

Designing Structured Tight Frames via an Alternating Projection Method

Joel A. Tropp, *Student Member, IEEE*, Inderjit S. Dhillon, *Member, IEEE*,
Robert W. Heath Jr., *Member, IEEE* and Thomas Strohmer

Abstract—Tight frames, also known as general Welch-Bound-Equality sequences, generalize orthonormal systems. Numerous applications—including communications, coding, and sparse approximation—require finite-dimensional tight frames that possess additional structural properties. This paper proposes an alternating projection method that is versatile enough to solve a huge class of inverse eigenvalue problems, which includes the frame design problem. To apply this method, one only needs to solve a matrix nearness problem that arises naturally from the design specifications. Therefore, it is fast and easy to develop versions of the algorithm that target new design problems. Alternating projection will often succeed even if algebraic constructions are unavailable.

To demonstrate that alternating projection is an effective tool for frame design, the article studies some important structural properties in detail. First, it addresses the most basic design problem: constructing tight frames with prescribed vector norms. Then, it discusses equiangular tight frames, which are natural dictionaries for sparse approximation. Last, it examines tight frames whose individual vectors have low peak-to-average-power ratio (PAR), which is a valuable property for CDMA applications. Numerical experiments show that the proposed algorithm succeeds in each of these three cases. The appendices investigate the convergence properties of the algorithm.

Index Terms—Algorithms, code division multiaccess, eigenvalues and eigenfunctions, extremal problems, frames, geometry, inverse problems, sequences

I. INTRODUCTION

TIGHT FRAMES provide a natural generalization of orthonormal systems, and they arise in numerous practical and theoretical contexts [1]. There is no shortage of tight frames, and applications will generally require that the vectors comprising the frame have some kind of additional structure. For example, it might be necessary for the vectors to have specific Euclidean norms, or perhaps they should have small

mutual inner products. Thus arises a design problem: How do you build a structured tight frame?

A. Contributions

To address the design question, this article proposes a numerical method based on alternating projection that builds on our work in [2], [3]. The algorithm alternately finds the nearest tight frame to a given ensemble of structured vectors; then it finds the ensemble of structured vectors nearest to the tight frame; and it repeats the process *ad infinitum*. This technique is analogous to the method of projection on convex sets (POCS) [4], [5], except that the class of tight frames is nonconvex, which complicates the analysis significantly. Nevertheless, our alternating projection algorithm affords simple implementations, and it provides a quick route to solve difficult frame design problems. We argue that similar techniques apply to a huge class of inverse eigenvalue problems.

This article demonstrates the elegance and effectiveness of the alternating projection approach with several examples that are motivated by applications. First, we address the most basic frame design problem: building tight frames with prescribed vector norms. This problem arises when constructing signature signatures for direct-spread, synchronous code division multiaccess (DS-CDMA) systems [6]–[8]. Second, we discuss equiangular tight frames, which have the property that each pair of distinct vectors meets at the same (acute) angle. These frames have many applications in coding and communications [9]–[12], and they also form natural dictionaries for sparse approximation [13]–[15]. Third, we examine tight frames whose individual vectors have low peak-to-average-power ratio (PAR), which is another valuable property for DS-CDMA signatures [3]. Our experiments show that alternating projection outperforms some algorithms that were specifically designed to solve these problems.

The appendices investigate the convergence properties of the algorithm. Although alternating projection between subspaces and convex sets has been studied in detail, very few results are available for an alternating projection between two nonconvex sets. This paper provides a rigorous treatment of the algorithm's behavior in this general setting by means of the theory of point-to-set maps. In particular, we establish a weak global convergence result, and we show that, under additional hypotheses, the algorithm exhibits stronger local convergence properties.

Note that there is a major conceptual difference between the use of finite models in the numerical calculation of infinite-dimensional frames and the design of finite-dimensional

Typeset on October 8, 2004.

J. A. Tropp was with the Inst. for Computational Engineering and Sciences, The University of Texas, Austin, TX 78705. He is currently with the Dept. of Mathematics, The University of Michigan, Ann Arbor, MI 48104 USA, <jtropp@umich.edu>. He was supported by an NSF Graduate Fellowship.

I. S. Dhillon is with the Dept. of Computer Sciences, The University of Texas, Austin, TX 78712 USA, <inderjit@cs.utexas.edu>. He acknowledges support by NSF CAREER Award No. ACI-0093404.

R. W. Heath is with the Dept. of Electrical and Computer Engineering, The University of Texas, Austin, TX 78712 USA, <rheath@ece.utexas.edu>. He acknowledges support by the Texas Advanced Technology Program under Grant No. 003658-0614-2001, the Samsung Advanced Institute of Technology, and the National Instruments Foundation.

T. Strohmer is with the Dept. of Mathematics, The University of California, Davis, CA 95616 USA, <strohmer@math.ucdavis.edu>. He was supported in part by NSF grant DMS 0208568.

frames. In the former case, the finite model plays only an auxiliary role in the approximate computation of an infinite-dimensional tight frame [1]. In the latter case, the problem under consideration is already finite-dimensional, thus it does not involve discretization issues. In this paper, we consider only finite-dimensional tight frames.

B. Previous Work

At root, finite-dimensional frame design is an algebraic problem. It boils down to producing a structured matrix with certain *spectral* properties, which may require elaborate discrete and combinatorial mathematics. In the past, most design methods have employed these techniques. To appreciate the breadth of this literature, one might peruse Sarwate's recent survey paper about tight frames comprised of unit-norm vectors [16]. The last few years have also seen some essentially algebraic algorithms that can construct tight frames with nonconstant vector norms [7], [17], [18].

When algebraic methods work, they work brilliantly. A numerical approach like alternating projection can hardly hope to compete with the most profound insights of engineers and mathematicians. On the other hand, algebraic and combinatoric tools are not always effective. For example, we might require a structured tight frame for a vector space whose dimension is not a prime-power. Even in these situations, alternating projection will often succeed. Moreover, it can help researchers develop the insight necessary for completing an algebraic construction. The power of alternating projection comes from replacing the difficult algebra with a simple analytic question: How does one find an ensemble of structured vectors nearest to a given tight frame? This minimization problem can usually be dispatched with standard tools, such as differential calculus or Karush–Kuhn–Tucker theory.

The literature does not offer many other numerical approaches to frame design. It appears that most current algorithms can be traced to the discovery by Rupf–Massey [6] and Viswanath–Anantharam [7] that tight frames with prescribed column norms are the optimal sequences for DS-CDMA systems. This application prompted a long series of papers, including [19]–[23], that describe iterative methods for constructing tight frames with prescribed column norms. These techniques are founded on an oblique characterization of tight frames as the minimizers of a quantity called *total squared correlation* (TSC). It is not clear how one could generalize these methods to solve different types of frame design problems. Moreover, the alternating projection approach that we propose significantly outperforms at least one of the TSC-minimization algorithms. Two of the algebraic methods that we mentioned above, [7] and [18], were also designed with the DS-CDMA application in mind, while the third algebraic method [17] comes from the *soi-disant* frame community. We are not aware of any other numerical methods for frame design.

C. Outline

Section II continues with a short introduction to tight frames. Then, in Section III, we state two formal frame design problems. Connections among frame design problems,

inverse eigenvalue problems, and matrix nearness problems are established. This provides a natural segue to the alternating projection algorithm. Afterward, we apply the basic algorithm to design three different types of structured frames, in order of increasing implementation difficulty. Section IV discusses tight frames with prescribed column norms; Section V covers equiangular tight frames; and Section VI constructs tight frames whose individual vectors have low peak-to-average-power ratio. Each of these sections contains numerical experiments. The body of the paper concludes with Section VII, which discusses the method, its limitations, and its extensions.

The back matter contains the bulk of the analysis. Appendix I offers a tutorial on point-to-set maps, and Appendix II applies this theory to obtain a rigorous characterization of the algorithm's convergence behavior. The first appendix also contains a brief survey of the alternating projection literature.

II. TIGHT FRAMES

This section offers a terse introduction to the properties of tight frames that are essential for our method. For more details, see [1], for example.

A. Frames

Let α and β be positive constants. A *finite frame* for the complex¹ Hilbert space \mathbb{C}^d is a sequence of N vectors $\{\mathbf{x}_n\}_{n=1}^N$ drawn from \mathbb{C}^d that satisfies a generalized Parseval condition:

$$\alpha \|\mathbf{v}\|_2^2 \leq \sum_{n=1}^N |\langle \mathbf{v}, \mathbf{x}_n \rangle|^2 \leq \beta \|\mathbf{v}\|_2^2 \quad \text{for all } \mathbf{v} \in \mathbb{C}^d. \quad (1)$$

We denote the Euclidean inner product with $\langle \cdot, \cdot \rangle$, and we write $\|\cdot\|_2$ for the associated norm. The numbers α and β are called the *lower* and *upper frame bounds*. The number of vectors in the frame may be no smaller than the dimension of the space (i.e., $N \geq d$).

If it is possible to take $\alpha = \beta$, then we have a *tight frame* or an α -*tight frame*. When the frame vectors all have unit norm, i.e., $\|\mathbf{x}_n\|_2 \equiv 1$, the system is called a *unit-norm frame*. Unit-norm tight frames are also known as *Welch-Bound-Equality sequences* [12], [24]. Tight frames with nonconstant vector norms have also been called *general Welch-Bound-Equality sequences* [7].

B. Associated Matrices

Form a $d \times N$ matrix with the frame vectors as its columns:

$$X = [\mathbf{x}_1 \quad \mathbf{x}_2 \quad \mathbf{x}_3 \quad \dots \quad \mathbf{x}_N].$$

This matrix is referred to as the *frame synthesis operator*, but we will usually identify the synthesis operator with the frame itself.

Two other matrices arise naturally in connection with the frame. We define the *Gram matrix* as $G \stackrel{\text{def}}{=} X^* X$. (The symbol $*$ indicates conjugate transposition of matrices and vectors.)

¹We work with complex vectors for complete generality. The adaptations for real vectors are transparent.

The diagonal entries of the Gram matrix equal the squared norms of the frame vectors, and the off-diagonal entries of the Gram matrix equal the inner products between distinct frame vectors. The Gram matrix is always Hermitian and positive semi-definite, and it has rank d .

The positive-definite matrix XX^* is usually called the *frame operator*. Since

$$\mathbf{v}^*(XX^*)\mathbf{v} = \sum_{n=1}^N |\langle \mathbf{v}, \mathbf{x}_n \rangle|^2$$

we can rewrite (1) as

$$\alpha \leq \frac{\mathbf{v}^*(XX^*)\mathbf{v}}{\mathbf{v}^*\mathbf{v}} \leq \beta. \quad (2)$$

That is, any Rayleigh quotient of XX^* must lie between α and β . It follows from the Rayleigh–Ritz Theorem [25, p. 176] that each eigenvalue of the frame operator lies in the interval $[\alpha, \beta]$.

When the frame is α -tight, the condition (2) is equivalent to the statement that $XX^* = \alpha \mathbf{I}_d$. Three other characterizations of an α -tight frame follow immediately.

Proposition 1: A $d \times N$ frame X is α -tight if and only if it satisfies one (hence all) of the following conditions.

- 1) All d nonzero singular values of X equal $\sqrt{\alpha}$.
- 2) All d nonzero eigenvalues of the Gram matrix X^*X equal α .
- 3) The rows of $\alpha^{-1/2}X$ form an orthonormal family.

These properties undergird our method for constructing tight frames. It is now clear that the being a tight frame is a *spectral* requirement on the matrix X .

C. Norms of Frame Vectors

Throughout this article, we will denote the squared norms of the frame vectors as

$$c_n \stackrel{\text{def}}{=} \|\mathbf{x}_n\|_2^2.$$

There is an intimate relationship between the tightness parameter of an α -tight frame and the norms of its columns. The computation is straightforward:

$$\alpha d = \text{Tr } XX^* = \text{Tr } X^*X = \sum_{n=1}^N \|\mathbf{x}_n\|_2^2 = \sum_{n=1}^N c_n. \quad (3)$$

The notation $\text{Tr}(\cdot)$ represents the matrix trace operator, which sums the diagonal entries of its argument.

A related point is that one cannot construct a tight frame with an arbitrary set of column norms. According to the Schur–Horn Theorem, a Hermitian matrix can exist if and only if its diagonal entries majorize² its eigenvalues [25, pp. 193–197]. If X is a $d \times N$ tight frame, the d nonzero eigenvalues of its Gram matrix all equal $\sum_n c_n/d$. Meanwhile, the diagonal entries of the Gram matrix are c_1, \dots, c_N . In this case, the majorization condition is equivalent to the system of inequalities

$$0 \leq c_k \leq \frac{1}{d} \sum_{n=1}^N c_n \quad \text{for each } k = 1, \dots, N. \quad (4)$$

²The literature equivocates about the direction of the majorization relation. We adopt the sense used by Horn and Johnson [25, p. 192].

It follows that a tight frame with squared column norms c_1, \dots, c_N exists if and only if (4) holds. For an arbitrary set of column norms, the frames that are “closest” to being tight have been characterized in [7], [26].

III. DESIGN VIA ALTERNATING PROJECTIONS

This section begins with formal statements of two frame design problems. Next, we establish a connection with inverse eigenvalue problems. It becomes clear that an alternating projection algorithm offers a simple and natural approach to general inverse eigenvalue problems, including both frame design problems. We then solve the matrix nearness problems that arise when implementing the proposed algorithm. The section concludes with a discussion of the algorithm’s convergence properties.

A. Structured Tight Frames

Define the collection of $d \times N$ α -tight frames:

$$\mathcal{X}_\alpha \stackrel{\text{def}}{=} \{X \in \mathbb{C}^{d \times N} : XX^* = \alpha \mathbf{I}_d\}. \quad (5)$$

We fix the tightness parameter α for simplicity. It is easy to extend our methods to situations where the tightness is not predetermined, and one can apply similar ideas to construct frames with prescribed upper and lower frame bounds, viz. the parameters α and β in (1). It is worth noting that \mathcal{X}_α is essentially a Stiefel manifold, which consists of all sets of d orthonormal vectors in \mathbb{C}^N [27].

Let \mathcal{S} denote a closed³ collection of $d \times N$ matrices that possess some structural property. In the sequel, we will explore several different structural constraints that have arisen in electrical engineering contexts. Section IV considers tight frames with specified column norms, and Section VI shows how to construct tight frames whose individual vectors have a low peak-to-average-power ratio. Many other structural properties are possible.

Each constraint set \mathcal{S} raises a basic question.

Problem 1: Find a matrix in \mathcal{S} that is minimally distant from \mathcal{X}_α with respect to a given norm.

If the two sets intersect, any solution to this problem is a structured tight frame. Otherwise, the problem requests a structured matrix that is “most nearly tight.” A symmetric problem is to find a tight frame that is “most nearly structured.”

B. Structured Gram Matrices

If the structural constraints restrict the inner products between frame vectors, it may be more natural to work with Gram matrices. Define a collection that contains the Gram matrices of all $d \times N$ α -tight frames:

$$\mathcal{G}_\alpha \stackrel{\text{def}}{=} \{G \in \mathbb{C}^{N \times N} : G = G^* \text{ and } G \text{ has eigenvalues } (\underbrace{\alpha, \dots, \alpha}_d, 0, \dots, 0)\}. \quad (6)$$

³We equip $\mathbb{C}^{d \times N}$ and $\mathbb{C}^{N \times N}$ with the topology induced by the Frobenius norm, which is identical with every other norm topology [25, p. 273].

The set \mathcal{G}_α is essentially a Grassmannian manifold, which consists of d -dimensional subspaces of \mathbb{C}^N [28]. One may also identify the matrices in \mathcal{G}_α as rank- d orthogonal projectors, scaled by α . (An orthogonal projector can be defined as an idempotent, Hermitian matrix. The rank of a projector equals the dimension of its range.)

Let \mathcal{H} be a closed collection of $N \times N$ Hermitian matrices that possess some structural property. In Section V, for example, we will consider equiangular tight frames. The Gram matrices of these frames have off-diagonal entries with identical moduli, and it is an important challenge to construct them.

Once again, a fundamental question arises.

Problem 2: Find a matrix in \mathcal{G}_α that is minimally distant from \mathcal{H} with respect to a given norm.

If the two sets intersect, any solution to this problem will lie in the intersection. Otherwise, the problem requests a tight frame whose Gram matrix is “most nearly structured.” We do not mention the problem of producing a matrix in \mathcal{H} that is nearest to \mathcal{G}_α because it is not generally possible to factor a matrix in \mathcal{H} to obtain a frame with dimensions $d \times N$.

C. Inverse Eigenvalue Problems

We view Problems 1 and 2 as *inverse eigenvalue problems* (IEPs). As Chu explains in [29], an IEP is an inquiry about structured matrices with prescribed spectral properties. These spectral properties may include restrictions on eigenvalues, eigenvectors, singular values, or singular vectors. According to Proposition 1, the defining characteristic of a tight frame is its spectrum, so frame design is an IEP.

In the study of IEPs, the two fundamental questions are *solvability* and *computability*. The former problem is to find necessary or sufficient conditions under which a given IEP has a solution. The latter problem is how to produce a matrix that has given spectral properties and simultaneously satisfies a structural constraint. The solvability and computability of some classical IEPs have been studied extensively by the matrix analysis community, although many open problems still remain. The articles [29], [30] survey this literature.

Although specific IEPs may require carefully tailored numerical methods, there are a few general tools available. One scheme is the coordinate-free Newton method, which has been explored in [31]–[33]. Newton-type algorithms do not apply to all problems, and they only converge locally. Another general method is the projected gradient approach developed by Chu and Driessel in [34]. This technique involves numerical integration of a matrix differential equation, which relies on advanced software that may not be readily available. Another problem with Newton methods and projected gradient methods is that they may not handle repeated singular values well. This shortcoming makes them a poor candidate for constructing tight frames, which have only two distinct singular values.

This article concentrates on alternating projection, which has occasionally been used to solve inverse eigenvalue problems (in [35] and [36], for example). But alternating projection has not been recognized as a potential method for solving

any type of inverse eigenvalue problem. The most general treatment of alternating projection in the IEP literature is probably [37], but the authors do not offer a serious analysis of their algorithm’s behavior.

Here is the basic idea behind alternating projection. We seek a point of intersection between the set of matrices that satisfy a structural constraint and the set of matrices that satisfy a spectral constraint. An alternating projection begins at a matrix in the first set, from which it computes a matrix of minimal distance in the second set. Then the algorithm reverses the roles of the two sets and repeats the process indefinitely. Alternating projection is easy to apply, and it is usually globally convergent in a weak sense, as we show later.

D. Alternating Projections

Let us continue with a formal presentation of the generic alternating projection method for solving inverse eigenvalue problems. Suppose that we have two collections, \mathcal{Y} and \mathcal{Z} , of matrices with identical dimensions. Of course, we are imagining that one collection of matrices incorporates a spectral constraint while the other collection incorporates a structural constraint. To ensure that the algorithm is well-posed, assume that one collection is closed and the other is compact.

Algorithm 1 (Alternating Projection):

INPUT:

- An (arbitrary) initial matrix Y_0 with appropriate dimensions
- The number of iterations, J

OUTPUT:

- A matrix \bar{Y} in \mathcal{Y} and a matrix \bar{Z} in \mathcal{Z}

PROCEDURE:

- 1) Initialize $j = 0$.
- 2) Find a matrix Z_j in \mathcal{Z} such that

$$Z_j \in \arg \min_{Z \in \mathcal{Z}} \|Z - Y_j\|_{\mathbb{F}}.$$

We use $\|\cdot\|_{\mathbb{F}}$ to indicate the Frobenius norm.

- 3) Find a matrix Y_{j+1} in \mathcal{Y} such that

$$Y_{j+1} \in \arg \min_{Y \in \mathcal{Y}} \|Y - Z_j\|_{\mathbb{F}}.$$

- 4) Increment j by one.
- 5) Repeat Steps 2–4 until $j = J$.
- 6) Let $\bar{Y} = Y_J$ and $\bar{Z} = Z_{J-1}$.

A solution to the optimization problem in Step 2 is called a *projection* of Y_j onto \mathcal{Z} in analogy with the case where \mathcal{Z} is a linear subspace. Step 3 computes the projection of Z_j onto \mathcal{Y} . In a Hilbert space, it can be shown geometrically that a given point has a unique projection onto each closed, convex set. Projections onto general sets may not be uniquely determined, which fiercely complicates the analysis of Algorithm 1.

von Neumann, in 1933, was the first to consider alternating projection methods. He showed that if \mathcal{Y} and \mathcal{Z} are closed, linear subspaces of a Hilbert space, then alternating projection converges to the point in $\mathcal{Y} \cap \mathcal{Z}$ nearest to Y_0 [38]. In 1959, Cheney and Goldstein demonstrated that alternating projection between two compact, convex subsets of a Hilbert

space always yields a point in one set at minimal distance from the opposite set [4]. These two results inspired the application of Algorithm 1 to the inverse eigenvalue problems, Problems 1 and 2. Of course, the constraint sets that we consider are generally *not* convex. For a more extensive discussion of the literature on alternating projection, turn to Appendix I-G.

To implement the alternating projection algorithm, one must first solve the minimization problems in Steps 2 and 3. For obvious reasons, these optimizations are called the *matrix nearness problems* associated with \mathcal{Y} and \mathcal{Z} . Already there is an extensive literature on matrix nearness problems. See, for example, the articles [39]–[41], the survey [42], and many sections of the book [25]. Section III-F exhibits solutions to the nearness problems associated with the spectral constraints \mathcal{X}_α and \mathcal{G}_α . Even when it is necessary to solve a new nearness problem, this task often reduces to an exercise in differential calculus. This is one of the great advantages of Algorithm 1. In this article, we will always measure the distance between matrices using the Frobenius norm $\|\cdot\|_F$ because it facilitates the solution of matrix nearness problems. Of course, one could develop a formally identical algorithm using other norms, metrics, or divergences.

Since the constraint sets are generally nonconvex, alternating projection may not converge as well as one might wish. This explains why we have chosen to halt the algorithm after a fixed number of steps instead of waiting for $\|Y_j - Y_{j+1}\|_F$ to decline past a certain threshold. Indeed, it is theoretically possible that the sequence of iterates will not converge in norm. In practice, it appears that norm convergence always occurs. Section III-G provides a skeletal discussion of the theoretical convergence of alternating projection. We do not flesh out the analysis until Appendices I and II because a proper treatment requires some uncommon mathematics.

E. Application to Problems 1 and 2

To solve Problem 1, we propose an alternating projection between the structural constraint set \mathcal{S} and the spectral constraint set \mathcal{X}_α . Two matrix nearness problems arise. In the next subsection, we demonstrate how to find a tight frame in \mathcal{X}_α nearest to an arbitrary matrix. Sections IV and VI contain detailed treatments of two different structural constraints.

To solve Problem 2, we alternate between the spectral constraint \mathcal{G}_α and the structural constraint \mathcal{H} . In the next subsection, we show how to produce a matrix in \mathcal{G}_α that is nearest to an arbitrary matrix. In Section V, we analyze a specific structural constraint \mathcal{H} . After performing the alternating projection, it may still be necessary to extract a tight frame from the output Gram matrix. This is easily accomplished with a rank-revealing QR factorization or with an eigenvalue decomposition. Refer to [43, Sec. 5.4 and Ch. 8] for details.

F. Nearest Tight Frames

Standard tools of numerical linear algebra can be used to calculate an α -tight frame that is closest to an arbitrary matrix in Frobenius norm.

Theorem 2: Let $N \geq d$, and suppose that the $d \times N$ matrix Z has singular value decomposition $U\Sigma V^*$. With respect to

the Frobenius norm, a nearest α -tight frame to Z is given by αUV^* . Note that UV^* is the unitary part of a polar factorization of Z .

Assume in addition that Z has full row-rank. Then αUV^* is the *unique* α -tight frame closest to Z . Moreover, one may compute UV^* using the formula $(ZZ^*)^{-1/2}Z$.

Proof: The proof of this well-known result is similar to that of Theorem 3, which appears below. See also pp. 431–432 of [25]. Classical references on related problems include [44], [45]. The formula for the polar factor may be verified with a direct calculation. ■

It is also straightforward to compute a matrix in \mathcal{G}_α nearest to an arbitrary Hermitian matrix. This theorem appears to be novel, so we provide a short demonstration.

Theorem 3: Suppose that Z is an $N \times N$ Hermitian matrix with an eigenvalue decomposition $U\Lambda U^*$, where the entries of Λ are arranged in algebraically nonincreasing order. Let U_d be the $N \times d$ matrix formed from the first d columns of U . Then $\alpha U_d U_d^*$ is a matrix in \mathcal{G}_α that is closest to Z with respect to the Frobenius norm. This closest matrix is unique if and only if λ_d strictly exceeds λ_{d+1} .

Proof: Given an Hermitian matrix A , let $\lambda(A)$ denote the vector of its eigenvalues, arranged in algebraically nonincreasing order.

We must minimize $\|Z - \alpha G\|_F$ over all rank- d orthogonal projectors G . In consequence of the Wielandt–Hoffman Theorem [25, p. 368], the objective function is bounded below.

$$\|Z - \alpha G\|_F \geq \|\lambda(Z) - \alpha \lambda(G)\|_2.$$

Equality obtains if and only if Z and G are simultaneously diagonalizable by the same unitary matrix.

Suppose that Z has eigenvalue decomposition $U(\text{diag } \lambda(Z))U^*$. A rank- d orthogonal projector of dimension N has d eigenvalues equal to one and $(N - d)$ eigenvalues equal to zero. Therefore, one minimizer of the objective function is the matrix

$$G = U(I_d \oplus 0_{N-d})U^*.$$

Form a matrix U_d by extracting the first d columns of U . Then our minimizer may be expressed more simply as $G = U_d U_d^*$. That is, G is the orthogonal projector onto a d -dimensional subspace spanned by eigenvectors corresponding to the d algebraically largest eigenvalues of Z . This subspace is uniquely determined if and only if the d -th and $(d + 1)$ -st eigenvalues of Z are distinct. The orthogonal projector onto a subspace is unique, and the uniqueness claim follows. ■

It may be valuable to know that there are specialized algorithms for performing the calculations required by Theorems 2 and 3. For example, Higham has developed stable numerical methods for computing the polar factor of a matrix [46], [47] that are more efficient than computing a singular value decomposition or applying the formula $(ZZ^*)^{-1/2}Z$.

G. Basic Convergence Results

It should be clear that alternating projection never increases the distance between successive iterates. This does not mean that it will locate a point of minimal distance between the

constraint sets. It can be shown, however, that Algorithm 1 is globally convergent in a weak sense.

Define the distance between a point M and a set \mathcal{Y} via

$$\text{dist}(M, \mathcal{Y}) = \inf_{Y \in \mathcal{Y}} \|Y - M\|_F.$$

Theorem 4 (Global Convergence of Algorithm 1): Let \mathcal{Y} and \mathcal{Z} be closed sets, one of which is bounded. Suppose that alternating projection generates a sequence of iterates $\{(Y_j, Z_j)\}$. This sequence has at least one accumulation point.

- Every accumulation point lies in $\mathcal{Y} \times \mathcal{Z}$.
- Every accumulation point (\bar{Y}, \bar{Z}) satisfies

$$\|\bar{Y} - \bar{Z}\|_F = \lim_{j \rightarrow \infty} \|Y_j - Z_j\|_F.$$

- Every accumulation point (\bar{Y}, \bar{Z}) satisfies

$$\|\bar{Y} - \bar{Z}\|_F = \text{dist}(\bar{Y}, \mathcal{Z}) = \text{dist}(\bar{Z}, \mathcal{Y}).$$

For a proof of Theorem 4, turn to Appendix II-A. In some special cases, it is possible to develop stronger convergence results and characterizations of the fixed points. We will mention these results where they are relevant. The convergence of Algorithm 1 is geometric at best [48]–[51]. This is the major shortfall of alternating projection methods.

Note that the sequence of iterates may have many accumulation points. Moreover, the last condition does not imply that the accumulation point (\bar{Y}, \bar{Z}) is a fixed point of the alternating projection. So what are the potential accumulation points of a sequence of iterates? Since the algorithm never increases the distance between successive iterates, the set of accumulation points includes every pair of matrices in $\mathcal{Y} \times \mathcal{Z}$ that lie at minimal distance from each other. It is therefore reasonable to claim that the algorithm tries to solve Problems 1 and 2.

IV. PRESCRIBED COLUMN NORMS

As a first illustration of alternating projection, let us consider the most basic frame design problem: How does one build a tight frame with prescribed column norms?

This question has arisen in the context of constructing optimal signature sequences for direct-spread synchronous code-division multiple-access (DS-CDMA) channels. There are some finite algorithms available that yield a small number of solutions to the problem [7], [18]. These methods exploit the connection between frames and the Schur–Horn Theorem. They work by applying plane rotations to an initial tight frame to adjust its column norms while maintaining its tightness. Casazza and Leon have also developed a finite method that seems different in spirit [17].

To construct larger collections of frames, some authors have proposed iterative algorithms [19]–[23]. These techniques attempt to minimize the *total squared correlation* (TSC) of an initial matrix subject to constraints on its column norms. The TSC of a matrix is defined as

$$\text{TSC}(S) \stackrel{\text{def}}{=} \|S^* S\|_F^2 = \sum_{m,n} |\langle s_m, s_n \rangle|^2.$$

If we fix the squared column norms of S to be c_1, \dots, c_N , a short algebraic manipulation shows that minimizing the TSC is equivalent to solving

$$\min_S \|SS^* - \alpha I_d\|_F$$

where $\alpha = \sum_n c_n/d$. In words, minimizing the TSC is equivalent to finding a frame with prescribed column norms that is closest in Frobenius norm to a tight frame [52].

In comparison, alternating projection affords an elegant way to produce many tight frames with specified column norms. It focuses on the essential property of a tight frame—its singular values—to solve the problem. In this special case, we provide a complete accounting of the behavior of the alternating projection. Moreover, experiments show that it outperforms some of the other iterative algorithms that were developed specifically for this problem.

A. Constraint Sets and Nearness Problems

The algorithm will alternate between the set of matrices with fixed column norms and the set of tight frames with an appropriate tightness parameter.

Let the positive numbers c_1, \dots, c_N denote the squared column norms that are desired. We do not require that these numbers satisfy the majorization inequalities given in (4), although one cannot hope to find a tight frame if these inequalities fail. In that case, we would seek a matrix with the prescribed column norms that is closest to being a tight frame. In the DS-CDMA application, the column norms depend on the users' power constraints [6], [7].

In light of (3), the relation between the tightness parameter and the column norms, it is clear that α must equal $\sum_n c_n/d$. The spectral constraint set becomes

$$\mathcal{X}_\alpha \stackrel{\text{def}}{=} \{X \in \mathbb{C}^{d \times N} : XX^* = (\sum_n c_n/d) I_d\}.$$

Given an arbitrary $d \times N$ matrix, one may compute the closest tight frame in \mathcal{X}_α using Theorem 2.

The structural constraint set contains matrices with the correct column norms.

$$\mathcal{S} \stackrel{\text{def}}{=} \{S \in \mathbb{C}^{d \times N} : \|s_n\|_2^2 = c_n\}.$$

It is straightforward to solve the matrix nearness problem associated with this collection.

Proposition 5: Let Z be an arbitrary matrix with columns $\{z_n\}$. A matrix in \mathcal{S} is closest to Z in Frobenius norm if and only if it has the columns

$$s_n = \begin{cases} c_n z_n / \|z_n\|_2, & z_n \neq \mathbf{0} \\ c_n \mathbf{u}_n, & z_n = \mathbf{0} \end{cases}$$

where \mathbf{u}_n represents an arbitrary unit vector. If the columns of Z are all nonzero, then the solution to the nearness problem is unique.

Proof: We must minimize $\|S - Z\|_F$ over all matrices S from \mathcal{S} . Square and rewrite this objective function:

$$\|S - Z\|_F^2 = \sum_{n=1}^N \|s_n - z_n\|_2^2.$$

We can minimize each summand separately. Fix an index n , and expand the n -th term using $\|s_n\|_2^2 = c_n$.

$$\|s_n - z_n\|_2^2 = c_n + \|z_n\|_2^2 - 2\sqrt{c_n} \operatorname{Re} \left\langle \frac{s_n}{\|s_n\|_2}, z_n \right\rangle.$$

If $z_n \neq \mathbf{0}$, the unique maximizer of $\operatorname{Re} \langle \mathbf{u}, z_n \rangle$ over all unit vectors is $\mathbf{u} = z_n / \|z_n\|_2$. If $z_n = \mathbf{0}$, then every unit vector \mathbf{u} maximizes the inner product. ■

B. Convergence Results

In this setting, alternating projection converges in a fairly strong sense.

Theorem 6: Let S_0 have full rank and nonzero columns, and suppose that the alternating projection generates a sequence of iterates $\{(S_j, X_j)\}$. This sequence possesses at least one accumulation point, say (\bar{S}, \bar{X}) .

- Both \bar{S} and \bar{X} have full rank and nonzero columns.
- The pair (\bar{S}, \bar{X}) is a fixed point of the alternating projection. In other words, if we applied the algorithm to \bar{S} or to \bar{X} every pair of iterates would equal (\bar{S}, \bar{X}) .
- The accumulation point satisfies

$$\|\bar{S} - \bar{X}\|_F = \lim_{j \rightarrow \infty} \|S_j - X_j\|_F.$$

- The component sequences are asymptotically regular, i.e.,

$$\|S_{j+1} - S_j\|_F \rightarrow 0 \quad \text{and} \quad \|X_{j+1} - X_j\|_F \rightarrow 0.$$

- *Either* the component sequences both converge in norm,

$$\|S_j - \bar{S}\|_F \rightarrow 0 \quad \text{and} \quad \|X_j - \bar{X}\|_F \rightarrow 0,$$

or the set of accumulation points forms a continuum.

Proof: See Appendix II-C. ■

In the present case, it is also possible to characterize completely the fixed points of the algorithm that lie in \mathcal{S} .

Proposition 7: A full-rank matrix S from \mathcal{S} is a fixed point of the alternating projection between \mathcal{S} and \mathcal{X}_α if and only if the columns of S are all eigenvectors of SS^* . That is, $SS^*S = S\Lambda$, where $\Lambda \in \mathbb{C}^{N \times N}$ is diagonal and positive with no more than d distinct entries.

Proof: Refer to Appendix II-D. ■

Many of the fixed points in \mathcal{S} do not lie at minimal distance from \mathcal{X}_α , so they are not solutions to Problem 1. Nevertheless, the fixed points still have a tremendous amount of structure. Each fixed point can be written as a union of tight frames for mutually orthogonal subspaces of \mathbb{C}^d , and the set of fixed points is identical with the set of critical points of the TSC functional subject to the column norm constraint [23], [52]. The Ulukus–Yates algorithm, another iterative method for designing tight frames with specified column norms, has identical fixed points [20].

C. Numerical Examples

We offer a few simple examples to illustrate that the algorithm succeeds, and we provide some comparisons with the Ulukus–Yates algorithm.

Suppose first that we wish to construct a unit-norm tight frame for \mathbb{R}^3 consisting of five vectors. Initialized with a 3×5

matrix whose columns are chosen uniformly at random from the surface of the unit sphere, the algorithm returns

$$\bar{S} = \begin{bmatrix} 0.1519 & 0.4258 & -0.7778 & 0.0160 & -0.9258 \\ 0.9840 & -0.6775 & 0.1882 & 0.3355 & -0.3024 \\ -0.0926 & 0.5998 & 0.5997 & 0.9419 & -0.2269 \end{bmatrix}.$$

Each column norm of the displayed matrix equals one to machine precision, and the singular values are identical in their first eight digits. In all the numerical examples, the algorithm was terminated on the condition that $\|S_{j+1} - S_j\|_F < 10^{-8}$. Implemented in Matlab, the computation took 65 iterations, which lasted 0.0293 seconds on a 1.6 GHz machine.

Now let us construct a tight frame for \mathbb{R}^3 whose five vectors have norms 0.75, 0.75, 1, 1.25, and 1.25. With random initialization, we obtain

$$\bar{S} = \begin{bmatrix} -0.1223 & 0.1753 & -0.7261 & 0.0128 & -1.0848 \\ 0.7045 & -0.6786 & 0.6373 & 0.0972 & -0.6145 \\ -0.2263 & 0.2670 & 0.2581 & 1.2461 & -0.0894 \end{bmatrix}.$$

The column norms are correct to machine precision, and the singular values are identical to seven digits. The computation took 100 iterations, which lasted 0.0487 seconds.

Next we examine a case where the column norms do not satisfy the majorization condition. Suppose that we seek a “nearly tight” frame with column norms 0.5, 0.5, 1, 1, and 2. Random initialization yields

$$\bar{S} = \begin{bmatrix} -0.1430 & 0.1353 & -0.4351 & -0.0941 & -1.8005 \\ 0.4293 & -0.4213 & 0.7970 & -0.2453 & -0.7857 \\ -0.2127 & 0.2329 & 0.4189 & 0.9649 & -0.3754 \end{bmatrix}.$$

The column norms are all correct, but, as predicted, the frame is not tight. Nevertheless, the last vector is orthogonal to the first four vectors, which form a tight frame for their span. This is an optimal solution to the frame design problem. The calculation required 34 iterations over 0.0162 seconds.

Of course, alternating projection can produce complex-valued tight frames, as well as larger frames in higher-dimensional spaces. Such ensembles are too large to display in these columns. To give a taste of the algorithm’s general performance, we have compared it with our implementation of the Ulukus–Yates algorithm [20]. To construct unit-norm tight frames of various sizes, we initialized each algorithm with the same random matrix. Then we plotted the comparative execution times. Figure 1 shows the results for 64 real dimensions, and Figure 2 shows the results for 64 complex dimensions. Note the different scales on the time axes.

Both algorithms perform slowly when N is small because tight frames are relatively scarce, which makes them difficult to find. Indeed, it is known that (modulo rotations) there exists a unique tight frame of $(d+1)$ vectors in d -dimensional space [53, p. 13]. Another reason that the alternating projection algorithm performs better as the problem grows is that a collection of N uniformly random unit-vectors converges almost surely to a tight frame as N tends to infinity [54, Thm. 1]. It is therefore perplexing that the Ulukus–Yates algorithm performs more and more slowly. One might attribute this




Fig. 1. Comparison of alternating projection with the Ulukus-Yates algorithm in 64 real dimensions.




Fig. 2. Comparison of alternating projection with the Ulukus-Yates algorithm in 64 complex dimensions.

behavior to the fact that the algorithm does not act to equalize the singular values of the frame.

V. EQUIANGULAR TIGHT FRAMES

In this section, we will consider a frame design problem that leads to a simple structural constraint on the Gram matrix. The goal of the alternating projection will be to design a suitable Gram matrix, from which the frame may be extracted afterward.

A tight frame is a generalization of an orthonormal basis because they share the Parseval property. But orthonormal bases have other characteristics that one may wish to extend. In particular, every orthonormal basis is *equiangular*. That is, each pair of distinct vectors has the same inner product, namely zero. This observation suggests that one seek out

equiangular tight frames. The underlying intuition is that these frames will contain vectors maximally separated in space.

Define an *equiangular tight frame* to be a unit-norm tight frame in which each pair of columns has the same *absolute* inner product. Since we are considering unit-norm tight frames, the absolute inner product between two frame vectors equals the cosine of the acute angle between the one-dimensional subspaces spanned by the two vectors. For this reason are the frames called equiangular. One can show that each inner product in an equiangular tight frame has modulus

$$\mu = \sqrt{\frac{N-d}{d(N-1)}}. \quad (7)$$

It is a remarkable fact that *every* ensemble of N unit vectors in d dimensions contains a pair whose absolute inner product strictly exceeds μ , *unless* the vectors form an equiangular tight frame. Unfortunately, equiangular tight frames only exist for rare combinations of d and N . In particular, a real equiangular tight frame can exist only if $N \leq \frac{1}{2}d(d+1)$, while a complex equiangular tight frame requires that $N \leq d^2$ [12, Thm. 2.3]. The paper [55] also contains detailed necessary conditions on real equiangular tight frames and on equiangular tight frames over finite alphabets.

One can view equiangular tight frames as a special type of *Grassmannian frame*. In finite dimensions, Grassmannian frames are unit-norm frames whose largest absolute inner product is minimal for a given d and N [12]. Their name is motivated by the fact that they correspond with sphere packings in the Grassmannian manifold of all one-dimensional subspaces of a vector space [28]. Grassmannian frames have applications in coding theory and communications engineering [9]–[12]. They also provide a natural set of vectors to use for sparse approximation [13]–[15].

In general, it is torturous to design Grassmannian frames. Not only is the optimization difficult, but there is no general procedure for deciding when a frame solves the optimization problem unless it meets a known lower bound. Most of the current research has approached the design problem with algebraic tools. A notable triumph of this type is the construction of Kerdock codes over \mathbb{Z}_2 and \mathbb{Z}_4 due to Calderbank et al. [56]. Other explicit constructions are discussed in the articles [10], [12]. In the numerical realm, Sloane has used his Gosset software to produce and study sphere packings in real Grassmannian spaces [57]. Sloane's algorithms have been extended to complex Grassmannian spaces in [58]. We are not aware of any other numerical methods.

In this article, we will construct equiangular tight frames for real and complex vector spaces using alternating projection. The method can easily be extended to compute Grassmannian frames and packings in Grassmannian manifolds, but that is another paper for another day [59, Ch. 7].

A. Constraint Sets and Nearness Problems

The signal of an equiangular tight frame is that each inner product between distinct vectors has the same modulus. Since the Gram matrix of a tight frame displays all of the inner products, it is more natural to construct the Gram matrix of an

equiangular tight frame than to construct the frame synthesis matrix directly. Therefore, the algorithm will alternate between the collection of Hermitian matrices that have the correct spectrum and the collection of Hermitian matrices that have sufficiently small off-diagonal entries.

Since we are working with unit-norm tight frames, the tightness parameter α must equal N/d . This leads to the spectral constraint set

$$\mathcal{G}_\alpha \stackrel{\text{def}}{=} \{G \in \mathbb{C}^{N \times N} : G = G^* \text{ and } G \text{ has eigenvalues } \underbrace{(N/d, \dots, N/d, 0, \dots, 0)}_d\}.$$

Theorem 3 shows how to find a matrix in \mathcal{G}_α nearest to an arbitrary Hermitian matrix.

In an equiangular tight frame, each vector has unit norm but no two vectors have inner product larger than μ . Therefore, we define the structural constraint set

$$\mathcal{H}_\mu \stackrel{\text{def}}{=} \{H \in \mathbb{C}^{N \times N} : H = H^*, \text{diag } H = \mathbf{1} \text{ and } \max_{m \neq n} |h_{mn}| \leq \mu\}.$$

It may seem more natural to require that the off-diagonal entries have modulus exactly equal to μ , but our experience indicates that the present formulation works better, perhaps because \mathcal{H}_μ is convex. The following proposition shows how to produce the nearest matrix in \mathcal{H}_μ .

Proposition 8: Let Z be an arbitrary matrix. With respect to Frobenius norm, the unique matrix in \mathcal{H}_μ closest to Z has a unit diagonal and off-diagonal entries that satisfy

$$h_{mn} = \begin{cases} z_{mn} & \text{if } |z_{mn}| \leq \mu \text{ and} \\ \mu e^{i \arg z_{mn}} & \text{otherwise.} \end{cases}$$

We use i to denote the imaginary unit.

Proof: The argument is straightforward. ■

B. Convergence Results

The general convergence result, Theorem 4, applies to the alternating projection between \mathcal{G}_α and \mathcal{H}_μ . We also obtain a local convergence result.

Theorem 9: Assume that the alternating projection between \mathcal{G}_α and \mathcal{H}_μ generates a sequence of iterates $\{(G_j, H_j)\}$, and suppose that there is an iteration J during which $\|G_J - H_J\|_F < N/(d\sqrt{2})$. Then the sequence of iterates possesses at least one accumulation point, say (\bar{G}, \bar{H}) .

- The accumulation point lies in $\mathcal{G}_\alpha \times \mathcal{H}_\mu$.
- The pair (\bar{G}, \bar{H}) is a fixed point of the alternating projection. In other words, if we applied the algorithm to \bar{G} or to \bar{H} , every iterate would equal (\bar{G}, \bar{H}) .
- The accumulation point satisfies

$$\|\bar{G} - \bar{H}\|_F = \lim_{j \rightarrow \infty} \|G_j - H_j\|_F.$$

- The component sequences are asymptotically regular, i.e.,

$$\|G_{j+1} - G_j\|_F \rightarrow 0 \quad \text{and} \quad \|H_{j+1} - H_j\|_F \rightarrow 0.$$

- Either the component sequences both converge in norm,

$$\|G_j - \bar{G}\|_F \rightarrow 0 \quad \text{and} \quad \|H_j - \bar{H}\|_F \rightarrow 0,$$

or the set of accumulation points forms a continuum.

Proof: See Appendix II-B. ■

C. Numerical Examples

First, let us illustrate just how significant a difference there is between vanilla tight frames and equiangular tight frames. Here is the Gram matrix of a six-vector, unit-norm tight frame for \mathbb{R}^3 :

$$\begin{bmatrix} 1.0000 & 0.2414 & -0.6303 & 0.5402 & -0.3564 & -0.3543 \\ 0.2414 & 1.0000 & -0.5575 & -0.4578 & 0.5807 & -0.2902 \\ -0.6303 & -0.5575 & 1.0000 & 0.2947 & 0.3521 & -0.2847 \\ 0.5402 & -0.4578 & 0.2947 & 1.0000 & -0.2392 & -0.5954 \\ -0.3564 & 0.5807 & 0.3521 & -0.2392 & 1.0000 & -0.5955 \\ -0.3543 & -0.2902 & -0.2847 & -0.5954 & -0.5955 & 1.0000 \end{bmatrix}.$$

Notice that the inner-products between vectors are quite disparate, ranging in magnitude between 0.2392 and 0.6303. These inner products correspond to acute angles of 76.2° and 50.9° . In fact, this tight frame is pretty tame; the largest inner products in a unit-norm tight frame can be arbitrarily close to one⁴. The Gram matrix of a six-vector, equiangular tight frame for \mathbb{R}^3 looks quite different:

$$\begin{bmatrix} 1.0000 & 0.4472 & -0.4472 & 0.4472 & -0.4472 & -0.4472 \\ 0.4472 & 1.0000 & -0.4472 & -0.4472 & 0.4472 & -0.4472 \\ -0.4472 & -0.4472 & 1.0000 & 0.4472 & 0.4472 & -0.4472 \\ 0.4472 & -0.4472 & 0.4472 & 1.0000 & -0.4472 & -0.4472 \\ -0.4472 & 0.4472 & 0.4472 & -0.4472 & 1.0000 & -0.4472 \\ -0.4472 & -0.4472 & -0.4472 & -0.4472 & -0.4472 & 1.0000 \end{bmatrix}.$$

Every pair of vectors meets at an acute angle of 63.4° . The vectors in this frame can be interpreted as the diagonals of an icosahedron [28, Table 1].

We have used alternating projection to compute equiangular tight frames, both real and complex, in dimensions two through six. The algorithm performed poorly when initialized with random vectors, which led us to adopt a more sophisticated approach. We begin with many random vectors and winnow this collection down by repeatedly removing whatever vector has the largest inner product against another vector. It is fast and easy to design starting points in this manner, yet the results are impressive. These calculations are summarized in Table I.

Alternating projection can locate every real, equiangular tight frame in dimensions two through six; algebraic considerations eliminate all the remaining values of N [55, Thm. A and Thm. 6.1]. Moreover, the method computes these ensembles very efficiently. For example, the algorithm produced a six-vector, equiangular tight frame for \mathbb{R}^3 after a single trial. In this case, 70 iterations lasting 0.4573 seconds were sufficient to determine the first eight decimal places of the inner products.

In the complex case, the algorithm was able to compute every equiangular tight frame that we know of. Unfortunately, no one has yet developed necessary conditions on the existence of complex, equiangular tight frames aside from the upper bound, $N \leq d^2$, and so we have been unable to rule out the existence of other ensembles. Some of the computations progressed quite smoothly. After 1000 iterations and 18.75 seconds, alternating projection delivered a collection of 25

⁴To see this, consider a tight frame that contains two copies of an orthonormal basis, where one copy is rotated away from the other by an arbitrarily small angle.

N	d					N	d				
	2	3	4	5	6		2	3	4	5	6
3	R	R	20
4	C	R	R	R	..	21	C	..
5	R	R	..	22
6	R	R	23
7	..	C	C	..	R	24
8	C	25	C	..
9	..	C	C	26
10	R	..	27
11	C	C	28
12	C	29
13	C	30
14	31	C
15	32
16	C	..	R	33
17	34
18	35
19	36	C

TABLE I
EQUIANGULAR TIGHT FRAMES

The notations \mathbb{R} and \mathbb{C} respectively indicate that alternating projection was able to compute a real, or complex, equiangular tight frame. Note that every real equiangular tight frame is automatically a complex equiangular tight frame. One period (.) means that no real equiangular tight frame exists, and two periods (..) mean that no equiangular tight frame exists at all.

vectors in five dimensions whose inner products were identical in the first eight decimal places. On the other hand, it took 5000 iterations and 85.75 seconds to produce 21 vectors in five dimensions whose inner products reached the same level of accuracy. Even worse, we were unable to locate the 31-vector equiangular tight frame in \mathbb{C}^6 until we had performed two dozen random trials that lasted several minutes each. It is some consolation that the authors of [58, Table I] indicate their algorithm could not compute this ensemble at all.

It seems clear that some equiangular tight frames are much easier to find than others. We have encountered less success at constructing equiangular tight frames in higher dimensions. But we have neither performed extensive experiments nor have we attempted to fine-tune the method.

VI. PEAK-TO-AVERAGE-POWER RATIO

Finally, let us present a situation in which the matrix nearness problem is much more difficult.

As we have mentioned, tight frames with prescribed vector norms coincide with signature sequences that maximize sum capacity in the uplink of direct-spread, synchronous code division multiple access (DS-CDMA) systems [6]–[8]. Unfortunately, general tight frames can have properties that are undesirable in practice. In particular, the individual frame vectors may have large *peak-to-average-power ratio* (PAR).

The PAR of an analog signal measures how the largest value of the signal compares with its average power. Signals with large PAR require higher dynamic range on the analog-to-digital converters and the digital-to-analog converters. They may also require more linear (and thus higher cost) power amplifiers. In DS-CDMA systems, the PAR is normally of concern only in the downlink (see, e.g., [60]), where linear combinations of signatures can conspire to have tremendous peak power. On the uplink, the PAR problem is fundamentally

different because it only involves individual signatures. Conventionally, the uplink PAR has not received attention because most systems use binary spreading sequences, which always have unit PAR. If general sum-capacity-optimal sequences are to be used in real systems, then PAR side constraints should be included in the design problem. Therefore, we will consider how to construct tight frames whose columns have prescribed norms and low peak-to-average-power ratios.

As discussed in Section IV, many algorithms have already been developed for constructing tight frames with prescribed vector norms, such as [7], [19], [20], [22]. Unfortunately, these methods cannot accept additional constraints on the vectors, and thus they are not suitable for finding tight frames whose vectors have low PAR. We show that alternating projection provides a way to produce these ensembles. The PAR problem makes an interesting test case because it induces a matrix nearness problem that is considerably more challenging than those we have examined in previous sections.

A. Constraint Sets and Matrix Nearness Problems

The PAR in a digital communication system is fundamentally related to the analog waveforms that are generated. From the perspective of sequence design, it usually suffices to consider the PAR defined directly from the discrete sequence. The discrete PAR of a vector \mathbf{z} is the quantity

$$\text{PAR}(\mathbf{z}) \stackrel{\text{def}}{=} \frac{\max_m |z_m|^2}{\sum_m |z_m|^2 / d}.$$

Note that $1 \leq \text{PAR}(\mathbf{z}) \leq d$. The lower extreme corresponds to a vector whose entries have identical modulus, while the upper bound is attained only by (scaled) canonical basis vectors.

Suppose that we require the columns of the frame to have squared norms c_1, \dots, c_N . In the DS-CDMA application, these numbers depend on the users' power constraints [6], [7]. It follows from (3) that $\alpha = \sum_n c_n / d$. The spectral constraint set becomes

$$\mathcal{X}_\alpha \stackrel{\text{def}}{=} \{\mathbf{X} \in \mathbb{C}^{d \times N} : \mathbf{X}\mathbf{X}^* = (\sum_n c_n / d) \mathbf{I}_d\}.$$

Theorem 2 delivers the solution to the associated matrix nearness problem.

Let ρ denote the upper bound on the PAR of the frame elements. Then the structural constraint set will be

$$\mathcal{S} \stackrel{\text{def}}{=} \{\mathbf{S} \in \mathbb{C}^{d \times N} : \text{PAR}(\mathbf{s}_n) \leq \rho \text{ and } \|\mathbf{s}_n\|_2^2 = c_n\}.$$

Given an arbitrary matrix \mathbf{Z} , we must compute the nearest element of \mathcal{S} . Since the structural constraint on each column is independent and the Frobenius norm is separable, each column yields an independent optimization problem. For each column \mathbf{z}_n of the input matrix, we claim that the following algorithm returns \mathbf{s}_n , the corresponding column of a nearest matrix \mathbf{S} from \mathcal{S} .

Algorithm 2 (Nearest Vector with Low PAR):

INPUT:

- An input vector \mathbf{z} from \mathbb{C}^d
- A positive number c , the squared norm of the solution vector

- A number ρ from $[1, d]$, which equals the maximum permissible PAR

OUTPUT:

- A vector \mathbf{s} from \mathbb{C}^d that solves

$$\min_{\mathbf{s}} \|\mathbf{s} - \mathbf{z}\|_2 \quad \text{subj. to} \quad \text{PAR}(\mathbf{s}) \leq \rho \text{ and } \|\mathbf{s}\|_2^2 = c.$$

PROCEDURE:

- 1) Scale \mathbf{z} to have unit norm; define $\delta = \sqrt{c\rho/d}$; and initialize $k = 0$.
- 2) Let \mathcal{M} index $(d - k)$ components of \mathbf{z} with least magnitude. If this set is not uniquely determined, increment k and repeat Step 2.
- 3) If $z_m = 0$ for each m in \mathcal{M} , a solution vector is

$$\mathbf{s} = \begin{cases} \sqrt{\frac{c-k\delta^2}{d-k}} & \text{when } m \in \mathcal{M}, \text{ and} \\ \delta e^{i \arg z_m} & \text{when } m \notin \mathcal{M}. \end{cases}$$

- 4) Otherwise, let

$$\gamma = \sqrt{\frac{c - k\delta^2}{\sum_{m \in \mathcal{M}} |z_m|^2}}.$$

- 5) If $\gamma z_m > \delta$ for any m in \mathcal{M} , increment k and return to Step 2.
- 6) The unique solution vector is

$$\mathbf{s} = \begin{cases} \gamma z_m & \text{when } m \in \mathcal{M}, \text{ and} \\ \delta e^{i \arg z_m} & \text{when } m \notin \mathcal{M}. \end{cases}$$

When $\rho = 1$, the output of the algorithm is a vector with unimodular entries that have the same phase as the corresponding entries of \mathbf{z} . On the other hand, when $\rho = d$, the output vector equals \mathbf{z} . We now prove that the algorithm is correct.

Proof: We must solve the optimization problem

$$\min_{\mathbf{s}} \|\mathbf{s} - \mathbf{z}\|_2^2 \quad \text{subject to} \quad \text{PAR}(\mathbf{s}) \leq \rho \text{ and } \|\mathbf{s}\|_2^2 = c.$$

Let us begin with some major simplifications. First, rewrite the PAR constraint by enforcing the norm requirement and rearranging to obtain the equivalent condition

$$\max_m |s_m| \leq \sqrt{c\rho/d}.$$

In the rest of the argument, the symbol δ will abbreviate the quantity $\sqrt{c\rho/d}$. The PAR constraint becomes $|s_m| \leq \delta$ for each $m = 1, \dots, d$.

Now, expand the objective function and enforce the norm constraint again to obtain

$$\min_{\mathbf{s}} \left[c - 2 \operatorname{Re} \langle \mathbf{s}, \mathbf{z} \rangle + \|\mathbf{z}\|_2^2 \right].$$

Observe that it is necessary and sufficient to minimize the second term. It follows that the optimizer does not depend on the scale of the input vector \mathbf{z} . So take $\|\mathbf{z}\|_2 = 1$ without loss of generality.

Next, note that the PAR constraint and the norm constraint do not depend on the phases of the components in \mathbf{s} . Therefore, the components of an optimal \mathbf{s} must have the same phases as the components of the input vector \mathbf{z} . In consequence, we may assume that both \mathbf{s} and \mathbf{z} are nonnegative real vectors.

We have reached a much more straightforward optimization problem. Given a nonnegative vector \mathbf{z} with unit norm, we must solve

$$\max_{\mathbf{s}} \langle \mathbf{s}, \mathbf{z} \rangle \quad \text{subject to} \quad \langle \mathbf{s}, \mathbf{s} \rangle = c \text{ and } 0 \leq s_m \leq \delta,$$

Observe that every point of the feasible set is a regular point, i.e., the gradients of the constraints are linearly independent. Therefore, Karush-Kuhn-Tucker (KKT) theory will furnish necessary conditions on an optimizer [61, Sec. 28].

We form the Lagrangian function

$$L(\mathbf{s}, \lambda, \boldsymbol{\mu}, \boldsymbol{\nu}) = -\langle \mathbf{s}, \mathbf{z} \rangle + \frac{1}{2} \lambda (\langle \mathbf{s}, \mathbf{s} \rangle - c) - \langle \mathbf{s}, \boldsymbol{\mu} \rangle + \langle \mathbf{s} - \delta \mathbf{1}, \boldsymbol{\nu} \rangle.$$

The Lagrange multipliers $\boldsymbol{\mu}$ and $\boldsymbol{\nu}$ are non-negative because they correspond to the lower and upper bounds on \mathbf{s} . Meanwhile, the multiplier λ is unrestricted because it is associated with the equality constraint.

The first-order KKT necessary condition on a regular local maximum \mathbf{s}^* is that

$$\begin{aligned} \mathbf{0} &= (\nabla_{\mathbf{s}} L)(\mathbf{s}^*, \lambda^*, \boldsymbol{\mu}^*, \boldsymbol{\nu}^*) \\ &= -\mathbf{z} + \lambda^* \mathbf{s}^* - \boldsymbol{\mu}^* + \boldsymbol{\nu}^*, \end{aligned} \quad (8)$$

where $\mu_m^* > 0$ only if $s_m^* = 0$ and $\nu_m^* > 0$ only if $s_m^* = \delta$. Notice that one of μ_m^* or ν_m^* must be zero because they correspond to mutually exclusive constraints. The second-order KKT necessary condition on a regular local maximum is that

$$\begin{aligned} 0 &\leq \mathbf{y}^T (\nabla_{\mathbf{s}}^2 L)(\mathbf{s}^*, \lambda^*, \boldsymbol{\mu}^*, \boldsymbol{\nu}^*) \mathbf{y} \\ &= \lambda^* \mathbf{y}^T \mathbf{y} \end{aligned}$$

for every vector \mathbf{y} in the subspace of first-order feasible variations. This subspace is nontrivial, so $\lambda^* \geq 0$.

Solve (8) to obtain

$$\lambda^* \mathbf{s}^* = \mathbf{z} + \boldsymbol{\mu}^* - \boldsymbol{\nu}^*.$$

Whenever $\mu_m^* > 0$, both $s_m^* = 0$ and $\nu_m^* = 0$. This combination is impossible because $z_m \geq 0$. Therefore, we may eliminate $\boldsymbol{\mu}^*$ to reach

$$\lambda^* \mathbf{s}^* = \mathbf{z} - \boldsymbol{\nu}^*.$$

The cases $\lambda^* = 0$ and $\lambda^* > 0$ require separate consideration.

If $\lambda^* = 0$, it is clear that $\boldsymbol{\nu}^* = \mathbf{z}$. Since $\nu_m^* > 0$ only if $s_m^* = \delta$, we must have $s_m^* = \delta$ whenever $z_m > 0$. Suppose that k components of \mathbf{s}^* equal δ . The remaining $(d - k)$ components are not uniquely determined by the optimization problem. From the many solutions, we choose one such that

$$s_m^* = \sqrt{\frac{c - k\delta^2}{d - k}} \quad \text{for } m \text{ where } z_m = 0.$$

This formula ensures that \mathbf{s}^* has the correct norm and that none of its entries exceeds δ .

When $\lambda^* > 0$, the solution has the form

$$\mathbf{s}^* = [\gamma \mathbf{z}]_{\delta},$$

where γ is positive and the operator $[\cdot]_{\delta}$ truncates to δ the components of its argument that exceed δ . It is clear that the

largest components of \mathbf{z} are all truncated at the same time. We only need to determine *which* components these are.

To that end, observe that $\gamma \mapsto \|\lceil \gamma \mathbf{z} \rceil_\delta\|_2$ is a strictly increasing function on $[0, \delta/z_{\min}]$, where z_{\min} is the least positive component of \mathbf{z} . For at most one value of γ , therefore, does the vector $\lceil \gamma \mathbf{z} \rceil_\delta$ have norm \sqrt{c} . If this norm value were not attained, then λ^* would equal zero. Let k be the number of entries of \mathbf{s}^* that equal δ , and suppose that \mathcal{M} indexes the remaining $(d - k)$ components. Then

$$c = \|\mathbf{s}^*\|_2^2 = k\delta^2 + \gamma^2 \sum_{m \in \mathcal{M}} |z_m|^2.$$

Recall that γ is positive. Therefore, is impossible that $k\delta^2 > c$. When $k\delta^2 = c$, it follows that $z_m = 0$ for each m in \mathcal{M} . Otherwise, z_m must be nonzero for some m in \mathcal{M} . Then the value of γ must be

$$\gamma = \sqrt{\frac{c - k\delta^2}{\sum_{m \in \mathcal{M}} |z_m|^2}}.$$

■

B. Convergence

For the alternating projection between the PAR constraint set and the set of α -tight frames, we have not proven a more elaborate convergence theorem than the basic result, Theorem 4, because it is not easy to guarantee that the solution to the PAR matrix nearness problem is unique. We have been able to provide a sufficient condition on the fixed points of the iteration that lie in the PAR constraint set \mathcal{S} . Note that similar fixed points arose in Section IV.

Theorem 10: A sufficient condition for a full-rank matrix S from \mathcal{S} to be a fixed point of the alternating projection between \mathcal{S} and \mathcal{X}_α is that the columns of S are all eigenvectors of SS^* . That is, $SS^*S = S\Lambda$, where $\Lambda \in \mathbb{C}^{N \times N}$ is diagonal and positive, with no more than d distinct entries.

Proof: Refer to Appendix II-E. ■

C. Numerical Examples

Let us demonstrate that alternating projection can indeed produce tight frames whose columns have specified PAR and specified norm. We will produce complex tight frames because, in the real case, PAR constraints can lead to a discrete optimization problem. The experiments all begin with the initial 3×6 matrix

$$\begin{bmatrix} .0748 + .3609i & .0392 + .4558i & .5648 + .3635i \\ .5861 - .0570i & -.2029 + .8024i & -.5240 + .4759i \\ -.7112 + .1076i & -.2622 - .1921i & -.1662 + .1416i \\ -.2567 + .4463i & .7064 + .6193i & .1586 + .6825i \\ -.1806 - .1015i & -.1946 - .1889i & .5080 + .0226i \\ .0202 + .8316i & .0393 - .2060i & .2819 + .4135i \end{bmatrix}.$$

The respective PAR values of its columns are 1.5521, 2.0551, 1.5034, 2.0760, 2.6475, and 1.4730.

Unit-PAR tight frames are probably the most interesting example. In each column of a unit-PAR tight frame, the entries share an identical modulus, which depends on the norm of the

column. Let us apply our algorithm to calculate a unit-PAR, unit-norm tight frame:

$$\begin{bmatrix} .1345 + .5615i & .1672 + .5526i & .4439 + .3692i \\ .5410 - .2017i & -.0303 + .5766i & -.5115 + .2679i \\ -.5768 + .0252i & -.2777 - .5062i & -.2303 + .5294i \\ -.3358 + .4696i & .4737 + .3300i & .0944 + .5696i \\ -.5432 - .1956i & -.3689 - .4442i & .5747 + .0554i \\ .1258 + .5635i & -.0088 - .5773i & .4132 + .4033i \end{bmatrix}.$$

Indeed, each of the columns has unit PAR and unit norm. The singular values of the matrix are identical to eight decimal places. The calculation required 78 iterations lasting 0.1902 seconds.

Alternating projection can also compute tight frames whose columns have unit PAR but different norms. For example, if we request the column norms 0.75, 0.75, 1, 1, 1.25, and 1.25, the algorithm yields

$$\begin{bmatrix} .3054 + .3070i & .1445 + .4082i & .3583 + .4527i \\ .4295 - .0549i & .1235 + .4150i & -.5597 + .1418i \\ -.4228 - .0936i & -.0484 - .4303i & .0200 + .5770i \\ -.4264 + .3893i & .4252 + .5831i & .3622 + .6242i \\ -.5393 - .2060i & -.4425 - .5701i & .7165 - .0863i \\ .2585 + .5162i & -.2894 - .6611i & .1291 + .7101i \end{bmatrix}.$$

One can check that the column norms, PAR and singular values all satisfy the design requirements to eight or more decimal places. The computation took 84 iterations over 0.1973 seconds.

Less stringent constraints on the PAR pose even less trouble. For example, we might like to construct a tight frame whose PAR is bounded by two and whose columns have norms 0.75, 0.75, 1, 1, 1.25 and 1.25. It is

$$\begin{bmatrix} .0617 + .1320i & .0184 + .2764i & .4299 + .3593i \\ .4256 - .1031i & -.0558 + .5938i & -.5920 + .4974i \\ -.5912 + .0025i & -.1304 - .3363i & -.0807 + .2857i \\ -.1382 + .2511i & .6847 + .7436i & .2933 + .6939i \\ -.4306 - .2650i & -.2095 - .3072i & .7317 + .0928i \\ .0852 + .8093i & -.3504 - .5289i & .2918 + .6048i \end{bmatrix}.$$

The computer worked for 0.0886 seconds, during which it performed 49 iterations. As usual, the singular values match to eight decimal places. It is interesting to observe that the frame exceeds the design specifications. The respective PAR values of its columns are 1.8640, 1.8971, 1.7939, 1.9867, 1.9618, and 1.0897.

VII. DISCUSSION

As advertised, we have developed an alternating projection method for solving frame design problems, and we have provided ample evidence that it succeeds. In this section, we discuss some implementation issues and some of the limitations of the algorithm. We conclude with a collection of related problems that one can also solve with alternating projection.

A. The Starting Point

For alternating projection to succeed, it is essential to choose a good starting point. Here are a few general strategies that may be useful.

The simplest method is to select N vectors uniformly at random from the surface of the unit sphere in \mathbb{C}^d and form them into an initial matrix. Although this technique sometimes works, it is highly probable that there will be pairs of strongly correlated vectors, and it is usually preferable for the frame to contain dissimilar vectors. Nevertheless, a collection of random vectors converges almost surely to a tight frame as more vectors are added [54, Thm. 1].

A more practical idea is to select many vectors, say $2dN$, and then use a clustering algorithm—such as Lloyd–Max [62], spherical k -means [63] or diametrical clustering [64]—to separate these vectors into N clusters. The cluster representatives will usually be much more diverse than vectors chosen at random. A related approach would select many random vectors and then greedily remove vectors that are highly correlated with the remaining vectors. This method seems to furnish excellent starting points for constructing equiangular tight frames. One might also build up a collection of random vectors by allowing a new vector to join only if it is weakly correlated with the current members.

Another technique is to start with a tight frame that has been developed for another application. By rotating the frame at random, it is possible to obtain many different starting points that retain some of the qualities of the original frame. In particular, equiangular tight frames make excellent initializers.

It is also possible to choose a collection of N vectors from a larger frame for \mathbb{C}^d . Similarly, one might truncate some coordinates from a frame in a higher-dimensional space. In particular, one might truncate an orthonormal basis for \mathbb{C}^N to retain only d coordinates. See [65], for example, which uses the Fourier transform matrix in this manner.

B. Limitations

Alternating projection cannot alleviate all the pain of frame design. While preparing this report, we encountered several difficulties.

A theoretical irritation is the lack of a proof that alternating projection converges in norm. No general proof is possible, as the counterexample in [66] makes clear. Nevertheless, it would be comforting to develop sufficient conditions that guarantee the convergence of alternating projections between nonconvex sets. The results of [66] are the best that we know of. We would also like to develop conditions that can ensure convergence to a pair of points at minimal distance. Here, the most general results are probably due to Csiszár and Tusnády [67].

Another major inconvenience is that alternating projection converges at a geometric rate (or worse) [48]–[51]. For large problems, it can be painful to wait on the solution. An interesting topic for future research would be a method of acceleration.

A more specific disappointment was the inability of alternating projection to construct tight frames over small finite

alphabets. It is straightforward to solve the matrix nearness problem associated with a finite alphabet, and it can be shown that the algorithm always converges in norm to a fixed point. But the algorithm never once yielded a tight frame. This failure is hardly surprising; discrete constraints are some of the most difficult to deal with in optimization. It may be possible to use annealing to improve the performance of the algorithm. This would be a valuable direction for future research.

C. Related Problems

We have permitted a great deal of freedom in the selection of the structural constraint set, but we only considered the spectral constraints that arise naturally in connection with tight frames. Nevertheless, alternating projection offers a straightforward method for addressing other inverse eigenvalue problems. For example, one might try to construct general frames with prescribed lower and upper frame bounds, α and β . Instead of forcing the Gram matrix to be a rank- d orthogonal projector, one might impose only a rank constraint or a constraint on its condition number. To implement the algorithm, it would only be necessary to solve the matrix nearness problem associated with these spectral constraints.

One can also use alternating projection to construct positive semi-definite (PSD) matrices that have certain structural properties. Higham, for example, has used a corrected alternating projection to produce the correlation matrix nearest to an input matrix [36]. (A *correlation matrix* is a PSD matrix with a unit diagonal.) Since the PSD matrices form a closed, convex set, it is possible to prove much more about the behavior of alternating algorithms.

We have also had good success using alternating projection to compute sphere packings in real and complex projective spaces. These methods can be extended to produce sphere packings in real and complex Grassmannian manifolds [59]. It seems clear that alternating projection has a promising future for a new generation of problems.

APPENDIX I POINT-TO-SET MAPS

To understand the convergence of the algorithms, we rely on some basic results from the theory of point-to-set maps. Zangwill's book [68] is a good basic reference with applications to mathematical programming. More advanced surveys include [69], [70]. de Leeuw presents statistical applications in [50]. We have drawn from all these sources here.

A. Point-to-Set Maps

Let \mathcal{Y} and \mathcal{Z} be arbitrary sets. The *power set* of \mathcal{Z} is the collection of all subsets of \mathcal{Z} , and it is denoted by $2^{\mathcal{Z}}$. A *point-to-set map* from \mathcal{Y} to \mathcal{Z} is a function $\Omega : \mathcal{Y} \rightarrow 2^{\mathcal{Z}}$. In words, Ω maps each point of \mathcal{Y} to a subset of \mathcal{Z} .

There are several different ways of combining point-to-set maps. Take two maps $\Omega_{yz} : \mathcal{Y} \rightarrow 2^{\mathcal{Z}}$ and $\Omega_{zw} : \mathcal{Z} \rightarrow 2^{\mathcal{W}}$. The composition of these maps carries a point y to a subset of \mathcal{W} via the rule

$$(\Omega_{zw} \circ \Omega_{yz})(y) = \bigcup_{z \in \Omega_{yz}(y)} \Omega_{zw}(z).$$

This definition can be extended in the obvious way to a longer composition of maps. Now, suppose Ω_{uv} maps \mathcal{U} to $2^{\mathcal{V}}$. The Cartesian product of Ω_{uv} and Ω_{yz} is the point-to-set map from $\mathcal{U} \times \mathcal{Y}$ to $\mathcal{V} \times \mathcal{Z}$ given by

$$(\Omega_{uv} \times \Omega_{yz})(u, y) = \Omega_{uv}(u) \times \Omega_{yz}(y).$$

B. Topological Properties

Suppose that the underlying sets are endowed with topologies so that we may speak of convergence. A map $\Omega : \mathcal{Y} \rightarrow 2^{\mathcal{Z}}$ is *closed* at the point \bar{y} in \mathcal{Y} whenever the statements $y_j \rightarrow \bar{y}$, $z_j \in \Omega(y_j)$, and $z_j \rightarrow \bar{z}$ together imply that $\bar{z} \in \Omega(\bar{y})$. One may interpret this definition as saying that the set $\Omega(\bar{y})$ is “bigger” than the sets in the sequence $\{\Omega(y_j)\}$. On the other hand, the map Ω is *open* at \bar{y} in \mathcal{Y} whenever the statements $y_j \rightarrow \bar{y}$ and $\bar{z} \in \Omega(\bar{y})$ together imply the existence of a number J and a sequence of points $\{z_j\}$ such that $z_j \rightarrow \bar{z}$ and $z_j \in \Omega(y_j)$ for all $j \geq J$. More or less, this statement means that the set $\Omega(\bar{y})$ is “smaller” than the sets in the sequence $\{\Omega(y_j)\}$. A map that is both open and closed at \bar{y} is said to be *continuous* at \bar{y} . We call Ω an *open map*, *closed map*, or *continuous map* whenever it has the corresponding property for every point in \mathcal{Y} .

Finite Cartesian products and finite compositions of open maps are open. Finite Cartesian products of closed maps are closed. If $\Omega_{yz} : \mathcal{Y} \rightarrow 2^{\mathcal{Z}}$ and $\Omega_{zw} : \mathcal{Z} \rightarrow 2^{\mathcal{W}}$ are closed and \mathcal{Z} is compact, then the composition $(\Omega_{zw} \circ \Omega_{yz})$ is closed.

C. Fixed Points

Suppose that Ω is a point-to-set map from \mathcal{Y} to itself. Let y be a point of \mathcal{Y} for which $\Omega(y) = \{y\}$. Then y is called a *fixed point* of the map Ω . In contrast, a *generalized fixed point* of Ω is a point for which $y \in \Omega(y)$. When we wish to emphasize the distinction, we may refer to a regular fixed point as a *strong* or *classical* fixed point.

D. Infimal Maps

Minimizing functions leads to a special type of point-to-set map. Suppose that $f : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ is a real-valued function of two variables, and let Ω be a point-to-set map from \mathcal{Y} to \mathcal{Z} . Associated with f and Ω is an *infimal map* defined by

$$M^z(y) \stackrel{\text{def}}{=} \arg \min_{z \in \Omega(y)} f(y, z).$$

If $f(y, \cdot)$ attains no minimal value on $\Omega(y)$, then $M^z(y) = \emptyset$, the empty set. Under mild conditions, infimal maps are closed.

Theorem 11 (Dantzig-Folkman-Shapiro [71]): If Ω is continuous at \bar{y} and $f(\bar{y}, \cdot)$ is continuous on $\Omega(\bar{y})$, then M^z is closed at \bar{y} .

In particular, the constant map $\Omega : y \mapsto \mathcal{Z}$ is continuous whenever \mathcal{Z} is closed. So minimizing a continuous function over a fixed, closed set always yields a closed infimal map.

E. Iterative Algorithms

Zangwill was apparently the first to recognize that many procedures in mathematical programming find their most natural expression in the language of point-to-set maps [68]. An *algorithmic map* or *algorithm* is simply a function $\Omega : \mathcal{Y} \rightarrow 2^{\mathcal{Y}}$. Given an initial point y_0 of \mathcal{Y} , an algorithmic map generates a sequence of iterates according to the rule

$$y_{j+1} \in \Omega(y_j).$$

Suppose that $f : \mathcal{Y} \rightarrow \mathbb{R}_+$ is a continuous, nonnegative function. We say that the algorithm Ω is *monotonic* with respect to f when

$$z \in \Omega(y) \quad \text{implies} \quad f(z) \leq f(y).$$

An algorithm *strictly monotonic* with respect to f is a monotonic algorithm for which

$$z \in \Omega(y) \text{ and } f(z) = f(y) \quad \text{imply} \quad z = y.$$

Zangwill showed that a closed, monotonic algorithm converges in a weak sense to a generalized fixed point. We present a streamlined version of his result.

Theorem 12 (Zangwill [68, p. 91]): Let Ω be a closed algorithmic map on a compact set \mathcal{Y} , and assume that Ω is monotonic with respect to a continuous, nonnegative function f . Suppose that the algorithm generates a sequence of iterates $\{y_j\}$.

- The sequence has at least one accumulation point in \mathcal{Y} .
- Each accumulation point \bar{y} satisfies $f(\bar{y}) = \lim_j f(y_j)$.
- Each accumulation point \bar{y} is a generalized fixed point of the algorithm.

R. R. Meyer subsequently extended Zangwill’s Theorem to provide a more satisfactory convergence result for strictly monotonic algorithms. One version of his result follows. For reference, a sequence $\{y_j\}$ in a normed space is called *asymptotically regular* when $\|y_{j+1} - y_j\| \rightarrow 0$.

Theorem 13 (Meyer [66]): Let \mathcal{Y} be a compact subset of a normed space, and assume that Ω is a closed algorithm on \mathcal{Y} that it is strictly monotonic with respect to the continuous, nonnegative function f . Suppose that Ω generates a sequence of iterates $\{y_j\}$. In addition to the conclusions of Zangwill’s Theorem, the following statements hold.

- Each accumulation point of the sequence is a (strong) fixed point of the algorithm.
- The sequence of iterates is asymptotically regular. In consequence, it has a continuum of accumulation points, or it converges in norm.
- In case that the fixed points of Ω on each isocontour of f form a discrete set, then the sequence of iterates converges in norm.

Under additional (burdensome) hypotheses, it is possible to prove norm convergence. A simple consequence of Meyer’s Theorem is that an algorithm Ω , strictly monotonic with respect to f , converges in norm if the fixed points of Ω on each isocontour of f form a discrete set. There is also a striking convergence theorem due to Opial [72], a proof of which appears in the recent paper [73].

F. Alternating Projection

An alternating projection can be interpreted as a kind of monotonic algorithm. Suppose that $f : \mathcal{Y} \times \mathcal{Z} \rightarrow \mathbb{R}_+$ is a continuous function. Then f induces two natural infimal maps,

$$M_y(z) \stackrel{\text{def}}{=} \arg \min_{y \in \mathcal{Y}} f(y, z) \quad \text{and}$$

$$M^z(y) \stackrel{\text{def}}{=} \arg \min_{z \in \mathcal{Z}} f(y, z).$$

If \mathcal{Y} and \mathcal{Z} are closed, then Theorem 11 shows that the maps M_y and M^z are both closed.

We interpret alternating projection as an algorithm on the product space $\mathcal{Y} \times \mathcal{Z}$ equipped with the usual product topology. Given an initial iterate y_0 from \mathcal{Y} , alternating projection generates a sequence of iterates $\{(y_j, z_j)\}$ via the rules

$$z_j \in M^z(y_j) \quad \text{and} \quad y_{j+1} \in M_y(z_j)$$

for each $j \geq 0$. Formally, this algorithm can be written as the composition of two sub-algorithms, Ω_{to} and Ω_{fro} , that are defined as

$$\Omega_{\text{to}} : (y, z) \mapsto \{y\} \times M^z(y) \quad \text{and}$$

$$\Omega_{\text{fro}} : (y, z) \mapsto M_y(z) \times \{z\}.$$

It follows that $\Omega \stackrel{\text{def}}{=} \Omega_{\text{fro}} \circ \Omega_{\text{to}}$ is a closed algorithm whenever \mathcal{Y} and \mathcal{Z} are compact. Both sub-algorithms decrease the value of f , so it should also be clear that Ω is monotonic with respect to f . Zangwill's Theorem tenders a basic convergence result.

Corollary 14: Let \mathcal{Y} and \mathcal{Z} be compact. Suppose that the alternating projection between \mathcal{Y} and \mathcal{Z} generates a sequence of iterates $\{(y_j, z_j)\}$.

- The sequence has at least one accumulation point.
- Each accumulation point of the sequence lies in $\mathcal{Y} \times \mathcal{Z}$.
- Each accumulation point is a generalized fixed point of the algorithm.
- Each accumulation point (\bar{y}, \bar{z}) satisfies $f(\bar{y}, \bar{z}) = \lim_j f(y_j, z_j)$.

If the infimal maps M_y and M^z are single-valued, we can achieve a much more satisfactory result.

Corollary 15: Let \mathcal{Y} and \mathcal{Z} be compact subsets of a normed space, and assume that the infimal maps M_y and M^z are single-valued. Suppose that the alternating projection between \mathcal{Y} and \mathcal{Z} generates a sequence of iterates $\{(y_j, z_j)\}$. In addition to the conclusions of Corollary 14, we have the following.

- Each accumulation point is a classical fixed point of the alternating projection.
- The sequence of iterates is asymptotically regular.
- The sequence of iterates either converges in norm or it has a continuum of accumulation points.

Proof: We just need to show that the algorithm is strictly monotonic with respect to f . Suppose that $f(y, z) = f(\Omega(y, z))$. Since the infimal maps never increase the value of f , we have the equalities

$$f(y, z) = f(y, M^z(y))$$

$$= f((M_y \circ M^z)(y), M^z(y)) = f(\Omega(y, z)).$$

Since M^z yields the *unique* minimizer of f with its first argument fixed, the first equality implies that $M^z(y) = \{z\}$. Likewise, the second equality yields $(M_y \circ M^z)(y) = \{y\}$. That is, $\Omega(y, z) = \{(y, z)\}$. An application of Meyer's Theorem completes the argument. ■

This result is a special case of a theorem of Fiorot and Huard [74]. In Appendix II, we will translate the language of these corollaries into more familiar terms.

G. Literature on Alternating Projection

Like most good ideas, alternating projection has a long biography and several aliases, including *alternating minimization*, *successive approximation*, *successive projection*, and *projection on convex sets*. This section offers a résumé of the research on alternating projection, but it makes no pretension to be comprehensive. Deutsch has written more detailed surveys, including [51], [75], [76].

According to Deutsch [75], alternating projection first appeared in a set of mimeographed lecture notes, written by John von Neumann in 1933. von Neumann proved that the alternating projection between two closed subspaces of a Hilbert space converges pointwise to the orthogonal projector onto their intersection [38]. Apparently, this theorem was not very well advertised, because many other authors have discovered it independently, including Aronszajn [48] and Wiener [77]. It was shown by Aronszajn [48] and Kayalar–Weinert [49] that both sequences of iterates converge geometrically with a rate exactly equal to the squared cosine of the (Friedrichs) principal angle between the two subspaces.

It is natural to extend the alternating projection between two subspaces by cyclically projecting onto several subspaces. Halperin demonstrated that, in a Hilbert space, the cyclic projection among a finite number of closed subspaces converges pointwise to the orthogonal projector onto their intersection [78]. The convergence is geometric [79]. Optimal bounds on the rate of convergence can be computed with techniques of Xu and Zikatonov [80]. Bauschke *et al.* study methods for accelerating cyclic projection in the recent paper [81].

It will come as no surprise that researchers have also studied alternating projection between subspaces of a Banach space. Unaware of von Neumann's work, Diliberto and Straus introduced an alternating method for computing the best sup-norm approximation of a bivariate continuous function as the sum of two univariate continuous functions, and they proved some weak convergence results [82]. The norm convergence of the sequence of iterates remained open until the work of Aumann [83]. M. Golomb extended the Diliberto–Straus algorithm to other best-approximation problems [84]. For more information on alternating algorithms in Banach spaces, see the monograph of Cheney and Light [85].

Another fruitful generalization is to consider projection onto convex subsets. The projector—or *proximity map*—onto a closed, convex subset of a Hilbert space is well-defined, because each point has a unique best approximation from that set. The basic result, due to Cheney and Goldstein, is that the alternating projection between two closed, convex subsets of a Hilbert space will converge to a pair of points

at minimal distance from each other, so long as one set is compact [4]. Dykstra [86], [87] and Han [88] independently developed a cyclic projection technique that, given a point, can compute its best approximation from the intersection of a finite number of closed, convex sets in a Hilbert space. Their algorithm requires a correction to each projection. Their method is closely related to earlier optimization techniques of Hildreth [89] and Bregman [5]. Details of this connection appear in the book [90] and the article [91]. Tseng develops a very important extension of the Dykstra–Han algorithm in [92]. To date, the most detailed treatment of projection on convex sets is probably the survey article [93].

Most of the work on alternating projection has involved the Euclidean distance, but it is possible to develop results for other divergence measures. In particular, Csiszár and Tusnády have shown that alternate minimization of the Kullback–Leibler divergence can be used to find a pair of minimally distant points contained within two convex sets of probability measures [67].

There has been some research on alternating projection between nonconvex sets, but the theoretical results so far are limited. Fiorot and Huard have applied the theorems of Zangwill and Meyer to obtain weak convergence results for a class of block relaxation schemes that include alternating and cyclic projection onto non-convex sets [74]. Combettes and Trussell have developed a technique which inflates the non-convex sets into convex sets; they offer some qualified convergence results [94]. Cadzow has also demonstrated empirically that cyclic projections among nonconvex sets can effectively solve some signal enhancement problems [95]. More research in this direction would be valuable.

Alternating projection has found application to many different problems, of which we offer a (small) selection. The most famous example from these pages must be the Blahut–Arimoto algorithm for computing channel capacity and rate-distortion functions [96], [97]. In the field of signal restoration and recovery, we mention the work of Landau–Miranker [98], Gerchberg [99], Youla–Webb [100], Cadzow [95], and Donoho–Stark [101]. Çetin, Gerek, and Yardimci show that projection on convex sets can compute multi-dimensional equiripple filters [102]. Xu and Zikatonov discuss how alternating projection can be used to solve the linear systems that arise in the discretization of partial differential equations [80]. In the matrix analysis community, alternating projection has been used as a computational method for solving inverse eigenvalue problems [35], [37] and for solving matrix nearness problems [36], [103]. In statistics, one may view the Expectation Maximization (EM) algorithm as an alternating projection [104]. de Leeuw has discussed other statistical applications in [50].

APPENDIX II CONVERGENCE AND FIXED POINTS

Armed with the theory of the last appendix, we are finally girded to attack the convergence of Algorithm 1. The results on point-to-set maps will allow us to dispatch this dragon quickly. Then we will turn our attention to the convergence of the algorithm in the special case that the frame vectors have

prescribed norms. This problem will require a longer siege, but it too will yield to our onslaught. The convergence results that we develop here are all novel.

A. Basic Convergence Proof

In this section, we establish the convergence of the basic alternating projection algorithm that appears in Section III-D. Our main burden is to translate the language of point-to-set maps into more familiar terms.

Theorem 16 (Global Convergence): Let \mathcal{Y} and \mathcal{Z} be closed sets, one of which is bounded. Suppose that alternating projection generates a sequence of iterates $\{(Y_j, Z_j)\}$. This sequence possesses at least one accumulation point, say (\bar{Y}, \bar{Z}) .

- The accumulation point lies in $\mathcal{Y} \times \mathcal{Z}$.
- The accumulation point satisfies

$$\|\bar{Y} - \bar{Z}\|_F = \lim_{j \rightarrow \infty} \|Y_j - Z_j\|_F.$$

- The accumulation point satisfies

$$\|\bar{Y} - \bar{Z}\|_F = \text{dist}(\bar{Y}, \mathcal{Z}) = \text{dist}(\bar{Z}, \mathcal{Y}).$$

Proof: Assume without loss of generality that \mathcal{Y} is the compact set, while \mathcal{Z} is merely closed. We must establish that we have all the compactness necessary to apply Corollary 14.

Without loss of generality, assume that $Y_0 \in \mathcal{Y}$. If $\delta = \|Y_0 - Z_0\|_F$, then subsequent iterates always satisfy

$$\begin{aligned} \|Y_j - Z_j\|_F &\leq \delta & \text{and} \\ \|Y_{j+1} - Z_j\|_F &\leq \delta. \end{aligned}$$

Thus, we may restrict our attention to the sets

$$\begin{aligned} \mathcal{Y}_1 &= \{Y \in \mathcal{Y} : \text{dist}(Y, \mathcal{Z}) \leq \delta\} & \text{and} \\ \mathcal{Z}_1 &= \{Z \in \mathcal{Z} : \text{dist}(Z, \mathcal{Y}) \leq \delta\}. \end{aligned}$$

Since \mathcal{Y} is compact, \mathcal{Y}_1 is compact because it is a closed subset of a compact set. On the other hand, \mathcal{Z}_1 is compact because it is the intersection of the closed set \mathcal{Z} with a compact set, namely the collection of matrices within a fixed distance of \mathcal{Y} .

We may apply Corollary 14. Each of the conclusions of the corollary has a straightforward analogue among the conclusions of the present theorem. The only question that may remain is what it means for a pair of matrices (\bar{Y}, \bar{Z}) to be a generalized fixed point of the alternating projection. A generalized fixed point of an algorithm is a point which is a possible successor of itself. In the present case, a pair of matrices can succeed itself if and only if the second component is a potential successor of the first *and* the first component is a potential successor of the second. The matrix \bar{Z} can succeed the matrix \bar{Y} if and only if

$$\|\bar{Z} - \bar{Y}\|_F = \text{dist}(\bar{Y}, \mathcal{Z}).$$

Likewise, \bar{Y} can succeed \bar{Z} if and only if

$$\|\bar{Y} - \bar{Z}\|_F = \text{dist}(\bar{Z}, \mathcal{Y}).$$

This observation completes the proof. ■

Since the collection of α -tight frames and the collection of their Gram matrices are both compact, the theorem has two immediate corollaries.

Corollary 17: If \mathcal{X}_α is the collection of α -tight frames, and \mathcal{S} is a closed set of matrices, then Theorem 16 applies with $\mathcal{Y} = \mathcal{S}$ and $\mathcal{Z} = \mathcal{X}_\alpha$.

Corollary 18: If \mathcal{G}_α contains the Gram matrices of all α -tight frames, and \mathcal{H} is a closed set of Hermitian matrices, then Theorem 16 applies with $\mathcal{Y} = \mathcal{G}_\alpha$ and $\mathcal{Z} = \mathcal{H}$.

B. Stronger Convergence Results

Meyer's Theorem suggests that it might be possible to provide a stronger convergence result for Algorithm 1 if we can ensure that the matrix nearness problems have unique solutions. In many cases, the nearness problems are uniquely soluble whenever the iterates get sufficiently close together. This provides a local convergence result that is much stronger than Zangwill's Theorem allows. First, we prove a general version of this result. Afterward, we show that it applies to an alternating projection that involves one of the spectral constraint sets \mathcal{X}_α or \mathcal{G}_α .

Recall that the distance between a matrix M and a set \mathcal{Y} is defined as

$$\text{dist}(M, \mathcal{Y}) \stackrel{\text{def}}{=} \inf_{Y \in \mathcal{Y}} \|M - Y\|_F.$$

Theorem 19: Let \mathcal{Y} and \mathcal{Z} be closed sets of matrices, one of which is compact. Suppose that the alternating projection between \mathcal{Y} and \mathcal{Z} generates a sequence of iterates $\{(Y_j, Z_j)\}$, and assume that the matrix nearness problems

$$\begin{aligned} \min_{Y \in \mathcal{Y}} \|Y - M\|_F \\ \min_{Z \in \mathcal{Z}} \|Z - M\|_F \end{aligned}$$

have unique solutions for any matrix M in the sequence of iterates. Then we reach the following conclusions.

- The sequence of iterates possesses at least one accumulation point, say (\bar{Y}, \bar{Z}) .
- The accumulation point lies in $\mathcal{Y} \times \mathcal{Z}$.
- The pair (\bar{Y}, \bar{Z}) is a fixed point of the alternating projection. In other words, if we applied the algorithm to \bar{Y} or to \bar{Z} every iterate would equal (\bar{Y}, \bar{Z}) .
- The accumulation point satisfies

$$\|\bar{Y} - \bar{Z}\|_F = \lim_{j \rightarrow \infty} \|Y_j - Z_j\|_F.$$

- The component sequences are asymptotically regular, i.e.,

$$\|Y_{j+1} - Y_j\|_F \rightarrow 0 \quad \text{and} \quad \|Z_{j+1} - Z_j\|_F \rightarrow 0.$$

- Either the component sequences both converge in norm,

$$\|Y_j - \bar{Y}\|_F \rightarrow 0 \quad \text{and} \quad \|Z_j - \bar{Z}\|_F \rightarrow 0,$$

or the set of accumulation points forms a continuum.

Proof: The argument in the proof of Theorem 16 shows that we are performing an alternating minimization between two compact sets. The hypotheses of the theorem guarantee that each iterate is uniquely determined by the previous iterate. Corollaries 14 and 15 furnish the stated conclusions.

The only point that may require clarification is what it takes for a pair of matrices (\bar{Y}, \bar{Z}) to be a classical fixed point of the alternating projection. A classical fixed point of an algorithm is the *only* possible successor of itself. In the case of alternating projection, the matrix \bar{Z} must be the unique successor of the \bar{Y} , and the matrix \bar{Y} must be the unique successor of \bar{Z} . This observation completes the argument. ■

Due to the peculiar structure of the spectral constraint sets \mathcal{X}_α and \mathcal{G}_α , the solutions to the associated matrix nearness problems are often unique. Therefore, the alternating projection algorithms that we have considered in this paper sometimes have better performance than the basic convergence result, Theorem 16, would predict.

We remind the reader that

$$\begin{aligned} \mathcal{X}_\alpha &\stackrel{\text{def}}{=} \{X \in \mathbb{C}^{d \times N} : XX^* = \alpha I_d\}, \quad \text{and} \\ \mathcal{G}_\alpha &\stackrel{\text{def}}{=} \{G \in \mathbb{C}^{N \times N} : G = G^*, \\ &\quad \text{and } G \text{ has eigenvalues } \underbrace{(\alpha, \dots, \alpha, 0, \dots, 0)}_d\}. \end{aligned}$$

The uniqueness of the matrix nearness problems will follow from the Wielandt–Hoffman Theorem, a powerful result from matrix analysis.

Theorem 20 (Wielandt–Hoffman [25, p. 368 and p. 419]):

Suppose that A and B are $N \times N$ Hermitian matrices, and let the vectors $\lambda(A)$ and $\lambda(B)$ list the eigenvalues of A and B in algebraically nonincreasing order. Then

$$\|\lambda(A) - \lambda(B)\|_2 \leq \|A - B\|_F.$$

Suppose instead that A and B are $d \times N$ rectangular matrices with $d \leq N$, and let $\sigma(A)$ and $\sigma(B)$ list the largest d singular values of A and B in nonincreasing order. Then

$$\|\sigma(A) - \sigma(B)\|_2 \leq \|A - B\|_F. \quad (9)$$

Note that if we solving matrix nearness problems with respect to the spectral norm, Weyl's Theorem would allow us to provide stronger bounds [25, p. 367].

Corollary 21 (Local Convergence with Constraint \mathcal{X}_α):

Let \mathcal{S} be a closed set of $d \times N$ matrices for which the associated matrix nearness problem

$$\min_{S \in \mathcal{S}} \|S - M\|_F$$

has a unique solution whenever $\text{dist}(M, \mathcal{S}) < \varepsilon$. Imagine that the alternating projection between \mathcal{S} and \mathcal{X}_α generates a sequence of iterates $\{(S_j, X_j)\}$ in which

$$\|S_j - X_j\|_F < \min\{\varepsilon, \alpha\} \quad \text{for some index } J.$$

Then the conclusions of Theorem 19 are in force.

Proof: According to Theorem 2, the matrix in \mathcal{X}_α nearest to a matrix M is unique so long as M has full rank. A $d \times N$ matrix is rank-deficient only if its d -th largest singular value is zero. Observe that the largest d singular values of each matrix in \mathcal{X}_α all equal $\alpha > 0$. According to the Wielandt–Hoffman Theorem, any matrix sufficiently close to \mathcal{X}_α cannot be rank-deficient. More precisely, $\text{dist}(M, \mathcal{X}_\alpha) < \alpha$ implies that M has full rank, which in turn shows that M has a unique best approximation in \mathcal{X}_α .

Define the constraint sets

$$\begin{aligned}\mathcal{Y} &\stackrel{\text{def}}{=} \mathcal{S} \cap \text{closure}\{S_j : j \geq J\} \quad \text{and} \\ \mathcal{Z} &\stackrel{\text{def}}{=} \mathcal{X}_\alpha \cap \text{closure}\{X_j : j \geq J\}.\end{aligned}$$

Note that \mathcal{Y} is closed and that \mathcal{Z} is compact. We will apply Theorem 19 to the tail of the sequence of iterates, beginning with index J . For $j \geq J$, each matrix S_j is close enough to \mathcal{Z} and each matrix X_j is close enough to \mathcal{Y} that we can ensure the matrix nearness problems have unique solutions. ■

Corollary 22 (Local Convergence with Constraint \mathcal{G}_α):

Let \mathcal{H} be a closed set of $N \times N$ matrices for which the associated matrix nearness problem

$$\min_{H \in \mathcal{H}} \|H - M\|_{\text{F}}$$

has a unique solution whenever $\text{dist}(M, \mathcal{H}) < \varepsilon$. Imagine that the alternating projection between \mathcal{G}_α and \mathcal{H} generates a sequence of iterates $\{(G_j, H_j)\}$ in which

$$\|G_j - H_j\|_{\text{F}} < \min\{\varepsilon, \alpha/\sqrt{2}\} \quad \text{for some index } J.$$

Then the conclusions of Theorem 19 are in force.

Proof: Theorem 3 indicates that the matrix in \mathcal{G}_α nearest to a matrix M is unique so long as its d -th and $(d+1)$ -st eigenvalues are distinct. Imagine that M is a matrix whose d -th and $(d+1)$ -st eigenvalues both equal τ . Since the d -th and $(d+1)$ -st eigenvalue of a matrix in \mathcal{G}_α are α and zero, the Wielandt–Hoffman Theorem shows that

$$\text{dist}(M, \mathcal{G}_\alpha)^2 \geq (\alpha - \tau)^2 + \tau^2.$$

Varying τ , the minimum value of the right-hand side is $\alpha^2/2$. Therefore, $\text{dist}(M, \mathcal{G}_\alpha) < \alpha/\sqrt{2}$ implies that the d -th and $(d+1)$ -st eigenvalues of M are distinct. In consequence, M has a unique best approximation from \mathcal{G}_α .

As before, define the constraint sets

$$\begin{aligned}\mathcal{Y} &\stackrel{\text{def}}{=} \mathcal{H} \cap \text{closure}\{H_j : j \geq J\} \quad \text{and} \\ \mathcal{Z} &\stackrel{\text{def}}{=} \mathcal{G}_\alpha \cap \text{closure}\{G_j : j \geq J\}.\end{aligned}$$

The set \mathcal{Y} is closed, and \mathcal{Z} is compact. We will apply Theorem 19 to the tail of the sequence of iterates, beginning with index J . For $j \geq J$, each matrix H_j is close enough to \mathcal{Z} and each matrix G_j is close enough to \mathcal{Y} that we can ensure the matrix nearness problems have unique solutions. ■

C. Specified Column Norms

This section offers a detailed analysis of the alternating projection between the set of α -tight frames and the collection of matrices with specified column norms.

Let c_1, \dots, c_N be positive numbers that denote the squared column norms we desire in the frame. Without loss of generality, we assume that $\alpha = \sum_n c_n/d = 1$ to streamline the proofs. Then the structural constraint set is

$$\mathcal{S} \stackrel{\text{def}}{=} \{S \in \mathbb{C}^{d \times N} : \|s_n\|_2^2 = c_n\}. \quad (10)$$

Since the tightness parameter α of the frame equals one, we define the set of 1-tight frames as

$$\mathcal{X}_1 \stackrel{\text{def}}{=} \{X \in \mathbb{C}^{d \times N} : XX^* = I_d\}.$$

Suppose that S_0 is a full-rank matrix drawn from \mathcal{S} , and perform an alternating projection between the sets \mathcal{S} and \mathcal{X}_1 to obtain sequences $\{S_j\}$ and $\{X_j\}$. Proposition 23 of the sequel shows that the sequence $\{S_j\}$ lies in a compact subset of \mathcal{S} whose elements have full rank, while the sequence $\{X_j\}$ lies in a compact subset of \mathcal{X}_1 whose elements have nonzero columns. By an appeal to the matrix nearness results, Theorem 2 and Proposition 5, we see that each iterate is uniquely determined by its predecessor. We may therefore apply Corollary 15.

In this subsection, we complete the foregoing argument by demonstrating that the iterates are well-behaved. In the next subsection, we classify the full-rank fixed points of the alternating projection between \mathcal{S} and \mathcal{X}_1 .

Set $c_{\min} = \min_n c_n$, and define the diagonal matrix C whose entries are $\sqrt{c_1}, \dots, \sqrt{c_N}$.

Proposition 23: Assume that the initial iterate S_0 is a full rank matrix from \mathcal{S} . For every positive index j ,

- 1) the Euclidean norm of each column of X_j is at least $\sqrt{c_{\min}}/\|C\|_{\text{F}}$; and
- 2) the smallest singular value of S_j is at least $\sqrt{c_{\min}}$.

The matrices that satisfy these bounds form compact subsets of \mathcal{X}_1 and \mathcal{S} .

Proof: Assume that $j \geq 0$, and make the inductive assumption that S_j has full rank. First, we bound the top singular value of S_j by exploiting the relationship between the singular values of a matrix and its Frobenius norm. Since C lists the column norms of S_j , it follows that $\|S_j\|_{\text{F}}^2 = \|C\|_{\text{F}}^2$. The squared Frobenius norm of S_j also equals the sum of its squared singular values. It is immediate that the maximum singular value of S_j satisfies

$$\sigma_{\max}(S_j)^2 \leq \|C\|_{\text{F}}^2. \quad (11)$$

Next we use this relation to estimate the column norms of X_j . Let S_j have singular value decomposition $U\Sigma V^*$, and write the n -th columns of S_j and X_j as s_n and x_n . On account of the fact that $X_j = (S_j S_j^*)^{-1/2} S_j$, we have

$$\begin{aligned}\|x_n\|_2 &= \left\| (S_j S_j^*)^{-1/2} s_n \right\|_2 \\ &= \left\| U\Sigma^{-1} U^* s_n \right\|_2 \\ &\geq \sqrt{c_{\min}}/\sigma_{\max}(S_j),\end{aligned} \quad (12)$$

since the norm of s_n is at least $\sqrt{c_{\min}}$. Introducing the bound (11) into inequality (12) yields the first part of the proposition.

Now, we show that the smallest singular value of S_{j+1} remains well away from zero. Define T_j to be the diagonal matrix that lists the column norms of X_j . Note that, since X_j is a submatrix of a unitary matrix, its column norms cannot exceed one, and so every entry of T_j^{-1} must be at least one. We can now express the matrix in \mathcal{S} nearest to X_j with the formula $S_{j+1} = X_j T_j^{-1} C$. It is well known that the smallest singular value of S_{j+1} is equal to the square root of the smallest eigenvalue of $S_{j+1} S_{j+1}^*$. We will apply the Rayleigh–Ritz Theorem [25, p. 176] to provide a lower bound on this eigenvalue. Let v be any nonzero, d -dimensional vector, and form the Rayleigh quotient

$$\frac{v^*(S_{j+1} S_{j+1}^*) v}{v^* v} = \frac{v^*(X_j T_j^{-1} C^2 T_j^{-1} X_j^*) v}{v^* v}.$$

Make the substitution $w = X_j^* v$, and observe that the $w^* w = v^* (X_j X_j^*) v = v^* v$ since the matrix X_j has orthonormal rows. Therefore, the Rayleigh quotient satisfies

$$\frac{v^* (S_{j+1} S_{j+1}^*) v}{v^* v} = \frac{w^* (T_j^{-1} C^2 T_j^{-1}) w}{w^* w}.$$

It follows that the smallest eigenvalue of S_{j+1} is no less than the smallest eigenvalue of $T_j^{-1} C^2 T_j^{-1}$. But this matrix is diagonal with entries no smaller than c_{\min} . Its smallest eigenvalue is no less than c_{\min} . We conclude

$$\sigma_{\min}(S_{j+1}) \geq \sqrt{c_{\min}},$$

and the second part of the proposition is complete.

Finally, we must make the compactness argument. We have shown that the squared singular values of an iterate S_j must lie in the closed interval $[\sqrt{c_{\min}}, \|C\|_F]$. The minimum squared singular value of a matrix is a continuous function of the matrix entries, which follows from conclusion (9) of the Wielandt–Hoffman Theorem. Therefore, the matrices whose smallest singular value lies in this interval form a closed set. We discover that the intersection of this set with the compact set \mathcal{S} is compact. The same type of argument implies that the sequence $\{X_j\}$ lies in a compact subset of \mathcal{X}_1 whose matrices have column norms bounded away from zero. ■

D. Fixed Points I

It remains to characterize the fixed points of the alternating projection between the set of matrices with fixed column norms and the set of α -tight frames.

Proposition 24: The full-rank stationary points of an alternating projection between \mathcal{S} and \mathcal{X}_α are precisely those full-rank matrices S from \mathcal{S} whose columns are all eigenvectors of SS^* . That is, $SS^* S = S\Lambda$, where $\Lambda \in \mathbb{C}^{N \times N}$ is diagonal and positive with at most d distinct values.

Proof: As before, we take $\alpha = 1$ for simplicity. Define the diagonal matrix $T = T(S)$ whose entries are the column norms of $(SS^*)^{-1/2} S$.

Suppose that S is a full-rank fixed point of the algorithm. Thus projecting S onto \mathcal{X}_1 and projecting back to \mathcal{S} returns S . Symbolically,

$$S = ((SS^*)^{-1/2} S) (T^{-1} C).$$

Define $\Lambda = T^{-1} C$. Then the equation becomes $(SS^*)^{-1/2} S = S\Lambda^{-1}$. Due to the joint eigenstructure of a positive-definite matrix and its positive-definite roots [25, p. 405], it follows that $(SS^*) S = S\Lambda^2$.

Conversely, suppose that S has full rank and that $(SS^*) S = S\Lambda^2$ for some positive diagonal matrix Λ . Equivalently, $(SS^*) s_n = \lambda_n^2 s_n$ for each n . It follows that

$$(SS^*)^{-1/2} s_n = \lambda_n^{-1} s_n \quad \text{for each } n.$$

Take norms to see that $t_n = \lambda_n^{-1} \sqrt{c_n}$. Combine these equations into the matrix equation $\Lambda = T^{-1} C$. Thus, S is a fixed point of the algorithm. ■

E. Fixed Points II

Proposition 24 allows us to provide a partial characterization of the fixed points of any alternating projection between the set of α -tight frames \mathcal{X}_α and any structural constraint set \mathcal{Z} that contains matrices with fixed column norms. This result applies even if the matrices in \mathcal{Z} have additional properties.

Proposition 25: Suppose that the column norms of matrices in \mathcal{Z} are fixed. A sufficient condition for a full-rank matrix Z in \mathcal{Z} to be a fixed point of the alternating projection between \mathcal{Z} and \mathcal{X}_α is that the columns of Z are all eigenvectors of ZZ^* . That is, $ZZ^* Z = Z\Lambda$, where Λ is a positive, diagonal matrix with no more than d distinct entries.

Proof: Let \mathcal{S} be the set defined in (10), and let \mathcal{Z} be a closed subset of \mathcal{S} . Suppose that the matrix Z in \mathcal{Z} is a fixed point of the alternating projection between \mathcal{S} and \mathcal{X}_α , and let X be the matrix in \mathcal{X}_α closest to Z . It follows that Z is the (unique) matrix in \mathcal{S} closest to X . In particular, Z is the matrix in \mathcal{Z} closest to X . Therefore, Z is also a fixed point of the alternating projection between \mathcal{Z} and \mathcal{X}_α . An appeal to Proposition 24 completes the proof. ■

REFERENCES

- [1] O. Christensen, *An introduction to frames and Riesz bases*. Boston: Birkhäuser, 2003.
- [2] J. A. Tropp, R. W. Heath, and T. Strohmer, "Optimal CDMA signature sequences, inverse eigenvalue problems and alternating projection," in *Proceedings of the 2003 IEEE International Symposium on Information Theory*, Yokohama, July 2003, p. 407.
- [3] J. A. Tropp, I. S. Dhillon, R. W. Heath, and T. Strohmer, "CDMA signature sequences with low peak-to-average-power ratio via alternating projection," in *Proceedings of the 37th Annual Asilomar Conference on Signals, Systems and Computers*, Monterey, Nov. 2003.
- [4] E. W. Cheney and A. A. Goldstein, "Proximity maps for convex sets," *Proc. Amer. Math. Soc.*, vol. 10, no. 3, pp. 448–450, June 1959.
- [5] L. M. Bregman, "The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming," *USSR Comput. Math. Math. Phys.*, vol. 7, no. 3, pp. 200–217, 1967.
- [6] M. Rupf and J. L. Massey, "Optimum sequence multisets for synchronous code-division multiple-access channels," *IEEE Trans. Inform. Theory*, vol. 40, no. 4, pp. 1261–1266, July 1994.
- [7] P. Viswanath and V. Anantharam, "Optimal sequences and sum capacity of synchronous CDMA systems," *IEEE Trans. Inform. Theory*, vol. 45, no. 6, pp. 1984–1991, Sept. 1999.
- [8] P. Viswanath and A. Anantharam, "Optimal sequences for CDMA under colored noise: A Schur-saddle function property," *IEEE Trans. Inform. Theory*, vol. 48, no. 6, pp. 1295–1318, June 2002.
- [9] L. Zheng and D. N. Tse, "Communication on the Grassmann manifold: A geometric approach to the non-coherent multiple-antenna channel," *IEEE Trans. Inform. Theory*, vol. 48, no. 2, pp. 359–383, Feb. 2002.
- [10] R. W. Heath, T. Strohmer, and A. J. Paulraj, "On quasi-orthogonal signatures for CDMA systems," Sept. 2002, submitted to *IEEE Trans. Inform. Theory*.
- [11] V. Tarokh and I.-M. Kim, "Existence and construction of noncoherent unitary space-time codes," *IEEE Trans. Inform. Theory*, vol. 48, no. 12, pp. 3112–3117, Dec. 2002.
- [12] T. Strohmer and R. W. Heath, "Grassmannian frames with applications to coding and communication," *Appl. Comp. Harmonic Anal.*, vol. 14, no. 3, pp. 257–275, May 2003.
- [13] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Trans. Inform. Th.*, vol. 49, no. 12, pp. 3320–3325, Dec. 2003.
- [14] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inform. Theory*, vol. 50, pp. 2231–2242, Oct. 2004.
- [15] D. L. Donoho and M. Elad, "Maximal sparsity representation by ℓ_1 minimization," *Proc. Natl. Acad. Sci.*, vol. 100, pp. 2197–2202, Mar. 2003.

- [16] D. V. Sarwate, "Meeting the Welch Bound with equality," in *Sequences and their Applications*. London: Springer, 1998, pp. 79–102.
- [17] P. G. Casazza and M. T. Leon, "Existence and construction of finite tight frames," 2002, unpublished.
- [18] J. A. Tropp, I. S. Dhillon, and R. W. Heath, "Finite-step algorithms for constructing optimal CDMA signature sequences," *IEEE Trans. Inform. Theory*, vol. 50, no. 11, Nov. 2004, to appear.
- [19] C. Rose, "CDMA codeword optimization: interference avoidance and class warfare," *IEEE Trans. Inform. Theory*, vol. 47, no. 6, pp. 2368–2382, Sept. 2001.
- [20] S. Ulukus and R. D. Yates, "Iterative construction of optimum signature sequence sets in synchronous CDMA systems," *IEEE Trans. Inform. Theory*, vol. 47, no. 5, pp. 1989–1998, 2001.
- [21] C. Rose, S. Ulukus, and R. D. Yates, "Wireless systems and interference avoidance," *IEEE Trans. Wireless Comm.*, vol. 1, no. 3, pp. 415–428, Jul. 2002.
- [22] S. Ulukus and A. Yener, "Iterative joint optimization of CDMA signature sequence and receiver filters," in *2002 Conference on Information Sciences and Systems*, Princeton Univ., Mar. 2002.
- [23] P. Anigstein and V. Anantharam, "Ensuring convergence of the MMSE iteration for interference avoidance to the global optimum," *IEEE Trans. Inform. Th.*, vol. 49, no. 4, pp. 873–885, Apr. 2003.
- [24] L. R. Welch, "Lower bounds on the maximum cross-correlation of signals," *IEEE Trans. Inform. Theory*, vol. 20, pp. 397–399, 1974.
- [25] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 1985.
- [26] P. G. Casazza, M. Leon, and J. C. Tremain, "Custom built tight frames," 2002, unpublished.
- [27] W. M. Boothby, *An Introduction to Differentiable Manifolds and Riemannian Geometry*. Academic Press, 2003.
- [28] J. H. Conway, R. H. Hardin, and N. J. A. Sloane, "Packing lines, planes, etc.: Packings in Grassmannian spaces," *Experimental Math.*, vol. 5, no. 2, pp. 139–159, 1996.
- [29] M. T. Chu, "Inverse eigenvalue problems," *SIAM Rev.*, vol. 40, no. 1, pp. 1–39, 1998.
- [30] M. T. Chu and G. H. Golub, "Structured inverse eigenvalue problems," *Acta Numerica*, 2001.
- [31] S. Friedland, J. Nocedal, and M. L. Overton, "The formulation and analysis of numerical methods for inverse eigenvalue problems," *SIAM J. Numer. Anal.*, pp. 634–667, 1987.
- [32] M. T. Chu, "Numerical methods for inverse singular value problems," *SIAM J. Numer. Anal.*, vol. 29, pp. 885–903, 1992.
- [33] —, "On the refinement of a Newton method for the inverse Toeplitz eigenvalue problem," North Carolina State University, Tech. Rep., 1994.
- [34] M. T. Chu and K. R. Driessel, "The projected gradient method for least squares matrix approximations with spectral constraints," *SIAM J. Numer. Anal.*, vol. 27, no. 4, pp. 1050–1060, Aug. 1990.
- [35] M. T. Chu, "Constructing a Hermitian matrix from its diagonal entries and eigenvalues," *SIAM J. Matrix Anal. Appl.*, vol. 16, pp. 207–217, 1995.
- [36] N. J. Higham, "Computing the nearest correlation matrix—a problem from finance," *IMA J. Numer. Anal.*, vol. 22, pp. 329–343, 2002.
- [37] M. T. Chu, R. Funderlic, and R. J. Plemmons, "Structured low-rank approximation," *Linear Algebra Appl.*, vol. 366, pp. 157–172, 2003.
- [38] J. von Neumann, *Functional Operators, Vol. II*, ser. Annals of Mathematics Studies. Princeton Univ. Press, 1950, no. 22.
- [39] G. H. Golub, A. Hoffman, and G. W. Stewart, "A generalization of the Eckart-Young-Mirsky matrix approximation theorem," *Linear Algebra Appl.*, vol. 88/89, pp. 317–327, 1987.
- [40] A. Ruhe, "Closest normal matrix finally found!" *BIT*, vol. 27, pp. 585–598, 1987.
- [41] N. J. Higham, "Computing a nearest symmetric positive semi-definite matrix," *Linear Algebra Appl.*, vol. 103, pp. 103–118, 1988.
- [42] —, *Applications of Matrix Theory*. Oxford Univ. Press, 1989, ch. Matrix nearness problems and applications, pp. 1–27.
- [43] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Johns Hopkins University Press, 1996.
- [44] B. F. Green, "The orthogonal approximation of an oblique structure in factor analysis," *Psychometrika*, vol. 17, pp. 429–440, 1952.
- [45] P. H. Schönemann, "A generalized solution of the orthogonal Procrustes problem," *Psychometrika*, vol. 31, pp. 1–10, 1966.
- [46] N. J. Higham, "Computing the polar decomposition—with applications," *SIAM J. Sci. Stat. Comput.*, no. 7, pp. 1160–1174, 1986.
- [47] N. J. Higham and P. Papadimitriou, "A parallel algorithm for computing the polar decomposition," *Parallel Comput.*, vol. 20, no. 8, pp. 1161–1173, Aug. 1994.
- [48] N. Aronszajn, *Introduction to the Theory of Hilbert Spaces, Vol. I*. Stillwater, OK: Research Foundation of Oklahoma A & M College, 1950.
- [49] S. Kayalar and H. Weinert, "Error bounds for the method of alternating projections," *Math. Control Signal Systems*, pp. 43–59, 1988.
- [50] J. de Leeuw, *Information Systems and Data Analysis*. Berlin: Springer, 1994, ch. Block relaxation algorithms in statistics, pp. 308–325.
- [51] F. Deutsch, *Best Approximation in Inner-Product Spaces*. New York: Springer-Verlag, 2001.
- [52] J. J. Benedetto and M. Fickus, "Finite normalized tight frames," *Adv. Comp. Math.*, vol. 18, no. 2–4, pp. 357–385, 2003.
- [53] P. Casazza and J. Kovačević, "Equal-norm tight frames with erasures," *Adv. Comp. Math.*, vol. 18, no. 2–4, pp. 387–430, 2003.
- [54] V. K. Goyal, M. Vetterli, and N. T. Thao, "Quantized overcomplete expansions in \mathbb{R}^n : analysis, synthesis and algorithms," *IEEE Trans. Inform. Theory*, vol. 44, no. 1, pp. 16–31, 1998.
- [55] M. Sustik, J. A. Tropp, I. S. Dhillon, and R. W. Heath, "On the existence of equiangular tight frames," University of Texas at Austin, Dept. of Computer Sciences TR 04-32, August 2004.
- [56] A. R. Calderbank, P. J. Cameron, W. M. Kantor, and J. J. Seidel, " \mathbb{Z}_4 -Kerdock codes, orthogonal spreads and extremal Euclidean line sets," *Proc. London Math. Soc.*, vol. 75, no. 2, pp. 436–480, 1997.
- [57] N. J. A. Sloane, "Packing planes in four dimensions and other mysteries," in *Proceedings of the Conference on Algebraic Combinatorics and Related Topics*, Yamagata, Japan, Nov. 1997.
- [58] D. Agrawal, T. J. Richardson, and R. L. Urbanke, "Multiple-antenna signal constellations for fading channels," *IEEE Trans. Inform. Theory*, vol. 47, no. 6, pp. 2618–2626, Sept. 2001.
- [59] J. A. Tropp, "Topics in sparse approximation," Ph.D. dissertation, Computational and Applied Mathematics, The University of Texas at Austin, August 2004.
- [60] V. K. N. Lau, "On the analysis of peak-to-average ratio (PAR) for IS95 and CDMA2000d systems," *IEEE Trans. Vehicular Tech.*, vol. 49, no. 6, pp. 2174–2188, Nov. 2000.
- [61] R. T. Rockafellar, *Convex Analysis*. Princeton Univ. Press, 1970.
- [62] S. P. Lloyd, "Least squares quantization in PCM," Bell Laboratories, Technical Note, 1957.
- [63] I. S. Dhillon and D. S. Modha, "Concept decompositions for large sparse text data using clustering," *Machine Learning*, vol. 42, no. 1, pp. 143–175, 2001.
- [64] I. S. Dhillon, E. Marcotte, and U. Roshan, "Diametrical clustering for identifying anti-correlated gene clusters," *Bioinformatics*, vol. 19, pp. 1612–1619, 2003.
- [65] B. M. Hochwald, T. L. Marzetta, T. J. Richardson, W. Sweldens, and R. Urbanke, "Systematic design of unitary space-time constellations," *IEEE Trans. Inform. Theory*, vol. 46, no. 6, pp. 1962–1973, 2000.
- [66] R. R. Meyer, "Sufficient conditions for the convergence of monotonic mathematical programming algorithms," *J. Comp. Sys. Sci.*, vol. 12, pp. 108–121, 1976.
- [67] I. Csizsár and G. Tusnády, "Information geometry and alternating minimization procedures," *Statistics and Decisions*, vol. 1, pp. 205–237, 1984, supplemental issue.
- [68] W. I. Zangwill, *Nonlinear Programming: A Unified Approach*. Englewood Cliffs: Prentice-Hall, 1969.
- [69] W. W. Hogan, "Point-to-set maps in mathematical programming," *SIAM Rev.*, vol. 15, no. 3, pp. 591–603, July 1973.
- [70] P. Huard, Ed., *Point-to-Set Maps and Mathematical Programming*, ser. Mathematical Programming Study. Amsterdam: North-Holland Publishing Co., 1979, no. 10.
- [71] G. B. Dantzig, J. Folkman, and N. Shapiro, "On the continuity of the minimum set of a continuous function," *J. Math. Anal. Appl.*, vol. 17, pp. 519–548, 1967.
- [72] Z. Opial, "Weak convergence of the sequence of successive approximations for nonexpansive mappings," *Bull. Amer. Math. Soc.*, vol. 73, pp. 591–597, 1967.
- [73] I. Daubechies, M. Defrise, and C. D. Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Comm. Pure Appl. Math.*, 2004, to appear.
- [74] J. C. Fiorot and P. Huard, "Composition and union of general algorithms of optimization," in *Point-to-Set Maps and Mathematical Programming*, ser. Mathematical Programming Studies. North Holland, 1979, vol. 10, pp. 69–85.
- [75] F. Deutsch, *Multivariate Approximation Theory*. Basel: Birkhäuser Verlag, 1979, ch. The alternating method of von Neumann, pp. 83–96.
- [76] —, *Approximation Theory, Spline Functions and Applications*. The Netherlands: Kluwer Academic Publishers, 1992, ch. The method of alternating orthogonal projections, pp. 105–121.

- [77] N. Wiener, "On the factorization of matrices," *Comment. Math. Helv.*, vol. 29, pp. 97–111, 1955.
- [78] I. Halperin, "The product of projection operators," *Acta Sci. Math.*, vol. 23, pp. 96–99, 1962.
- [79] K. T. Smith, D. C. Solmon, and S. L. Wagner, "Practical and mathematical aspects of the problem of reconstructing objects from radiographs," *Bull. Amer. Math. Soc.*, vol. 83, pp. 1227–1270, 1977.
- [80] J. Xu and L. Zikatonov, "The method of alternating projections and the method of subspace corrections in Hilbert space," *J. Amer. Math. Soc.*, vol. 15, no. 3, pp. 573–597, Apr. 2002.
- [81] H. H. Bauschke, F. Deutsch, H. Hundal, and S.-H. Park, "Accelerating the convergence of the method of alternating projections," *Trans. Amer. Math. Soc.*, vol. 355, no. 9, pp. 3433–3461, May 2003.
- [82] S. P. Diliberto and E. G. Straus, "On the approximation of a function of several variables by the sum of functions of fewer variables," *Pacific J. Math.*, vol. 1, pp. 195–210, 1951.
- [83] G. Aumann, *Über approximative Nomographie I, II and III*. Bayer. Akad. Wiss. Math.-Nat. Kl. S.B., 1958–1960.
- [84] M. Golomb, *On Numerical Approximation*. Madison: Univ. of Wisconsin Press, 1959, ch. Approximation by functions of fewer variables, pp. 275–327.
- [85] E. W. Cheney and W. A. Light, *Approximation Theory in Tensor Product Spaces*, ser. Lecture Notes in Mathematics. Berlin: Springer-Verlag, 1985, no. 1169.
- [86] R. L. Dykstra, "An algorithm for restricted least squares regression," *J. Amer. Stat. Assoc.*, vol. 78, pp. 837–842, 1983.
- [87] J. P. Boyle and R. L. Dykstra, *Advances in Order Restricted Inference 37*, ser. Lecture notes in statistics. Springer, 1985, ch. A method for finding projections onto the intersection of convex sets in Hilbert spaces, pp. 28–47.
- [88] S.-P. Han, "A successive projection method," *Math. Prog.*, vol. 40, pp. 1–14, 1988.
- [89] C. Hildreth, "A quadratic programming procedure," *Naval Res. Logist. Quart.*, vol. 4, 1957.
- [90] Y. Censor and S. A. Zenios, *Parallel Optimization: Theory, Algorithms, and Applications*, ser. Numerical Mathematics and Scientific Computation. Oxford University Press, 1997.
- [91] L. M. Bregman, Y. Censor, and S. Reich, "Dykstra's Algorithm as the Nonlinear Extension of Bregman's Optimization Method," *Journal of Convex Analysis*, vol. 6, no. 2, pp. 319–333, 1999.
- [92] P. Tseng, "Dual coordinate ascent methods for non-strictly convex minimization," *Mathematical Programming*, vol. 59, pp. 231–247, 1993.
- [93] H. H. Bauschke and J. Borwein, "On projection algorithms for solving convex feasibility problems," *SIAM Rev.*, vol. 38, no. 3, pp. 367–426, 1996.
- [94] P. L. Combettes and H. J. Trussell, "Method of successive projections for finding a common point of sets in metric spaces," *J. Optim. Theory and Appl.*, vol. 67, no. 3, pp. 487–507, 1990.
- [95] J. A. Cadzow, "Signal enhancement—a composite property mapping algorithm," *IEEE Trans. Acoustics, Speech and Signal Processing*, vol. 36, no. 1, pp. 49–62, Jan. 1988.
- [96] R. Blahut, "Computation of channel capacity and rate distortion functions," *IEEE Trans. Inform. Theory*, vol. 18, pp. 460–473, 1972.
- [97] S. Arimoto, "An algorithm for calculating the capacity of an arbitrary discrete memoryless channel," *IEEE Trans. Inform. Theory*, vol. 18, pp. 14–20, 1972.
- [98] H. J. Landau and W. L. Miranker, "The recovery of distorted band-limited signals," *J. Math. Anal. Appl.*, vol. 2, pp. 97–104, 1961.
- [99] R. W. Gerchberg, "Super-resolution through error energy reduction," *Optica Acta*, vol. 21, pp. 709–720, 1974.
- [100] D. C. Youla and H. Webb, "Image restoration by the method of convex projections. I—theory," *IEEE Trans. Medical Imaging*, vol. 1, pp. 81–94, 1982.
- [101] D. L. Donoho and P. B. Stark, "Uncertainty principles and signal recovery," *SIAM J. Appl. Math.*, vol. 49, no. 3, pp. 906–931, Jun. 1989.
- [102] A. E. Çetin, Ö. N. Gerek, and Y. Yardimci, "Equiripple FIR filter design by the FFT algorithm," *IEEE Signal Process. Mag.*, Mar. 1997.
- [103] W. Glunt and T. L. Hayden, "The nearest "doubly stochastic" matrix to a real matrix with the same first moment," *Numer. Linear Algebra Appl.*, vol. 5, no. 6, Dec. 1998.
- [104] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Royal Stat. Soc.*, vol. 39, pp. 1–38, 1977.

Joel A. Tropp (S'03) earned the B.A. degree in Plan II Liberal Arts Honors and the B.S. degree in Mathematics from the University of Texas at Austin (UT-Austin) in May 1999, graduating *magna cum laude*. He continued his graduate studies in Computational and Applied Mathematics (CAM) at UT-Austin where he completed the M.S. degree in May 2001 and the Ph.D. degree in August 2004.

Dr. Tropp is presently an Assistant Professor in the Mathematics Department of the University of Michigan at Ann Arbor.

As an undergraduate, Dr. Tropp was a recipient of the Barry M. Goldwater National Science Scholarship and a semi-finalist for the British Marshall. He won a CAM Graduate Fellowship in 1999 and a National Science Foundation Graduate Fellowship in 2000.

Inderjit S. Dhillon (M'01) is an assistant professor in the Department of Computer Sciences at UT-Austin. He received his Ph.D. from the EECS department at UC Berkeley in 1997, where his dissertation work led to the first stable $O(n^2)$ algorithm for the symmetric tridiagonal eigenproblem. Inderjit's main research interests are in numerical linear algebra, scientific computing, data mining and machine learning. Prior to joining UT-Austin, Inderjit has held research positions at AT&T Bell Laboratories, Murray Hill and IBM Almaden. He holds a B.Tech. degree in Computer Science from the Indian Institute of Technology, Bombay. Inderjit is a member of ACM, IEEE and SIAM, and is on the Editorial Board for the SIAM Journal on Matrix Analysis and Applications. His work is funded by an NSF CAREER award, an NSF ITR award, and a Texas Advanced Research Program award.

Robert W. Heath Jr. (S'96, M'01) received the B.S. and M.S. degrees from the University of Virginia, Charlottesville, VA, in 1996 and 1997 respectively, and the Ph.D. from Stanford University, Stanford, CA, in 2002, all in electrical engineering.

From 1998 to 1999, he was a Senior Member of the Technical Staff at Iospan Wireless Inc, San Jose, CA where he played a key role in the design and implementation of the physical and link layers of the first commercial MIMO-OFDM communication system. From 1999 to 2001 he served as a Senior Consultant for Iospan Wireless Inc. In 2003 he founded MIMO Wireless Inc, a consulting company dedicated to the advancement of MIMO technology. Since January 2002, he has been with the Department of Electrical and Computer Engineering at The University of Texas at Austin where he serves as an Assistant Professor as part of the Wireless Networking and Communications Group. His research interests include interference management in wireless networks, sequence design, and all aspects of MIMO communication including antenna design, practical receiver architectures, limited feedback techniques, and scheduling algorithms.

Dr. Heath serves as an Associate Editor for the IEEE Transactions on Vehicular Technology.

Thomas Strohmer got his M.S. and Ph.D. in Mathematics in 1991 and 1994 respectively from the University of Vienna, Austria. He was a Research Assistant at the Department of Mathematics, University of Vienna from 1991 to 1997. He spent one year as Erwin Schrodinger Fellow at the Department of Statistics at the Stanford University and then joined the Department of Mathematics at the University of California in Davis in 1998, where he is now Full Professor.

His general research interests are in harmonic analysis, numerical analysis, digital signal processing, and wireless communications. He is co-editor of two books and on the board of several journals. He also serves as consultant for the telecommunications industry.