

DESIGN AND ANALYSIS OF MULTI-COSET ARRAYS

James D. Krieger^{*}, Yuval Kochman[†], and Gregory W. Wornell^{*}

^{*} Dept. EECS, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

[†] School of CSE, Hebrew University of Jerusalem, Jerusalem, Israel

^{*} {jameskrieger, gww}@mit.edu, [†] yuvalko@cs.huji.ac.il

ABSTRACT

An efficient sparse antenna array architecture is developed for coherent imaging of sparse but otherwise unknown scenes. In this architecture, the array elements form a periodic nonuniform pattern. Using analysis that explicitly takes into account the presence of noise, we develop an efficient pattern design procedure based on co-arrays, describe an efficient scene support recovery algorithm as part of image reconstruction in the form of a modification to the MUSIC algorithm, and discuss a failure detection technique based on evaluating “back-projection” error. Since our development exploits a close connection to multi-coset sampling of bandlimited waveforms, our results may in turn may also be useful in the design of those systems.

Index Terms— antenna arrays, sparse sampling, compressed sensing, MUSIC algorithm, millimeter-wave imaging

1. INTRODUCTION

CMOS circuits now operating at terahertz frequencies are enabling host of low-cost millimeter wave imaging applications. This, in turn, is creating important opportunities and new signal processing challenges in sensor and array processing.

One significant challenge is the number of array elements typically required. In a vehicle collision avoidance system, for example, obtaining sufficient resolution might require an aperture of roughly 2m. But in this case a traditional phased array operating at 100 GHz with half-wavelength element spacing would require roughly 1000 antennas, which is daunting to system designers.

In this work, we focus on significantly reducing the number of antenna elements required for a given aperture. There is a long history of such development in both the antenna and array processing communities. General-purpose approaches have tended to lead to limited gains; see, e.g., [1] and references therein. However, in some cases, large reductions are possible.

Consider, for example, the classical problem of direction-finding with multiple sources, to which, e.g., the MUSIC algorithm [2] is often applied. In this case, it is possible to achieve high-resolution with relatively few antenna elements because of the simple structure of scene. Indeed, the number of elements required is typically on the order of the number of sources. Hence, the presence of structure in the environment allows the number of elements to be reduced.

This general insight can be exploited in a variety of ways in emerging systems. In particular, when the scene being imaged is sparse in an appropriate sense—even without knowing *where* it is

sparse—then it is possible to commensurately reduce the number of elements in an imaging array. Moreover, such sparseness is quite common in typical applications.¹

For antennas of just a few elements, the array design and image reconstruction can often be fairly straightforward and exploit classical techniques. However, for arrays of even a few dozen elements, such direct approaches quickly become computationally infeasible to design, and impractical to implement. As a result, there is a need to impose useful structure on the array, to enable efficient design and processing; in particular, the class of structured nonuniform arrays has received increased attention in recent work within the frame of sparse sensing (see, e.g., [3]).

As is well known in the community, there is a close mathematical relationship between the problem of imaging from a discrete array, and that of reconstructing a bandlimited time-domain waveform from samples. We exploit this connection in our development. In particular, there has been much work on the reconstruction of waveforms whose frequency content is limited to subbands of the full bandwidth from samples taken at the Landau rate. The Landau rate is proportional to the actual spectral support and thus often significantly lower than the classical Shannon-Nyquist rate determined by the full bandwidth. Moreover, in such work, the location of the frequency bands often not be known a priori.

In our development, we focus on the antenna-array counterpart of a particular structured approach based on multi-coset sampling [4], which uses a quasi-periodic (recurrent nonuniform) sampling pattern. In [5], we described the adaptation of multi-coset sampling to imaging arrays, and emphasized the finite-window effects that play a more prominent role in the latter. In this paper, we more fully develop the system design, implementation, and signal processing, explicitly taking noise into account in our analysis. Our analysis yields: 1) a novel class of effective multi-coset patterns; 2) an efficient support recovery method based on a simple modification to the MUSIC algorithm; and 3) a methodology for detecting reconstruction failures when the scene density exceeds the level for which an array is designed.

2. MULTI-COSET IMAGING

We begin by summarizing those aspects of multi-coset array image reconstruction required for our development. We focus on the regime where the number of array elements N is large, and defer noise considerations. For further details, see [5].

This work was supported in part by NSF under Grant No. CCF-1017772, by SRC FCRP C2S2, and by USAF under Contract No. FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the authors and are not necessarily endorsed by the US Government.

¹Note that in a typical scene while there are objects at some range in any particular direction, when we use enough bandwidth to sufficiently resolve range as well, we find significant sparseness in the range-azimuth plane.

2.1. Scene Model, Array Structure, and Notation

Let $x[n]$ be the complex-valued response at array element n of a half-wavelength spaced uniform array. Its Fourier transform $X(\psi)$ is the value of the corresponding far-field image in the direction θ , where $\psi = \sin \theta/2$.

Of interest are far-field scenes that are sparse. For any pair of integers $Q \leq L$, we say that a scene is (Q, L) -sparse if $X(\psi) = 0$ for all $\psi \notin \Psi_\Omega$, where

$$\Psi_\Omega = \bigcup_{k=0}^{Q-1} \left[\frac{q_k}{L}, \frac{q_k + 1}{L} \right), \quad (1)$$

where $\Omega = \{q_0, q_1, \dots, q_{Q-1}\}$ is referred to as the (*block*) *support* of the scene, with $q_k \in \mathbb{Z}$, for all k , and $0 \leq q_0 < q_1 < \dots < q_{Q-1} \leq L - 1$.

For sparse scenes it is possible to reconstruct $x[n]$ from a subset of the array elements. One such class of subsets, is multi-coset arrays. We view such array as being composed of L overlapping uniformly spaced subarrays, termed cosets. The coset response is then

$$x^{(p)}[n] = x[n] \sum_m \delta[n - (mL + p)], \quad 0 \leq p \leq L - 1. \quad (2)$$

For any pair of integers $P \leq L$, a (P, L) multi-coset array is the union of some choice of P of these cosets. The particular choice of cosets $\mathcal{P} = \{p_0, p_1, \dots, p_{P-1}\}$ with $0 \leq p_0 < p_1 < \dots < p_{P-1} \leq L - 1$ is referred to as the *coset pattern* of the array.

2.2. Image Reconstruction Given Known Support

The image formed from a single coset response $X^{(p)}(\psi)$ is the Fourier transform of (2). Due to the excess spacing between the coset elements, the coset image suffers from aliasing (grating lobes), appearing as L identical sectors. As a result, we may ignore all but one of the sectors. Following [4, 5], we define

$$\begin{aligned} Y_p(\psi) &\triangleq X^{(p)}(\psi)H(\psi), \\ X_q(\psi) &\triangleq X(\psi + q/L)H(\psi), \end{aligned} \quad H(\psi) \triangleq \begin{cases} 1 & \psi \in [0, 1/L) \\ 0 & \text{otherwise,} \end{cases}$$

whence

$$Y_p(\psi) = \sum_{q=0}^{L-1} A_{pq} X_q(\psi) \quad \text{with} \quad A_{pq} \triangleq \frac{1}{L} e^{j2\pi pq/L}. \quad (3)$$

When the scene is (Q, L) -sparse as in (1), the summation need only be taken over $q \in \Omega$. The system of equations relating the coset images for \mathcal{P} to the original image in the support Ω may be written in the form $\mathbf{Y}_{\mathcal{P}}(\psi) = \mathbf{A}_{\mathcal{P}\Omega} \mathbf{X}_\Omega(\psi)$. When $\mathbf{A}_{\mathcal{P}\Omega}$ has full rank, we can reconstruct $\mathbf{X}_\Omega(\psi)$ via $\mathbf{X}_\Omega(\psi) = \mathbf{A}_{\mathcal{P}\Omega}^+ \mathbf{Y}_{\mathcal{P}}(\psi)$, where $\mathbf{A}_{\mathcal{P}\Omega}^+ = (\mathbf{A}_{\mathcal{P}\Omega} \mathbf{A}_{\mathcal{P}\Omega}^H)^{-1} \mathbf{A}_{\mathcal{P}\Omega}^H$ is the pseudoinverse of the measurement matrix. A pattern \mathcal{P} which guarantees full rank for any support Ω of size up to Q is called a *universal pattern*. It is shown in [4] that such patterns exist as long as $P \geq Q$.

2.3. Blind Support Recovery

When the support Ω is not known a priori, as is the case of interest in this work, an estimate $\hat{\Omega}$ is obtained from the available measurements; this is referred to as *blind support recovery*. In this case, we rewrite (3) as

$$\mathbf{Y}_{\mathcal{P}}(\psi) = \mathbf{A}_{\mathcal{P}} \mathbf{X}(\psi), \quad (4)$$

where the omission of the subscript Ω implies the inclusion of all sectors in the expression. Eq. (4) is an underdetermined system, but since $\mathbf{X}(\psi)$ has at most Q nonzero entries, it can be solved as a *compressed sensing (CS)* problem. Furthermore, since the location of these nonzero entries is the same for all $\psi \in [0, 1/L)$, this is an instance of CS with so-called *multiple measurement vectors (MMV)*. In terms of the scene correlation matrix

$$\mathbf{S}_{\mathbf{X}} = \int_0^{1/L} \mathbf{X}(\psi) \mathbf{X}^H(\psi) d\psi, \quad (5)$$

support recovery is guaranteed whenever [6, 7]

$$P \geq 2Q - \text{rank}(\mathbf{S}_{\mathbf{X}}) + 1. \quad (6)$$

When, in addition, $\{X_q(\psi)\}$ form a linearly independent set, (6) becomes $P \geq Q + 1$, which is much more favorable than the worst-case $2Q$. While such linear independence holds for most scenes of interest, this may not be the case for finite arrays (where the effective number of measurements is reduced).

3. SUPPORT RECOVERY ALGORITHMS

One approach to support recovery is based on ℓ_1 -minimization. With this approach, the coset correlation matrix

$$\mathbf{S}_{\mathbf{Y}_{\mathcal{P}}} = \int_0^{1/L} \mathbf{Y}_{\mathcal{P}}(\psi) \mathbf{Y}_{\mathcal{P}}^H(\psi) d\psi = \mathbf{A}_{\mathcal{P}} \mathbf{S}_{\mathbf{X}} \mathbf{A}_{\mathcal{P}}^H \quad (7)$$

is expressed as $\mathbf{S}_{\mathbf{Y}_{\mathcal{P}}} = \mathbf{V} \mathbf{V}^H$, where $\mathbf{V} = \mathbf{A}_{\mathcal{P}} \mathbf{W}$. Since $\mathbf{A}_{\mathcal{P}}$ is $P \times L$ with $P \leq L$, the matrix \mathbf{W} is not unique. While we seek the solution \mathbf{W}_0 that minimizes the number of rows having nonzero entries, CS replaces this ℓ_0 -minimization with a computationally more convenient ℓ_1 -minimization. Specifically, defining the length- L vector \mathbf{w} with entries equal to the ℓ_2 -norm of the corresponding rows of \mathbf{W} , we solve

$$\mathbf{W}_1 \triangleq \arg \min_{\mathbf{w}} \|\mathbf{w}\|_1 \quad \text{s.t.} \quad \mathbf{V} - \mathbf{A}_{\mathcal{P}} \mathbf{W} = \mathbf{0}.$$

A computationally simpler alternative to the ℓ_1 -minimization, originally proposed in [4], is based on the MUSIC algorithm. Specifically, in the absence of noise, the correlation matrix (7) has Q nonzero eigenvalues, so we partition the eigenvector matrix as $\mathbf{U} = [\mathbf{U}_S \mathbf{U}_N]$, where the $P \times (P - Q)$ matrix \mathbf{U}_N containing the eigenvectors corresponding to the zero-valued eigenvalues represents the noise subspace of \mathbf{U} . Each of the L columns of $\mathbf{A}_{\mathcal{P}}$ is projected onto the noise subspace. Since the columns corresponding to the active sectors contained within the support Ω lie in the subspace orthogonal to \mathbf{U}_N , their projections will be zero, and thus we recover the support $\hat{\Omega}$ as the indices of these columns.

3.1. Recovery in the Presence of Noise

Additional care is required when noise is present. In the ℓ_1 minimization approach, the constraint must be relaxed in order to obtain a sparse solution, while with the MUSIC-based approach, difficulties arise due to the finite array lengths that must be employed in practice— \mathbf{U}_S and \mathbf{U}_N no longer accurately partition the signal and noise subspaces, and the eigenvalues associated with \mathbf{U}_N take on nonzero values.

The situation is complicated further when the support size Q is unknown. In these cases, an intermediate step must be included to estimate the signal subspace dimension \hat{Q} . In the case of MUSIC,

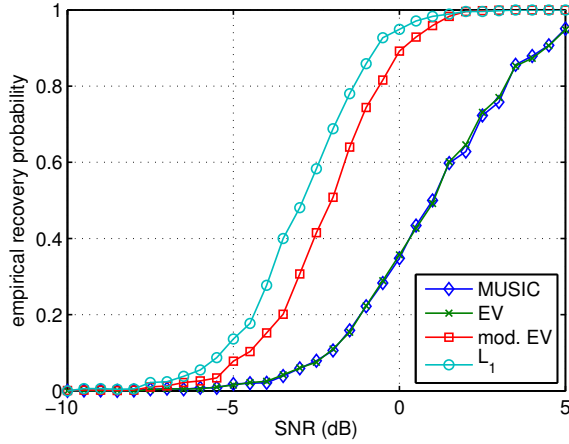


Fig. 1. Comparative performance of the support recovery algorithms in noise for a scene with $Q = 7$ active sectors using an array with coset period $L = 19$, and $P = 9$ active cosets. The coset pattern is $\mathcal{P} = \{0, 1, 2, 3, 5, 7, 12, 13, 16\}$.

this may be accomplished through direct thresholding of the eigenvalues; more sophisticated approaches are described in [8]. However, when $\hat{Q} > Q$, two issues arise: 1) the estimate of the dimensionality of the noise subspace $P - \hat{Q}$ is reduced making the recovery estimate less accurate; and 2) the matrix $\mathbf{A}_{\mathcal{P}\hat{Q}}$ used in the reconstruction is more poorly conditioned, resulting in increased noise amplification.

A weighted version of the MUSIC algorithm based on the eigenvalue (EV) method [9] accounts for errors due to finite sample sets by weighing the projections onto each subspace direction more heavily for smaller eigenvalues via the *null spectrum*

$$D_{\text{EV}}(q) = \left(\sum_{m=Q+1}^P \lambda_m^{-1} |\mathbf{a}_q^\dagger \mathbf{u}_m|^2 \right)^{-1}. \quad (8)$$

However, this continues to require an estimate of Q . To obviate the need such an estimate, we further modify the EV method, weighing the entire column space of \mathbf{U} via

$$\tilde{D}_{\text{EV}}(q) = \left(\sum_{m=1}^P \left(\lambda_m^{-1/2} - \lambda_1^{-1/2} \right)^2 |\mathbf{a}_q^\dagger \mathbf{u}_m|^2 \right)^{-1}. \quad (9)$$

Using numerical simulation, we obtain that the modified EV retains the performance of the ℓ_1 -minimization approach while preserving the computational complexity of the MUSIC approach, as depicted in the representative results of Fig. 1. In the simulations, the recovery algorithms were applied randomly generated sparse scenes, with white Gaussian noise added to the measurements.

4. ARRAY DESIGN

We now describe a design procedure for the multi-coset array.

4.1. Array Parameters

First, the choice of the coset period L involves a compromise. Large values of L leads to a finer partitioning of scene support via (1),

allowing sparsity to be exploited more completely. However, small values of L increase the number of periods in the array, which keeps side lobes small [5], and improves the rank of (5), facilitating support recovery.² In practice, it is generally best to choose the smallest L for which the block density has largely converged to its value in the $L \rightarrow \infty$ limit.

For the number of active cosets P , choosing $P = Q + 1$ will be sufficient for most scenes, though in low SNR settings, support recovery is generally much better with at least some additional margin; e.g., at least $P = Q + 2$.

4.2. Coset Pattern Selection

Although there are many universal coset patterns that guarantee scene recovery in the absence of noise (such as the “bunched” pattern $\mathcal{P}_B = \{0, 1, \dots, P-1\}$), universal patterns (like \mathcal{P}_B) for which $\mathbf{A}_{\mathcal{P}Q}$ is ill-conditioned will not be good when noise is present. And while coset patterns with low condition numbers $\kappa(\mathbf{A}_{\mathcal{P}Q})$ for all Q will ensure good reconstruction SNR (RSNR), a brute force search for such patterns is impractical for all but the smallest values of L .

Instead, we develop a computationally much simpler approach inspired by the “minimum redundancy linear array (MRLA)” designs introduced in [10]. MRLAs minimize the number of sensor pairs having identical spacings in an attempt to obtain the best representation of the full correlation matrix with the fewest elements. Our approach can be viewed as a key modification to this methodology in which we take into account the periodic structure the multi-coset array. For this purpose, we define a distance between two cosets as the minimum distance between the corresponding coset elements within either the same, or the neighboring, coset periods; specifically, this distance is $t_{lk} = \min\{|l - k|, L - |l - k|\}$.

To develop the methodology, consider the correlation matrix associated with all L cosets, i.e., $\mathbf{S}_Y = \mathbf{A} \mathbf{S}_X \mathbf{A}^\dagger$. Moreover, for the moment, consider the model in which the Q functions $\{X_q(\psi)\}$ in (5) comprise a linearly independent set such that $\{\mathbf{S}_X\}_{m,n} = \sigma_m^2 \delta_{mn}$, where $\sigma_m^2 = \{\mathbf{S}_X\}_{m,m}$ is the signal energy from block m . In this case, \mathbf{S}_Y will have Hermitian-circulant structure, i.e.,

$$\{\mathbf{S}_Y\}_{lk} = \frac{1}{L^2} \sum_{q=0}^{L-1} \sigma_q^2 e^{j2\pi(l-k)q/L}. \quad (10)$$

In (10), the dependence of the matrix entries on the relative spacing between elements indicates the importance of the pairwise spacings as represented by the $2(L-1)$ off-diagonals. Specifically, the information contained in \mathbf{S}_Y can be obtained by representing each of the possible spacings a single time. The symmetries in the Hermitian-circulant structure reduce the number of parameters by a further factor of four, suggesting the entire matrix could be represented by only $\lceil (L-1)/2 \rceil$ parameters. In reality, there will be measurement noise and some correlation between the different sectors, and thus \mathbf{S}_Y will vary to some extent along each diagonal. As such, multiple occurrences of a particular pairwise spacing can be interpreted as multiple samples of noisy data, suggesting that evenly distributed spacings are desirable.

Such designs can be found quite efficiently. In particular, to determine the number of occurrences of each pairwise spacing in a given pattern \mathcal{P} , we let $s_k = 1$ if $k \in \mathcal{P}$ (and zero otherwise) and

²The number of periods plays a role similar to the number of “snapshots” in traditional direction-finding problems.

L	P	\mathcal{P}^*	c^*	γ_O	γ_B
7	3	{013}	1	-2.0	4.0
7	4	{0124}	2	0	4.4
11	5	{01247}	2	0	12.3
11	6	{012457}	3	0	11.8
13	4	{0139}	1	-1.5	13.7
13	9	{0123457910}	6	0	14.2
19	9	{012357121316}	4	0	37.9
19	10	{01235712131516}	5	0	38.4

Table 1. Some modified co-array coset patterns \mathcal{P}^* .

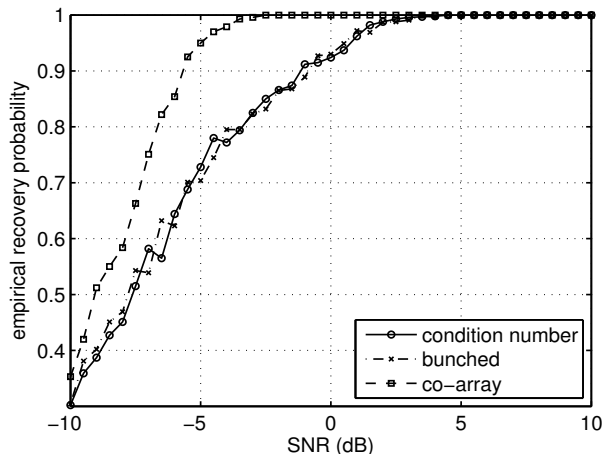


Fig. 2. Support recovery performance for different coset patterns, with $L = 7$, $P = 3$, and $Q = 2$.

compute

$$c(t) = \sum_{\{(l,k): t_{lk}=t\}} s_l s_k, \quad 1 \leq t \leq \frac{(L-1)}{2}, \quad (11)$$

which is a modification of the co-array in MRLA design. Since the total number of spacings is identical for all patterns of length P , we select the pattern having the co-array with smallest ℓ_2 -norm.

Examples of the modified co-array coset patterns \mathcal{P}^* are provided in Table 1 for several choices of (prime) L and P . Our choices of L and P in this table have the special quality that the number of unique element pairs $P(P-1)/2$ is an integer multiple of the number of possible spacings $(L-1)/2$, so that for all spacings t we have $c(t) = c^*$ for some c^* . For reference the loss γ_O in condition number (in dB) relative to the minimum possible is shown, along with the gain γ_B (in dB) relative to the “bunched” coset pattern.

To evaluate their effectiveness in support recovery, the modified co-array designs were evaluated using numerical simulation. Remarkably, the modified co-array designs were consistently superior to the minimum condition number designs (when they differed), as well as to all other designs to which we compared them. Fig. 2 illustrates a representative scenario.

5. FAILURE DETECTION

In applications, it is often critical to be able to detect when the scene density exceeds the level for which a given multi-coset array is designed, to prevent false images from being produced. In

principle, this capability is straightforward to incorporate into the multi-coset architecture based on evaluating what can be viewed as a back-projection error (BPE).

Consider a (P, L) multi-coset array with coset pattern \mathcal{P} and a (Q, L) -sparse scene, where Q is unknown. Using a support estimate \hat{Q} obtained from the measurements $\mathbf{Y}_{\mathcal{P}}(\psi)$, the image estimate is $\hat{\mathbf{X}}_{\hat{Q}}(\psi) = \mathbf{A}_{\hat{Q}}^+ \mathbf{Y}_{\mathcal{P}}(\psi)$. If \hat{Q} is correct,

$$\hat{\mathbf{Y}}_{\mathcal{P}\hat{Q}}(\psi) \triangleq \mathbf{A}_{\mathcal{P}\hat{Q}} \hat{\mathbf{X}}_{\hat{Q}}(\psi) = \mathbf{A}_{\mathcal{P}\hat{Q}} \mathbf{A}_{\hat{Q}}^+ \mathbf{Y}_{\mathcal{P}}(\psi) \approx \mathbf{Y}_{\mathcal{P}}(\psi) \quad (12)$$

where the degree to which approximation holds depends on the SNR. Since $\mathbf{A}_{\mathcal{P}\hat{Q}} \mathbf{A}_{\hat{Q}}^+$ is the projection matrix onto the range of $\mathbf{A}_{\mathcal{P}\hat{Q}}$, (12) projects the available data $\mathbf{Y}_{\mathcal{P}}(\psi)$ back onto the subspace from which it was estimated to have come. As such, the mean-square back-projection error

$$\sigma_{\text{BPE}}^2 = \int_0^1 \|\mathbf{Y}_{\mathcal{P}}(\psi) - \hat{\mathbf{Y}}_{\mathcal{P}\hat{Q}}(\psi)\|_2^2 d\psi \quad (13)$$

is a measure of the integrity of the reconstruction.

As discussed in Section 2, at sufficiently high SNR, a multi-coset array with a (P, L) -universal pattern is able to recover the support \mathcal{Q} of a (Q, L) -sparse scene in most cases provided the scene is sufficiently sparse, i.e., $Q \geq P - 1$. In such cases, since the support estimate is (or contains) the correct support ($\mathcal{Q} \subseteq \hat{\mathcal{Q}}$), the back-projection error is zero. However, when the scene is too dense, the recovery stage will fail to fully determine the support ($\hat{\mathcal{Q}} \subset \mathcal{Q}$). In this case, much of the energy contained in the unidentified support blocks $\mathcal{Q} \setminus \hat{\mathcal{Q}}$ will vanish during back-projection, allowing the failure to be detected via (13). In practice, the BPE threshold at which we declare a failure depends on SNR; we omit the details due to space constraints.

6. REFERENCES

- [1] Y.T. Lo and S.W. Lee, *Antenna Handbook: Antenna Theory*, Van Nostrand Reinhold, New York, NY, 1993.
- [2] R. O. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE Trans. Antennas Propag.*, vol. 34, no. 3, pp. 276–280, Mar. 1986.
- [3] P.P. Vaidyanathan and P. Pal, “Sparse sensing with co-prime samplers and arrays,” *Signal Processing, IEEE Transactions on*, vol. 59, no. 2, pp. 573–586, Feb.
- [4] P. Feng and Y. Bresler, “Spectrum-blind minimum-rate sampling and reconstruction of multiband signals,” in *Proc. ICASSP*, Atlanta, GA, 1996, vol. 3, pp. 1688–1691.
- [5] Y. Kochman and G. W. Wornell, “Finite multi-coset sampling and sparse arrays,” in *Proc. ITA*, La Jolla, CA, 2011, pp. 1–7.
- [6] Y. M. Lu and M. N. Do, “A theory for sampling signals from a union of subspaces,” *IEEE Trans. Signal Process.*, vol. 56, no. 6, pp. 2334–2345, 2008.
- [7] M. Mishali and Y. C. Eldar, “Blind multiband signal reconstruction: Compressed sensing for analog signals,” *IEEE Trans. Signal Process.*, vol. 57, no. 3, pp. 993–1009, 2009.
- [8] H. L. Van Trees, *Optimum Array Processing (Detection, Estimation, and Modulation Theory, Part IV)*, Wiley-Interscience, New York, NY, 2002.

- [9] D. Johnson and S. DeGraaf, "Improving the resolution of bearing in passive sonar arrays by eigenvalue analysis," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 30, no. 4, pp. 638–647, 1982.
- [10] A. Moffet, "Minimum-redundancy linear arrays," *IEEE Trans. Antennas Propag.*, vol. 16, no. 2, pp. 172–175, 1968.