

Compressive Adaptation of Large Steerable Arrays

Dinesh Ramasamy, Sriram Venkateswaran, Upamanyu Madhow
Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106, USA

Abstract—We consider the problem of adapting very large antenna arrays (e.g., with 1000 elements or more) for tasks such as beamforming and nulling, motivated by emerging applications at very high carrier frequencies in the millimeter (mm) wave band and beyond, where the small wavelengths make it possible to pack a very large number of antenna elements (e.g., realized as a printed circuit array) into nodes with compact form factors. Conventional least squares techniques, which rely on access to baseband signals for individual array elements, do not apply. Hence the preferred approach is to perform radio frequency (RF) beamsteering, with a single complex baseband signal emerging from a receive array, or going into a transmit array. Further, we are interested in what can be achieved with coarse-grained control of individual elements (e.g., four-phase, or even binary phase, control). In this paper, we propose an adaptation architecture matched to these hardware constraints. Our approach comprises the following two steps. The first step is *compressive estimation* of a sparse spatial channel using a small number of measurements, each using a different set of randomized weights. However, unlike the standard compressive sensing formulation, we are interested in estimating continuous-valued parameters such as the angles of arrivals of various paths. The second step is *quantized beamsteering*, where weights for beamforming and nulling, subject to the constraint of severe quantization, are computed using the channel estimates from the first step. We provide promising preliminary results illustrating the efficacy of this approach.

I. INTRODUCTION

We begin with the following question: how does one effectively adapt a very large array (e.g., 1000 elements or more) for tasks such as beamforming and nulling, while accounting for natural hardware constraints? The motivating application is communication using very high carrier frequencies in the millimeter (mm) wave band and beyond, where the small wavelengths make it possible to pack a very large number of antenna elements (e.g., realized as a printed circuit array) into nodes with compact form factors. Using a separate RF chain for each antenna element is out of the question in such settings, hence it is not possible to employ standard least squares style adaptation, which requires access to the complex baseband signal corresponding to each antenna element. Thus, we limit ourselves to RF beamforming: at the transmitter, a single complex baseband waveform is upconverted to RF and distributed to the antenna elements, with digital control of the amplitude and phase of each element; at the receiver, the RF signals at different elements are combined after digitally controlling the

This work was supported by the National Science Foundation through the grant CNS-0520335, and by the Institute for Collaborative Biotechnologies through the grant W911NF-09-0001 from the U.S. Army Research Office. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

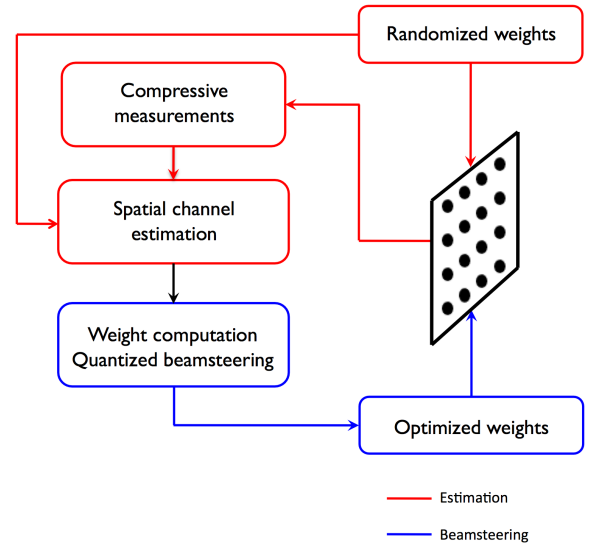


Fig. 1. Architecture to adapt large steerable arrays with coarse control: Explicit estimation of the steering direction using a few random projections, followed by beamsteering with only four phases.

amplitude and phase at each element, and downconverted to obtain a single complex baseband waveform.

Approach and Summary of Results: In this paper, we propose a new approach for adapting large arrays, shown in Figure 1, comprising the following steps. The first step is *compressive estimation*, in which a small set of measurements using RF beamforming with pseudorandom weights are employed for estimation of continuous-valued parameters characterizing a sparse spatial channel. However, rather than employing the inherently discrete ℓ_1 optimization framework of standard compressive sensing, we employ a version of orthogonal matching pursuit to obtain coarse estimates of the spatial frequencies, followed by sequential Newton refinements that provide accuracies far better than would be possible by optimizing over a discrete grid. The second step is *quantized beamsteering*, in which these explicit channel estimates are employed for computing weights for beamforming and nulling, subject to the severe quantization (e.g., phase-only control with a small number of phases). Starting with an unquantized zero-forcing solution, we show that sequential optimization provides effective solutions for heavily quantized weights.

Our goal here is to provide preliminary results for canonical problems in order to validate our overall approach, with

comprehensive system design and performance evaluation left for future work. For example, we abstract out the cross-layer protocols and low layer signal processing (e.g., correlation against training sequences, compensating for carrier offsets) required to implement compressive estimation. For quantized beamsteering, we consider the problem of steering one beam and a few nulls, without specifying whether the nulls correspond to undesired multipath or interference. A number of interesting theoretical issues require further investigation, as discussed in the conclusions.

Related work: Our approach of *explicit* estimation followed by weight computation is a stark contrast to *implicit* adaptation using classical least squares techniques. It also differs from recent codebook-based approaches to 60 GHz RF beamforming [1], [2]. The latter do not, for example, provide enough information for interference suppression. An alternative approach for implicit adaptation, which could potentially provide beamforming as well as interference suppression gains, is randomized linear ascent, proposed by Widrow and McCool more than three decades ago [3]. However, this algorithm does not scale to large arrays (at least not for rapid initial training), since its convergence time is proportional to the *square* of the number of array elements. The millimeter wave channel is quite sparse, with the number of dominant multipaths being small compared to the large number of array elements of interest to us (e.g., see [4], [5] for modeling of outdoor 60 GHz links). It is therefore natural to invoke ideas from compressive sensing [6]–[9], where a signal which is sparse with respect to a fixed basis, is reconstructed (e.g., using ℓ_1 optimization) from projections onto a small number of vectors that are “incoherent” with respect to the original basis, in terms of satisfying the so-called restricted isometry property (RIP). Roughly speaking, RIP means that the most of the energy of the sparse signal is captured by these “compressive measurements.” In our compressive estimation approach, we leverage the first part of the compressive sensing framework, in terms of capturing the required information using a small number of measurements. However, standard ℓ_1 reconstruction does not work well under basis mismatch [10] (e.g., for estimation of a continuous-valued frequency using a DFT basis), hence it is necessary to develop alternative techniques for estimating continuous-valued parameters. Theoretical frameworks for the latter are emerging [11], [12], but their implications for specific scenarios, and the development of effective algorithms, requires further work. In particular, our problem of spatial frequency estimation maps exactly to the standard frequency estimation problem recently revisited by Duarte and Baraniuk in a compressive setting [13]. We are, however, interested in both one- and two-dimensional frequencies (corresponding to linear and two-dimensional antenna arrays), and our Newton-based technique provides more refined estimates than the grid-based algorithm in [13]. Also, we are ultimately interested not just in the abstract problem of spatial frequency estimation, but also in cross-layer designs for beamsteering which use it as a building block.

Prior work related to our approach to quantized beam-

steering includes [14], [15]. A naive approach is to compute complex-valued weights without quantization constraints, followed by scalar quantization of each weight; however, this does not perform well with drastic quantization. In [14], it was shown that sequential update of quantized weights can be used to obtain effective interference suppression. We use similar ideas here, although our goal now is a multiobjective cost function including both beamforming and nulling. Beamforming patterns for large arrays with heavily quantized weights have been explored recently in [15], but direct optimization of cost functions related to communication goals is not considered there.

Outline: The system model is described in Section II along with an overview of our solution. Section III contains a detailed description of our compressive estimation algorithm, and provides numerical results showing its efficacy. Section IV describes our approach to quantized beamsteering in detail, and provides example numerical results that illustrate both its effectiveness and its suboptimality. Finally, Section V contains our conclusions.

II. SYSTEM MODEL AND OVERVIEW

Consider a uniform d spaced linear array with N elements. Typical values of interest are $d = \lambda/2$ or $d = \lambda/3$, where λ denotes the carrier wavelength. For angle of arrival (AoA) θ , the normalized array response is given by

$$\mathbf{a}(\omega) = (1, e^{j\omega}, e^{j2\omega}, \dots, e^{j(n-1)\omega})^T / \sqrt{N}, \quad (1)$$

where $\omega = 2\pi d \sin \theta / \lambda$ is the corresponding *spatial frequency*. Thus, AoA estimation can be equivalently viewed as estimation of spatial frequency. For a multipath channel with K dominant paths, the net array response is a linear combination of the form:

$$\mathbf{a}_{net} = \sum_{k=1}^{k=K} g_k \mathbf{a}(\omega_k), \quad (2)$$

where g_k denotes the complex amplitude of the k th multipath component, and ω_k , its spatial frequency.

Our goal is to estimate $\{g_k, \omega_k\}$ using M measurements of the form:

$$y_i = \mathbf{w}_i^T \mathbf{a}_{net} + z_i \quad i = 1, 2, \dots, M,$$

where \mathbf{w}_i are the *receive beamforming* weights and z_i is the measurement noise. Implementation considerations with large arrays restrict the entries of \mathbf{w}_i to a small set: for example, to two bits of phase precision, with entries chosen from $\{\pm 1, \pm j\}$. Concatenating the measurements, we get

$$\mathbf{y} = \mathbf{W}^T \mathbf{a}_{net} + \mathbf{z}, \quad (3)$$

where $\mathbf{W} = [\mathbf{w}_1 \ \mathbf{w}_2 \ \dots \ \mathbf{w}_M]$.

We begin with the following question: given that the number of AoAs K is much smaller than the number of antenna elements N , can we reduce the number of measurements M substantially? To gain some intuition, consider a special case of the problem where the all the spatial frequencies are integral

multiples of $2\pi/N$. In this case, \mathbf{a}_{net} can be expressed as a linear combination of the columns of

$$\mathbf{A} = \left[\mathbf{a}(0) \mathbf{a}\left(\frac{2\pi}{N}\right) \dots \mathbf{a}\left(\frac{2\pi(N-1)}{N}\right) \right], \quad (4)$$

so that $\mathbf{a}_{net} = \mathbf{A}\mathbf{x}$ where $\mathbf{x} \in \mathbb{C}^N$. Note, however, that only K entries of \mathbf{x} are nonzero, since there are only K AoAs. The observations \mathbf{y} can be written as

$$\mathbf{y} = \mathbf{W}^T \mathbf{A}\mathbf{x} + \mathbf{z}. \quad (5)$$

By substituting (1) in (4), we observe that the matrix \mathbf{A} is the Discrete Fourier Transform (DFT) matrix. Since the DFT matrix is orthonormal, standard Compressive Sensing theory guarantees that with (a) measurements of the form (5), (b) the elements of \mathbf{W} chosen i.i.d. from a class of distributions (including the Bernoulli distribution) and (c) the number of measurements $M = O(K \log N)$, \mathbf{x} (and thus the spatial frequencies) can be recovered with high probability. This suggests that when the spatial frequencies come from the DFT grid, roughly $K \log N$ ($\ll N$) measurements, with the beamforming weights \mathbf{W} chosen i.i.d. from $\{\pm 1, \pm j\}$, suffice to estimate them.

However, in reality, the spatial frequencies ω_k are not restricted to integer multiples of $2\pi/N$ and come from a continuum. We could ignore this fact and apply standard CS recovery algorithms assuming that spatial frequencies are integer multiples of $2\pi/N$. However, this limits the estimation accuracy to the spacing between the frequencies (even when we oversample the frequencies). We now outline our approach for overcoming this problem.

First, we estimate the AoAs coarsely by quantizing the spatial frequencies to integer multiples of $2\pi/R$ ($R \geq N$) and then applying a variant of orthogonal matching pursuit to obtain coarse estimates of $\{\omega_m\}$. These estimates are further refined using Newton-style local search; this is inspired by recent work [16] showing that Newton-based approaches are very effective in attaining fundamental bounds on timing estimation in a classical setting. We describe the algorithm in detail in Section III.

Quantized Beamsteering: Each node can use the compressive estimation algorithm outlined above on a slow time scale to track the spatial channel to/from each of its neighbors. It can then use these estimates to steer its beam towards a desired communication partner, while steering nulls towards other neighbors and undesired multipath components. Our compressive estimation framework does not require tight amplitude and phase control for the array elements. Thus, an important question is whether we can also perform tasks such as beamforming and nulling with such coarse control of array elements. We explore this in Section IV.

We begin with a standard zero-forcing solution, whose phases are quantized to the available alphabet. This is then refined by a sequential update process as in [14]. We are unable to make any claims of optimality at this point, and simply provide illustrative numerical results which show that we can effectively utilize a large array with phase-only control.

III. COMPRESSIVE ESTIMATION

In this section, we describe an algorithm to estimate multiple spatial frequencies using compressive measurements from an array. We assume that the number of frequencies K is known. The algorithm has two stages: first, we discretize the set of frequencies $[0, 2\pi]$ and use a variant of the traditional Orthogonal Matching Pursuit (OMP) to coarsely estimate the K frequencies on the discrete grid. We then refine each frequency sequentially using Newton's method. For simplicity of exposition, we first discuss estimation of a single frequency. We then use these ideas as building blocks in developing an algorithm to estimate K frequencies.

Define $\mathbf{S}(\omega)$ to be $\mathbf{W}^T \mathbf{a}(\omega)$, $\omega \in [0, 2\pi]$; the i th entry of $\mathbf{S}(\omega)$ is the DTFT of \mathbf{w}_i evaluated at ω . The measurement model in (3) reduces to

$$\mathbf{y} = \sum_{k=1}^K g_k \mathbf{S}(\omega_k) + \mathbf{z}. \quad (6)$$

Single Frequency: When there is only one frequency, the measurements \mathbf{y} are given by

$$\mathbf{y} = g_1 \mathbf{S}(\omega_1) + \mathbf{z}, \quad (7)$$

where the noise $\mathbf{z} \sim CN(0, \sigma^2 \mathbb{I}_M)$.

The estimation procedure is built on the following observation: the generalized likelihood ratio test (GLRT) estimate of ω_1 , treating the unknown gain g_1 as a nuisance parameter, is the frequency at which the *normalized* array response $\mathbf{S}(\omega)/\|\mathbf{S}(\omega)\|$ correlates best with the observations \mathbf{y} :

$$\hat{\omega}_1 = \arg \max_{\omega} J(\omega) = \arg \max_{\omega} \frac{|\mathbf{S}^H(\omega) \mathbf{y}|^2}{\|\mathbf{S}(\omega)\|^2}. \quad (8)$$

The first step is to estimate ω_1 coarsely by picking the maximum of the cost function $J(\omega)$ from among R discrete frequencies $\{0, \frac{2\pi}{R}, \dots, \frac{2\pi(R-1)}{R}\}$. Here, R is typically chosen at least as large as the number of antenna elements N .

Next, we refine the coarse estimate by using Newton's method to optimize the cost function $J(\omega)$. Let us denote the estimate of ω_1 in the i th iteration of Newton's method by $\hat{\omega}_1^{(i)}$. Starting off at $\hat{\omega}_1^{(i)}$, an additional Newton refinement step produces the estimate $\hat{\omega}_1^{(i+1)}$, given by

$$\hat{\omega}_1^{(i+1)} = \hat{\omega}_1^{(i)} - \frac{J'(\hat{\omega}_1^{(i)})}{J''(\hat{\omega}_1^{(i)})}. \quad (9)$$

However, since Newton's method only solves for $J'(\omega) = 0$, it could lead to a minimum of $J(\omega)$, instead of the desired maximum, when the function is convex at $\hat{\omega}_1^{(i)}$, or $J''(\hat{\omega}_1^{(i)}) \geq 0$. We therefore use (9) to refine the estimate only when $J(\omega)$ is concave at $\hat{\omega}_1^{(i)}$.

When the function is convex at $\hat{\omega}_1^{(i)}$, we use the following observation to design an alternate refinement rule: if the GLRT estimate $\hat{\omega}_1$ in (8) is close to the true frequency ω_1 ("reasonable" number of measurements and noise), then $\mathbf{y} = \mathbf{S}(\omega_1) + \mathbf{z} \approx \mathbf{S}(\hat{\omega}_1)$. Thus, the maximum of the cost

function $J(\hat{\omega}_1) \approx \|\mathbf{y}^2\|$. Therefore, to get closer to the optimum, we design a Newton step that solves $J(\omega) = \|\mathbf{y}\|^2$. This gives us the update rule

$$\hat{\omega}_1^{(i+1)} = \hat{\omega}_1^{(i)} + \frac{\|\mathbf{y}\|^2 - J(\hat{\omega}_1^{(i)})}{J'(\hat{\omega}_1^{(i)})}. \quad (10)$$

Multiple Frequencies: When there is more than one frequency to be estimated, we use each part of the algorithm described above recursively. However, there is one key modification: while estimating the l th frequency ω_l , we project out the contributions from the frequencies that we have already estimated to the observations \mathbf{y} (and also the template $\mathbf{S}(\omega)$ for the refinement stage), thereby roughly reducing it to a single frequency problem at each stage. We briefly describe the coarse estimation and the Newton refinement steps.

Coarse estimation: Suppose that we have estimated $l-1$ frequencies $\hat{\omega}_1, \hat{\omega}_2, \dots, \hat{\omega}_{l-1}$ coarsely. The contribution from these frequencies to the observations \mathbf{y} lies in the span of the matrix

$$\mathbf{B} = [\mathbf{S}(\hat{\omega}_1) \quad \mathbf{S}(\hat{\omega}_2) \quad \dots \quad \mathbf{S}(\hat{\omega}_{l-1})].$$

Thus, we can eliminate these contributions by operating on the observations \mathbf{y} with the orthogonal projection matrix $\mathbf{B}^\perp = \mathbb{I} - \mathbf{B}(\mathbf{B}^H\mathbf{B})^{-1}\mathbf{B}^H$ to obtain the residual vector,

$$\mathbf{r}_{l-1} = \mathbf{B}^\perp \mathbf{y}.$$

We now pretend that the l th frequency ω_l is the only one left in \mathbf{r}_{l-1} and estimate it exactly as in the single frequency setting: by sampling the function

$$J_l(\omega) = \frac{|\mathbf{r}_{l-1}^H \mathbf{S}(\omega)|^2}{\|\mathbf{S}(\omega)\|^2}$$

at $\{0, \frac{2\pi}{R}, \dots, \frac{2\pi(R-1)}{R}\}$, $R \geq N$ and picking the maximum. We denote this estimate by $\hat{\omega}_l$, add it to the set of estimated frequencies and continue with the coarse estimation procedure until we estimate K frequencies.

Newton refinement: We refine our estimates of the K frequencies by applying Newton's method to them sequentially. The refinement process for each frequency is very similar to the single frequency case; the only difference is that the contributions from the other frequencies are projected out (from both \mathbf{y} and $\mathbf{S}(\omega)$) while defining the cost function.

Suppose that at some stage of the refinement process, our estimates of the K frequencies are given by $\hat{\omega}_1, \hat{\omega}_2, \dots, \hat{\omega}_K$ respectively and we want to further refine the l th frequency. Just as before, the contributions from the other $K-1$ frequencies can be eliminated from the observations \mathbf{y} by using the orthogonal projection matrix \mathcal{H}^\perp , which is obtained from

$$\mathcal{H} = [\mathbf{S}(\hat{\omega}_1) \quad \dots \quad \mathbf{S}(\hat{\omega}_{l-1}) \quad \mathbf{S}(\hat{\omega}_{l+1}) \quad \dots \quad \mathbf{S}(\hat{\omega}_K)]$$

as $\mathcal{H}^\perp = \mathbb{I} - \mathcal{H}(\mathcal{H}^H\mathcal{H})^{-1}\mathcal{H}^H$. The residual is given by $\mathcal{H}^\perp \mathbf{y}$ and we denote this by \mathbf{y}_l . Assuming that the $K-1$ frequency estimates are reasonably good, we have,

$$\mathbf{y}_l \approx g_l \mathcal{H}^\perp \mathbf{S}(\omega_l) + \tilde{\mathbf{z}}$$

This resembles the model (7) for single frequency estimation. We can therefore use the Newton updates specified before to refine $\hat{\omega}_l$ by replacing \mathbf{y} with \mathbf{y}_l and $\mathbf{S}(\omega)$ with $\mathcal{H}^\perp \mathbf{S}(\omega)$ in the definition of the cost function (8), (9) and (10).

One round of Newton updates consists of applying the above refinements to each of the K frequencies, and we run multiple such rounds.

Results: We begin by simulating the problem of estimating $K = 2$ spatial frequencies with an $N = 32$ element array, with the elements spaced $\lambda/2$ apart. We find that $M = 12$ measurements typically suffice to estimate the two AoAs. Figure 2 shows that, in the absence of measurement noise, the proposed Newton refinements improve the estimates of the spatial frequencies to any desired accuracy. Figure 3 is a scatterplot of the estimation errors before and after refinement in the presence of measurement noise over 1000 trials. We see that the errors after refinement are considerably smaller (most of them lie well below the line with slope 1 shown in blue). For this scenario, the algorithm succeeds in 96% of the trials (failures are due to gross errors in the coarse estimation stage) and the Newton refinements improve the median estimation error from 0.4015° to 0.1376° .

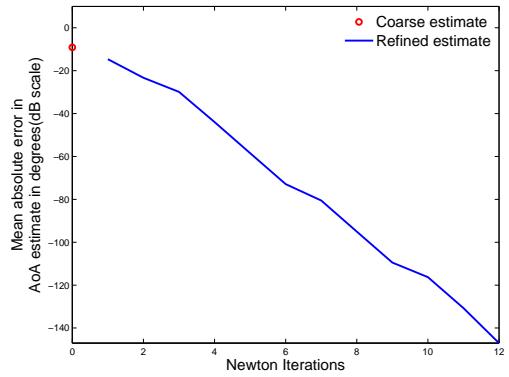


Fig. 2. $10 \log_{10}$ (Average AoA error) is plotted as a function of refinement rounds. We take $M = 12$ noiseless measurements from a $N = 32$ element $\lambda/2$ spaced array when $K = 2$ beams impinge the antenna array. A coarse estimate of the frequencies is made using a $4N = 128$ grid.

To illustrate that we need significantly fewer measurements than the number of elements N as the array gets larger, we consider the problem of estimating $K = 2$ spatial frequencies (one at 12dB and another at 9dB) from a linear array of $N = 1024$ elements spaced $\lambda/2$ apart. We find that $M = 24$ measurements (only 2.3% of the number of elements) suffice to estimate the two AoAs, with the refinement process improving the median error from $9.6^\circ \times 10^{-3}$ to $2.7^\circ \times 10^{-3}$.

We are currently investigating extensions of the proposed algorithm to two dimensional arrays. Preliminary simulations are encouraging: we find that $M = 24$ measurements once again suffice to estimate 2 AoAs (12dB and 9dB above noise level) from an $N = 1024(32 \times 32)$ element array. Details will be reported in later publications.

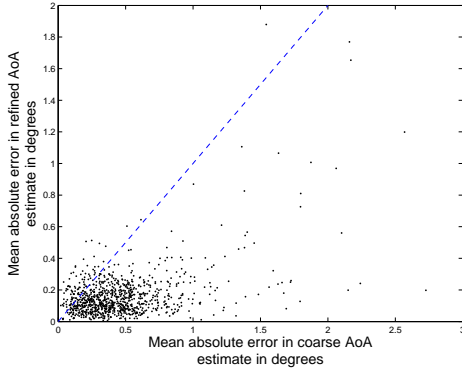


Fig. 3. Comparison of AoA estimation errors in the presence of noise before and after refinements of the spatial frequencies. $M = 12$ measurements made from a $N = 32$ element $\lambda/2$ spaced array when $K = 2$ beams, one at 12dB and another at 9dB fall on the array. $4N = 128$ grid points used to arrive at the coarse estimates. The blue dashed line is the slope one line.

IV. QUANTIZED BEAMSTEERING

The AoA estimates from Section III give us the directions in which we need to steer transmissions and place nulls (to combat multipath or avoid interference). However, hardware design imposes severe constraints: we would like to steer beams by only changing the antenna phases that are heavily quantized (say, to two bits of precision). We now explain how we can do this effectively.

Consider a two dimensional square array consisting of N (typically $32 \times 32 = 1024$) elements. The separation between the elements along either side is d . Denoting the transmit gain and phase at the (m, n) th element by g_{mn} and β_{mn} respectively, the received power at an azimuth ϕ and elevation θ is given by

$$P = \left| \sum_{m,n} g_{mn} e^{j\beta_{mn}} \exp(j(m\omega_x + n\omega_y)) \right|^2, \quad (11)$$

where $\omega_x = (2\pi d/\lambda) \sin \theta \cos \phi$ and $\omega_y = (2\pi d/\lambda) \sin \theta \sin \phi$ are the spatial frequencies associated with the elevation and azimuth. For convenience, we denote the combined spatial frequencies (ω_x, ω_y) by $\boldsymbol{\omega}$ and the power at $\boldsymbol{\omega}$ by $P(\boldsymbol{\omega})$.

Our goals are twofold: (a) The received power in the direction of intended receiver $P(\boldsymbol{\omega}_0)$ must be as high as possible. (b) The interference caused in Q other directions (these may represent other users or undesired multipath components), given by $P(\boldsymbol{\omega}_i), i = 1, \dots, Q$, must be as small as possible. We try to achieve these goals simultaneously by maximizing the *signal-to-null* ratio γ , given by

$$\gamma = \frac{P(\boldsymbol{\omega}_0)}{\sum_{i=1}^Q P(\boldsymbol{\omega}_i)}. \quad (12)$$

We abstract the hardware constraints as follows: (a) The gains g_{mn} are fixed once and for all to a two dimensional Chebyshev window (in order to control undesired sidelobes). (b) The phase shifts β_{mn} which we use to steer the transmission can only take one of the four values $\{\pm 1, \pm j\}$.

Before describing our approach, we introduce some notation. We separate the terms in (11) into the phases that we control and the rest that are fixed. Let $\boldsymbol{\psi}$ be a vectorized version of the phases $[e^{j\beta_{mn}}]$ and $\mathbf{a}(\boldsymbol{\omega})$ be the corresponding vectorized version of the steering matrix $[g_{mn} \exp(-j(m\omega_x + n\omega_y))]$. Then the power $P(\boldsymbol{\omega}) = |\mathbf{a}(\boldsymbol{\omega})^H \boldsymbol{\psi}|^2$. We also denote the l th entry of $\boldsymbol{\psi}$ by $\psi[l]$.

The basic idea behind our algorithm is as follows: given a feasible solution for the phases $\boldsymbol{\psi}$, we can improve it with low-complexity until we settle at a local optimum. To see this, note that $\psi[l]$ takes only one of four values $\{\pm 1, \pm j\}$. Therefore, given the phases at the other elements, we can easily find which of these choices for $\psi[l]$ maximizes γ : hold the other phases fixed, try out each of the four phases in the l th position of $\boldsymbol{\psi}$, use (12) to compute γ for each of these candidates and pick the maximum. Let us denote the maximizing phase by $\alpha^* \in \{\pm 1, \pm j\}$. If $\psi[l]$ is different from α^* , we can improve the solution by simply replacing $\psi[l]$ with α^* . By repeating this procedure for different choices of l , we can improve γ , without ever lowering it, until we settle at a local optimum. We refer to the phases at the end of this procedure as the sequentially optimized phases.

Since the powers are computed at specific frequencies in (12), the solution thus obtained could be sensitive to the errors in estimating these frequencies from Section III. To reduce this sensitivity, we replace the metric γ with a modified version $\tilde{\gamma}$, where the powers are computed in a small band around each frequency $\boldsymbol{\omega}_i, i = 0, 1, \dots, Q$:

$$\tilde{\gamma} = \frac{\int P(\boldsymbol{\omega}_0 - \mathbf{h}) d\mathbf{h}}{\sum_{i=1}^Q \int P(\boldsymbol{\omega}_i - \mathbf{h}) d\mathbf{h}}.$$

The width of the band (the region of integration) is dictated by the estimation error in Section III.

Initialization: The only thing that we need to specify now is a starting point that is both feasible and “reasonably good” (so that we do not settle at a bad local optimum). We compute the starting point by relaxing the constraints on the entries of $\boldsymbol{\psi}$, allowing them to take any value whatsoever. We can now cast our problem as a standard zero-forcing problem of choosing weights to null out transmissions in Q directions perfectly, while maximizing the power in the intended direction. Specifically, we solve:


$$\begin{aligned} & \text{maximize} && |\mathbf{a}(\boldsymbol{\omega}_0)^H \boldsymbol{\psi}| \\ & \text{subject to} && \|\boldsymbol{\psi}\| = 1 \\ & && \mathbf{a}(\boldsymbol{\omega}_i)^H \boldsymbol{\psi} = 0 \quad \forall i \in \{1, 2, \dots, Q\} \end{aligned}$$

We then quantize these phases to the closest among $\{\pm 1, \pm j\}$, thus giving us a feasible solution to the problem that is also reasonably good. We term these phases the naïvely quantized phases and use them to initialize our sequential optimization procedure.

Results: We steer the transmission towards a receiver while simultaneously placing nulls in $Q = 2$ directions, using a 32×32 element array. Figure 4(a) plots histograms of γ obtained with the naïvely quantized phases and the sequentially

optimized phases. We see that γ is tightly clustered around its mean value of 58 dB for the sequentially optimized phases. The naïvely quantized phases have a lower γ on the average (48 dB) and also exhibit greater variance.

Figure 4(b) plots the gains provided by sequential optimization against the signal-to-null ratio obtained with the naïvely quantized phases. We see that the improvements are large when naive quantization does poorly and decrease as naive quantization become better. Thus, the sequential optimization procedure gives us the largest gains exactly when we need them.



(a) Histogram of γ optimized (red) and naïvely quantized (blue) phases (b) Improvements provided by sequential optimization plotted against the naïve signal to null ratio

Fig. 4. Results for quantized beamsteering with a 32×32 element array with the elements placed $\lambda/2$ apart. We place $Q = 2$ nulls while steering transmissions towards a receiver.

V. CONCLUSIONS

The preliminary results in this paper show the feasibility of our approach to compressive adaptation of large arrays with drastically simplified hardware control: the architecture can be realized using RF beamsteering with coarse-grained control of the phases of the array elements. The proposed approach of explicit estimation and weight computation differs fundamentally from implicit adaptation using classical least squares, which is incompatible with RF beamforming, as well as from codebook-based techniques, which enable coarse beamforming but not nulling.

An important topic for future work is the design of cross-layer protocols and signal processing for compressive adaptation in specific settings of interest, such as for packetized 60 GHz backhaul mesh networks [17]. At a fundamental level, it is important to develop a theoretical understanding of the limits of compressive estimation of continuous-valued parameters, as well as algorithms for attaining these limits. This problem has received far less attention than the “discrete” compressive sensing problem. Similarly, we would like to develop a theoretical understanding of the limits of quantized beamsteering, as well as improved algorithms for computing optimal or near-optimal solutions at reasonable complexity.

Acknowledgements

This work was informed and motivated by discussions with Prof. Mark Rodwell on large arrays and the associated hardware constraints.

REFERENCES

- [1] J. Wang, Z. Lan, C. woo Pyo, T. Baykas, C. sean Sum, M. Rahman, J. Gao, R. Funada, F. Kojima, H. Harada, and S. Kato, “Beam codebook based beamforming protocol for multi-gbps millimeter-wave wpan systems,” *Selected Areas in Communications, IEEE Journal on*, vol. 27, no. 8, pp. 1390–1399, october 2009.
- [2] S. Lin, K. Ng, H. Wong, K. Luk, S. Wong, and A. Poon, “A 60ghz digitally controlled rf beamforming array in 65nm cmos with off-chip antennas,” in *Radio Frequency Integrated Circuits Symposium (RFIC), 2011 IEEE*, june 2011, pp. 1–4.
- [3] B. Widrow and J. McCool, “A comparison of adaptive algorithms based on the methods of steepest descent and random search,” *Antennas and Propagation, IEEE Transactions on*, vol. 24, no. 5, pp. 615–637, sep 1976.
- [4] H. Zhang, S. Venkateswaran, and U. Madhow, “Channel modeling and MIMO capacity for outdoor millimeter wave links,” in *Proc. IEEE WCNC 2010*, April 2010.
- [5] H. Zhang and U. Madhow, “Statistical modeling of fading and diversity for outdoor 60GHz channels,” in *Proc. International Workshop on mmWave Communications (mmCom 2010)*, September 2010.
- [6] E. Candes, J. Romberg, and T. Tao, “Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information,” *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489–509, feb. 2006.
- [7] E. Candes and T. Tao, “Near-optimal signal recovery from random projections: Universal encoding strategies?” *Information Theory, IEEE Transactions on*, vol. 52, no. 12, pp. 5406–5425, dec. 2006.
- [8] —, “Decoding by linear programming,” *Information Theory, IEEE Transactions on*, vol. 51, no. 12, pp. 4203–4215, dec. 2005.
- [9] D. Donoho, “Compressed sensing,” *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, april 2006.
- [10] Y. Chi, L. Scharf, A. Pezeshki, and A. Calderbank, “Sensitivity to basis mismatch in compressed sensing,” *Signal Processing, IEEE Transactions on*, vol. 59, no. 5, pp. 2182–2195, may 2011.
- [11] M. Duarte and Y. Eldar, “Structured compressed sensing: From theory to applications,” *Signal Processing, IEEE Transactions on*, vol. 59, no. 9, pp. 4053–4085, sept. 2011.
- [12] T. Blumensath and M. Davies, “Sampling theorems for signals from the union of finite-dimensional linear subspaces,” *Information Theory, IEEE Transactions on*, vol. 55, no. 4, pp. 1872–1882, april 2009.
- [13] M. Duarte and R. Baraniuk, “Recovery of frequency-sparse signals from compressive measurements,” in *Communication, Control, and Computing (Allerton), 2010 48th Annual Allerton Conference on*, 29 2010-oct. 1 2010, pp. 599–606.
- [14] O. Bakr, M. Johnson, R. Mudumbai, and U. Madhow, “Interference suppression in the presence of quantization errors,” in *Communication, Control, and Computing, 2009. Allerton 2009. 47th Annual Allerton Conference on*, 30 2009-oct. 2 2009, pp. 1161–1168.
- [15] C.-P. Yeang, G. Wornell, and L. Zheng, “Oversampling transmit and receive antenna arrays,” in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, march 2010, pp. 2522–2525.
- [16] P. Bidigare, U. Madhow, R. Mudumbai, and D. Scherber, “Attaining fundamental bounds on timing synchronization,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2012)*, Kyoto, Japan, March 2012, to appear.
- [17] S. Singh, R. Mudumbai, and U. Madhow, “Distributed coordination with deaf neighbors: efficient medium access for 60 GHz mesh networks,” in *Proc. IEEE Infocom 2010*, San Diego, CA, USA, March 2010.