

Unit norm tight frames in finite-dimensional spaces

Dustin G. Mixon*

November 3, 2014

Abstract

There are several problems in signal processing which motivate the study of unit norm tight frames in finite-dimensional spaces. We discuss some of these problems and then observe various constructions and properties that have been identified in the last decade.

1 Introduction

A **frame** is a countable collection of vectors $\Phi = \{\varphi_i\}_{i \in I}$ in a separable Hilbert space \mathcal{H} with the property that there exist **frame bounds** $0 < A \leq B < \infty$ such that

$$A\|x\|^2 \leq \sum_{i \in I} |\langle x, \varphi_i \rangle|^2 \leq B\|x\|^2 \quad \forall x \in \mathcal{H}.$$

Here, $\|\cdot\|$ denotes the norm induced by the inner product $\langle \cdot, \cdot \rangle$ corresponding to \mathcal{H} . Intuitively, a frame allows the mapping $x \mapsto \{\langle x, \varphi_i \rangle\}_{i \in I}$ to capture the energy of any $x \in \mathcal{H}$, and this property enables any such x to be reconstructed with the help of some **dual frame**. In particular, for every frame $\Phi = \{\varphi_i\}_{i \in I} \subseteq \mathcal{H}$, there exists a dual frame $\Psi = \{\psi_i\}_{i \in I} \subseteq \mathcal{H}$ such that

$$x = \sum_{i \in I} \langle x, \varphi_i \rangle \psi_i \quad \forall x \in \mathcal{H}.$$

For example, any orthonormal basis Φ of \mathcal{H} is a frame with frame bounds $A = B = 1$, and a corresponding dual frame is simply $\Psi = \Phi$. In general, we say a frame is **tight** if $A = B$. In this case, choosing the dual frame $\Psi = (1/A)\Phi$ allows for a particularly painless reconstruction formula:

$$x = \frac{1}{A} \sum_{i \in I} \langle x, \varphi_i \rangle \varphi_i \quad \forall x \in \mathcal{H}.$$

Finally, we say a frame Φ is **unit norm** if each of the frame elements satisfies $\|\varphi_i\| = 1$. In this way, a unit norm tight frame (UNTF) is a natural generalization of orthonormal basis, and as we will see, there is an assortment of UNTFs which are not orthonormal bases.

These notes consider the case where \mathcal{H} has finite dimension, meaning \mathcal{H} is isomorphic to either \mathbb{R}^M or \mathbb{C}^M for some positive integer M . In this setting, the finitude of the upper frame bound forces the number of frame elements to be finite, say N , and it is convenient to express the frame as an $M \times N$ matrix whose columns are the frame elements. By abuse of notation, we will denote this matrix by $\Phi = [\varphi_1 \cdots \varphi_N]$. The frame bounds are then lower and upper bounds on the quantity

$$\frac{\|\Phi^*x\|^2}{\|x\|^2} = \frac{\langle \Phi^*x, \Phi^*x \rangle}{\|x\|^2} = \frac{\langle x, \Phi\Phi^*x \rangle}{\|x\|^2}, \quad x \neq 0,$$

*Department of Mathematics and Statistics, Air Force Institute of Technology, Wright-Patterson AFB, OH

which is minimized and maximized at the smallest and largest eigenvalues of $\Phi\Phi^*$, respectively. As such, the optimal frame bounds are precisely these extreme eigenvalues. At this point, we note that $A > 0$ precisely when the frame elements span, and so any frame of N elements in M dimensions must satisfy $N \geq M$. Also, it is clear that Φ is tight precisely when $\Phi\Phi^*$ is a multiple A of the $M \times M$ identity matrix I . Equivalently, the rows of Φ are orthogonal, each with norm \sqrt{A} . In addition, we may easily determine the necessary value of A by appealing to the cyclic property of the trace:

$$MA = \text{Tr}[AI] = \text{Tr}[\Phi\Phi^*] = \text{Tr}[\Phi^*\Phi] = N,$$

where the last step follows from the fact that each frame element has norm 1. Overall, a UNTF in finite-dimensional space is an $M \times N$ matrix such that

- the rows are orthogonal,
- each row has norm $\sqrt{N/M}$, and
- each column has norm 1.

As one might expect, the simultaneous conditions on the rows and columns make the construction of UNTFs somewhat challenging.

In the next section, we will provide a series of applications that motivate the study of UNTFs in finite-dimensional spaces. In Section 3, we will then provide a list of examples to build up an intuition for what UNTFs look like. This will prompt two main questions. First, how do UNTFs correspond to the notion of equidistribution on the unit sphere? This is answered in Section 4, in which we discuss the frame potential. Second, how might one construct all possible UNTFs? This is answered in Section 5, where we discuss the theory of eigensteps.

2 Motivating applications

In this section, we discuss several applications of UNTFs to problems in signal processing. This will motivate the rest of these notes, which consider various constructions and properties of UNTFs.

2.1 Robustness to noise or erasures

Suppose Alice has a message that she wishes to transmit to Bob, but the channel through which she can send her message will invariably corrupt the transmission. The goal is to develop an encoding–decoding scheme that will ensure Bob will receive the message.

To model this scenario, suppose the channel acts by adding random noise to each entry of the transmitted signal. For example, perhaps a sequence of i.i.d. $N(0, \sigma^2)$ random variables is added. Then Alice might redundantly encode her signal $x \in \mathbb{R}^M$ using a frame, meaning $\Phi^*x \in \mathbb{R}^N$ is sent over the channel. What Bob then receives has the form $y = \Phi^*x + e$, where e denotes the channel’s random noise vector. Considering the isotropy of the Gaussian distribution, Bob can easily determine the maximum likelihood estimate \hat{x} of x by solving a least squares problem:

$$\hat{x} := \arg \min_z \|\Phi^*z - y\|.$$


If we take Ψ to be the pseudoinverse of Φ , i.e., $\Psi = (\Phi\Phi^*)^{-1}\Phi$, then $\hat{x} = \Psi y$. Notice that Ψ is a dual frame of Φ since $\Psi\Phi^* = (\Phi\Phi^*)^{-1}\Phi\Phi^* = I$; in fact, this Ψ is the **canonical dual frame** of Φ .

Now that an optimal dual frame (decoder) has been identified for each frame (encoder), we will optimize the encoder frame. For this, we measure performance by mean-squared error:

$$\text{MSE} := \mathbb{E}\|\hat{x} - x\|^2 = \mathbb{E}\|\Psi(\Phi^*x + e) - x\|^2 = \mathbb{E}\|\Psi e\|^2.$$

Intuitively, this is a measure of average distortion, and we wish to minimize this quantity. To this end, it has been shown that of all unit norm frames Φ , the minimizers of MSE are UNTFs; see Theorem 3.1 in [14].

For another model of the channel between Alice and Bob, suppose that after noise is added to the signal, some arbitrary entry of the signal is set to zero. This is called an erasure channel, as depicted in the following:



Notice that in this setup, Bob intends to blindly decode by applying the canonical dual frame Ψ of Φ regardless of what erasure he may perceive. Here, if we let D_n denote the diagonal matrix of all 1s, but with a 0 at the n th diagonal entry, then the following quantity measures performance in the event that the n th entry is erased:

$$\text{MSE}_n := \mathbb{E}\|\hat{x} - x\|^2 = \mathbb{E}\|\Psi D_n(\Phi^*x + e) - x\|^2.$$

At this point, we can consider two extreme models for how an entry is selected for erasing. If an entry is selected at random for erasure, then we would consider the average MSE over all n . On the other hand, if an adversary decides which entry to erase after observing x , then we would maximize MSE over all n . Interestingly, both average- and worst-case MSE are minimized when Φ is a UNTF; see Theorem 4.4 in [14].

2.2 Fingerprinting to defeat piracy

Suppose you are a famous recording artist, and you are about to release a new single. Since you are aware that media piracy is a cause for concern, you decide to send a slightly different copy of your song file to each recipient—that way, if a particular version of your song becomes popular on the internet, you will know which recipient of your song was the original culprit. If $x \in \mathbb{R}^M$ denotes your song file, then you make different copies of the file by adding a personalized fingerprint to the file, namely $x + \varphi_n$, where $\{\varphi_n\}_{n=1}^N$ is the total collection of fingerprints you will use.

Unfortunately, some of the recipients might conspire to produce a forgery of your song file. Suppose $\mathcal{K} \subseteq \{1, \dots, N\}$ denotes the indices of these conspiring culprits. Then they might attempt to forge your file by making a noisy linear combination of their personalized fingerprinted versions:

$$\hat{x} = \sum_{n \in \mathcal{K}} c_n(x + \varphi_n) + e, \quad \sum_{n \in \mathcal{K}} c_n = 1, \quad e \sim N(0, I).$$

Here, the purpose of the noise is to “cover their tracks” so as to keep you from determining \mathcal{K} .

Suppose a forgery \hat{x} surfaces on the internet. Then its difference from the true file has the form

$$\hat{x} - x = \sum_{n \in \mathcal{K}} c_n \varphi_n + e.$$

To test whether the n th recipient is a possible culprit, you can take the inner product $\langle \hat{x} - x, \varphi_n \rangle$. Then the largest such inner product determines the most guilty-looking recipient. If this recipient is a member of \mathcal{K} , then you can interrogate him to determine the remainder of the culprits. Knowing this, the culprits might select the scalars $\{c_n\}_{n \in \mathcal{K}}$ in their linear combination so as to frame an innocent recipient. Theorem 7 in [18] shows that this is actually impossible provided the fingerprints form a certain type of UNTF, called an **equiangular** UNTF; this is a UNTF with the additional property that there is some $c \geq 0$ such that $|\langle \varphi_i, \varphi_j \rangle| = c$ for every $i, j \in \{1, \dots, N\}$ with $i \neq j$. In particular, the probability of accidentally accusing the most guilty-looking innocent recipient is always small, regardless of how the culprits choose $\{c_n\}_{n \in \mathcal{K}}$.

2.3 Sparse decomposition

Consider a radar antenna that transmits a ping $\varphi \in \mathbb{C}^M$, which is then re-radiated from various aircraft in a region of interest. Then the radar antenna can listen for the superposition of these different echoes, each being a different translation and modulation of the original ping (corresponding to a time delay due to distance and a Doppler shift due to velocity). As such, one can consider a frame whose frame elements are the original ping φ along with all translates and modulates of the ping: $\Phi = \{T^a E^b \varphi\}_{a,b \in \mathbb{Z}/M\mathbb{Z}}$; here, we use cyclic translates and modulates for simplicity:

$$(Tx)(m) := x(m-1), \quad (Ex)(m) := e^{2\pi im/M} x(m),$$

which is reasonable if we zero-pad the original ping appropriately. Such a frame Φ is called a **Gabor frame**, and provided the ping φ has unit norm, this frame will necessarily be a UNTF. With this notation, each target corresponds to a different pair (a, b) according to its distance and velocity, and we receive the superposition

$$y = \sum_{(a,b) \in \mathcal{K}} x_{(a,b)} T^a E^b \varphi = \Phi x,$$

where $x \in \mathbb{C}^{M^2}$ has $|\mathcal{K}|$ nonzero entries. The purpose of radar is to determine information about the various objects in the region of interest, so we would like to determine x from y , since x encodes the distances and velocities of the $|\mathcal{K}|$ objects that echoed the ping. However, $y = \Phi x$ is a severely underdetermined linear system, and so to do this, we would need to exploit the fact that x has few nonzero entries.

This is an instance of a more general problem called **sparse decomposition**, where we seek to express y as a sparse combination of columns of Φ . This problem of solving a linear system with a sparsity prior has received a lot of attention recently under the name **compressed sensing**, and today, we have several algorithms of reconstructing x from y provided x is sufficiently sparse and the columns of Φ are sufficiently incoherent, meaning no two columns of Φ look alike. More explicitly, the various algorithms tend to perform better when the **worst-case coherence**

$$\mu := \max_{\substack{i,j \in \{1, \dots, N\} \\ i \neq j}} |\langle \varphi_i, \varphi_j \rangle|$$

is small. For the application of radar, this suggests the use of a ping that looks very little like its various translates and modulates, as in [15]. As established in [23], the general problem of sparse decomposition is particularly solvable when Φ is a UNTF with small worst-case coherence.

3 Examples

The previous section illustrated how UNTFs can be used in a variety of signal processing applications. Recognizing the importance of UNTFs, this leaves one longing for a nontrivial example. From the introduction, the reader is already aware that every orthonormal basis is an example of a UNTF. This section offers an extensive list of additional examples for the reader's satisfaction. We start with UNTFs in two-dimensional real space:

Example 1. Pick $N \geq 3$ and consider the N th roots of unity. These are N complex numbers, each of unit modulus, and we can view each as a vector in \mathbb{R}^2 whose coordinates are the real and imaginary parts. Then the resulting N vectors form a unit norm tight frame in \mathbb{R}^2 .

This example is intuitively pleasing because it suggests that UNTFs satisfy some notion of equidistribution on the unit sphere in \mathbb{R}^M , which is discussed at length in the next section. It also leads one to wonder whether there are any other UNTFs in \mathbb{R}^2 . This is answered in the following theorem:

Theorem 2 (Theorem 2.7 in [14]). *The vectors $\{\varphi_n\}_{n=1}^N \subseteq \mathbb{R}^2$ form a unit norm tight frame if and only if the corresponding complex numbers $\{z_n\}_{n=1}^N \subseteq \mathbb{C}$ satisfy*

$$\sum_{n=1}^N z_n^2 = 0.$$

The proof of this theorem amounts to a straightforward exercise in trigonometry. Not only can we use this result to verify the previous example, it also gives us an intuition for the entire set of $2 \times N$ UNTFs for any fixed $N \geq 2$. Indeed, we can add $\{z_n^2\}_{n=1}^N$ head-to-tail as if they are unit vectors, and then the above characterization gives that the sum goes back to the origin, meaning the unit vectors form a closed chain. One can imagine deforming this chain as if it were a necklace made of sticks, and by the above theorem, the set of all possible configurations corresponds to the set of all possible UNTFs in \mathbb{R}^2 . The fact that two unit vectors sum to zero precisely when they are negatives of each other corresponds to the fact that orthonormal bases are the only UNTFs in \mathbb{R}^2 with $N = 2$. Also, three unit vectors sum to zero precisely when they form the legs of an equilateral triangle, indicating that the cube roots of unity form the only UNTF of 3 vectors in \mathbb{R}^2 up to rotation and negation; for obvious reasons, this is called the **Mercedes-Benz frame**. When $N \geq 4$, the necklace of sticks has a lot more freedom. For example, when $N = 4$, unit vectors sum to zero precisely when they form the legs of a rhombus, and intuitively, we can push or pull at opposite corners to deform the rhombus into a continuous ensemble of rhombi, which correspond to a continuous ensemble of distinct UNTFs. Intuitively, increasing N produces additional degrees of freedom, leading to even more UNTFs.

Now that we have a complete characterization of UNTFs in \mathbb{R}^2 , we continue by considering examples in \mathbb{R}^3 :

Example 3. Consider a Platonic solid which is circumscribed by the unit sphere in \mathbb{R}^3 (for examples, see Figure 1). Then the vertices of this solid form a unit norm tight frame.

This example was verified in [10], and it can be viewed as a 3-dimensional analog of the roots of unity, again exhibiting an intuitive notion of equidistribution on the sphere. In general, if vectors exhibit enough symmetry, you can expect them to form a UNTF. The following result makes this more precise:




Figure 1: Platonic solids (from [19]). When circumscribed by the unit sphere, each vertex is a vector of unit norm, and the collection of vertices forms a unit norm tight frame in \mathbb{R}^3 . The fact that these vertices are well-distributed on the sphere suggests that UNTFs might satisfy some notion of equidistribution, which is investigated further in Section 4.

Theorem 4 (Theorem 6.3 in [24]). *Let \mathcal{H} denote either \mathbb{R}^M or \mathbb{C}^M , and let G be a finite group of norm-preserving linear operators over \mathcal{H} . Suppose G is irreducible, that is, $\{U\varphi\}_{U \in G}$ spans \mathcal{H} for every $\varphi \in \mathcal{H}$ of unit norm. Then every such spanning set is a unit norm tight frame.*

The symmetry group of each platonic solid is irreducible, as is the (dihedral) symmetry group of the N th roots of unity, and so the fact that these correspond to UNTFs is a consequence of this result. In general, a finite group need not be irreducible for there to exist a vector whose orbit under the group forms a UNTF:

Example 5. Consider the group generated by a diagonal matrix whose diagonal entries are distinct N th roots of unity. Then the orbit of the first identity basis fails to span, whereas the orbit of the normalized all-ones vector forms a unit norm tight frame.

For instance, denoting $\omega = e^{2\pi i/7}$, we might have

$$D = \begin{bmatrix} \omega^3 & 0 & 0 \\ 0 & \omega & 0 \\ 0 & 0 & \omega^6 \end{bmatrix}, \quad \Phi = \frac{1}{\sqrt{3}} \begin{bmatrix} 1 & \omega^3 & \omega^6 & \omega^2 & \omega^5 & \omega & \omega^4 \\ 1 & \omega & \omega^2 & \omega^2 & \omega^4 & \omega^5 & \omega^6 \\ 1 & \omega^6 & \omega^5 & \omega^4 & \omega^3 & \omega^2 & \omega \end{bmatrix}.$$

Notice that the orbit of the normalized all-ones vector under such a group will always produce a matrix whose rows come from the discrete Fourier transform matrix (suitably scaled). With this identification, it is easy to verify that such matrices, called **harmonic frames**, are necessarily UNTFs. The construction of harmonic frames establishes that a complex UNTF exists for every pair (M, N) such that $M \leq N$. For the real case, one can leverage the real and imaginary parts of a harmonic frame [26]. The following example gives another general construction of real UNTFs:

Example 6. Given (M, N) such that $N \geq 2M$, construct each row of an $M \times N$ UNTF using a greedy process known as **Spectral Tetris**¹. In particular, design the matrix one row at a time, each time putting as much energy as possible on the left. For example, $(M, N) = (3, 7)$ corresponds to the following UNTF:

$$\begin{bmatrix} 1 & 1 & \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{6}} & 0 & 0 & 0 \\ 0 & 0 & \frac{\sqrt{5}}{\sqrt{6}} & -\frac{\sqrt{5}}{\sqrt{6}} & \frac{\sqrt{2}}{\sqrt{6}} & \frac{\sqrt{2}}{\sqrt{6}} & 0 \\ 0 & 0 & 0 & 0 & \frac{\sqrt{2}}{\sqrt{3}} & -\frac{\sqrt{2}}{\sqrt{3}} & 1 \end{bmatrix}.$$

¹Spectral Tetris was named after the tile-matching puzzle video game *Tetris* because each frame element can be viewed as a localized contribution to the spectrum of $\Phi\Phi^*$, which we want to be flat so that Φ is tight.

To get this, we first note that each entry must have size at most 1 (by the unit-norm constraint), and so we populate the first two entries of the first row with 1s accordingly. Since each row must have norm $\sqrt{7/3}$, we may not put a third 1 in the first row. We also need to allow for the second row to be orthogonal to the first, and so we evenly distribute the remaining energy across the next two entries. For the second row, the unit-norm constraint forces the first two entries to be 0, and then the orthogonality constraint (along with the greedy objective of putting as much energy on the left as possible) forces the next two entries to be $\pm\sqrt{5/6}$. The energy remaining for the second row is less than 1, and so we evenly distribute it over the next two entries. Finally, the last row is forced to have 0s in the first four entries by the unit-norm constraint, and then the next two entries become $\pm\sqrt{2/3}$, leaving a unit of energy that we put in the last entry.

Spectral Tetris is guaranteed to produce a UNTF provided $N \geq 2M$ [6], and the generated UNTFs necessarily enjoy two important properties. First, the greedy method ensures that each row has at most $N/M + 2$ nonzero entries, meaning the entire matrix is particularly sparse, having $O(N)$ nonzero entries, thereby allowing for efficient matrix-vector multiplication. In fact, the UNTFs which come from Spectral Tetris are known to be the sparsest possible [7]. The second important property to note is that columns which are far apart tend to have disjoint support, and so they are orthogonal to each other; this is useful to the construction of a generalization of UNTFs called **tight fusion frames** or tight “frames of subspaces” [6].

Having observed several examples of UNTFs, we are led to two natural questions. First, is there a notion of equidistribution on the unit sphere that UNTFs necessarily satisfy? We have seen several examples of UNTFs for which this seems plausible, but then the fact that orthonormal bases are also examples of UNTFs seems to suggest otherwise. The next section will illustrate that, indeed, there is a natural notion of equidistribution which completely characterizes UNTFs, even despite orthonormal bases. Second, Theorem 2 corresponds to a method of constructing all possible UNTFs in \mathbb{R}^2 , which leads one to wonder: Is there a more general construction available for \mathbb{R}^M and \mathbb{C}^M ? Indeed there is, as we detail in Section 5.

4 The frame potential

Imagine a collection of N electrons which are free to move within a spherical shell. By Coulomb’s law, each electron imposes a force on each of the other $N - 1$ electrons, and the magnitude of this force is inversely proportional to the square of the distance between the two electrons. As such, if all of the electrons are initialized in a small region of the sphere, they will push each other away and eventually settle into some arrangement of positions that attempts to minimize potential energy. Of course, the electrons may fail to converge to an arrangement of minimal potential because their pursuit is greedy, thereby leading to a local minimizer, which we will say exhibits **equilibrium** under Coulomb’s force. Still, there exists a global minimizer by the extreme value theorem, which we will say exhibits **equidistribution** under Coulomb’s force. Finding equidistributed arrangements under Coulomb’s force is known as the Thomson problem [22], and is a special case of Smale’s 7th problem [21].

With sufficient imagination, one may conjure any number of force laws and observe which arrangements on the sphere minimize potential energy. For our purposes, it is convenient to restrict one’s imagination so that the force that $y \in \mathbb{R}^M$ applies to $x \in \mathbb{R}^M$ has direction $(x - y)/\|x - y\|$ and magnitude determined by $\|x - y\|$:

$$F(x, y) := f(\|x - y\|) \frac{x - y}{\|x - y\|}.$$

(We allow $f(\cdot)$ to take negative outputs, as is the case for gravitation.) The potential energy corresponding to points at x and y is then given by

$$P(x, y) = - \int_0^1 F(tx + (1-t)y, y) \cdot (x - y) dt = - \int_0^{\|x-y\|} f(u) du,$$

meaning $P(x, y) = p(\|x - y\|)$, where $p(\cdot)$ satisfies $p(0) = 0$ and $p'(r) = -f(r)$. Given a force, which in turn determines a potential, we seek arrangements of points $\{\varphi_n\}_{n=1}^N$ on the sphere which minimize the total potential energy:

$$\sum_{n=1}^N \sum_{\substack{n'=1 \\ n' \neq n}}^N P(\varphi_n, \varphi_{n'}). \quad (1)$$

In pursuit of global minimizers, we first consider the arrangements which satisfy the Lagrange equations (the ‘‘critical points’’). For the class of forces we are considering, Theorem 4.6 in [1] gives that these arrangements necessarily have the property that each φ_n is a scalar multiple of

$$\sum_{\substack{n'=1 \\ n' \neq n}}^N F(\varphi_n, \varphi_{n'}). \quad (2)$$

We note that due to their symmetries, the vertices of each Platonic solid satisfies this condition, regardless of f , and so one is compelled to find a particular f for which every UNTF is equidistributed. If there were such an f , it must also allow orthonormal bases to be equidistributed, which is certainly not the case for Coulomb’s force. However, we do have $\Phi\Phi^* = (N/M)I$ for any UNTF Φ , which implies that

$$\left(\frac{N}{M} - 1\right)\varphi_n = \sum_{\substack{n'=1 \\ n' \neq n}}^N \langle \varphi_n, \varphi_{n'} \rangle \varphi_{n'}. \quad (3)$$

As such, we seek a force F that allows us to use (3) to show that (2) is a constant multiple of φ_n . To this end, we define the **frame force** to be

$$\text{FF}(x, y) := \langle x, y \rangle (x - y).$$

Notice that this choice of force takes $f(r) = r - r^3/2$. Then we can use (3) to get

$$\begin{aligned} \sum_{\substack{n'=1 \\ n' \neq n}}^N \text{FF}(\varphi_n, \varphi_{n'}) &= \sum_{\substack{n'=1 \\ n' \neq n}}^N \langle \varphi_n, \varphi_{n'} \rangle (\varphi_n - \varphi_{n'}) \\ &= \sum_{\substack{n'=1 \\ n' \neq n}}^N \langle \varphi_n, \varphi_{n'} \rangle \varphi_n - \sum_{\substack{n'=1 \\ n' \neq n}}^N \langle \varphi_n, \varphi_{n'} \rangle \varphi_{n'} = \left(\sum_{\substack{n'=1 \\ n' \neq n}}^N \langle \varphi_n, \varphi_{n'} \rangle - \frac{N}{M} + 1 \right) \varphi_n, \end{aligned}$$

which is a scalar multiple of φ_n . As such, UNTFs satisfy the Lagrange equations for the frame force, as is necessary for equidistribution. To test equidistribution, we consider the corresponding potential: Integrating and negating $f(r) = r - r^3/2$ gives $p(r) = -r^2/2 + r^4/8$, and so

$$P(x, y) = p(\|x - y\|) = -\frac{\|x - y\|^2}{2} \left(1 - \frac{\|x - y\|^2}{4} \right) = \frac{1}{2} (|\langle x, y \rangle|^2 - 1).$$

Ignoring constant terms, this means that minimizing total potential energy (1) is equivalent to minimizing the **frame potential**:

$$\text{FP}(\Phi) := \sum_{n=1}^N \sum_{n'=1}^N |\langle \varphi_n, \varphi_{n'} \rangle|^2 = \|\Phi^* \Phi\|_{\text{HS}}^2.$$

With this, the following establishes the requirements for equidistribution:

$$0 \leq \left\| \Phi \Phi^* - \frac{N}{M} I \right\|_{\text{HS}}^2 = \text{Tr}[(\Phi \Phi^*)^2] - 2 \frac{N}{M} \text{Tr}[\Phi \Phi^*] + \frac{N^2}{M^2} \text{Tr}[I] = \text{FP}(\Phi) - \frac{N^2}{M}.$$

Indeed, rearranging gives that the frame potential satisfies $\text{FP}(\Phi) \geq N^2/M$ with equality precisely when $\Phi \Phi^* = (N/M)I$, i.e., Φ is a UNTF. This inequality is sometimes called the zeroth-order Welch bound [25]. By construction, we know UNTFs exist whenever (M, N) satisfies $M \leq N$, and so UNTFs are precisely the global minimizers of the frame potential, i.e., UNTFs are the arrangements of points on the sphere which exhibit equidistribution under the frame force. Next, we turn to the more difficult problem of equilibrium. Indeed, for general forces (such as the Coulomb force), points often settle for a local minimizer of total potential which is not global. Perhaps surprisingly, the following result establishes that this is not an issue when applying the frame force:

Theorem 7 (Theorem 7.1 in [1]). *Consider the frame potential of all arrangements of N points in the unit sphere in \mathbb{R}^M or \mathbb{C}^M . Then every local minimizer is also a global minimizer.*

Note that while the motivation for the frame potential came from a frame force defined between points in real space, the above theorem is valid for both real and complex spaces. The proof of this theorem uses the contrapositive: Given any arrangement which is not tight, find a family of arbitrarily close arrangements, each having strictly smaller frame potential; then the original arrangement is not a local minimizer. The utility of this result is demonstrated in part by its application to the so-called Paulsen problem.

4.1 Application: The Paulsen problem

Given a frame which is nearly unit norm and nearly tight, how far is the closest UNTF? This question was first posed by Vern Paulsen [2], and has since been dubbed **the Paulsen problem**. To evaluate what it means to be nearly unit norm and nearly tight, one must first choose a notion of “distance” from unit-norm tightness. Consider any continuous gauge function $\rho: \mathbb{C}^{M \times N} \rightarrow [0, \infty)$ such that $\rho(\Phi) = 0$ precisely when Φ is a UNTF and furthermore $\rho(\Phi)$ is bounded away from zero whenever $\|\Phi\|_{\text{HS}}$ is sufficiently large. For example, one may take

$$\rho(\Phi) = \|\text{diag}(\Phi^* \Phi) - \mathbf{1}\|_2 + \left\| \Phi \Phi^* - \frac{N}{M} I \right\|_{\text{HS}}.$$

Then the Paulsen problem asks for a function $\delta = \delta(\epsilon)$ such that $\rho(\Phi) < \delta$ implies the existence of a UNTF that is within ϵ of Φ .

There is a nice argument due to Don Hadwin [2] which guarantees the existence of such a function: Suppose otherwise—that there is a sequence $\{\Phi_i\}_{i=1}^{\infty}$ for which $\rho(\Phi_i) \rightarrow 0$ but the distance between each Φ_i and each UNTF is bounded away from zero. Then a tail of this sequence lies in a compact set of the form $\{\Phi : \rho(\Phi) \leq \eta\}$, and so Bolzano–Weierstrass guarantees a convergent subsequence. Let Φ_{∞} denote the limit point of this subsequence. By the continuity of ρ , we then

have $\rho(\Phi_\infty) = 0$, meaning Φ_∞ is a UNTF, thereby contradicting the assumption that the distance between each Φ_i and each UNTF is bounded away from zero.

(We can establish more about the function $\delta(\epsilon)$ by observing that set of UNTFs forms an algebraic variety. Indeed, one may leverage a powerful result from algebraic geometry called Lojewicz’s inequality [9] to conclude that for certain choices of ρ , the function $\delta(\epsilon)$ should have the form $\delta(\epsilon) = C\epsilon^\alpha$, where C and α possibly depend on M and N .)

Note that it is easy to find a nearby unit norm frame by normalizing the frame elements, and in fact, one can control how much tightness is lost in this process [2]. As such, one may seek a solution to the Paulsen problem by assuming without loss of generality that the original frame is a unit norm frame. This leads to the following problem:

Problem 8. *For every (M, N) such that $M \leq N$, find δ , C and α such that for every unit norm frame Φ satisfying $\|\Phi\Phi^* - (N/M)I\|_{\text{HS}} \leq \delta$, there exists a UNTF Φ^* such that*

$$\|\Phi^* - \Phi\|_{\text{HS}} \leq C \left\| \Phi\Phi^* - \frac{N}{M}I \right\|_{\text{HS}}^\alpha.$$

Having distilled the problem of interest, we now identify the applicability of the frame potential: Since local minimizers of the frame potential are UNTFs, it makes sense to attempt gradient descent to approach a nearby UNTF from an initial unit norm frame. Unfortunately, gradient descent is not guaranteed to converge from any arbitrary initial arrangement; indeed, the gradient vanishes at any arrangement which satisfies the Lagrange equations (the “critical points”), which occurs precisely when the arrangement can be partitioned into mutually orthogonal sub-arrangements which form UNTFs for their spans [1].

Call an arrangement **orthogonally partitionable (OP)** if there is a nontrivial partition of the arrangement into mutually orthogonal sub-arrangements, and ϵ -OP if one can partition into sub-arrangements such that inner products between members of different sub-arrangements are less than ϵ in absolute value. Intuitively, gradient descent produces a sequence of unit norm frames which are not ϵ -OP, then the gradient at these arrangements will never vanish, and so we should expect good performance; indeed, in this case, we get convergence at a linear rate, and furthermore, the limiting UNTF is within a factor of $\|\Phi\Phi^* - (N/M)I\|_{\text{HS}}$ away from the initial arrangement Φ (i.e., $\alpha = 1$ in Problem 8, which is optimal) [5]. Additionally, if $\|\Phi\Phi^* - (N/M)I\|_{\text{HS}}$ is sufficiently small, then Φ is necessarily not ϵ -OP, provided M and N are relatively prime; this makes intuitive sense since a UNTF cannot be OP when M and N are relatively prime (why not?). As such, one may perform gradient descent of the frame potential to solve the Paulsen problem when M and N are relatively prime.

In the remaining case where M and N are not relatively prime, there is no ϵ for which we can guarantee that the gradient descent iterations never become ϵ -OP (see Example 9 in [5]), and so we need a different approach. In this case, if ever the arrangement becomes ϵ -OP, then we can “jump” to a nearby arrangement which is exactly OP, and if the original was sufficiently tight, we can ensure that the new arrangement’s partitions have equal redundancies (that is, for each sub-arrangement, the ratio of number of points to the dimension of their span is constant). We can then run gradient descent on each sub-arrangement individually, keeping its span fixed, and jumping again when necessary. Each time we need to jump, we incur a penalty in our estimate of the distance to the limiting UNTF. In particular, $\alpha = 1/7^J$ in Problem 8 if J jumps are used, and since $J \leq M$, we can at least guarantee $\alpha = 1/7^M$ in this non-relatively-prime case [5]. Improving on this value of α remains an open problem.

5 Eigensteps

The previous section completely characterized the set of UNTFs as local minimizers of the frame potential. We also saw that one can perform gradient descent on the frame potential (with intermittent “jumps” as necessary) to find a UNTF close to a given unit norm frame which is nearly tight. While this characterization is useful (both for intuition and for solving the Paulsen problem), it is not explicit. This section provides an explicit construction of every UNTF of N vectors in M dimensions, real or complex.

Before explaining the theory behind this general construction, it is instructive to make an attempt without the theory. The following is the beginning of an arbitrary real construction with $(M, N) = (3, 5)$:

$$\begin{bmatrix} 1 & \cos \theta & ? & ? & ? \\ 0 & \sin \theta & ? & ? & ? \\ 0 & 0 & ? & ? & ? \end{bmatrix}.$$

Indeed, the first two frame elements are completely determined (up to orthogonal rotation) by the angle θ between them. We should keep track of the spectrum of $\Phi\Phi^*$ since the addition of columns will only increase the eigenvalues, and in the end, we need all of the eigenvalues to be $5/3$. At the moment, the spectrum is completely determined by θ :

$$\lambda(\Phi_2\Phi_2^*) = (1 + \cos \theta, 1 - \cos \theta, 0).$$

(We are denoting the first n columns of Φ by Φ_n .) Also, the eigenspaces of $\Phi_2\Phi_2^*$ are the bisector of the frame elements, the line in the xy -plane which is orthogonal to this bisector, and the z -axis. The corresponding eigenvalues indicate how much of the frame elements’ energy lie in these subspaces. In particular, since the eigenvalues only increase with the addition of more frame elements, we know that $1 + \cos \theta \leq 5/3$, and so θ cannot be too small. We now need to carefully pick a third frame element so as to not have too much energy pointing in the direction of the leading eigenvector. But what choices are available? And are the choices parameterized by some parameter like θ ?

At this point, we recall that it is difficult to find analytic expressions for eigenvalues of $M \times M$ matrices when $M \geq 5$. As such, one might suspect that keeping track of the spectrum of $\Phi\Phi^*$ is infeasible. This suspicion happens to be wrong; while it may be difficult to find a closed-form expression for eigenvalues in terms of matrix entries, it is easy to analytically express the entries of a diagonalizable matrix in terms of its eigenvalues and eigenvectors. In similar spirit, if we know how we want the spectrum of $\Phi_n\Phi_n^*$ to evolve as n ranges from 1 to N , then we can use this to produce analytic expressions for the frame elements. In particular, if we want the spectrum to change from $(\alpha_1, \dots, \alpha_M)$ to $(\beta_1, \dots, \beta_M)$ by adding a frame element φ_{n+1} , then by Theorem 2 in [3], we can determine the necessary size of the projection of φ_{n+1} onto each of the eigenspaces of $\Phi_n\Phi_n^*$:

$$\|P_{n;\lambda}\varphi_{n+1}\|^2 = -\lim_{x \rightarrow \lambda} (x - \lambda) \frac{(x - \beta_1) \cdots (x - \beta_M)}{(x - \alpha_1) \cdots (x - \alpha_M)},$$

where $P_{n;\lambda}$ denotes the projection onto the eigenspace of $\Phi_n\Phi_n^*$ corresponding to eigenvalue λ . Given these sizes of projections, we have the freedom to choose the orientations of the projections, and any such choice will uniquely determine the new frame element φ_{n+1} and update the spectrum as desired. Once this choice has been made, there is a corresponding analytic expression for how the eigenspaces of $\Phi_{n+1}\Phi_{n+1}^*$ are updated from those of $\Phi_n\Phi_n^*$; this expression is a bit more complicated, as it depends on the dimensions of the eigenspaces, and so the interested reader should consult Theorem 7 in [3].

Overall, we have a method of constructing all UNTFs with a prescribed sequence of spectra $\{\lambda(\Phi_n \Phi_n^*)\}_{n=1}^N$, and so it remains to determine all possible sequences of spectra. Since $\Phi_1 \Phi_1^*$ is a rank-1 projection, we know that the first spectrum is simply $(1, 0, \dots, 0)$. Also, since $\Phi_N \Phi_N^* = \Phi \Phi^* = (N/M)I$, the last spectrum must be $(N/M, \dots, N/M)$. Every time we add a frame element, the energy of that frame element contributes to the spectrum:

$$\text{Tr}[\Phi_{n+1} \Phi_{n+1}^*] = \text{Tr}[\Phi_n \Phi_n^* + \varphi_{n+1} \varphi_{n+1}^*] = \text{Tr}[\Phi_n \Phi_n^*] + \text{Tr}[\varphi_{n+1} \varphi_{n+1}^*] = \text{Tr}[\Phi_n \Phi_n^*] + 1.$$

In words, the sum of the eigenvalues increases by 1 with each frame element. The least obvious observation to make is a classical result that the addition of a positive semidefinite rank-1 matrix to a matrix with spectrum $\alpha = (\alpha_1, \dots, \alpha_M)$ produces a new matrix whose spectrum $\beta = (\beta_1, \dots, \beta_M)$ **interlaces** the original:

$$\alpha_N \leq \beta_N \leq \alpha_{N-1} \leq \beta_{N-1} \leq \dots \leq \alpha_2 \leq \beta_2 \leq \alpha_1 \leq \beta_1.$$

As shorthand, we write $\alpha \sqsubseteq \beta$. With these four observations, we define **eigensteps**² to be any collection of sequences $\{\{\lambda_{m;n}\}_{m=1}^M\}_{n=1}^N$ such that

- the first sequence is $(\lambda_{1;1}, \lambda_{2;1}, \dots, \lambda_{M;1}) = (1, 0, \dots, 0)$,
- the last sequence is $(\lambda_{1;N}, \dots, \lambda_{M;N}) = (N/M, \dots, N/M)$,
- the trace is consistent with unit-norm frame elements, i.e., $\sum_{m=1}^M \lambda_{m;n} = n$ for each n , and
- each sequence interlaces the previous: $\{\lambda_{m;n}\}_{m=1}^M \sqsubseteq \{\lambda_{m;n+1}\}_{m=1}^M$.

By design, we know that for every UNTF Φ , the sequence of spectra $\{\lambda(\Phi_n \Phi_n^*)\}_{n=1}^N$ is necessarily a sequence of eigensteps. It turns out that the converse also holds: For every sequence of eigensteps, there exists a UNTF with a matching sequence of spectra (this is the main result in [3]). In summary, every UNTF can be constructed by first selecting eigensteps, and then iteratively selecting frame elements in terms of these eigensteps using the process described earlier. For the sake of example, we return to the case of 5 frame elements in \mathbb{R}^3 :

Example 9. Instead of attempting to write down a general family of UNTFs directly, we start by finding eigensteps and then use these eigensteps as instructions for producing all UNTFs. In the following, we express eigensteps in an $M \times N$ matrix Λ whose (m, n) th entry is $\lambda_{m;n}$. The first and second properties of eigensteps determine the first and last columns:

$$\Lambda = \begin{bmatrix} 1 & ? & ? & ? & \frac{5}{3} \\ 0 & ? & ? & ? & \frac{5}{3} \\ 0 & ? & ? & ? & \frac{5}{3} \end{bmatrix}.$$

Next by interlacing, we know that equal eigenvalues in one sequence determines an eigenvalue in the next/previous:

$$\Lambda = \begin{bmatrix} 1 & ? & \frac{5}{3} & \frac{5}{3} & \frac{5}{3} \\ 0 & ? & ? & \frac{5}{3} & \frac{5}{3} \\ 0 & 0 & ? & ? & \frac{5}{3} \end{bmatrix}.$$

²There are two meanings behind the name *eigensteps*. First, each spectrum in the sequence is another step toward the desired UNTF. Second, when working with the interlacing constraint, it is helpful to visualize each spectrum in the sequence in terms of a bar graph, which resembles a staircase.




Figure 2: Pairs of parameters (x, y) for which the eigensteps (4) satisfy the interlacing constraints (5). Specifically, (a) illustrates the various half-spaces implicated by the interlacing inequalities, whereas (b) provides their intersection. Picking any (x, y) in this pentagon and plugging into (4) will produce a valid sequence of eigensteps, which can then be used as instructions for constructing a 3×5 unit norm tight frame. Also, for every such unit norm tight frame, there is a corresponding point in this pentagon from which it can be constructed. In this way, eigensteps provide an explicit construction of every unit norm tight frame.

The trace condition then forces the entries in each column to sum to 1:

$$\Lambda = \begin{bmatrix} 1 & 2 - y & \frac{5}{3} \\ 0 & y & \frac{4}{3} - x \\ 0 & 0 & x \end{bmatrix}. \quad (4)$$

The variables x and y are not completely free—one must ensure that the eigensteps they generate satisfy interlacing (the other three properties of eigensteps are already satisfied at this point). In this case, the interlacing inequalities involving x and y are

$$\begin{aligned} x &\leq \frac{2}{3} \leq \frac{4}{3} - x \leq \frac{5}{3}, \\ 0 &\leq x \leq y \leq \frac{4}{3} - x \leq 2 - y \leq \frac{5}{3}, \\ 0 &\leq y \leq 1 \leq 2 - y. \end{aligned} \quad (5)$$

Overall, the (x, y) which produce eigensteps form a convex polygon, as depicted in Figure 2. In general, the set of eigensteps for a given (M, N) form a convex polytope, and this convex polytope is parameterized in general in [13]. Now that we have identified every possible choice of eigensteps, we can use them as instructions to construct all possible UNTFs. Intuitively, if we were to standardize the directions along eigenspaces that we select in each step, then varying x and y would produce a 2-dimensional submanifold of the set of UNTFs; indeed, for this example, one such standardized selection produces a continuum of 3×5 UNTFs provided in Table 1 in [3], all the entries of which are relatively simple functions of x and y .

5.1 Application: Schur–Horn theorem

Eigensteps can be used to provide a constructive proof of the Schur–Horn theorem. To state this theorem, we require a definition: $\{\lambda_n\}_{n=1}^N$ is said to **majorize** $\{\mu_n\}_{n=1}^N$ if

$$\begin{aligned} \sum_{n'=1}^n \lambda_{n'} &\geq \sum_{n'=1}^n \mu_{n'} & \forall n \in \{1, \dots, N-1\}, \\ \sum_{n'=1}^N \lambda_{n'} &= \sum_{n'=1}^N \mu_{n'}. \end{aligned}$$

As shorthand, we write $\{\lambda_n\}_{n=1}^N \succeq \{\mu_n\}_{n=1}^N$.

Theorem 10 (Schur–Horn theorem, [16, 20]). *There exists a self-adjoint matrix with spectrum $\{\lambda_n\}_{n=1}^N$ and diagonal entries $\{\mu_n\}_{n=1}^N$ if and only if $\{\lambda_n\}_{n=1}^N \succeq \{\mu_n\}_{n=1}^N$.*

To prove this theorem, first note that it suffices to prove it in the special case where the spectrum is nonnegative. Indeed, for any α , we obviously have that the self-adjoint matrix G exists if and only if $G + \alpha I$ exists. Also, $\{\lambda_n\}_{n=1}^N \succeq \{\mu_n\}_{n=1}^N$ is almost immediately equivalent to $\{\lambda_n + \alpha\}_{n=1}^N \succeq \{\mu_n + \alpha\}_{n=1}^N$. As such, we may pick α so that the spectrum is nonnegative without loss of generality. When it exists, the resulting positive semidefinite matrix will enjoy a Cholesky factorization $\Phi^* \Phi$, in which case μ_n corresponds to the norm squared of the n th column of Φ . As such, the Schur–Horn theorem reduces to a generalization of the UNTF existence problem: When does there exist a matrix Φ with prescribed norms such that $\Phi^* \Phi$ has a prescribed spectrum? (Note that in the case where Φ is a UNTF, the spectrum of $\Phi^* \Phi$ is simply a zero-padded version of that of $\Phi \Phi^* = (M/N)I$.) This can be determined by generalizing the theory of eigensteps (i.e., change the final spectrum and the trace condition appropriately); this more general form is provided in [3]. From this perspective, majorization essentially follows from a repeated application of interlacing [17]. The proof then follows from the fact that a particular greedy choice of eigensteps, called **Top Kill**³, will always exist when the spectrum majorizes the squared lengths, and that no eigensteps exist otherwise.

5.2 Application: Optimal frame completion

Consider a scenario in which you currently have a certain collection of measurement vectors, and you are given a budget to construct additional measurement vectors. How should you select the additional measurement vectors? In order to enjoy robust reconstruction, you are inclined to use the additional measurement vectors to complete your current system to a tight frame; this idea was first studied in [8]. Sometimes, it is impossible to complete to a tight frame. For example, if your current frame is

$$\Phi = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 \end{bmatrix},$$

then the spectrum of $\Phi \Phi^*$ is $(3, 2, 2)$, and so a budget of one additional vector is insufficient thanks to interlacing: the largest eigenvalue after completion must be ≥ 3 , whereas the smallest eigenvalue must be 2. In such situations, one is inclined to find a frame completion which is the “best possible.” This problem of **optimal frame completion** was first posed in [12].

³Top Kill was named after the procedure used in an attempt to seal an oil well in response to the 2010 Deepwater Horizon oil spill, which was prominently featured in the news at the time of its development.

Recently, this problem was completely solved using eigensteps [11]. The solution has two parts. First, [11] characterizes every possible completion of a given frame with a budget of additional frame elements of prescribed lengths; in the special case where the additional frame elements each have length zero, this result recovers the Schur–Horn theorem, and so this characterization is a generalization of sorts. Second, [11] explicitly identifies the completion whose spectrum is majorized by the spectrum of every other completion; the classical theory of Schur convexity then guarantees that, for any function of the spectrum which is both convex and symmetric (e.g., the frame potential, the sum of the reciprocals of the eigenvalues, etc.), this majorization-minimal completion also minimizes this function of the spectrum. As such, the identified completion is optimal, essentially independent of the notion of optimality.

5.3 Application: UNTF connectivity

For a fixed (M, N) , is the set of real (complex) UNTFs path-connected? This question was first posed by David Larson in a Research Experiences for Undergraduates summer program in 2002. UNTFs are natural generalizations of orthonormal bases, and so this question is a natural analog to the facts that the orthogonal group is not connected, while the unitary group is. However, only special cases of this **frame homotopy problem** were solved until recently, when [4] provided a complete solution using eigensteps.

To understand how eigensteps play a role in the solution of this problem, it is instructive to recall the proof that the unitary group is connected. Every unitary matrix U is normal (since $UU^* = I = U^*U$), and so it is diagonalized by some unitary matrix S . Since U acts as an isometry, its eigenvalues necessarily have unit modulus, and so we have a factorization $U = SDS^*$, where the diagonal entries of D have unit modulus. At this point, we observe that D can be continuously deformed to the identity matrix by letting the diagonal entries traverse the complex unit circle at appropriate rates; denote this matrix path by $D(t)$ for $0 \leq t \leq 1$. Then $SD(t)S^*$ is unitary for every t since it is a product of unitary matrices, and $SD(1)S^* = SS^* = I$. Since every unitary matrix is path-connected to the identity matrix, every pair of unitary matrices then enjoys a connecting path by transitivity.

Now consider what the theory of eigensteps provides: a two-step procedure to construct every possible UNTF. The first step is to choose eigensteps, and the second step is to choose directions along each iteration’s eigenspaces. One may think of this as a sort of polar decomposition of UNTFs. Just like the unitary group’s connectivity is proved by fixing certain factors in the requisite unitary diagonalization, the collection of complex UNTFs is proved to be connected by fixing directionalities and then fixing eigensteps in the polar decomposition. Indeed, given two complex UNTFs Φ and Ψ , one may continuously deform the eigensteps of Φ to match those of Ψ (this is possible since the set of all eigensteps is a convex polytope), and then continuously deform the directionalities by leveraging the unitary group’s connectivity. Actually, the argument is a little more subtle if either frame’s eigensteps lies on the boundary of the polytope, and the argument is a bit more complicated in the real case since the orthogonal group is not connected; the interested reader is invited to find more details in [4].

Acknowledgments

The author thanks Matt Fickus and Pete Casazza for introducing him to frames with contagious enthusiasm. This work was supported by NSF Grant No. DMS-1321779. The views expressed in these notes are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

References

- [1] J. J. Benedetto, M. Fickus, Finite normalized tight frames, *Adv. Comput. Math.* 18 (2003) 357–385.
- [2] B. G. Bodmann, P. G. Casazza, The road to equal-norm Parseval frames, *J. Funct. Anal.* 258 (2010) 397–420.
- [3] J. Cahill, M. Fickus, D. G. Mixon, M. J. Poteet, N. Strawn, Constructing finite frames of a given spectrum and set of lengths, *Appl. Comput. Harmon. Anal.* 35 (2013) 52–73.
- [4] J. Cahill, D. G. Mixon, N. Strawn, Connectivity and irreducibility of algebraic varieties of finite unit norm tight frames, Available online: arXiv:1311.4748
- [5] P. G. Casazza, M. Fickus, D. G. Mixon, Auto-tuning unit norm frames, *Appl. Comput. Harmon. Anal.* 32 (2012) 1–16.
- [6] P. G. Casazza, M. Fickus, D. G. Mixon, Y. Wang, Z. Zhou, Constructing tight fusion frames, *Appl. Comput. Harmon. Anal.* 20 (2011) 175–187.
- [7] P. G. Casazza, A. Heinecke, F. Krahmer, G. Kutyniok, Optimally sparse frames, *IEEE Trans. Inform. Theory* 57 (2011) 7279–7287.
- [8] D.-J. Feng, L. Wang, Y. Wang, Generation of finite tight frames by Householder transformations, *Adv. Comput. Math.* 24 (2006) 297–309.
- [9] J. F. Fernando, J. M. Gamboa, On Lojasiewicz inequality and the nullstellenstanz for rings of semialgebraic functions, *Math. Res. Lett.* 16 (2010) 10001–10015.
- [10] M. Fickus, Constructions of normalized tight frames, preprint.
- [11] M. Fickus, J. Marks, M. J. Poteet, A generalized Schur-Horn theorem and optimal frame completions, Available online: arXiv:1408.2882
- [12] M. Fickus, D. G. Mixon, M. J. Poteet, Frame completions for optimally robust reconstruction, *Proc. SPIE* 8138 (2011) 81380Q/1–8.
- [13] M. Fickus, D. G. Mixon, M. J. Poteet, N. Strawn, Constructing all self-adjoint matrices with prescribed spectrum and diagonal, *Adv. Comput. Math.* 39 (2013) 585–609.
- [14] V. K. Goyal, J. Kovačević, J. A. Kelner, Quantized frame expansions with erasures, *Appl. Comput. Harmon. Anal.* 10 (2001) 203–233.
- [15] M. A. Herman, T. Strohmer, High-resolution radar via compressed sensing, *IEEE Trans. Signal Process.* 57 (2009) 2275–2284.
- [16] A. Horn, Doubly stochastic matrices and the diagonal of a rotation matrix, *Amer. J. Math.* 76 (1954) 620–630.
- [17] R. A. Horn, C. R. Johnson, *Matrix Analysis*, Cambridge University Press, Cambridge, 1985.
- [18] D. G. Mixon, C. J. Quinn, N. Kiyavash, M. Fickus, Fingerprinting with equiangular tight frames, *IEEE Trans. Inform. Theory* 59 (2013) 1855–1865.

- [19] Platonic solid, Wikipedia, Available online: http://en.wikipedia.org/wiki/Platonic_solid
- [20] I. Schur, Über eine klasse von mittelbildungen mit anwendungen auf die determinationentheorie, Sitzungsber. Berl. Math. Ges. 22 (1923) 9–20.
- [21] S. Smale, Mathematical problems for the next century, Math. Intell. 20 (1998) 7–15.
- [22] J. J. Thomson, On the structure of the atom: An investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; With application of the results to the theory of atomic structure, Philosophical Mag. 6, vol. 7, (1904) 237–265.
- [23] J. A. Tropp, On the conditioning of random subdictionaries, Appl. Comput. Harmon. Anal. 25 (2008) 1–24.
- [24] R. Vale, S. Waldron, Tight frames and their symmetries, Constr. Approx. 21 (2005) 83–112.
- [25] L. R. Welch, Lower bounds on the maximum cross correlation of signals, IEEE Trans. Inform. Theory 20 (1974) 397–399.
- [26] G. Zimmermann, Normalized tight frames in finite dimensions, In: Recent Progress in Multivariate Approximation, Eds., K. J. W. Haussmann, M. Reimer, Birkhäuser, Basel, 2001, pp. 249–252.