FULL SPARK FRAMES

BORIS ALEXEEV, JAMESON CAHILL, AND DUSTIN G. MIXON

ABSTRACT. Finite frame theory has a number of real-world applications. In applications like sparse signal processing, data transmission with robustness to erasures, and reconstruction without phase, there is a pressing need for deterministic constructions of frames with the following property: every size-M subcollection of the M-dimensional frame elements is a spanning set. Such frames are called full spark frames, and this paper provides constructions using Vandermonde matrices and discrete Fourier transforms. Later, we prove that full spark Parseval frames are dense in the entire set of Parseval frames, meaning full spark frames are abundant, even if one imposes an additional tightness constraint. Finally, we prove that testing whether a given matrix is full spark is hard for NP under randomized polynomial-time reductions, indicating that deterministic full spark constructions are particularly significant because they guarantee a property which is otherwise difficult to check.

1. INTRODUCTION

A *frame* over a Hilbert space $\mathcal{H}$ is a collection of vectors $\{f_i\}_{i \in \mathcal{I}} \subseteq \mathcal{H}$ with *frame bounds* $0 < A \leq B < \infty$ such that

$$A\|x\|^2 \leq \sum_{i \in \mathcal{I}} |\langle x, f_i \rangle|^2 \leq B\|x\|^2 \qquad \forall x \in \mathcal{H}. \tag{1}$$

For a finite frame $\{f_n\}_{n=1}^N$ in an M-dimensional Hilbert space $\mathbb{H}_M$, the upper frame bound B which satisfies (1) trivially exists. In this finite-dimensional setting, the notion of being a frame solely depends on the lower frame bound A being strictly positive, which is equivalent to having the frame elements $\{f_n\}_{n=1}^N$ span the Hilbert space.

In practice, frames are chiefly used in two ways. The *synthesis operator* $F \colon \mathbb{C}^N \to \mathbb{H}_M$ of a finite frame $F = \{f_n\}_{n=1}^N$ is defined by $Fy := \sum_{n=1}^N y[n]f_n$. As such, the $M \times N$ matrix representation F of the synthesis operator has the frame elements $\{f_n\}_{n=1}^N$ as columns, with the natural identification $\mathbb{H}_M \cong \mathbb{C}^M$. Note that here and throughout, we make no notational distinction between a frame F and its synthesis operator. The adjoint of the synthesis operator is called the *analysis operator* $F^* \colon \mathbb{H}_M \to \mathbb{C}^N$, defined by $(F^*x)[n] := \langle x, f_n \rangle$.

Classically, frames are used to redundantly decompose signals, $y = F^*x$, before synthesizing the corresponding frame coefficients, $z = Fy = FF^*x$, and so the *frame operator*

$FF^*\colon \mathbb{H}_M \to \mathbb{H}_M$ is often analyzed to determine how well this process preserves information about the original signal x. In particular, if the frame bounds are equal, the frame operator has the form $FF^* = AI_M$, and so signal reconstruction is rather painless: $x = \frac{1}{A}FF^*x$; in this case, the frame is called *tight*. Oftentimes, it is additionally desirable for the frame elements to have unit norm, in which case the frame is a *unit norm frame*. Moreover, the *worst-case coherence* between unit norm frame elements $\mu := \max_{i \neq j} |\langle f_i, f_j \rangle|$ satisfies $\mu^2 \geq \frac{N-M}{M(N-1)}$, and when equality is achieved, the frame is necessarily tight with $|\langle f_i, f_j \rangle| = \mu$ for all distinct pairs $i, j \in \{1, \ldots, N\}$ [34]; in this case, the frame is called an *equiangular tight frame* (ETF).

The utility of unit norm tight frames and ETFs is commonly expressed in terms of a scenario in which frame coefficients $\{(F^*x)[n]\}_{n=1}^N$ are transmitted over a noisy or lossy channel before reconstructing the signal:

$$y = \mathcal{D}(F^*x), \qquad \tilde{x} = (FF^*)^{-1}Fy, \tag{2}$$

where $\mathcal{D}(\cdot)$ represents the channel's random and not-necessarily-linear deformation process. Using an additive white Gaussian noise model, Goyal [16] established that, of all unit norm frames, unit norm tight frames minimize mean squared error in reconstruction. For the case of a lossy channel, Holmes and Paulsen [18] established that, of all tight frames, unit norm tight frames minimize worst-case error in reconstruction after one erasure, and that equiangular tight frames minimize this error after two erasures. We note that the reconstruction process in (2), namely the application of $(FF^*)^{-1}F$, is inherently blind to the effect of the deformation process of the channel. This contrasts with Püschel and Kovačević's more recent work [26], which describes an adaptive process for reconstruction after multiple erasures; we will return to this idea later.

Another application of frames is sparse signal processing. This field concerns signal classes which are sparse in some basis. As an example, natural images tend to be nearly sparse in the wavelet basis [9]. Some applications have signal sparsity in the identity basis [23]. Given a signal x, let Ψ represent its sparsifying basis and consider

$$y = F\Psi^*x + e, \tag{3}$$

where F is the $M \times N$ synthesis operator of a frame, Ψ^*x has at most K nonzero entries, and e is some sort of noise. When given measurements y, one typically wishes to reconstruct the original vector x; this process is commonly referred to as *compressed sensing* since we often take $M \ll N$. Note that y can be viewed as a noisy linear combination of a few unknown frame elements, and the goal of compressed sensing is to determine which frame elements are active in the combination, and further estimate the scalar multiples used in this combination.

In order to have any sort of inversion process for (3), even in the noiseless case where $e = 0$, F must map K-sparse vectors injectively. That is, we need $Fx_1 \neq Fx_2$ for any distinct K-sparse vectors x_1 and x_2. Subtraction then gives that $2K$-sparse vectors like $x_1 - x_2$ cannot be in the nullspace of F. This identification led to the following definition [11]:

Definition 1. *The* spark *of a matrix F is the size of the smallest linearly dependent subset of columns, i.e.,*

$$\mathrm{Spark}(F) = \min\{\|x\|_0 : Fx = 0, \ x \neq 0\}.$$

Using the above analysis, Donoho and Elad [11] showed that a signal x with sparsity level $< \mathrm{Spark}(F)/2$ is necessarily the unique sparsest solution to $y = Fx$, and furthermore, there exist signals x of sparsity level $\geq \mathrm{Spark}(F)/2$ which are not the unique sparsest solution to $y = Fx$. This demonstrates that matrices with larger spark are naturally equipped to support signals with larger sparsity levels. One is naturally led to consider matrices that support the largest possible sparse signal class; we say an $M \times N$ matrix F is *full spark* if its spark is as large as possible, i.e., $\mathrm{Spark}(F) = M + 1$. Equivalently, $M \times N$ full spark matrices have the property that every $M \times M$ submatrix is invertible; as such, a full spark matrix is necessarily full rank, and therefore a frame.

Consider a matrix whose entries are independent continuous random variables. Certainly, the matrix is full spark with probability one, but this random process does not allow one to control certain features of the matrix, such as unit-norm tightness or equiangularity. Indeed, ETFs are notoriously difficult to construct, but they appear to be particularly well-suited for sparse signal processing [1, 14, 23]. For example, Bajwa et al. [1] uses Steiner ETFs to recover the support of $\Psi^* x$ in (3); given measurements y, the largest entries in the back-projection $F^* y$ often coincide with the support of $\Psi^* x$. However, Steiner ETFs have particularly small spark [15], and back-projection will correctly identify the support of x even when the corresponding columns in F are dependent; in this case, there is no way to estimate the nonzero entries of $\Psi^* x$. This illustrates one reason to build deterministic full spark frames: their submatrices are necessarily invertible, making it possible to estimate these nonzero entries.

For another application of full spark frames, we return to the problem (2) of reconstructing a signal from distorted frame coefficients. Specifically, we consider Püschel and Kovačević's work [26], which focuses on an Internet-like channel that is prone to multitudes of erasures. In this context, they reconstruct the signal after first identifying which frame coefficients were not erased; with this information, the signal can be estimated provided the corresponding frame elements span. In this sense, full spark frames are maximally robust to erasures. In particular, an $M \times N$ full spark frame is robust to $N - M$ erasures since any M of the frame coefficients will uniquely determine the original signal.

Yet another application of full spark frames is phaseless reconstruction, which can be viewed in terms of a channel, as in (2); in this case, $\mathcal{D}(\cdot)$ is the entrywise absolute value function. Phaseless reconstruction has a number of real-world applications including speech processing [4], X-ray crystallography [6], and quantum state estimation [29]. As such, there has been a lot of work to reconstruct an M-dimensional vector (up to an overall phase factor) from the magnitudes of its frame coefficients, most of which involves frames in operator space, which inherently require $N = \Omega(M^2)$ measurements [3, 29]. However, Balan et al. [4] show that if an $M \times N$ real frame F is full spark with $N \geq 2M - 1$, then $\mathcal{D} \circ F^*$ is injective, meaning an inversion process is possible with only $N = O(M)$ measurements. This result prompted an ongoing search for efficient phaseless reconstruction processes [2, 6], but no reconstruction process can succeed without a good family of frames, such as full spark frames.

Despite the fact that full spark frames have a multitude of applications, to date, there has not been much progress in constructing deterministic full spark frames, let alone full spark frames with additional desirable properties. A noteworthy exception is Püschel and Kovačević's work [26], in which real full spark tight frames are constructed using polynomial

transforms. In the present paper, we start by investigating Vandermonde frames, harmonic frames, and modifications thereof. We find some classes of full spark ETFs, and characterize the $M \times N$ full spark harmonic frames in which N is a prime power. In Section 3, we show that $M \times N$ full spark Parseval frames form a dense subset of the entire collection of $M \times N$ Parseval frames. As such, full spark frames are abundant, even after imposing the additional condition of tightness. This result is balanced with Section 4, in which we prove that verifying whether a matrix is full spark is hard for NP under randomized polynomial-time reductions. In other words, assuming NP $\not\subseteq$ BPP (a computational complexity assumption slightly stronger than P $\neq$ NP and nearly as widely believed), then there is no method by which one can efficiently test whether matrices are full spark. As such, the deterministic constructions in Section 2 are significant in that they guarantee a property which is otherwise difficult to check.

2. Deterministic constructions of full spark frames

A square matrix is invertible if and only if its determinant is nonzero, and in our quest for deterministic constructions of full spark frames, this characterization will reign supreme. One class of matrices has a particularly simple determinant formula: Vandermonde matrices. Specifically, Vandermonde matrices have the following form:

$$V = \begin{bmatrix} 1 & 1 & \cdots & 1 \\ \alpha_1 & \alpha_2 & \cdots & \alpha_N \\ \vdots & \vdots & \vdots & \vdots \\ \alpha_1^{M-1} & \alpha_2^{M-1} & \cdots & \alpha_N^{M-1} \end{bmatrix}, \tag{4}$$

and square Vandermonde matrices, i.e., with $N = M$, have the following determinant:

$$\det(V) = \prod_{1 \leq i < j \leq M} (\alpha_j - \alpha_i). \tag{5}$$

Consider (4) in the case where $N \geq M$. Since every $M \times M$ submatrix of V is also Vandermonde, we can modify the indices in (5) to calculate the determinant of the submatrices. These determinants are nonzero precisely when the bases $\{\alpha_n\}_{n=1}^N$ are distinct, yielding the following result:

Lemma 2. *A Vandermonde matrix is full spark if and only if its bases are distinct.*

When viewed in terms of frame theory, Vandermonde matrices naturally point to the discrete Fourier transform:

Theorem 3. *The only $M \times N$ Vandermonde matrices that are equal norm and tight have bases in the complex unit circle. Among these, the frames with the smallest worst-case coherence have bases that are equally spaced in the complex unit circle, provided $N \geq 2M$.*

Proof. Suppose a Vandermonde matrix is equal norm and tight. Note that a zero base will produce the zeroth identity basis element δ_0. Letting $\mathcal{P}$ denote the indices of the nonzero bases, the fact that the matrix is full rank implies $|\mathcal{P}| \geq M - 1$. Also, equal norm gives that

the frame element length

$$\|f_n\|^2 = \sum_{m=0}^{M-1} |f_n[m]|^2 = \sum_{m=0}^{M-1} |\alpha_n^m|^2 = \sum_{m=0}^{M-1} |\alpha_n|^{2m}$$

is constant over $n \in \mathcal{P}$. Since $\sum_{m=0}^{M-1} x^{2m}$ is strictly increasing over $0 < x < \infty$, there exists $c > 0$ such that $|\alpha_n|^2 = c$ for all $n \in \mathcal{P}$. Next, tightness gives that the rows have equal norm, implying that the first two rows have equal norm, i.e., $|\mathcal{P}|c = |\mathcal{P}|c^2$. Thus $c = 1$, and so the nonzero bases are in the complex unit circle. Furthermore, since the zeroth and first rows have equal norm by tightness, we have $|\mathcal{P}| = N$, and so every base is in the complex unit circle.

Now consider the inner product between Vandermonde frame elements whose bases $\{e^{2\pi i x_n}\}_{n=1}^N$ come from the complex unit circle:

$$\langle f_n, f_{n'} \rangle = \sum_{m=0}^{M-1} (e^{2\pi i x_n})^m \overline{(e^{2\pi i x_{n'}})^m} = \sum_{m=0}^{M-1} e^{2\pi i (x_n - x_{n'})m}.$$

We will show that the worst-case coherence comes from the two closest bases. Consider the following function:

$$g(x) := \left| \sum_{m=0}^{M-1} e^{2\pi i x m} \right|^2.$$

We will prove two things about this function:

    (i) $\frac{d}{dx} g(x) < 0$ for every $x \in (0, \frac{1}{2M})$,
    (ii) $g(x) \leq g(\frac{1}{2M})$ for every $x \in (\frac{1}{2M}, 1 - \frac{1}{2M})$.

First, we claim that (i) and (ii) are sufficient to prove our result. To establish this, we first show that the two closest bases $e^{2\pi i x_{n'}}$ and $e^{2\pi i x_{n''}}$ satisfy $|x_{n'} - x_{n''}| \leq \frac{1}{2M}$. Without loss of generality, the n's are ordered in such a way that $\{x_n\}_{n=0}^{N-1} \subseteq [0, 1)$ are nondecreasing. Define

$$d(x_n, x_{n+1}) := \begin{cases} x_{n+1} - x_n, & n = 0, \ldots, N-2 \\ x_0 - (x_{N-1} - 1), & n = N-1, \end{cases}$$

and let n' be the n which minimizes $d(x_n, x_{n+1})$. Since the minimum is less than the average, we have

$$d(x_{n'}, x_{n'+1}) \leq \frac{1}{N} \left((x_0 - (x_{N-1} - 1)) + \sum_{n=0}^{N-1} (x_{n+1} - x_n) \right) = \frac{1}{N} \leq \frac{1}{2M}, \qquad (6)$$

provided $N \geq 2M$. Note that if we view $\{x_n\}_{n \in \mathbb{Z}_N}$ as members of $\mathbb{R}/\mathbb{Z}$, then $d(x_n, x_{n+1}) = x_{n+1} - x_n$. Since $g(x)$ is even, then (i) implies that $|\langle f_{n'+1}, f_{n'} \rangle|^2 = g(x_{n'+1} - x_{n'})$ is larger than any other $g(x_p - x_{p'}) = |\langle f_p, f_{p'} \rangle|^2$ in which $x_p - x_{p'} \in [0, \frac{1}{2M}] \cup [1 - \frac{1}{2M}, 1)$. Next, (6) and (ii) together imply that $|\langle f_{n'+1}, f_{n'} \rangle|^2 = g(x_{n'+1} - x_{n'}) \geq g(\frac{1}{2M})$ is larger than any other $g(x_p - x_{p'}) = |\langle f_p, f_{p'} \rangle|^2$ in which $x_p - x_{p'} \in (\frac{1}{2M}, 1 - \frac{1}{2M})$, provided $N \geq 2M$. Combined, (i) and (ii) give that $|\langle f_{n'+1}, f_{n'} \rangle|$ achieves the worst-case coherence of $\{f_n\}_{n \in \mathbb{Z}_N}$. Additionally, (i) gives that the worst-case coherence $|\langle f_{n'+1}, f_{n'} \rangle|$ is minimized when $x_{n'+1} - x_{n'}$ is maximized, i.e., when the x_n's are equally spaced in the unit interval.

To prove (i), note that the geometric sum formula gives

$$g(x) = \left| \sum_{m=0}^{M-1} e^{2\pi ixm} \right|^2 = \left| \frac{e^{2M\pi ix} - 1}{e^{2\pi ix} - 1} \right|^2 = \frac{2 - 2\cos(2M\pi x)}{2 - 2\cos(2\pi x)} = \left(\frac{\sin(M\pi x)}{\sin(\pi x)} \right)^2, \qquad (7)$$

where the final expression uses the identity $1 - \cos(2z) = 2\sin^2 z$. To show that g is decreasing over $(0, \frac{1}{2M})$, note that the base of (7) is positive on this interval, and performing the quotient rule to calculate its derivative will produce a fraction whose denominator is nonnegative and whose numerator is given by

$$M\pi \sin(\pi x) \cos(M\pi x) - \pi \sin(M\pi x) \cos(\pi x). \qquad (8)$$

This factor is zero at $x = 0$ and has derivative:

$$-(M^2 - 1)\pi^2 \sin(\pi x) \sin(M\pi x),$$

which is strictly negative for all $x \in (0, \frac{1}{2M})$. Hence, (8) is strictly negative whenever $x \in (0, \frac{1}{2M})$, and so $g'(x) < 0$ for every $x \in (0, \frac{1}{2M})$.

For (ii), note that for every $x \in (\frac{1}{2M}, 1 - \frac{1}{2M})$, we can individually bound the numerator and denominator of what the geometric sum formula gives:

$$g(x) = \left| \sum_{m=0}^{M-1} e^{2\pi ixm} \right|^2 = \frac{|e^{2M\pi ix} - 1|^2}{|e^{2\pi ix} - 1|^2} \leq \frac{|e^{\pi i} - 1|^2}{|e^{\pi i/M} - 1|^2} = \left| \sum_{m=0}^{M-1} e^{\pi im/M} \right|^2 = g(\tfrac{1}{2M}). \qquad \square$$

Consider the $N \times N$ discrete Fourier transform (DFT) matrix, scaled to have entries of unit modulus:

$$\begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \cdots & \omega^{N-1} \\ 1 & \omega^2 & \omega^4 & \cdots & \omega^{2(N-1)} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \omega^{N-1} & \omega^{2(N-1)} & \cdots & \omega^{(N-1)(N-1)} \end{bmatrix},$$

where $\omega = e^{-2\pi i/N}$. The first M rows of the DFT form a Vandermonde matrix of distinct bases $\{\omega^n\}_{n=0}^{N-1}$; as such, this matrix is full spark by Lemma 2. In fact, the previous result says that this is in some sense an optimal Vandermonde frame, but this might not be the best way to pick rows from a DFT. Indeed, several choices of DFT rows could produce full spark frames, some with smaller coherence or other desirable properties, and so the remainder of this section focuses on full spark DFT submatrices. First, we note that not every DFT submatrix is full spark. For example, consider the 4×4 DFT:

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & -1 & i \\ 1 & -1 & 1 & -1 \\ 1 & i & -1 & -i \end{bmatrix}.$$

Certainly, the zeroth and second rows of this matrix are not full spark, since the zeroth and second columns of this submatrix form the all-ones matrix, which is not invertible. So what can be said about the set of permissible row choices? The following result gives some necessary conditions on this set:

Theorem 4. *Take an $N \times N$ discrete Fourier transform matrix, and select the rows indexed by $\mathcal{M} \subseteq \mathbb{Z}_N$ to build the matrix F. If F is full spark, then so is the matrix built from rows indexed by*

    (i) *any translation of $\mathcal{M}$,*
    (ii) *any $A\mathcal{M}$ with A relatively prime to N,*
    (iii) *the complement of $\mathcal{M}$ in $\mathbb{Z}_N$.*

Proof. For (i), we first define D to be the $N \times N$ diagonal matrix whose diagonal entries are $\{\omega^n\}_{n=0}^{N-1}$. Note that, since $\omega^{(m+1)n} = \omega^n \omega^{mn}$, translating the row indices $\mathcal{M}$ by 1 corresponds to multiplying F on the right by D. For some set $\mathcal{K} \subseteq \mathbb{Z}_N$ of size $M := |\mathcal{M}|$, let $F_\mathcal{K}$ denote the $M \times M$ submatrix of F whose columns are indexed by $\mathcal{K}$, and let $D_\mathcal{K}$ denote the $M \times M$ diagonal submatrix of D whose diagonal entries are indexed by $\mathcal{K}$. Then since $D_\mathcal{K}$ is unitary, we have

$$|\det((FD)_\mathcal{K})| = |\det(F_\mathcal{K} D_\mathcal{K})| = |\det(F_\mathcal{K})||\det(D_\mathcal{K})| = |\det(F_\mathcal{K})|.$$

Thus, if F is full spark, $|\det((FD)_\mathcal{K})| = |\det(F_\mathcal{K})| > 0$, and so FD is also full spark. Using this fact inductively proves (i) for all translations of $\mathcal{M}$.

For (ii), let G denote the submatrix of rows indexed by $A\mathcal{M}$. Then for any set $\mathcal{K} \subseteq \mathbb{Z}_N$ of size M,

$$\det(G_\mathcal{K}) = \det(\omega^{(Am)k})_{m \in \mathcal{M}, k \in \mathcal{K}} = \det(\omega^{m(Ak)})_{m \in \mathcal{M}, k \in \mathcal{K}} = \det(F_{A\mathcal{K}}).$$

Since A is relatively prime to N, multiplication by A permutes the elements of $\mathbb{Z}_N$, and so $A\mathcal{K}$ has exactly M distinct elements. Thus, if F is full spark, then $\det(G_\mathcal{K}) = \det(F_{A\mathcal{K}}) \neq 0$, and so G is also full spark.

For (iii), we let G be the $(N-M) \times N$ submatrix of rows indexed by $\mathcal{M}^{\mathrm{c}}$, so that

$$N I_N = \begin{bmatrix} F^* & G^* \end{bmatrix} \begin{bmatrix} F \\ G \end{bmatrix} = F^* F + G^* G. \tag{9}$$

We will use contraposition to show that F being full spark implies that G is also full spark. To this end, suppose G is not full spark. Then G has a collection of $N-M$ linearly dependent columns $\{g_i\}_{i \in \mathcal{K}}$, and so there exists a nontrivial sequence $\{\alpha_i\}_{i \in \mathcal{K}}$ such that

$$\sum_{i \in \mathcal{K}} \alpha_i g_i = 0.$$

Considering $g_i = G\delta_i$, where δ_i is the ith identity basis element, we can use (9) to express this linear dependence in terms of F:

$$0 = G^* 0 = G^* \sum_{i \in \mathcal{K}} \alpha_i g_i = \sum_{i \in \mathcal{K}} \alpha_i G^* G \delta_i = \sum_{i \in \mathcal{K}} \alpha_i (N I_N - F^* F)\delta_i.$$

Rearranging then gives

$$x := N \sum_{i \in \mathcal{K}} \alpha_i \delta_i = \sum_{i \in \mathcal{K}} \alpha_i F^* F \delta_i. \tag{10}$$

Here, we note that x is nonzero since $\{\alpha_i\}_{i \in \mathcal{K}}$ is nontrivial, and that $x \in \text{Range}(F^*F)$. Furthermore, whenever $j \notin \mathcal{K}$, we have from (10) that

$$\langle x, F^*F\delta_j \rangle = \langle F^*Fx, \delta_j \rangle = N \left\langle F^*F \sum_{i \in \mathcal{K}} \alpha_i \delta_i, \delta_j \right\rangle = N^2 \left\langle \sum_{i \in \mathcal{K}} \alpha_i \delta_i, \delta_j \right\rangle = 0,$$

and so $x \perp \text{Span}\{F^*F\delta_j\}_{j \in \mathcal{K}^c}$. Thus, the containment $\text{Span}\{F^*F\delta_j\}_{j \in \mathcal{K}^c} \subseteq \text{Range}(F^*F)$ is proper, and so

$$M = \text{Rank}(F) = \text{Rank}(F^*F) > \text{Rank}(F^*F_{\mathcal{K}^c}) = \text{Rank}(F_{\mathcal{K}^c}).$$

Since the $M \times M$ submatrix $F_{\mathcal{K}^c}$ is rank-deficient, it is not invertible, and therefore F is not full spark. $\qquad\square$

We note that our proof of (iii) above uses techniques from Cahill et al. [5], and can be easily generalized to prove that the Naimark complement of a full spark tight frame is also full spark. Theorem 4 tells us quite a bit about the set of permissible choices for DFT rows. For example, not only can we pick the first M rows of the DFT to produce a full spark Vandermonde frame, but we can also pick any consecutive M rows, by Theorem 4(i). We would like to completely characterize the choices that produce full spark harmonic frames. The following classical result does this in the case where N is prime:

Theorem 5 (Chebotarëv, see [31]). *Let N be prime. Then every square submatrix of the $N \times N$ discrete Fourier transform matrix is invertible.*

As an immediate consequence of Chebotarëv's theorem, every choice of rows from the DFT produces a full spark harmonic frame, provided N is prime. Note that each of these are equal-norm tight frames by construction. Harmonic frames can also be designed to have minimal coherence; Xia et al. [38] produces harmonic equiangular tight frames by selecting row indices which form a difference set in $\mathbb{Z}_N$. Interestingly, most known families of difference sets in $\mathbb{Z}_N$ require N to be prime [19], and so the corresponding harmonic equiangular tight frames are guaranteed to be full spark by Chebotarëv's theorem. In the following, we use Chebotarëv's theorem to demonstrate full spark for a class of frames which contains harmonic frames, namely, frames which arise from concatenating harmonic frames with any number of identity basis elements:

Theorem 6 (cf. [35, Theorem 1.1]). *Let N be prime, and pick any $M \leq N$ rows of the $N \times N$ discrete Fourier transform matrix to form the harmonic frame H. Next, pick any $K \leq M$, and take D to be the $M \times M$ diagonal matrix whose first K diagonal entries are $\sqrt{\frac{N+K-M}{MN}}$, and whose remaining $M - K$ entries are $\sqrt{\frac{N+K}{MN}}$. Then concatenating DH with the first K identity basis elements produces an $M \times (N+K)$ full spark unit norm tight frame.*

Proof. Let F denote the resulting $M \times (N+K)$ frame. We start by verifying that F is unit norm. Certainly, the identity basis elements have unit norm. For the remaining frame elements, the modulus of each entry is determined by D, and so the norm squared of each frame element is

$$K\left(\tfrac{N+K-M}{MN}\right) + (M-K)\left(\tfrac{N+K}{MN}\right) = 1.$$

To demonstrate that F is tight, it suffices to show that $FF^* = \frac{N+K}{M}I_M$. The rows of DH are orthogonal since they are scaled rows of the DFT, while the rows of the identity portion are orthogonal because they have disjoint support. Thus, FF^* is diagonal. Moreover, the norm squared of each of the first K rows is $N(\frac{N+K-M}{MN}) + 1 = \frac{N+K}{M}$, while the norm squared of each of the remaining rows is $N(\frac{N+K}{MN}) = \frac{N+K}{M}$, and so $FF^* = \frac{N+K}{M}I_M$.

To demonstrate that F is full spark, first note that every $M \times M$ submatrix of DH is invertible since

$$|\det((DH)_{\mathcal{K}})| = |\det(DH_{\mathcal{K}})| = |\det(D)||\det(H_{\mathcal{K}})| > 0,$$

by Chebotarëv's theorem. Also, in the case where $K = M$, we note that the $M \times M$ submatrix of F composed solely of identity basis elements is trivially invertible. The only remaining case to check is when identity basis elements and columns of DH appear in the same $M \times M$ submatrix $F_{\mathcal{K}}$. In this case, we may shuffle the rows of $F_{\mathcal{K}}$ to have the form

$$\begin{bmatrix} A & 0 \\ B & I_K \end{bmatrix}.$$

Since shuffling rows has no impact on the size of the determinant, we may further use a determinant identity on block matrices to get

$$|\det(F_{\mathcal{K}})| = \left| \det \begin{bmatrix} A & 0 \\ B & I_K \end{bmatrix} \right| = |\det(A)\det(I_K)| = |\det(A)|.$$

Since A is a multiple of a square submatrix of the $N \times N$ DFT, we are done by Chebotarëv's theorem. $\qquad\square$

As an example of Theorem 6, pick N to be a prime congruent to 1 mod 4, and select $\frac{N+1}{2}$ rows of the $N \times N$ DFT according to the index set $\mathcal{M} := \{k^2 : k \in \mathbb{Z}_N\}$. If we take $K = 1$, the process in Theorem 6 produces an equiangular tight frame of redundancy 2, which can be verified using quadratic Gauss sums. Note that this corresponds to a special case of a construction in Zauner's thesis [39], which was later studied by Renes [28] and Strohmer [33]. Theorem 6 says that this construction is full spark.

Maximally sparse frames have recently become a subject of active research [8, 15]. We note that when $K = M$, Theorem 6 produces a maximally sparse $M \times (N + K)$ full spark frame, having a total of $M(M - 1)$ zero entries. To see that this sparsity level is maximal, we note that if the frame had any more zero entries, then at least one of the rows would have M zero entries, meaning the corresponding $M \times M$ submatrix would have a row of all zeros and hence a zero determinant. Similar ideas were studied previously by Nakamura and Masson [24].

Another interesting case is where $K = M = N$, i.e., when the frame constructed in Theorem 6 is a union of the unitary DFT and identity bases. Unions of orthonormal bases have received considerable attention in the context of sparse approximation [11, 36]. In fact, when N is a perfect square, concatenating the DFT with an identity basis forms the canonical example F of a dictionary with small spark [11]. To be clear, the Dirac comb of $\sqrt{N}$ spikes is an eigenvector of the DFT, and so concatenating this comb with the negative of its Fourier transform produces a $2\sqrt{N}$-sparse vector in the nullspace of F. In stark contrast, when N is prime, Theorem 6 shows that F is full spark.

The vast implications of Chebotarëv's theorem leads one to wonder whether the result admits any interesting generalization. Certainly, if one could characterize the full spark submatrices of a general DFT, this would provide ample freedom to optimize full spark frames for additional considerations. While we do not have a characterization for the general case, we do have one for the case where N is a prime power. Before stating the result, we require a definition:

Definition 7. *We say a subset $\mathcal{M} \subseteq \mathbb{Z}_N$ is* uniformly distributed over the divisors of N *if, for every divisor d of N, the d cosets of $\langle d \rangle$ partition $\mathcal{M}$ into subsets, each of size $\lfloor \frac{|\mathcal{M}|}{d} \rfloor$ or $\lceil \frac{|\mathcal{M}|}{d} \rceil$.*

At first glance, this definition may seem rather unnatural, but we will discover some important properties of uniformly distributed rows from the DFT. As an example, we take a short detour by considering the restricted isometry property (RIP), which has received considerable attention recently for its use in compressed sensing. We say a matrix F is (K, δ)-RIP if

$$(1 - \delta)\|x\|^2 \leq \|Fx\|^2 \leq (1 + \delta)\|x\|^2 \qquad \text{whenever } \|x\|_0 \leq K.$$

Candès and Tao [7] demonstrated that the sparsest $\Psi^* x$ which satisfies (3) can be found using ℓ_1-minimization, provided F is $(2K, \sqrt{2} - 1)$-RIP. Later, Rudelson and Vershynin [30] showed that a matrix of random rows from a DFT and normalized columns is RIP with high probability. We will show that harmonic frames satisfy RIP only if the selected row indices are nearly uniformly distributed over sufficiently small divisors of N.

To this end, recall that for any divisor d of N, the Fourier transform of the d-sparse normalized Dirac comb $\frac{1}{\sqrt{d}} \chi_{\langle \frac{N}{d} \rangle}$ is the $\frac{N}{d}$-sparse normalized Dirac comb $\sqrt{\frac{d}{N}} \chi_{\langle d \rangle}$. Let F be the $N \times N$ unitary DFT, and let H be the harmonic frame which arises from selecting rows of F indexed by $\mathcal{M}$ and then normalizing the columns. In order for H to be (K, δ)-RIP, $\mathcal{M}$ must contain at least one member of $\langle d \rangle$ for every divisor d of N which is $\leq K$, since otherwise

$$H \frac{1}{\sqrt{d}} \chi_{\langle \frac{N}{d} \rangle} = \sqrt{\frac{N}{|\mathcal{M}|}} (F \frac{1}{\sqrt{d}} \chi_{\langle \frac{N}{d} \rangle})_{\mathcal{M}} = \sqrt{\frac{N}{|\mathcal{M}|}} \left(\sqrt{\frac{d}{N}} \chi_{\langle d \rangle} \right)_{\mathcal{M}} = \sqrt{\frac{d}{|\mathcal{M}|}} \chi_{\mathcal{M} \cap \langle d \rangle} = 0,$$

which violates the lower RIP bound at $x = \frac{1}{\sqrt{d}} \chi_{\langle \frac{N}{d} \rangle}$. In fact, the RIP bounds indicate that

$$\|Hx\|^2 = \|H \frac{1}{\sqrt{d}} \chi_{\langle \frac{N}{d} \rangle}\|^2 = \left\| \sqrt{\frac{d}{|\mathcal{M}|}} \chi_{\mathcal{M} \cap \langle d \rangle} \right\|^2 = \frac{d}{|\mathcal{M}|} |\mathcal{M} \cap \langle d \rangle|$$

cannot be more than δ away from $\|x\|^2 = 1$. Similarly, taking x to be $\frac{1}{\sqrt{d}} \chi_{\langle \frac{N}{d} \rangle}$ modulated by a, i.e., $x[n] := \frac{1}{\sqrt{d}} \chi_{\langle \frac{N}{d} \rangle}[n] e^{2\pi i a n / N}$ for every $n \in \mathbb{Z}_N$, gives that $\|Hx\|^2 = \frac{d}{|\mathcal{M}|} |\mathcal{M} \cap (a + \langle d \rangle)|$ is also no more than δ away from 1. This observation gives the following result:

Theorem 8. *Select rows indexed by $\mathcal{M} \subseteq \mathbb{Z}_N$ from the $N \times N$ discrete Fourier transform matrix and then normalize the columns to produce the harmonic frame H. Then H satisfies the (K, δ)-restricted isometry property only if*

$$\left| |\mathcal{M} \cap (a + \langle d \rangle)| - \frac{|\mathcal{M}|}{d} \right| \leq \frac{|\mathcal{M}|}{d} \delta$$

for every divisor d of N with $d \leq K$ and every $a = 0, \ldots, d - 1$.

Now that we have an intuition for uniform distribution in terms of modulated Dirac combs and RIP, we take this condition to the extreme by considering uniform distribution over all divisors. Doing so produces a complete characterization of full spark harmonic frames when N is a prime power:

Theorem 9. *Let N be a prime power, and select rows indexed by $\mathcal{M} \subseteq \mathbb{Z}_N$ from the $N \times N$ discrete Fourier transform matrix to build the submatrix F. Then F is full spark if and only if $\mathcal{M}$ is uniformly distributed over the divisors of N.*

Note that, perhaps surprisingly, an index set $\mathcal{M}$ can be uniformly distributed over p but not over p^2, and vice versa. For example, $\mathcal{M} = \{0, 1, 4\}$ is uniformly distributed over 2 but not 4, while $\mathcal{M} = \{0, 2\}$ is uniformly distributed over 4 but not 2.

Since the first M rows of a DFT form a full spark Vandermonde matrix, let's check that this index set is uniformly distributed over the divisors of N. For each divisor d of N, we partition the first M indices into the d cosets of $\langle d \rangle$. Write $M = qd + r$ with $0 \le r < d$. The first qd of the M indices are distributed equally amongst all d cosets, and then the remaining r indices are distributed equally amongst the first r cosets. Overall, the first r cosets contain $q + 1 = \lfloor \frac{M}{d} \rfloor + 1$ indices, while the remaining $d - r$ cosets have $q = \lfloor \frac{M}{d} \rfloor$ indices; thus, the first M indices are indeed uniformly distributed over the divisors of N. Also, when N is prime, *every* subset of $\mathbb{Z}_N$ is uniformly distributed over the divisors of N in a trivial sense. In fact, Chebotarëv's theorem follows immediately from Theorem 9. In some ways, portions of our proof of Theorem 9 mirror recurring ideas in the existing proofs of Chebotarëv's theorem [10, 13, 31, 35]. For the sake of completeness, we provide the full argument and save the reader from having to parse portions of proofs from multiple references. We start with the following lemmas, whose proofs are based on the proofs of Lemmas 1.2 and 1.3 in [35].

Lemma 10. *Let N be a power of some prime p, and let $P(z_1, \ldots, z_M)$ be a polynomial with integer coefficients. Suppose there exists Nth roots of unity $\{\omega_m\}_{m=1}^M$ such that $P(\omega_1, \ldots, \omega_M) = 0$. Then $P(1, \ldots, 1)$ is a multiple of p.*

Proof. Denoting $\omega := e^{-2\pi i/N}$, then for every $m = 1, \ldots, M$, we have $\omega_m = \omega^{k_m}$ for some $0 \le k_m < N$. Defining the polynomial $Q(z) := P(z^{k_1}, \ldots, z^{k_M})$, we have $Q(\omega) = 0$ by assumption. Also, $Q(z)$ is a polynomial with integer coefficients, and so it must be divisible by the minimal polynomial of ω, namely, the cyclotomic polynomial $\Phi_N(z)$. Evaluating both polynomials at $z = 1$ then gives that $p = \Phi_N(1)$ divides $Q(1) = P(1, \ldots, 1)$. $\qquad\square$

Lemma 11. *Let N be a power of some prime p, and pick $\mathcal{M} = \{m_i\}_{i=1}^M \subseteq \mathbb{Z}_N$ such that*

$$\frac{\displaystyle\prod_{1 \le i < j \le M} (m_j - m_i)}{\displaystyle\prod_{m=0}^{M-1} m!} \tag{11}$$

is not a multiple of p. Then the rows indexed by $\mathcal{M}$ in the $N \times N$ discrete Fourier transform form a full spark frame.

Proof. We wish to show that $\det(\omega_n^m)_{m \in \mathcal{M}, 1 \le n \le M} \ne 0$ for all M-tuples of distinct Nth roots of unity $\{\omega_n\}_{n=1}^M$. Define the polynomial $D(z_1, \ldots, z_M) := \det(z_n^m)_{m \in \mathcal{M}, 1 \le n \le M}$. Since columns i and j of $(z_n^m)_{m \in \mathcal{M}, 1 \le n \le M}$ are identical whenever $z_i = z_j$, we know that D vanishes in each of these instances, and so we can factor:

$$D(z_1, \ldots, z_M) = P(z_1, \ldots, z_M) \prod_{1 \le i < j \le M} (z_j - z_i)$$

for some polynomial $P(z_1, \ldots, z_M)$ with integer coefficients. By Lemma 10, it suffices to show that $P(1, \ldots, 1)$ is not a multiple of p, since this implies $D(\omega_1, \ldots, \omega_M)$ is nonzero for all M-tuples of distinct Nth roots of unity $\{\omega_n\}_{n=1}^M$.

To this end, we proceed by considering

$$A := \left(z_1 \frac{\partial}{\partial z_1} \right)^0 \left(z_2 \frac{\partial}{\partial z_2} \right)^1 \cdots \left(z_M \frac{\partial}{\partial z_M} \right)^{M-1} D(z_1, \ldots, z_M) \bigg|_{z_1 = \cdots = z_M = 1}. \tag{12}$$

To compute A, we note that each application of $z_j \frac{\partial}{\partial z_j}$ produces terms according to the product rule. For some terms, a linear factor of the form $z_j - z_i$ or $z_i - z_j$ is replaced by z_j or $-z_j$, respectively. For each the other terms, these linear factors are untouched, while another factor, such as $P(z_1, \ldots, z_M)$, is differentiated and multiplied by z_j. Note that there are a total of $M(M-1)/2$ linear factors, and only $M(M-1)/2$ differentiation operators to apply. Thus, after expanding every product rule, there will be two types of terms: terms in which every differentiation operator was applied to a linear factor, and terms which have at least one linear factor remaining untouched. When we evaluate at $z_1 = \cdots = z_M = 1$, the terms with linear factors vanish, and so the only terms which remain came from applying every differentiation operator to a linear factor. Furthermore, each of these terms before the evaluation is of the form $P(z_1, \ldots, z_M) \prod_{1 \le i < j \le M} z_j$, and so evaluation at $z_1 = \cdots = z_M = 1$ produces a sum of terms of the form $P(1, \ldots, 1)$; to determine the value of A, it remains to count these terms. The $M-1$ copies of $z_M \frac{\partial}{\partial z_M}$ can only be applied to linear factors of the form $z_M - z_i$, of which there are $M-1$, and so there are a total of $(M-1)!$ ways to distribute these operators. Similarly, there are $(M-2)!$ ways to distribute the $M-2$ copies of $z_{M-1} \frac{\partial}{\partial z_{M-1}}$ amongst the $M-2$ linear factors of the form $z_{M-1} - z_i$. Continuing in this manner produces an expression for A:

$$A = (M-1)!(M-2)! \cdots 1!0! \, P(1, \ldots, 1). \tag{13}$$

For an alternate expression of A, we substitute the definition of $D(z_1, \ldots, z_M)$ into (12). Here, we exploit the multilinearity of the determinant and the fact that $(z_n \frac{\partial}{\partial z_n}) z_n^m = m z_n^m$ to get

$$A = \det(m^{n-1})_{m \in \mathcal{M}, 1 \le n \le M} = \prod_{1 \le i < j \le M} (m_j - m_i), \tag{14}$$

where the final equality uses the fact that $(m^{n-1})_{m \in \mathcal{M}, 1 \le n \le M}$ is the transpose of a Vandermonde matrix. Equating (13) to (14) reveals that (11) is an expression for $P(1, \ldots, 1)$. Thus, by assumption, $P(1, \ldots, 1)$ is not a multiple of p, and so we are done. $\quad\square$

Proof of Theorem 9. ($\Leftarrow$) We will use Lemma 11 to demonstrate that F is full spark. To apply this lemma, we need to establish that (11) is not a multiple of p, and to do this, we

will show that there are as many p-divisors in the numerator of (11) as there are in the denominator. We start by counting the p-divisors of the denominator:

$$\prod_{m=0}^{M-1} m! = \prod_{m=1}^{M-1}\prod_{\ell=1}^{m} \ell = \prod_{\ell=1}^{M-1}\prod_{m=1}^{M-l} \ell. \tag{15}$$

For each pair of integers $k, a \geq 1$, there are $\max\{M - ap^k,\ 0\}$ factors in (15) of the form $\ell = ap^k$. By adding these, we count each factor ℓ as many times as it can be expressed as a multiple of a power of p, which equals the number of p-divisors in ℓ. Thus, the number of p-divisors of (15) is

$$\sum_{k=1}^{\lfloor \log_p M \rfloor} \sum_{a=1}^{\lfloor \frac{M}{p^k} \rfloor} (M - ap^k). \tag{16}$$

Next, we count the p-divisors of the numerator of (11). To do this, we use the fact that $\mathcal{M}$ is uniformly distributed over the divisors of N. Since N is a power of p, the only divisors of N are smaller powers of p. Also, the cosets of $\langle p^k \rangle$ partition $\mathcal{M}$ into subsets $S_{k,b} := \{m_i \equiv b \mod p^k\}$. We note that $m_j - m_i$ is a multiple of p^k precisely when m_i and m_j belong to the same subset $S_{k,b}$ for some $0 \leq b < p^k$. To count p-divisors, we again count each factor $m_j - m_i$ as many times as it can be expressed as a multiple of a prime power:

$$\sum_{k=1}^{\lfloor \log_p M \rfloor} \sum_{b=0}^{p^k-1} \binom{|S_{k,b}|}{2}. \tag{17}$$

Write $M = qp^k + r$ with $0 \leq r < p^k$. Then $q = \lfloor \frac{M}{p^k} \rfloor$. Since $\mathcal{M}$ is uniformly distributed over p^k, there are r subsets $S_{k,b}$ with $q+1$ elements and $p^k - r$ subsets with q elements. We use this to get

$$\sum_{b=0}^{p^k-1} \binom{|S_{k,b}|}{2} = \binom{q+1}{2} r + \binom{q}{2}(p^k - r) = \frac{q}{2}\Big((q-1)p^k + 2r + (qp^k - qp^k)\Big).$$

Rearranging and substituting $M = qp^k + r$ then gives

$$\sum_{b=0}^{p^k-1} \binom{|S_{k,b}|}{2} = \frac{q}{2}\Big(2M - (q+1)p^k\Big) = Mq - \binom{q+1}{2}p^k = \sum_{a=1}^{\lfloor \frac{M}{p^k} \rfloor}(M - ap^k).$$

Thus, there are as many p-divisors in the numerator (17) as there are in the denominator (16), and so (11) is not divisible by p. Lemma 11 therefore gives that F is full spark.

($\Rightarrow$) We will prove that this direction holds regardless of whether N is a prime power. Suppose $\mathcal{M} \subseteq \mathbb{Z}_N$ is not uniformly distributed over the divisors of N. Then there exists a divisor d of N such that one of the cosets of $\langle d \rangle$ intersects $\mathcal{M}$ with $\leq \lfloor \frac{M}{d} \rfloor - 1$ or $\geq \lceil \frac{M}{d} \rceil + 1$ indices. Notice that if a coset of $\langle d \rangle$ intersects $\mathcal{M}$ with $\leq \lfloor \frac{M}{d} \rfloor - 1$ indices, then the complement $\mathcal{M}^c$ intersects the same coset with $\geq \lceil \frac{N-M}{d} \rceil + 1 = \lceil \frac{|\mathcal{M}^c|}{d} \rceil + 1$ indices. By Theorem 4(iii), $\mathcal{M}$ produces a full spark harmonic frame precisely when $\mathcal{M}^c$ produces a full spark harmonic frame, and so we may assume without loss of generality that there exists a coset of $\langle d \rangle$ which intersects $\mathcal{M}$ with $\geq \lceil \frac{M}{d} \rceil + 1$ indices.

To prove that the rows with indices in $\mathcal{M}$ are not full spark, we find column entries which produce a singular submatrix. Writing $M = qd + r$ with $0 \leq r < d$, let $\mathcal{K}$ contain $q = \lfloor \frac{M}{d} \rfloor$ cosets of $\langle \frac{N}{d} \rangle$ along with r elements from an additional coset. We claim that the DFT submatrix with row entries $\mathcal{M}$ and column entries $\mathcal{K}$ is singular. To see this, shuffle the rows and columns to form a matrix A in which the row entries are grouped into common cosets of $\langle d \rangle$ and the column entries are grouped into common cosets of $\langle \frac{N}{d} \rangle$. This breaks A into rank-1 submatrices: each pair of cosets $a + \langle d \rangle$ and $b + \langle \frac{N}{d} \rangle$ produces a submatrix

$$\left(\omega^{(a+id)(b+j\frac{N}{d})} \right)_{i \in \mathcal{I}, j \in \mathcal{J}} = \omega^{ab} \left(\omega^{bdi} \omega^{a\frac{N}{d}j} \right)_{i \in \mathcal{I}, j \in \mathcal{J}}$$

for some index sets $\mathcal{I}$ and $\mathcal{J}$; this is a rank-1 outer product. Let $\mathcal{L}$ be the largest intersection between $\mathcal{M}$ and a coset of $\langle d \rangle$. Then $|\mathcal{L}| \geq \lceil \frac{M}{d} \rceil + 1$ is the number of rows in the tallest of these rank-1 submatrices. Define $A_{\mathcal{L}}$ to be the $M \times M$ matrix with entries $A_{\mathcal{L}}[i,j] = A[i,j]$ whenever $i \in \mathcal{L}$ and zero otherwise. Then

$$\mathrm{Rank}(A) = \mathrm{Rank}(A_{\mathcal{L}} + A - A_{\mathcal{L}}) \leq \mathrm{Rank}(A_{\mathcal{L}}) + \mathrm{Rank}(A - A_{\mathcal{L}}). \tag{18}$$

Since $A - A_{\mathcal{L}}$ has $|\mathcal{L}|$ rows of zero entries, we also have

$$\mathrm{Rank}(A - A_{\mathcal{L}}) \leq M - |\mathcal{L}| \leq M - (\lceil \tfrac{M}{d} \rceil + 1). \tag{19}$$

Moreover, since we can decompose $A_{\mathcal{L}}$ into a sum of $\lceil \frac{M}{d} \rceil$ zero-padded rank-1 submatrices, we have $\mathrm{Rank}(A_{\mathcal{L}}) \leq \lceil \frac{M}{d} \rceil$. Combining this with (18) and (19) then gives that $\mathrm{Rank}(A) \leq M - 1$, and so the DFT submatrix is not invertible. $\qquad \square$

Note that our proof of Theorem 9 establishes the necessity of having row indices uniformly distributed over the divisors of N in the general case. This leaves some hope for completely characterizing full spark harmonic frames. Naturally, one might suspect that the uniform distribution condition is sufficient in general, but this suspicion fails when $N = 10$. Indeed, the following DFT submatrix is singular despite the row indices being uniformly distributed over the divisors of 10:

$$\left(e^{-2\pi i m n/10} \right)_{m \in \{0,1,3,4\}, n \in \{0,1,2,6\}}.$$

Also, just as we used Chebotarëv's theorem to analyze the harmonic equiangular tight frames from Xia et al. [38], we can also use Theorem 9 to determine whether harmonic equiangular tight frames with a prime power number of frame elements are full spark. Unfortunately, none of the infinite families in [38] have the number of frame elements in the form of a prime power (other than primes). Luckily, there is at least one instance in which the number of frame elements happens to be a prime power: the harmonic frames that arise from Singer difference sets have $M = \frac{q^d - 1}{q-1}$ and $N = \frac{q^{d+1} - 1}{q-1}$ for a prime power q and an integer $d \geq 2$; when $q = 3$ and $d = 4$, the number of frame elements $N = 11^2$ is a prime power. In this case, the row indices we select are

$$\mathcal{M} = \{1, 2, 3, 6, 7, 9, 11, 18, 20, 21, 25, 27, 33, 34, 38, 41, 44, 47, 53, 54, 55, 56,$$
$$58, 59, 60, 63, 64, 68, 70, 71, 75, 81, 83, 89, 92, 99, 100, 102, 104, 114\},$$

but these are not uniformly distributed over 11, and so the corresponding harmonic frame is not full spark by Theorem 9.

3. Full Spark Frames are dense

In this section, we rigorously show that full spark frames are dense in the entire set of matrices. This corresponds to our intuition that a matrix whose entries are independent continuous random variables is full spark with probability one. Later, we apply the additional constraint of tightness, specifically Parseval tightness, where the frame bound is 1. With this constraint, we will show that full spark Parseval frames are dense in the entire set of Parseval frames.

In order to make our arguments rigorous, we view each entry F_{mn} of the $M \times N$ matrix F in terms of its real and imaginary parts: $x_{mn} + iy_{mn}$; this decomposition will become helpful later when we consider inner products between rows of F, which are not algebraic operations on complex vectors. Recall that F is full spark precisely when each $M \times M$ submatrix has nonzero determinant. We note that for each submatrix, the determinant is a polynomial in the x_{mn}'s and y_{mn}'s, and having this polynomial be nonzero is equivalent to having either its real or imaginary part be nonzero. This naturally leads us to the following definition from algebraic geometry: A *real algebraic variety* is the set of common zeros of a finite set of polynomials, that is, given polynomials $p_1, \ldots, p_r \in \mathbb{R}[x_1, \ldots, x_k]$, we define the corresponding real algebraic variety by

$$V(p_1, \ldots, p_r) := \{x \in \mathbb{R}^k : p_1(x) = \cdots = p_r(x) = 0\}.$$

Each submatrix determinant corresponds to a real algebraic variety $V \subseteq \mathbb{R}^{2MN}$ of two polynomials, and having this determinant be nonzero is equivalent to restricting to the complement of V.

In general, a variety is equipped with a topology known as the *Zariski topology*, in which subvarieties are the closed sets. As such, the set of matrices with a nonzero determinant forms a Zariski-open set, since it is the complement of a variety. In order to exploit this Zariski-openness, we require the additional concept of irreducibility: A variety is said to be *irreducible* if it cannot be written as a finite union of proper subvarieties. As an example, the entire space $\mathbb{R}^k$ is the variety which corresponds to the zero polynomial; in this case, every proper subvariety is lower-dimensional, and so $\mathbb{R}^k$ is trivially irreducible. On the other hand, the variety in $\mathbb{R}^2$ defined by $xy = 0$ is not irreducible because it can be expressed as the union of varieties defined by $x = 0$ and $y = 0$. Irreducibility is important because it says something about Zariski-open sets:

Theorem 12. *If V is an irreducible algebraic variety, then every nonempty Zariski-open subset of V is dense in V in the standard topology.*

For example, the hyperplane defined by $x_1 = 0$ is a subvariety of $\mathbb{R}^k$, and since $\mathbb{R}^k$ is irreducible, the complement of the hyperplane is dense in the standard topology. Going back to the variety defined by $xy = 0$, we know it can be expressed as a union of proper subvarieties, namely the x- and y-axes, and complementing one gives a subset of the other, neither of which is dense in the entire variety.

We are now ready to prove the following result:

Theorem 13. *Every matrix is arbitrarily close to a full spark frame.*

Proof. In an $M \times N$ matrix, there are $\binom{N}{M}$ submatrices of size $M \times M$, and the determinant of each of these submatrices corresponds to a variety of two real polynomials. Since the set of $M \times N$ full spark frames is defined as the (finite) intersection of the complements of these varieties, it is a Zariski-open subset of the irreducible variety $\mathbb{R}^{2MN}$ of all $M \times N$ matrices. Moreover, this set is nonempty since it contains the matrix formed by the first M rows of the $N \times N$ DFT, and so we are done by Theorem 12. $\square$

We now focus on the set of *Parseval frames*, that is, tight frames with frame bound 1. Note that $M \times N$ Parseval frames are characterized by the rows forming an orthonormal system of size M in N-dimensional space. The set of all such orthonormal systems is known as the *Stiefel manifold*, denoted $\mathrm{St}(M, N)$. In general, a *manifold* is a set of vectors with a well-defined tangent space at every point in the set. Since this manifold property is nice, we would like to think of varieties as manifolds, but there exist varieties with points at which tangent spaces are not well-defined; such points are called *singularities*. Note that our definition of singularity is geometric, i.e., where the variety fails to be a real differentiable manifold, as opposed to algebraic. For example, the variety defined by $xy = 0$ is the union of the x- and y-axes, and as such, has a singularity at $x = y = 0$. Certainly with different types of varieties, there are other types of singularities which may arise, but there are also many varieties which do not have singularities at all, and we call these varieties *nonsingular*. Nonsingularity is a useful property for the following reason:

Theorem 14. *An algebraic variety which is nonsingular and connected is necessarily irreducible.*

Note that the orthonormality conditions which characterize Parseval frames F can be expressed as polynomial equations in the real and imaginary parts of the entries F_{mn}. As such, we may view the Stiefel manifold as a real algebraic variety; this variety is nonsingular because it is a manifold. Moreover, the variety is connected because the (connected) unitary group $\mathrm{U}(N)$ acts transitively on $\mathrm{St}(M, N)$. Thus by Theorem 14, the variety of Parseval frames is irreducible. Having established this, we can now prove the following result:

Theorem 15. *Every Parseval frame is arbitrarily close to a full spark Parseval frame.*

Proof. Proceeding as in the proof of Theorem 13, we have that the set of $M \times N$ full spark Parseval frames is Zariski-open in the irreducible variety of all $M \times N$ Parseval frames. Again considering the first M rows of the $N \times N$ DFT, we know this set is nonempty, and so we are done by Theorem 12. $\square$

We note that Theorems 13 and 15 are also true when we further require the frames to be real. In this case, we cannot use the first M rows of the $N \times N$ DFT to establish that the Zariski-open sets are nonempty. For the real version of Theorem 13, we use an $M \times N$ Vandermonde matrix with distinct real bases; see Lemma 2. However, this construction must be modified to use it in the proof of the real version of Theorem 15, since such Vandermonde matrices will not be tight; see Theorem 3. Given an $M \times N$ full spark frame F, the modification $G := (FF^*)^{-1/2}F$ is full spark and Parseval; indeed, $GG^* = I_M$, and since F is a frame, $(FF^*)^{-1/2}$ is full rank, and so the columns of G are linearly independent precisely when the corresponding columns of F are linearly independent.

Another way that the proof of Theorem 15 changes in the real case is in verifying that the real Stiefel manifold is irreducible. Just as in the complex case, this follows from Theorem 14 since $\mathrm{St}(M, N)$ is connected [27], but the fact that $\mathrm{St}(M, N)$ is connected in the real case is not immediate. By analogy, the orthogonal group $\mathrm{O}(N)$ certainly acts transitively on $\mathrm{St}(M, N)$, but unlike the unitary group, the orthogonal group has two connected components. Intuitively, this is resolved by the fact that $N > M$, granting additional freedom of movement throughout $\mathrm{St}(M, N)$.

In addition to Theorems 13 and 15, we would like a similar result for unit norm tight frames, i.e., that every unit norm tight frame is arbitrarily close to a full spark unit norm tight frame. Certainly, the set of unit norm tight frames is a real algebraic variety, but it is unclear whether this variety is irreducible. Without knowing whether the variety is irreducible, we can follow the proofs of Theorems 13 and 15 to conclude that full spark unit norm tight frames are dense in the irreducible components in which they exist—a far cry from the density result we seek. This illustrates a significant gap in our current understanding of the variety of unit norm tight frames. It should be mentioned that Strawn [32] showed that the variety of $M \times N$ unit norm tight frames (over real or complex space) is nonsingular precisely when M and N are relatively prime. Additionally, Dykema and Strawn [12] proved that the variety of $2 \times N$ real unit norm tight frames is connected, and so by Theorem 14, this variety is irreducible when N is odd. It is unknown whether the variety of $M \times N$ unit norm tight frames is connected in general. Finally, we note that a Theorem 15 gives a weaker version of the result we would like: every unit norm tight frame is arbitrarily close to a full spark tight frame with frame element lengths arbitrarily close to 1.

4. The computational complexity of verifying full spark

In the previous section, we demonstrated the abundance of full spark frames, even after imposing the additional condition of tightness. But how would we check whether a particular $M \times N$ frame is full spark? Naturally, we would determine whether each of the $M \times M$ submatrices is invertible; computing $\binom{N}{M}$ determinants would do. However, this is not necessarily the fastest way to go about testing whether a matrix is full spark. For example, perhaps there is an easy-to-calculate expression for the product of the $\binom{N}{M}$ determinants; after all, this product is nonzero precisely when the matrix is full spark. If such an expression exists, then it would suffice to find it, but how might one disprove the existence of any such expression?

At the heart of the matter is computational complexity theory, which provides a rigorous playing field for expressing how hard certain problems are. In this section, we consider the complexity of the following problem:

Problem 16 (Full Spark). *Given a matrix, is it full spark?*

As an example of complexity, intuitively, doubling an integer is no harder than adding integers, since one can use addition to multiply by 2. To make this more precise, complexity theorists use what is called a *polynomial-time reduction*, that is, a polynomial-time algorithm that solves problem A by exploiting an oracle which solves problem B; the reduction indicates that solving problem A is no harder than solving problem B (up to polynomial factors in time), and we say "A reduces to B," or $A \leq B$. Since we can use the polynomial-time routine

$x + x$ to produce $2x$, we conclude that doubling an integer reduces to adding integers, as expected.

In complexity theory, problems are categorized into complexity classes according to the amount of resources required to solve them. For example, the complexity class P contains all problems which can be solved in polynomial time, while problems in EXPTIME may require as much as exponential time. Problems in NP have the defining quality that solutions can be verified in polynomial time given a certificate for the answer. Clearly, P $\subseteq$ NP, since we can ignore the certificate and still solve the problem in polynomial time. Finally, a problem B is called NP-*hard* if every problem A in NP reduces to B, that is, every problem in NP is no harder than B, and a problem is called NP-*complete* if it is both NP-hard and in NP. We note that $A \leq B$ and $B \leq C$ together imply $A \leq C$, and so to demonstrate that a problem C is NP-hard, it suffices to show that $B \leq C$ for some NP-hard problem B. We also note that an NP-hard problem can be solved in polynomial time only if P = NP; this is an open problem, but it is widely believed that P $\neq$ NP.

Sometimes, it can be difficult to find a deterministic reduction from one problem to another. One example is reducing the satisfiability problem (SAT) to the unique satisfiability problem (UNIQUE SAT). To be clear, SAT is an NP-hard problem [20] that asks whether there exists an input for which a given Boolean function returns "true," while UNIQUE SAT asks the same question with an additional promise: that the given Boolean function is satisfiable only if there is a *unique* input for which it returns "true." Intuitively, UNIQUE SAT is easier than SAT because we might be able to exploit the additional structure of uniquely satisfiable Boolean functions; thus, it could be difficult to find a reduction from SAT to UNIQUE SAT. Despite this intuition, there is a *randomized* polynomial-time reduction from SAT to UNIQUE SAT [37]. Defined over all Boolean functions of n variables, the reduction maps functions that are not satisfiable to other functions that are not satisfiable, and with probability $\geq \frac{1}{8n}$, it maps satisfiable functions to uniquely satisfiable functions. After applying this reduction to a given Boolean function, if a UNIQUE SAT oracle declares "uniquely satisfiable," then we know for certain that the original Boolean function was satisfiable. But the reduction will only map a satisfiable problem to a uniquely satisfiable problem with probability $\geq \frac{1}{8n}$, so what good is this reduction? The answer lies in something called *amplification*; since the success probability is, at worst, polynomially small in n (i.e., $\geq \frac{1}{p(n)}$), we can repeat our oracle-based randomized algorithm a polynomial number of times $np(n)$ and achieve an error probability $\leq (1 - \frac{1}{p(n)})^{np(n)} \sim e^{-n}$ which is exponentially small. In this section, we give a randomized polynomial-time reduction from the following problem:

Problem 17. *Given a bipartite graph, what is the girth of its transversal matroid?*

Before giving the reduction, we will show that Problem 17 is NP-hard, but first we briefly review some matroid theory. A *matroid* is a pair $(E, \mathcal{I})$, where E is a finite set and $\mathcal{I}$ is a collection of *independent* subsets of E such that

(i) $\emptyset \in \mathcal{I}$,
(ii) $A \subseteq B \in \mathcal{I}$ implies $A \in \mathcal{I}$,
(iii) $A, B \in \mathcal{I}$ and $|A| > |B|$ implies $\exists x \in A \setminus B$ such that $B \cup \{x\} \in \mathcal{I}$.

We see that matroids provide a natural generalization of the concept of linear independence in vector spaces. Just as spark is the size of the smallest linearly dependent set, the *girth* of a matroid is the size of the smallest subset of E that is not in $\mathcal{I}$. For different kinds of matroids, there are different ways to identify the independent sets. For a *transversal* matroid, the natural way to define $\mathcal{I}$ uses a bipartite graph between two disjoint sets of vertices in which the vertices in one of the disjoint sets are labeled by elements of E; in this case, $\mathcal{I}$ is defined to be all subsets of E whose vertices form the ends of a matching. Hall's marriage theorem [17] gives a remarkable characterization of the independent sets in a transversal matroid: $B \in \mathcal{I}$ if and only if every subset $A \subseteq B$ has $\geq |A|$ neighbors in the bipartite graph. A matroid $(E, \mathcal{I})$ is *representable over a field* $\mathbb{F}$ if, for some M, there exists a mapping $\varphi \colon E \to \mathbb{F}^M$ such that $\varphi(A)$ is linearly independent if and only if $A \in \mathcal{I}$; as such, the girth of $(E, \mathcal{I})$ is the spark of $\varphi(E)$. In our reduction, we make use of the fact that every transversal matroid is representable over $\mathbb{R}$ [25].

The following result comes from McCormick's thesis [22], which credits the proof to Stockmeyer; since [22] is difficult to access and the proof is instructive, we include it below:

Theorem 18. *Problem 17 is* NP-*hard.*

Proof. We will reduce from the NP-complete clique decision problem, which asks "Given a graph, does it contain a clique of K vertices?" [20]. First, we may assume $K \geq 4$ without loss of generality, since any such clique can otherwise be found in cubic time by an exhaustive search. Take a graph $G = (E, V)$, and consider the bipartite graph G' between disjoint sets E and $V \sqcup \{1, \ldots, \binom{K}{2} - K - 1\}$, in which $e \leftrightarrow v$ for every $e \in E$ and $v \in e$, and $e \leftrightarrow k$ for every $e \in E$ and $k \in \{1, \ldots, \binom{K}{2} - K - 1\}$. We claim that the girth of the transversal matroid of G' is $\binom{K}{2}$ precisely when there exists a K-clique in G.

We start by analyzing the girth of a transversal matroid. Consider any dependent set $C \subseteq E$ with $\leq |C| - 2$ neighbors in G'. Then removing any member x of C will produce a smaller set $C \setminus \{x\}$ with $\leq |C| - 2 = |C \setminus \{x\}| - 1$ neighbors in G', which is necessarily dependent by the pigeonhole principle. Now consider any dependent set $C \subseteq E$ with $\geq |C|$ neighbors in G'. By the Hall's marriage theorem, there exists a proper subset $C' \subseteq C$ with $< |C'|$ neighbors in G', meaning C' is a smaller dependent set. Thus, the girth c is the size of the smallest subset $C \subseteq E$ with $|C| - 1$ total neighbors in G'.

Suppose $c = \binom{K}{2}$. Then since C is adjacent to every vertex in $\{1, \ldots, \binom{K}{2} - K - 1\}$, C has $(c - 1) - (\binom{K}{2} - K - 1) = K$ neighbors in V. These are precisely the vertices D in G which are induced by the edges in C, and so C is contained in the set C' of edges induced by D, of which there are $\leq \binom{|D|}{2} = \binom{K}{2}$, with equality only if D induces a K-clique in G. Since $\binom{K}{2} = |C| \leq |C'| \leq \binom{|D|}{2} = \binom{K}{2}$ implies equality, there exists a K-clique in G.

Now suppose there exists a K-clique with edges C. Then C has $\binom{K}{2}$ elements and $K + (\binom{K}{2} - K - 1) = \binom{K}{2} - 1$ neighbors in G'. To prove that $c = \binom{K}{2}$, it suffices to show that there is no smaller subset $C' \subseteq E$ with $|C'| - 1$ total neighbors in G'. Suppose, to the contrary, that $c = \binom{K}{2} - \ell$ for some $\ell > 0$. Then there exists $C' \subseteq E$ with c elements and $(c - 1) - (\binom{K}{2} - K - 1) = K - \ell$ neighbors in V. Note that each $e \in E$ contains two vertices in G, and so by the definition of G', C' necessarily has $K - \ell \geq 2$ neighbors in V. Also, since the $K - \ell$ neighbors of C' in V arise from the subgraph of G induced by C', and since

those neighbors induce at most $\binom{K-\ell}{2}$ edges including C', we have $\binom{K}{2} - \ell = c \leq \binom{K-\ell}{2}$. This inequality simplifies to $\ell \geq 2K - 3$, which combines with $K - \ell \geq 2$ to contradict the fact that $\ell > 0$. $\qquad \square$

Having established that Problem 17 is NP-hard, we reduce from it the main problem of this section. Our proof is specifically geared toward the case where the matrix in question has integer entries; this is stronger than manipulating real (complex) numbers exactly as well as with truncations and tolerances.

Theorem 19. FULL SPARK *is hard for* NP *under randomized polynomial-time reductions.*

Proof. We will give a randomized polynomial-time reduction from Problem 17 to FULL SPARK. As such, suppose we are given a bipartite graph G, in which every edge is between the disjoint sets A and B. Take $M := |B|$ and $N := |A|$. Using this graph, we randomly draw an $M \times N$ matrix F using the following process: for each $i \in B$ and $j \in A$, pick the entry F_{ij} randomly from $\{1, \ldots, N2^{N+1}\}$ if $i \leftrightarrow j$ in G; otherwise set $F_{ij} = 0$. In Proposition 3.11 of [21], it is shown that the columns of F form a representation of the transversal matroid of G with probability $\geq \frac{1}{2}$. For the moment, we assume that F succeeds in representing the matroid.

Since the girth of the original matroid equals the spark of its representation, for each $K = 1, \ldots, M$, we test whether $\mathrm{Spark}(F) > K$. To do this, take H to be some $M \times P$ full spark frame. We will determine an appropriate value for P later, but for simplicity, we can take H to be the Vandermonde matrix formed from bases $\{1, \ldots, P\}$; see Lemma 2. We claim we can randomly select K indices $\mathcal{K} \subseteq \{1, \ldots, P\}$ and test whether $H_{\mathcal{K}}^* F$ is full spark to determine whether $\mathrm{Spark}(F) > K$. Moreover, after performing this test for each $K = 1, \ldots, M$, the probability of incorrectly determining $\mathrm{Spark}(F)$ is $\leq \frac{1}{2}$, provided P is sufficiently large.

We want to test whether $H_{\mathcal{K}}^* F$ is full spark and use the result as a proxy for whether $\mathrm{Spark}(F) > K$. For this to work, we need to have $\mathrm{Rank}(H_{\mathcal{K}}^* F_{\mathcal{K}'}) = K$ precisely when $\mathrm{Rank}(F_{\mathcal{K}'}) = K$ for every $\mathcal{K}' \subseteq \{1, \ldots, N\}$ of size K. To this end, it suffices to have the nullspace $\mathcal{N}(H_{\mathcal{K}}^*)$ of $H_{\mathcal{K}}^*$ intersect trivially with the column space of $F_{\mathcal{K}'}$ for every $\mathcal{K}'$. To be clear, it is always the case that $\mathrm{Rank}(H_{\mathcal{K}}^* F_{\mathcal{K}'}) \leq \mathrm{Rank}(F_{\mathcal{K}'})$, and so $\mathrm{Rank}(F_{\mathcal{K}'}) < K$ implies $\mathrm{Rank}(H_{\mathcal{K}}^* F_{\mathcal{K}'}) < K$. If we further assume that $\mathcal{N}(H_{\mathcal{K}}^*) \cap \mathrm{Span}(F_{\mathcal{K}'}) = \{0\}$, then the converse also holds. To see this, suppose $\mathrm{Rank}(H_{\mathcal{K}}^* F_{\mathcal{K}'}) < K$. Then by the rank-nullity theorem, there is a nontrivial $x \in \mathcal{N}(H_{\mathcal{K}}^* F_{\mathcal{K}'})$. Since $H_{\mathcal{K}}^* F_{\mathcal{K}'} x = 0$, we must have $F_{\mathcal{K}'} x \in \mathcal{N}(H_{\mathcal{K}}^*)$, which in turn implies $x \in \mathcal{N}(F_{\mathcal{K}'})$ since $\mathcal{N}(H_{\mathcal{K}}^*) \cap \mathrm{Span}(F_{\mathcal{K}'}) = \{0\}$ by assumption. Thus, $\mathrm{Rank}(F_{\mathcal{K}'}) < K$ by the rank-nullity theorem.

Now fix $\mathcal{K}' \subseteq \{1, \ldots, N\}$ of size K such that $\mathrm{Rank}(F_{\mathcal{K}'}) = K$. We will show that the vast majority of choices $\mathcal{K} \subseteq \{1, \ldots, P\}$ of size K satisfy $\mathcal{N}(H_{\mathcal{K}}^*) \cap \mathrm{Span}(F_{\mathcal{K}'}) = \{0\}$. To do this, we consider the columns $\{h_k\}_{k \in \mathcal{K}}$ of $H_{\mathcal{K}}$ one at a time, and we make use of the fact that $\mathcal{N}(H_{\mathcal{K}}^*) = \bigcap_{k \in \mathcal{K}} \mathcal{N}(h_k^*)$. In particular, since H is full spark, there are at most $M - K$ columns of H in the orthogonal complement of $\mathrm{Span}(F_{\mathcal{K}'})$, and so there are at least $P - (M - K)$ choices of h_{k_1} for which $\mathcal{N}(h_{k_1}^*)$ does not contain $\mathrm{Span}(F_{\mathcal{K}'})$, i.e.,

$$\dim\Big(\mathcal{N}(h_{k_1}^*) \cap \mathrm{Span}(F_{\mathcal{K}'})\Big) = K - 1.$$

Similarly, after selecting the first J h_k's, we have $\dim(S) = K - J$, where

$$S := \bigcap_{j=1}^{J} \mathcal{N}(h_{k_j}^*) \cap \mathrm{Span}(F_{\mathcal{K}'}).$$

Again, since H is full spark, there are at most $M - (K - J)$ columns of H in the orthogonal complement of S, and so the remaining $P - (M - (K - J))$ columns are candidates for $h_{k_{J+1}}$ that give

$$\dim\left(\bigcap_{j=1}^{J+1} \mathcal{N}(h_{k_j}^*) \cap \mathrm{Span}(F_{\mathcal{K}'}) \right) = \dim\left(\mathcal{N}(h_{k_{J+1}}^*) \cap S \right) = K - (J+1).$$

Overall, if we randomly pick $\mathcal{K} \subseteq \{1, \ldots, P\}$ of size K, then

$$\Pr\left(\mathcal{N}(H_{\mathcal{K}}^*) \cap \mathrm{Span}(F_{\mathcal{K}'}) = \{0\} \right) \geq (1 - \tfrac{M-K}{P})(1 - \tfrac{M-(K-1)}{P}) \cdots (1 - \tfrac{M-1}{P})$$
$$\geq (1 - \tfrac{M}{P})^K$$
$$\geq 1 - \tfrac{MK}{P},$$

where the final step is by Bernoulli's inequality. Taking a union bound over all choices of $\mathcal{K}' \subseteq \{1, \ldots, N\}$ and all values of $K = 1, \ldots, M$ then gives

$$\Pr\left(\begin{array}{c} \text{fail to determine} \\ \mathrm{Spark}(F) \end{array} \right) \leq \sum_{K=1}^{M} \binom{N}{K} \Pr\left(\mathcal{N}(H_{\mathcal{K}}^*) \cap \mathrm{Span}(F_{\mathcal{K}'}) \neq \{0\} \right)$$
$$\leq \sum_{K=1}^{M} \binom{N}{K} \frac{MK}{P}$$
$$\leq \frac{M^3 2^N}{P}.$$

Thus, to make the probability of failure $\leq \tfrac{1}{2}$, it suffices to have $P = M^3 2^{N+1}$.

In summary, we succeed in representing the original matroid with probability $\geq \tfrac{1}{2}$, and then we succeed in determining the spark of its representation with probability $\geq \tfrac{1}{2}$. The probability of overall success is therefore $\geq \tfrac{1}{4}$. Since our success probability is, at worst, polynomially small, we can apply amplification to achieve an exponentially small error probability. $\square$

Our use of random linear projections in the above reduction to FULL SPARK is similar in spirit to Valiant and Vazirani's use of random hash functions in their reduction to UNIQUE SAT [37]. Since their randomized reduction is the canonical example thereof, we find our reduction to be particularly natural.

REFERENCES

[1] Bajwa, W.U., Calderbank, R., Mixon, D.G.: Two are better than one: Fundamental parameters of frame coherence, Appl. Comput. Harmon. Anal. (in press)

[2] Balan, R., Bodmann, B.G., Casazza, P.G., Edidin, D.: Fast algorithms for signal reconstruction without phase, Proc. SPIE, 67011L, 1–9 (2007)

[3] Balan, R., Bodmann, B.G., Casazza, P.G., Edidin, D.: Painless reconstruction from magnitudes of frame coefficients, J. Fourier Anal. Appl. **15**, 488–501 (2009)

[4] Balan, R., Casazza, P., Edidin, D.: On signal reconstruction without phase, Appl. Comput. Harmon. Anal. **20**, 345–356 (2006)

[5] Cahill, J., Casazza, P.G., Heinecke, A.: A notion of redundancy for infinite frames, Proc. SampTA (2011). Available online at http://www.framerc.org/papers/ANotionOfRedundancyForInfiniteFrames.pdf

[6] Candès, E.J., Strohmer, T., Voroninski, V.: PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming. Available online: arXiv:1109.4499

[7] Candès, E.J., Tao, T.: Decoding by linear programming, IEEE Trans. Inform. Theory **51** 4203–4215 (2005)

[8] Casazza, P.G., Heinecke, A., Krahmer, F., Kutyniok, G.: Optimally sparse frames, IEEE Trans. Inform. Theory (in press)

[9] Davenport, M.A., Duarte, M.F., Eldar, Y.C., Kutyniok, G.: Introduction to Compressed Sensing. In: Eldar, Y.C., Kutyniok, G. (Eds.), Compressed Sensing: Theory and Applications, Cambridge University Press (2011)

[10] Delvaux, S., Van Barel, M.: Rank-deficient submatrices of Fourier matrices, Linear Algebra Appl. **429**, 1587–1605 (2008)

[11] Donoho, D.L., Elad, M.: Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization, Proc. Nat. Acad. Sci. **100**, 2197–2202 (2003)

[12] Dykema, K., Strawn, N.: Manifold structure of spaces of spherical tight frames, Int. J. Pure Appl. Math. **28**, 217–256 (2006)

[13] Evans, R.J., Isaacs, I.M.: Generalized Vandermonde determinants and roots of prime order, Proc. Amer. Math. Soc. **58**, 51–54 (1977)

[14] Fickus, M., Mixon, D.G.: Deterministic matrices with the restricted isometry property, Proc. SPIE (2011)

[15] Fickus, M., Mixon, D.G., Tremain, J.C.: Steiner equiangular tight frames, Linear Algebra Appl. (in press)

[16] Goyal, V.K.: Beyond Traditional Transform Coding. Ph.D. Thesis, University California, Berkeley (1998)

[17] Hall, P.: On Representatives of Subsets, J. London Math. Soc. **10**, 26–30 (1935)

[18] Holmes, R.B., Paulsen, V.I.: Optimal frames for erasures, Linear Algebra Appl. **377**, 31–51 (2004)

[19] Jungnickel, D., Pott, A., Smith, K.W.: Difference Sets. In: Colbourn, C.J., Dinitz, J.H. (Eds.), Handbook of Combinatorial Designs, 2nd ed., pp. 419–435 (2007)

[20] Karp, R.M.: Reducibility Among Combinatorial Problems. In: Miller, R.E., Thatcher, J.W. (Eds.), Complexity of Computer Computations, pp. 85–103. Plenum, New York (1972)

[21] Marx, D.: A parameterized view on matroid optimization problems, Theor. Comput. Sci. **410**, 4471–4479 (2009)

[22] McCormick, S.T.: A Combinatorial Approach to Some Sparse Matrix Problems. Ph.D. Thesis, Stanford University (1983)

[23] Mixon, D.G., Quinn, C., Kiyavash, N., Fickus, M.: Equiangular tight frame fingerprinting codes, Proc. IEEE Int. Conf. Acoust. Speech Signal Process., 1856–1859 (2011)

[24] Nakamura, S., Masson, G.M.: Lower bounds on crosspoints in concentrators, IEEE Trans. Comput. **C-31**, 1173–1179 (1982)

[25] Piff, M.J., Welsh, D.J.A.: On the vector representation of matroids, J. London Math. Soc. **2**, 284–288 (1970)

[26] Püschel, M., Kovačević, J.: Real, tight frames with maximal robustness to erasures, Proc. Data Compr. Conf., 63–72 (2005)

[27] Rapcsák, T.: On minimization on Stiefel manifolds, Eur. J. Oper. Res. **143**, 365–376 (2002)

[28] Renes, J.: Equiangular tight frames from Paley tournaments, Linear Algebra Appl. **426**, 497–501 (2007)

[29] Renes, J.M., Blume-Kohout, R., Scott, A.J., Caves, C.M.: Symmetric informationally complete quantum measurements, J. Math. Phys. **45**, 2171–2180 (2004)

[30] Rudelson, M., Vershynin, R.: On sparse reconstruction from Fourier and Gaussian measurements, Commun. Pure Appl. Anal. **61**, 1025–1045 (2008)

[31] Stevenhagen, P., Lenstra, H.W.: Chebotarëv and his density theorem, Math. Intelligencer **18**, 26–37 (1996)

[32] Strawn, N.: Finite frame varieties: Nonsingular points, tangent spaces, and explicit local parameterizations, J. Fourier Anal. Appl. **17**, 821–853 (2011)

[33] Strohmer, T.: A note on equiangular tight frames, Linear Algebra Appl. **429**, 326–330 (2008)

[34] Strohmer, T., Heath, R.W.: Grassmannian frames with applications to coding and communication, Appl. Comput. Harmon. Anal. **14**, 257–275 (2003)

[35] Tao, T.: An uncertainty principle for cyclic groups of prime order, Math. Research Letters **12**, 121–128 (2005)

[36] Tropp, J.A.: On the conditioning of random subdictionaries, Appl. Comput. Harmon. Anal. **25**, 1–24 (2008)

[37] Valiant, L., Vazirani, V.: NP is as easy as detecting unique solutions, Theor. Comput. Sci. **47**, 85–93 (1986)

[38] Xia, P., Zhou, S., Giannakis, G.B.: Achieving the Welch bound with difference sets, IEEE Trans. Inform. Theory **51**, 1900–1907 (2005)

[39] Zauner, G.: Quantendesigns: Grundzüge einer nichtkommutativen Designtheorie. Ph.D. thesis, University of Vienna (1999)

(Alexeev) Department of Mathematics, Princeton University, Princeton, New Jersey 08544, USA; E-mail: balexeev@math.princeton.edu

(Cahill) Department of Mathematics, University of Missouri, Columbia, Missouri 65211, USA; E-mail: jameson.cahill@gmail.com

(Mixon) Program in Applied and Computational Mathematics, Princeton University, Princeton, New Jersey 08544, USA; E-mail: dmixon@princeton.edu