Basis Pursuit for Robust
Passive Acoustic Beamforming

Ben Shapo and Chris Kreucher
Integrity Applications Incorporated
900 Victors Way, Suite 220
Ann Arbor, MI 48108
bshapo@integrity-apps.com, ckreuche@umich.edu

Abstract[1] — **Beamforming is a process that supplies directional gain to sensor array processing. One modality where beamforming adds great value is passive sonar, where real-aperture arrays receive signals emitted by acoustic sources. In passive sonar systems, the beamformer is the backbone of a processing structure that detects, localizes, and classifies external targets. Conventional beamformers use deterministic time-delays (often implemented as phase shifts) to arrange coherent addition of plane-wave signals at each sensor. Recently, adaptive beamformers take advantage of signal time history by imposing a model on the environment. Basis Pursuit is another reconstruction approach used in Compressed Sensing that also enforces a physics-based model – in this case a model of scene sparsity. This paper describes an application of this technique to the beamforming problem. The main benefit of the Basis Pursuit beamforming approach is that it is robust to missing array elements, providing nearly full-aperture performance in a reduced sensor environment. This result is advantageous in the case of processing with inoperative hydrophones. It may also provide cost savings by allowing array design with fewer hydrophones.**

I. INTRODUCTION

Recently, the signal processing community has shown great interest in Compressed Sensing ("CS") [1]. In CS, fewer (sometimes significantly fewer) samples of a physical phenomenon than traditional Nyquist bounds would mandate reconstruct the original signal via a technique such as Basis Pursuit (BP). Good reconstruction is possible under the assumption that the sampled signals are actually sparse. Passive sonar often obeys this assumption, with 180- (or 360)-degree field of view but few targets.

This paper describes a beamforming method that exploits the sparsity constraints used in BP. Like the Conventional Beamformer ("CBF"), we assume each signal impinging on a passive array of sensors arrives as a plane-wave [2]. In addition, we also assume the number of plane-wave signals is small compared to the number of arrival angles considered ("beams").

Our method produces results comparable to CBF in the case of a full-aperture array. However, when the number of sensors decreases (either by hyrdophones becoming unexpectedly inoperable or by design), our method produces higher-fidelity results than the CBF approach operating on the same reduced input data.

This paper proceeds as follows. Section II reviews a standard model of plane-wave physics and derives the CBF. A simulated example demonstrates the power of this classic approach, and highlights the benefits of using data in the wide frequency ranges available with broadband sensors.

Section III builds on the model presented in Section II, and gives a brief description of our BP approach to beamforming. This section shows results from a BP-beamformer operating on the same simulated data input from Section II.

Section IV demonstrates the benefits of the methodology on Navy collected at-sea data by comparing CBF and BP beamformer outputs in the case where the number of sensors has been significantly reduced.

II. PLANE WAVE MODEL AND CBF

A. Plane Wave Model

This section's contents are based on material from [2]. We assume an array of acoustic sensors with energy from acoustic sources impinging on the elements in the sensor array. The array's aperture refers to the fact that each element resides at a different spatial location. This diversity causes propagation times from a target source to each element to differ. In the frequency domain, these propagation times manifest themselves as different phase values. Figure 1 illustrates the situation for the special case of a linear array. A plane wave (red line) arrives at an angle θ with respect to the array's endfire direction. If the distance from the source to the array's center element is r, then the distance from the source to an element located at a distance D from the center element is: $r - \Delta r$, where $\Delta r = D \cos(\theta)$.

For a narrowband acoustic signal, or for a single Fourier component of a broadband signal (i.e., a plane wave), this relation provides the differential phase for a signal received at the nth sensor element:

$$\Delta\varphi(n) = \frac{2\pi}{\lambda}\Delta r(n) , \qquad (1)$$

where λ is the wavelength of the plane wave.

B. CBF

The received signals would add coherently across elements if their arrival times (or phases) were forced to align. The idea of the Conventional Beamformer ("CBF") is to hypothesize a large set of candidate arrival angles ("beams") and, on a per-beam basis, impart appropriate phase delays at

each element to force any signals arriving from that direction to add coherently. This processing causes signals in each hypothesized direction to add constructively, while signals emanating from other directions do not.
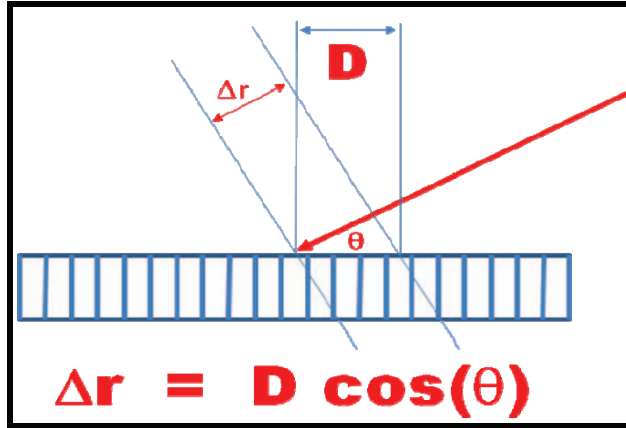

Figure 1: Sensor Array Geometry.

From the time-delay relationship at the bottom of Figure 1, time delay is a function of the cosine of arrival angle. Thus, in the figures that follow, the independent axis representation of arrival angle is $u = \cos(\theta)$ (rather than θ).

The quantity that describes a narrowband, plane-wave signal arriving from an angle u as it impinges on each element is a $N\times 1$ vector (where N is the number of elements in the array). Each term in the vector represents the phase rotation of the signal as it arrives at the corresponding element. This quantity is called a "steering vector" or a "manifold vector." Denote the steering vector for a 1D array as $\underline{a}$. We have (from Figure 1) this expression for a plane wave signal arriving from angle specified by u for the nth element in the vector:

$$a_n(u) = e^{-jkp(n)u} \qquad (2)$$

where $p(n)$ is the location of the nth element in 1D, and the wavenumber k is $2\pi/\lambda$. Thus, for a 1D array with elements uniformly spaced by a distance d, the entire vector becomes:

$$\underline{a}(u) = \begin{bmatrix} 1 & e^{-jkdu} & e^{-jk2du} & \cdots & e^{-jk(N-1)du} \end{bmatrix}^T \qquad (3)$$

Implementation of a CBF beamformer at a single bearing (or beam) simply involves rotating the phase of each hypothesized signal (i.e., a plane wave at arrival angle specified by u) by the appropriate amount to cause the signal at each element to combine coherently. This procedure amounts to simply pre-multiplying the signal by $\underline{a}^H$, where $()^H$ denotes the conjugate transpose operation.

By extension, we can represent the entire CBF operation, where we postulate a large number of arrival directions, as a linear filter with fixed coefficients. Recall that N is the number of elements, and defining N_θ as the number of hypothesized arrival angles (beams), a $N\times N_\theta$ "steering matrix"

$$\underline{\underline{A}} = \begin{bmatrix} \underline{a}(u_1) & \underline{a}(u_2) & \cdots & \underline{a}(u_{N_\theta}) \end{bmatrix}^T \qquad (4)$$

allows the operation

$$\underline{b} = \underline{\underline{A}}^H \underline{z} \qquad (5)$$

to perform the entire beamforming operation (all N_θ beams at a single frequency), where $\underline{z}$ is a $N\times 1$ vector representing the received data at each element. Here, the output $\underline{b}$ is a $N_\theta \times 1$ vector representing all output beams.

Figure 2 (blue line) shows a simulated CBF profile for two ideal (noise-free environment) targets: one at u=-0.1 and the other at u=0.1. As is evident from the figure, the targets have different SNR values. There are 51 sensor elements and CBF forms 201 beams, at the array's design frequency (i.e., the element spacing is $\lambda/2$ at the plane wave's frequency). The figure clearly indicates peaks at the arrival angles corresponding to the two targets, as well as a good deal of sidelobe structure. The red line shows a more realistic beam profile, where additive noise corrupts the arriving signals at each element.
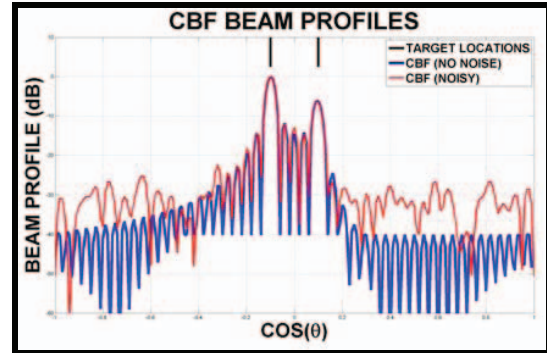

Figure 2: Standard CBF Processing (2 Targets).

C. Broadband

A powerful technique for improving target SNR and localization is to take advantage of the broad frequency bands afforded by many passive sensors. Simple incoherent combination of beam profiles (like those in Figure 2) over multiple frequencies results in improved performance. Figure 3 provides a simple example. Here, beam profiles for the same scenario as Figure 2 appear for different frequencies. The blue line indicates the profile for the wavelength corresponding to the element spacing, and the green line shows the profile for the wavelength that corresponds to one-fourth the element spacing. Note that operating the array at this frequency causes spatial aliasing (or "grating lobe") artifacts. The red line indicates simple incoherent addition of beam profiles over a band of frequencies from DC to four times the array's design frequency (i.e., element spacing of 2λ at that frequency). Note that the red line possesses narrower mainbeam width than many of the low-frequency bins, with sidelobe structure as well as grating lobes almost eliminated by the averaging process.
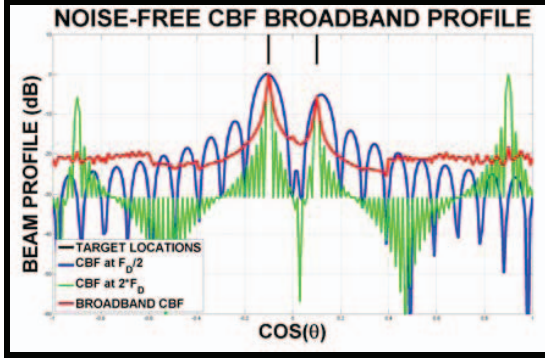
Figure 3: Broadband CBF.

III. BASIS PURSUIT

CBF processing is a very effective technique and it is employed in many real scenarios. However, it does not exploit physical characteristics of the environment that may lead to improved performance. This section describes a plane-wave model for beamforming, and develops an alternate formulation of the beamforming problem based on the Basis Pursuit ("BP") method.

A. Plane-wave signals

We have already written a single plane-wave signal arriving from arrival angle θ (where $u \equiv \cos(\theta)$) via Equation (3) above, and also written the expression for a collection of unit-amplitude plane-wave signals as $\underline{\underline{A}}$ in Equation (4). Given this notation, a simple expression describes a set of signals arriving at various angles θ_m:

$$\underline{z} = \underline{\underline{A}}\,\underline{s} + \underline{n} \qquad (6)$$

Here, $\underline{s}$ is the set of complex source signals ($N_\theta \times 1$); $\underline{\underline{A}}$ is the steering matrix, which maps signals to received data (a sensor model); $\underline{n}$ is an ($N_\theta \times 1$) noise vector at each of the potential source locations; and $\underline{z}$ is the data received at the sensor elements ($N \times 1$) The crucial assumption here is that there is a potential source of energy (target) at each of the N_θ arrival angles, and all entries in the $\underline{s}$ vector corresponding to angles where no target is present are zero. In other words, most of the N_θ entries in the complex signals amplitude vector $\underline{s}$ are zero and $\underline{s}$ is sparse.

B. Basis Pursuit Approach to Beamforming

Our application of BP to the beamforming problem relies on two basic assumptions. One is narrowband signals arrive at the array as plane-waves. This is a common assumption made in most beamforming algorithms. The second assumption is that there are a small number of signals compared to the number of available sensor elements. For example, the scenario that gave rise to Figure 2 and Figure 3 has two signals. In this case, most beams represent directions with no signal present.

Many authors have produced tools to solve matrix equations subject to the constraint that the solution be sparse. Some have demonstrated that, under certain conditions [3][5], solving

$$\arg\min_{s} \quad \|s\|_1 \quad \text{subject to} \quad \left\|\underline{\underline{A}}\underline{s} - \underline{z}\right\|_2 \leq \sigma \qquad (7)$$

is equivalent to solving

$$\arg\min_{s} \quad \|s\|_0 \quad \text{subject to} \quad \left\|\underline{\underline{A}}\underline{s} - \underline{z}\right\|_2 \leq \sigma \qquad (8)$$

i.e., the one-norm minimization generates the maximally sparse solution to the matrix equation (Equation (6)).

Returning to the simulations of Figure 2 and Figure 3, we compute the BP solution to the same input data. The particular BP implementation is publicly available [4] from the authors of [3]. The result appears in Figure 4, where the blue line reproduces the noise-free CBF solution from Figure 2, and the red line shows the sparse BP solution. We have picked sensible values for the optimization parameters required by the algorithm and error criterion. The BP approach is able to isolate the two plane-wave targets without any false targets in the sidelobe structure.
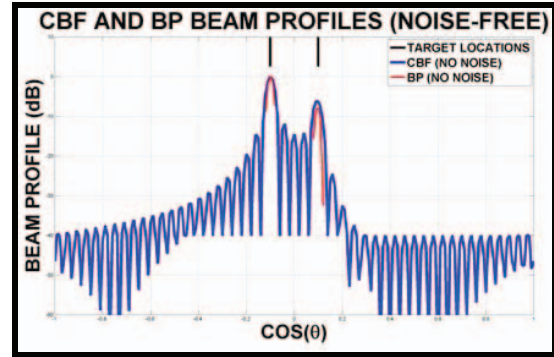

Figure 4: BP vs. CBF for two Plane Wave Signals.

A more realistic scenario is one with corrupting noise, as in Figure 2. We added Gaussian noise to the target signals, and the results appear in Figure 5. Here, the noise floor pushes the far-out, noise-free, 40dB CBF sidelobes (with respect to maximum signal) up by about 10dB to approximately 30dB. The BP solution, however, is virtually unaffected by the noise, and still produces two signal peaks at the correct arrival angles with no false signals.
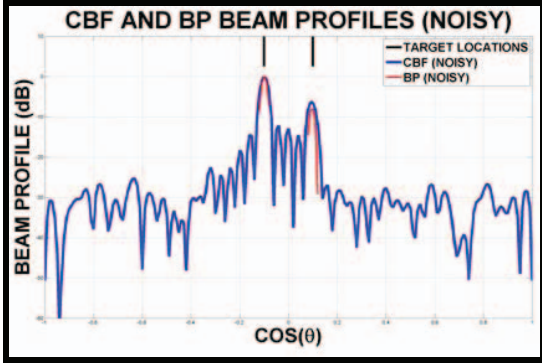
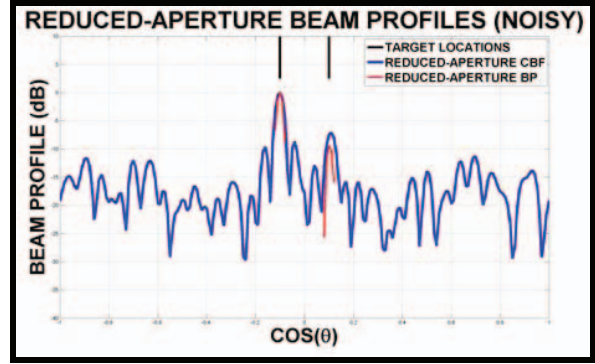**Figure 5: BP vs. CBF for Plane Wave Signals in Noise.**

Figure 7: CS-BP v. CS-CBF with Failed Hydrophones.

C. Applications

The BP approach produces solutions with simulated data characterized by much more favorable sidelobe profiles than CBF. In addition, another application for BP methods is operating a reduced number of sensor elements. This often occurs on deployed arrays in the presence of failed hydrophones. Additionally, it may assist in the design of future arrays which can deliver higher precision beamforming then the conventional approach.

The simulation shown in Figure 6 and Figure 7 illustrate the need for an effective approach. Figure 6 repeats the noisy CBF beam profile for the two targets in blue. The red line indicates CBF performance with a random set of 20 hydrophones removed (of the original 51). In this case, the processing is the same as Equation (5) but with the $\underline{\underline{A}}$ matrix reduced, each deleted column corresponding to a missing hydrophone. It is very clear that performance has declined significantly, with raised inner sidelobe levels and a 15 dB increase in the noise floor region.
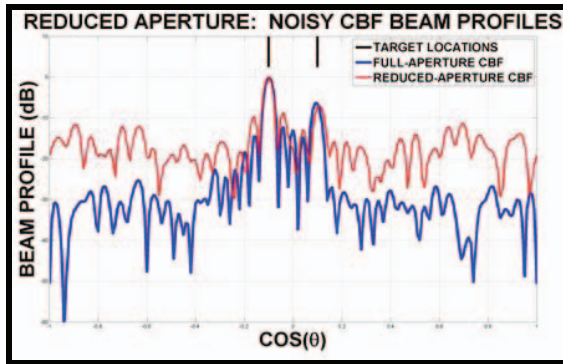
IV. REAL DATA RESULTS

A. Passive Sonar Display Modality

This section illustrates the performance of the BP approach using real data collected during an at-sea experiment. The single time graph in Figure 3 (red line) is useful because it provides information about where targets are located in space at a frozen time snapshot. However, a more useful display provides the bearing information present in that curve over time. Figure 8 shows an example of such a display. Each horizontal slice shows the equivalent of the red line in Figure 3 in grayscale (i.e., a bearing profile, in cosine-space). White pixels represent high-energy bearing-time events, and black pixels correspond to low-energy. The vertical axis represents time. Thus objects that move in the x-direction on these displays have bearing rate relative to the sensor over time. Vertical objects, like all the ones in this real-world scenario, do not have significant bearing changes over time. Because of its format, this display is known as a Bearing-Time Record ("BTR.").

Figure 6: CBF with Failed Hydrophones.

The BP approach is able to produce nearly full-fidelity results with the dramatically reduced number of sensors. Figure 7 shows the outcome in red, and the reduced-sensors CBF result appears for reference in blue. Comparing the red BP curve here with the full-aperture BP curve (Figure 5, red) indicates that the large number of missing hydrophones (40%) has not significantly impacted target-finding performance.
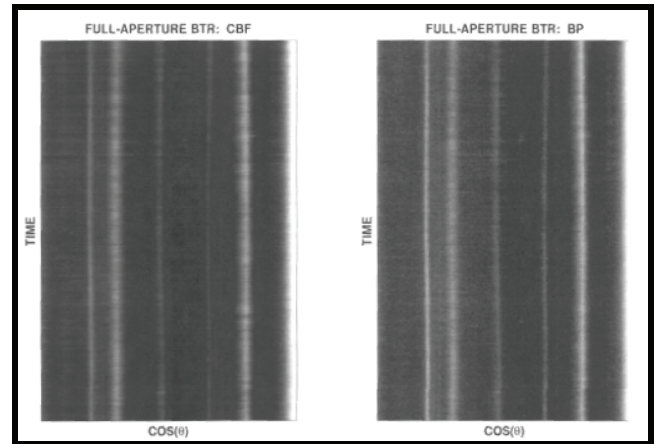

Figure 8: BTR Displays: Full-Aperture CBF v. BP.

This figure shows that the CBF (left-panel) and BP (rght-panel) approaches produce comparable results for the full-aperture case. The same targets are evident in both images.

B. BP in the Presence of Reduced Sensor Elements

Reproducing the conditions of Figure 6 (eliminating a significant number of the hydrophones) over all frequency bins of course has a negative impact on BTR quality. In the left panel of Figure 9, CBF processing seems to have lost or almost lost two of the 5 clearly visible, non-endfire targets. However, BP, operating on the same exact input data, is able to maintain these targets on the display.
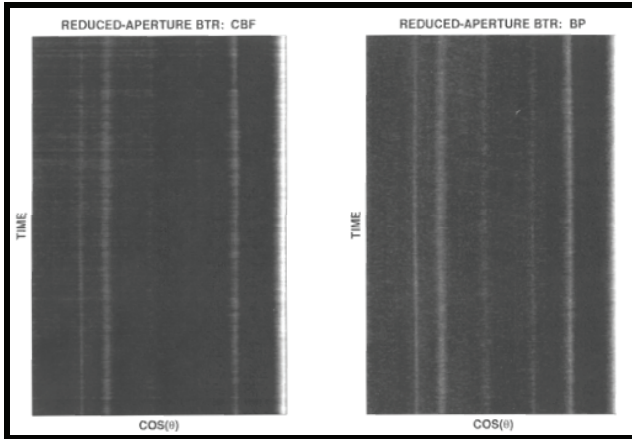

Figure 9: BTR Displays: CS-Aperture CBF v. BP.

The next two figures quantify the gains BP enjoys over CBF in this reduced sensor situation. In Figure 10, color overlays indicate target bearing. A simple maximum amplitude search over a manually specified window (both for CBF and BP) determines this truth data.

The target truth bearing-time trajectories in Figure 10 provide a method for estimating target SNR on the respective BTR displays. On a per-target, per-time-slice basis, signal is defined as the BTR amplitude in the signal beam, and noise is defined as the median amplitude over a local window (centered at the target truth for each target). Figure 11 shows the SNR for each target (corresponding to the five targets in Figure 10, numbered left to right) as a function of time for the entirety of the data that appears in Figure 8. With no exceptions, the BP-CS approach produces a higher SNR target that the CBF-CS processing.
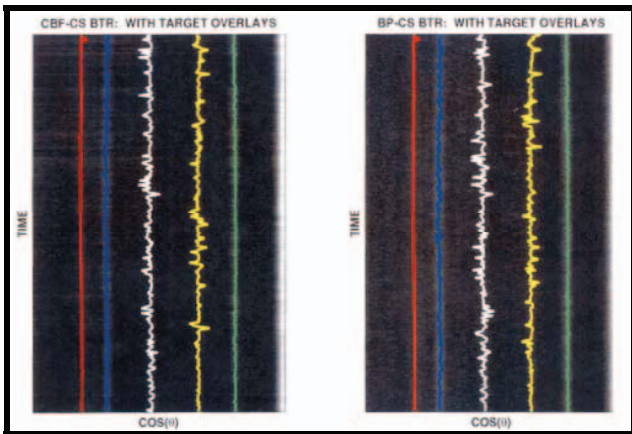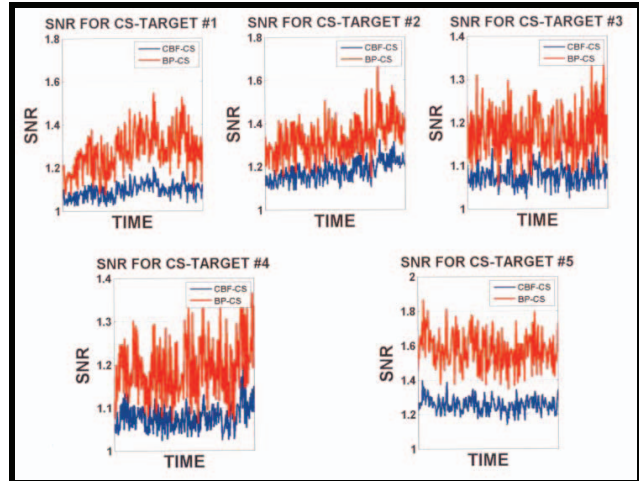

Figure 10: CS-Aperture BTRs: With Target Overlays.

Figure 11: Target-SNR curves: BP-CS vs. CBF-CS.

REFERENCES

[1]   Emmanuel J. Candes and Michael B. Wakin, "An Introduction to Compressive Sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21-30, March, 2008.

[2]   Harry L. Van Trees, *Detection, Estimation, and Modulation Theory, Part IV, Optimum Array Processing.* New York: John Wiley & Sons, 2002.

[3]   Ewout Van Den Berg and Michael P. Friedlander, Probing the Pareto Frontier for Basis Pursuit Solutions, Siam J. Sci. Computing, Vol. 31, No. 2: pp. 890–912.

[4]   http://www.cs.ubc.ca/labs/scl/spgl1/

[5]   E. J. Candes, J. Romberg, and T. Tao, Stable signal recovery from incomplete and inaccurate measurements, Comm. Pure Appl. Math., 59 (2006), pp. 1207–1223.