\mathbf{c}}(\mathbf{s}) = (\mathbf{s}^* \otimes \mathbf{I}_{NML})^H \mathbf{C}_V (\mathbf{s}^* \otimes \mathbf{I}_{NML})$. Incidentally, if the channel matrices $\mathbf{\Gamma}_q$ are random variables instead of deterministic, then $\mathbf{C}_V$ serves as the covariance matrix for $\operatorname{vec}(\mathbf{\Gamma}_q)$.

One might ask if the given clutter channel matrices form the "minimal" or "canonical" representation, as mentioned in by Choi. Thankfully, matrix decompositions and the nature of observability in apertures provide us with a simple solution. Recall that the rank of the clutter "subspace" (*i.e.*, the rank of $\mathbf{R}_{\mathbf{c}}(\mathbf{s})$ and thus $\mathbf{C}_V$) is limited by both the physical extent and nature of non-target scatterers and our overall ability to observe this state of nature. For side-looking airborne arrays, the well-known Brennan rule [82, 60] is a reasonable approximation if certain conditions hold, but as [83] showed, a more robust result is obtained by applying the Landau-Pollak theorem. In any case, let us assume that the "true" rank of the clutter is $Q_{\text{eff}} \leq Q \leq N^2ML$ (the subscript denoting the *effective* number of clutter patches). Then, we can use the economy eigendecomposition to find two equivalent

representations of $\mathbf{C}_V$, namely:

$$\mathbf{C}_V = \breve{\mathbf{U}}_C \mathbf{D}_C \breve{\mathbf{U}}_C^H = \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \breve{\mathbf{u}}_i \breve{\mathbf{u}}_i^H.$$

Here, $\breve{\mathbf{U}}_C \in \mathbb{C}^{N^2ML \times Q_{\text{eff}}}$ is the matrix that forms the basis for the Q_{eff}-dimensional vectorized clutter subspace, whose ith column is the eigenvector $\breve{\mathbf{u}}_i \in \mathbb{C}^{N^2ML}$, $i \in \{1, \ldots, Q_{\text{eff}}\}$. $\mathbf{D}_C \in \mathbb{C}^{Q_{\text{eff}} \times Q_{\text{eff}}}$ is a diagonal matrix whose (i,i)th element is the nonzero eigenvalue $\sigma_{\mathsf{c},i}^2$. We can also consider the full eigendecomposition $\mathbf{C}_V = \breve{\mathbf{U}} \mathbf{D} \breve{\mathbf{U}}$. Here, the unitary matrix $\breve{\mathbf{U}} = [\breve{\mathbf{U}}_C \;\; \breve{\mathbf{U}}_N]$, where $\breve{\mathbf{U}}_C$ is as above and $\breve{\mathbf{U}}_N$ collects the $N^2ML - Q_{\text{eff}}$ eigenvectors in the nullspace of $\mathbf{C}_V$, which we can regard as the vectors $\breve{\mathbf{u}}_i, i \in \{Q_{\text{eff}}+1, \ldots, N^2ML\}$. $\mathbf{D}$ is just the direct sum of $\mathbf{D}_C$ and a $N^2ML - Q_{\text{eff}} \times N^2ML - Q_{\text{eff}}$ all-zeros matrix. This decomposition also provides us with an alternative representation of the signal-dependent clutter covariance matrix $\mathbf{R}_{\mathsf{c}}(\mathbf{s})$. Let us assume that there is a set of N^2ML matrices $\mathbf{U}_i \in \mathbb{C}^{NML \times N}$ whose vectorizations are $\text{vec}(\mathbf{U}_i) = \breve{\mathbf{u}}_i$ – that is, they correspond to the eigenvectors of $\mathbf{C}_V$. Then, the signal-dependent clutter covariance can also be given by

$$\mathbf{R}_{\mathsf{c}}(\mathbf{s}) = \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \mathbf{U}_i \mathbf{s}\mathbf{s}^H \mathbf{U}_i^H. \tag{3.11}$$

We note that $\mathbf{U}_i^H \mathbf{U}_i = \frac{1}{N}\mathbf{I}_N$ for all i, not just those in the clutter indices. Thus, $\sqrt{N}\mathbf{U}_i$ is a rank-N partial isometry, and they form the the linearly independent matrix set we mentioned generically above.

3.1.2 The Quadratic Semidefinite Program

Previously, we alluded to the fact that a biquadratic program, while NP-hard and non-convex, could be efficiently relaxed to a convex problem, assuming certain structure. It is here that we introduce this convex problem type, known as the quadratic semidefinite

problem/program (QSDP) [2], which will form the basis of the rest of our contributions in this dissertation. A QSDP is, in some sense, the "simplest" extension of the well-known and studied linear semidefinite problem (SDP), which we saw previously as a signal or filter-only relaxation in the iterative techniques of Chapter 2. As we will soon see, the QSDP shares some properties with the standard SDP.[3] However, the literature is sparser on general QSDP theory and fewer guarantees made be made about the properties of their solutions, absent any other preidentified structure.

Traditionally, QSDPs have been studied for real matrices only, but most of the theory presented here can be extended to Hermitian matrices with little effort. We will leave discussion of this extension for later. For now, we follow the development of this theory as presented by Li [56] in his dissertation; earlier examinations of this problem type can be found in [84, 85]. The general form of a QSDP is given by

$$\begin{aligned} \min_{\mathbf{X}} \quad & \frac{1}{2}\mathbf{X} \bullet \mathcal{Q}(\mathbf{X}) + \mathbf{C} \bullet \mathbf{X} \\ \text{s.t.} \quad & \mathcal{A}_E(\mathbf{X}) = \mathbf{b}_E \qquad \mathcal{A}_I(\mathbf{X}) \leq \mathbf{b}_I \qquad \mathbf{X} \in \mathbb{S}^n_+ \end{aligned} \tag{3.12}$$

where $\mathcal{Q} : \mathbb{S}^n \to \mathbb{S}^n$ is a self-adjoint linear operator, and $\mathcal{A}_E : \mathbb{S}^n \to \mathbb{R}^{m_E}$ and $\mathcal{A}_I : \mathbb{S}^n \to \mathbb{R}^{m_I}$ are linear maps to real equality and inequality constraints, respectively, given by $\mathbf{b}_E \in \mathbb{R}^{m_E}$ and $\mathbf{b}_E \in \mathbb{R}^{m_I}$. Here, $\leq$ is shorthand for an inequality that is true for every element in the vector.

[2] More rarely, it is known as the semidefinite quadratic problem (SQP) (see, *e.g.*)

[3] Depending on the structure of a solution, the solution to a linear SDP and a QSDP may in fact coincide. This observation will come into play later in this chapter.

It can be shown that the dual of (3.12) is given by

$$\begin{aligned} \max_{\mathbf{W},\,\boldsymbol{\Sigma},\,\mathbf{y}_E,\,\mathbf{y}_I} \quad & -\frac{1}{2}\mathbf{W}\bullet\mathcal{Q}(\mathbf{W})+\mathbf{b}_E^T\mathbf{y}_E+\mathbf{b}_I^T\mathbf{y}_I \\ \text{s.t.} \quad & -\mathcal{Q}(\mathbf{W})+\boldsymbol{\Sigma}+\mathcal{A}_E^*(\mathbf{y}_E)+\mathcal{A}_I^*(\mathbf{y}_I)=\mathbf{C} \\ & \mathbf{W}\in\mathbb{S}^n \qquad \boldsymbol{\Sigma}\in\mathcal{S}_+^n \qquad \mathbf{y}_I\succeq 0 \end{aligned} \tag{3.13}$$

where $\mathbf{W}$ is a dual variable such that, at the optimum $\mathbf{W},\mathbf{X}$, $\mathcal{Q}(\mathbf{W})=\mathcal{Q}(\mathbf{X})$; $\boldsymbol{\Sigma}$ is the dual variable for the semidefiniteness; and $\mathbf{y}_E,\mathbf{y}_I$ and $\mathcal{A}_E^*,\mathcal{A}_I^*$ are the dual variables and adjoint operators for the equality and inequality constraints, respectively. Notice that the dual itself contains a quadratic semidefinite term that, at the optimum, is equal to that of the primal problem. While one might be tempted to believe this presents a serious challenge to the solvability of such problems, there are often other structural arguments one can make, particularly about the quadratic operator $\mathcal{Q}(\cdot)$, that can simplify the process. In our case, the operator relates to the clutter, and decouples things in a way that leads to a "nice" conclusion.

More importantly, we should ask when (3.12) and (3.13) are convex, since this (at the very least) permits the use of general convex optimization toolboxes like CVX [86, 87]. It turns out that (3.12) is convex when $\mathcal{Q}$ is completely positive (or *semidefinite*, as is often said in the optimization literature). It can be shown that, given an appropriate partitioning of $\mathbf{X}$, the inner product $\mathbf{X}\bullet\mathcal{Q}(\mathbf{X})$ is equivalent to the vector-quadratic form $\operatorname{vec}(\mathbf{X})^H\mathbf{Q}\operatorname{vec}(\mathbf{X})$, where $\mathbf{Q}=\Psi(\langle\mathcal{Q}\rangle)\in\mathbb{S}_+^{n^2}$ is the Choi matrix of $\mathcal{Q}$. We will commonly use the properties of $\mathbf{Q}$ to stand in for those of $\mathcal{Q}$, since there is a direct equivalence (isometric isomorphism) between the two [88].

Toh [85] observed that the computational complexity of such a problem can depend primarily on the number of linear constraints and the effective rank of the operator $\mathcal{Q}$, which is the same as the rank of $\mathbf{Q}$, say $\operatorname{rank}(\mathbf{Q})=r_Q\leq n^2$. More specifically, if reformulated

into a semidefinite-quadratic-linear program (SQLP), the computational complexity for an interior-point solver is no more than $\mathcal{O}((r_Q + m_E + m_I)^3)$. If $r_Q \ll n^2$, then this can represent a potentially significant speedup, but one must calculate the appropriate rank-revealing decomposition of $\mathbf{Q}$, which can be expensive if $\mathbf{Q}$ is dense. This is fundamentally the approach of interior-point solvers like SDPT3 [89, 90] and SeDuMi [91], though they can exploit the sparsity of $\mathcal{Q}, \mathcal{A}_E, \mathcal{A}_I$ if it exists. Let $m_T = m_E + m_I$ indicate the total number of affine constraints. In [85], Toh also proposed an interior point solver with complexity $\mathcal{O}(m_T^3 n)+\mathcal{O}(m_T^2 n^2)+\mathcal{O}(m_T^3)$, though there is currently no publicly-available implementation of it. Currently, the most promising solver for these problems is QSDPNAL [56, 92], a solver developed by X. D. Li in 2015 based on fast and compact solutions to the dual (3.13) . Recent works [93, 94, 95, 96] have shown that methods like QSDPNAL have excellent convergence properties, even if multiple solutions exist. This means that though they may be expensive, even a poor initialization can result in a quick solution. Unfortunately, QSDPNAL is also not currently publicly available for general use. We will see, however, that even the interior point solvers SDPT3 and SeDuMi can demonstrate the power of appropriately considering QSDPs.

Handling Complex Matrix Variables

As mentioned above, the primal (3.12) and dual (3.13) QSDP problems are technically defined for real variables only. Since radar signal processing, due to the consideration of electrical fields, deals mostly in complex variables, we require a notional mechanism to allow "conversion" between Hermitian and symmetric matrices.

Let $\mathbf{X} \in \mathbb{H}_+^n$ be the complex variable of interest. To make this a real QSDP, we need to introduce a new variable $\widehat{\mathbf{X}} \in \mathbb{S}_+^{2n}$, defined as

$$\widehat{\mathbf{X}} = \begin{bmatrix} \mathrm{Re}(\mathbf{X}) & -\mathrm{Im}(\mathbf{X}) \\ \mathrm{Im}(\mathbf{X}) & \mathrm{Re}(\mathbf{X}) \end{bmatrix}, \tag{3.14}$$

where we can obviously decompose $\mathbf{X} = \mathrm{Re}(\mathbf{X}) + \jmath\,\mathrm{Im}(\mathbf{X})$. Let us define the "real"-ization matrix $\mathbf{D}_{\mathrm{R},n} \in \mathbb{C}^{2n\times n}$ as $\mathbf{D}_{\mathrm{R},n} = \frac{1}{\sqrt{2}}[\mathbf{I}_n\ \jmath\mathbf{I}_n]$, such that $\mathbf{D}_{\mathrm{R},n}^H\mathbf{D}_{\mathrm{R},n} = \mathbf{I}_n$. Then, we can define a matrix mapping between the complex and real representations as $\mathbf{X} = \mathbf{D}_{\mathrm{R},n}\widehat{\mathbf{X}}\mathbf{D}_{\mathrm{R},n}^H$. Finally, we require "real" versions of complex linear mappings associated with a complex QSDP: a quadratic operator $\mathcal{Q}_C$ and equality and inequality operators $\mathcal{A}_{E,C}, \mathcal{A}_{I,C}$. While the latter two can be easily defined by substitution of the first matrix mapping, the operator in the quadratic term requires slightly more care to account for the inner product. Using the properties of trace, we can identify the real quadratic operator $\mathcal{Q}$ as

$$\mathcal{Q}(\widetilde{\mathbf{X}}) = \mathbf{D}_{\mathrm{R},n}^H \left(\mathcal{Q}_C(\mathbf{D}_{\mathrm{R},n}\breve{\mathbf{X}}\mathbf{D}_{\mathrm{R},n}^H)\right)\mathbf{D}_{\mathrm{R},n}. \tag{3.15}$$

In effect, this process means that any complex QSDP has double the variables, as expected.

3.1.3 Relaxing BQPs to QSDPs

With the preliminaries out of the way, we now turn to employing the theory of QSDPs to find a more tractable means of solving the joint BQP (2.30), or, more correctly, the joint quartic problem (2.35). The goal here is to use semidefinite relaxation (first mentioned in Chapter 2.3.2) to relax the quartic problem (2.35) into a QSDP. At that point, we can analytically characterize the set of solutions (and the potential tightness of the relaxation), as well as solve the problem numerically using the interior point methods [89, 90, 91] discussed above.

Our first task is to massage (2.35) into a form that permits a relaxation into the standard form of (3.12), but we will also identify a form that is more amenable to direct analysis. First, we can recast the complex Capon constraint $\mathbf{w}^H\mathbf{T}\mathbf{s} = \kappa$ as two real constraints that depend on the joint parameter vector $\mathbf{b}$:

$$\begin{aligned}\mathrm{Re}\{\mathbf{w}^H\mathbf{T}\mathbf{s}\} &= \mathbf{b}^H\widetilde{\mathbf{T}}_R\mathbf{b} = \kappa_R \\ \mathrm{Im}\{\mathbf{w}^H\mathbf{T}\mathbf{s}\} &= \mathbf{b}^H\widetilde{\mathbf{T}}_I\mathbf{b} = \kappa_I,\end{aligned}$$

where $\widetilde{\mathbf{T}}_R, \widetilde{\mathbf{T}}_I \in \mathbb{H}^J$ are the real and imaginary parts of $\widetilde{\mathbf{T}}$ uniquely determined as $\widetilde{\mathbf{T}}_R = \frac{1}{2}(\widetilde{\mathbf{T}} + \widetilde{\mathbf{T}}^H)$, $\widetilde{\mathbf{T}}_I = \frac{-\jmath}{2}(\widetilde{\mathbf{T}} - \widetilde{\mathbf{T}}^H)$. Observe that both constraints are required in order to correctly enforce the total Capon constraint.

Next, consider the objective function $f(\mathbf{b}) = \mathbf{b}^H\mathbf{R}_{\mathbf{u}}(\mathbf{b})\mathbf{b} = \mathbf{R}_{\mathbf{u}}(\mathbf{b}) \bullet \mathbf{b}\mathbf{b}^H$, where we have used the properties of the trace and the matrix inner product to obtain the second form. We could use the basic expansion of the undesired covariance $\mathbf{R}_{\mathbf{u}}(\mathbf{b}) = \widetilde{\mathbf{R}}_{\mathbf{ni}} + \widetilde{\mathbf{R}}_{\mathbf{c}}(\mathbf{b})$, but recall that there is a completely-positive linear operator $\widetilde{\mathcal{C}}$ such that $\widetilde{\mathbf{R}}_{\mathbf{c}}(\mathbf{b}) = \widetilde{\mathcal{C}}(\mathbf{b}\mathbf{b}^H) = \sum_{q=1}^{Q} \bar{R}_q^{\gamma}\widetilde{\mathbf{\Gamma}}_q\mathbf{b}\mathbf{b}^H\widetilde{\mathbf{\Gamma}}_q^H$. Thus, reformulating the objective to be amenable to semidefinite relaxation can be achieved in two ways. First, we can naively use the inner product form, which gives us

$$f(\mathbf{b}) = \widetilde{\mathbf{R}}_{\mathbf{ni}} \bullet \mathbf{b}\mathbf{b}^H + \mathbf{b}\mathbf{b}^H \bullet \widetilde{\mathcal{C}}(\mathbf{b}\mathbf{b}^H). \tag{3.16}$$

Second, we can use the secondary definition of $\widetilde{\mathcal{C}}(\mathbf{b}\mathbf{b}^H)$ to examine the second term of the inner product, that is, $\mathbf{b}\mathbf{b}^H \bullet \widetilde{\mathcal{C}}(\mathbf{b}\mathbf{b}^H)$. If we expand the inner product, we obtain

$$\mathbf{b}\mathbf{b}^H \bullet \widetilde{\mathcal{C}}(\mathbf{b}\mathbf{b}^H) = \sum_{q=1}^{Q} \overline{R}_q^{\gamma}\mathbf{b}^H\widetilde{\mathbf{\Gamma}}_q\mathbf{b}\mathbf{b}^H\widetilde{\mathbf{\Gamma}}_q^H\mathbf{b}.$$

Recognizing that $\mathbf{b}^H\widetilde{\mathbf{\Gamma}}_q\mathbf{b} = \mathrm{vec}(\mathbf{b}\mathbf{b}^H)^H\,\mathrm{vec}(\widetilde{\mathbf{\Gamma}}_q)$, we then have

$$\begin{aligned}\mathbf{b}\mathbf{b}^H \bullet \widetilde{\mathcal{C}}(\mathbf{b}\mathbf{b}^H) &= \mathrm{vec}(\mathbf{b}\mathbf{b}^H)^H \left(\sum_{q=1}^{Q} \bar{R}_q^{\gamma}\,\mathrm{vec}(\widetilde{\mathbf{\Gamma}}_q)\,\mathrm{vec}(\widetilde{\mathbf{\Gamma}}_q)^H\right)\mathrm{vec}(\mathbf{b}\mathbf{b}^H) \\ &= \mathrm{vec}(\mathbf{b}\mathbf{b}^H)^H\widetilde{\mathbf{C}}_V\,\mathrm{vec}(\mathbf{b}\mathbf{b}^H). \end{aligned} \tag{3.17}$$

The second vectorized form of the objective is therefore given by

$$f(\mathbf{b}) = \widetilde{\mathbf{R}}_{\mathbf{ni}} \bullet \mathbf{b}\mathbf{b}^H + \mathrm{vec}(\mathbf{b}\mathbf{b}^H)^H\widetilde{\mathbf{C}}_V\,\mathrm{vec}(\mathbf{b}\mathbf{b}^H), \tag{3.18}$$

where again $\widetilde{\mathbf{C}}_V$ is the Choi matrix of $\widetilde{\mathcal{C}}$. Finally, we turn to the signal constraint $\mathbf{s} \in \Omega_{\mathbf{s}}$, which is transformed to $\mathbf{\Psi}_S\mathbf{b} \in \Omega_{\mathbf{s}}$.

Hence, we recast (2.35) into two equivalent forms: the inner product-operator form

$$\begin{aligned} \min_{\mathbf{b}} \quad & \widetilde{\mathbf{R}}_{\mathbf{ni}} \bullet \mathbf{b}\mathbf{b}^H + \widetilde{\mathcal{C}}\left(\mathbf{b}\mathbf{b}^H\right) \bullet \mathbf{b}\mathbf{b}^H \\ \text{s.t.} \quad & \widetilde{\mathbf{T}}_R \bullet \mathbf{b}\mathbf{b}^H = \kappa_R \quad \widetilde{\mathbf{T}}_I \bullet \mathbf{b}\mathbf{b}^H = \kappa_I \\ & \mathbf{\Psi}_S\mathbf{b} \in \Omega_{\mathbf{s}}, \end{aligned} \tag{3.19}$$

and the vectorized form

$$\begin{aligned} \min_{\mathbf{b}} \quad & \widetilde{\mathbf{R}}_{\mathbf{ni}} \bullet \mathbf{b}\mathbf{b}^H + \mathrm{vec}(\mathbf{b}\mathbf{b}^H)^H\widetilde{\mathbf{C}}_V\,\mathrm{vec}(\mathbf{b}\mathbf{b}^H) \\ \text{s.t.} \quad & \widetilde{\mathbf{T}}_R \bullet \mathbf{b}\mathbf{b}^H = \kappa_R \quad \widetilde{\mathbf{T}}_I \bullet \mathbf{b}\mathbf{b}^H = \kappa_I \\ & \mathbf{\Psi}_S\mathbf{b} \in \Omega_{\mathbf{s}}. \end{aligned} \tag{3.20}$$

Following the same path to semidefinite relaxation we followed in Chapter 2, let us define the matrix variable $\mathbf{B} = \mathbf{b}\mathbf{b}^H \in \mathbb{H}_+^J$. It is clear that, as is, $\mathbf{B}$ is rank one, which imposes an obviously non-convex constraint. Instead, we permit $\mathbf{B}$ to be of any rank, which results in a

lifting of the problems (3.19) and (3.20) to the equivalent QSDPs in both operator form

$$\begin{aligned} \min_{\mathbf{B}} \quad & \widetilde{\mathbf{R}}_{\mathbf{ni}} \bullet \mathbf{B} + \left(\widetilde{\mathcal{C}}(\mathbf{B})\right) \bullet \mathbf{B} \\ \text{s.t.} \quad & \widetilde{\mathbf{T}}_R \bullet \mathbf{B} = \kappa_R \quad \widetilde{\mathbf{T}}_I \bullet \mathbf{B} = \kappa_I \\ & \mathbf{\Psi}_S^T \mathbf{B} \mathbf{\Psi}_S \in \mathcal{T}(\Omega_{\mathbf{s}}) \\ & \mathbf{B} \in \mathbb{H}_J^+, \end{aligned} \tag{3.21}$$

and vectorized form

$$\begin{aligned} \min_{\mathbf{B}} \quad & \widetilde{\mathbf{R}}_{\mathbf{ni}} \bullet \mathbf{B} + \operatorname{vec}(\mathbf{B})^H \widetilde{\mathbf{C}}_V \operatorname{vec}(\mathbf{B}) \\ \text{s.t.} \quad & \widetilde{\mathbf{T}}_R \bullet \mathbf{B} = \kappa_R \quad \widetilde{\mathbf{T}}_I \bullet \mathbf{B} = \kappa_I \\ & \mathbf{\Psi}_S^T \mathbf{B} \mathbf{\Psi}_S \in \mathcal{T}(\Omega_{\mathbf{s}}) \\ & \mathbf{B} \in \mathbb{H}_+^J \end{aligned} \tag{3.22}$$

where $\mathcal{T}(\cdot)$ is an appropriate transformation of the original constraint space to an equivalent relaxed constraint space. Notice we did not, at any point, progress through a matrix bilinear form in finding these semidefinite relaxations, as seen in [49], etc. This is because the constraint set in this problem prevents such a representation – namely, the Capon constraint is bilinear in the vector arguments already.

The problems in (3.21) and (3.22) can be shown to be convex if and only if the set $\mathcal{T}(\Omega_{\mathbf{s}})$ is convex. In the case of the power-constrained problem we will consider for the rest of this chapter, $\Omega_{\mathbf{s}} = \left\{\mathbf{s} : \|\mathbf{s}\|^2 \leq P_o\right\}$ and $\mathcal{T}(\Omega_{\mathbf{s}}) = \left\{\mathbf{B} : \operatorname{tr}(\mathbf{\Psi}_S \mathbf{B} \mathbf{\Psi}_S^T) \leq P_o\right\}$ for a given upper power bound P_o, both of which are convex in their respective variables. We will show later that while $\Omega_{\mathbf{s}}$ may not be convex, an appropriate parameterization will make $\mathcal{T}(\Omega_{\mathbf{s}})$ convex for the practical waveform constraints we consider in Chapter 4.

Choosing a Relaxation & Notation

While both (3.21) and (3.22) are clearly QSDPs, they are exactly equivalent only if there exists a rank-one optimal solution to both. In fact, (3.22) induces a different form of the clutter operator $\widetilde{\mathcal{C}}$ that we call $\widetilde{\mathcal{C}}_{\mathrm{V}}$, given by

$$\widetilde{\mathcal{C}}_{\mathrm{V}}(\mathbf{B}) = \frac{1}{2}\begin{bmatrix} \mathbf{0}_{NML\times NML} & \sum_{i=1}^{Q} \overline{R}_i^{\gamma} \operatorname{tr}(\widetilde{\mathbf{\Gamma}}_i^H \mathbf{B})\widetilde{\mathbf{\Gamma}}_i \\ \sum_{i=1}^{Q} \overline{R}_i^{\gamma} \operatorname{tr}(\widetilde{\mathbf{\Gamma}}_i \mathbf{B})\mathbf{\Gamma}_i^H & \mathbf{0}_{N\times N} \end{bmatrix}. \tag{3.23}$$

This operator is clearly Hermitian-preserving, but due to the zero diagonal blocks, it cannot be completely positive. However, this still produces a convex function in concert with the inner product, because for any $\mathbf{B}$, it can be shown that $\widetilde{\mathcal{C}}_{\mathrm{V}}(\mathbf{B}) \bullet \mathbf{B} = \operatorname{vec}(\mathbf{B})^H \widetilde{\mathbf{C}}_V \operatorname{vec}(\mathbf{B}) \geq 0$. Hence, it is regarded as a positive-semidefinite operator.

It is this relaxation that we will consider for the rest of this dissertation, because, as we will show, we can easily obtain the Karush-Kuhn-Tucker conditions (which are necessary and sufficient for optimality), guarantees on solution rank, and intuitive interpretations of this relaxation's behavior using this form. We will discuss potential directions for the other form later in Chapter 6.

To facilitate more compact expressions for the vectorized relaxed problem, we introduce the following notation: First, define the following partition of $\mathbf{B}$: $\mathbf{B}_1 = \mathbf{\Psi}_W \mathbf{B} \mathbf{\Psi}_W^T \in \mathbb{H}^{NML}, \mathbf{B}_2 = \mathbf{\Psi}_W \mathbf{B} \mathbf{\Psi}_S^T \in \mathbb{C}^{NML\times N}, \mathbf{B}_3 = \mathbf{\Psi}_S \mathbf{B} \mathbf{\Psi}_S^T \in \mathbb{H}^N$. In matrix form, this is

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_1 & \mathbf{B}_2 \\ \mathbf{B}_2^H & \mathbf{B}_3 \end{bmatrix}.$$

We can also define the vectorizations of each of these matrices, which require the definition of the following composite projection matrices

$$\mathbf{P}_{SS} = \mathbf{\Psi}_S \otimes \mathbf{\Psi}_S \qquad \mathbf{P}_{WW} = \mathbf{\Psi}_W \otimes \mathbf{\Psi}_W$$
$$\mathbf{P}_{SW} = \mathbf{\Psi}_S \otimes \mathbf{\Psi}_W \qquad \mathbf{P}_{WS} = \mathbf{\Psi}_W \otimes \mathbf{\Psi}_S.$$

Note that $\mathbf{P}_{SW} = \mathbf{K}_{NML,N}\mathbf{P}_{WS}\mathbf{K}_{NML,N}^T$, and that $\mathbf{P}_{xy}\mathbf{P}_{vw}^T$ is the appropriate identity matrix when the pair xy is the same as the pair vw and zero otherwise. Then, if $\boldsymbol{\beta} = \text{vec}(\mathbf{B})$, we have $\boldsymbol{\beta}_1 = \mathbf{P}_{WW}^T\boldsymbol{\beta}, \boldsymbol{\beta}_2 = \mathbf{P}_{WS}^T\boldsymbol{\beta}, \boldsymbol{\beta}_3 = \mathbf{P}_{SS}^T\boldsymbol{\beta}$. Additionally, we can define $\boldsymbol{\beta}_{2,H} = \mathbf{P}_{SW}^T\boldsymbol{\beta}$, but then $\boldsymbol{\beta}_{2,H} = \text{vec}(\mathbf{B}_2^H) = \mathbf{K}_{NML,N}\boldsymbol{\beta}_2^*$.

With these in mind, we can then provide a "cleaner" form of the vectorized relaxed problem (3.22) as applied to the simple power constraint on the signal:

$$\begin{aligned} \min_{\mathbf{B}\in\mathbb{H}_J^+} \quad & \text{vec}^H(\widetilde{\mathbf{R}}_{\mathbf{ni}})\boldsymbol{\beta} + \boldsymbol{\beta}^H\widetilde{\mathbf{C}}_V\boldsymbol{\beta} \\ \text{s.t.} \quad & \text{vec}^H(\widetilde{\mathbf{T}}_R)\boldsymbol{\beta} = \kappa_R \\ & \text{vec}^H(\widetilde{\mathbf{T}}_I)\boldsymbol{\beta} = \kappa_I \\ & \text{vec}^H(\mathbf{\Psi}_S^T\mathbf{\Psi}_S)\boldsymbol{\beta} \leq P_o. \end{aligned} \tag{3.24}$$

Synthesis of Rank-One Solutions

If we can obtain a solution to either relaxed problem (3.21), (3.22), a natural question is how to apply this solution to the original problems (2.35) and (2.30). Obviously, when $\text{rank}(\mathbf{B}) = 1$, the relaxation is tight and we can directly recover the pair $(\mathbf{w}, \mathbf{s})$ from $\mathbf{B}$. If the rank is greater than one, we can use a rank-one approximation that still satisfies the constraints. We will see in Chapter 3.5 below that for the power-constrained case, we can obtain solutions that are nearly rank-one in a numerical sense. In any rank case, one simple heuristic approximation of a rank-one solution is a weighted sum of the eigenvectors of $\mathbf{B}$.

We consider primarily the approximate sum of basis-length K, given by $\mathbf{b}_K^a = \sum_{i=1}^{K} \sqrt{\eta_i}\mathbf{u}_i$ where η_i is the ith non-zero eigenvalue of $\mathbf{B}$ in decreasing order, $\mathbf{u}_i$ is its paired eigenvector, and $K \in \{1, 2, \ldots, \text{rank}(\mathbf{B})\}$. From this, we can recover the approximate pair $(\mathbf{w}_K^a, \mathbf{s}_K^a) = (\mathbf{\Psi}_W \mathbf{b}_K^a, \mathbf{\Psi}_S \mathbf{b}_K^a)$. If a solution is rank-one, or nearly so, then $(\mathbf{w}_1^a, \mathbf{s}_1^a)$ is usually feasible and nearly optimal. If it is not feasible, then it can be made feasible – for the power-constrained case, this amounts to the rescaling process used for EigenAM in Chapter 2.3.

3.2 Constraint Qualifications and KKT Conditions

With the general form of our relaxation now developed, we now turn to examining how best to solve for the relaxed matrix $\mathbf{B}$ and its relation to an optimal beamformer-signal pair $(\mathbf{w}, \mathbf{s})$. While numerical methods will be the primary driver of obtaining such solutions, we can use some analytic techniques to sketch, at the very least, the notional contours of the relaxed solution and its relation to the channel state parameters and signal-independent noise-and-interference covariance. In this section, we begin the task of analyzing the optimal solutions of (3.21) by discussing the conditions for optimality. We first consider some constraint qualifications – namely, Slater's condition for both the primal and dual problems, and the potential degeneracy of the dual solution, inspired by [97] – that guarantee certain optimality properties that are required for the Karush-Kuhn-Tucker conditions to be necessary and sufficient. In this process, we demonstrate that there is no duality gap (and thus have a quite strong convergence guarantee) under a very reasonable, physically-realizable restriction. Finally, we summarize the necessary and sufficient conditions for optimality and provide explicit linkage of these conditions to the power-constrained STAP problem.

Before we begin, we recall that we have vectorized our semidefinite variable $\boldsymbol{\beta} = \text{vec}(\mathbf{B})$. So long as we use the appropriate rules of differentiation (see [98]) and remember that there is an isomorphism between Hermitian matrices of size n and the real vector space of size n^2

(which manifests as a matrix operation on $\boldsymbol{\beta}$), this is an appropriate representation for our variable of interest. Additionally, all of our usual inner products, etc. are maintained, so there is no mathematical conflict with this choice.

3.2.1 Constraint Qualifications

We begin with an analysis of constraint qualifications. To perform this analysis, we need to first identify the equivalent of the QSDP dual problem (3.13) for the vectorized QSDP (3.22), which amounts to defining the dual variables and identifying the operators involved.

Clearly, we can let $\boldsymbol{\Sigma} \in \mathbb{H}_+^J$ be the dual variable for the semidefiniteness constraint. Next, let $\mu_R, \mu_I \in \mathbb{R}$ be the dual variables for the real and imaginary part of the Capon constraints, respectively. Since we generally assume in this work that the set $\mathcal{T}(\Omega_{\mathbf{s}})$ represents an affine constraint (usually a stacked trace), let $\boldsymbol{\lambda}_T \in \mathbb{R}^{|\mathcal{T}(\Omega_{\mathbf{s}})|}$ be the dual variable for the signal-only constraints, where $|\mathcal{T}(\Omega_{\mathbf{s}})|$ indicates the cardinality of the set. Finally, let $\boldsymbol{\Delta} \in \mathbb{H}^J$ be the equivalent of the variable $\mathbf{W}$ in (3.13). That is, $\boldsymbol{\Delta}$ is a dual variable such that $\widetilde{\mathcal{C}}_{\mathrm{V}}(\mathbf{B}) = \widetilde{\mathcal{C}}_{\mathrm{V}}(\boldsymbol{\Delta})$ or, if $\boldsymbol{\delta} = \mathrm{vec}(\boldsymbol{\Delta})$, $\widetilde{\mathbf{C}}_V \boldsymbol{\beta} = \widetilde{\mathbf{C}}_V \boldsymbol{\delta}$.

Because $\boldsymbol{\Sigma}$ is of the same dimensions as $\mathbf{B}$, let us define an identical partitioning for $\boldsymbol{\Sigma}$ which will be useful:

$$\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_1 & \boldsymbol{\Sigma}_2 \\ \boldsymbol{\Sigma}_2^H & \boldsymbol{\Sigma}_3 \end{bmatrix},$$

where the dimensions of each submatrix match those of the submatrices of $\mathbf{B}$ with an identical subscript. By solving the equality constraint in (3.13) for our case (and adding a factor of 2 to the quadratic component), we can see the optimal $\boldsymbol{\Sigma}$ is a function of the other dual

variables as follows:

$$\boldsymbol{\Sigma} = \widetilde{\mathbf{R}}_{\mathbf{ni}} + 2\widetilde{\mathcal{C}}_{\mathrm{V}}(\boldsymbol{\Delta}) - (\mu_R \widetilde{\mathbf{T}}_R + \mu_I \widetilde{\mathbf{T}}_I) + \boldsymbol{\Psi}_S^T \boldsymbol{\Sigma}_3(\boldsymbol{\lambda}_T) \boldsymbol{\Psi}_S. \tag{3.25}$$

In fact, we can simplify this expression slightly by taking advantage of our reduction definitions above. First, the two Capon constraint matrices can be reduced to

$$\mu_R \widetilde{\mathbf{T}}_R + \mu_I \widetilde{\mathbf{T}}_I = \tilde{\mu}_C^* \widetilde{\mathbf{T}} + \tilde{\mu}_C \widetilde{\mathbf{T}}^H = \begin{bmatrix} \mathbf{0}_{NML \times NML} & \tilde{\mu}_C^* \mathbf{T} \\ \tilde{\mu}_C \mathbf{T} & \mathbf{0}_{N \times N} \end{bmatrix} \tag{3.26}$$

where we define the new Capon constraint dual variable $\tilde{\mu}_C = \frac{\mu_R + \jmath\, \mu_I}{2} \in \mathbb{C}$. Next, assuming $\boldsymbol{\Delta}$ is partitioned identically to $\mathbf{B}, \boldsymbol{\Sigma}$, we can show

$$2\widetilde{\mathcal{C}}_{\mathrm{V}}(\boldsymbol{\Delta}) = \begin{bmatrix} \mathbf{0}_{NML \times NML} & \sum_{q=1}^{Q} \overline{R}_q^{\gamma} \operatorname{tr}(\boldsymbol{\Gamma}_q^H \boldsymbol{\Delta}_2) \boldsymbol{\Gamma}_q \\ \sum_{q=1}^{Q} \overline{R}_q^{\gamma} \operatorname{tr}(\boldsymbol{\Delta}_2^H \boldsymbol{\Gamma}_q) \boldsymbol{\Gamma}_q^H & \mathbf{0}_{N \times N} \end{bmatrix} \tag{3.27}$$

Hence, a final expression for the general optimal semidefinite dual variable is

$$\boldsymbol{\Sigma} = \begin{bmatrix} \mathbf{R}_{\mathbf{ni}} & \sum_{q=1}^{Q} \overline{R}_q^{\gamma} \operatorname{tr}(\boldsymbol{\Gamma}_q^H \boldsymbol{\Delta}_2) \boldsymbol{\Gamma}_q - \tilde{\mu}_C^* \mathbf{T} \\ \left(\sum_{q=1}^{Q} \overline{R}_q^{\gamma} \operatorname{tr}(\boldsymbol{\Delta}_2^H \boldsymbol{\Gamma}_q) \boldsymbol{\Gamma}_q - \tilde{\mu}_C^* \mathbf{T} \right)^H & \boldsymbol{\Sigma}_3(\boldsymbol{\lambda}_T) \end{bmatrix}. \tag{3.28}$$

If we consider just the power constrained problem (3.24), then $\boldsymbol{\lambda}_T$ collapses to one scalar, $\lambda_P \in \mathbb{R}^+$, and hence the dual semidefiniteness variable is

$$\boldsymbol{\Sigma} = \begin{bmatrix} \mathbf{R}_{\mathbf{ni}} & \sum_{q=1}^{Q} \overline{R}_q^{\gamma} \operatorname{tr}(\boldsymbol{\Gamma}_q^H \boldsymbol{\Delta}_2) \boldsymbol{\Gamma}_q - \tilde{\mu}_C^* \mathbf{T} \\ \left(\sum_{q=1}^{Q} \overline{R}_q^{\gamma} \operatorname{tr}(\boldsymbol{\Delta}_2^H \boldsymbol{\Gamma}_q) \boldsymbol{\Gamma}_q - \tilde{\mu}_C^* \mathbf{T} \right)^H & \lambda_P \mathbf{I}_N \end{bmatrix}. \tag{3.29}$$

Since $\mathbf{\Delta}$ is such that it identically matches the variable of interest $\mathbf{B}$ on the clutter operator, we will subsequently ignore this variable in our analysis, as we can directly solve for the necessary parts of $\mathbf{B}$ later.

Slater's Condition

While there are many types of constraint qualifications for QSDP (see, *e.g.*, [97]), the one that has the most utility is also the one that can be the easiest to compute. Recall that *strong duality* is said to hold for a given optimization problem if the primal problem is convex and it satisfies Slater's condition [76, pp. 240, 265] (*i.e.*, the problem is *strictly feasible*).

We begin by defining this condition. Assume that we have a convex objective function $f_o(\mathbf{x})$. Consider, then, the following general optimization problem:

$$\begin{aligned} \min_{\mathbf{x}\in\Omega} \quad & f_o(\mathbf{x}) \\ \text{s.t.} \quad & g_i(\mathbf{x}) \leq 0 \quad i = 1, \ldots, m \\ & h_j(\mathbf{x}) = 0 \quad j = 1, \ldots, n \end{aligned} \tag{3.30}$$

where Ω is a convex set, g_i, h_j are convex functions, and we define the set $\mathcal{D} = \Omega \cap \operatorname{dom}(f_o) \cap (\cap_{i=1}^{m} \operatorname{dom}(g_i))$ to be the total feasible domain of the problem.

Definition 3.4 (Slater Condition)**.** *The optimization problem above satisfies the* Slater condition *if there exists at least one point in the relative interior (*i.e.*, not on the boundary) of the problem's feasibility set that satisfies all of the equality constraints and strictly satisfies any inequality constraints; that is,* $\exists\ \mathbf{x}_s \in \operatorname{relint}(\Omega)$ *such that* $h_j(\mathbf{x}_s) = 0 \quad j = 1, \ldots, n$ *and* $g_i(\mathbf{x}) < 0 \quad i = 1, \ldots, m$.

For our semidefinite programs,the feasibility set is the positive semidefinite matrices $\mathbb{H}_{+}^{J}$, whose relative interior is the positive-definite matrices $\mathbb{H}_{++}^{J}$.

If the Slater condition is satisfied, then Bonnans and Shapiro [99] offer the following theorems for convex semidefinite programs of any kind, one about the relationship between the primal and dual problems and the other about the existence of Lagrange multipliers (which will become necessary in the next section). For these purposes, let $\nu^\star$ denote the optimal value of a given optimization problem.

Theorem 3.3 (Theorem 5.81, [99]). *Suppose the problem* (3.30) *is convex. If the Slater condition for* (3.30) *holds, then there is no duality gap between the primal and dual problems (*i.e., $\nu_P^\star = \nu_P^\star$*). Moreover, if their common optimal value $\nu^\star$ is finite, then the set of optimal solutions of the dual problem is non-empty and bounded.*

Theorem 3.4 (Theorem 5.83, [99]). *Suppose that an SDP problem is convex and let $\mathbf{X}^\star$ be its optimal solution. If the Slater condition for this problem holds, then the set of Lagrange multipliers for $\mathbf{X}^\star$ is non-empty and bounded, and is the same for any optimal solution of the SDP problem.*

While we could consider a derivation of the Slater condition for the general primal problem (3.22), it is actually quite straightforward:

Theorem 3.5 (General RBQP Slater Condition). *The Slater condition for* (3.22) *is satisfied if there exists a $\mathbf{B} \succ 0$ such that $\widetilde{\mathbf{T}} \bullet \mathbf{B} = \kappa$ and $\mathbf{\Psi}_S^T \mathbf{B} \mathbf{\Psi}_S \in \text{relint}(\mathcal{T}(\Omega_\mathbf{s}))$.*

The relative interior of $\mathcal{T}(\Omega_\mathbf{s})$ is extremely problem dependent, and we will see in Chapter 4 it is not even needed in some cases. However, to start, let us consider the power-constrained case (3.24). In this case, $\text{relint}(\mathcal{T}(\Omega_\mathbf{s})) = \{\mathbf{B} : \text{tr}(\mathbf{\Psi}_S^T \mathbf{B} \mathbf{\Psi}_S) < P_o\}$ We assert that the Slater condition is satisfied for this relaxed BQP given an even simpler reasonable condition, which we describe in the following theorem.

Theorem 3.6 (Power-constrained RBQP Primal Slater Condition). *The relaxed joint trans-receive design problem in* (3.24) *satisfies the Slater condition if $\|\mathbf{T}\|_F^2 > \frac{|\kappa|^2}{P_o}$.*

Proof: See Appendix A. ■

This is indeed a reasonable assumption, because it says that the Capon constraint, *i.e.* the target-matching goal, cannot exceed the overall reflected power from the target.

In the side-looking STAP case, we can make a more precise relation to the dimensionality of the problem, using the fact that the target matrix has a Kronecker structure – that is, $\mathbf{T} = \mathbf{v}_t \otimes \mathbf{I}_N \otimes \mathbf{a}_t$. This means that $\mathbf{T}^H\mathbf{T} = \|\mathbf{v}_t\|^2\|\mathbf{a}_t\|^2\mathbf{I}_N$. If we assume that the spatial response comes from a uniform linear array and that the Doppler response is similarly uniform, then $\|\mathbf{a}_t\|^2 = M$ and $\|\mathbf{v}_t\|^2 = L$. This further implies that $\mathbf{T}^H\mathbf{T} = ML\mathbf{I}_N$ and $\|\mathbf{T}\|_F^2 = NML$. Hence, via the theorem, the Slater condition is satisfied so long as $NML > \frac{|\kappa|^2}{P_o}$.

In either case, since most reasonable STAP problems (or, indeed, any joint trans-receive problem in active sensing) satisfy this condition, we can be assured that almost any numerical method used to solve (3.21) will converge to a solution that is both primally and dually optimal (or nearly so) by Theorem 3.3.

It is here we note that the condition in Theorem 3.6 was essentially foreshadowed by the Slater condition for the projected QCQP problem (2.48). Recall that the inequality in that case required $\left\|\mathbf{P}_{\mathbf{w}}^{\perp}\mathbf{q}\right\|_2^2 \leq P_o - \frac{|\kappa|^2}{\|\mathbf{y}_{\mathbf{w}}\|_2^2}$. This is strictly satisfied by, for example, $\mathbf{q} = \mathbf{0}_{N\times 1}$ if and only if $\frac{|\kappa|^2}{P_o} < \left\|\mathbf{T}^H\mathbf{w}_k\right\|_2^2$, which is upper bounded by $\|\mathbf{T}\|_F^2\|\mathbf{w}_k\|_2^2$. If we assume the norm of the filter iterate $\|\mathbf{w}_k\|_2^2 \leq 1$, then we exactly recover the condition in Theorem 3.6!

3.2.2 Obtaining the KKTs

In the subsequent analysis, let us assume that Theorem 3.6 holds for our specific problem. It is well-known that, for an optimization problem that satisfies the Slater condition, the Karush-Kuhn-Tucker (KKT) conditions are necessary first-order conditions for the existence and optimality of a set of primal and dual variables (see, *e.g.* [76, 99]). Furthermore, if

such a problem is convex, then the KKTs become necessary and sufficient conditions, which implies that solving them directly may provide insight into the solutions of the problem. In this section, we summarize and analyze the KKTs for the problem (3.24); further details of their derivation can be found in [65] and Appendix B.

Let $\mathcal{H}[\cdot]$ indicate the Hessian of a given function and $\mathcal{L}(\cdot)$ be the Lagrangian function for (3.24). Then, using the notation above, the KKT conditions for the relaxed biquadratic problem can be written as

1. $\nabla_{\mathbf{B}}\mathcal{L}(\mathbf{B}^o, \tilde{\mu}_C^o, \lambda_P^o, \mathbf{\Sigma}^o) = \mathbf{0}_{J\times J}$ (Stationarity of Lagrangian)
2. $\lambda_P^o \geq 0,\ \mathbf{\Sigma}^o \succeq 0$ (Positivity of inequality/conic constraints)
3. $\lambda_P^o(\mathrm{tr}(\mathbf{\Psi}_S^T\mathbf{B}^o\mathbf{\Psi}_S) - P_o) = 0,\ \mathbf{B}^o\mathbf{\Sigma}^o = \mathbf{0}_{J\times J}$ (Complementary slackness)
4. $\mathrm{tr}(\mathbf{\Psi}_S^T\mathbf{B}^o\mathbf{\Psi}_S) - P_o \leq 0,\ \mathbf{B} \succeq 0$ (Inequality/conic constraint feasibility)
5. $\mathrm{tr}(\mathbf{B}^o\widetilde{\mathbf{T}}) = \kappa$ (Equality constraint feasibility)
6. $\mathcal{H}[\mathcal{L}(\mathbf{B}^o, \tilde{\mu}_C^o, \lambda_P^o, \mathbf{\Sigma}^o)] \succeq 0$ (Positive semidefiniteness of the Hessian)

where the optimal value of a variable is denoted by a superscript o and $\nabla_{\mathbf{B}}$ indicates a gradient on the manifold defined by the domain of the variable $\mathbf{B}$. That is, for a matrix $\mathbf{B}^o$ to be a regular minimizer of the related optimization problem, it is necessary that it satisfies these conditions. If the problem is convex, these are necessary and sufficient conditions for the optimal minimizer and associated Lagrange multipliers. For future notational simplicity, we will drop the superscript o until absolutely necessary (*i.e.*, a final statement of the optimal solution). According to [98], the gradient in the first KKT condition is $\nabla_{\mathbf{B}} = \frac{\partial}{\partial B^*}$ or, in other words, $\mathcal{D}_{\mathbf{B}^*}\mathcal{L} = \mathrm{vec}^T(\frac{\partial \mathcal{L}}{\partial B^*})$. Thus, this vectorized derivative form can take the place of the gradient, which we take advantage of in Appendix B.

Since we have assumed the Slater condition is satisfied, the above KKTs are assured to have meaning because, by Theorem 3.4, the set of Lagrange multipliers $(\tilde{\mu}_C, \lambda_P, \mathbf{\Sigma})$ is non-empty, bounded, and identical for any primal solution $\mathbf{B}$. Furthermore, Lagrange parameters and dual variables coincide for any convex SDP problem [99, p. 484], which will become apparent soon.

A brief summary of the salient elements of the KKTs follows. First, we note that the Hessian of the Lagrangian $\mathcal{H}[\mathcal{L}(\mathbf{B}, \tilde{\mu}_C, \lambda_P, \mathbf{\Sigma})] = \widetilde{\mathbf{C}}_V \succeq 0$ for all possible KKT points $(\mathbf{B}, \tilde{\mu}_C, \lambda_P, \mathbf{\Sigma})$ by definition, and thus the relevant condition is always satisfied. Next, by rearranging and solving for the gradient condition, we obtain an expression for the semidefiniteness multiplier $\mathbf{\Sigma}$ (which itself is PSD):

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{R_{ni}} & \mathbf{\Sigma}_2 \\ \mathbf{\Sigma}_2^H & \lambda_P \mathbf{I}_N \end{bmatrix}, \tag{3.31}$$

where $\mathbf{\Sigma}_2 = \sum_{q=1}^{Q} \overline{R}_\gamma^q \operatorname{tr}(\mathbf{\Gamma}_q^H \mathbf{B}_2)\mathbf{\Gamma}_q - \tilde{\mu}_C^* \mathbf{T}$ and, as above, we have dropped the explicit dependence on the dual variable $\mathbf{\Delta}$. Observe that this is precisely the same as (3.29) due to the coincidence of primal Lagrange and dual variables. If we vectorize $\mathbf{\Sigma}_2$, we obtain a more obvious connection to the original cost function:

$$\boldsymbol{\sigma}_2 = \mathbf{C}_V \boldsymbol{\beta}_2 - \tilde{\mu}_C^* \boldsymbol{\tau}, \tag{3.32}$$

where $\boldsymbol{\tau} = \operatorname{vec}(\mathbf{T})$. Hence, the magnitude of this dual variable can be regarded as a "remainder" of sorts – that is, the difference between how much the optimal primal solution must align with the clutter (the first term) and the target itself, scaled by the remaining gain from the equality constraint (the second term).

Let $\mathbf{P}_{\mathbf{C}_V} = \mathbf{C}_V \mathbf{C}_V^\dagger$ be the orthogonal projection matrix onto $\mathcal{R}(\mathbf{C}_V)$, and $\mathbf{P}_{\mathbf{C}_V}^\perp = \mathbf{I}_{N^2ML} - \mathbf{P}_{\mathbf{C}_V}$. In order for (3.32) to have a solution for $\boldsymbol{\beta}_2$ (which a non-trivial set of optimal variables

will have), we also require

$$\mathbf{P}_{\mathbf{C}_V}^{\perp}\boldsymbol{\sigma}_2 = -\tilde{\mu}_C^*\mathbf{P}_{\mathbf{C}_V}^{\perp}\boldsymbol{\tau}, \tag{3.33}$$

that is, the portion of $\boldsymbol{\Sigma}_2$ with columns orthogonal to each clutter patch is a scaled version of the portion of the target similarly situated. Furthermore, due to the positive semidefiniteness of $\boldsymbol{\Sigma}$, the clutter-target remainder $\boldsymbol{\Sigma}_2$ must be contained within signal-independent noise-and-interference spectrum, *i.e.*, $\mathcal{R}(\boldsymbol{\Sigma}_2) \subseteq \mathcal{R}(\mathbf{R_{ni}})$.

Next, we summarize some results from the application of the matrix complementarity conditions in Appendix B. Namely, we can find two equivalent representations of the power constraint multiplier λ_P:

$$\lambda_P = \frac{\mathrm{tr}(\mathbf{B}_1\mathbf{R_{ni}})}{P_o} = -\frac{\mathrm{tr}(\mathbf{B}_2^H\boldsymbol{\Sigma}_2)}{P_o}, \tag{3.34}$$

as well as a characterization of the objective function in terms of a dual variable,

$$\tilde{\mu}_C^*\kappa = \boldsymbol{\beta}_2^H\mathbf{C}_V\boldsymbol{\beta}_2 + \mathrm{tr}(\mathbf{B}_1\mathbf{R_{ni}}). \tag{3.35}$$

From (3.34), we have an interpretation of $\frac{1}{\lambda_P}$ (when λ_P is positive) as an SINR of sorts; that is, we can regard it as the ratio of the transmit power to the power received from non-signal dependent sources after filtering. More interestingly, we can see from (3.35) that the optimal rejection of undesired interference and noise is ultimately dependent on the Capon constraint and its dual variable. Hence, we also have $\nu_{RBQP-PC}^o = (\tilde{\mu}_C^o)^*\kappa$! The left hand side of (3.35) also tells us that the optimal phase of $\tilde{\mu}_C$ is that of κ, since the right hand side is always real and non-negative.

As it happens, the equivalence in (3.35) comes about irrespective of particular constraints on the transmitted signal. In fact, many of the above conditions are generic to the general

problem (3.22), which will be useful in the future analyses of more realistically-constrained design problems in Chapter 4.

3.3 Implications of the KKT Conditions

The optimality conditions shown above and in Appendix B are, at first blush, a complicated set of matrix equations to solve. Furthermore, to our knowledge, there are no generic conditions in the literature that would allow us to immediately identify properties of the relaxed solution to (3.24) or its relationship to high-quality solutions of the particular BQP (2.30). Indeed, most authors investigating QSDPs (see [56] and citations within) have eschewed such guarantees or only provide them for quadratic operators very much unlike the clutter tensor $\mathcal{C}$. However, we can derive some problem-specific insight by examining and solving for certain KKT conditions that have physical meaning.

In this section, we will first prove a variety of generic properties for any solution $\mathbf{B}^o$ of (3.21) and then relate these properties to the ability to directly recover high-quality solutions to (2.30). In particular, we will show that if the signal-independent noise-and-interference covariance matrix $\mathbf{R_{ni}}$ is full rank, then the solution is both power-bounded and more likely to have simple solution recovery.

3.3.1 General properties of the relaxed solution

Since, in our problem, the power constraint is an inequality, the first major question to ask is: under what conditions the solution reaches that bound – that is, when does $\mathrm{tr}(\mathbf{B}_3^o) = P_o$? It turns out that this is a pivotal component of characterizing any solution of (3.21). Not only does this allow us to anticipate the role of clutter whitening in the solution, but it also allows us to predict a bound on the solution matrix's rank.

Before we begin, we note that due to the complementary slackness condition (KKT condition 3 above), the power bound is attained if and only if the dual variable $\lambda_P > 0$. Therefore, this condition will be shorthand in our proofs for attaining the power bound.

First, we show that the solution does not reach the power bound P_o iff the slackness matrix $\mathbf{\Sigma}_2 = \mathbf{0}_{NML\times N}$ by proving the following lemma.

Lemma 3.1. *If it exists, a solution to* (3.24) *does not reach the power bound if and only if the filter-signal dual slackness variable* $\mathbf{\Sigma}_2$ *is zero; that is,* $\lambda_P = 0 \iff \mathbf{\Sigma}_2 = \mathbf{0}_{NML\times N}$

Proof: First, we proceed in the forward direction. If $\lambda_P = 0$, then the slackness matrix becomes

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{R_{ni}} & \mathbf{\Sigma}_2 \\ \mathbf{\Sigma}_2^H & \mathbf{0}_{N\times N} \end{bmatrix}.$$

To be part of a feasible solution, this must be positive semidefinite, which is only possible if $\mathbf{\Sigma}_2 = \mathbf{\Sigma}_2(\mathbf{0}_{N\times N})^\dagger(\mathbf{0}_{N\times N}) = \mathbf{0}_{NML\times N}$ (see [100, Theorem $I_{a''}$]). Hence, the forward direction is proved.

In the reverse direction, we prove via contradiction. Assume that $\mathbf{\Sigma}_2 = \mathbf{0}_{NML\times N}$ and $\lambda_P > 0$. The matrix slackness condition (B.8) dictates that $\lambda_P\mathbf{B}_2 = -\mathbf{B}_1\mathbf{\Sigma}_2$. Under our first assumption, this becomes $\lambda_P\mathbf{B}_2 = \mathbf{0}_{NML\times N}$, which simplifies to $\mathbf{B}_2 = \mathbf{0}_{NML\times N}$ under the second assumption. However, any feasible solution must also satisfy the Capon constraint $\mathrm{tr}(\mathbf{B}_2^H\mathbf{T}) = \kappa \neq 0$. Clearly, $\mathbf{B}_2 = \mathbf{0}_{NML\times N}$ violates this constraint, which leads to our contradiction and completes the proof. ■

Observe that if $\mathbf{\Sigma}_2 = \mathbf{0}_{NML\times N}$, then, via (3.32), recovery of an optimal $\mathbf{B}_2$ would merely be a whitening of the target with respect to the clutter matrix $\mathbf{C}_V$ (since then $\mathbf{C}_V\boldsymbol{\beta}_2 = \tilde{\mu}_C^*\boldsymbol{\tau}$). However, as shown above, this is not possible in a power-bounded solution, which means that additional resources in the joint spectrum must have been allocated to some other whitening process.

It turns out that this other whitening process is related to signal-*independent* interference and noise, as we will show below. Indeed, every feasible solution reaches the power bound if and only if the noise-and-interference correlation matrix $\mathbf{R_{ni}}$ is full rank, the proof of which makes use of Lemma 3.1.

Theorem 3.7. *If it exists, a solution to* (3.24) *does reach the power bound if and only if the signal-independent disturbance covariance is positive-definite/full-rank. In other words,* $\lambda_P > 0 \iff \text{rank}(\mathbf{R_{ni}}) = NML$.

Proof of $\text{rank}(\mathbf{R_{ni}}) = NML \Rightarrow \lambda_P > 0$*:* We will prove this by contradiction as well. Assume that $\mathbf{R_{ni}}$ is full rank and $\lambda_P = 0$. Given the second condition, Lemma 3.1 requires that $\mathbf{\Sigma}_2 = \mathbf{0}_{NML \times N}$. If we apply this to the slackness condition (B.7), then $\mathbf{B}_1\mathbf{R_{ni}} = \mathbf{0}_{NML \times NML}$. However, since $\mathbf{R_{ni}}$ is full rank, this implies $\mathbf{B}_1 = \mathbf{0}_{NML \times NML}$. This violates our non-triviality, but we will continue with the proof to show we reach a further contradiction. Since the overall solution matrix must be PSD, $\mathbf{B}_2 = \mathbf{0}_{NML \times N}$ as well. As in the proof of Lemma 3.1, we have reached a contradiction because the Capon constraint is violated, which completes the proof of sufficiency.

Proof of $\lambda_P > 0 \Rightarrow \text{rank}(\mathbf{R_{ni}}) = NML$*:* First, observe that for a non-trivial PSD solution matrix, $\mathcal{R}(\mathbf{B}_2^H) \subseteq \mathcal{R}(\mathbf{B}_3)$ by [101, Theorem 7.7.9(a,b)]. Next, we turn to slackness condition (B.10), which states $\lambda_P\mathbf{B}_3 = -\mathbf{B}_2^H\mathbf{\Sigma}_2$. If $\lambda_P > 0$, then clearly $\mathcal{R}(\mathbf{B}_3) \subseteq \mathcal{R}(\mathbf{B}_2^H)$. By the standard rules of subset inclusion, then, $\mathcal{R}(\mathbf{B}_2^H) = \mathcal{R}(\mathbf{B}_3)$ and $\text{rank}(\mathbf{B}_2^H) = \text{rank}(\mathbf{B}_3)$. This fact will become useful later.

We will now use the results of [102] on another slackness condition (B.9) and use a substitution of (B.10) to get to our destination. Using [102, Theorem 2.2] on (B.9) to solve for $\mathbf{R_{ni}}$, we have the following requirements for $\mathbf{R_{ni}}$ to be at least PSD (which it is):

1. $-\mathbf{B}_3\mathbf{\Sigma}_2^H\mathbf{B}_2 \succeq 0$: If we substitute (B.10) into this, we obtain $\lambda_P\mathbf{B}_3^2 \succeq 0$, which is satisfied because $\lambda_P > 0$ and any square of a Hermitian matrix is PSD.

2. $\mathcal{R}(\mathbf{B}_3\boldsymbol{\Sigma}_2^H) \subseteq \mathcal{R}(\mathbf{B}_2^H)$: This is satisfied because $\mathbf{B}_3\boldsymbol{\Sigma}_2^H = -\frac{1}{\lambda_P}\mathbf{B}_2^H\boldsymbol{\Sigma}_2\boldsymbol{\Sigma}_2^H$ via (B.10), and the range inclusion follows directly.

3. $\text{rank}(-\mathbf{B}_3\boldsymbol{\Sigma}_2^H\mathbf{B}_2) = \text{rank}(\mathbf{B}_3\boldsymbol{\Sigma}_2^H)$: This is not immediately satisfied by the other conditions, but it does imply that $\text{rank}(\mathbf{B}_3) = \text{rank}(\mathbf{B}_3\boldsymbol{\Sigma}_2^H)$.

The more interesting requirement comes from [103], which adds the following: $\mathbf{R_{ni}}$ is positive definite (and thus full rank) if and only if $\text{rank}(-\mathbf{B}_3\boldsymbol{\Sigma}_2^H\mathbf{B}_2) = \text{rank}(\mathbf{B}_2^H)$. We know from the third PSD requirement above that $\text{rank}(-\mathbf{B}_3\boldsymbol{\Sigma}_2^H\mathbf{B}_2) = \text{rank}(\mathbf{B}_3)$, and thus $\mathbf{R_{ni}}$ is full rank if and only if $\text{rank}(\mathbf{B}_3) = \text{rank}(\mathbf{B}_2^H)$. However, as seen from above, this condition is already satisfied if $\lambda_P > 0$, and so our proof is complete. ■

In cases where there is perfect state information (that is, full knowledge of the noise and interference covariance *a priori*), the noise-and-interference covariance is *always* full-rank, so Theorem 3.7 allows us to conclude that most theoretical solutions will achieve the power bound. This would also be true if there is sufficient sample support to provide good estimates of the non-clutter/non-target processes. However, computation or insufficient training data might result in a rank-deficient estimate of $\mathbf{R_{ni}}$, which would mean the power bound is not attained. The effects of such a scenario were initially examined in [65] and will be further explored in Chapter 3.3.2. Furthermore, this means that if a rank-one solution is obtained, then the relaxation will have provided us a solution that is directly comparable to power-equality-constrained problems, like those in [44].

Since the proof of Theorem 3.7 provides us with the evidence to show that power-bounded solutions occur only in non-singular noise-and-interference environments, we can also demonstrate an additional property of any power-bounded solution: namely, the rank of the overall solution matrix. Recall that for $\mathbf{B}$ to be PSD, $\mathcal{R}(\mathbf{B}_2) \subseteq \mathcal{R}(\mathbf{B}_1)$. However, the slackness condition (B.7) provides that $\mathbf{B}_1\mathbf{R_{ni}} = -\mathbf{B}_2\boldsymbol{\Sigma}_2^H$. Since $\mathbf{R_{ni}}$ is full rank, (B.7) becomes $\mathbf{B}_1 = -\mathbf{B}_2\boldsymbol{\Sigma}_2^H\mathbf{R_{ni}}^{-1}$, which implies $\mathcal{R}(\mathbf{B}_1) \subseteq \mathcal{R}(\mathbf{B}_2)$. Thus, $\mathcal{R}(\mathbf{B}_1) = \mathcal{R}(\mathbf{B}_2)$, and

$\text{rank}(\mathbf{B}_1) = \text{rank}(\mathbf{B}_2) = \text{rank}(\mathbf{B}_3) \leq N$, where the last equality is implied by Theorem 3.7's proof and the inequality is obvious.

In fact, we can strengthen this inequality with the following lemma, which relates the pseudo-SINR λ_P to the optimal power-bounded rank.

Lemma 3.2. *In a power-bounded solution,* $\text{rank}(\mathbf{B}) \leq N - \text{rank}(\lambda_P \mathbf{I}_N - \boldsymbol{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{-1} \boldsymbol{\Sigma}_2)$.

Proof: Our above note on Theorem 3.7 is our starting point. Since the matrix product $\mathbf{B}\boldsymbol{\Sigma}$ is zero, so is its rank. This further implies $\text{rank}(\mathbf{B}) \leq J - \text{rank}(\boldsymbol{\Sigma})$. In general, since $\boldsymbol{\Sigma}$ is PSD, $\text{rank}(\boldsymbol{\Sigma}) = \text{rank}(\mathbf{R}_{\mathbf{ni}}) + \text{rank}(\lambda_P \mathbf{I}_N - \boldsymbol{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{\dagger} \boldsymbol{\Sigma}_2)$. Therefore, generally, $\text{rank}(\mathbf{B}) \leq (NML - \text{rank}(\mathbf{R}_{\mathbf{ni}})) + (N - \text{rank}(\lambda_P \mathbf{I}_N - \boldsymbol{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{\dagger} \boldsymbol{\Sigma}_2))$. In a power bounded solution, $\text{rank}(\mathbf{R}_{\mathbf{ni}}) = NML$ and the inverse exists, thus $\text{rank}(\mathbf{B}) \leq N - \text{rank}(\lambda_P \mathbf{I}_N - \boldsymbol{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{-1} \boldsymbol{\Sigma}_2)$. ■

Indeed, this is a specific version of the following theorem, which is generic for any of the QSDPs we will consider in this dissertation:

Theorem 3.8. *The rank of any optimal solution* $\mathbf{B}$ *of* (3.22) *is bounded by the inequality*

$$\text{rank}(\mathbf{B}) \leq (NML - \text{rank}(\mathbf{R}_{\mathbf{ni}})) + (N - \text{rank}(\boldsymbol{\Sigma}_3(\boldsymbol{\lambda}_T) - \boldsymbol{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{\dagger} \boldsymbol{\Sigma}_2)).$$

As we will show in Chapter 3.3.3, the condition in Lemma 3.2 can be tightened even further, though it requires a small trick of linear algebra. Hence, the only way to guarantee rank-one solutions (and thus obtain an immediate solution to (2.30)) is for $\text{rank}(\lambda_P \mathbf{I}_N - \boldsymbol{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{-1} \boldsymbol{\Sigma}_2) = N - 1$. That said, as we will show below in Chapter 3.5, if there is a single dominant eigenvalue, then the effective rank will be one, and a high quality power-bounded solution may be obtained. This proof also implies that if the solution is not power-bounded, then recovery of a high-quality approximation may be more difficult, as the relaxed solution potentially spans a much larger space (up to $NML - \text{rank}(\mathbf{R}_{\mathbf{ni}})$ more directions!).

Finally, one might ask if the optimal solution produces non-physical SINR guarantees, since the objective function is the undesired signal variance (and thus the denominator of SINR). For power-bounded solutions, the following lemma confirms that this is impossible; for context, recall from (3.35) that the optimal value of (3.24) $\nu^o_{RBQP-PC} = (\tilde{\mu}^o_C)^* \kappa$.

Lemma 3.3. *Any optimal power-bounded solution to* (3.24) *has a non-zero optimal value; that is,* $\lambda_P > 0 \implies (\tilde{\mu}_C \neq 0 \iff \nu^o_{RBQP-PC} > 0)$.

Proof: We prove by contradiction. Assume $\tilde{\mu}_C = 0$. When applied to (3.32), we have $\boldsymbol{\sigma}_2 = \mathbf{C}_V \boldsymbol{\beta}_2$. Premultiplying with $\boldsymbol{\beta}_2^H$, we have $\boldsymbol{\beta}_2^H \boldsymbol{\sigma}_2 = \boldsymbol{\beta}_2^H \mathbf{C}_V \boldsymbol{\beta}_2 \geq 0$ as the quadratic form of a PSD matrix. But we have already established that for $\lambda_P > 0$, $\boldsymbol{\beta}_2^H \boldsymbol{\sigma}_2 < 0$, which is a contradiction and our proof is complete. ■

Hence, a power-bounded solution will necessarily allow some remaining noise and interference energy, small as it might hopefully be. Further insights into possible "complete" nulling for non-power bounded solutions will be characterized next.

3.3.2 Non-Power-Bounded Solutions

To continue our characterization of the solutions of (3.24) by examining the behavior of non-power-bounded solutions. As mentioned above, there may be situations when $r_{\texttt{NI}} = \text{rank}(\mathbf{R}_{\mathbf{ni}}) < NML$ – for example, if we actually use an estimate $\hat{\mathbf{R}}_{\mathbf{ni}}$ of the noise-and-interference covariance, or if we consider a noise-free case for analysis. In these cases, we know from Theorem 3.7 that $\lambda_P = 0$. This implies that the solution is not power bounded, and $\text{tr}(\mathbf{B}_3) < P_o$. Furthermore, due to Lemma 3.1, $\boldsymbol{\Sigma}_2 = \mathbf{0}_{NML \times N}$. With these in mind, we can produce, at the very least, a flowchart of solution properties in the non-power-bounded case.

As a simple beginning, we note that the matrix complementary slackness conditions reduce to two:

$$\mathbf{R_{ni}}\mathbf{B}_1 = \mathbf{0}_{NML\times NML} \tag{3.36}$$

$$\mathbf{R_{ni}}\mathbf{B}_2 = \mathbf{0}_{NML\times N}. \tag{3.37}$$

This implies that if the solution is not power-bounded, it must null the entire noise-and-interference spectrum. Hence, we know that the matrices $\mathbf{B}_1$ and $\mathbf{B}_2$ have the general form:

$$\mathbf{B}_1 = \mathbf{P}^{\perp}_{\mathbf{R_{ni}}}\mathbf{V}_1\mathbf{P}^{\perp}_{\mathbf{R_{ni}}} \tag{3.38}$$

$$\mathbf{B}_2 = \mathbf{P}^{\perp}_{\mathbf{R_{ni}}}\mathbf{V}_2 \tag{3.39}$$

where $\mathbf{V}_1 \in \mathbb{H}^{NML}_{++}$ and $\mathbf{V}_2 \in \mathbb{C}^{NML\times N}$ are arbitrary matrices, and $\mathbf{P}^{\perp}_{\mathbf{R_{ni}}} = \mathbf{I}_{NML} - \mathbf{R_{ni}}\mathbf{R_{ni}}^{\dagger}$ is the orthogonal projection matrix onto the nullspace of $\mathbf{R_{ni}}$. Additionally, because of the first slackness condition, we know that $\text{rank}(\mathbf{B}_1) \leq NML - r_{\texttt{NI}}$.

This form of $\mathbf{B}_2$ has a feasibility consequence. Recall that a solution is feasible only if the Capon constraint $\text{tr}(\mathbf{B}_2^H\mathbf{T}) = \kappa \neq 0$ is satisfied. If we substitute our new form of $\mathbf{B}_2$ into the constraint, it becomes

$$\text{tr}(\mathbf{B}_2^H\mathbf{T}) = \text{tr}(\mathbf{V}_2^H\mathbf{P}^{\perp}_{\mathbf{R_{ni}}}\mathbf{T}) = \kappa. \tag{3.40}$$

Since $\mathbf{V}_2$ is arbitrary, but not trivial, feasibility is violated when $\mathbf{P}^{\perp}_{\mathbf{R_{ni}}}\mathbf{T} = \mathbf{0}_{NML\times N}$, which only occurs when $\mathcal{R}(\mathbf{T}) \subseteq \mathcal{R}(\mathbf{R_{ni}})$. If we let $\breve{\mathbf{R}}_{\mathbf{ni}} = \mathbf{I}_N \otimes \mathbf{R_{ni}}$, this condition can be vectorized as $\boldsymbol{\tau} \in \mathcal{R}(\breve{\mathbf{R}}_{\mathbf{ni}})$). Hence, there is no feasible solution if the target is embedded in the noise & interference spectrum. We state this directly in the following Lemma.

Lemma 3.4. *If $\lambda_P = 0$ is the only feasible dual variable and the target is embedded in the signal-independent disturbances $\mathcal{R}(\mathbf{T}) \subseteq \mathcal{R}(\mathbf{R_{ni}})$, there is no feasible solution to* (3.24) *or its equivalent problems.*

Continuing on, our Lagrangian minimization becomes

$$\mathbf{C}_V \boldsymbol{\beta}_2 = \tilde{\mu}_C^* \boldsymbol{\tau} \tag{3.41}$$

or, given the form of $\mathbf{B}_2$ above,

$$\mathbf{C}_V \mathbf{P}^{\perp}_{\breve{\mathbf{R}}_{\mathbf{ni}}} \mathbf{v}_2 = \tilde{\mu}_C^* \boldsymbol{\tau} \tag{3.42}$$

where $\mathbf{P}^{\perp}_{\breve{\mathbf{R}}_{\mathbf{ni}}} = \mathbf{I}_N \otimes \mathbf{P}^{\perp}_{\mathbf{R_{ni}}}$ and $\mathbf{v}_2 = \mathrm{vec}(\mathbf{V}_2)$.

Let us momentarily consider the generally unrealistic case when $\mathbf{C}_V$ is full rank (*i.e.*, $Q_{\text{eff}} = N^2ML$). Under the $\lambda_P = 0$ hypothesis, we can directly solve for $\boldsymbol{\beta}_2$:

$$\boldsymbol{\beta}_2 = \tilde{\mu}_C^* \mathbf{C}_V^{-1} \boldsymbol{\tau}. \tag{3.43}$$

However, since $\boldsymbol{\beta}_2 \in \mathcal{N}(\breve{\mathbf{R}}_{\mathbf{ni}})$, $\mathbf{P}_{\breve{\mathbf{R}}_{\mathbf{ni}}} \boldsymbol{\beta}_2 = \mathbf{0}_{N^2ML \times 1}$, so we can also say that

$$\tilde{\mu}_C^* \mathbf{P}_{\breve{\mathbf{R}}_{\mathbf{ni}}} \mathbf{C}_V^{-1} \boldsymbol{\tau} = \mathbf{0}_{N^2ML \times 1} \tag{3.44}$$

If we want this to be a feasible result (*i.e.*, $\tilde{\mu}_C \neq 0$), then the whitened target $\mathbf{C}_V^{-1}\boldsymbol{\tau} \in \mathcal{N}(\breve{\mathbf{R}}_{\mathbf{ni}})$ (and by extension, $\boldsymbol{\tau} \in \mathcal{N}(\breve{\mathbf{R}}_{\mathbf{ni}})$). Hence, if the clutter is full rank, the target must be clear of the noise and interference (and the dimensionality of the available resources must be such that this is possible) for a feasible solution to exist. This is effectively a further restriction of Lemma 3.4.

Assuming $\mathbf{C}_V$ is less than full rank, we can find another condition on the target response for the Lagrange multipliers by premultiplying the above equation by the clutter nullspace projection matrix $\mathbf{P}^{\perp}_{\mathbf{C}_V}$, which becomes $\tilde{\mu}^*_C \mathbf{P}^{\perp}_{\mathbf{C}_V} \boldsymbol{\tau} = \mathbf{0}_{N^2ML\times 1}$. This means that either $\tilde{\mu}_C = 0$ or $\boldsymbol{\tau} \in \mathcal{R}(\mathbf{C}_V)$. Thus, if $\boldsymbol{\tau} \notin \mathcal{R}(\mathbf{C}_V)$ (that is, it has a non-zero component outside of the clutter), then $\tilde{\mu}_C = 0$. The non-embedded scenario also simplifies things considerably, because now:

$$\mathbf{C}_V \mathbf{P}^{\perp}_{\breve{\mathbf{R}}_{\mathbf{ni}}} \mathbf{v}_2 = \mathbf{0}_{N^2ML\times 1} \tag{3.45}$$

Hence, either $\mathbf{v}_2$ (and thus $\boldsymbol{\beta}_2$) lies in $\mathcal{N}(\mathbf{C}_V \mathbf{P}^{\perp}_{\breve{\mathbf{R}}_{\mathbf{ni}}})$ or $\mathbf{v}_2 = \boldsymbol{\beta}_2 = \mathbf{0}_{N^2ML\times 1}$. The second possibility violates the Capon constraint, therefore $\mathbf{v}_2,\ \boldsymbol{\beta}_2 \in \mathcal{N}(\mathbf{C}_V \mathbf{P}^{\perp}_{\breve{\mathbf{R}}_{\mathbf{ni}}})$. However, if (for whatever reason), the target does not lie in this space, then we have reached another infeasiblity result via the Capon constraint. Furthermore, a non-embedded target implies the following lemma:

Lemma 3.5. *Assume the clutter is less than full rank (i.e.,* $\mathrm{rank}(\mathbf{C}_V) < N^2ML$*) and the target is not fully embedded within it (i.e.,* $\boldsymbol{\tau} \notin \mathcal{R}(\mathbf{C}_V)$*). Then, any non-power-bounded solution has an optimal value* $\nu^o_{RBQP-PC} = 0$*.*

This "complete nulling" scenario seems non-physical at first, because an optimal value of zero implies the obtained SINR is infinity. However, recall that the inciting requirement for a non-power-bounded solution is a singular signal-independent disturbance covariance matrix $\mathbf{R}_{\mathbf{ni}}$. This usually arises not from any particular physical characteristics of the observed scattering environment, but from an incomplete estimation or specification of $\mathbf{R}_{\mathbf{ni}}$. Application of this solution to the "true" noise-and-interference situation will result in residual noise-and-interference power emanating from the portion of the spectrum that is not canceled by the solution described by Lemma 3.5. Hence, there are two possible paths a designer can take: accept this uncanceled noise as a consequence of this design, or perform the optimization on a diagonally-loaded equivalent $\hat{\mathbf{R}}_{\mathbf{ni}} = \epsilon \mathbf{I}_{NML} + \mathbf{R}_{\mathbf{ni}}$ (where ϵ is some sufficiently small loading

factor) guaranteed to be positive definite to compensate for the unestimated portion of the spectrum.

If $\boldsymbol{\tau}$ *does* lie entirely within the clutter, then there are a few complex scenarios. First, consider a situation where $\mathcal{N}(\breve{\mathbf{R}}_{\mathbf{ni}}) \subseteq \mathcal{N}(\mathbf{C}_V)$, which means the clutter is subsumed entirely into interference and noise (and thus $\mathcal{R}(\mathbf{C}_V) \subseteq \mathcal{R}(\breve{\mathbf{R}}_{\mathbf{ni}})$ & $Q_{\text{eff}} \leq Nr_{\texttt{NI}}$). In this case, $\mathbf{C}_V \mathbf{P}^{\perp}_{\breve{\mathbf{R}}_{\mathbf{ni}}} = \mathbf{0}_{N^2ML \times N^2ML}$. However, this also means that the *target* is now also embedded in the noise and interference. By Lemma 3.4, there is no feasible solution! In any other case, a solution exists if and only if $\tilde{\mu}^*_C \boldsymbol{\tau} \in \mathcal{R}(\mathbf{C}_V \mathbf{P}^{\perp}_{\breve{\mathbf{R}}_{\mathbf{ni}}})$.

We conclude this portion of our analysis with Figure 3.1, which is a flowchart which summarizes our developments on non-power-bounded solutions.

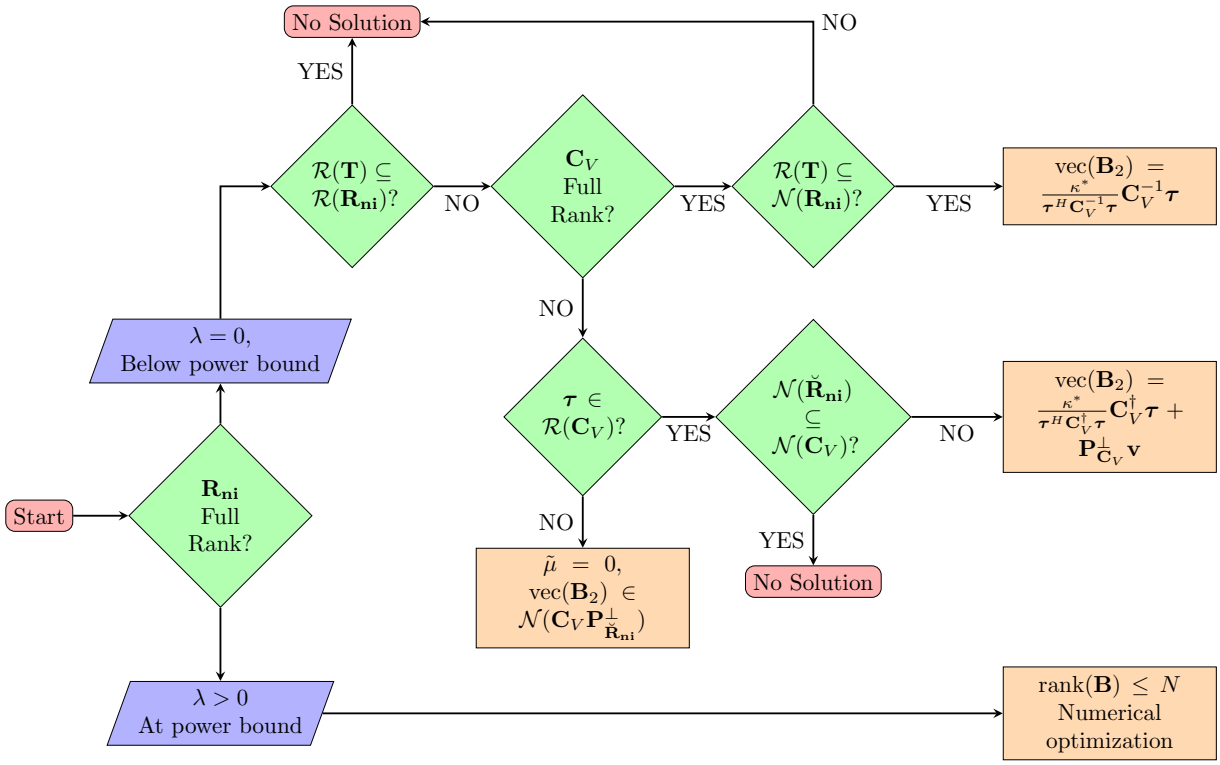

Figure 3.1: Flowchart characterizing non-power-bounded solutions to the RBQP (3.24).

3.3.3 Power-Bounded Solutions

Next, we dig deeper into characterizing power-bounded solutions, which we identified as the most common/practical set of solutions to the relaxed problem (3.24). Hence, for the rest of this section, assume that we have obtained a power-bounded solution to our power-constrained QSDP and the conditions necessary for that to be true are satisfied.

We begin by demonstrating a method to further tighten the rank bound described in Lemma 3.2. In this case, it is obvious that $\lambda_P \mathbf{I}_N \succ 0$ and $\mathbf{R_{ni}} \succ 0$. Furthermore, recall that the matrix slackness multiplier $\mathbf{\Sigma}$, defined in (3.31), is a positive-semidefinite matrix. Let $\eta_{\max}(\mathbf{A})$ indicate the maximum eigenvalue of a Hermitian matrix $\mathbf{A}$. Using these facts about the slackness matrix, we can identify a useful form of the optimal power slackness variable λ_P:

Lemma 3.6. *In a power-bounded solution,* $\lambda_P^o = \eta_{\max}(\mathbf{\Sigma}_2^{oH} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{\Sigma}_2^o)$.

Proof. We prove this directly. Since $\mathbf{\Sigma} \succeq 0$ and both $\lambda_P \mathbf{I}_N \succ 0$ and $\mathbf{R_{ni}} \succ 0$ are full rank, then $\lambda_P \mathbf{I}_N - \mathbf{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{\Sigma}_2 \succeq 0$ by [101, Theorem 7.7.9(e)]. Additionally, by [101, Theorem 7.7.9(c)], $\eta_{\max}(\frac{1}{\lambda_P} \mathbf{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{\Sigma}_2) \leq 1$, or $\lambda_P \geq \eta_{\max}(\mathbf{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{\Sigma}_2)$. Both of these are biconditional statements.

The inequalities above are strict if and only if $\mathbf{\Sigma}$ is *also* positive definite. However, by complementary slackness, we know that $\mathbf{\Sigma} \succ 0$ is part of an optimal KKT tuple if and only if $\mathbf{B}^o = \mathbf{0}_{J \times J}$. We know, based on the equality constraints, that such a primal solution will never be feasible. Thus, based on feasibility, we know that these inequalities cannot be strict. Hence, the only remaining possibility is $\lambda_P^o = \eta_{\max}(\mathbf{\Sigma}_2^{oH} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{\Sigma}_2^o)$. □

Next, consider the following well-known theorem about eigenvalues and their multiplicities:

Theorem 3.9 (Theorem 1.2.18, [101]). *Let* $\mathbf{A} \in \mathbb{C}^{n\times n}$ *and suppose one of its eigenvalues* δ *has algebraic multiplicity* k*. Then* $\operatorname{rank}(\mathbf{A} - \delta\mathbf{I}_n) \geq n - k$ *with equality for* $k = 1$*.*

We can see an immediate connection with our problem and with Lemma (3.2). Since λ_P^o is an eigenvalue of $\boldsymbol{\Sigma}_2^{oH}\mathbf{R}_{\mathbf{ni}}^{-1}\boldsymbol{\Sigma}_2^o$, Theorem 3.9 tells us that $\operatorname{rank}(\lambda_P^o\mathbf{I}_N - \boldsymbol{\Sigma}_2^{oH}\mathbf{R}_{\mathbf{ni}}^{-1}\boldsymbol{\Sigma}_2^o) \geq N - k_{\lambda_P}$, where k_{λ_P} is the algebraic multiplicity of λ_P^o. Combining this and the previous lemmas, we arrive at a far tighter rank bound for power-bounded solutions:

Theorem 3.10. *In a power-bounded solution,* $\operatorname{rank}(\mathbf{B}) \leq k_{\lambda_P}$*, with equality if* $k_{\lambda_P} = 1$*.*

Proof. By combining Lemmas (3.2) and (3.6) and Theorem 3.9, we have

$$\begin{aligned}\operatorname{rank}(\mathbf{B}) &\leq N - \operatorname{rank}(\lambda_P^o\mathbf{I}_N - \boldsymbol{\Sigma}_2^{oH}\mathbf{R}_{\mathbf{ni}}^{-1}\boldsymbol{\Sigma}_2^o)\\ &\leq N - (N - k_{\lambda_P}) = k_{\lambda_P}.\end{aligned}$$

Equality when $k_{\lambda_P} = 1$ is, therefore, rather obvious. □

This theorem provides us with a basis for examining two potential (semi-) closed form power-bounded solutions to (3.24). We name and characterize these solutions primarily by their action on the clutter, through the signal-independent representation $\mathbf{C}_V$.

Nulling-Only Solutions

First, we examine the "nulling-only" solution, which occurs when $\mathbf{B} \in \mathcal{N}(\widetilde{\mathcal{C}}_{\mathrm{V}})$ or, in terms of the vectorized form, $\boldsymbol{\beta}_2^o \in \mathcal{N}(\mathbf{C}_V)$. Such a solution is feasible only if the target is not embedded in the clutter, *i.e.*, $\boldsymbol{\tau} \notin \mathcal{R}(\mathbf{C}_V)$, and if the clutter does not span the entire product space of degrees of freedom, *i.e.*, $\operatorname{rank}\mathbf{C}_V < N^2ML$.

Assuming this is true, we can find an immediate solution for the variable $\mathbf{\Sigma}_2$: $\mathbf{\Sigma}_2^o = -\tilde{\mu}_C^{o*}\mathbf{T}$. Additionally, since a nulling-only solution ensures $\mathbf{C}_V\boldsymbol{\beta}_2 = \mathbf{0}_{N^2ML\times 1}$, (3.35) reduces to $\tilde{\mu}_C^*\kappa = \lambda_P P_o$. These equivalences lead us to the following lemma, which provides an exact solution of λ_P:

Lemma 3.7. *In a power-bounded solution, if* $\boldsymbol{\beta}_2^o \in \mathcal{N}(\mathbf{C}_V)$, $\lambda_P^o = \frac{|\kappa|^2}{P_o}(P_o\ \eta_{\max}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}))^{-1}$ *and* $\tilde{\mu}_C^o = \kappa(P_o\ \eta_{\max}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}))^{-1}$.

Proof. The proof is fairly straightforward. Applying $\mathbf{\Sigma}_2^o = -\tilde{\mu}_C^{o*}\mathbf{T}$ to Lemma (3.6), $\lambda_P = \eta_{\max}(\mathbf{\Sigma}_2^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{\Sigma}_2) = |\tilde{\mu}_C|^2\eta_{\max}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T})$. Substituting this into the optimal value form above then implies λ_P^o is the solution of the quadratic equation:

$$\lambda_P^2\frac{P_o^2\ \eta_{\max}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T})}{|\kappa|^2} - \lambda_P = 0 \tag{3.46}$$

The only roots of this equation are zero, which is infeasible, and $\frac{|\kappa|^2}{P_o^2\eta_{\max}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T})}$. The optimal $\tilde{\mu}_C$ is a direct consequence of λ_P. □

The natural corollary to this is that the optimal value of a nulling-only solution, regardless of the exact vector $\boldsymbol{\beta}_2$, is always

$$\nu_{\text{null}}^o = \frac{|\kappa|^2}{P_o\ \eta_{\max}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T})}, \tag{3.47}$$

implying the optimal relaxed SINR is

$$\text{SINR}_{\text{null}}^o = \frac{|\kappa|^2}{\nu_N^o} = P_o\ \eta_{\max}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}). \tag{3.48}$$

This corresponds exactly with the SINR upper bound proposed by [35], augmented by a factor of P_o, and somewhat surprisingly, does not depend on the Capon gain κ whatsoever. Hence, a nulling-only solution produces a generalized whiten-and-match joint "filtering" of

the scattering environment where the whitening operation corresponds entirely to the signal independent sources.

So far, though, we have only demonstrated that the optimal value of a nulling-only solution is fixed. What does the nulling-only solution $\mathbf{B}_{\text{null}}$ look like? We originally derived this form in [65], but we repeat it here for completeness: If we apply the above analysis to the KKT condition (B.8), then $\mathbf{B}_2 = \frac{P_o}{\kappa}\mathbf{B}_1\mathbf{T}$. In order for this to be completely orthogonal to the clutter, $\mathbf{C}_V \operatorname{vec}(\mathbf{B}_1\mathbf{T}) = \mathbf{C}_V(\mathbf{T}^T \otimes \mathbf{I}_{NML}) \operatorname{vec}(\mathbf{B}_1) = \mathbf{0}_{N^2ML\times 1}$.

Returning to the KKTs ((B.8), in particular), we also have $\mathbf{B}_3 = \frac{P_o^2}{|\kappa|^2}\mathbf{T}^H\mathbf{B}_1\mathbf{T}$. Since scalar feasibility requires that $\operatorname{tr}(\mathbf{B}_3) = P_o$, this implies that $\operatorname{tr}(\mathbf{T}^H\mathbf{B}_1\mathbf{T}) = \frac{|\kappa|^2}{P_o}$, which also satisfies the equality constraints. These results can also be shown to satisfy the conditions (B.8) and (B.9), as well as all of the positive semidefiniteness conditions, so long as $\operatorname{tr}(\mathbf{B}_1\mathbf{R}_{\mathbf{ni}}) \leq \frac{|\kappa|^2}{\|\mathbf{T}\|_F^2}\frac{\operatorname{tr}(\mathbf{R}_{\mathbf{ni}})}{P_o}$. This is necessarily satisfied if the Slater condition in Theorem 3.6 is satisfied. Hence, we have as a possible incomplete solution

$$\mathbf{B} = \begin{bmatrix} \mathbf{B}_1 & \frac{P_o}{\kappa}\mathbf{B}_1\mathbf{T} \\ (\frac{P_o}{\kappa}\mathbf{B}_1\mathbf{T})^H & \frac{P_o^2}{|\kappa|^2}\mathbf{T}^H\mathbf{B}_1\mathbf{T} \end{bmatrix} \tag{3.49}$$

where $\mathbf{B}_1$ minimizes $\operatorname{tr}(\mathbf{B}_1\mathbf{R}_{\mathbf{ni}})$ subject to the above requirements.

The eagle-eyed reader will note correctly that we can reframe the implicit matrix completion in (3.49) as the solution to a linear SDP:

$$\begin{aligned} \min_{\mathbf{B}_1 \in \mathbb{H}_+^{NML}} \quad & \operatorname{tr}(\mathbf{B}_1\mathbf{R}_{\mathbf{ni}}) \\ \text{s.t.} \quad & \mathbf{C}_V(\mathbf{T}^T \otimes \mathbf{I}_{NML}) \operatorname{vec}(\mathbf{B}_1) = \mathbf{0}_{N^2ML\times 1} \\ & \operatorname{tr}(\mathbf{T}\mathbf{T}^H\mathbf{B}_1) = \frac{|\kappa|^2}{P_o} \end{aligned} \tag{3.50}$$

or, perhaps, after collapsing the nullspace requirement into its necessary constituent parts,

$$\begin{aligned} \min_{\mathbf{B}_1 \in \mathbb{H}_+^{NML}} \quad & \mathrm{tr}(\mathbf{B}_1 \mathbf{R}_{\mathbf{ni}}) \\ \text{s.t.} \quad & \mathrm{tr}(\mathbf{T}\boldsymbol{\Gamma}_q^H \mathbf{B}_1) = 0 \quad \forall q \in \{1, \ldots, Q\} \cdot \\ & \mathrm{tr}(\mathbf{T}\mathbf{T}^H \mathbf{B}_1) = \frac{|\kappa|^2}{P_o} \end{aligned} \tag{3.51}$$

Since $\mathbf{B}_1$ is effectively a relaxation of $\mathbf{w}\mathbf{w}^H$, it is now clear that the "responsibility" in a nulling-only solution relies entirely on the receive spectrum. Ironically, the solution of this linear SDP can be much harder to obtain numerically than a direct implementation of (3.24), because $\mathbf{R}_{\mathbf{ni}}$ is often quite dense, and there are at least $Q_{\text{eff}} + 1$ equality constraints to satisfy, also in dense matrices. Thus, we will not pursue this solution any further, but note that efficient large-scale SDP solvers are becoming more popular as a topic of research.

Whiten-and-Null Solutions

Next, we consider a second class of power-bounded solutions we define as "whiten-and-null" solutions, because the operation of $\boldsymbol{\beta}_2$ is to whiten the target w.r.t. the clutter and then also null whatever is left over. As with the nulling-only solution, we require the target to not be wholly embedded in the clutter; that is, $\boldsymbol{\tau} \notin \mathcal{N}(\mathbf{C}_V)$

We begin our construction of the whiten-and-null solution by assuming that the optimal $\boldsymbol{\sigma}_2 = -\tilde{\mu}_C^* \boldsymbol{\tau}^\perp$, where $\boldsymbol{\tau}^\perp = \mathbf{P}_{\mathbf{C}_V}^\perp \boldsymbol{\tau}$ is the part of the target orthogonal to the clutter. This assumption, when substituted into (3.32) leads to the following form for $\boldsymbol{\beta}_2$:

$$\boldsymbol{\beta}_2 = \tilde{\mu}_C^* \mathbf{C}_V^\dagger \boldsymbol{\tau} + \mathbf{z}_2^\perp. \tag{3.52}$$

where $\mathbf{z}_2^{\perp} \in \mathcal{N}(\mathbf{C}_V) \subseteq \mathbb{C}^{N^2ML}$ is an arbitrary clutter-nulling vector. This implicitly also means that $\mathbf{C}_V\boldsymbol{\beta}_2 = \tilde{\mu}_C^*\boldsymbol{\tau}^{\|}$, where $\boldsymbol{\tau}^{\|} = \mathbf{P}_{\mathbf{C}_V}\boldsymbol{\tau}$ is the part of the target embedded in the clutter, which will become useful later.

Now, any feasible solution requires $\boldsymbol{\beta}_2^H\boldsymbol{\tau} = \kappa$. Expanding this equality constraint and solving for $\tilde{\mu}_C$ results in

$$\tilde{\mu}_C = \frac{\kappa - \mathbf{z}_2^{\perp H}\boldsymbol{\tau}}{\boldsymbol{\tau}^H\mathbf{C}_V^+\boldsymbol{\tau}}. \tag{3.53}$$

An interim form of the resulting objective value for this potential solution is therefore

$$\tilde{\mu}_C^*\kappa = \frac{\kappa(\kappa - \mathbf{z}_2^{\perp H}\boldsymbol{\tau})^*}{\boldsymbol{\tau}^H\mathbf{C}_V^+\boldsymbol{\tau}}. \tag{3.54}$$

After substituting this and the other forms into (3.35), we therefore have

$$\lambda_P P_o = \frac{\mathbf{z}_2^{\perp H}\boldsymbol{\tau}(\kappa - \mathbf{z}_2^{\perp H}\boldsymbol{\tau})^*}{\boldsymbol{\tau}^H\mathbf{C}_V^+\boldsymbol{\tau}}. \tag{3.55}$$

In order for this to mean anything, we require another expression for λ_P in terms of the alignment factor $\mathbf{z}_2^{\perp H}\boldsymbol{\tau}$, which we have by combining the assumed form of $\boldsymbol{\sigma}_2$ and $\tilde{\mu}_C$ with Lemma 3.6:

$$\lambda_P = \frac{|\kappa - \mathbf{z}_2^{\perp H}\boldsymbol{\tau}|^2}{(\boldsymbol{\tau}^H\mathbf{C}_V^+\boldsymbol{\tau})^2}\eta_{\max}(\mathbf{T}^{\perp H}\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}^{\perp}). \tag{3.56}$$

After some algebra, we obtain a final expression for the nulling-target alignment $\mathbf{z}_2^{\perp H}\boldsymbol{\tau}$:

$$\mathbf{z}_2^{\perp H}\boldsymbol{\tau} = \kappa\frac{P_o\ \eta_{\max}(\mathbf{T}^{\perp H}\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}^{\perp})}{\boldsymbol{\tau}^H\mathbf{C}_V^+\boldsymbol{\tau} + P_o\ \eta_{\max}(\mathbf{T}^{\perp H}\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}^{\perp})}. \tag{3.57}$$

We can then substitute this expression into every other form we showed above. The final optimal value is, for example,

$$\nu_{WN}^{o} = \frac{|\kappa|^2}{\boldsymbol{\tau}^H \mathbf{C}_V^{+} \boldsymbol{\tau} + P_o \, \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})}, \tag{3.58}$$

which corresponds to an output $\mathrm{SINR}_{WN}^{o} = \boldsymbol{\tau}^H \mathbf{C}_V^{+} \boldsymbol{\tau} + P_o \, \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})$. Other optimal KKT parameters are

$$\tilde{\mu}_C = \frac{\kappa}{\boldsymbol{\tau}^H \mathbf{C}_V^{+} \boldsymbol{\tau} + P_o \, \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})} \tag{3.59}$$

$$\lambda_P = \frac{\eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})}{(\boldsymbol{\tau}^H \mathbf{C}_V^{+} \boldsymbol{\tau} + P_o \, \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp}))^2} \tag{3.60}$$

$$\boldsymbol{\Sigma}_2 = -\frac{\kappa^*}{\boldsymbol{\tau}^H \mathbf{C}_V^{+} \boldsymbol{\tau} + P_o \, \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})} \mathbf{T}^{\perp} \tag{3.61}$$

$$\mathbf{B}_2 = \frac{\kappa^*}{\boldsymbol{\tau}^H \mathbf{C}_V^{+} \boldsymbol{\tau} + P_o \, \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})} \mathbf{T}_w + \mathbf{Z}_2^{\perp} \tag{3.62}$$

where $\mathbf{T}_w$ is the matrix equivalent of the whitened target vector $\mathbf{C}_V^{\dagger} \boldsymbol{\tau}$. It is then possible to further derive the rest of the solution from here, but we leave this for future investigations. Without delving too deeply into the algebra, we note that this is essentially an eigenvector solution of an equation that appears after significant manipulations of the KKTs.

It is possible that this is not optimal for a particular choice of $\boldsymbol{\tau}$, depending on how "embedded" the target is in the clutter and the relative strength of the clutter-target components with respect to the noise, etc. Since we have another power-bounded solution, it is worthwhile to compare when once is more "optimal" than the other in these terms. Clearly, if the target is well-separated from the clutter ($\boldsymbol{\tau} \in \mathcal{N}(\mathbf{C}_V)$), then the optimal values (and hence the optimal solutions) must coincide, because there is no more target-in-clutter to whiten. If $\boldsymbol{\tau}$ has any components in the clutter, however, the nulling-only solution is "more optimal"

than the whiten-and-null solution above if and only if

$$\nu_N^o < \nu_{WN}^o. \tag{3.63}$$

This corresponds to the following data-dependent inequality

$$P_o \left(\eta_{\max}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}) - \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp}) \right) > \boldsymbol{\tau}^H \mathbf{C}_V^+ \boldsymbol{\tau} \tag{3.64}$$

where the inequality has flipped due to considering the denominators. It would seem that $\eta_{\max}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}) - \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})$ should be non-negative, since $\mathbf{T}^{\perp}$ is "lesser" than $\mathbf{T}$ (in a norm sense) if any of the target resides in the clutter. Thus, it seems reasonable to say that a nulling-only solution is optimal if and only if

$$P_o > \frac{\boldsymbol{\tau}^H \mathbf{C}_V^+ \boldsymbol{\tau}}{\eta_{\max}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}) - \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})}. \tag{3.65}$$

Hence, for a nulling-only solution to be more optimal than a whiten-and-null solution, the power must exceed (essentially) the target-to-clutter ratio over the in-clutter target-to-noise-and-interference ratio.

This can be more clearly seen if we assume that the signal-independent sources are statistically white; that is, let $\mathbf{R}_{\mathbf{ni}} = \sigma_n^2 \mathbf{I}_{NML}$. Then, the ratio in question is

$$\begin{aligned} \frac{\boldsymbol{\tau}^H \mathbf{C}_V^+ \boldsymbol{\tau}}{\eta_{\max}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}) - \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}^{\perp})} &= \frac{\sum_{i=1}^{Q_{\text{eff}}} \frac{|\operatorname{tr}(\mathbf{T}^H \mathbf{U}_i)|^2}{\sigma_{\mathsf{c},i}^2}}{\frac{\eta_{\max}(\mathbf{T}^H \mathbf{T}) - \eta_{\max}(\mathbf{T}^{\perp H} \mathbf{T}^{\perp})}{\sigma_n^2}} \\ &= \sum_{i=1}^{Q_{\text{eff}}} \frac{|\operatorname{tr}(\mathbf{T}^H \mathbf{U}_i)|^2}{\sigma_{max}^2(\mathbf{T}) - \sigma_{max}^2(\mathbf{T}^{\perp})} \frac{\sigma_n^2}{\sigma_{\mathsf{c},i}^2} \end{aligned}$$

We can clearly see, then, that a nulling-only solution is "more optimal" if

$$\frac{P_o}{\sigma_n^2} > \sum_{i=1}^{Q_{\text{eff}}} \frac{|\operatorname{tr}(\mathbf{T}^H \mathbf{U}_i)|^2}{\sigma_{\mathsf{c},i}^2(\sigma_{max}^2(\mathbf{T}) - \sigma_{max}^2(\mathbf{T}^\perp))}. \tag{3.66}$$

This would seem to make sense, because if the power is sufficiently large (or the expected maximum SNR is, at least), then nulling is optimal, regardless of the relative alignment of the target with the clutter. Conversely, if the target is sufficiently clutter aligned, then merely nulling the clutter directions would provide unacceptable performance.

Finally, we note that these two power-bounded solutions are not the only ones that satisfy the KKTs, but they are the most intuitively informative. Future research would do well to continue this kind of characterization and the particulars of the degrees of freedom

3.4 Waterfilling-like Behaviors

Using some of these general results from above, we can show that the optimal solution to the KKTs at the power bound satisfies equations that resemble the well-known "waterfilling" concept, which is informed by the work of Bell [16, 17], Kay [40], and Setlur, et al. [42], among others.

For the remainder of this section, we consider the "canonical" representation of $\mathbf{C}_V$ as examined in 3.1.1. Now, given this form of the clutter matrix, we can show a waterfilling-like effect by combining the matrix slackness condition (B.9) and one of the Lagrangian conditions. If we recast (3.32) in matrix form, and use the set of partial isometries, it expands to

$$\mathbf{\Sigma}_2 = \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \operatorname{tr}(\mathbf{U}_i^H \mathbf{B}_2)\, \mathbf{U}_i - \tilde{\mu}_C^* \mathbf{T}. \tag{3.67}$$

Substituting this into (B.9) gives us

$$\mathbf{R_{ni}}\mathbf{B}_2 = \tilde{\mu}_C^* \mathbf{T}\mathbf{B}_3 - \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \operatorname{tr}(\mathbf{U}_i^H \mathbf{B}_2)\, \mathbf{U}_i \mathbf{B}_3. \tag{3.68}$$

Let us assume that we are power-bounded and everything that implies from the lemmas above. Thus, we can "directly" find $\mathbf{B}_2$ by applying $\mathbf{R}_{\mathbf{ni}}^{-1}$ above to obtain:

$$\mathbf{B}_2 = \tilde{\mu}_C^* \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}\mathbf{B}_3 - \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \operatorname{tr}(\mathbf{U}_i^H \mathbf{B}_2)\, \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{U}_i \mathbf{B}_3 \tag{3.69}$$

Premultiplying both sides with $\mathbf{U}_j^H$ and taking the trace gives us

$$\begin{aligned} \operatorname{tr}(\mathbf{U}_j^H \mathbf{B}_2) = \tilde{\mu}_C^* \operatorname{tr}(\mathbf{U}_j^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}\mathbf{B}_3) \\ - \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \operatorname{tr}(\mathbf{U}_i^H \mathbf{B}_2)\ \operatorname{tr}(\mathbf{U}_j^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{U}_i \mathbf{B}_3). \end{aligned} \tag{3.70}$$

We can extract the jth term from the sum and collect it on the left hand side to obtain

$$\begin{aligned} &\operatorname{tr}(\mathbf{U}_j^H \mathbf{B}_2)(1 + \sigma_{\mathsf{c},j}^2\ \operatorname{tr}(\mathbf{U}_j^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{U}_j \mathbf{B}_3)) \\ &\qquad = \tilde{\mu}_C^* \operatorname{tr}(\mathbf{U}_j^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T}\mathbf{B}_3) \\ &\qquad\quad - \sum_{\substack{i=1 \\ i \neq j}}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \operatorname{tr}(\mathbf{U}_i^H \mathbf{B}_2)\ \operatorname{tr}(\mathbf{U}_j^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{U}_i \mathbf{B}_3). \end{aligned} \tag{3.71}$$

Now, the second term on the left-hand side of the above equation could be divided out as long as there was a guarantee it was always positive. First, if $\sigma_{\mathsf{c},j}^2 = 0$, this term is 1, which is positive. Otherwise, $\sigma_{\mathsf{c},j}^2 > 0$ since they are the non-zero eigenvalues of a positive semidefinite matrix. Next, we need to examine the trace statement. $\mathbf{B}_3$ is positive semidefinite by construction, and $\mathbf{U}_j^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{U}_j$ is positive definite because each $\mathbf{U}_j$ is a full rank scaled partial isometry and $\mathbf{R}_{\mathbf{ni}}$ is positive definite (see [101, Theorem 7.7.2]). Thus, the trace of this

matrix product is always positive and real. With this in hand, we can now say

$$\begin{aligned}\operatorname{tr}(\mathbf{U}_j^H\mathbf{B}_2) = \tilde{\mu}_C^* &\frac{\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}\mathbf{B}_3)}{1+\sigma_{\mathsf{c},j}^2\,\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_j\mathbf{B}_3)} \\ &- \sum_{\substack{i=1\\ i\neq j}}^{Q_{\text{eff}}} \frac{\sigma_{\mathsf{c},i}^2\,\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_i\mathbf{B}_3)}{1+\sigma_{\mathsf{c},j}^2\,\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_j\mathbf{B}_3)}\operatorname{tr}(\mathbf{U}_i^H\mathbf{B}_2). \end{aligned} \tag{3.72}$$

Since $\mathbf{B}_3$ is a relaxed version of $\mathbf{s}\mathbf{s}^H$, we can regard the expression $\sigma_{\mathsf{c},j}^2\,\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_j\mathbf{B}_3)$ to be a relaxed form of $\sigma_{\mathsf{c},j}^2\,\mathbf{s}^H\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_j\mathbf{s}$, which is (effectively) a clutter-to-noise-and-interference ratio for the jth basis matrix. Similarly, $\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}\mathbf{B}_3)$ is a joint target-and-clutter to noise-and-interference ratio, and the cross term $\sigma_{\mathsf{c},i}^2\,\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_i\mathbf{B}_3)$ captures the ratio of patch-to-patch interaction and the noise-and-interference level. Thus, the first term on the right hand side is effectively a normalized measure of the clutter-matched target spectrum given a signal basis $\mathbf{B}_3$, and the right hand side is a normalized measure of the intraclutter spectrum given that same basis.

Traditional waterfilling dictates that power is preferentially injected to bands where the overall signal-to-noise ratio is high, proceeding to "worse"-off bands until the available power is exhausted. This is somewhat *inverted* in (3.72) because the left-hand side is an unscaled version of the subspace alignment between the solution $\mathbf{B}_2$ and the jth canonical clutter transfer matrix $\mathbf{U}_j$. Naively, we would like to *minimize* this alignment for $j \in \{1,\ldots,Q_{\text{eff}}\}$, since aligning with the clutter would nominally degrade our matching of the target. However, the coupled nature of (3.72) requires more nuance than that. First, observe that for the non-clutter directions, *i.e.*, $j \in \{Q_{\text{eff}}+1, ..., N^2ML\}$, $\sigma_{\mathsf{c},j}^2 = 0$ and (3.72) becomes

$$\begin{aligned}\operatorname{tr}(\mathbf{U}_j^H\mathbf{B}_2) = \tilde{\mu}_C^*&\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}\mathbf{B}_3) \\ &- \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2\,\operatorname{tr}(\mathbf{U}_j^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_i\mathbf{B}_3)\operatorname{tr}(\mathbf{U}_i^H\mathbf{B}_2). \end{aligned} \tag{3.73}$$

This means that the solution's alignment in the non-clutter directions follows that of the whitened target's, less the combined crossover between non-clutter and clutter in the whitened spectrum. If the target is strong in these directions, the solution will align towards them. Otherwise, the solution must match the target in-spectrum as near as possible in directions both where the whitened clutter power is the lowest and the alignment with other directions is minimized. This corresponds to the findings in [42], which showed a similar behavior in two-step mutual information waveform design over consecutive transmit epochs. In this case, the "two steps" can be regarded as the trans-receive pair instead of sequential temporal designs, and the "separation" is between target and clutter instead of different targets.

But what of the value $\tilde{\mu}_C$? This clearly relates to the total available resources – in this case, κ and P_o – and provides us with the "water" in waterfilling. We can find a form of $\tilde{\mu}_C$ by reexamining (3.69) as follows. First, we premultiply by the target matrix $\mathbf{T}^H$ and take the trace of both sides. Recognizing that a feasible solution satisfies $\mathrm{tr}(\mathbf{T}^H\mathbf{B}_2) = \kappa^*$, we now have

$$\begin{aligned}\kappa^* = \tilde{\mu}_C^* \,&\mathrm{tr}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}\mathbf{B}_3) \\ &- \sum_{i=1}^{Q_{\mathrm{eff}}} \sigma_{\mathsf{c},i}^2 \,\mathrm{tr}(\mathbf{U}_i^H\mathbf{B}_2)\ \mathrm{tr}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_i\mathbf{B}_3).\end{aligned} \tag{3.74}$$

This is where the traditional waterfilling appears, since we are saying that the gain across the target (which we know in this case to be κ^*) is the upper bound on the available resources (given by the first term) minus the overall impact of the clutter weighted by its alignment (the second term). According to the above argument, $\mathrm{tr}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}\mathbf{B}_3)$ is never zero. Therefore, after rearrangement, a final form for $\tilde{\mu}_C^*$ is:

$$\tilde{\mu}_C^* = \frac{\kappa^* + \sum_{i=1}^{Q_{\mathrm{eff}}} \sigma_{\mathsf{c},i}^2 \,\mathrm{tr}(\mathbf{U}_i^H\mathbf{B}_2)\ \mathrm{tr}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{U}_i\mathbf{B}_3)}{\mathrm{tr}(\mathbf{T}^H\mathbf{R}_{\mathbf{ni}}^{-1}\mathbf{T}\mathbf{B}_3)} \tag{3.75}$$

Note that the second term in the numerator must have the same collective phase as κ^*, because this is $\tilde{\mu}_C^*$'s phase as well. This may be inserted into (3.72) and then solved for a matrix equation that we will leave for later study.

Furthermore, it implies that the optimal value of any solution to the KKTs is therefore

$$\nu^o_{RBQP-PC} = \tilde{\mu}_C^* \kappa = \frac{|\kappa|^2 + |\kappa| \left| \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \operatorname{tr}(\mathbf{U}_i^H \mathbf{B}_2) \operatorname{tr}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{U}_i \mathbf{B}_3) \right|}{\operatorname{tr}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T} \mathbf{B}_3)}, \tag{3.76}$$

and hence the optimal $\text{SINR}_{RBQP-PC} = \frac{|\kappa|^2}{\nu^o_{RBQP-PC}}$ is

$$\text{SINR}^o_{RBQP-PC} = \frac{\operatorname{tr}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T} \mathbf{B}_3)}{1 + \frac{1}{|\kappa|} \left\| \sum_{i=1}^{Q_{\text{eff}}} \sigma_{\mathsf{c},i}^2 \operatorname{tr}(\mathbf{U}_i^H \mathbf{B}_2) \operatorname{tr}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{U}_i \mathbf{B}_3) \right\|}. \tag{3.77}$$

Since $\operatorname{tr}(\mathbf{B}_3) \leq P_o$ and the denominator is greater than one, we also have the upper bound $\text{SINR}^o_{RBQP-PC} \leq P_o \operatorname{tr}(\mathbf{T}^H \mathbf{R}_{\mathbf{ni}}^{-1} \mathbf{T})$. These directly coincide with the expressions for the power-bounded cases above, and can provide the basis for limiting arguments on the magnitude of $|\kappa|$ which we will not discuss here. Thus, the optimal solution to the relaxed problem (particularly when power-bounded) describes a generalized whiten-and-match trans-receive filter process that exhibits waterfilling behavior, shaping the transmit process to match the target's response in clutter as much as possible while still minimizing the clutter response.

3.5 Numerical Analysis

In this section, we demonstrate the utility of our proposed relaxation scheme through comparative numerical simulation. While we have preliminarily reduced solving the QSDP (3.24) to a matrix completion problem, it is possible to solve the problem numerically with commercial solvers. Most of the analysis presented here is enabled by the modeling toolbox CVX [87, 86] and the solvers SDPT3 [89, 90] and SeDuMi [91]. Unless otherwise stated, the sce-

narios presented to the solver were as follows: As noted in [65], our available computational resources and current simulation environment constrain our overall problem size. Therefore, in this case, we assume $N = 5$, $M = 5$, and $L = 16$. However, we do not believe that this is an inherent limitation which precludes the applicability of our technique to larger problem size – as we have discussed above, there lingers the possibility of using purpose-built QSDP solvers like QSDPNAL to explore larger problem sizes.

The radar operates on a carrier frequency of $f_o = 1$ GHz and transmits pulses at a pulse repetition frequency $f_p = 20$ KHz. The receive array has an interelement spacing of $d = \lambda_o/2$. The noise covariance matrix was a scaled correlation matrix with correlation function $\exp(-0.05|n|)$ for $n \in \{1, \ldots, NML\}$. Interferers were placed at the azimuth-elevation pairs (0.3941, 0.3) radians and (-0.4951, 0.3) radians, with correlation matrices given by the Toeplitz matrix associated with the correlation function $\exp(0.2|l|)$ for $l \in \{1, \ldots, NL\}$. The clutter was modeled with $Q = 25$ patches of $P = 5$ scatterers each, equally spaced over the azimuth interval $(-\pi/2, \pi/2)$ at an elevation angle of $\frac{\pi}{4}$ radians. Adjacent scatterers within each patch are correlated with coefficient -0.2. The target was placed at the angle coordinates $(\theta_t, \phi_t) = (0.3, \frac{\pi}{3})$ and travels at a relative normalized Doppler frequency of -0.255. In all cases, we have made the Capon constraint real.

As mentioned in Chapter 3.1.3, we can obtain approximate rank-one solutions of (2.24) from our relaxed problem by considering the approximate sum vector of basis-length K, $\mathbf{b}_K^a$, generated by the K largest non-zero eigenvalues of $\mathbf{B}$. From this, we can recover the approximate pair $(\mathbf{w}_K^a, \mathbf{s}_K^a) = (\mathbf{\Psi}_W \mathbf{b}_K^a, \mathbf{\Psi}_S \mathbf{b}_K^a)$, which is then rescaled as necessary to meet the constraints. In practice, violations of the equality constraint are minimal and do not affect the overall performance of our method. In most cases, we consider $K = 1$, though $K = N$ often reduces the need for rescaling.

Some of the following results involve comparisons with competing algorithms previously described in Chapter 2, the details of which we recall below. First, we consider the signal-

filter EigenAM (henceforth, AM) scheme of Setlur & Rangaswamy [8], which Tang and Tang [47] later reported as their "Algorithm 1" despite operating on a different cost function. This technique initializes the problem in (2.24) with a fixed signal $\mathbf{s}_{init}$ and solves for $\mathbf{w}$ (the $\mathbf{w}$-step), then uses that optimal $\mathbf{w}$ in (2.24) to solve for a new $\mathbf{s}$ (the $\mathbf{s}$-step), and so on until convergence. We also compare against a modified version of the spectral AM scheme of Aubry, et al. [44] (AA2), which attempts to directly maximize the SINR through a similar procedure to AM, but uses a linear SDP in the $\mathbf{s}$-step. Our modification removes the similarity constraint, which necessitates a small change in the rank-1 signal decomposition – the similarity constraint matrix is replaced by the target gain matrix. It is here we note two primary advantages of our method: it requires no problem-dependent initialization, other than what the available solving method dictates, and it is not an explicitly iterative scheme.

3.5.1 Characterization of the Relaxed Solution

We begin by briefly examining the resulting solution matrix $\mathbf{B}^o$ of the relaxed design problem in (3.24) and related formulations for the common scenario. Despite QSDPs lacking the solution rank guarantees of linear SDPs, we have found that, depending on the parameter scaling, the given solver produces solutions of rank no greater than N, as predicted by Lemma 3.2. This rank property of $\mathbf{B}^o$ is demonstrated in Figure 3.2, which displays the solution's first $2N$ eigenvalues on a logarithmic scale over variation in κ and P_o for the common scenario above. First, we examine variation in κ. Here, κ ranges from 10^{-2} to $NML \times 10^4$ while $P_o = 10^7$. This range is dictated by the Slater condition (Theorem 3.6), and hence choosing $|\kappa|^2 < NMLP_o = 4 \times 10^9$. Figure 3.2(a) demonstrates that generally, as κ decreases for a fixed power level (here, $P_o = 10^7$), the effective numerical rank of $\mathbf{B}_o$ approaches N. This indicates a transition region must exist where the desired gain falls below some threshold dictated by the noise, interference, and clutter characteristics. In contrast, as shown in Figure 3.2(b), varying the power P_o for a fixed κ (here, $\kappa = 100$) only affects the

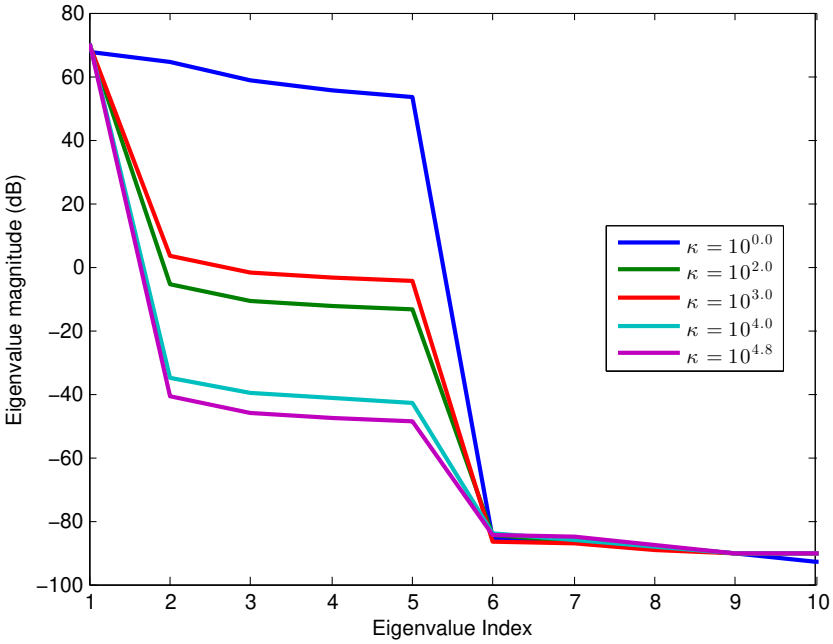

(a) Capon constraint κ, $P_o = 10^7$

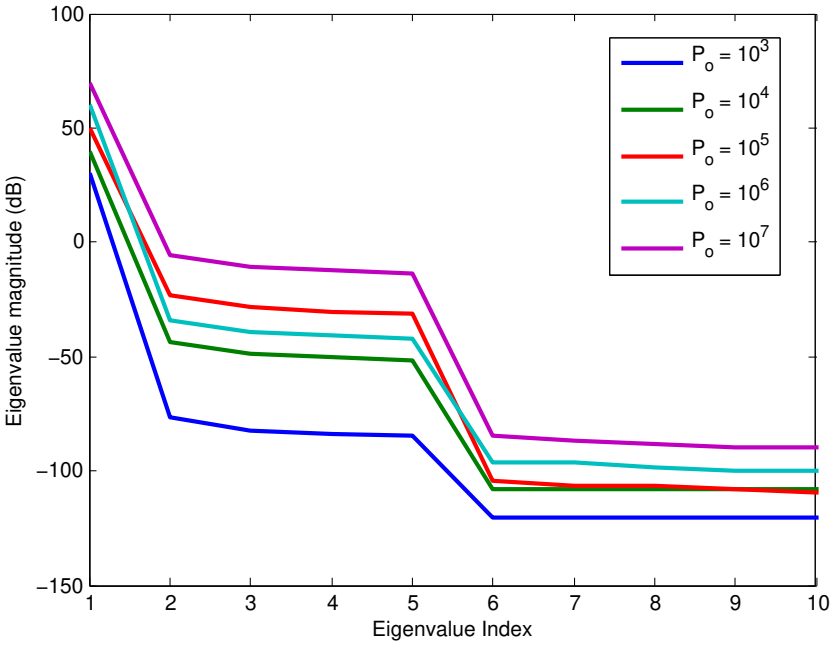

(b) Power constraint P_o, $\kappa = 100$.

Figure 3.2: First $2N$ eigenvalues of power-constrained relaxed solution for varying (a) Capon constraint κ, (b) power constraint P_o. ©2018 IEEE

overall size of the eigenvalues, not the rank. We note again that our variation in P_o meets the Slater condition, as the lower bound for power is $|\kappa|^2/NML = 25$. As demonstrated in [65], this also has implications for the utility of constructed rank one approximations. Additionally, we note that even if $\text{rank}(\mathbf{B}) \geq 1$, the individual eigenvectors themselves are often not feasible, which strengthens the validity of choosing the approximate vector $\mathbf{b}_K^a$. In summary, we have demonstrated that across most parameter choices, the relaxed solution is effectively rank one, which validates our choice of approximate retrieval scheme.

3.5.2 Interference Effects

In this subsection, we consider the impact of interference on the relaxed solution only. Our goal here is to determine if the relaxation produces some unknown or unexpected bias or inability to null or mitigate the impact of interference due to the relaxation itself.

To examine this effect, we consider the traditional adapted pattern for STAP, which plots the following function

$$\mathcal{P}(f_d, \theta;\ \phi) = |\mathbf{w}_o^H(\mathbf{v}(f_d) \otimes \mathbf{s}_o \otimes \mathbf{a}(\theta, \phi)|^2. \tag{3.78}$$

That is, the adapted pattern for a given beamformer-signal pair $(\mathbf{w}_o, \mathbf{s}_o)$ is a function of the Doppler frequency f_d and the azimuth θ at a given elevation ϕ. We consider the same simulation parameters as above, with and without interference, and provide the overall adapted pattern at an elevation cut matched to the target.

With no interference, the naive implementation performs relatively well, as illustrated in Figure 3.3. The target is well localized, with the peak of the pattern at its location, and a deep null is steered along the clutter ridge (represented by an X at each patch's angle-doppler phase center).

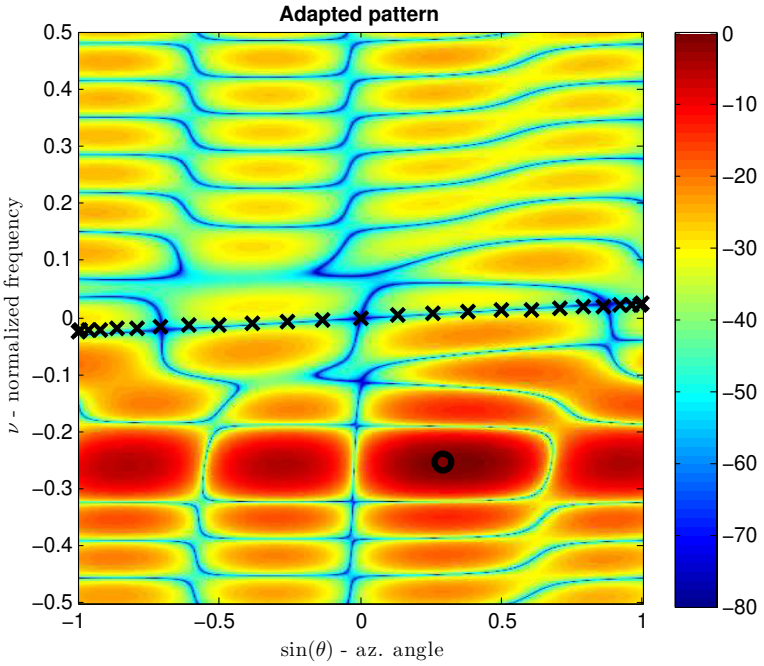

Figure 3.3: Adapted Pattern (dB relative to peak), RBQP Solver. Target at ◯, Clutter phase centers at ×. No Interference. ©2018 IEEE

Further evidence of the excellent localization of the target can be seen in cuts of the adapted pattern along the target azimuth and Doppler frequency, shown in Figures 3.4 and 3.5 below. In each case, the target position (shown by the dashed line) aligns with the peak of the adapted pattern cut. The broad mainlobe width and high sidelobe peaks evident in Figure 3.4 are a function of the relatively small number of antennas considered. Performance could be improved by a multiplicative spatial taper or adding more antennas. We also see in Figure 3.5 that the clutter suppression extends beyond the deep null along the primary clutter ridge, eliminating additional sidelobes "behind" the ridge at higher Doppler frequencies.

Now suppose we inject a broadband interferer "close" to the target – as a reminder, the target is at $(\theta_t, \phi_t) = (0.3, \frac{\pi}{3})$ radians and the interferer is at $(\theta_I, \phi_I) = (0.3941, 0.3)$ radians. Figure 3.6 shows the adapted pattern under these conditions. Note that a clear bias is introduced by the interferer, since in this case we lack the spatial receive resources necessary to null it exactly and resolve the target simultaneously. Observe, however, that an aliased null

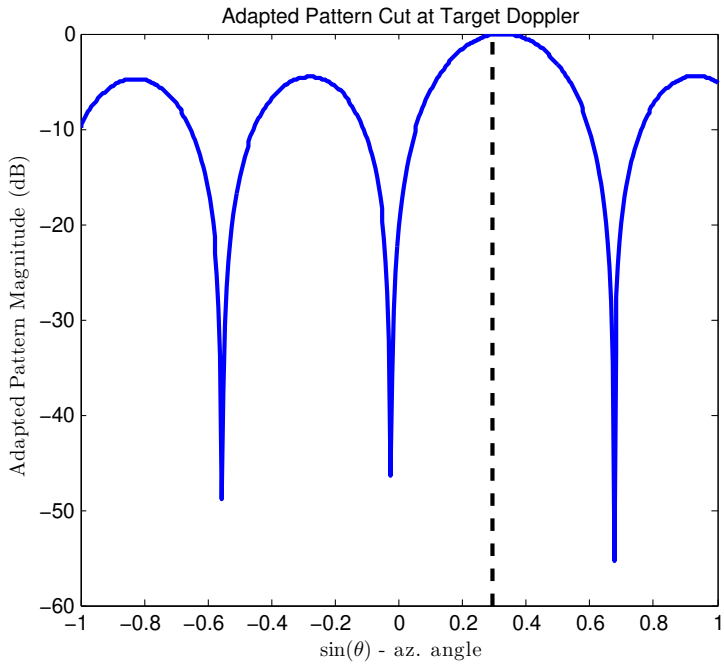

Figure 3.4: Adapted Pattern (dB scale), cut at target Doppler, RBQP Solver. Target azimuth at dashed line.

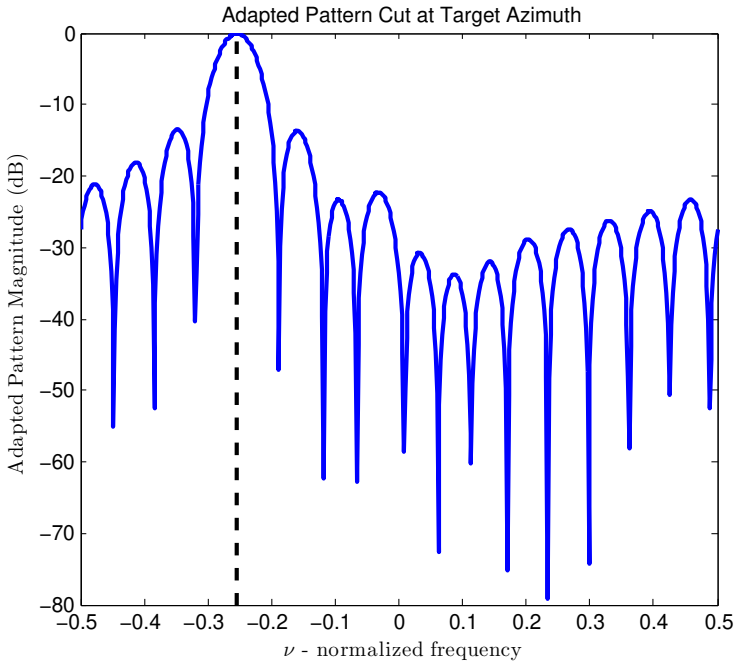

Figure 3.5: Adapted Pattern (dB scale), cut at target azimuth, RBQP Solver. Target Doppler at dashed line.

is present in the spectrum at approximately $\theta = \arcsin(0.2) = 2\theta_I \sin(\phi_I)\sin(\phi_t)$, indicating that it is merely a sampling issue and not a failure of the scheme to identify the interference. Furthermore, we have found that if the elevation cut of the adapted pattern is taken at the elevation of the interferer, a clear null at the appropriate azimuth appears.

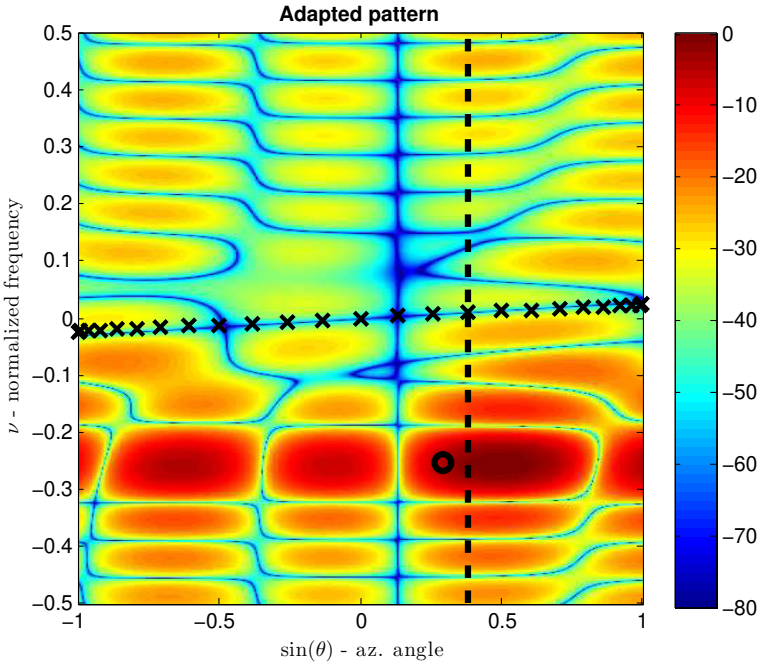

Figure 3.6: Adapted Pattern (dB relative to peak), RBQP Solver. Target at ◯, Clutter phase centers at ×, Interferer (dashed line) at $(\theta, \phi) = (0.3941, 0.3)$ radians

For completeness, we also show the adapted pattern cuts under the interference scenario as well. Again, in all figures, the dashed line represents the target location. Figure 3.7 show the cut along the target Doppler. The significant bias seen above is more clear here, which would result in a rather large angle estimation error. This also, indirectly, demonstrates a disconnect between signal-filter design intended for position estimation versus detection.

The cut along the target azimuth, shown in Figure 3.8, still appropriately localizes the target in some sense. However,the peak gain at the target Doppler has dropped about 5 dB from the gain in Figure 3.5. Thus, there is a clear loss in both dimensions due to the close interferer.

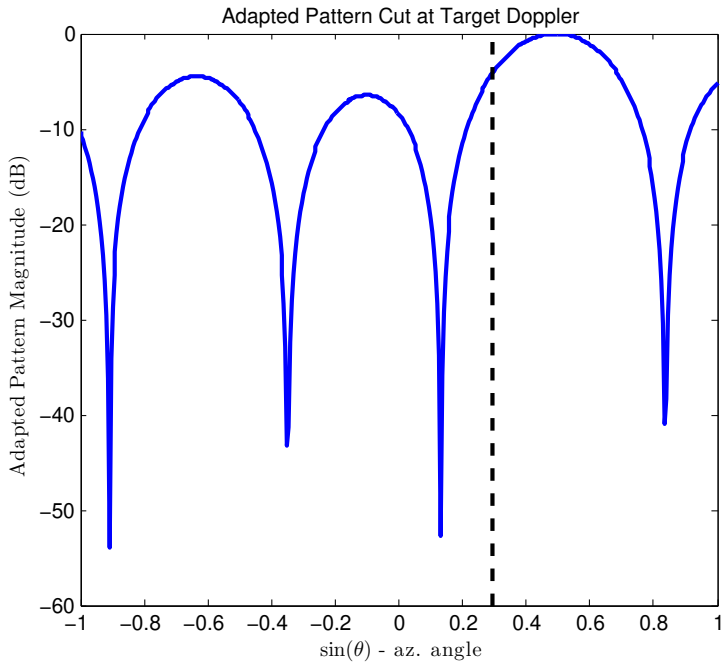

Figure 3.7: Adapted Pattern (dB scale), cut at target Doppler, RBQP Solver. Interference scenario.

We also observe more clearly the additional suppression of the "behind" clutter sidelobes, which relate to the nulling of the aliased location of the interferer.

We note that this interference impact (in all cases) is mitigated by increasing M in our simulations. This is not surprising, since more antenna elements over the same aperture increases the degrees of freedom available to null interference and localize the target. It is also entirely eliminated if the interference is well-separated from the target in space.

3.5.3 Comparison with Alternating Minimization Schemes

In this subsection, we compare the performance of our approximate relaxed BQP (RBQP) solution to the existing methodologies in the literature described above.

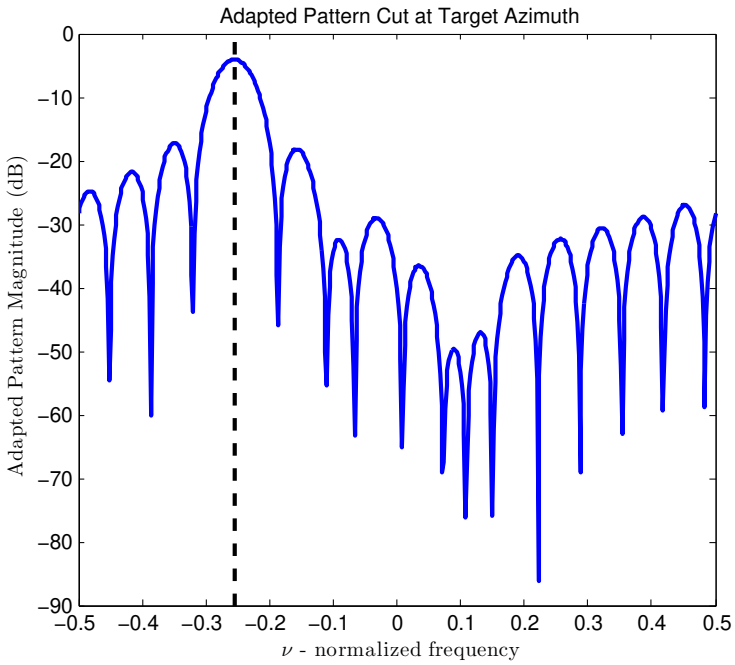

Figure 3.8: Adapted Pattern (dB scale), cut at target azimuth, RBQP Solver. Interference scenario.

SINR Comparisons

One means of comparison between these algorithms is the robustness of the output SINR for each scheme to variations in the target Doppler frequency f_d. For ease of comparison, we presented the two iterative solvers with parameters we found to assist their convergence: each was initialized with the same chirp waveform having a time-bandwidth product of 50, but the waveform used to initialize AA2 was scaled to have $P_o = 1$, and the resulting SINR was then rescaled to match the actual power constraint $P_o = 10^7$. This is because the authors in [44] claim the algorithm is invariant to scaling[4]. To further facilitate equal comparison, we set $\kappa = \sqrt{P_o}$, which is what AA2 natively produces after rescaling. The convergence threshold for AA2 is set to $\epsilon = 10^{-3}$, and AM is terminated after 20 iterations (which provided a comparable convergence factor). Fig. 3.9 shows this metric for each of the

[4]As we will see, however, this is not true for the formulation in [44], which neglects the power entirely. The authors of [47] corrected this mistake, but as we will see in Chapter 4, this does not change the overall contours of our findings.

considered algorithms – our scheme (labeled RBQP), AM, and AA2 – when applied to the common scenario mentioned previously.

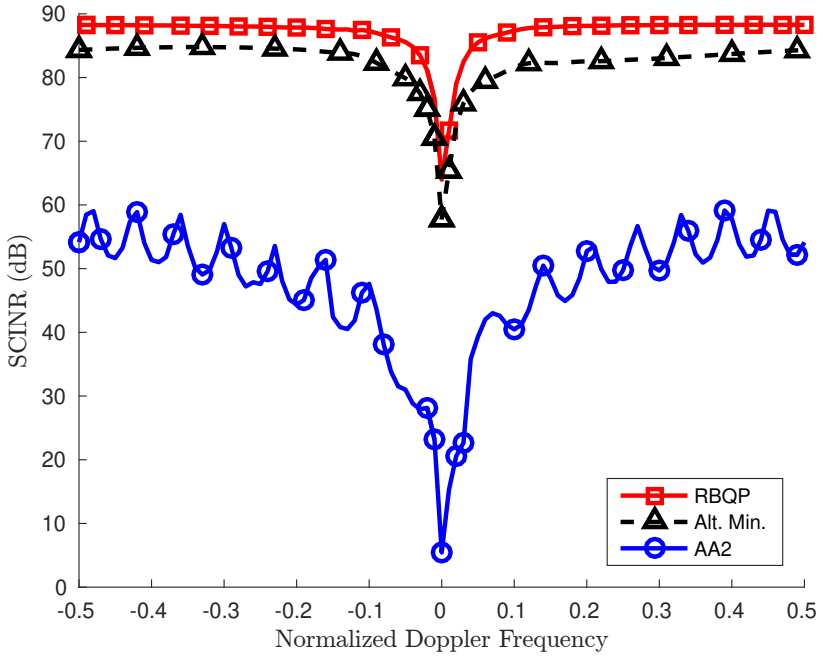

Figure 3.9: Comparison of SINR as a function of target Doppler frequency for multiple power-constrained algorithms. ©2018 IEEE

First, it is clear that the solutions produced by RBQP and AM are extremely robust to changes in target Doppler at scale, while AA2 has significant variation. Furthermore, in this scenario, our method generally outperforms both competing algorithms in terms of overall SINR and minimum detectable velocity. The relative gap between RBQP and AM at larger Doppler shifts makes sense, given that our method is a relaxation of that in [8] and that AM is terminated rather quickly.

We note here that additional simulations have shown AA2 can outperform both the RBQP procedure and AM, but this requires a number of highly controlled and generally unrealistic assumptions. Namely, the convergence threshold must be fairly high (say, $\epsilon = 1$, if not higher) and $P_o = 1$. Even under these conditions, AA2 often merely produces a permuted

version of the initializing waveform, which does not significantly change the initial objective value.

Adapted Pattern Comparisons

The performance benefits provided by RBQP can also be observed in the adapted patterns mentioned in the subsection above. Consider a case where only the second interferer, located at the azimuth-elevation pair (-0.4951, 0.3) radians, is active. Furthermore, assume that the clutter and target are both at the same elevation angle as the interferer. The target remains at an azimuth of 0.3 radians but, due to the elevation change, the target relative Doppler is now -0.087. We note that the only major difference from the previous scenarios is the change in elevation angle for the clutter patches and target.

Figure 3.10 shows the adapted pattern at the target elevation cut for each algorithm, with the 0 dB reference set at the notional peak for each algorithm – by construction, this is $|\kappa|^2 = P_o$. However, as we showed in Figure 3.9, this does not mean that the algorithms have no difference in detection properties, since the SINR obtained by RBQP is superior to the other two algorithms. Therefore, the adapted patterns shown should not be viewed as directly comparable in terms of SINR, but in terms of relative features and cancellation near the interference, clutter, and target. In all patterns, we can see that the target is well-resolved, deep nulls are placed along the clutter ridge, and a null is placed near the azimuth of the interferer. We can see, though, that both our method (Figure 3.10(a)) and AM (Figure 3.10(b)) further shape the spectrum by shifting the phase centers of the target response's Doppler sidelobes away from that of the known target (demarcated by the $\bigcirc$) and reducing other sidelobes nearer to the ridge. The relative similarity between these two adapted patterns is generally unsurprising, given the similar SINR performance shown above.

In contrast, AA2 (Figure 3.10(c)) retains a more typical grating lobe structure, with no additional nulling or shaping beyond the clutter ridge and interference null. This, in turn, corresponds to the lower SINR seen in Figure 3.9. There are a variety of possible explanations for this behavior, the specifics of which are the subject of ongoing research that will be reported in the future.

Detection Performance Comparisons

Next, we demonstrate more specifically the impact of the obtained SINR seen in Figure 3.9 on overall detection performance, which has been recently adapted by Setlur and Rangaswamy [104]. For a Gauss-Gauss problem, the general performance of the detection problem (2.20) can be defined by the relationship between the probability of detection (deciding $\mathscr{H}_1$ with a target present) P_d, the probability of false alarm (deciding $\mathscr{H}_1$ with no target present) P_{fa}, and the *deflection coefficient* d^2, which is dependent on the SINR. Mathematically, we have [105]:

$$P_d = Q(Q^{-1}(P_{fa}) - \sqrt{d^2}) \tag{3.79}$$

where $Q(x)$ is the complementary cumulative distribution function (CCDF) of the standard normal distribution, and $Q^{-1}(p)$ is its inverse error function. In this case, the deflection coefficient is given by the SINR in linear terms as evaluated at a particular filter-signal combination; that is, given a (hopefully optimal) pair $(\mathbf{w}^\star, \mathbf{s}^\star)$, the deflection coefficient is

$$d^2(\mathbf{w}^\star, \mathbf{s}^\star) = 2|\rho_t|^2 \frac{|\mathbf{w}^{\star H}\mathbf{T}\mathbf{s}^\star|^2}{\mathbf{w}^{\star H}\mathbf{R}_\mathbf{u}(\mathbf{s}^\star)\mathbf{w}^\star}. \tag{3.80}$$

We now consider the deflection coefficients generated by each algorithm – $d^2_{RBQP}, d^2_{AM}, d^2_{AA2}$ – and examine their resulting behavior under different levels of target scaling. Let d^2_{base} be the baseline deflection coefficient. If the target's amplitude is at a gain of G relative

to the baseline deflection coefficient (corresponding to a particular SINR), then the target amplitude is $|\rho_t|^2 = G/d_{base}^2$. We consider the most challenging scenario for every algorithm: a target at zero Doppler. Furthermore, we vary G to be 0, 3, and 6 dB above the baseline deflection, which we assume to be d_{RBQP}^2.

Our findings are presented in Figure 3.11 below. To put it lightly, there is no contest, as AA2's scaling issue results in detection performance that is functionally identical to chance at scale. Due to the smaller gap in SINR at zero Doppler between AM and RBQP, the performance hit is not as drastic, but the gain of RBQP is readily apparent. As expected, in all cases, detection performance improves as the target gain increases, even for AA2. We note that if P_o is smaller, *e.g.*, $P_o = 1$, this advantage can collapse, because the scaling error in AA2 is no longer an issue. However, this is also dependent on the other factors (convergence criterion, relative scaling of the signal-independent disturbances, etc.) we mentioned above and, as such, we feel confident that our method's performance is legitimate.

Runtime Comparisons

Not only is a performance improvement obtained, but the RBQP algorithm has a reduced computational load compared with AA2 in [44], as shown in Table 3.1, which lists the median runtime and standard deviation for each algorithm across the Doppler frequency sweep above. While our algorithm is slower than AM (for a fixed number of iterations),

Table 3.1: Runtime Comparison of Power-Constrained Algorithms ©2018 IEEE

Alg.	Med. runtime (sec)	St. Dev. (sec)
RBQP	93.09	4.394
AM	30.95	N/A
AA2	286.53	793.17

AA2 is inconsistent and slower than either RBQP or AM. In fact, AA2's execution time

varied significantly with the target Doppler's relative alignment to the clutter. Again, the restrictive cases mentioned above can improve the runtime of AA2, but at the cost of the aforementioned caveats.

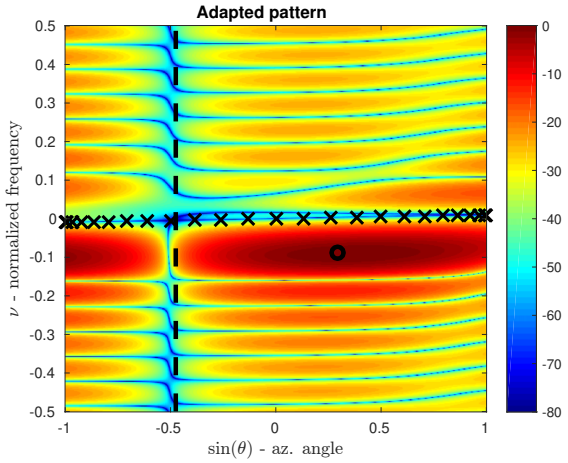

(a) Relaxed BQP (RBQP)

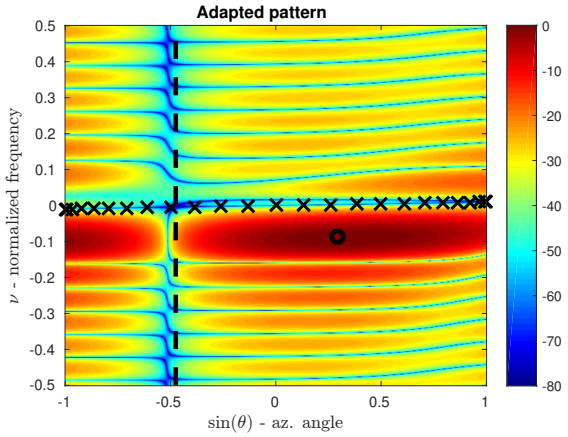

(b) Alt. Min. (AM)

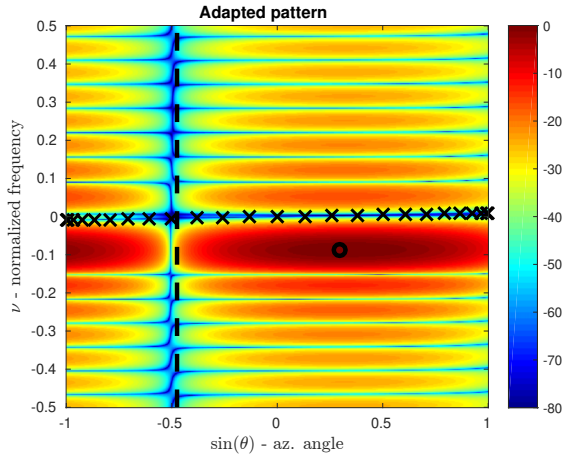

(c) Alg. 2 of [44] (AA2)

Figure 3.10: Adapted Patterns (dB rel. to peak) for (a) RBQP, (b) AM, (c) AA2. Target at ◯, Clutter phase centers at ×, Interferer (dashed line) at $(\theta, \phi) = (-0.4951, 0.3)$ radians.

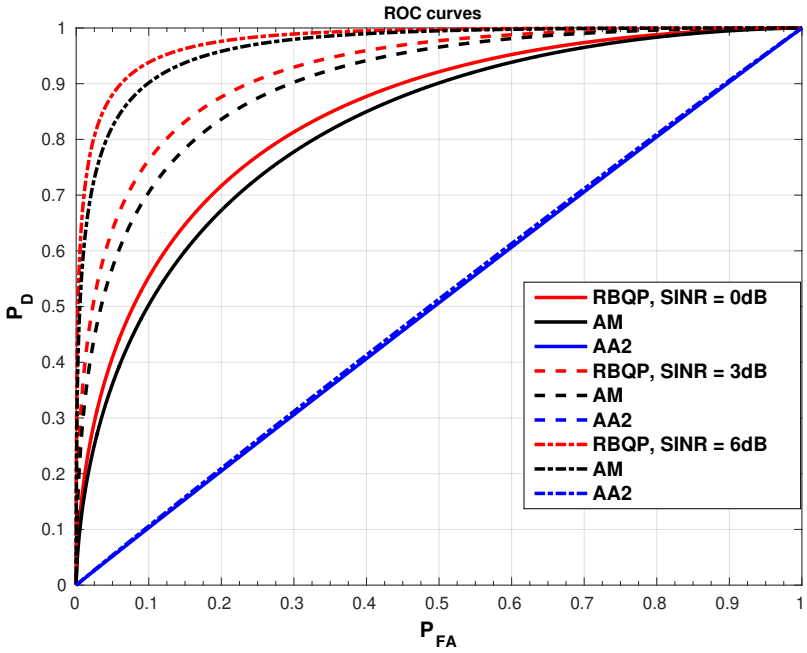

Figure 3.11: Receiver operating characteristic (ROC) curves for multiple power-constrained algorithms as a function of target gain.

Chapter 4

Joint Relaxation Design under Realistic Waveform Constraints

As noted in Chapter 1, optimal signal-filter pairs generated under simple power constraints often have excellent SINR/SCINR performance, but are also typically unsuitable for realistic implementation. This is for a variety of reasons, including significant variation in amplitude and phase, poor ambiguity properties, and transmission in undesired frequency bands. In short, both hardware constraints and target separability concerns drive the need for further design constraints, particularly on the transmitted waveform.

In this chapter, we will investigate the feasibility of extending our joint convex relaxation technique to two particularly salient and relevant waveform constraints: constant modulus/envelope (CM/CE) and the "similarity" constraint (SC). Each of these has specific motivations for their application and utility. The constant modulus constraint is employed to ensure the waveform is suitable for efficient generation by linear amplifiers and balanced RF chains, while the "similarity" constraint was initially motivated by so-called "ambiguity inheritance" properties. We will see that while both constraint sets can be individually

(and jointly) accommodated by by our relaxation, the ability of existing numerical solvers to directly recover good solutions can be somewhat limited, particularly in constant modulus scenarios. Thankfully, we will also propose a refinement technique that employs gradient descent to mitigate this shortcoming, though any one of the signal-filter iterative techniques of Chapter 2 would provide the same results. Indeed, this refinement process may be applied to any such joint relaxation process, but we demonstrate its utility first with the CM-constrained problem.

Many of the solution paths and processes we will investigate in this chapter are directly translatable from those generically described in Chapter 3. Therefore, we will only replicate what is absolutely necessary for clarity's sake.

The rest of the chapter is organized as follows. To introduce further the peculiarities of adding these constraints, Chapter 4.1 will discuss both types of constraints and their motivations in further detail than previous chapters. The rest of the chapter will be generally split along constraint lines. The CM problem, its KKTs, and the refinement strategies will form the basis of Chapter 4.2. Similarly, the elements of the SC problem will be considered in Chapter 4.3. Simulations for both CM and SC scenarios as applied to WA-STAP will be explored in Chapter 4.4, where we will definitively show the extensive gain that can come from our technique. We note that portions of this chapter constitute the recently-submitted work [106].

4.1 Applying Realistic Waveform Constraints

Regardless of the specific radar model under consideration, many initial attempts at waveform design (by itself or jointly with a receiver) neglect practical constraints to simplify the analysis, as we have above. Unfortunately, neglecting realistic constraints in the design

process can limit both realizability and performance of the designed waveform-filter pair when deployed in an actual system. Several design constraints have been considered in the literature to address these issues, each with a particular impact or goal that affects realistic usage.

First, we consider the constant modulus/constant envelope constraint (henceforth CM), also known as a phase-only design. A constant modulus constraint limits the complex envelope of each signal chip to a certain value ρ; for the vector $\mathbf{s}$, this amounts to

$$|s_i| = \rho$$

for every element s_i. This is a feature of a variety of well-studied waveform classes, like frequency-modulated chirps, phase-coded pulse trains, and the venerable rectangular pulse, primarily because it has favorable ambiguity function properties. This constraint is obviously non-convex, but as we will show below, it can easily be accommodated into our framework with a small modification.

A related constraint is the peak-to-average power ratio (PAPR/PAR). For the aforementioned signal $\mathbf{s}$, it manifests mathematically as

$$\frac{\max_i |s_i|^2}{\|\mathbf{s}\|^2} \leq \varsigma$$

where $\varsigma \in [1, N]$, and ensures that amplitude modulation (intentional or otherwise) does not induce extreme spikes in instantaneous power that might strain the transmit amplifier's dynamic range and lead to clipping. The relationship between PAPR and CM is fairly simple, as any CM pulse has a PAPR of 1, regardless of the given modulus parameter ρ. In that sense, a PAPR constraint is a looser constraint than CM, and is used if the implementing hardware on either end of the processing chain can permit mild deviations. While it has been used elsewhere in the literature (see, e.g. [36]), we will not discuss it further here.

Second, we consider what is popularly known as the "similarity" constraint (SC), which limits the deviation between the designed waveform $\mathbf{s}$ and some known waveform $\mathbf{s}_g$ to a (usually small) upper bound in a norm sense. This constraint has manifested itself in a variety of ways in the literature (see above), but it is usually applied through a p-norm, i.e.,

$$\|\mathbf{s} - \mathbf{s}_g\|_p^2 \leq \epsilon$$

Typically, $p = 2$ or ∞.

The usual justification for this constraint is exemplified by [45], whose authors state that the known signal $\mathbf{s}_g$ acts "as a benchmark, which allows the designed waveform to share some of the good ambiguity properties of the known waveform." A further example of this meme is circularly defined in [44], who use it for pulse code design: "[I]mposing [the constraint], it is possible to indirectly control the ambiguity function of the considered coded pulse train: the smaller [ϵ], the higher the degree of similarity between the ambiguity functions of the desired radar code and [$\mathbf{s}_g$]." However, a precise description of the properties inherited by the perturbed signal is, to our knowledge, not available in the literature. At most, as Patton [78] notes in his dissertation, the work of Sussman in [107] demonstrates that a least-squares approximation of ambiguity functions leads to a least-squares approximation of signal vectors under an energy constraint.

Even with this in mind, it is entirely possible for a "small" perturbation to lead to a significant deviation from a "known good" ambiguity property. We address this contradiction with a small continuous-time example. Assume we wish to design a complex signal function $s(t)$ to minimize a cost functional similar to (2.24). Furthermore, let us assume that $s(t)$ is "similarly" constrained in the L_2-norm sense by a goal signal function $s_g(t)$, i.e.

$$\int_{-\infty}^{\infty} |s(t) - s_g(t)|^2 \, \mathrm{d}t \;\leq\; \epsilon. \tag{4.1}$$

It is reasonable to assume that an optimal $s(t)$ will take the form $s_o(t) = s_g(t) + p(t)$, where $p(t)$ is some perturbation function orthogonal to $s_g(t)$. Clearly, $\int_{-\infty}^{\infty} |p(t)|^2 \, \mathrm{d}t \leq \epsilon$ for $s_o(t)$ to be optimal (or even feasible). Now, consider the ambiguity function [108] $\chi_o(\tau, \nu)$ of the optimal signal at delay τ and frequency ν, defined as

$$\chi_o(\tau,\ \nu) = \int_{-\infty}^{\infty} s_o(t) s_o^*(t+\tau) e^{\jmath 2\pi\nu t} \ \mathrm{d}t \tag{4.2}$$

and let $\chi_g(\tau, \nu)$ and $\chi_p(\tau, \nu)$ be the similarly-formed ambiguity functions of $s_g(t)$ and $p(t)$, respectively. By expanding $s_o(t)$ in its ambiguity function and collecting terms, we see that $\chi_o(\tau, \nu)$ is also equal to

$$\begin{aligned} \chi_o(\tau,\nu) &= \chi_g(\tau,\nu) + \chi_p(\tau,\nu) \\ &\quad + \psi(\tau,\nu;p,s_g) + \psi(\tau,\nu;s_g,p) \end{aligned} \tag{4.3}$$

where $\psi(\nu, \tau; s_g, p)$ is the cross-ambiguity function defined by $\psi(\nu, \tau; s_g, p) = \int_{-\infty}^{\infty} s_g(t) p^*(t+\tau) e^{\jmath 2\pi\nu t} \ \mathrm{d}t$, and vice versa for $\psi(\nu, \tau; p, s_g)$. Now, it is well known [108] that the maximum of any ambiguity function occurs at the pair $(\tau, \nu) = (0, 0)$. Based on the "similarity" constraint, we know that $\chi_p(0, 0) \leq \epsilon$, and the cross-term $\psi(0, 0; s_g, p) = 0$ due to the orthogonality. While the peak of the ambiguity function is not significantly perturbed, at any (τ, ν) pair, we have

$$\begin{aligned} |\chi_o(\tau,\nu)| &\leq |\chi_g(\tau,\nu)| + |\chi_p(\tau,\nu)| \\ &\quad + |\psi(\tau,\nu;p,s_g)| + |\psi(\tau,\nu;s_g,p)| \qquad (4.4) \\ &\leq |\chi_g(\tau,\nu)| + \epsilon + |\psi(\tau,\nu;p,s_g)| + |\psi(\tau,\nu;s_g,p)|. \qquad (4.5) \end{aligned}$$

Then, due to the Cauchy-Schwartz inequality,

$$|\psi(\tau,\nu;s_g,p)|^2 \leq \|s_g(t)\|^2 \|p(t)\|^2 = \|s_g(t)\|^2 \epsilon, \tag{4.6}$$

which is also true for $\psi(\tau, \nu; p, s_g)$. Hence, at any (τ, ν) pair, we have

$$|\chi_o(\tau, \nu)| - |\chi_g(\tau, \nu)| \leq \epsilon + 2\sqrt{\epsilon}\|s_g(t)\|^2 \tag{4.7}$$

where equality is possible for certain poor choices of $p(t)$, τ, and ν. Under appropriate scaling, both $|\chi_o(\tau, \nu)|$ and $|\chi_g(\tau, \nu)|$ are less than one, but the right hand side of (4.7) could approach one, e.g., for $\epsilon = 0.1$. Furthermore, even this bound does not guarantee *where* such variations will occur in the time-frequency domain. This example supports the view in [78], which determined that unless ϵ is exceedingly small, the two-norm similarity constraint's overall control of the ambiguity function is, at best, tenuous. While many authors who consider the similarity constraint use the infinity norm (rather than the two-norm) to impose phase-only control [45, 38, 37], the functional interpretation does not change much, if at all.

With this in mind, we propose a more realistic characterization of the motivation for using a "similarity" constraint. Namely, it is a "proximal" constraint, whereby the designed waveform is kept within some radius of the "goal" waveform for regularization purposes. To motivate this proximal interpretation, consider that all signal-dependent clutter mitigation schemes require accurate prior knowledge of the non-target scattering phenomena to be observed. One way to obtain this information is by probing the environment with pilot signals – for example, the strategy proposed in [3]. It is reasonable to assume that the validity of the estimated channel information for subsequent adaptation periods depends, in the small sample regime, upon the proximality of a newly-designed waveform-filter pair to the original pilot signals. That is, the signal we design must span a similar portion of the clutter "space". This is more reasonable than "ambiguity inheritance" since, as we will show, the 2-norm and its relaxed counterparts admit a subspace interpretation. Another possible interpretation of the proximal constraint comes from a control theory perspective. Consider a realistic radar system with actuation lag, i.e. the desired rate of adaptation/control is faster than the hardware is able to adapt. Such a rate limitation imposes an upper bound on, say, the

transmit energy that can be changed from one adaptation period to another. This naturally translates to a proximal operator between the current adaptation period k and the previous period $k-1$. Whatever the interpretation, we will show later that even existing methods that purport to use this constraint appropriately can violate this constraint quite severely.

4.2 Constant Modulus Joint Design

In this section, we analyze the application of the QSDP relaxation technique to the constant-modulus constrained joint design problem. First, we describe the initial problem and then explain how it can be relaxed into a convex QSDP. Then, the conditions for strong duality and the KKTs are examined for insights into the solvability and the tightness of the relaxation to the original problem. We then consider multiple methods to obtain good rank-one approximate solutions from the QSDP, inspired by common practice in linear SDP, via eigendecomposition and refinement.

4.2.1 Relaxation Paths

Under a constant modulus constraint, where we keep each chip in the pulse at a constant amplitude ρ, the unrelaxed problem is:

$$\begin{aligned} \min_{\mathbf{w},\mathbf{s}} \quad & \mathbf{w}^H \mathbf{R}_{\mathbf{u}}(\mathbf{s}) \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{T} \mathbf{s} = \kappa \\ & |s_i| = \rho \quad \forall i = 1, \ldots, N \end{aligned} \qquad . \tag{4.8}$$

or, alternatively, in vectorized form:

$$\begin{aligned} \min_{\mathbf{w},\mathbf{s}} \quad & \mathbf{w}^H \mathbf{R}_{\mathbf{u}}(\mathbf{s}) \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{T} \mathbf{s} = \kappa \quad \cdot \\ & \mathbf{s} \odot \mathbf{s}^* = \rho^2 \mathbf{1}_N \end{aligned} \tag{4.9}$$

Using the notation from (2.24), we can see that the signal constraint set $\Omega_{\mathbf{s}}$ is given by $\Omega_{\mathbf{s}} = \{\mathbf{s} | \mathbf{s} \odot \mathbf{s}^* = \rho^2 \mathbf{1}_N\}$.

In order to get to the appropriate relaxation of the biquadratic problem (4.9) in the form of the QSDP (3.22), we must identify the transform $\mathcal{T}_{CM}(\cdot)$ that maps the modulus constraint to its matrix equivalent. We can achieve this by first defining the diagonal selection matrix $\mathbf{L}_{d,p} \in \mathbb{Z}^{p^2 \times p}$ as $\mathbf{L}_{d,p} = [\text{vec}(\mathbf{E}_{11}) \ \text{vec}(\mathbf{E}_{22}) \ \cdots \ \text{vec}(\mathbf{E}_{pp})]$ where $\mathbf{E}_{kk}$ is the $N \times N$ single-entry matrix with a one on the kth element of its diagonal and zero elsewhere. It can be shown [109] that this matrix has the following property in regards to the Hadamard product:

Property 4.1. $\mathbf{L}_{d,p}^T (\mathbf{A} \otimes \mathbf{B}) \mathbf{L}_{d,q} = \mathbf{A} \odot \mathbf{B}$*, for any* $p \times q$ *matrices* $\mathbf{A}, \mathbf{B}$*..*

Noting that the Kronecker product $\mathbf{s} \otimes \mathbf{s}^*$ is equivalent to $\text{vec}(\mathbf{s}\mathbf{s}^H)$, we can apply the property to the modulus constraint and obtain $\mathbf{L}_{d,N}^T \text{vec}(\mathbf{s}\mathbf{s}^H) = \rho^2 \mathbf{1}_N$. As we have shown previously in Chapter 3, $\mathbf{s}\mathbf{s}^H$ relaxes to $\mathbf{\Psi}_S \mathbf{B} \mathbf{\Psi}_S^T$, and therefore our relaxed mapping for constant modulus signals is

$$\mathcal{T}_{CM}(\Omega_{\mathbf{s}}) = \left\{ \mathbf{B} : \mathbf{L}_{d,N}^T \text{vec}(\mathbf{\Psi}_S \mathbf{B} \mathbf{\Psi}_S^T) = \rho^2 \mathbf{1}_N \right\}. \tag{4.10}$$

Alternatively, one might consider this set to be the collection of trace constraints

$$\mathcal{T}_{CM}(\Omega_{\mathbf{s}}) = \left\{ \mathbf{B} : \text{tr}(\mathbf{E}_{kk} \mathbf{\Psi}_S \mathbf{B} \mathbf{\Psi}_S^T) = \rho^2, \forall k \in \{1, 2, \ldots, N\} \right\} \tag{4.11}$$

with $\mathbf{E}_{kk}$ defined as above. We note that [47] showed a similar mapping for constant-modulus relaxation in the signal-design step of an alternating-minimization type algorithm for "joint" design. In any case, this is a set of N affine equality constraints, which is a convex polytope – hence, our relaxation is convex. In terms of the generic QSDP (3.12), we can define the affine equality-constraint operator $\mathcal{A}_{CM} : \mathbb{H}^J \to \mathbb{R}^N$ as

$$\mathcal{A}_{CM}(\mathbf{B}) = \begin{bmatrix} \mathrm{tr}(\mathbf{\Psi}_S^T \mathbf{E}_{11} \mathbf{\Psi}_S \mathbf{B}) \\ \mathrm{tr}(\mathbf{\Psi}_S^T \mathbf{E}_{22} \mathbf{\Psi}_S \mathbf{B}) \\ \vdots \\ \mathrm{tr}(\mathbf{\Psi}_S^T \mathbf{E}_{NN} \mathbf{\Psi}_S \mathbf{B}) \end{bmatrix}. \tag{4.12}$$

4.2.2 The Slater Condition for Constant Modulus

While we have shown that the constant modulus constraint set under a joint relaxation is convex, we would also like to guarantee reasonable conditions for optimal solutions. As in Chapter 3.2, we look to the Slater condition as an easily computable/identifiable constraint qualification that ensures there is no duality gap and that the Lagrange multipliers/dual variables we rely on for the analysis below exist and are bounded. In the constant modulus case, this amounts to finding a $\mathbf{B} \succ 0$ such that

$$\mathbf{L}_{d,N}^T \boldsymbol{\beta}_3 = \rho^2 \mathbf{1}_N \tag{4.13a}$$

$$\boldsymbol{\beta}_2^H \boldsymbol{\tau} = \kappa \tag{4.13b}$$

where $\boldsymbol{\tau}$, $\boldsymbol{\beta}_2$, and $\boldsymbol{\beta}_3$ are defined in Chapter 3 above.

While there are many possible solutions that satisfy this condition, as in Chapter 3 and Appendix B, we wish to find a simpler necessary condition that would be satisfied by a

reasonably implemented radar design. This lead us to the following theorem for the primal constant-modulus Slater condition:

Theorem 4.1 (Primal CM Slater's Condition)**.** *The constant modulus QSDP satisfies the Slater condition if* $N\rho^2 > \frac{|\kappa|^2}{\|\mathbf{T}\|_F^2}$.

Proof: Details can be found in Appendix C. ■

This mirrors Theorem 3.6, which states that so long as the Capon constraint does not exceed the maximum gain from the target, there will always be at least one interior point and no duality gap for a finite primal objective. Again, since violating the Slater condition implies violating the principle of conservation of energy, we will assume from here on that it is satisfied and therefore can characterize primal solutions exactly through the dual variables.

4.2.3 KKT Conditions

Since we have assumed that the Slater condition holds for all cases of note, we now turn to investigating the KKTs of the constant modulus QSDP, which are necessary and sufficient for any optimal solution $\mathbf{B}^o$.

Recall that the Lagrange multiplier associated with the Capon constraint is $\tilde{\mu}_C$. Furthermore, define a vector $\tilde{\boldsymbol{\mu}}_M \in \mathbb{R}^N$ to be the Lagrange multiplier associated with the modulus constraint – replacing the generic set inclusion dual variable $\boldsymbol{\lambda}_T$. This also means that the semidefiniteness slackness variable $\boldsymbol{\Sigma}$ from (3.28) becomes

$$\boldsymbol{\Sigma} = \begin{bmatrix} \mathbf{R_{ni}} & \boldsymbol{\Sigma}_2(\tilde{\mu}_C) \\ \boldsymbol{\Sigma}_2(\tilde{\mu}_C)^H & \mathrm{diag}(\tilde{\boldsymbol{\mu}}_M) \end{bmatrix}$$

where the operator $\mathrm{diag}(\cdot)$ places the elements of a length-N vector down the diagonal of an $N \times N$ matrix. This slackness variable must be positive semidefinite to be part of a feasible

solution, which requires:

$$\mathcal{R}(\boldsymbol{\Sigma}_2^H) \subseteq \mathcal{R}(\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M)) \tag{4.14a}$$

$$\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M) \succeq \boldsymbol{\Sigma}_2^H \mathbf{R}_{\mathbf{ni}}^{\dagger} \boldsymbol{\Sigma}_2 \tag{4.14b}$$

where we have dropped the explicit dependence on $\tilde{\mu}_C$ for the cross-slackness variable $\boldsymbol{\Sigma}_2 = \sum_{i=1}^{Q_{\mathrm{eff}}} \nu_i \,\mathrm{tr}(\mathbf{U}_i^H \mathbf{B}_2)\, \mathbf{U}_i - \tilde{\mu}_C^* \mathbf{T}$. The first semidefiniteness condition provides us with our first feasibility lemma. Define the l_0 "norm" as the number of nonzero components of a real vector, denoted by $\|\cdot\|_0$.

Lemma 4.1. *For an optimal primal-dual set of variables,* $\mathrm{rank}(\boldsymbol{\Sigma}_2) \leq \|\tilde{\boldsymbol{\mu}}_M\|_0$

The proof is immediate due to the range inclusion. One corollary stemming from this lemma and the range inclusion is that for a feasible solution, the ith column of $\boldsymbol{\Sigma}_2$ is zero if and only if the corresponding multiplier $\tilde{\mu}_{M,i} = 0$. Carrying this further, this implies an overall feasibility theorem for the dual variables $\boldsymbol{\Sigma}_2$ and $\tilde{\boldsymbol{\mu}}_M$:

Theorem 4.2. $\tilde{\boldsymbol{\mu}}_M = \mathbf{0}_{N\times 1}$ *is a feasible dual variable if and only if* $\boldsymbol{\Sigma}_2 = \mathbf{0}_{NML\times N}$.

This is similar to Lemma 3.1, since the modulus dual variable $\tilde{\boldsymbol{\mu}}_M$ plays a similar role to the power dual variable λ_P in the power constrained case. When this theorem's conditions are satisfied, the only "action" of the cross-spectrum $\mathbf{B}_2$ is to whiten the target in clutter, since it implies that $\mathbf{C}_V \boldsymbol{\beta}_2 = \tilde{\mu}_C^* \boldsymbol{\tau}$.

But when and how do these dual variables impact the primal solution in other ways? We can simplify the necessary and sufficient matrix complementarity condition ($\mathbf{B}\boldsymbol{\Sigma} = \mathbf{0}_{J\times J}$)

into four joint conditions:

$$\mathbf{B}_1\mathbf{R}_{\mathbf{ni}} = -\mathbf{B}_2\boldsymbol{\Sigma}_2^H \tag{4.15}$$

$$\mathbf{B}_1\boldsymbol{\Sigma}_2 = -\mathbf{B}_2\,\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M) \tag{4.16}$$

$$\mathbf{B}_2^H\mathbf{R}_{\mathbf{ni}} = -\mathbf{B}_3\boldsymbol{\Sigma}_2^H \tag{4.17}$$

$$\mathbf{B}_2^H\boldsymbol{\Sigma}_2 = -\mathbf{B}_3\,\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M) \tag{4.18}$$

These mirror the conditions (B.7)–(B.10), with an appropriate change of signal constraint dual variables. It is here we recall the idea of *non-trivial* solutions; that is, solutions where no component of $\mathbf{B}$ is zero. When combined with Theorem 4.2, the condition (4.15) indicates that the dual variable tuple $(\tilde{\boldsymbol{\mu}}_M, \boldsymbol{\Sigma}_2) = (\mathbf{0}_{N\times 1}, \mathbf{0}_{NML\times N})$ is characteristic of non-trivial primal solutions only if $\mathbf{R}_{\mathbf{ni}}$ is less than full rank, which mirrors the converse of Theorem 3.7. If this is so, then the solution takes on the "complete clutter nulling" character seen in Chapter 3.3.2, which we will not discuss further here. Thus, in all subsequent analyses, we assume that $\mathbf{R}_{\mathbf{ni}}$ is indeed full-rank.

Along a similar line, we can use the conditions (4.16) and (4.18) to generate an optimal solution for $\tilde{\boldsymbol{\mu}}_M$, which conforms with the lemma and theorem above. If we premultiply (4.18) by $\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M)$ and satisfy (4.16), then we have

$$\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M)\mathbf{B}_3\,\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M) = \boldsymbol{\Sigma}_2^H\mathbf{B}_1\boldsymbol{\Sigma}_2. \tag{4.19}$$

It is clear that any feasible $\mathbf{B}_3$ can be represented as $\mathbf{B}_3 = \rho^2\mathbf{I}_N + \mathbf{Z}_3$, where $\mathbf{Z}_3$ is a Hermitian matrix with an all-zero diagonal. Due to this construct, $\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M)\mathbf{Z}_3\,\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M)$ also has an all-zero diagonal, which implies

$$\rho^2\,\mathrm{diag}(\tilde{\boldsymbol{\mu}}_M)^2 = \mathrm{diag}_M(\boldsymbol{\Sigma}_2^H\mathbf{B}_1\boldsymbol{\Sigma}_2), \tag{4.20}$$

where the operator $\text{diag}_M(\cdot)$ extracts the diagonal elements of the argument matrix. Let $\sqrt{\mathbf{x}}$ denote the elementwise square root of the vector $\mathbf{x}$. Then, the optimal form of $\tilde{\boldsymbol{\mu}}_M$ is

$$\tilde{\boldsymbol{\mu}}_M = \frac{1}{\rho}\sqrt{\text{vec}(\text{diag}_M(\boldsymbol{\Sigma}_2^H \mathbf{B}_1 \boldsymbol{\Sigma}_2))}. \tag{4.21}$$

In turn, this means that the overall rank of the solution is mostly dependent on the interaction of the relaxed beamformer matrix $\mathbf{B}_1$ with the dual clutter-target variable $\boldsymbol{\Sigma}_2$.

4.2.4 Recovering Rank-1 Solutions

Given that the rank of the solution is, at best, upper-bounded by N in full-rank $\mathbf{R_{ni}}$ cases, it is important to discuss methods of recovering an effective rank-one solution from the solution matrix $\mathbf{B}$. As we will soon see in Chapter 4.4 below, there are many situations where the relaxed solution to the CM problem does not admit even a numerically rank-one solution. Recall the approximate sum vector of basis-length K given by $\mathbf{b}_K^a$, and the resulting filter-signal pair $(\mathbf{w}_K^a, \mathbf{s}_K^a) = (\boldsymbol{\Psi}_W \mathbf{b}_K^a, \boldsymbol{\Psi}_S \mathbf{b}_K^a)$. This method's construction was discussed initially in Chapter 3.1.3, and it was shown there to be highly effective if there was a clear dominant eigenvalue. If this pair is not feasible, we can make it so for the CM problem by generating the pair $(\hat{\mathbf{w}}_K^a, \hat{\mathbf{s}}_K^a)$ as follows:

$$\hat{\mathbf{s}}_K^a = \rho e^{\jmath \angle \mathbf{s}_K^a} \tag{4.22}$$

$$\hat{\mathbf{w}}_K^a = \frac{\kappa}{(\mathbf{s}_K^a)^H \mathbf{T}^H \mathbf{w}_K^a} \mathbf{w}_K^a \tag{4.23}$$

where $\angle \mathbf{s}_K^a$ is the vector of phases of the elements of $\mathbf{s}_K^a$ and the exponentiation is elementwise. As we will soon show, if the rank of the solution is greater than one *and* subsequent eigenvalues are close or equal to the "dominant" eigenvalue, then this approximation can significantly degrade the quality of the solution, feasible or not. In these cases, which we term

"rank-inflated" solutions, we can obtain a refined solution by initializing an existing iterative solver (i.e., gradient descent or the alternating minimization of [8]) with either $(\mathbf{w}_K^a, \mathbf{s}_K^a)$ or $(\hat{\mathbf{w}}_K^a, \hat{\mathbf{s}}_K^a)$. We will see below that even one step of refinement can provide significant improvement over the approximate sum when rank-inflation occurs.

Another possibility is randomization. This idea is inspired by a similar approach used for linear semidefinite programs that are relaxations of QCQPs [71], the process of which was indirectly discussed in Chapter 2.3. We showed there that a stochastic version of this QCQP (where objective and constraints are satisfied in expectation) is exactly equivalent to the LSDP, provided the solution to the QCQP is drawn from a zero-mean Gaussian distribution whose covariance is given by the solution to the LSDP. In the case of biquadratic or quartic problems, such an equivalence will not be exact. Consider the constant-modulus problem with one clutter patch; that is,

$$g(\mathbf{w}, \mathbf{s}) = \mathbf{w}^H \mathbf{R}_{\mathbf{ni}} \mathbf{w} + \mathbf{w}^H \mathbf{\Gamma} \mathbf{s}\mathbf{s}^H \mathbf{\Gamma}^H \mathbf{w}, \tag{4.24}$$

or, in the joint variable,

$$\tilde{g}(\mathbf{b}) = \mathbf{b}^H \widetilde{\mathbf{R}}_{\mathbf{ni}} \mathbf{b} + \mathbf{b}^H \tilde{\mathbf{\Gamma}} \mathbf{b}\mathbf{b}^H \tilde{\mathbf{\Gamma}}^H \mathbf{b}. \tag{4.25}$$

Obviously, if $\mathbf{b} \sim \mathbb{C}\mathcal{N}(0, \mathbf{B})$, then the linear first term of the relaxed problem is, of course, equivalent. However, the quartic term requires more nuance. It can be shown that given our random variable assumption

$$\mathbf{E}\{\mathbf{b}^H \tilde{\mathbf{\Gamma}} \mathbf{b}\mathbf{b}^H \tilde{\mathbf{\Gamma}}^H \mathbf{b}\} = \operatorname{tr}(\mathbf{B}\tilde{\mathbf{\Gamma}}\mathbf{B}\tilde{\mathbf{\Gamma}}^H) + |\operatorname{tr}(\mathbf{B}\tilde{\mathbf{\Gamma}})|^2 \tag{4.26}$$

which contains, but is not solely composed of, the relaxed cost $|\operatorname{tr}(\mathbf{B}\tilde{\mathbf{\Gamma}})|^2$. Thus, there is not an exact one-to-one mapping between a stochastic biquadratic program and its QSDP relaxation. However, it stands to reason that since the relaxed cost function is embedded in

the expectation above and our constraint sets are exactly satisfied in expectation, a feasible random vector that minimized this cost might be useful.

In the case of the constant modulus constraint, we follow a version of the procedure described in [47]: first, draw n_{rand} samples from the aforementioned distribution. Next, construct a feasible pair for each sample – create a waveform that follows the phase of each sample, but meeting the constant modulus constraint, and then construct a filter from the sample scaled so that the Capon constraint is met. This is similar to the process mentioned above for the approximate-basis solution. Then, the reported pair from the method is the one which minimizes the original cost function. We will see in Chapter 4.4, however, that this method has severe limitations because of the aforementioned additional cost.

4.3 Similarity-Constrained Joint Design

As mentioned previously, the "similarity" constraint has been typically considered in the literature as a means to mitigate undesirable properties of the waveform in a jointly optimal pair designed under power-only constraints or the constant-modulus constraint considered in the previous section. While we have shown in Chapter 4.1 that it is inappropriate to view this mitigation under the lens of the memetic "ambiguity inheritance" concept, the potential performance and system-limitation interpretations still provide sufficient technical reason to explore such a constraint. In addition, we will demonstrate through the KKT conditions for the relaxed problem that a subspace alignment interpretation of the similarity constraint is warranted.

Let $\mathbf{s}_g$ be the reference waveform under consideration to be modified. Under the similarity/proximal constraint, the waveform constraint set Ω_s in (2.24) is

$$\Omega_s = \left\{ \mathbf{s} : \|\mathbf{s}\|^2 \leq P_o, \|\mathbf{s} - \mathbf{s}_g\|^2 \leq \varepsilon \right\}, \tag{4.27}$$

where P_o is the previously considered power constraint and ε is a constant used to control how similar the designer wishes the new waveform to be to the reference. For simplicity, we set the power of the reference waveform to the overall power constraint, i.e. $\|\mathbf{s}_g\|^2 = P_o$. This means that $\varepsilon \in [0, 2P_o]$.

With this in place, we can write the joint optimization problem under consideration as

$$\begin{aligned} \min_{\mathbf{w},\mathbf{s}} \quad & \mathbf{w}^H \mathbf{R}_{\mathbf{u}}(\mathbf{s}) \mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H \mathbf{T} \mathbf{s} = \kappa \\ & \|\mathbf{s}\|^2 \leq P_o \\ & \|\mathbf{s} - \mathbf{s}_g\|^2 \leq \epsilon. \end{aligned} \tag{4.28}$$

Notice that unlike in a power-equality-constrained problem, there is the potential for the upper bound on ε to vary with $\|\mathbf{s}\|^2$ if the power constraint is not on the boundary. However, we will show that at least the solution to the relaxation will achieve this bound under most physically realizable conditions, and so we shall defer resolving this issue for later research.

The following subsections will describe the method to cast the joint similarity problem in the form of a QSDP (including an additional step previously described in the literature), as well as the KKT conditions of the QSDP and necessary conclusions that can be formed from them.

4.3.1 Relaxation Paths

The problem described in (4.28) does not immediately lend itself to a semidefinite relaxation, but, as shown in [110], there is a means to make the similarity constraint amenable to it. First, note that the similarity constraint is effectively equivalent to $\Re\{\mathbf{s}^H\mathbf{s}_g\} \geq P_o - \varepsilon/2$. Now, consider the following variant of (4.28), which we term the "alignment form":

$$\begin{aligned} \min_{\mathbf{w},\mathbf{s}} \quad & \mathbf{w}^H\mathbf{R}_{\mathrm{u}}(\mathbf{s})\mathbf{w} \\ \text{s.t.} \quad & \mathbf{w}^H\mathbf{T}\mathbf{s} = \kappa \\ & \|\mathbf{s}\|^2 \leq P_o \\ & |\mathbf{s}^H\mathbf{s}_g|^2 \geq \delta \end{aligned} \tag{4.29}$$

where $\delta = (P_o - \varepsilon/2)^2$.

For a fixed $\mathbf{w}$, it has been shown [110, 47] that, given an optimal solution $\hat{\mathbf{s}}$ to a problem similar to (4.29) and the angle $\psi = \arctan(\Im\{\mathbf{s}^H\mathbf{s}_g\}/\Re\{\mathbf{s}^H\mathbf{s}_g\})$, $\mathbf{s} = \hat{\mathbf{s}}e^{\jmath\psi}$ is an optimal solution to similar to (4.28). Since our aim is to jointly optimize $\mathbf{w}, \mathbf{s}$ in one step, we show how to harness the above observation for the joint problem. Let $\breve{\mathbf{w}}, \breve{\mathbf{s}}$ be optimal for (4.29). Then, consider the constructed potential point $\mathbf{s}_? = \breve{\mathbf{s}}e^{\jmath\psi}$, where ψ is now the angle between $\breve{\mathbf{s}}$ and $\mathbf{s}_g$. It is clear that, as in the fixed beamformer problem, $\mathbf{s}_?$ is a feasible point in both waveform constraint sets. To satisfy the Capon constraint, we need to find a potential $\mathbf{w}_?$. By inspection, a feasible choice is $\mathbf{w}_? = \breve{\mathbf{w}}e^{\jmath\psi}$. Finally, since the phases necessarily cancel in the objective function, the objective is the same when evaluated at either pair for both problems. This therefore proves that if $(\breve{\mathbf{w}}, \breve{\mathbf{s}})$ is optimal for (4.29), then $(\mathbf{w}_?, \mathbf{s}_?) = (\breve{\mathbf{w}}e^{\jmath\psi}, \breve{\mathbf{s}}e^{\jmath\psi})$ is optimal for (4.28), and we can continue our relaxation process by considering the alignment form only.

Once again using the selection matrices $\mathbf{\Psi}_S, \mathbf{\Psi}_W$, we can clearly recast the new optimization problem in a relaxed form by moving to the joint variable $\mathbf{b}$ and its relaxation $\mathbf{B}$. The relaxed equivalent constraint set is now

$$\mathcal{T}_{SC}(\Omega_{\mathbf{s}}) = \left\{ \mathbf{B} : \mathbf{s}_g^H \mathbf{\Psi}_S \mathbf{B} \mathbf{\Psi}_S^T \mathbf{s}_g \geq \delta \right\}. \tag{4.30}$$

In vectorized form, this problem is

$$\begin{aligned} \min_{\mathbf{B} \in \mathcal{H}_+^{N(ML+1)}} \quad & \boldsymbol{\beta}_2^H \mathbf{C}_V \boldsymbol{\beta}_2 + \boldsymbol{\beta}_1^H \mathbf{r}_{\mathbf{ni}} \\ \text{s.t.} \quad & \boldsymbol{\beta}_2^H \boldsymbol{\tau} = \kappa \\ & \boldsymbol{\beta}_3^H \operatorname{vec}(\mathbf{I}_N) \leq P_o \\ & \boldsymbol{\beta}_3^H (\mathbf{s}_g^* \otimes \mathbf{s}_g) \geq \delta \end{aligned} . \tag{4.31}$$

If the optimal solution to this problem is rank 1, then we can perform the correction above and obtain an optimal result to (4.28). Otherwise, the decomposition techniques of Chapter 3.1.3 and Chapter 4.2.4 can be used to obtain a feasible but suboptimal pair that is then phase corrected. We will demonstrate below that the KKTs and additional information lead us to believe that most achievable solutions are rank one.

4.3.2 The Slater Condition

We can attack the Slater condition for this class of problem much the same way we have for the original problem and its constant modulus variation. In this case, the relative interior of the constraint set is all positive definite matrices $\mathbf{B}$ where $\operatorname{tr}(\mathbf{B}_3) < P_o$ and $\mathbf{s}_g^H \mathbf{B}_3 \mathbf{s}_g > \delta$ that satisfy the equality constraint.

An achievable version of the Slater condition is given in the theorem below, followed by a simple proof.

Theorem 4.3 (Primal SC Slater's Condition)**.** *The similarity-constrained QSDP* (4.31) *satisfies Slater's condition if both of these inequalities hold:*

$$\|\mathbf{T}\|_F^2 > \frac{|\kappa|^2}{P_o} \tag{4.32}$$

$$\delta < P_o^2 \tag{4.33}$$

Proof: The first inequality is proven in Appendix A, since the effect of the power constraint is identical to the lesser constrained problem (3.24).

The second inequality comes from a simple observation about the trace operation. Recall that for two positive semidefinite matrices $\mathbf{C}, \mathbf{D}$, $\mathrm{tr}(\mathbf{CD}) \leq \mathrm{tr}(\mathbf{C})\,\mathrm{tr}(\mathbf{D})$. Obviously, $\mathbf{s}_g^H \mathbf{B}_3 \mathbf{s}_g = \mathrm{tr}(\mathbf{B}_3 \mathbf{s}_g \mathbf{s}_g^H) \leq P_o\,\mathrm{tr}(\mathbf{B}_3)$. However, since $\mathrm{tr}(\mathbf{B}_3) < P_o$, we have $\mathbf{s}_g^H \mathbf{B}_3 \mathbf{s}_g < P_o^2$. This is possible if and only if $\delta < P_o^2$, and so we have the final necessary aspect to fulfill Slater's condition. ■

Observe that if $\delta = P_o^2$, then not only is there effectively no design problem, the Slater condition can never be satisfied. Hence, we assume that these conditions are satisfied for all practical problems below.

4.3.3 KKT Conditions & Implications

In this subsection, we examine the KKTs of the relaxed problem (4.31) above to analytically determine properties of the relaxed solution, as well as any additional implications for a potential rank-one solution.

From above, let us retain the Lagrange multipliers $\tilde{\mu}_C \in \mathbb{C}$ for the Capon constraint and $\boldsymbol{\Sigma}$ for the semidefiniteness, and define the Lagrange multipliers $\lambda_P, \lambda_S \in \mathbb{R}^+$ for the **p**ower and **s**imiliarity constraints, respectively, in lieu of $\boldsymbol{\lambda}_T$.

By differentiating the Lagrangian of (4.31), we find the optimal slackness variable has the form

$$\boldsymbol{\Sigma} = \begin{bmatrix} \mathbf{R}_{\mathbf{ni}} & \boldsymbol{\Sigma}_2(\tilde{\mu}_C) \\ \boldsymbol{\Sigma}_2(\tilde{\mu}_C)^H & \lambda_P \mathbf{I}_N - \lambda_S \mathbf{s}_g \mathbf{s}_g^H \end{bmatrix} \tag{4.34}$$

where the definition of $\boldsymbol{\Sigma}_2(\tilde{\mu}_C)$ is identical to the constant modulus case. Defining the projection matrix onto $\mathbf{s}_g$ as $\mathbf{P}_{\mathbf{s}_g} = \frac{\mathbf{s}_g \mathbf{s}_g^H}{P_o}$, an equivalent definition of the lower diagonal block of $\boldsymbol{\Sigma}$ is $\boldsymbol{\Sigma}_3 = \lambda_P \mathbf{I}_N - \lambda_S P_o \mathbf{P}_{\mathbf{s}_g}$. Since the similarity constraint effectively enforces a minimum alignment of the transmitted signal basis with the goal signal, the appearance of a rank-one projection is unsurprising.

Due to the inequality constraints, there are two scalar complementary slackness conditions:

$$\lambda_P \operatorname{tr}(\mathbf{B}_3) = \lambda_P P_o, \quad \lambda_P \geq 0 \tag{4.35}$$

$$\lambda_S \mathbf{s}_g^H \mathbf{B}_3 \mathbf{s}_g = \lambda_S \delta, \quad \lambda_S \geq 0 \tag{4.36}$$

As with the power-bounded problem in Chapter 3, a critical task will be to determine when solutions are explicitly power-bounded ($\lambda_P > 0$), as this determines when KKT points of our relaxation coincide with relaxations of equality-constrained problems considered elsewhere. It is also important to determine when a solution is alignment-bounded ($\lambda_S > 0$), since it has important implications for its analytical structure.

In the SC case, the matrix complementarity condition $\mathbf{B}\boldsymbol{\Sigma} = \mathbf{0}_{J\times J}$ devolves, as before, into four conditions: (4.15), (4.17), and the "new" conditions

$$\mathbf{B}_1\boldsymbol{\Sigma}_2 = -\mathbf{B}_2(\lambda_P\mathbf{I}_N - \lambda_S P_o\mathbf{P}_{\mathbf{s}_g}) \tag{4.37}$$

$$\mathbf{B}_2^H\boldsymbol{\Sigma}_2 = -\mathbf{B}_3(\lambda_P\mathbf{I}_N - \lambda_S P_o\mathbf{P}_{\mathbf{s}_g}) \tag{4.38}$$

The semidefiniteness conditions for $\boldsymbol{\Sigma}$ are

$$\mathcal{R}(\boldsymbol{\Sigma}_2^H) \subseteq \mathcal{R}(\lambda_P\mathbf{I}_N - \lambda_S P_o\mathbf{P}_{\mathbf{s}_g}) \tag{4.39}$$

$$\lambda_P\mathbf{I}_N - \lambda_S P_o\mathbf{P}_{\mathbf{s}_g} \succeq \boldsymbol{\Sigma}_2^H\mathbf{R}_{\mathbf{ni}}^{\dagger}\boldsymbol{\Sigma}_2 \tag{4.40}$$

This final condition leads us to our first major theorem, which we present with minimal proof:

Theorem 4.4. *For a similarity-constrained relaxed joint design problem, any non-power-bounded optimal solution will also be non-alignment-bounded; that is,* $\lambda_P = 0 \implies \lambda_S = 0$.

Proof. Assume $\lambda_P = 0$ is part of a feasible, optimal solution. This would require (4.40) to have the sum of two PSD matrices be *negative* semidefinite, which means no feasible solution could also have $\lambda_S > 0$. □

This theorem inspires us to find a more expressive representation of part of the primal solution; namely, $\mathbf{B}_3$. Since $\mathbf{s}_g$ spans a particular subspace of $\mathbb{C}^N$, then we can subdivide $\mathbf{B}_3$ through the projections onto and away from $\mathbf{s}_g$, or

$$\begin{aligned}\mathbf{B}_3 &= (\mathbf{P}_{\mathbf{s}_g} + \mathbf{P}_{\mathbf{s}_g}^{\perp})\mathbf{B}_3(\mathbf{P}_{\mathbf{s}_g} + \mathbf{P}_{\mathbf{s}_g}^{\perp}) \\ &= (\mathbf{B}_3)_{\mathbf{s}_g}^{\perp} + \mathbf{B}_{3,CT} + \alpha\mathbf{s}_g\mathbf{s}_g^H,\end{aligned} \tag{4.41}$$

where $(\mathbf{B}_3)^{\perp}_{\mathbf{s}_g} = \mathbf{P}^{\perp}_{\mathbf{s}_g}\mathbf{B}_3\mathbf{P}^{\perp}_{\mathbf{s}_g}$, $\mathbf{B}_{3,CT} = \mathbf{P}_{\mathbf{s}_g}\mathbf{B}_3\mathbf{P}^{\perp}_{\mathbf{s}_g} + \mathbf{P}^{\perp}_{\mathbf{s}_g}\mathbf{B}_3\mathbf{P}_{\mathbf{s}_g}$ are the crossterms, and α is the fractional alignment of the solution such that $\mathbf{s}_g^H\mathbf{B}_3\mathbf{s}_g = \alpha P_o^2$. Notice that $\mathbf{B}_{3,CT}$ has zero trace, due to the projections.

Given this form, we can find, among other things, a relevant interval for the parameter α, and a specific refinement given the theorem above. First, recall a feasible solution has $\mathrm{tr}(\mathbf{B}_3) \leq P_o$. Direct manipulation of this inequality tells us that

$$\mathrm{tr}((\mathbf{B}_3)^{\perp}_{\mathbf{s}_g}) \leq (1-\alpha)P_o \tag{4.42}$$

which, to be feasible, requires $\alpha \leq 1$. Furthermore, since a feasible solution also requires $\mathbf{s}_g^H\mathbf{B}_3\mathbf{s}_g \geq \delta$, we have $\alpha \geq \frac{\delta}{P_o^2}$. Thus, overall,

$$\alpha \in [\delta_{\mathrm{f}}, 1] \tag{4.43}$$

where we have defined the fractional alignment $\delta_{\mathrm{f}} = \frac{\delta}{P_o^2}$. Notice that this is a valid interval iff Theorem 4.3 is true and the Slater condition is satisfied. If the conditions for Theorem 4.4 hold, then this interval becomes open on both sides.

In fact, this generic form for $\mathbf{B}_3$ results in a simpler form of (4.36) as well. Since $\lambda_S\mathbf{s}_g^H\mathbf{B}_3\mathbf{s}_g = \lambda_S\delta$, then, clearly

$$\lambda_S\left(\alpha - \delta_{\mathrm{f}}\right) = 0. \tag{4.44}$$

This means that in alignment-bounded solutions (i.e. $\lambda_S > 0$), $\alpha_{AB} = \delta_{\mathrm{f}}$. Furthermore, since all alignment-bounded solutions are necessarily power-bounded, $\mathrm{tr}((\mathbf{B}_3)^{\perp}_{\mathbf{s}_g}) = (1-\delta_{\mathrm{f}})P_o$.

Observe that at no time did we involve $\mathbf{\Sigma}_2$ in our assumptions; that said, the PSD requirement for $\mathbf{\Sigma}$ requires $\mathbf{\Sigma}_2 = \mathbf{0}_{NML\times N}$ in a non-power bounded scenario. As in the power-

constrained problem of Chapter 3, a necessary condition for a non-power bounded optimal solution is $\text{rank}(\mathbf{R_{ni}}) < NML$. Continuing from there, we have the following lemma:

Lemma 4.2. *If* $\text{rank}(\mathbf{R_{ni}}) < NML$*, then the tuple* $(\lambda_P, \lambda_S, \mathbf{\Sigma}_2) = (0, 0, \mathbf{0}_{NML\times N})$ *is a feasible set of dual variables.*

The proof of this lemma is implicit in the proof of Theorem 4.4. If we assume instead that $\mathbf{\Sigma}_2 = \mathbf{0}_{NML\times N}$, we again come to the conclusion that $\mathbf{R_{ni}}$ must be singular for this to represent a feasible solution, as before. This time, we look to (4.38) for additional information. Since $\mathbf{B}_3 = \mathbf{0}_{N\times N}$ is not a feasible solution, the only way our assumption is part of a feasible solution is if either a) $\lambda_P \mathbf{I}_N - \lambda_S \mathbf{s}_g \mathbf{s}_g^H = \mathbf{0}_{N\times N}$ or b) $0 < \text{rank}(\lambda_P \mathbf{I}_N - \lambda_S \mathbf{s}_g \mathbf{s}_g^H) < N$. Condition a) only occurs, as established above, if $\lambda_P = \lambda_S = 0$. Thus condition b) implies $\lambda_P > 0$.

The rest of condition b) requires a little more work. We do this by investigating the possible ranks that $\mathbf{\Sigma}_3$ could take, using some intuition about semidefiniteness.

Lemma 4.3. *The rank of* $\mathbf{\Sigma}_3^o$ *is either zero,* $N-1$*, or* N*.*

Proof. The aforementioned rank conditions correspond to $\lambda_P = 0$ (and the consequences of Theorem 4.4); $\lambda_P = \lambda_S P_o > 0$, which makes $\mathbf{\Sigma}_3 = \lambda_P \mathbf{P}_{\mathbf{s}_g}^{\perp}$; and $\lambda_P > \lambda_S P_o$, respectively. □

Now we can proceed to our analysis of Condition b). Recall that Condition b) requires $\text{rank}(\mathbf{\Sigma}_3) < N$ *and* $\lambda_P > 0$, but the only scenario from Lemma 4.3 that satisfies these requirements is $\mathbf{\Sigma}_3 = \lambda_P \mathbf{P}_{\mathbf{s}_g}^{\perp}$, $\lambda_S = \frac{\lambda_P}{P_o}$.

Condition b) also has a few consequences on the solution. Under this assumption, (4.38) becomes

$$\mathbf{\Sigma}_3 \mathbf{B}_3 = \lambda_P \mathbf{P}_{\mathbf{s}_g}^{\perp} \mathbf{B}_3 = \mathbf{0}_{N\times N}. \tag{4.45}$$

Since $\lambda_P > 0$, this is equivalent to saying $\mathcal{R}(\mathbf{B}_3) \subseteq \mathcal{R}(\mathbf{P}_{\mathbf{s}_g})$, or $\mathbf{B}_3 = \alpha \mathbf{s}_g \mathbf{s}_g^H$ for some $\alpha > 0$. In fact, we can find the exact value of α. Since $\lambda_P > 0$, $\mathrm{tr}(\mathbf{B}_3) = \alpha P_o = P_o$, and thus $\alpha = 1$.

However, since $\lambda_S > 0$, we also know that $\mathbf{s}_g^H \mathbf{B}_3 \mathbf{s}_g = \delta$. Hence, a necessary condition for this scenario to be possible is $\delta = P_o^2$. This is a fairly restrictive condition, because it means there is no signal design problem, as we are asking the signal to be exactly aligned to the original signal. Given the above, we have effectively proved the following lemma:

Lemma 4.4. *If* $\mathrm{rank}(\mathbf{R_{ni}}) < NML$*, then* $\mathbf{\Sigma}_2 = \mathbf{0}_{NML \times N}$ *implies one of two mutually exclusive alternatives for a feasible solution. Either:*

1. $\delta = P_o^2$*, and therefore* $\lambda_P > 0, \lambda_S = 0$*, and* $\mathbf{B}_3 = \mathbf{s}_g \mathbf{s}_g^H$*, or,*
2. $\lambda_P = \lambda_S = 0$*, and the optimal solutions are those of the original power-constrained problem* (3.24)*.*

By contrapositive, we also therefore have a result analogous to Theorem 3.7 for the similarity-constrained problem:

Lemma 4.5. *A full rank noise-and-interference covariance matrix implies a power bounded solution; that is,* $\mathrm{rank}(\mathbf{R_{ni}}) = NML \implies \lambda_P > 0$ *and* $\mathbf{\Sigma}_2 \neq \mathbf{0}_{NML \times N}$

This equivalence is unsurprising, since this problem adds only a hyperplane constraint to the original power-constrained problem.

Since we now realize that full rank signal-independent interference again leads to a power-bounded solution, let us focus solely on these solutions for now. First, we can identify an exact value of the alignment parameter α in this case under a certain condition. By combining (4.37) and (4.38), we have $\mathbf{\Sigma}_2^H \mathbf{B}_1 \mathbf{\Sigma}_2 = \mathbf{\Sigma}_3 \mathbf{B}_3 \mathbf{\Sigma}_3$. By premultiplying both sides

with $\mathbf{P}_{\mathbf{s}_g}^{\perp}$ and performing some algebra, we find the optimal α is therefore

$$\alpha = 1 - \frac{\operatorname{tr}(\mathbf{P}_{\mathbf{s}_g}^{\perp} \boldsymbol{\Sigma}_2^H \mathbf{B}_1 \boldsymbol{\Sigma}_2)}{\lambda_P^2 P_o}. \tag{4.46}$$

Clearly, $\alpha = 1$ is optimal (and thus our goal signal is the optimal signal) iff $\mathbf{P}_{\mathbf{s}_g}^{\perp} \boldsymbol{\Sigma}_2^H \mathbf{B}_1 \boldsymbol{\Sigma}_2 = \mathbf{0}_{N \times N}$. Equivalently, $\mathbf{s}_g$ must be an eigenvector of $\boldsymbol{\Sigma}_2^H \mathbf{B}_1 \boldsymbol{\Sigma}_2$ associated with the eigenvalue $\sqrt{P_o}$.

We can additionally find appropriate expressions for λ_P and λ_S in each alignment scenario. If we take the trace of (4.15) and (4.38) and equate them, we have

$$\operatorname{tr}(\mathbf{B}_1 \mathbf{R}_{\mathbf{ni}}) = \lambda_P \operatorname{tr}(\mathbf{B}_3) - \lambda_S \operatorname{tr}(\mathbf{B}_3 \mathbf{s}_g \mathbf{s}_g^H). \tag{4.47}$$

Furthermore, by applying (4.36) and (4.35) and simplifying for λ_P:

$$\lambda_P = \frac{\operatorname{tr}(\mathbf{B}_1 \mathbf{R}_{\mathbf{ni}}) + \lambda_S \delta}{P_o}. \tag{4.48}$$

To wit, then, if λ_S is ever zero, then at the very least, the power constraint multiplier follows the original power-constrained problem. If we substitute this expression into the alignment expression above, we have

$$\alpha = 1 - P_o \frac{\operatorname{tr}(\mathbf{P}_{\mathbf{s}_g}^{\perp} \boldsymbol{\Sigma}_2^H \mathbf{B}_1 \boldsymbol{\Sigma}_2)}{(\operatorname{tr}(\mathbf{B}_1 \mathbf{R}_{\mathbf{ni}}) + \lambda_S \delta)^2}. \tag{4.49}$$

When not alignment-bounded ($\lambda_S = 0$), then we have a "simple" expression for the alignment fraction as a function of the other variables, suitable for a purpose-built solver, like the augmented Lagrangian methods employed by [56]. When the solution *is* alignment-bounded,

we can directly solve for λ_S, since here $\alpha = \delta_{\mathrm{f}}$. After some tedious algebra, we find

$$\lambda_S = \delta^{-1}\left(\sqrt{\frac{P_o}{1-\delta_{\mathrm{f}}}\,\mathrm{tr}(\mathbf{P}_{\mathbf{s}_g}^{\perp}\boldsymbol{\Sigma}_2^H\mathbf{B}_1\boldsymbol{\Sigma}_2)} - \mathrm{tr}(\mathbf{B}_1\mathbf{R}_{\mathbf{ni}})\right). \tag{4.50}$$

Therefore, in alignment-bounded cases, λ_P is given by

$$\lambda_P = \sqrt{\frac{\mathrm{tr}(\mathbf{P}_{\mathbf{s}_g}^{\perp}\boldsymbol{\Sigma}_2^H\mathbf{B}_1\boldsymbol{\Sigma}_2)}{P_o(1-\delta_{\mathrm{f}})}}. \tag{4.51}$$

We conclude this section with a theorem that encapsulates most of the derivations so far, and directly connects the power-constrained problem with the similarity-constrained one:

Theorem 4.5. *Let $\mathbf{B}_{SC}^o$ be an optimal solution to* (4.31)*, and $\mathbf{B}_{PC}^o$ be an optimal solution to the power-constrained problem* (3.24)*. Then, when the solution is non-alignment-bounded ($\lambda_S = 0$), $\mathbf{B}_{SC}^o = \mathbf{B}_{PC}^o$, i.e. the similarity-constrained solution is identical to the power-constrained solution. Furthermore, this equivalence occurs for all δ such that*

$$\mathbf{s}_g^H\boldsymbol{\Psi}_S\mathbf{B}_{PC}^o\boldsymbol{\Psi}_S^T\mathbf{s}_g \geq \delta,$$

that is, $\mathbf{B}_{PC}^o$ is an alignment-bounded solution only when $\mathbf{s}_g^H\boldsymbol{\Psi}_S\mathbf{B}_{PC}^o\boldsymbol{\Psi}_S^T\mathbf{s}_g$ meets the bound. Equivalently, this is true when $\epsilon \geq 2\left(P_o - \sqrt{\mathbf{s}_g^H\boldsymbol{\Psi}_S\mathbf{B}_{PC}^o\boldsymbol{\Psi}_S^T\mathbf{s}_g}\right)$.

Proof: There are two major cases to consider in this theorem – non-power-bounded ($\lambda_P = 0$) and power-bounded ($\lambda_P > 0$) – but they are both motivated by the same principle. Namely, if the alignment constraint is inactive, the KKTs obtained are exactly identical to those of the power-constrained problem of (3.24). Since this is a convex problem, KKTs are necessary and sufficient conditions for any solution – hence, identical KKTs imply equivalent solutions in this case.

The second element of the theorem is a natural consequence of our assumption that the solution is not alignment-bounded, and therefore the alignment constraint must be inactive. Since we have established above that $\mathbf{B}_{PC}^{o}$ is a solution, then the constraint is inactive when $\mathbf{s}_g^H \mathbf{\Psi}_S \mathbf{B}_{PC}^{o} \mathbf{\Psi}_S^T \mathbf{s}_g > \delta$. The alignment-bounded equivalence is trivial, since if $\mathbf{B}_{PC}^{o}$ is a solution and $\mathbf{s}_g^H \mathbf{\Psi}_S \mathbf{B}_{PC}^{o} \mathbf{\Psi}_S^T \mathbf{s}_g = \delta$, then $\mathbf{B}_{PC}^{o}$ is therefore alignment-bounded. ■

4.4 Numerical Analysis

In this section, we present numerical results to verify the KKT assertions made above and to compare the performance of our proposed methods with the existing state of the art. We subdivide this section as above, with an analysis of the constant-modulus problem followed by one of the similarity-constrained problem.

In many of the simulations presented here, we consider a common set of parameters for ease of comparison. Specifically, unless otherwise specified, the radar under consideration transmits $L = 8$ pulses of $N = 8$ samples each, and receives the scattered energy with a uniform linear array of $M = 5$ antennas at a half-wavelength per-element spacing. Each pulse is modulated to a carrier frequency $f_o = 1$ GHz and the time between each pulse is $T_p = 50\mu$s. If there is a target under consideration, it is located at the azimuth-elevation pair $(\theta, \phi) = (0.3, 0.3)$, and moves with a relative Doppler of -0.2545 relative to the platform if the Doppler is not otherwise mentioned. Clutter patches are modeled as above, each at an elevation of 0.3 radians and with an azimuth uniformly spaced across the interval $[-\frac{\pi}{2}, \frac{\pi}{2}]$ dictated by the number of clutter patches. "Light" clutter denotes a scenario with $Q_c = 25$ clutter patches, while "heavy" clutter denotes $Q_c = 180$ patches. In any case, the rank of the clutter covariance in the $(\mathbf{w}, \mathbf{s})$ product space is less than $ML = 40$, due to the structure.

The simulations presented below, and our methods in particular, are again enabled by the convex optimization toolbox CVX [87, 86] and its interfaces to the solvers SDPT3 [89, 90] and SeDuMi [91]. In each implementation of our algorithms, the performance shown is from the solver that performed the "best" in a given metric.

4.4.1 Constant-modulus Evaluations

We begin by analyzing the sensitivity of the rank to the modulus parameter in a heavy clutter scenario, set $\kappa = 1$, and consider a noise-only Toeplitz matrix for $\mathbf{R_{ni}}$. In Figure 4.1, we can see that for low values of ρ, we obtain rank-N solutions with one dominant eigenvalue (as with the power-constrained problem). However, as ρ increases, the spectrum "fills out," resulting in rank-N solutions with nearly equal non-zero eigenvalues. This behavior, when it happens, can have a significant effect on the ability to recover quality approximate solutions, as we will soon see below. We also note that if the noise is white, then the high-ρ rank scenario will play out even at low modulus values, but this is not detrimental as the approximate sum mentioned above produces feasible and near-optimal solutions!

Next, we consider the performance of our methods across variations in target Doppler, first in comparison with each other, and then with other algorithms in the literature. We begin by comparing the performance of our relaxed constant modulus method under the two naive rank-one recovery strategies: the SVD-based approximate sum (henceforth, RBQP-CM-SVD) and the randomization-based procedure (henceforth, RBQP-CM-Rand) where we evaluate $n_{rand} = 1000$ random vectors. For RBQP-CM-SVD, we choose the maximum eigenvalue vector, and hence $K = 1$. We also consider the SINR predicted by the optimal value obtained from the solution of the QSDP, to see how much the approximation schemes deviate from the optimal point. Our findings are presented in Figure 4.2. It is immediately clear that the randomization technique RBQP-CM-Rand produces beyond suboptimal re-

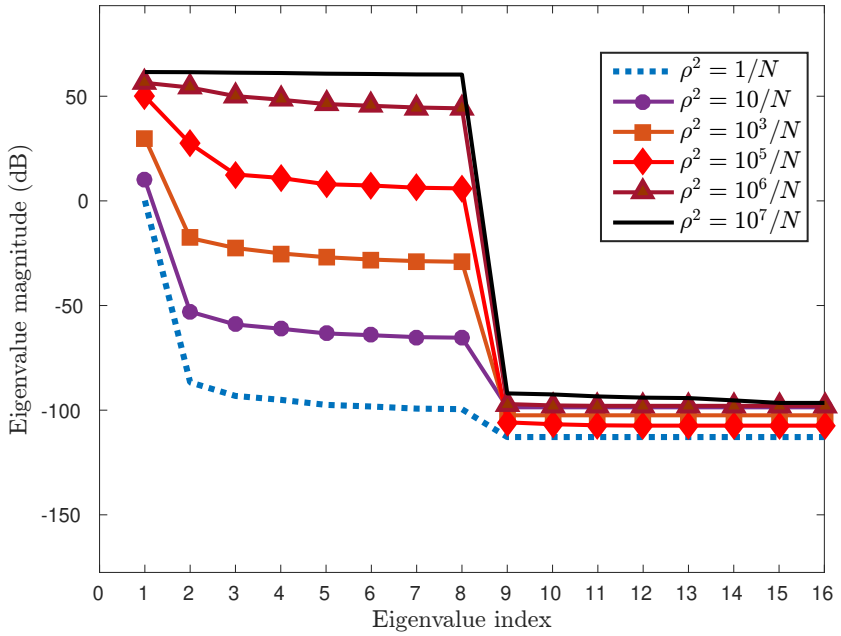

Figure 4.1: First $2N$ eigenvalues of relaxed con.-mod. solution $\mathbf{B}_{CM}$ as a function of ρ.

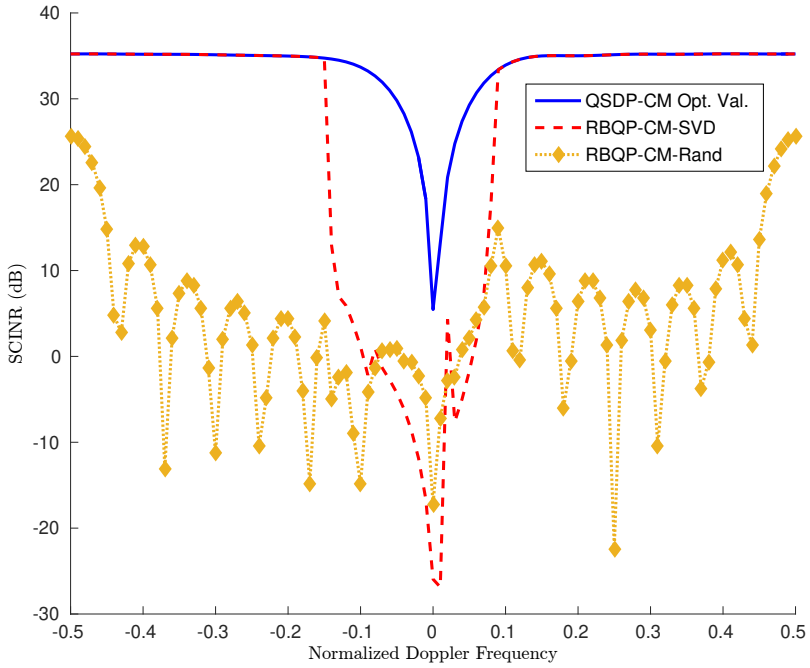

Figure 4.2: Comparison in SINR w.r.t. normalized Doppler of SVD- and randomization-based rank-one recovery techniques for QSDP-based con.-mod. algorithms

sults, even at high Doppler frequencies where the target should be well-separated from the clutter. This is due to the aforementioned mismatch between the expectation in (4.26) and the cost function in (3.22). Additionally, increasing n_{rand} further (we considered 5000 and 10000 random points, as well) provides no noticeable gain, pointing again to the cost-function mismatch as the culprit. It may be possible to refine these randomized points further, similar to what we will investigate below, but given the difficulties at higher Dopplers, we believe this is a mostly fruitless endeavor.

In contrast, at high Doppler frequencies, the naive recovery technique RBQP-CM-SVD performs well, and even generally approximates the power-bounded performance. However, it is clear that, for a range of lower Dopplers, this performance suffers significantly and does not come anywhere close to the expected bound. This is due to the aforementioned rank inflation – at these Dopplers, $\text{rank}(\mathbf{B}_{CM}^{opt}) \approx 2$ with each direction having approximately equal strength, and hence the naive approximation is insufficient, even if made feasible. Additional simulations have shown that the value of K (i.e., how many eigenvectors of $\mathbf{B}$ we add) does not change the overall performance in the low-Doppler regime much, which indicates that additional directions do not add further target matching or clutter/noise suppression – they often merely define feasibility in the constraint set.

Since RBQP-CM-SVD showed significant promise at higher Dopplers, we now examine the impact of refining those approximate solutions in some fashion. In particular, we consider a feasible refinement of RBQP-CM-SVD solutions by using them to initialize a signal-only gradient descent of (4.9) enabled by MATLAB's `fminunc` (henceforth, RBQP-CM-FR), and then obtain a filter from (2.36a). The cost function used in the gradient descent is similar to [36, Eq. (18)], with appropriate modifications similar to that in [34, Appendix B]. As seen in Figure 4.5, under feasible refinement, the gain lost by the initial approximation is immediately recovered for very little computational penalty (indeed, the gradient descent never needs more than 5 iterations), and our overall performance exactly matches the expected

gain. Of course, the procedure could be reversed, with a filter-only gradient descent and

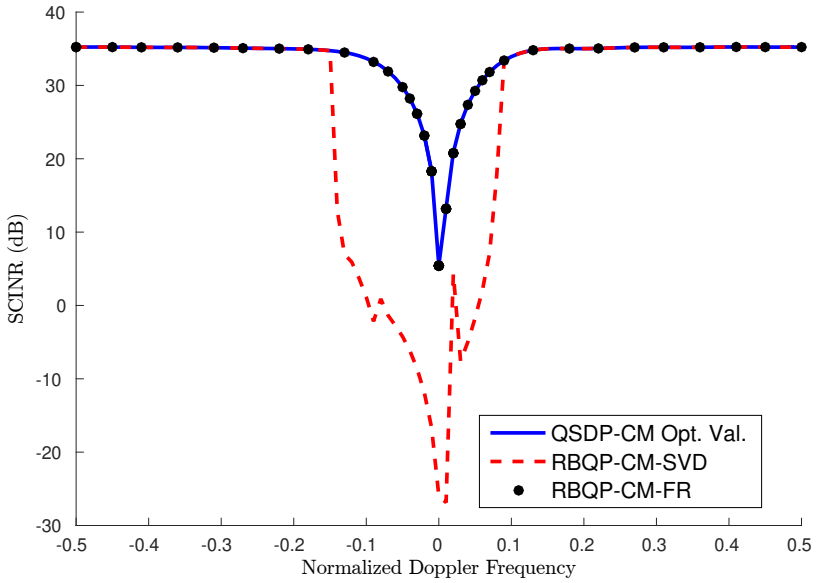

Figure 4.3: Comparison in SINR w.r.t. normalized Doppler of SVD- and feasible refinement (FR)-based rank-one recovery techniques from QSDP-based con.-mod. algorithms

obtaining a signal via (2.36b). We note that the overall variation in the refined filter-signal pairs mostly comes from the filter, with the largest variation occurring at zero Doppler – this can be attributed to the magnitudes of the two dominant eigenvalues getting closer as the target moves into the clutter ridge.

The performance hit experienced by RBQP-CM-SVD can be somewhat predicted by comparison of the optimal values of the QSDPs for the constant modulus problem (RBQP-CM) and that of the power-constrained problem (RBQP-PC) from [66]. This result is documented in Figure 4.4, which depicts the log ratio of RBQP-CM's optimal value to RBQP-PC's optimal value. Obviously, there is some loss at all Dopplers between the more constrained and less constrained problem, but never more than about 0.65 dB. However, what is noticeable are significant changes around the points where performance degradation of the SVD approach appeared in Figure 4.3. This would indicate that the target-clutter environment is producing some sort of significant change to the structure of the constant-modulus solution that

necessitates a higher rank. We leave speculation as to the exact mechanics of this behavior for future research.

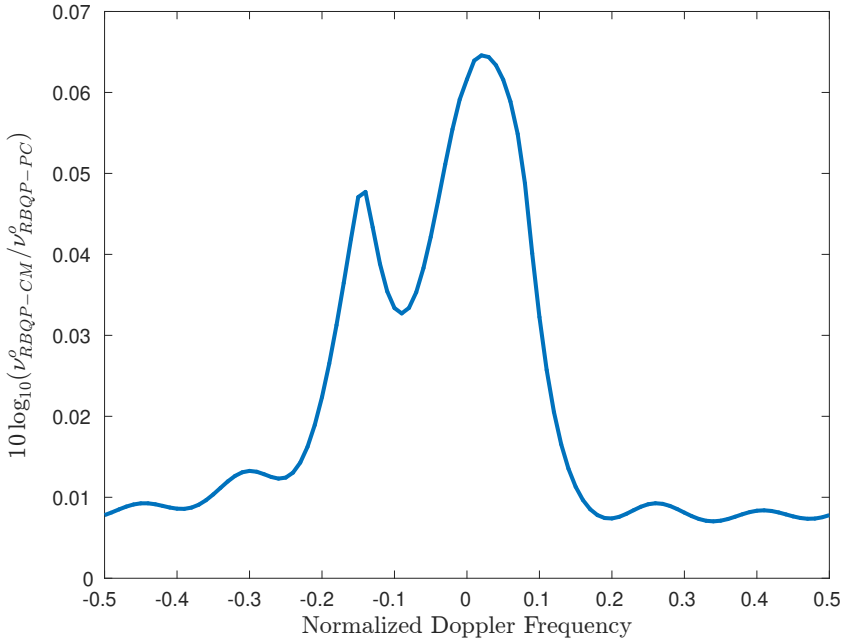

Figure 4.4: Ratio of RBQP-CM to RBQP-PC objective values vs. Doppler

Next, we consider comparison of the superior RBQP algorithm, RBQP-CM-FR, with existing algorithms from the literature. The algorithms considered are Algorithms 2 and 3 of [47] (denoted T&T-A2 and T&T-A3, respectively), our previous technique from [66] (denoted RBQP-PC-SVD), and the constant modulus alternating minimization of [8] (denoted CMAM), each initialized by feasible signals of random phase. Algorithms T&T-A2, T&T-A3, and CMAM all sequentially optimize $\mathbf{w}, \mathbf{s}$ pairs in alternating steps and choose the minimum eigenvector of a given matrix for their $\mathbf{w}$-step. What distinguishes them is the $\mathbf{s}$-step: T&T-A2 solves a linear SDP similar to the relaxation and randomization process we discussed above, T&T-A3 uses the power method-like iterations of [48] to solve a fractional program, and CMAM forms a Capon-like "temporal" beamformer. As seen in Figure 4.5, the feasible refinement of the RBQP solution outperforms every other technique by at least 2 dB across most Doppler values, though TT-A2 approaches RBQP-CM-FR at certain

points. Finer detail of low Doppler performance can be seen in Figure 4.6, which shows the continued improved performance of RBQP-CM-FR over all others.

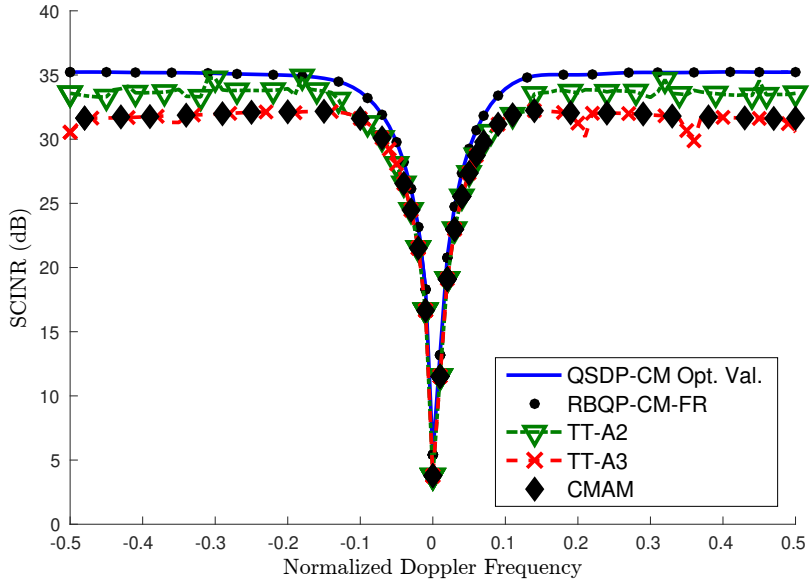

Figure 4.5: Constant modulus algorithm variation in SINR w.r.t. normalized Doppler

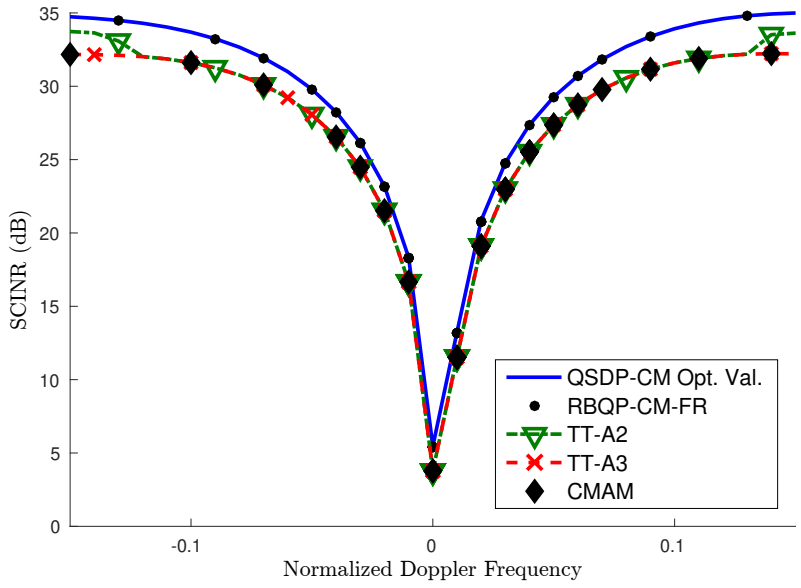

Figure 4.6: Constant modulus algorithm variation in SINR w.r.t. normalized Doppler in the interval [-0.1, 0.1]

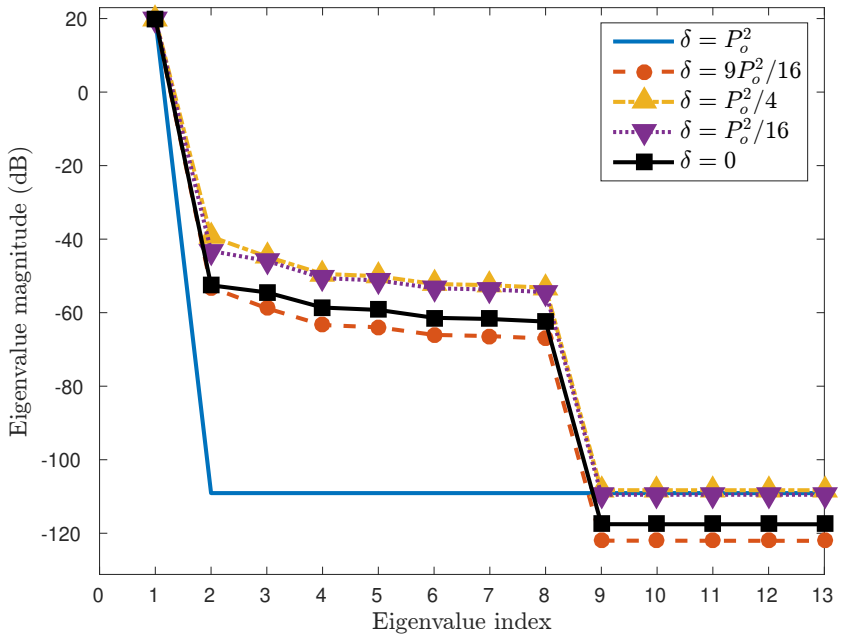

Figure 4.7: First $N + 5$ eigenvalues of relaxed sim.-con. solution $\mathbf{B}_{SC}$ as a function of similarity parameter δ.

4.4.2 Similarity-constrained Evaluations

Our analysis continues with the similarity constrained problem. Again, we begin with an examination of the expected rank of the solution matrix $\mathbf{B}$, to deterimine the viability of a simple decomposition for obtaining approximate solutions. In Figure 4.7, we examine the variation in rank over the alignment parameter δ, for a light clutter scenario with $P_o = \kappa = 1$. Clearly, when $\delta = P_o^2$, we must have a rank-one solution. For all other δ, we have what appear to be nominally rank-N, but effectively rank-1 scenarios, as the dominant eigenvalue far exceeds the other $N - 1$ nonzero values. Furthermore, it seems that as $\delta \to 0$, the overall spectrum lifts and then recedes to some degree. This does not change the practical result of the relaxation or its approximate decomposition, however.

Since we now know that we can obtain representative approximate rank-1 solutions, we can examine, for example, how these solutions satisfy the constraints. Typically, these solutions

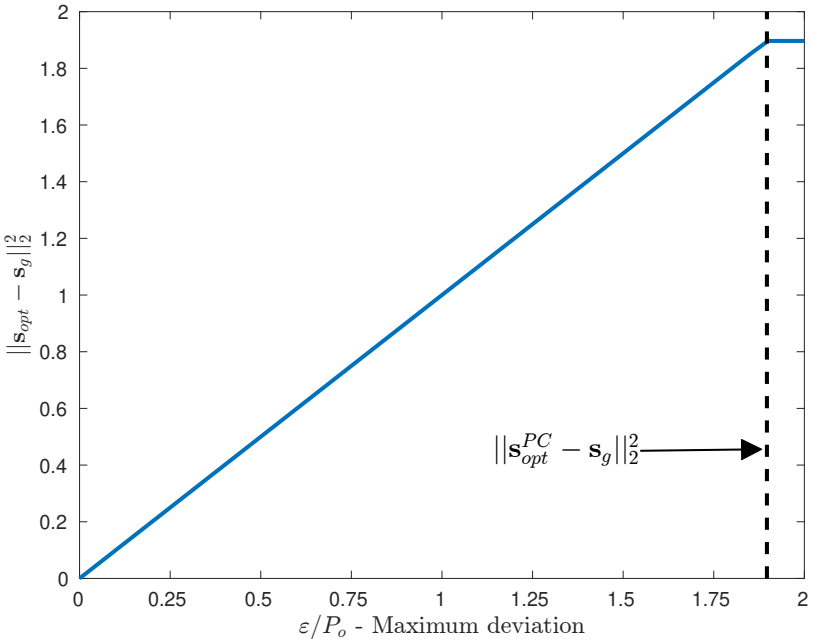

Figure 4.8: Obtained vs. expected variation of optimal $\mathbf{s}$ from $\mathbf{s}_g$.

exactly satisfy the power and Capon constraints regardless of the choice of the similarity parameter ε/alignment parameter δ. However, depending on the choice of ε, the similarity of the approximate waveform solution $\mathbf{s}_{opt}$ to the reference waveform can vary significantly. This is demonstrated in Figure 4.8, which shows that for most values of ε, the evaluated difference $\|\mathbf{s}_{opt} - \mathbf{s}_g\|^2 \approx \varepsilon$. As ε approaches its maximum value, however, $\|\mathbf{s}_{opt} - \mathbf{s}_g\|^2$ saturates to some maximum difference which, as predicted in Theorem 4.5 above, is when the solution is exactly identical to the power-bounded solution. This can be regarded in the context of the alignment formulation, in that any further deviations from $\mathbf{s}_g$ would lead to worse alignment with either the target or the clutter responses.

Next, we consider the comparative performance of our joint approximate solution to other similar algorithms in an operationally relevant context for STAP, Doppler variation. This result is obtained for $P_o = \kappa = 1$ in a heavy clutter scenario, where $\epsilon = 0.1$. For comparison, we also examined Algorithms 4 and 5 of [47] (denoted T&T-A4 and T&T-A5, respectively) and Algorithm 2 of [44] (denoted AA2). T&T-A4 and -A5 follow their constant modulus

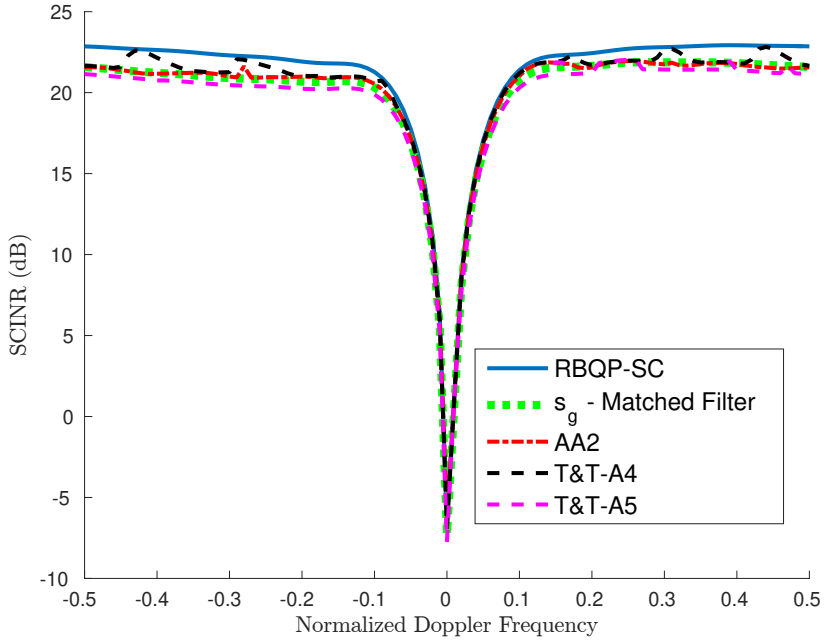

Figure 4.9: Similarity constraint algorithm variation in SINR w.r.t. normalized Doppler

counterparts in design intent, with A4 using a SDP relaxation in the **s**-step, and A5 relying on the power method-like iterations. AA2 also uses an SDP relaxation in the **s**-step, but enforces the similarity constraint in a more complex fashion than T&T-A4. The algorithm we explored in Chapter 4.3 is termed RBQP-SC. As can be seen in Figure 4.9, our method significantly outperforms the known-good waveform & its matched filter, AA2, and T&T-A5. This is true in both high Doppler and low Doppler instances, as RBQP-SC has a much larger minimum detectable velocity at any threshold. Additionally, while there are cases where [47]'s Algorithm 4 approaches RBQP-SC's performance, this is evidence that our relaxation does, in fact, act as a limiting asymptote of performance for any possible algorithm, as we have proposed in [74].

Finally, we conclude our analysis by directly demonstrating the proposition made in Chapter 4.1 in regards to ambiguity functions. Figures 4.10, 4.11, and 4.12 show the ambiguity function of the reference chirp function, the zero-Doppler solution waveform of RBQP-SC, and the zero-Doppler solution waveform of T&T-A4, respectively. These solutions are obtai-

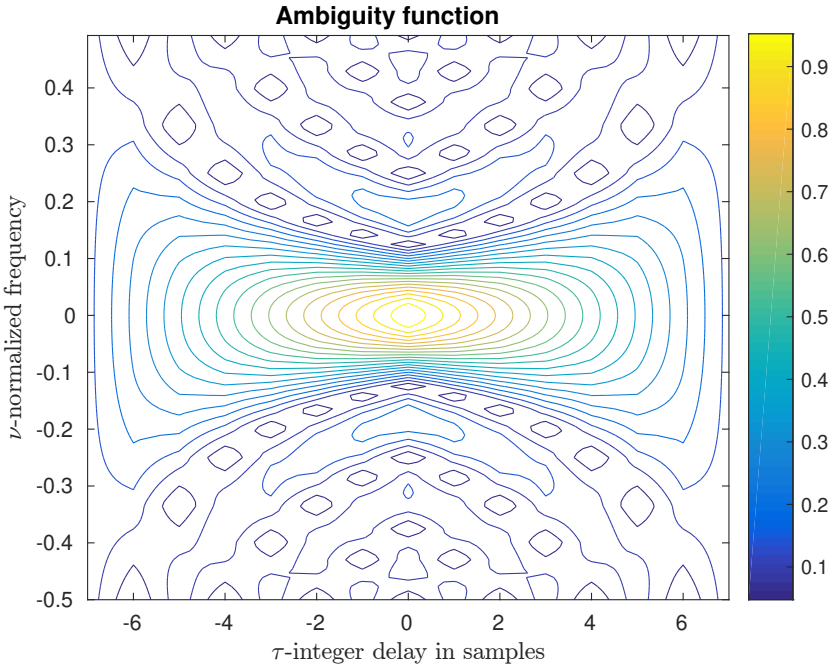

Figure 4.10: Ambiguity function of $\mathbf{s}_g$

ned from the simulation above, where performance at this particular Doppler is, in general, equally poor for all algorithms. Clearly, despite "only" varying from the reference waveform by $\varepsilon = 0.1P_o$, the solution ambiguity functions are significantly different from both the reference and each other. This is to be expected, since two-norm similarity, as theorized above, does not guarantee ambiguity inheritance.

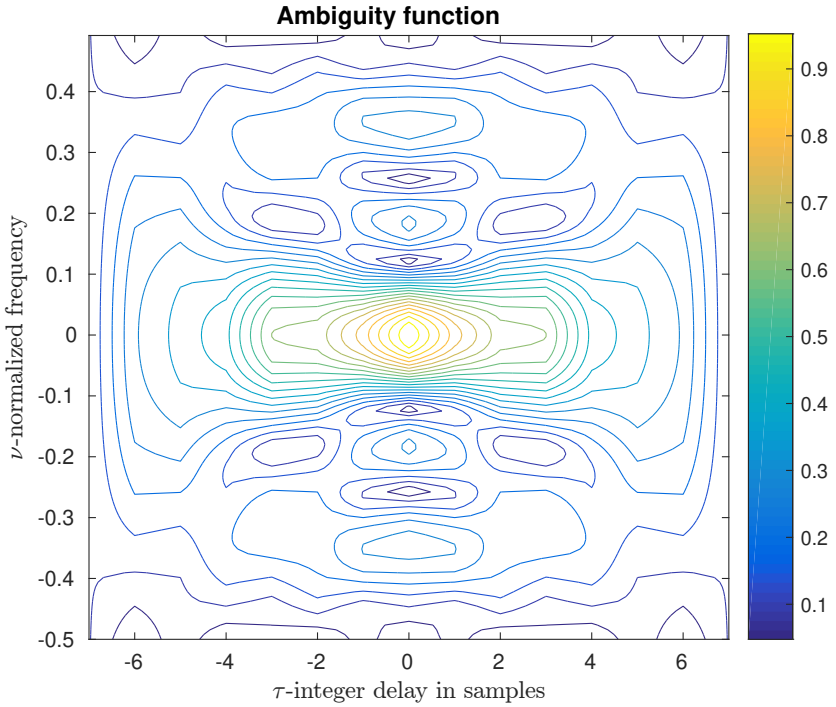

Figure 4.11: Ambiguity function of optimal $\mathbf{s}$ from RBQP-SC at zero Doppler frequency

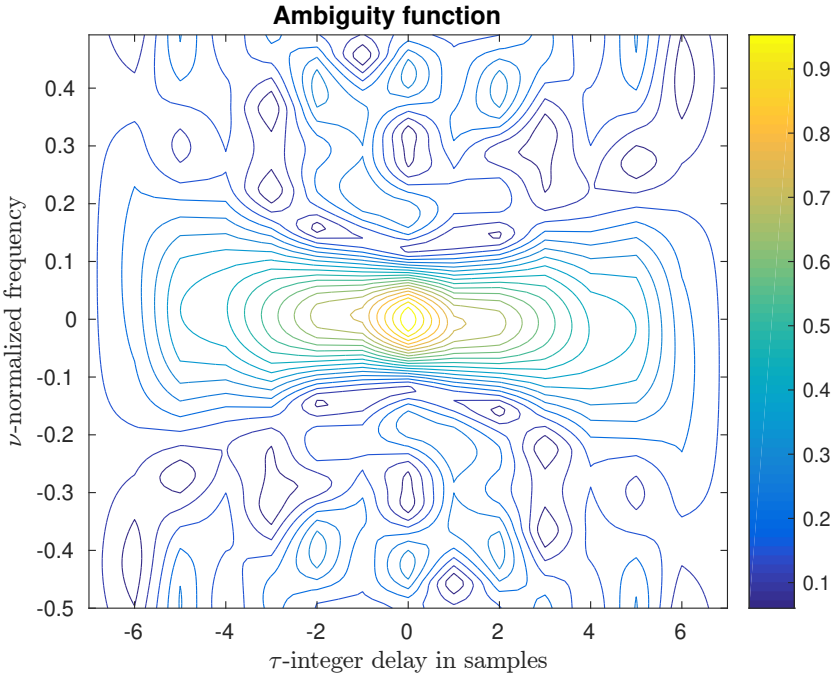

Figure 4.12: Ambiguity function of optimal $\mathbf{s}$ from T&T-A4 at zero Doppler frequency

Chapter 5

Applications to Other Sensing Modalities

So far, our work in the previous chapters has been primarily motivated by the waveform adaptive SIMO STAP application advanced by Setlur and Rangaswamy [8]. However, the majority of our findings are not restricted to this application alone. Indeed, the QSDP techniques are applicable to *any* active sensing problem with a plurality of (nominally orthogonal) inputs and outputs, as well as some known or estimated characterization of a desired target response, non-target input-dependent responses, and input-independent disturbances.

Many of these scenarios are, as the title of this dissertation implies, characteristic of radar systems. In this chapter, we shall demonstrate the applicability of our technique to alternative channel and system models. A generalized overview of the channel concept will be provided in Chapter 5.1. In Chapter 5.2, we will consider a challenging, if not necessarily realistic, alternative model: the simplified reverberant channel model of Stoica, He, and Li [36], which is nominally intended for sonar systems, but also describes a radar channel where the primary clutter returns considered are from range bins other than that of the target.

We will conclude with comparative simulations between our approach and that of [36] in Chapter 5.3.

5.1 Channel Models for Other Modalities

As noted in Chapter 2.1.2, Guerci and others [63, 3] have advocated for a matrix transfer function approach to modeling active sensing systems. Assuming a sensing system with M_T transmit degrees of freedom and M_R receive degrees of freedom, the model in (2.18) describes a vector of received complex amplitudes $\mathbf{y} \in \mathbb{C}^{M_R}$ in terms of the transmit resources $\mathbf{s} \in \mathbb{C}^{M_T}$ and channel matrices $\mathbf{H}_s \in \mathbb{C}^{M_R \times M_T}$ which act on those resources. The (i, j)th element of $\mathbf{H}_s$ describes the coupling between the jth transmit resource and the ith receive resource, induced by the resolvable scatterer(s) and described by transmit-independent stochastic Green's functions. Since the examples we consider below follow the illustrative single-target scenario in the previous chapters, we will repurpose the notation for the target channel matrix $\mathbf{T}$ and the clutter channel matrix $\mathbf{\Gamma}_q$ seen in (2.19) for the different models we will subsequently consider.

In order to resolve or appropriately process the return $\mathbf{y}$, the number of receive resources typically equals or exceeds that of the transmit resources; that is, $M_R \geq M_T$. For the WA-STAP scenario, this was absolutely satisfied, since the filter $\mathbf{w} \in \mathbb{C}^{NML}$ processes the fast time samples as well as the slow-time and spatial samples. In the scenario we consider in Chapter 5.2 below, however, the purely fast-time nature of the channel forces $M_T = M_R = N$. This may lead to identifiability concerns, which we will see below.

5.2 Reverberant Channel Joint Design

In this section, we consider the reverberant channel model of Stoica, He, and Li [36], a simpler SISO radar/sonar model that accounts only for the fast-time response of the channel and ignores Doppler shifts. We will demonstrate the connection to our joint optimization approach and compare the relaxed solution to the trio of techniques considered in [36]

5.2.1 The Reverberant Channel & Mismatched Filters

We begin by describing the reverberant channel, which can be understood as the simplest discrete-time instantiation of a signal-dependent clutter model. For the most part, we follow the details and notation of [36] and references therein, but where necessary, we retain our own notation for consistency. Let $M_T = M_R = N$. Then, the transmit resource vector $\mathbf{s} \in \mathbb{C}^N$ is a vector of N complex weights that modulate a series of subpulses transmitted into the scattering environment. A basic power constraint is enforced on this signal; namely, $\|\mathbf{s}\|^2 = N$. The authors in [36] consider two additional constraints on $\mathbf{s}$ previously discussed in Chapter 4.1, constant modulus and PAPR. For simplicity of future comparison, we will only consider the constant modulus constraint $|\mathbf{s}_i| = 1$, which will permit us to employ the techniques seen in Chapter 4.2.

When aligned to the target range delay of interest, the received vector $\mathbf{y} \in \mathbb{C}^N$ is

$$\mathbf{y} = \alpha_0 \mathbf{s} + \sum_{\substack{i=-(N-1) \\ i \neq 0}}^{N-1} \alpha_i \mathbf{J}_i \mathbf{s} + \mathbf{y}_{\mathtt{ni}} \tag{5.1}$$

where α_0 is the target reflectivity in the given bin, the other α_i are the reflectivities of the scatterers in the adjacent range bins, and $\mathbf{y}_{\mathtt{ni}}$ is a zero-mean random vector of signal-independent disturbances, including noise and jamming, with a covariance of $\mathbf{R}_{\mathtt{ni}}$. Due

to the sampling structure inherent in this channel model, the authors in [36] assume $\mathbf{R_{ni}}$ is Toeplitz, which will come into play later. Here, the clutter channel matrices are the $2(N-1)$ shift matrices $\mathbf{J}_i$, which shift the elements of $\mathbf{s}$ down for positive values of i and up for negative values of i, replacing the shifted elements with zeros as necessary. Furthermore, transposition of the matrix reverses the sign of the shift index; that is, $\mathbf{J}_{-i} = \mathbf{J}_i^T$. The target channel matrix, clearly, is the identity matrix $\mathbf{I}_N$.

We assume that the clutter reflectivities α_i are statistically independent of each other, the target reflectivity α_0 and the signal-independent disturbance $\mathbf{y}_{\mathtt{ni}}$; and that they have equal variance $\sigma_{\mathtt{c}}^2$. As with the WA-STAP problem, we assume that $\sigma_{\mathtt{c}}^2$ and $\mathbf{R_{ni}}$ are known *a priori*.

Unlike the detection problem considered for WA-STAP, the primary goal in this problem is to obtain a minimum mean square error (MSE) estimate of the target reflectivity α_0. Assuming that $\mathbf{y}$ is linearly processed by a filter $\mathbf{w} \in \mathbb{C}^N$, the estimate $\hat{\alpha}_0$ is given by

$$\hat{\alpha}_0 = \frac{\mathbf{w}^H \mathbf{y}}{\mathbf{w}^H \mathbf{s}}. \tag{5.2}$$

If the filter used is the signal itself, $\mathbf{w} = \mathbf{s}$, we have a matched filter estimate. However, this choice of filter is only optimal if there is no clutter/reverberation, since the signal (obviously) correlates with itself under time shifts, which inflates the error. If any other filter $\mathbf{w}$ is used, it is what is known as a mismatched filter (MMF), which constitutes additional degrees of freedom to cancel the clutter, interference, and noise.

If we wish to find an optimal pair $(\mathbf{w}, \mathbf{s})$, then we must define the cost function of interest. The MSE of the estimate $\hat{\alpha}_0$ is given by

$$\mathrm{MSE}(\hat{\alpha}_0) = \mathop{\mathbf{E}}_{\alpha_0} |\hat{\alpha}_0 - \alpha_0| = \frac{\mathbf{w}^H \mathbf{R}_{\mathtt{u}}(\mathbf{s}) \mathbf{w}}{|\mathbf{w}^H \mathbf{s}|^2}, \tag{5.3}$$

where $\mathbf{R}_{\mathsf{u}}(\mathbf{s})$ is the signal-dependent covariance defined as

$$\mathbf{R}_{\mathsf{u}}(\mathbf{s}) = \mathbf{R}_{\mathbf{ni}} + \sigma_{\mathsf{c}}^2 \sum_{\substack{i=-(N-1) \\ i\neq 0}}^{N-1} \mathbf{J}_i \mathbf{s}\mathbf{s}^H \mathbf{J}_i^T. \tag{5.4}$$

This is identical to our previous definitions of $\mathbf{R}_{\mathsf{u}}(\mathbf{s})$, just with the clutter channel substitution noted above. Notice that this is the reciprocal of the SINR form (2.23), and clearly, minimizing the MSE is equivalent to maximizing the SINR. With the cost function defined, we can now describe the general optimization problem we consider in this section:

$$\begin{aligned} \min_{\mathbf{w},\mathbf{s}} \quad & \frac{\mathbf{w}^H \mathbf{R}_{\mathsf{u}}(\mathbf{s}) \mathbf{w}}{|\mathbf{w}^H \mathbf{s}|^2} \\ \text{s.t.} \quad & \|\mathbf{s}\|^2 = N \\ & |\mathbf{s}_k| = 1,\ k = 1, \ldots, N. \end{aligned} \tag{5.5}$$

5.2.2 The CREW Approaches

With the optimization problem now defined, we describe the three general techniques advanced by [36] to find an optimal $(\mathbf{w}, \mathbf{s})$ pair. These techniques, which they denote as Cognitive Receiver and Waveform design (CREW), are examples of the signal-filter AM paradigm described in Chapter 2.3.1. Two of these techniques use a form similar to EigenAM's – namely, (2.36a) – in the $\mathbf{w}$-step to generate the mismatched filter, while the other relies on the matched filter and a different design of $\mathbf{s}$ to minimize the undesired responses. We describe the general contours of their design below and refer the interested reader to [36] for further details of their implementation.

The first algorithm, CREW(gra), starts by assuming the minimum eigenvector in the **w**-step, producing the filter

$$\mathbf{w} = \mathbf{R}_{\mathtt{u}}^{-1}(\mathbf{s}). \tag{5.6}$$

Then, the MSE at the **s**-step is MSE $= (\mathbf{s}^H \mathbf{R}_{\mathtt{u}}^{-1}(\mathbf{s})\mathbf{s})^{-1}$. Minimizing this is equivalent to maximizing the denominator, and hence we have the optimization problem

$$\begin{aligned} \max_{\mathbf{s}} \quad & \mathbf{s}^H \mathbf{R}_{\mathtt{u}}^{-1}(\mathbf{s})\mathbf{s} \\ \text{s.t.} \quad & \|\mathbf{s}\|^2 = N \\ & |\mathbf{s}_k| = 1,\ k = 1, \ldots, N. \end{aligned} \tag{5.7}$$

Any feasible signal will have the form $\mathbf{s} = \exp(\jmath\boldsymbol{\phi})$, where $\boldsymbol{\phi} \in \mathbb{R}^N$ is a vector of phases. Hence, they obtain an unconstrained maximization problem, which they solve via **gra**dient descent, giving the algorithm its name. To improve performance, they manually supply the optimization technique (MATLAB's `fminunc`) with the phase-dependent gradients.

The next technique, CREW(fre), also uses the **w**-step form $\mathbf{w} = \mathbf{R}_{\mathtt{u}}^{-1}\mathbf{s}$, but it does not come into play until the signal portion of the algorithm has converged. The signal portion takes advantage of the assumed Toeplitz structure of $\mathbf{R}_{\mathbf{ni}}$ to identify the optimal spectral allocation of the signal **s**, and then attempts to synthesize a signal with that spectrum by sequentially updating an initialized sequence of signals. The frequency allocation process is performed with a bisection algorithm and is waterfilling-like, preferentially placing power in bands where the signal-independent interference is low. The frequency bands used are determined by the $2N-1$ point discrete Fourier transform of the distinct elements (first row and column) of $\mathbf{R}_{\mathbf{ni}}$.

Finally, the last technique, CREW(mat), is similar to CREW(fre), but assumes that the matched filter $\mathbf{w} = \mathbf{s}$ is used after the signal synthesis process. This necessitates the use of

a slightly different frequency allocation process (and a different bisection algorithm), but it ultimately is waterfilling-like as well.

5.2.3 Applying the Relaxed BQP

With the competing techniques described, we now discuss how we might employ the relaxed BQP process for constant-modulus constrained problems to attack this problem, as well as note some specific results that are implied by the theorems in Chapter 4.2. Much of this analysis is nearly identical to that of Chapters 3 and 4, so we will be economical in our explanations.

First, we must massage the original problem (5.5) into a form similar to (2.30), so we can perform the relaxation. As noted above, we have the following relationships between the various parameters in the constant modulus problem. The total number of (complex) parameters is $2N$, the target channel matrix $\mathbf{T} = \mathbf{I}_N$, the modulus constraint is $\rho = 1$, and the clutter channel matrices $\mathbf{\Gamma}_q$ are given by

$$\mathbf{\Gamma}_q = \begin{cases} \mathbf{J}_{-q} & q \in \{1, \ldots, N-1\} \\ \mathbf{J}_{q-N+1} & q \in \{N, \ldots, 2(N-1)\} \end{cases} . \tag{5.8}$$

However, the original problem in does not explicitly specify an "optimal" denominator $\mathbf{w}^H\mathbf{s}$ and thus does not directly provide us with a particular Capon constraint κ.

At the very least, we can use the Slater condition of Theorem 4.1 to identify bounds for κ where the duality gap is guaranteed to be zero and the Lagrange parameters are guaranteed to exist. Recall that the theorem states the Slater condition is satisfied if $N\rho^2 > \frac{|\kappa|^2}{\|\mathbf{T}\|_F^2}$. When

applied to this problem, and after some algebra, this condition becomes

$$|\kappa| < N. \tag{5.9}$$

This is a sensible constraint, since the matched filter implies $|\kappa| = N$, and an optimal mismatched filter would likely have a smaller alignment than the matched filter itself. With that in mind, we also note that since the target matrix is the identity, the relaxed version of the Capon constraint becomes

$$\mathrm{tr}(\mathbf{B}_2^H) = \kappa \tag{5.10}$$

$$\mathrm{tr}(\mathbf{B}_2) = \kappa^*. \tag{5.11}$$

Thus, if $\kappa \in \mathbb{R}$, we can assume that $\mathbf{B}_2$ is a Hermitian matrix. This is the limit of what we can do theoretically; thus, we will investigate the implications of changing κ in the simulations below.

Additionally, since we can reasonably assume that $\mathrm{rank}(\mathbf{R_{ni}}) = N$, we can assume that the dual variable $\tilde{\boldsymbol{\mu}}_M \neq \mathbf{0}_{N\times 1}$ and any solution will take on the spectrum-filling character we saw previously. We can also find an explicit form of $\boldsymbol{\Sigma}_2$ as well. Recall that $\boldsymbol{\Sigma}_2 = \sigma_{\mathsf{c}}^2 \sum_{q=1}^{Q} \mathrm{tr}(\Gamma_q^H \mathbf{B}_2)\Gamma_q - \tilde{\mu}_C * \mathbf{T}$. Applying this to the reverberant channel parameters, we see

that $\mathbf{\Sigma}_2$ actually becomes Toeplitz itself, with the form

$$\mathbf{\Sigma}_2 = \begin{bmatrix} \tilde{\mu}_C^* & \mathrm{tr}(\mathbf{J}_{-1}\mathbf{B}_2) & \mathrm{tr}(\mathbf{J}_{-2}\mathbf{B}_2) & \dots & \dots & \mathrm{tr}(\mathbf{J}_{-(N-1)}\mathbf{B}_2) \\ \mathrm{tr}(\mathbf{J}_1\mathbf{B}_2) & \tilde{\mu}_C^* & \mathrm{tr}(\mathbf{J}_{-1}\mathbf{B}_2) & \ddots & & \vdots \\ \mathrm{tr}(\mathbf{J}_2\mathbf{B}_2) & \mathrm{tr}(\mathbf{J}_1\mathbf{B}_2) & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \mathrm{tr}(\mathbf{J}_{-1}\mathbf{B}_2) & \mathrm{tr}(\mathbf{J}_{-2}\mathbf{B}_2) \\ \vdots & & \ddots & \mathrm{tr}(\mathbf{J}_1\mathbf{B}_2) & \tilde{\mu}_C^* & \mathrm{tr}(\mathbf{J}_{-1}\mathbf{B}_2) \\ \mathrm{tr}(\mathbf{J}_{(N-1)}\mathbf{B}_2) & \dots & \dots & \mathrm{tr}(\mathbf{J}_2\mathbf{B}_2) & \mathrm{tr}(\mathbf{J}_1\mathbf{B}_2) & \tilde{\mu}_C^* \end{bmatrix}. \tag{5.12}$$

Clearly, if κ is real, then $\tilde{\mu}_C$ is real and $\mathbf{\Sigma}_2$ is also Hermitian. We could, potentially, use this structure to more efficiently solve the relaxed problem, but we leave this for a later time.

Next, we turn to a discussion of the clutter matrix $\mathbf{C}_V$ and its properties in this application. We begin by observing that for any two shift matrices, we have

$$\mathrm{tr}(\mathbf{J}_i^T\mathbf{J}_j) = \mathrm{vec}(\mathbf{J}_i))^T\,\mathrm{vec}(\mathbf{J}_j) = (N-|i|)\delta_{ij} \tag{5.13}$$

where it is assumed that i and j share the same sign. If i and j do not share the same sign, then the trace above is also zero. These facts imply the vector set $\left\{\frac{1}{N-|i|}\,\mathrm{vec}(\mathbf{J}_i)\right\}_{i=-(N-1),i\neq 0}^{N-1}$ is orthonormal, and spans a $2(N-1)$ subset of $\mathbb{C}^{N^2}$. Since the clutter matrix is given by

$$\mathbf{C}_V = \sigma_{\mathsf{c}}^2 \sum_{\substack{i=-(N-1) \\ i\neq 0}}^{N-1} \mathrm{vec}(\mathbf{J}_i)\,\mathrm{vec}(\mathbf{J}_i)^T, \tag{5.14}$$

we can surmise that the spanning subset above is, in fact, the range $\mathcal{R}(\mathbf{C}_V)$! Hence, for the reverberant channel, $\mathrm{rank}(\mathbf{C}_V) = 2(N-1)$, which takes up a mere fraction of its total dimension. We might imagine that this means an optimal solution has many degrees of freedom to choose from in nulling the clutter, but it also limits the possibility of unique

solutions. However, this does not guarantee that a given solver will be able to attain such nulling under the constraints.

5.3 Simulations

With the general theory described, we now turn to comparative simulations to determine if the RBQP approach is viable for the reverberant channel, or at least comparable to the techniques demonstrated in [36]. In this section, we will consider varying the total length of the signal and the type of signal-independent jamming encountered to demonstrate some potential gains and challenges of blindly applying our technique to other models.

5.3.1 Common Simulation Parameters & Methods

In these simulations, we use most of the common parameters from [36]. In particular, we assume that the signal-independent covariance matrix $\mathbf{R}_{\mathbf{ni}}$ has the following form

$$\mathbf{R}_{\mathbf{ni}} = \sigma_{\mathtt{n}}^2 \mathbf{I}_N + \sigma_{\mathtt{i}}^2 \mathbf{R}_{\mathtt{i}}, \tag{5.15}$$

where $\sigma_{\mathtt{n}}^2, \sigma_{\mathtt{i}}^2$ are the noise and jammer powers, respectively, and $\mathbf{R}_{\mathtt{i}}$ is the normalized correlation matrix of the interference/jamming, which is also assumed Toeplitz. Assume that the discrete frequency spectrum of the jammer is given by the coefficients $\{\eta_p\}_{p=1}^{2N-1}$, each

evaluated at the frequency $f_p = \frac{p-1}{2N-1}$. Then, the jammer correlation matrix is given by

$$\mathbf{R}_{\mathtt{i}} = \begin{bmatrix} 1 & \frac{q_1^*}{q_0} & \dots & \frac{q_{N-1}^*}{q_0} \\ \frac{q_1}{q_0} & 1 & & \frac{q_{N-2}^*}{q_0} \\ \vdots & & \ddots & \vdots \\ \frac{q_{N-1}}{q_0} & \frac{q_{N-2}}{q_0} & \dots & 1 \end{bmatrix} \tag{5.16}$$

where the coefficients q_i are the elements of the inverse DFT of $\boldsymbol{\eta} = [\eta_1 \ \dots \ \eta_{2N-1}]$, given by IDFT($\boldsymbol{\eta} = [q_0 \ q_1 \ \dots \ q_{N-1} \ q_{N-1}^* \ \dots \ q_1]$, and we have normalized the matrix to have an all-ones diagonal.

For consistency with [36], we choose $\sigma_{\mathtt{n}}^2 = 0.1, \sigma_{\mathtt{i}}^2 = 100$, and $\sigma_{\mathtt{c}}^2 = 1$ throughout all simulations. We consider a large variety of signal lengths, namely $N = [5, 8, 10, 20, 25, 50, 75, 100, 150, 200]$. While CREW(fre) and CREW(mat) have been shown to efficiently compute signal lengths in the thousands, we limit the maximum problem size to 200 because CREW(gra) (owing to the need for first-order derivatives) and RBQP (owing to the limitations of certain interior-point solvers) can require significant amounts of memory to operate. As before, to enable computation of the RBQP method, we implement the problem in CVX [87, 86] and use its interfaces to the solvers SDPT3 [89, 90] and SeDuMi [91]. In all cases, if we compare RBQP with CREW in terms of MSE, we use the CREW algorithm with the lowest MSE value for a given signal length. All CREW solvers were initialized with the Golomb polyphase sequence of length N (see [108, p. 128]), as they were in [36].

5.3.2 Spot Jamming

We begin by considering the behavior of these algorithms in response to a spot jammer placing power at a single frequency $f_0 = 0.2$ Hz. The power spectrum of such a jammer is

given by

$$\eta_p = \begin{cases} 1 & p = \lfloor (2N-1)f_0 \rfloor \\ 0 & \text{otherwise} \end{cases} \tag{5.17}$$

where the range of p is defined above.

First, we consider if the "optimal" value of $|\kappa|$ obtained through any of the CREW algorithms indeed satisfies the Slater condition (5.9). Figure 5.1 shows the obtained value of $|\kappa|$ from CREW(fre) in the spot jamming scenario as a function of the signal length N, with the dotted line denoting the Slater condition boundary. We see that, indeed, the iterative technique satisfies the relaxed Capon constraint – we observed similar behavior from both CREW(mat) and CREW(gra), so we will not consider them here. With this in mind, we now note that

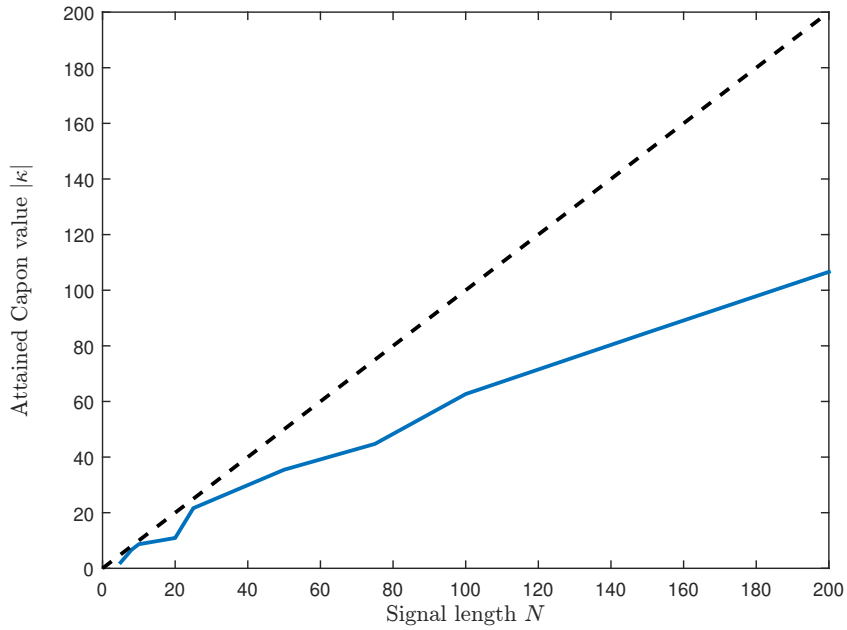

Figure 5.1: Obtained Capon constraint $|\kappa|$ as a function of associated signal length N for CREW(fre), under spot jamming.

we set the Capon constraint $\kappa = 1$ for all instantiations of the RBQP, because, as noted

previously, so long as the Slater condition is satisfied, the relaxed cost function scales with κ to keep the resultant SINR/MSE consistent.

With this in mind, we now attempt to compare the CREW methods and the RBQP process. In Figure 5.2, we compare the best MSE obtained from the CREW algorithms (denoted CREW(best)), the optimal value of the QSDP (QSDP Opt. Val.), and the best MSE obtained from the approximate-sum solution $\mathbf{b}_K^a$ made feasible (denoted RBQP-Feas). It is clear

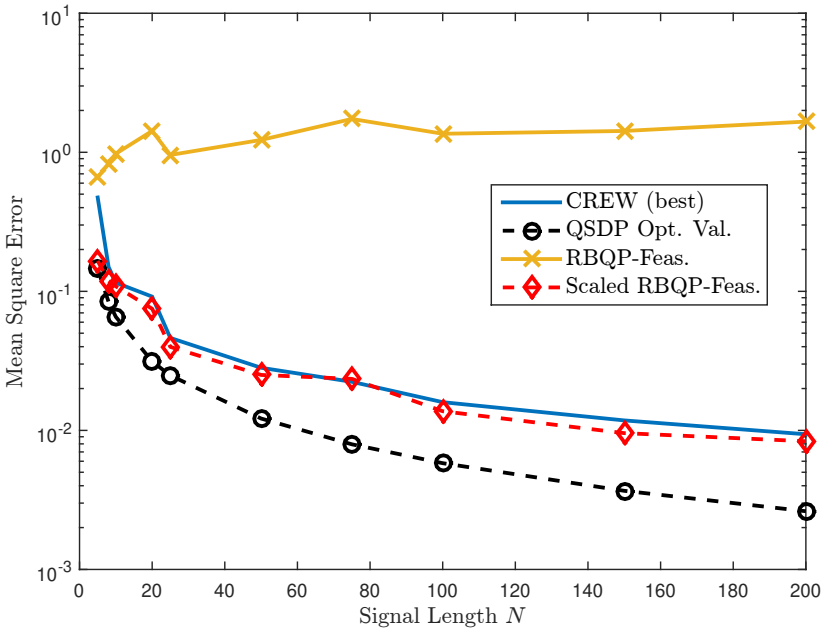

Figure 5.2: MSE comparison of CREW algorithms, the QSDP bound, and two versions of RBQP-Feas. under spot jamming.

that while our QSDP relaxation does provide a lower bound on the MSE obtained by the CREW methods, RBQP-Feas. appears to struggle with approaching either the bound or the CREW-obtained MSE. Interestingly, however, if RBQP-Feas.'s MSE is scaled by $N-1$ (shown as Scaled RBQP-Feas.), then RBQP-Feas.'s performance is quite close to (and in some cases, bests) that of CREW. This is our first clue that the primary cause of the performance loss may lie in a scaling issue in either the synthesis or the relaxed solver itself.

A natural question, as it was in Chapter 4, is if refinement of the feasible solution can improve performance. The results of multiple refinement processes are shown in Figure 5.3 below, where we refine the RBQP-Feas. signal/filter pair with one-step alternating minimization (AM) or one of the CREW techniques (RBQP-(fre)-R, RBQP-(gra)-R, RBQP-(mat)-R). While all of the refinement techniques certainly improve upon the unrefined RBQP-Feas.,

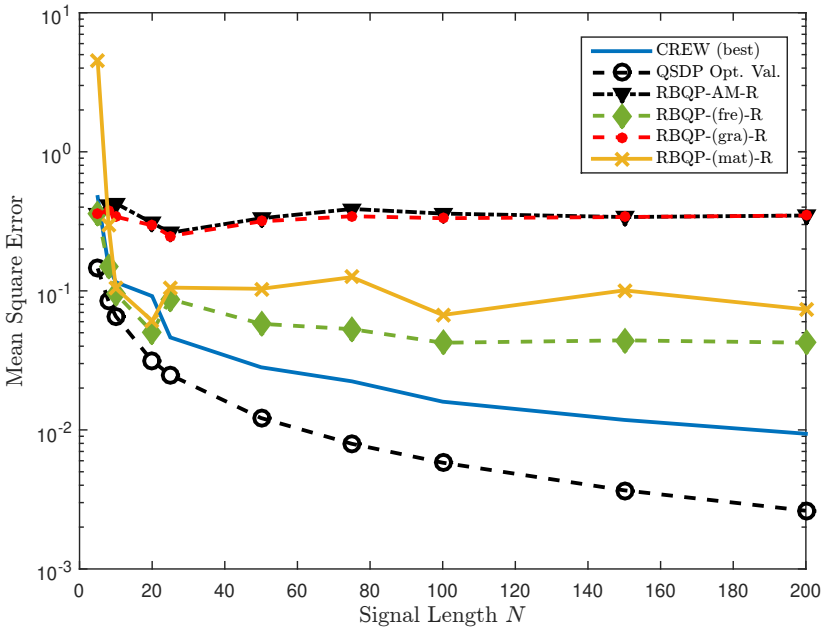

Figure 5.3: MSE comparison of refinement of RBQP techniques and baselines under spot jamming.

refinement through CREW(fre) (RBQP-(fre)-R) shows the most promise, though for longer signals, the performance gains taper off. This gives us further reason to believe that while the current solver techniques for QSDP determine the actual optimal value for the original program, they may not be able to obtain good-quality solutions for larger problem sizes, which in turn affects the capacity for refinement.

We can diagnose this performance issue in multiple ways, but the most instructive is determining if the conditions for the rank bound given by Theorem 3.8 are satisfied. Since our simulations assume a signal-independent covariance $\mathbf{R_{ni}}$ that is full rank, the theorem tells us

that the rank of the optimal relaxed solution rank($\mathbf{B}$) is bounded above by $N - \text{rank}(\mathbf{\Sigma}/\mathbf{R_{ni}})$ where the matrix $\mathbf{\Sigma}/\mathbf{R_{ni}}$ is the Schur complement of the upper diagonal block of the optimal dual variable Σ. A comparison of rank($\mathbf{B}$), the rank of the dual Schur complement rank($\mathbf{\Sigma}/\mathbf{R_{ni}}$) and the expected rank bound above for the spot jamming scenario is shown in Figure 5.4. We have estimated the ranks by determining the point at which the singular values of the respective matrices drops below 1e-12. If the relevant solvers were obtaining

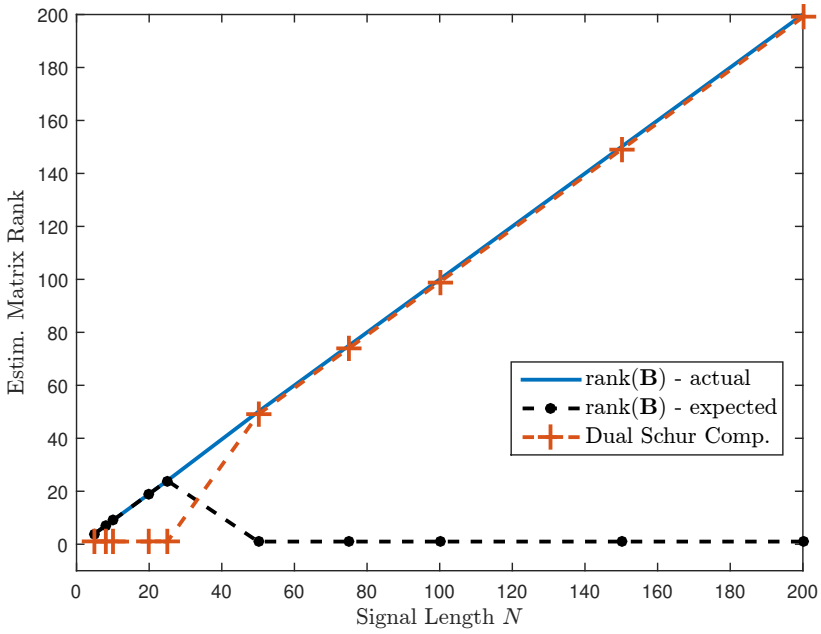

Figure 5.4: Estimated rank comparison of QSDP optimal solution, optimal dual Schur complement, and rank bound of Theorem 3.8 under spot jamming.

exact KKT-satisfying solutions to the relaxed problem, then the actual rank curve would either be below or along the expected rank curve. However, past $N = 50$, this is not the case, indicating that the interior point solvers of SDPT3/SeDuMi do not produce exact solutions to the desired QSDP, which certainly can lead to the recovery problems noted above. Note that while the curves for rank($\mathbf{B}$) and rank($\mathbf{\Sigma}/\mathbf{R_{ni}}$) appear to coincide past $N = 50$, they are actually separated by a constant factor of one. This correspondence is further evidence of the limitations of the current interior-point solvers, since ideally $\text{rank}(\mathbf{B}) + \text{rank}(\mathbf{\Sigma}/\mathbf{R_{ni}}) \leq N$.

We therefore conclude that, for this particular problem, increasing the problem size past a certain point can lead to issues in recovering an acceptable rank-one approximate solution from our relaxation. As we will see below, however, this is not necessarily true across all environmental scenarios.

5.3.3 Barrage Jamming

We now repeat the above analysis for the case of a barrage jammer emitting in the frequency band $[f_1, f_2] = [0.2, 0.3]$ Hz, which corresponds to a power spectrum of

$$\eta_p = \begin{cases} 1 & p \in [\lfloor (2N-1)f_1 \rfloor, \lfloor (2N-1)f_2 \rfloor] \\ 0 & \text{otherwise} \end{cases}. \tag{5.18}$$

First, we return to the comparison of CREW(best) with the QSDP optimal value and the feasible recovery RBQP-Feas, which is shown in Figure 5.5. As expected, the increase in jamming has raised the MSE for all of the techniques, and RBQP-Feas. still struggles as the signal length N increases. While we were unable to determine an exact scaling to "return" RBQP-Feas. to a respectable value, we do believe that it a function of the jammer frequency interval and N.

As before, we then examined if refinement of RBQP-Feas. could restore performance, if not to the bound, then at least to the level attained by the CREW techniques. Figure 5.6 demonstrates clearly that all of the refinement techniques are able to restore performance to a far higher degree than in the spot jamming scenario. In particular, RBQP-CREW(fre)-R refinement mostly coincides with the CREW (best) performance, with RBQP-(mat)-R coming in a close second for larger signal lengths. The relatively poorer performance of RBQP-AM-R and RBQP-(gra)-R can be attributed to how ill-conditioned these procedures

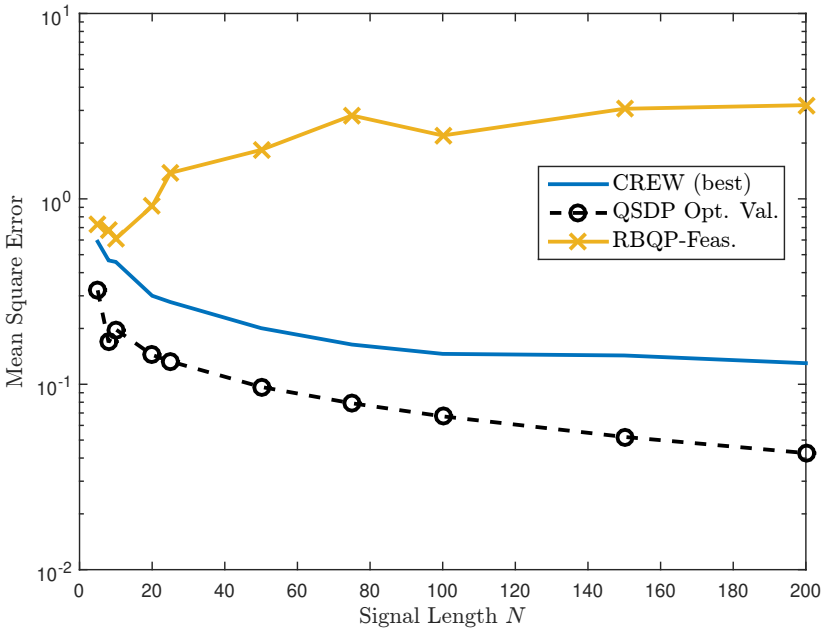

Figure 5.5: MSE comparison of CREW algorithms, the QSDP bound, and two versions of RBQP-Feas. under barrage jamming.

can become for large variable sizes. Similar to the constant-modulus problem in Chapter 4,

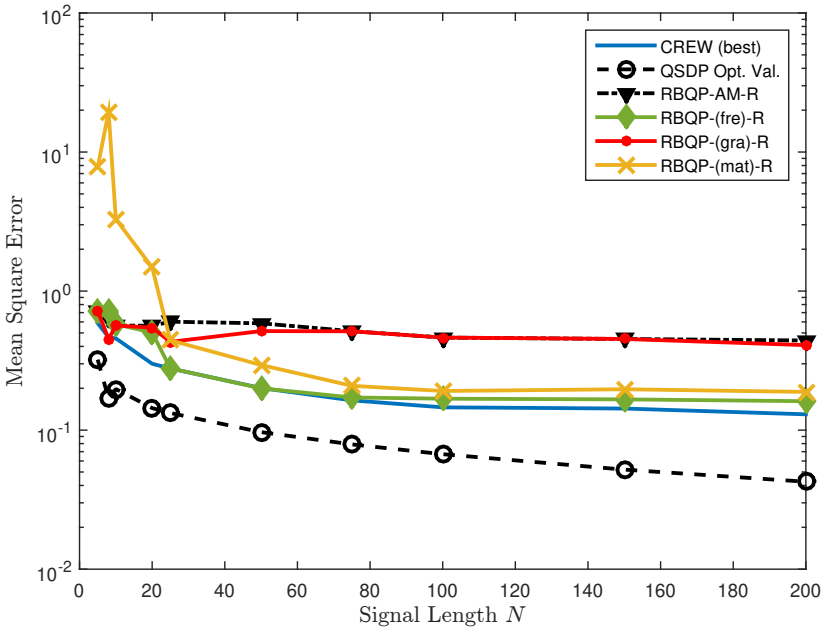

Figure 5.6: MSE comparison of refinement of RBQP techniques and baselines under barrage jamming.

we attribute the performance gain in this "more-interference" case to the additional known structure in the problem that allows the initial refinement solution to be "closer" to the optimal than in the spot jamming scenario.

We can confirm this by once again examining the ranks of the optimal solution $\mathbf{B}$, the dual Schur complement ($\mathbf{\Sigma}/\mathbf{R_{ni}}$), and the expected rank bound, which is shown in Figure 5.7 below. Again, the rank is estimated to be the point where each matrix's singular values drop below 1e-12. Here, the rank of $\mathbf{B}$ coincides with the predicted rank bound up to $N = 100$

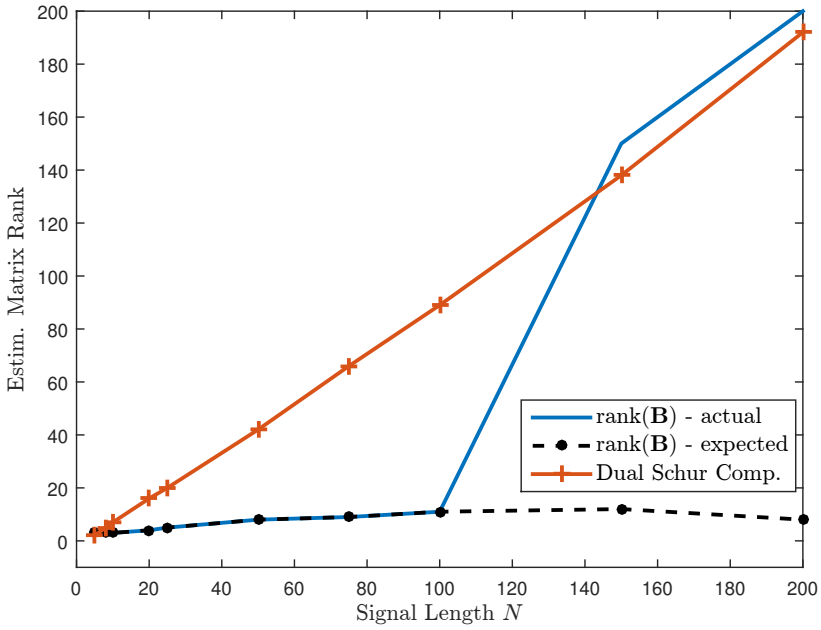

Figure 5.7: Estimated rank comparison of QSDP optimal solution, optimal dual Schur complement, and rank bound of Theorem 3.8 under barrage jamming.

(which is never more than 11). This indicates again that the additional known structure in $\mathbf{R_{ni}}$ permits the solver to more accurately identify the optimal relaxed solution, but only up to a point. Hence, we believe that our technique is, in fact, still valid even for this very challenging (but nominally simpler) model, and the limiting process is the solver. We leave the more detailed work necessary to complete such an analysis for future research.

Chapter 6

Conclusions & Future Work

In this dissertation, we examined the utility of a variety of convex relaxations for joint signal-beamformer design for radar systems. We demonstrated both the limitations and potential for various iterative methods, as well as developed quadratic semidefinite programming techniques to directly solve for both parameters simultaneously across a number of constraints and radar paradigms. This chapter summarizes the most salient aspects of our contributions and provides directions and questions for future research.

6.1 Conclusions

First, in Chapter 2, we demonstrated that under most typical design paradigms, the joint problem is a complex biquadratic program (BQP) with, at minimum, coupled bilinear constraints. Since the BQP is known to be NP-hard and non-convex, historical methods focused on either explicitly quartic procedures in the signal variable or an iterative tradeoff between signal and beamformer. We showed that a variety of new methods may be employed in the latter case – in particular, ones that consider simple projection of the eigenspace of the previ-

ous iterate onto the cost function and other constraints. These methods were further shown to typically have superior convergence properties to other methods extant in the literature and, as a preview for later chapters, converged to the optimal value of the relaxed BQP in Chapter 3. Furthermore, we proved that monotonic convergence for any bi-convex program with per-iteration dependence on the constraint sets is guaranteed, with certain conditions ensuring superlinear convergence. This proof also demonstrated that, for our particular problem, appropriately considering the Capon constraint at each iteration guarantees identical limit points are reached, regardless of implementation. Therefore, our contributions showed potential simultaneous optimization would provide both theoretical and practical benefits.

Next, in Chapter 3, we outlined the theory of quadrative semidefinite programs in relation to the joint design problem, both in general and for the the least restrictive power-only constraint. By using the primal KKTs, we were able to completely characterize all feasible solutions to the relaxation. We first demonstrated an immediate linkage between the sequential feasibility of the projected QCQP methods of Chapter 2 and the Slater condition of the relaxed problem, which is fundamentally achievable in a practical radar system.

Furthermore, we showed that a solution to the relaxed problem is always power-bounded if the rank of the covariance matrix of the signal-independent sources is full. In these cases, the rank of the relaxed problem is guaranteed to be no greater than the length of the signal, N, which bodes well for potential recovery of a high-quality rank-one solution for the initial problem. In fact, we also showed that the rank of the relaxed solution is at most the multiplicity of the maximum eigenvalue of the dual matrix, which elegantly links the target, the clutter operator, and the noise-interference covariance matrix elegantly. While most theoretical cases assume a full rank covariance, it is possible that in a realistic scenario to have an incomplete estimate that is less than full rank. In these cases, all solutions are not power-bounded, and attempt to encroach on this assumed nullspace. Hence, the rank

bound is inflated by the size of the nullspace of the covariance. We additionally showed a comprehensive characterization of all solutions in this non-power-bounded scenario.

Regardless of the rank of the signal-independent sources, our theoretical discoveries were confirmed through computational experiments. These showed that, in fact, most parameter cases lead to an effectively rank-one solution, particularly in the power-bounded solutions. Hence, optimal recovery of a good filter-signal pair is more than possible; it is probable. Further numerical analysis revealed that across both angle and Doppler variation of the target, the QSDP/RBQP method can outperform all of the alternating methods discussed in Chapter 2. This performance gain typically occurs when the power P_o is large, since overall convergence depends on appropriate parameter scaling. This boost in SINR also leads to performance increases in detection, which is a boon for our motivating waveform adaptive STAP paradigm. We also saw that the adapted patterns seen in traditional STAP are functionally preserved, subject to the availability of degrees of freedom to obtain the desired amount of resolution and source separability.

In Chapter 4, we extended the QSDP formulation to account for more restrictive waveform constraints: constant modulus and similarity. To begin, however, we demonstrated that the common interpretation of the similarity constraint is mostly fallacious through a simple counterexample. Rather than interpreting this constraint as permitting inheritance of "good" waveform properties, the similarity constraint should instead be regarded as "merely" a proximal term induced by other systemic constraints, like estimation validity or transmitter actuation. With this in hand, we then reapplied our methods of Chapter 3 to the newly constrained versions of the design problem.

First, we considered the constant modulus formulation, where there are historically many approaches to casting the design problem. Since our relaxation approach essentially requires design of both amplitude and phase components of the signal (due to the spectral nature of the relaxation), we did not attempt an explicitly phase only design. Again using the

KKTs, we showed the general observations on rank and feasibility observed in the power-constrained problem of Chapter 3 still mostly hold, with rank guarantees dependent on the environment's capacity to meet the modulus constraint. As before, in full-rank noise-and-interference scenarios, the rank of the optimal relaxed solution is no more than N. Unlike in the power constrained case, however, we did not find that conventional convex solvers were able to recover numerically rank-one feasible solutions across all scenarios. As we showed, in medium-to-low Doppler frequencies (or indeed anywhere where the effective alignment of the target with the clutter was high), the effective rank of the constant modulus relaxed solution inflates to at least two. Furthermore, the magnitude of the secondary eigenvalue can be a significant fraction of the overall spectral power, which meant that not only were maximum eigenvector synthesis solutions suboptimal, they were often non-feasible! Attempting to use the approximate basis sum procedure to synthesize solutions performed no better, even when forced into feasibility. Introducing signal-independent interference led to the seemingly paradoxical reduction of the region over which the rank-inflation occurred, which we attribute mostly to introducing additional known structure into the signal-independent portion of the received signal covariance matrix, as well as possibly "smoothing out" difficult regions of clutter. We anticipate that if the interference is more significantly aligned with the target regions, we anticipate that some of this supposed gain will, in fact, be lost.

Regardless of the environmental conditions that may mitigate the impact of this rank-inflation, this loss of performance is not insignificant. This difficulty prompted our investigation into other methods of solution synthesis from the relaxed solution spectrum. First, we considered a randomization strategy, inspired by the literature for linear semidefinite programs and its use in some of the iterative techniques of Chapter 2. However, at least in our QSDP case, the variance of the cost function under consideration actually has an excess term that is effectively non-convex. Therefore, any approximation arising from this method minimized this new total cost, which may not minimize the particular cost function that we desire. This inconsistency's impact was borne out in our simulations, which showed that

randomization provided very poor solutions, even when compared to the aforementioned extremely suboptimal approximations above.

Our second strategy considered applying one of the alternating techniques from Chapter 2 to refine the approximate basis solutions after they had been made feasible. In principle, either the signal or the filter obtained from an approximate basis solution can be used to initialize any of the algorithms discussed previously. For fairness, we chose to use a common gradient descent technique to perform the refinement process.

Even with very few iterations, this refinement technique restored all gain lost from the "expected" relaxed optimal value by the approximate basis solutions in rank-inflated scenarios. There are possible cases where additional refinement iterations would be valuable, among other variations that might be considered, but we will leave such speculation to the next section. Furthermore, not only did this refinement process provide a rank-one "solution" to the relaxed problem, it also outperformed the competing techniques currently in the literature by a wide margin (up to 3 dB at certain angle-Doppler locations). Additionally, the overall variation of the refined waveform-filter pair from the initializer is quite small, indicating both that the initial approximation is quite close to the optimal point (and the cost function can vary significantly even slightly away from the minima) and that a different QSDP solver could produce this solution on its own. On this latter point, we note that this is possible because the general version of the rank prediction described in Lemma 3.2 is merely an upper bound! Ultimately, we are convinced that, with the addition of the refinement step, our design philosophy fares well for constant-modulus problems.

Far more immediately successful was the application of the joint relaxation to similarity-constrained problems. Due to the added structure of the goal signal, we were able to propose a nearly-closed-form solution in power-bounded cases, and also identified succinctly when solutions are alignment-bounded. In particular, solutions are alignment-bounded if and only if the signal-independent disturbance covariance is full rank and the goal signal is not

already sufficiently aligned with the power-constrained solutions detailed in Chapter 3. This was provided in Theorem 4.5 and shown numerically to hold generically. Hence, we propose the concept of performing the less-constrained optimization prior to the more constrained optimization, identify when the power-constrained solution's alignment exceeds that of the goal signal's, and only perform the similarity-constrained optimization when this exceedance does not occur.

We showed that in most cases, the rank of the relaxed solution is again upper-bounded by the signal length N and numerically produces nearly rank-one optimal solutions. This enabled us to eschew any further refinement of the approximate basis recovery technique, as explored for the constant modulus case. Additionally, we showed that, once again, our relaxation technique outperforms those available in the literature by several dB across multiple scenarios. This is true even when solver tolerances for the competing techniques are set to rather low values. We also directly demonstrated, through multiple proofs, our assertion that the ambiguity "inheritance" interpretation of the similarity constraint is, in fact, severely misguided.

In Chapter 5, we intended to show that our general technique has utility beyond the waveform-adaptive STAP applications considered in Chapters 3 and 4, though there is need for further investigation To do this, we discussed another sensing model that admits the vector-matrix channel representation: the basic reverberant channel as investigated by Stoica, He, and Li [36]. In this case, the channel model advanced in Chapter 2 applies directly, and so our techniques would nominally be a one-to-one map. While we found that the QSDP process did produce an effective lower bound for the optimal value of the mean square error, recovery of good rank-one solutions was challenging, even with refinement through the CREW techniques. This was especially for the less-difficult spot jamming case. However, the performance recovery attained by refinement in the barrage jamming case, as well as our speculative investigation of the ranks of the obtained primal and dual relaxed soluti-

ons, indicated that numerical difficulties encountered by the interior-point solvers used to obtain the relaxed solution were the primary driver of the performance loss. Additionally, we note that CREW(fre) and CREW(mat) take advantage of additional structure inherent in the problem to both speed up calculation (through the use of FFTs) and provide better numerical stability. We believe that a purpose-built QSDP solver could mitigate some of these losses and attain better overall performance. Finally, we note that this channel model has limitations of its own, which contributes to the numerical issues we encountered in the QSDP solver. Namely, the reverberant channel in [36] is a SISO model and therefore does not have any other degrees of freedom to separate the multipath scattering from the direct scattering, other than the given time shifts. Additional degrees of freedom in either space or time would radically change the complexion of the channel matrices, as in the WA-STAP case, and can lead to far better performance in general. Furthermore, we believe that the model itself incorrectly states the overall resolvability of the multipath components – it is hard to believe that any reverberant channel would have resolvable fluctuations up to the signal length within a short overall pulse length. Ultimately, a joint signal-filter design process is only as good as the structure of the model with which it is presented. In that regard, irrespective of the signal model, we believe our technique does quite well indeed.

6.2 Limitations and Future Work

While, in a selfish way, we would love to assume that this work is the final word on all things QSDP for active sensing optimization, it is inescapable fact that there are both limitations to what we have presented here, as well as significant room for future growth and research along these lines. Since we are our own harshest critic, we start with the shortcomings.

First, we recognize that the bulk of our work hinges on the ability of the system designer to create a linear channel model that sufficiently captures the interaction between independent

transmit and receive degrees of freedom. Thus, more complex parameter choices, like field polarization or carrier frequency, are functionally not supported by our technique, at least not directly. It is our belief, however, that some of these design choices – automated or not – might be better addressed by higher levels of control and information tracking, see, e.g. feedback control in target tracking [11, 12]. This also extends to the choice of constraints considered; we did not consider spectral co-existence, limited transmission alphabets, or direct ambiguity control, though our technique may reasonably be extended to include them.

Continuing on the modeling limitations, all of our developments were predicated on complete prior knowledge of the signal-independent source covariance and the channel responses of the clutter and target. There are two possible directions that are worth investigating. First, ironically, due to the form of our relaxation, uncertainty in the clutter can be handled exactly. Our reasoning is thus: as with any modeling, we can assume the clutter patches $\mathbf{\Gamma}_q$ come from some expected distribution. Assuming the vectors $\text{vec}(\mathbf{\Gamma}_q)$ come from a zero-mean distribution, we might consider a design where $\mathbf{E}\,\mathbf{C}_V = \mathbf{E}\{\text{vec}(\mathbf{\Gamma}_q)\,\text{vec}(\mathbf{\Gamma}_q)^H\}$ is used instead of the random variable $\mathbf{C}_V$. Therefore, so long as we assumed a known target channel matrix, the optimization techniques shown in this dissertation functionally do not change, though a certain loss of performance should be expected from using the clutter covariance information $\mathbf{E}\{\mathbf{C}_V\}$ instead of the clairvoyantly known $\mathbf{C}_V$.

Next, we might consider uncertainty in the target channel matrix as well; that is the target channel matrix comes from some distribution $\text{vec}(\mathbf{T}) \sim \mathcal{CN}(\mathbf{0}_{N^2ML\times 1}, \mathbf{R}_\tau)$, where the covariance matrix $\mathbf{R}_\tau$ is known a priori. Some earlier work in this vein was performed by [35], but it (obviously) has yet to be explored in the RBQP framework. However, we note that this would essentially result in biquadratic constraint, due to the structure of $\mathbf{R}_\tau$. If relaxed, this becomes a quadratically-constrained QSDP (QC-QSDP), which some of the solvers discussed above cannot handle and others handle poorly. Fortunately, there has been

some research into purpose-built QC-QSDP solvers which could be employed for this task [111].

Finally, we note that, throughout this dissertation, we used traditional interior-point solvers, which functionally limited the overall size of the problems we could reasonably consider with our available computing power. We believe strongly that employing purpose-built QSDP solvers like QSDPNAL [56, 92] would lead to significant solver speedup and also permit the investigation into realistic problem sizes (tens-to-hundreds of spatial, Doppler, and temporal samples). A purpose-built augmented Lagrangian solver would also be extremely useful in our specific case, because the block separability and "obvious" optimal dual values allow for extreme simplification of the process. However, we did not attack this specifically in this dissertation – our goal was to investigate the technique, not directly create the optimization tools.

In summary, what we have shown in this dissertation is merely the potential of joint relaxation procedures for simultaneous design of transmission and reception in active sensing systems. We recognize that there is far more work yet to be done, and we hope that this work can be an effective starting point.

Bibliography

[1] J. F. O'Rourke and M. G. Amin, "Root map characterization of interference location in a wide aperture antenna array," in *Proceedings of the 32nd Midwest Symposium on Circuits and Systems*, Aug 1989, pp. 590–593 vol.1.

[2] J. F. O'Rourke, "The Use of Root Maps in Antenna Arrays," Master's thesis, Villanova University, Villanova, PA, May 1989.

[3] J. R. Guerci, J. S. Bergin, R. J. Guerci, M. Khanin, and M. Rangaswamy, "A new MIMO clutter model for cognitive radar," in *2016 IEEE Radar Conference (RadarConf)*, May 2016, pp. 1–6.

[4] M. I. Skolnik, Ed., *Radar Handbook*, 3rd ed. McGraw Hill, 2008.

[5] P. Setlur and M. Rangaswamy, "Signal dependent clutter waveform design for radar STAP," in *In Proc. IEEE Radar Conference*, 2014.

[6] ——, "Joint filter and waveform design for radar STAP in signal dependent interference," https://arxiv.org/abs/1510.00055, US Air Force Res. Lab., Sensors Directorate, WPAFB, Dayton, OH, Tech. Rep. *DTIC, available at : https://arxiv.org/abs/1510.00055* , 2014.

[7] ——, "Projected Gradient Waveform Design for Fully Adaptive Radar STAP," in *In Proc. IEEE Radar Conference*, 2015.

[8] ——, "Waveform design for radar STAP in signal dependent interference," *IEEE Transactions on Signal Processing*, vol. 64, no. 1, pp. 19–34, January 2016.

[9] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part III: Radar-Sonar Signal Processing and Gaussian Signals in Noise.* John Wiley & Sons, Inc., 2001.

[10] G. E. Smith, Z. Cammenga, A. Mitchell, K. L. Bell, M. Rangaswamy, J. T. Johnson, and C. J. Baker, "Experiments with cognitive radar," in *2015 IEEE 6th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, Dec 2015, pp. 293–296.

[11] D. Kershaw and R. Evans, "Optimal waveform selection for tracking systems," *IEEE Trans. on Information Theory*, vol. 40, no. 5, pp. 1536 – 1550, Sep. 1994.

[12] ——, "Waveform selective probabilistic data association," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 33, no. 4, pp. 1180–1188, 1997.

[13] J. Guerci and E. Baranoski, "Knowledge-aided adaptive radar at DARPA: an overview," *IEEE Signal Processing Magazine*, vol. 23, no. 1, pp. 41–50, Jan 2006.

[14] J. Chalk, "The optimum pulse-shape for pulse communication," *Proceedings of the IEE - Part III: Radio and Communication Engineering*, vol. 97, pp. 88–92(4), March 1950. [Online]. Available: http://digital-library.theiet.org/content/journals/10.1049/pi-3.1950.0015

[15] P. Woodward, *Probability and information theory, with applications to radar.* Pergamon Press London, 1953.

[16] M. Bell, "Information theory and radar: mutual information and the design and analysis of radar waveforms and systems," Ph.D. dissertation, California Institute of Technology, 1988.

[17] ——, "Information theory and radar waveform design," *IEEE Transactions on Information Theory*, vol. 39, no. 5, pp. 1578–1597, 1993.

[18] R. Manasse, "The use of pulse coding to discriminate against clutter," Massachusetts Institute of Technology, Lincoln Laboratory, Tech. Rep. 312-12, June 1961. [Online]. Available: http://www.dtic.mil/dtic/tr/fulltext/u2/260230.pdf

[19] H. L. Van Trees, "Optimum signal design and processing for reverberation-limited environments," *IEEE Transactions on Military Electronics*, vol. 9, no. 3, pp. 212–229, July 1965.

[20] T. Kooij, "Optimum signals in noise and reverberation," in *Signal processing with emphasis on underwater acoustics. Vol. II. (Proceedings of NATO Advanced Study Institute, Twente Institute of Technology, Enschede, The Netherlands, 12-23 Aug 1968)*, vol. 1. NATO Advanced Study Institute, 1968, pp. 17–1.

[21] M. Ares, "Optimum burst waveforms for detection of targets in uniform range-extended clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 3, p. 138, Jan. 1967.

[22] S. Lee and J. J. Uhran, "Optimum Signal and Filter Design in Underwater Acoustic Echo Ranging Systems," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 9, no. 5, pp. 679–687, Sep. 1973.

[23] J. S. Thompson and E. L. Titlebaum, "The design of optimal radar waveforms for clutter rejection using the maximum principle," *Supplement to IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-3, pp. 581–589, Nov. 1967.

[24] T. Grettenberg, "Signal selection in communication and radar systems," *IEEE Transactions on Information Theory*, vol. 9, no. 4, pp. 265–275, 1963.

[25] S. Applebaum, "Adaptive arrays," Syracuse University Research Corp., Tech. Rep. SURC SPL TR 66-001, August 1966, reprinted in IEEE Transactions on Antennas and Propagation, vol. 24, no. 5, pp. 585–598, Sep. 1976.

[26] A. W. Rihaczek, "Optimal Filters for Signal Detection in Clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-1, no. 3, pp. 297–299, Dec. 1965.

[27] D. DeLong and E. Hofstetter, "On the design of optimum radar waveforms for clutter rejection," *IEEE Transactions on Information Theory*, vol. 13, no. 3, pp. 454–463, July 1967.

[28] ——, "Optimum Radar Signal-Filter Pairs in a Cluttered Environment," *IEEE Transactions on Information Theory*, vol. 16, no. 1, pp. 89–90, Jan. 1970.

[29] W. D. Rummler, "A technique for improving the clutter performance of coherent pulse train signals," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 3, pp. 898–906, Nov. 1967.

[30] S. U. Pillai, D. C. Youla, H. S. Oh, and J. R. Guerci, "Optimum transmit-receiver design in the presence of signal-dependent interference and channel noise," in *Conference Record of the Thirty-Third Asilomar Conference on Signals, Systems, and Computers (Cat. No.CH37020)*, vol. 2, 1999, pp. 870–875. [Online]. Available: http://ieeexplore.ieee.org.wrs.idm.oclc.org/stamp/stamp.jsp?tp=&arnumber=831834&isnumber=17893

[31] S. U. Pillai, H. S. Oh, D. C. Youla, and J. R. Guerci, "Optimal transmit-receiver design in the presence of signal-dependent interference and channel noise," *IEEE Transactions on Information Theory*, vol. 46, no. 2, pp. 577–584, Mar 2000.

[32] S. U. Pillai and H. S. Oh, "Optimum MIMO transmit-receiver design in presence of interference," in *Circuits and Systems, 2003. ISCAS '03. Proceedings of the 2003 International Symposium on*, vol. 4, May 2003, pp. IV–436–IV–439 vol.4.

[33] B. Friedlander, "A subspace framework for adaptive radar waveform design," in *Conference Record of the Thirty-Ninth Asilomar Conference on Signals, Systems and Computers, 2005.*, Pacific Grove, CA, October 2005, pp. 1135–1139.

[34] ——, "Waveform design for MIMO radars," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 43, no. 3, pp. 1227–1238, July 2007.

[35] C. Y. Chen and P. P. Vaidyanathan, "MIMO radar waveform optimization with prior information of the extended target and clutter," *IEEE Transactions on Signal Processing*, vol. 57, no. 9, pp. 3533–3544, 2009.

[36] P. Stoica, H. He, and J. Li, "Optimization of the receive filter and transmit sequence for active sensing," *IEEE Trans. Signal Processing*, vol. 60, no. 4, pp. 1730–1740, April 2012.

[37] O. Aldayel, V. Monga, and M. Rangaswamy, "Successive QCQP refinement for MIMO radar waveform design under practical constraints," *IEEE Transactions on Signal Processing*, vol. 64, no. 14, pp. 3760–3774, July 2016.

[38] G. Cui, X. Yu, V. Carotenuto, and L. Kong, "Space-time transmit code and receive filter design for colocated mimo radar," *IEEE Transactions on Signal Processing*, vol. 65, no. 5, pp. 1116–1129, March 2017.

[39] S. M. Kay and J. H. Thanos, "Optimal transmit signal design for active sonar/radar," in *2002 IEEE Conference on Acoustics, Speech, and Signal Processing (ICASSP '02)*, Orlando, FL, 2002, pp. II–1513–II–1516.

[40] S. Kay, "Optimal signal design for detection of gaussian point targets in stationary gaussian clutter/reverberation," *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 1, pp. 31–41, June 2007.

[41] R. Romero and N. A. Goodman, "Information-theoretic matched waveform in signal dependent interference," in *2008 IEEE Radar Conference*, May 2008, pp. 1–6.

[42] P. Setlur, N. Devroye, and M. Rangaswamy, "Radar waveform design with the two step mutual information," in *2014 IEEE Radar Conference*, May 2014, pp. 1317–1322.

[43] A. Aubry, A. De Maio, M. Piezzo, A. Farina, and M. Wicks, "Cognitive design of the receive filter and transmitted phase code in reverberating environment," *IET Radar, Sonar and Navigation*, vol. 6, no. 9, pp. 822–833, December 2012.

[44] A. Aubry, A. De Maio, A. Farina, and M. Wicks, "Knowledge-aided (potentially cognitive) transmit signal and receive filter design in signal-dependent clutter," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 1, pp. 93–117, January 2013.

[45] G. Cui, H. Li, and M. Rangaswamy, "MIMO radar waveform design with constant modulus and similarity constraints," *IEEE Transactions on Signal Processing*, vol. 62, no. 2, pp. 343–353, Jan 2014.

[46] W. Zhu and J. Tang, "Robust Design of Transmit Waveform and Receive Filter for Colocated MIMO Radar," *IEEE Signal Processing Letters*, vol. 22, no. 11, pp. 2112–2116, November 2015.

[47] B. Tang and J. Tang, "Joint design of transmit waveforms and receive filters for MIMO radar space-time adaptive processing," *IEEE Trans. Signal Processing*, vol. 64, no. 18, pp. 4707–4722, 2016.

[48] J. Li, J. Guerci, and L. Xu, "Signal Waveform's Optimal-under-Restriction Design for Active Sensing," *IEEE Signal Processing Letters*, vol. 13, no. 9, pp. 565–568, September 2006.

[49] C. Ling, J. Nie, L. Qi, and Y. Ye, "Biquadratic optimization over unit spheres and semidefinite programming relaxations," *SIAM Journal of Optimization*, vol. 20, pp. 1286–1310, 2009.

[50] C. Ling, X. Zhang, and L. Qi, "Semidefinite relaxation approximation for multivariate bi-quadratic optimization with quadratic constraints," *Numer. Linear Algebra Appl.*, vol. 19, no. 1, pp. 113–131, apr 2011. [Online]. Available: http://dx.doi.org/10.1002/nla.781

[51] S.-L. Hu and Z.-H. Huang, "Alternating direction method for bi-quadratic programming," *Journal of Global Optimization*, vol. 51, no. 3, pp. 429–446, November 2011.

[52] X. Zhang, C. Ling, and L. Qi, "Semidefinite relaxation bounds for bi-quadratic optimization problems with quadratic constraints," *J. Glob. Optim.*, vol. 49, pp. 293–311, 2011.

[53] Y. Yang and Q. Yang, "On solving biquadratic optimization via semidefinite relaxation," *Computational Optimization and Applications*, vol. 53, no. 3, pp. 845–867, 2012. [Online]. Available: http://dx.doi.org/10.1007/s10589-012-9462-2

[54] I. M. Bomze, C. Ling, L. Qi, and X. Zhang, "Standard bi-quadratic optimization problems and unconstrained polynomial reformulations," *J. Glob. Optim.*, vol. 52, no. 4, pp. 663–687, 2012. [Online]. Available: http://dx.doi.org/10.1007/s10898-011-9710-5

[55] Y. Yang, Q. Yang, and L. Qi, "Approximation bounds for trilinear and biquadratic optimization problems over nonconvex constraints," *Journal of Optimization Theory and Applications*, vol. 163, no. 3, pp. 841–858, December 2014.

[56] X. Li, "A Two-Phase Augmented Lagrangian Method for Convex Composite Quadratic Programming," Ph.D. dissertation, National University of Singapore, 2015.

[57] P. Stoica, J. Li, and Y. Xie, "On Probing Signal Design For MIMO Radar," *IEEE Transactions on Signal Processing*, vol. 55, no. 8, pp. 4151–4161, August 2007.

[58] G. Healey, "The New Moneyball: How Ballpark Sensors Are Changing Baseball," *Proceedings of the IEEE*, vol. 105, no. 11, pp. 1999–2002, Nov 2017.

[59] R. Klemm, *Principles of Space-Time Adaptive Processing.* Institution of Electrical Engineers, 2002.

[60] J. Ward, "Space-time adaptive processing for airborne radar," Massachusetts Institute of Technology, Lincoln Laboratory, Tech. Rep., 1994.

[61] D. Madurasinghe and A. P. Shaw, "Mainlobe Jammer Nulling via TSI Finders: A Space Fast-Time Adaptive Processor," *EURASIP Journal on Applied Signal Processing*, vol. 2006, no. 1, p. 048789, 2006.

[62] Y. Seliktar, D. B. Williams, and E. J. Holder, "A space/fast-time adaptive monopulse technique," *EURASIP Journal on Applied Signal Processing*, vol. 2006, p. 14510, 2006.

[63] J. R. Guerci, *Cognitive Radar: The Knowledge-Aided Fully Adaptive Approach.* Norwood, MA: Artech House, 2010.

[64] J. Capon, "High-resolution frequency-wavenumber spectrum analysis," *Proceedings of the IEEE*, vol. 57, no. 8, pp. 1408–1418, January 1969.

[65] S. M. O'Rourke, P. Setlur, M. Rangaswamy, and A. L. Swindlehurst, "Relaxed Bi-Quadratic Optimization for Joint Filter-Signal Design in Signal-Dependent Space-Time Adaptive Processing (STAP)," Air Force Research Laboratory, Sensors Directorate, Dayton, OH, Interim Tech. Report AFRL-RY-WP-TR-2016-0197, December 2016. [Online]. Available: https://arxiv.org/abs/1703.08115

[66] ——, "Relaxed bi-quadratic optimization for joint filter-signal design in signal-dependent STAP," *IEEE Transactions on Signal Processing*, vol. 66, no. 5, pp. 1300 –1315, March 2018.

[67] J. Gorski, F. Pfeuffer, and K. Klamroth, "Biconvex sets and optimization with biconvex functions: a survey and extensions," *Mathematical Methods of Operations Research*, vol. 66, no. 3, pp. 373–407, December 2007.

[68] U. Niesen, D. Shah, and G. W. Wornell, "Adaptive Alternating Minimization Algorithms," *IEEE Transactions on Information Theory*, vol. 55, no. 3, pp. 1423–1429, March 2009.

[69] K. T. Arasu, M. T. Mohan, S. M. O'Rourke, M. Rangaswamy, and P. Setlur, "Adaptive Alternating Minimization Algorithms on Disjoint Metric Spaces," 2017, unpublished.

[70] B. Jiang, Z. Li, and S. Zhang, "On Cones of Nonnegative Quartic Forms," *Foundations of Computational Mathematics*, vol. 17, no. 1, pp. 161–197, February 2017.

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

[72] Y. Huang and D. P. Palomar, "Rank-Constrained Separable Semidefinite Programming With Applications to Optimal Beamforming," *IEEE Transactions on Signal Processing*, vol. 58, no. 2, pp. 664–678, Feb. 2010.

[73] W. Ai, Y. Huang, and S. Zhang, "New results on hermitian rank-one decomposition," *Mathematical Programming*, vol. 128, no. 1–2, pp. 253–283, 2011.

[74] P. Setlur, S. M. O'Rourke, and M. Rangaswamy, "Constrained least squares, SDP, and QCQP perspectives on joint biconvex radar receiver and waveform design," in *Radar 2017 – International Conference on Radar Systems*, Belfast, UK, October 2017.

[75] N. Z. Shor, "Quadratic optimization problems," *Soviet Journal of Circuits and Systems Sciences*, vol. 25, no. 6, pp. 1–11, 1987.

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

[77] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Athena Scientific, 1999.

[78] L. K. Patton, "On the satisfaction of modulus and ambiguity function constraints in radar waveform optimization for detection," Ph.D. dissertation, Wright State Univ,, Dayton, OH, USA, Jun. 2009.

[79] M. D. Choi, "Completely Positive Linear Maps on Complex Matrices," *Linear Algebra and its Applications*, vol. 10, pp. 285–290, 1975.

[80] J. A. Poluikis and R. D. Hill, "Completely Positive and Hermitian-Preserving Linear Transformations," *Linear Algebra and its Applications*, vol. 35, pp. 1–10, 1981.

[81] J. R. Magnus and H. Neudecker, "The Commutation Matrix: Some Properties and Applications," *The Annals of Statistics*, vol. 7, no. 2, pp. 381–394, Mar. 1979.

[82] L. E. Brennan and F. M. Staudaher, "Subclutter visibility demonstration," Adaptive Sensors Incorporated, Tech. Rep. RL-TR-92-21, March 1992.

[83] N. A. Goodman and J. M. Stiles, "On clutter rank observed by arbitrary arrays," *IEEE Trans. Signal Processing*, vol. 55, no. 1, pp. 178 – 186, January 2007.

[84] M. Kojima, S. Shindoh, and S. Hara, "Interior-point methods for the monotone semidefinite linear complementarity problem in symmetric matrices," *SIAM Journal on Optimization*, vol. 7, no. 1, pp. 86–125, 1997. [Online]. Available: http://dx.doi.org/10.1137/S1052623494269035

[85] K. Toh, R. Tütüncü, and M. Todd, "Inexact primal-dual path-following algorithms for a special class of convex quadratic sdp and related problems," *Pacific Journal of Optimization*, vol. 3, pp. 135–164, 2007.

[86] M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110, http://stanford.edu/~boyd/graph_dcp.html.

[87] ——, "CVX: Matlab software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, June 2015.

[88] G. P. Barker, R. D. Hill, and R. D. Haertel, "On the Completely Positive and Positive-Semidefinite-Preserving Cones," *Linear Algebra and its Applications*, vol. 56, pp. 221–229, 1984.

[89] K. Toh, M. Todd, and R. Tütüncü, "SDPT3 — a Matlab software package for semidefinite programming," *Optimization Methods and Software*, vol. 11, pp. 545–581, 1999.

[90] R. Tütüncü, K. Toh, and M. Todd, "Solving semidefinite-linear-quadratic programs using SDPT3," *Mathematical Programming Ser. B*, vol. 95, pp. 189–217, 2003.

[91] J. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11–12, pp. 625–653, 1999, version 1.05 available from `http://fewcal.kub.nl/sturm`.

[92] X. Li, D. Sun, and K.-C. Toh, "QSDPNAL: A two-phase augmented Lagrangian method for convex quadratic semidefinite programming," `arXiv:1512.08872v2 [math.OC]`, Dec. 2015. [Online]. Available: http://arxiv.org/pdf/1512.08872v2

[93] D. Han, D. Sun, and L. Zhang, "Linear rate convergence of the alternating direction method of multipliers for convex composite quadratic and semidefinite programming," `arXiv:1508.02134v1 [math.OC]`. [Online]. Available: http://arxiv.org/pdf/1508.02134v1

[94] Y. Cui, C. Ding, and X. Zhao, "Quadratic growth conditions for convex matrix optimization problems associated with spectral functions," `arXiv:1702.03262v3 [math.OC]`, February 2017. [Online]. Available: http://arxiv.org/pdf/1702.03262v3

[95] Y. Cui, D. Sun, and K.-C. Toh, "On the asymptotic superlinear convergence of the augmented lagrangian method for semidefinite programming with multiple solutions," `arXiv:1610.00875v1 [math.OC]`, October 2016. [Online]. Available: http://arxiv.org/pdf/1610.00875v1

[96] ——, "On the r-superlinear convergence of the kkt residues generated by the augmented lagrangian method for convex composite conic programming," `arXiv:1706.08800v1 [math.OC]`, June 2017. [Online]. Available: http://arxiv.org/pdf/1706.08800v1

[97] H. Qi, "Local duality of nonlinear semidefinite programming," *Mathematics of Operations Research*, vol. 34, no. 1, pp. 124–141, 2009. [Online]. Available: http://www.jstor.org/stable/40538370

[98] A. Hjørungnes, *Complex-Valued Matrix Derivatives*, 1st ed. Cambridge, UK: Cambridge University Press, 2011.

[99] J. F. Bonnans and A. Shapiro, *Perturbation Analysis of Optimization Problems*. New York, NY: Springer-Verlag Limited, 2000.

[100] E. Kreindler and A. Jameson, "Conditions for nonnegativeness of partitioned matrices," *IEEE Trans. Automat. Contr.*, vol. 17, no. 1, pp. 147–148, Feb. 1972.

[101] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. New York, NY: Cambridge University Press, 2013.

[102] C. G. Khatri and S. K. Mitra, "Hermitian and nonnegative definite solutions of linear matrix equations," *SIAM Journal on Applied Mathematics*, vol. 31, no. 4, pp. 579–585, 12 1976.

[103] J. K. Baksalary, "Nonnegative definite and positive definite solutions to the matrix equation AXA* = B," *Linear and Multilinear Algebra*, vol. 16, no. 1-4, pp. 133–139, 1984. [Online]. Available: http://dx.doi.org/10.1080/03081088408817616

[104] P. Setlur and M. Rangaswamy, "Flawed Waveform Design of Augusto Aubry, Antonio DeMaio, et al." Air Force Research Laboratory, Sensors Directorate, Dayton, OH, Tech. Rep. AFRL-RY-WP-TR-2018-0018, February 2018.

[105] S. M. Kay, *Fundamentals of Statistical Signal Processing, vol. II: Detection Theory*. Upper Saddle River, NJ: Prentice Hall, 1998.

[106] S. M. O'Rourke, P. Setlur, M. Rangaswamy, and A. L. Swindlehurst, "Quadratic Semidefinite Programming for Waveform-Constrained Joint Filter-Signal Design in STAP," *IEEE Transactions on Signal Processing*, 2018, submitted.

[107] S. M. Sussman, "Least-square synthesis of radar ambiguity functions," *IRE Transactions on Information Theory*, vol. 8, no. 3, pp. 246–254, April 1962.

[108] N. Levanon and E. Mozeson, *Radar Signals*. Hoboken, NJ: John Wiley & Sons, Inc., 2004.

[109] H. Neudecker, S. Liu, and W. Polasek, "The hadamard product and some of its applications in statistics," *Statistics*, vol. 26, no. 4, pp. 365–373, 1995. [Online]. Available: https://doi.org/10.1080/02331889508802503

[110] A. De Maio, S. De Nicola, Y. Huang, D. P. Palomar, S. Zhang, and A. Farina, "Code design for radar stap via optimization theory," *IEEE Transactions on Signal Processing*, vol. 58, no. 2, pp. 679–694, Feb 2010.

[111] J. Sun and S. Zhang, "A modified alternating direction method for convex quadratically constrained quadratic semidefinite programs," *European Journal of Operational Research*, vol. 207, no. 3, pp. 1210 – 1220, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0377221710005308

Appendix A

Proof of Theorem 3.6

In this appendix, we provide an abbreviated proof of Theorem 3.6. In order to prove that the Slater Condition is satisfied, we first find the generic matrix that satisfies the equality constraints, then ensure the resulting matrix is both positive definite and strictly satisfies the power inequality. As a combined matrix-vector equation, the complex equality constraints are given by

$$\begin{bmatrix} \tilde{\boldsymbol{\tau}}^T \mathbf{K}_{J,J} \\ \tilde{\boldsymbol{\tau}}^H \end{bmatrix} \boldsymbol{\beta} = \mathbf{E}\boldsymbol{\beta} = \begin{bmatrix} \kappa \\ \kappa^* \end{bmatrix} \tag{A.1}$$

where $\mathbf{E}$ is implicitly defined.

We first demonstrate that a generic matrix solution exists to this equation and provide its form. Since $\mathbf{E}$ is a fat matrix, its pseudoinverse is $\mathbf{E}^\dagger = \mathbf{E}^H(\mathbf{E}\mathbf{E}^H)^\dagger$. It can be shown that the Gramian matrix above is

$$\begin{bmatrix} \tilde{\boldsymbol{\tau}}^T \mathbf{K}_{J,J} \\ \tilde{\boldsymbol{\tau}}^H \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{\tau}}^T \mathbf{K}_{J,J} \\ \tilde{\boldsymbol{\tau}}^H \end{bmatrix}^H = \|\mathbf{T}\|_F^2 \mathbf{I}_2.$$

So long as we have a non-zero target matrix, this is always invertible. Furthermore, this implies that $\mathbf{P_E} = \mathbf{E}\mathbf{E}^\dagger = \mathbf{I}_2$ and hence a solution to (A.1) always exists. Therefore, a general solution to the equality constraints is the matrix

$$\mathbf{B} = \mathbf{Z} + \frac{1}{\|\mathbf{T}\|_F^2} \begin{bmatrix} \mathbf{0}_{NML\times NML} & (\kappa^* - \mathrm{tr}(\mathbf{T}^H\mathbf{Z}_2))\mathbf{T} \\ (\kappa - \mathrm{tr}(\mathbf{Z}_2^H\mathbf{T}))\mathbf{T}^H & \mathbf{0}_{N\times N} \end{bmatrix}.$$

where $\mathbf{Z}$ is a Hermitian matrix of identical dimension and partitioning to $\mathbf{B}$. This additional matrix is necessary because the min-norm solution obtained by setting $\mathbf{Z} = \mathbf{0}_{J\times J}$,

$$\mathbf{B}_{min} = \frac{1}{\|\mathbf{T}\|_F^2} \begin{bmatrix} \mathbf{0}_{NML\times NML} & \kappa^*\mathbf{T} \\ \kappa\mathbf{T}^H & \mathbf{0}_{N\times N} \end{bmatrix},$$

is an indefinite matrix, and hence not a Slater point. Hence, it is now sufficient to prove there exists a $\mathbf{Z}$ such that $\mathbf{B} \succ 0$ and $\mathrm{tr}(\mathbf{Z}_3) < P_o$. We make the judicious choice that $\mathbf{Z}_2 = \mathbf{0}_{NML\times N}$. Under this assumption, our solution matrix is

$$\mathbf{B} = \begin{bmatrix} \mathbf{Z}_1 & \frac{\kappa^*}{\|\mathbf{T}\|_F^2}\mathbf{T} \\ \frac{\kappa}{\|\mathbf{T}\|_F^2}\mathbf{T}^H & \mathbf{Z}_3 \end{bmatrix}.$$

The positive definiteness requirement can then be expressed as the matrix inequalities

$$\mathbf{Z}_1 \succ 0, \quad \mathbf{Z}_3 - \frac{|\kappa|^2}{\|\mathbf{T}\|_F^4}\mathbf{T}^H\mathbf{Z}_1^{-1}\mathbf{T} \succ 0$$

Here, we rely on another judicious choice, setting $\mathbf{Z}_1 = \mathbf{I}_{NML}$, which is clearly positive definite.[1]Now, it is merely sufficient to prove there exists $\mathbf{Z}_3$ such that $\mathbf{Z}_3 \succ \frac{|\kappa|^2}{\|\mathbf{T}\|_F^4}\mathbf{T}^H\mathbf{T}$ *and* $\mathrm{tr}(\mathbf{Z}_3) < P_o$. Let us assume that such a matrix exists. Following [101, Corollary 7.7.4(d)], since $\mathbf{Z}_3 \succ \frac{|\kappa|^2}{\|\mathbf{T}\|_F^2}\mathbf{T}^H\mathbf{T}$, then $\mathrm{tr}(\mathbf{Z}_3) > \frac{|\kappa|^2}{\|\mathbf{T}\|_F^4}\mathrm{tr}(\mathbf{T}^H\mathbf{T}) = \frac{|\kappa|^2}{\|\mathbf{T}\|_F^2}$. But we already know that

[1]We can make this choice because $NML > N$. If, for whatever reason, the number of transmit resources were greater than the number of receive resources, then we could start with $\mathbf{Z}_3$ and continue from there.

$\text{tr}(\mathbf{Z}_3) < P_o$. Hence, we have the chained inequality

$$P_o > \text{tr}(\mathbf{Z}_3) > \frac{|\kappa|^2}{\|\mathbf{T}\|_F^2}.$$

It therefore follows that there will exist such a matrix $\mathbf{Z}_3$ if and only if $P_o > \frac{|\kappa|^2}{\|\mathbf{T}\|_F^2}$. Hence, we have found the sufficent condition of Theorem 3.6 directly and the proof is complete. We note that a sufficently good choice is $\mathbf{Z}_3 = \frac{\alpha}{N}\mathbf{I}_N$, with α in the aforementioned interval. ■

Appendix B

Further Details of the Primal KKTs of the Power-Constrained RBQP

In this appendix, we summmarize the derivation of the KKT conditions alluded to in Chapter 3.2.2. Much of this material has been reproduced from [65, 66].

B.1 KKT Condition 1: Stationarity of Lagrangian

Let $\mathbf{\Sigma}$ be the slackness variable associated with the PSD condition on $\mathbf{B}$, and let $\boldsymbol{\sigma}$ be its vectorization. Furthermore, let $\boldsymbol{\rho} = \text{vec}(\widetilde{\mathbf{R}}_{\mathbf{ni}})$ and $\boldsymbol{\psi}_S = \text{vec}(\mathbf{\Psi}_S^T \mathbf{\Psi}_S)$. When vectorized, the Lagrangian of (3.24) under complex constraints is

$$\mathcal{L}(\mathbf{B}, \mathbf{\Sigma}, \tilde{\mu}_C, \lambda_P) = \boldsymbol{\beta}^H (\widetilde{\mathbf{C}}_V \boldsymbol{\beta} + \boldsymbol{\rho} - \tilde{\mu}_C^* \tilde{\boldsymbol{\tau}} + \lambda_P \boldsymbol{\psi}_S - \boldsymbol{\sigma}) + \tilde{\mu}_C^* \kappa - \lambda_P P_o. \tag{B.1}$$

After taking the appropriate derivative, KKT Condition 1 is

$$(\widetilde{\mathbf{C}}_V + \mathbf{K}_{J,J}\widetilde{\mathbf{C}}_V^*\mathbf{K}_{J,J})\boldsymbol{\beta} = \boldsymbol{\sigma} - (\boldsymbol{\rho} + \tilde{\mu}_C^*\tilde{\boldsymbol{\tau}} + \tilde{\mu}_C\tilde{\boldsymbol{\tau}}_H + \lambda_P\boldsymbol{\psi}_S). \tag{B.2}$$

where $\tilde{\boldsymbol{\tau}}_H = \text{vec}(\widetilde{\mathbf{T}}^H)$. Given the four disjoint projection matrices $\mathbf{P}_{WW}, \mathbf{P}_{SW}, \mathbf{P}_{WS}, \mathbf{P}_{SS}$ defined in Chapter 3.1.1, we can decompose the vectorizations of $\mathbf{B}$ and $\boldsymbol{\Sigma}$ into sums of vectorizations of each partition. For example,

$$\boldsymbol{\beta} = \mathbf{P}_{WW}^T\boldsymbol{\beta}_1 + \mathbf{P}_{SW}^T\boldsymbol{\beta}_2 + \mathbf{P}_{WS}^T\boldsymbol{\beta}_{2,H} + \mathbf{P}_{SS}^T\boldsymbol{\beta}_3$$

where the subscript indicates which submatrix is vectorized, and $\boldsymbol{\beta}_{2,H} = \text{vec}(\mathbf{B}_2^H)$. It can be shown [65], then, that (B.2) implies (3.31) so long as

$$\mathbf{C}_V\boldsymbol{\beta}_2 = \boldsymbol{\sigma}_2 + \tilde{\mu}_C^*\boldsymbol{\tau} \tag{B.3}$$

whose solution exists if and only if $\mathbf{P}_{\mathbf{C}_V}^{\perp}\boldsymbol{\sigma}_2 = -\tilde{\mu}_C^*\mathbf{P}_{\mathbf{C}_V}^{\perp}\boldsymbol{\tau}$ where $\mathbf{P}_{\mathbf{C}_V}^{\perp} = \mathbf{I}_{N^2ML} - \mathbf{C}_V\mathbf{C}_V^{\dagger}$ is the orthogonal projector onto the nullspace of $\mathbf{C}_V$. Additionally, the matrices $\mathbf{B}_1$ and $\mathbf{B}_3$ are "free" parameters under (B.2), but they are constrained by later conditions.

B.2 KKT Conditions 2-4: The Inequality Constraints

With the gradient condition exhausted, we turn to the inequality constraints (power bound, positive-semidefiniteness of the solution) and their related conditions. For convenience, we shall attack these somewhat independently in separate subsections, though they will interact.

B.2.1 The power constraint

Using the partition mentioned above, the conditions related to the power constraint are:

$$\lambda_P \geq 0 \tag{B.4}$$

$$\mathrm{Tr}(\mathbf{B}_3) - P_o \leq 0 \tag{B.5}$$

$$\lambda_P(\mathrm{Tr}(\mathbf{B}_3) - P_o) = 0. \tag{B.6}$$

A nice result of the slackness condition is $\lambda_P \,\mathrm{Tr}(\mathbf{B}_3) = \lambda_P P_o$, which we will use later. The other "result" is that $\lambda_P = 0$ when the solution does not reach the power bound, and $\lambda_P > 0$ when it does. This fact will inform interpretations of the solution found in Section 3.3.

B.2.2 Positive semidefiniteness of B

The conditions for semidefiniteness of the relaxed beamformer-signal basis are slightly more complex, but reveal a significant amount of structure to the solution. First, the direct form of these conditions are

$$\boldsymbol{\Sigma} \succeq 0, \quad \mathbf{B} \succeq 0, \quad \boldsymbol{\Sigma}\mathbf{B} = \mathbf{0}_{J\times J}.$$

Of course, in this form, they are not especially useful. However, given (3.31), we know the structure of $\boldsymbol{\Sigma}$ somewhat. The most useful characterization of semidefiniteness here is that $\boldsymbol{\Sigma}$ is PSD iff a contraction $\mathbf{X} \in \mathbb{C}^{NML\times N}$ exists such that $\boldsymbol{\Sigma}_2 = \sqrt{\lambda_P}\mathbf{R}_{\mathbf{ni}}^{1/2}\mathbf{X}$ (see [101, Theorem 7.7.9(a,b)]).

The matrix complementary slackness condition reduces to the following 4 equalities:

$$\mathbf{B}_1\mathbf{R}_{\mathbf{ni}} = -\mathbf{B}_2\boldsymbol{\Sigma}_2^H \tag{B.7}$$

$$\mathbf{B}_1\boldsymbol{\Sigma}_2 = -\lambda_P\mathbf{B}_2 \tag{B.8}$$

$$\mathbf{B}_2^H\mathbf{R}_{\mathbf{ni}} = -\mathbf{B}_3\boldsymbol{\Sigma}_2^H \tag{B.9}$$

$$\mathbf{B}_2^H\boldsymbol{\Sigma}_2 = -\lambda_P\mathbf{B}_3. \tag{B.10}$$

These conditions have significant consequences in Section 3.3, but they also result in equivalent forms for λ_P[1]. First, taking the trace of (B.10) and substituting in (B.6) on the right hand side, we have

$$\lambda_P = -\frac{\mathrm{tr}(\mathbf{B}_2^H\boldsymbol{\Sigma}_2)}{P_o}.$$

Since λ_P is both real and nonnegative, this means that $\mathrm{tr}(\mathbf{B}_2^H\boldsymbol{\Sigma}_2)$ is real and nonpositive. We can apply this to the trace of (B.7) to obtain another form, $\lambda_P = \frac{\mathrm{tr}(\mathbf{B}_1\mathbf{R}_{\mathbf{ni}})}{P_o}$.

We can also apply this logic to (B.3) to reveal an interesting consequence about the cost function. If premultiplied by $\boldsymbol{\beta}_2$ and rearranged,

$$\boldsymbol{\beta}_2^H\boldsymbol{\sigma}_2 = \boldsymbol{\beta}_2^H\mathbf{C}_V\boldsymbol{\beta}_2 - \tilde{\mu}_C^*\boldsymbol{\beta}_2^H\boldsymbol{\tau}$$

Applying the equality constraint, this is equivalent to

$$\boldsymbol{\beta}_2^H\boldsymbol{\sigma}_2 = \boldsymbol{\beta}_2^H\mathbf{C}_V\boldsymbol{\beta}_2 - \tilde{\mu}_C^*\kappa.$$

Incidentally, this implies $\tilde{\mu}_C^*\kappa$ is real, and hence the optimal phase of $\tilde{\mu}_C$ is that of κ.

[1] As we see in Chapter 4, conditions (B.7) and (B.9) are in fact general for any set of signal constraints.

From the above, however, we can see that $-\boldsymbol{\beta}_2^H \boldsymbol{\sigma}_2 = \lambda_P P_o = \mathrm{tr}(\mathbf{B}_1 \mathbf{R}_{\mathbf{ni}})$, and so

$$\tilde{\mu}_C^* \kappa = \boldsymbol{\beta}_2^H \mathbf{C}_V \boldsymbol{\beta}_2 + \mathrm{tr}(\mathbf{B}_1 \mathbf{R}_{\mathbf{ni}}). \tag{B.11}$$

The right hand side of (B.11) is immediately recognizable as our objective function, which implies a potential future analysis on the dual problem that is equivalent to our findings here. This is where we obtain that the optimal value of the RBQP (3.24) is wholly described by the dual variable $\tilde{\mu}_C$.

B.3 KKT Condition 5: Equality Constraints

This is the final major KKT condition left to examine, because KKT Condition 6 is trivially satisfied by $\mathbf{C}_V$ being positive semidefinite. Here, our primary concern is the equality constraint $\boldsymbol{\beta}_2^H \boldsymbol{\tau} = \kappa$. According to the first KKT condition, we know that

$$\boldsymbol{\beta}_2 = \mathbf{C}_V^{\dagger}(\boldsymbol{\sigma}_2 + \tilde{\mu}_C^* \boldsymbol{\tau}) + \mathbf{P}_{\mathbf{C}_V}^{\perp} \mathbf{z}_2. \tag{B.12}$$

for some vector $\mathbf{z}_2 \in \mathbb{C}^{N^2ML}$. Substituting this into the equality constraint gives us

$$\boldsymbol{\sigma}_2^H \mathbf{C}_V^{\dagger} \boldsymbol{\tau} + \tilde{\mu}_C \boldsymbol{\tau}^H \mathbf{C}_V^{\dagger} \boldsymbol{\tau} + \mathbf{z}_2^H \mathbf{P}_{\mathbf{C}_V}^{\perp} \boldsymbol{\tau} = \kappa. \tag{B.13}$$

More importantly, we recognize that the columns of $\mathbf{B}_2$ and $\mathbf{T}$ must align for there to be a non-trivial feasible solution.

Appendix C

Proof of Theorem 4.1

In this Appendix, we provide an abbreviated proof of Theorem 4.1.

We recognize that we must first solve the simultaneous equalities in (4.13). If we put them in terms of the joint vectorized matrix $\boldsymbol{\beta}$ and vectorized target matrix $\tilde{\boldsymbol{\tau}} = \text{vec}(\widetilde{\mathbf{T}})$, then the set of equations is

$$\mathbf{Y}\boldsymbol{\beta} = \begin{bmatrix} \rho^2 \mathbf{1}_N^T & \kappa & \kappa^* \end{bmatrix}^T \tag{C.1}$$

where $\mathbf{Y} = \begin{bmatrix} \mathbf{P}_{SS}^T \mathbf{L}_{d,N} & \mathbf{K}_{J,J}\tilde{\boldsymbol{\tau}} & \tilde{\boldsymbol{\tau}}^* \end{bmatrix}^T$. So long as there is a target constraint, the pseudoinverse $\mathbf{Y}^\dagger = \mathbf{Y}^H(\mathbf{Y}\mathbf{Y}^H)^{-1}$ is given by

$$\mathbf{Y}^\dagger = \begin{bmatrix} \mathbf{P}_{SS}^T \mathbf{L}_{d,N} & \frac{1}{\|\mathbf{T}\|_F^2}\mathbf{K}_{J,J}\tilde{\boldsymbol{\tau}}^* & \frac{1}{\|\mathbf{T}\|_F^2}\tilde{\boldsymbol{\tau}} \end{bmatrix} \tag{C.2}$$

Clearly, the projector onto the columns of $\mathbf{Y}$ is the identity and there is always a general solution to the equality constraints.

The "min norm" part of the solution, when reshaped into a matrix, is given by:

$$\mathbf{B}_{min} = \begin{bmatrix} \mathbf{0}_{NML \times NML} & \frac{\kappa^*}{\|\mathbf{T}\|_F^2}\mathbf{T} \\ \frac{\kappa}{\|\mathbf{T}\|_F^2}\mathbf{T}^H & \rho^2 \mathbf{I}_N \end{bmatrix} \tag{C.3}$$

Clearly, this can never be positive definite, so we need the "nullspace" element of the general solution to achieve our interior point. With this component added, it can be shown that the full general Hermitian solution to the equality constraints is therefore

$$\mathbf{B} = \begin{bmatrix} \mathbf{Z}_1 & \mathbf{Z}_2 + \frac{\kappa^* - \mathrm{tr}(\mathbf{T}^H \mathbf{Z}_2)}{\|\mathbf{T}\|_F^2}\mathbf{T} \\ (\mathbf{Z}_2 + \frac{\kappa^* - \mathrm{tr}(\mathbf{T}^H \mathbf{Z}_2)}{\|\mathbf{T}\|_F^2}\mathbf{T})^H & \rho^2 \mathbf{I}_N + \mathbf{Z}_3 \end{bmatrix} \tag{C.4}$$

where $\mathbf{Z}_3$ has an all-zero diagonal. In order to find a positive definite solution, we could make a judicious guess and set $\mathbf{Z}_2 = \mathbf{0}_{NML \times N}$ and $\mathbf{Z}_3 = \mathbf{0}_{N \times N}$. This means the potential solution $\mathbf{B}_{guess}$ is

$$\mathbf{B}_{guess} = \begin{bmatrix} \mathbf{Z}_1 & \frac{\kappa^*}{\|\mathbf{T}\|_F^2}\mathbf{T} \\ \frac{\kappa}{\|\mathbf{T}\|_F^2}\mathbf{T}^H & \rho^2 \mathbf{I}_N \end{bmatrix}. \tag{C.5}$$

As in our proof of Theorem 1 of [66], we can again judiciously guess $\mathbf{Z}_1 = \mathbf{I}_{NML}$. This ensures that the only necessary condition for an interior point to exist is

$$\rho^2 \mathbf{I}_N - \frac{|\kappa|^2}{\|\mathbf{T}\|_F^2}\mathbf{T}^H \mathbf{T} \succ 0. \tag{C.6}$$

This condition is satisfied if and only if the above matrix is a strict contraction, which implies

$$\frac{|\kappa|^2}{\|\mathbf{T}\|_F^2} < \rho^2 \frac{\|\mathbf{T}\|_F^2}{\|\mathbf{T}\|_2^2} \tag{C.7}$$

where $\|\mathbf{T}\|_2$ is the spectral norm, or largest singular value, of $\mathbf{T}$. It is well-known (see, e.g., [101]) that $\|\mathbf{T}\|_F^2 \leq N \|\mathbf{T}\|_2^2$. Applying this to (C.7), a necessary condition for the existence

of an interior point is therefore

$$\frac{|\kappa|^2}{\|\mathbf{T}\|_F^2} < N\rho^2 \tag{C.8}$$

which completes the proof. ■