Cluster analysis and robust use of full-field models for sonar beamforming

Brian Tracey, , Nigel Lee, , Srinivas Turaga, and , and

Citation: The Journal of the Acoustical Society of America **120**, 2635 (2006); doi: 10.1121/1.2346128
View online: http://dx.doi.org/10.1121/1.2346128
View Table of Contents: http://asa.scitation.org/toc/jas/120/5
Published by the Acoustical Society of America

Cluster analysis and robust use of full-field models for sonar beamforming

Brian Tracey,[a] Nigel Lee, and Srinivas Turaga[b]
MIT Lincoln Laboratory, 244 Wood Street, Lexington, MA 02420

(Received 17 February 2006; revised 1 August 2006; accepted 8 August 2006)

Multipath propagation in shallow water can lead to mismatch losses when single-path replicas are used for horizontal array beamforming. Matched field processing (MFP) seeks to remedy this by using full-field acoustic propagation models to predict the multipath arrival structure. Ideally, MFP can give source localization in range and depth as well as detection gains, but robustly estimating range and depth is difficult in practice. The approach described here seeks to collapse full-field replica outputs to bearing, which is robustly estimated, while retaining any signal gains provided by the full-field model. Cluster analysis is used to group together full-field replicas with similar responses. This yields a less redundant "sampled field" describing a set of representative multipath structures for each bearing. A detection algorithm is introduced that uses clustering to collapse beamformer outputs to bearing such that signal gains are retained while increases in the noise floor are minimized. Horizontal array data from SWELLEX-96 are used to demonstrate the detection benefits of sampled field as compared to single-path beamforming. © *2006 Acoustical Society of America.* [DOI: 10.1121/1.2346128]

PACS number(s): 43.30.Wi, 43.60.Mn, 43.60.Jn [EJS] Pages: 2635–2647

I. INTRODUCTION

Horizontal array beamforming for passive sonar is often based on the assumption that the signals impinging on the array arrive via a single propagation path. A model accounting for wave-front curvature across the array may be used near broadside, while a planewave model is used near endfire or for distant sources. However, multipath acoustic propagation is often important in shallow water. The use of a single-path model can therefore introduce a mismatch that reduces the output power of targets and may significantly limit the effectiveness of adaptive interference cancellation techniques.[1] The effects of multipath propagation are most noticeable at endfire, as the ability of horizontal arrays to resolve multipath is best for sources in this sector.[2] The approach described in this paper seeks to improve horizontal array performance near endfire in the presence of multipath.

Matched field processing (MFP) has been investigated as a means for accounting for multipath.[3] MFP uses full-field acoustic models to calculate replica vectors that represent the array response for a three-dimensional grid of possible source locations. Typically, MFP attempts to achieve simultaneous improvements over single-path beamformers in both detection performance (by using physics-based signal models to reduce signal mismatch) and source localization (by providing range and depth in addition to bearing).

Experimental use of MFP has shown that robust range and depth localization is difficult to achieve. In many cases this is due to uncertainty in the environmental inputs required by acoustic propagation models. Robustness can be increased by improving environmental calibration to give a more accurate set of replicas[4–6] or by explicitly accounting for uncertainty.[7–10] Another problem with MFP, especially for horizontal arrays, is that the array response may be very similar for sources at widely separated ranges and depths. Range and depth estimation therefore involves choosing the correct source location from a set of highly ambiguous beamformer outputs, which requires very accurate calibration and high signal-to-noise ratio.[11]

The signal processing approach described below is motivated by two observations. First, simulations described in Sec. III were used to investigate MFP localization accuracy for horizontal arrays when environmental information is uncertain. The results indicate that localization in bearing is much more tolerant of environmental uncertainty than is localization in range and depth. Second, experimental results in Sec. V demonstrate that use of MFP replicas can lead to reduced mismatch losses and therefore potentially better detection. These two observations lead to the following question: Is it possible to use MFP replicas in a way that achieves mismatch reduction, *without* attempting range and depth localization? Answering this question requires a method for collapsing the three-dimensional (3D) MFP beamformer output to bearing.

A simple approach is to maximize over MFP beamformer outputs for all ranges and depths along each bearing. This preserves any gains on signal provided by MFP, but leads to a significant increase in the displayed noise floor. As shown below, the resulting output shows little improvement in signal detectability.

An alternate procedure is developed here that includes a detection stage before collapsing the beamformer output to bearing. The procedure begins by analyzing the set of predicted MFP replicas and representing them by a set of *clusters*.[12] As described in Sec. II, each cluster corresponds to a set of spatial locations that have essentially the same

[a]Current address: Neurometrix, Inc. 62 Fourth St. Waltham, MA.
[b]Current address: Department of Brain and Cognitive Science, Massachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge MA.

multipath structure (i.e., replicas that are highly correlated with each other). In standard MFP these locations would appear as highly ambiguous regions in the output. The cluster analysis output is denoted here a "sampled field," as the number of clusters, or distinct multipath structures, is often much smaller than the number of possible three-dimensional source locations (from this point of view, the full set of MFP replicas can be considered an "over-sampled field," as it contains redundant information). Because the set of clusters captures the diversity of multipath predicted by the full-field model, any mismatch reductions provided by the full-field model are retained.

After cluster analysis is complete, beamforming is done using a representative replica vector from each cluster (this gives computational gains relative to MFP,[13] but computational gain is not the primary objective of the current work). The cluster outputs are then analyzed. If a cluster has higher output power than other clusters that are highly correlated with it, it is judged to be signal-like. Otherwise, the cluster is judged to represent either noise or sidelobe energy from a nearby source, and its output power is set to an estimate of the noise level. After this detection stage is complete, the data are collapsed to the bearing dimension. Used in this manner, clustering allows an automated analysis of ambiguous regions in the output so that decisions can be made as to which set of search locations may contain the signal. Similar strategies have been used in analysis of radar outputs.[14] Thus any signal gains in the proposed approach result from the use of full-field replicas, while clustering is used to collapse the beamformer outputs to a bearings-only display.

An alternative method for collapsing MFP output to a bearings-only display has been proposed by Yang and Yates.[15] Their work requires an iterative search over range and depth to refine estimates of source bearing, and therefore it implicitly requires 3D localization. By contrast, the sampled field approach presented here collapses full-field outputs to bearings only without requiring range and depth localization. In cases where the environmental knowledge is good, matched field range and depth localization could be attempted as a follow-on stage to the sampled field bearing localization described here.

Section II below describes cluster analysis in more detail, while Sec. III outlines simulations used to study localization errors due to environmental uncertainty. The cluster-based signal processing approach is detailed in Sec. IV and is labeled sampled field processing (SFP) to distinguish it from standard MFP. Section V analyzes data from the SWellEx-96 experiment, showing that the approach improves signal detectability at endfire as compared to planewave or MFP approaches, especially when adaptive beamforming is employed.

II. CLUSTER ANALYSIS OF FULL-FIELD REPLICAS

Full-field MFP replicas are based on the Green's function, or channel response, for a set of possible source locations in range, depth, and azimuth. The results below use a range-independent normal mode model to calculate the Green's function,[16] although clustering can be applied to the output of any ocean propagation model. Taking $(x,y)=(0,0)$ to be at the center of the array, with $z=0$ at the ocean surface, the Green's function for a sensor at $\mathbf{r}_i=(x_i,y_i,z_i)$ is given by

$$G(f,\mathbf{r}_i,\Theta)=\frac{ie^{-i\pi/4}}{\rho(z_s)\sqrt{8\pi}}\sum_{n=1}^{N}\psi_n(z_s)\psi_n(z_i)\frac{e^{-ik_n r_i}}{\sqrt{k_n}}, \tag{1}$$

where ψ and k_n are the mode shapes and mode horizontal wave numbers at the frequency f, z_s is the source depth, $r_i=\sqrt{(x_s-x_i)^2+(y_s-y_i)^2}$ is the source-receiver range, and $\Theta=(x_s,y_s,z_s)$ is the source location in three dimensions.

The full-field replica vectors are computed as normalized versions of the Green's function:

$$\mathbf{v}(f,\Theta)=\frac{\mathbf{G}}{|\mathbf{G}|}\sqrt{N}, \tag{2}$$

where $\mathbf{G}$ is a vector of calculated Green's functions for each receiver. By contrast, planewave or range-focued replicas are calculated using a single-path propagation model

$$\mathbf{v}_{\mathrm{pw}}(f,\Theta)=e^{-ik\mathbf{r}(\Theta)}, \tag{3}$$

where $k=\omega/c$ is the wave number for a wave traveling at speed c, and $\mathbf{r}(\Theta)$ is a vector of ranges between each array element and the hypothesized source location.

The set of MFP replica vectors can be analyzed using clustering, a technique that automatically breaks a set of observations into *clusters*, or groups of vectors that are similar to each other.[12] This similarity is defined by a distance metric specified by the user. For each cluster, a list of cluster members is returned along with a cluster center. The cluster center can be used as a representative of all the members of the cluster. The clustering algorithms used here require the user to specify a cluster *radius*, which is the maximum distance allowed between any point in a cluster and its center.

When applying clustering to beamforming, a natural distance metric is provided by the mismatch between pairs of replicas. The measure has been used by Cox[17] and is defined as

$$\mathrm{dist}(\mathbf{v}_i,\mathbf{v}_j)\triangleq 1-\frac{\|\mathbf{v}_i^H\cdot\mathbf{v}_j\|^2}{\|\mathbf{v}_i\|^2\|\mathbf{v}_j\|^2}, \tag{4}$$

where the frequency domain replica vectors are given by $\mathbf{v}_1,\mathbf{v}_2,\ldots,\mathbf{v}_n$. Replica pairs that are highly correlated will have low mismatch and will tend to be clustered together.

Using this distance metric, a cluster is defined as a set of replicas $\mathbf{v}_i$ that all have a distance from a cluster *center* that is less than some threshold η

$$\mathrm{dist}(\mathbf{v}_i,\tilde{\mathbf{v}}_j^{\mathrm{cen}})\leqslant\eta. \tag{5}$$

The cluster center $\tilde{\mathbf{v}}_j^{\mathrm{cen}}$ is chosen from the original set of full-field replicas.

Clustering of MFP replicas can be computationally intensive and requires the use of efficient algorithms. The method used here is a "greedy" algorithm that first computes a matrix containing the mismatch between every pair of vectors. The vector with the largest number of close matches (i.e., replicas with mismatch less than the threshold η) is picked as the first cluster center. All replica vectors closer

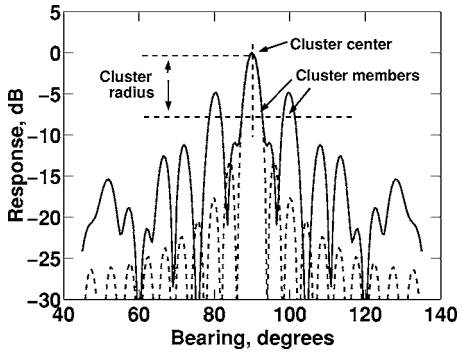

FIG. 1. Beam patterns for nominal well-sampled array (dashed line) vs sparse array (solid line). For a sparse array, steering directions that are highly correlated with the main axis of response may be widely separated in space. Clustering provides a way to identify these regions of similar response.

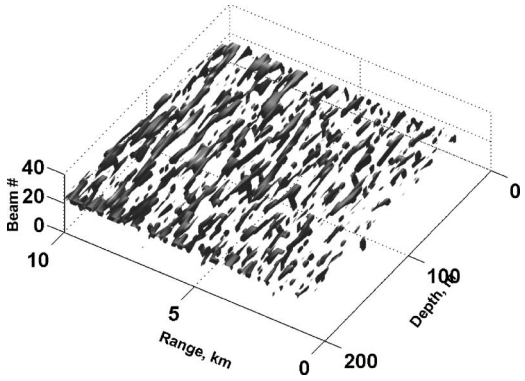

FIG. 2. Plot showing members of an example cluster formed from MFP replicas for a single-line array at 94 Hz. The cluster (with 1.5 dB cluster radius) is well contained in azimuth and does not spread over many azimuthal beams. For a given azimuth, high correlation between many range and depth locations is evident.

than the threshold are assigned to this center. The cluster center and its cluster members are removed from the full set of replicas, and the process is repeated until all replicas are clustered. A more complete description of clustering algorithms is given in the Appendix.

The application of clustering to beamforming can be made clearer by considering a planewave example. Figure 1 compares beampatterns for well-sampled and sparse line arrays steered to a bearing of 90°. The steering angle is equivalent to a cluster center. The beam extent is defined in terms of a maximum allowable loss with respect to the steering angle; this maximum allowable loss is the equivalent of the cluster radius. When high ambiguities are present, as occurs with sparse arrays, beams that are within the specified cluster radius may be spatially disjoint. Clustering provides an automated way to identify these disjoint regions of similar response.

The following simulation example shows how cluster analysis can be applied to MFP replicas. Simulations were performed for a 101 element, 200 m long horizontal line array situated at a depth of 120 m in the Santa Barbara Channel environment[18] (see Table I for details). Full-field MFP replicas were generated for a bearing sector around endfire, for depths across the full water column and for ranges up to 10 km.

A typical cluster generated from the set of MFP replicas is shown in Fig. 2. This plot shows the locations of all MFP replicas that are identified as belonging to the same cluster. The cluster radius in decibels ($10 \log_{10} \eta^2$) is 1.5 dB for this example, meaning that each replica included in the cluster has at most 1.5 dB mismatch with the cluster center.

From Fig. 2 one can observe some characteristics that are generally true of MFP replicas for horizontal arrays. First, bearing is much better resolved than range or depth. This is a result of the array geometry, which has good sampling of azimuthal arrivals but poor sampling of the multipath structure used to infer range and depth. Range and depth estimation therefore requires finding the peak of a set of points that are very highly correlated. Second, although the cluster shown in Fig. 2 extends over many range and depth cells, there are many gaps in its coverage along the bearing shown. Multiple replicas (or clusters) are needed to provide low-mismatch coverage for all source locations along a given azimuth.

The 1.5 dB cluster radius for Fig. 2 was chosen based on beam spacings often used in planewave beamforming. Decisions about how widely to space steering vectors are commonly made by setting a maximum allowable loss relative to the beam's steering angle. While beams can be overlapped at the −3 dB point (corresponding to a 3 dB cluster radius), they may be spaced more finely to reduce loss on targets that lie between beams.

TABLE I. Parameters for sound speed and geoacoustic model used in simulations: z=depth from surface; c_c =compressional sound speed; ρ=density; α_c=compressional wave attenuation. The water column extends to 209 m depth and is followed by a layered seabed.

z (m)	c_c (m/s)	ρ (g/cm^3)	α_c (dB/λ)
0	1500	1	0
25	1500	1	0
75	1494	1	0
209	1488	1	0
209	1607	1.95	0.37
309	1702	1.95	0.37
309	1862	1.98	0.035
609	2374	1.98	0.035
609	2374	2.03	0.04

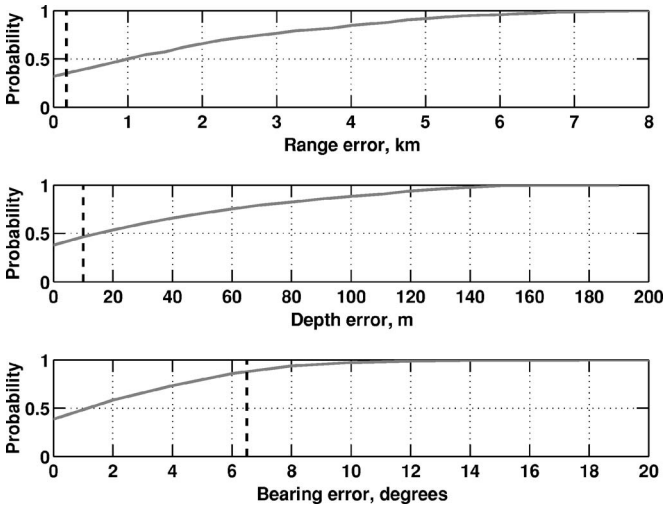

FIG. 3. Cumulative distribution functions for localization errors, single line HLA. Mainlobe widths for range, depth, and bearing are shown as dashed lines on each plot. Bearing errors are typically within the mainlobe, while median range and depth errors exceed the mainlobe size.

III. ARGUMENT FOR BEARINGS-ONLY LOCALIZATION

As will be discussed in Sec. IV, clustering enables efficient collapsing of MFP outputs to bearings-only localization. This section motivates why bearings-only localization using MFP replicas is of interest. Ambiguity surface patterns such as those shown in Fig. 2 suggest that bearing estimates found using MFP-derived replicas should be more robust than range or depth estimates. Monte Carlo studies of the effects of environmental mismatch are used here to verify this intuition.

Simulations were carried out for the line array and environment described above. Simulated data was generated for 200 random source locations, all at endfire bearing to the array but with ranges and depths uniformly distributed over the search space (0–200 m depth and 2–10 km range). No noise was added. MFP replica vectors were generated using mismatched environmental models, in which the initial compressional speed in the sediment was perturbed. Perturbations of ±10 and ±20 m/s were used for a series of runs whose results were accumulated. These mismatched replicas were used in beamforming. The 3D location of the peak beamformer output was found for each random source location.

Figure 3 shows the cumulative distribution function (CDF) of MFP localization errors caused by environmental uncertainty. This result shows that bearing errors were generally low, with the median error being roughly 1/3 of the beamwidth. In contrast, the median error in the range and depth dimensions exceeded the estimated mainlobe resolutions in these dimensions, which were estimated as described in.[18] This underscores the difficulty of the 3D localization problem.

Signal-to-noise ratio (SNR) is another important consideration for MFP localization.[11] As suggested by Fig. 2, many locations in the range/depth plane have very similar response. Estimating the correct range and depth from this flat parameter space therefore requires higher SNR than estimating bearing only.

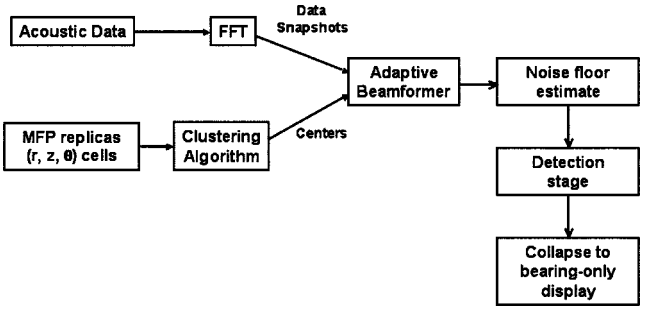

FIG. 4. Proposed signal processing flow for sampled field processing.

IV. USE OF CLUSTERS IN PROCESSING

Figure 4 shows a proposed signal processing flow that is motivated by the two previous sections. Those sections introduced the ideas of using clustering (to gain a nonredundant representation of the set of MFP replicas), and of bearings-only localization (to reduce sensitivity to environmental uncertainty and low SNR). The processing chain in Fig. 4 demonstrates how these two ideas can be combined. The steps in processing are as follows.

A. Beamforming stages

The five left-most boxes in Fig. 4 show steps related to beamforming. The steps involved in calculating and clustering the MFP replicas were described in Sec. II. Since the replicas depend on the environment and the array geometry, they need to be recalculated whenever either of these change. Clusters are generated for all frequencies of interest.

The acoustic data collected on the N array elements are transformed into the frequency domain using a short-time Fourier transform. This gives a set of $N\times 1$ data snapshots versus time and frequency. The snapshots are denoted $\mathbf{x}(f_0,t_0)$, where f_0 is the frequency and t_0 is the time of the snapshot. The data shown below were processed using a 2 s Hanning-windowed fast Fourier transform (FFT) with 50% overlap.

The nonadaptive or *conventional* beamformer (CBF) uses a weight vector that is simply a scaled version of the steering vector $\mathbf{w}_c(\cdot)=\mathbf{v}(\cdot)/N$. This normalization ensures that there is unity gain in the look direction. The conventional beamformer output for the direction Θ is given by

$$P_c(f_0,t_0,\Theta)=|\mathbf{w}_\mathbf{c}^H(f_0,\Theta)\mathbf{x}(f_0,t_0)|^2. \quad (6)$$

For simplicity, the dependence on the frequency f_0 and snapshot time t_0 is omitted below.

Adaptive beamforming (ABF) is implemented for the data analysis below using the dominant mode rejection (DMR) algorithm,[19] with a white noise gain constraint added to increase robustness to signal mismatch. DMR is a variant of the minimum variance distortionless response (MVDR) algorithm with diagonal loading. The loaded MVDR weight vector is given by

$$\mathbf{w}_d(\Theta) = \frac{(\hat{\mathbf{R}} + \delta(\Theta)\mathbf{I})^{-1}\mathbf{v}(\Theta)}{\mathbf{v}^H(\Theta)(\hat{\mathbf{R}} + \delta(\Theta)\mathbf{I})^{-1}\mathbf{v}(\Theta)}, \tag{7}$$

where Θ is the beam steering direction, $\delta(\Theta)$ is the direction-dependent diagonal load level, and $\hat{\mathbf{R}}$ is the estimated data covariance matrix. $\hat{\mathbf{R}}$ can be estimated from the data using a time-averaged sample covariance matrix of frequency-domain snapshots

$$\hat{\mathbf{R}} = \frac{1}{L}\mathbf{X}\mathbf{X}^H, \tag{8}$$

where $\mathbf{X}$ is a $N \times L$ data matrix, formed by accumulating L frequency-domain snapshots derived from short-time FFT data on the N sensors. Full-rank processing would require averaging over a long time period ($L \geq N$), during which motion effects may become significant. Therefore, DMR focuses on the D largest eigenvalues (and corresponding eigenvectors) and deemphasizes the remaining N-D eigenvalues/eigenvectors by constructing a modified sample covariance matrix $\tilde{\mathbf{R}}$

$$\tilde{\mathbf{R}} = \sum_{k=1}^{D} \lambda_k \mathbf{u}_k \mathbf{u}_k^H + \alpha \sum_{k=D+1}^{N} \mathbf{u}_k \mathbf{u}_k^H, \tag{9}$$

where α is an estimate of the input noise floor. $\tilde{\mathbf{R}}$ is then substituted for $\hat{\mathbf{R}}$ in Eq. (7) to calculate the adaptive weights $\mathbf{w}_d$. These weights are used in Eq. (6) (in place of $\mathbf{w}_c$) to find the adaptive beamformer output power P_a.

The load level δ in Eq. (7) is increased via an iterative process until it satisfies a white noise gain constraint[20]

$$\mathbf{w}_m^H(\cdot)\mathbf{w}_m(\cdot) \leq \frac{\beta}{N}. \tag{10}$$

The constant β above is a "relaxation" parameter, defined as the factor by which the weight norm in Eq. (10) is allowed to to exceed the "white noise gain" of $1/N$ (which is the value of the weight norm without diagonal loading and with white noise as input).

In passive processing applications, such as considered here, data from the signal of interest cannot be excluded when estimating the covariance matrix $\tilde{\mathbf{R}}$. An adaptive beamformer will (mistakenly) interpret any mismatched[21] portions of the signal of interest as noise and will (mistakenly) adapt to and reject this portion of the signal itself, an effect termed *signal self-cancellation* or *signal self-nulling*.[22] Diagonal loading via the white noise gain constraint provides a means for limiting this signal self-cancellation.[22] The white noise gain parameter β in Eq. (10) was set to 2 (3 dB) for the results below, yielding an adaptive processor that is relatively tolerant to signal model mismatch.

B. Noise estimation

After beamforming, the output powers are analyzed to estimate the noise level as a function of bearing. The harmonic mean noise estimator is used here.[23] The noise as a function of bearing θ is estimated as

$$P_N(\theta) = \left(\sum_{k \in \mathcal{K}} P(k)^{-1} \right)^{-1} \tag{11}$$

where $P(k)$ is the beamformer output power (either conventional or adaptive) for the kth cluster, and $\mathcal{K}$ is the set of clusters whose cluster center lies along a bearing within $\pm M$ azimuthal beams of θ. For data processing shown below, $M=7$. The harmonic mean acts to suppress high-power outputs due to signals so that an accurate noise estimate is obtained. A variety of noise estimation methods are available, and the choice of the harmonic mean estimator is not critical for cluster-based processing.

C. Collapsing to bearing

Two possible methods for collapsing the beamformer outputs to bearing are considered here.

1. Maximize over all outputs along each bearing (ORing)

The simplest way to collapse MFP outputs to bearing only is to maximize over all matched field outputs along each bearing:

$$P_{1-D,mfp}(\theta) = \max_{(r,z)} P(r,z,\theta). \tag{12}$$

A similar collapsing can be done for a cluster-based beamformer. In this case the maximization is over the set $\mathcal{I}$ of clusters whose centers lie on the given bearing

$$P_{\text{SFP-OR}}(\theta) = \max_{i \in \mathcal{I}} P_i(\theta). \tag{13}$$

Because clustering captures the multipath structures described by the MFP replicas, the outputs of Eqs. (12) and (13) are expected to be very similar.

Signal power is preserved when using maximization to collapse beamformer output to bearing only. Unfortunately, maximization has the effect of increasing the noise level. As a thought experiment, imagine that all outputs along a given bearing correspond to noise. Maximizing the outputs gives a result that corresponds to the highest observed noise value, which may be significantly higher than the average noise level. As a result, a signal on an adjacent bearing will have a lower observed signal-to-background ratio. Maximization also tends to emphasize sidelobe energy from nearby sources, which again reduces the signal-to-background ratio.

The operation shown in Eqs. (12) and (13) has been labeled "ORing" in the sonar literature, and the SNR losses it causes are known as "ORing loss."[24,25] This terminology is adopted below. Data results generated using Eq. (12) will be labeled "ORed MFP," while cluster-based results generated using Eq. (13) will be denoted "SFP-OR."

2. Cluster-based detection approach

Clustering provides additional information about the relationship among points in the search space that is not available in standard MFP. The distance in terms of mismatch between cluster centers can be calculated, giving a measure

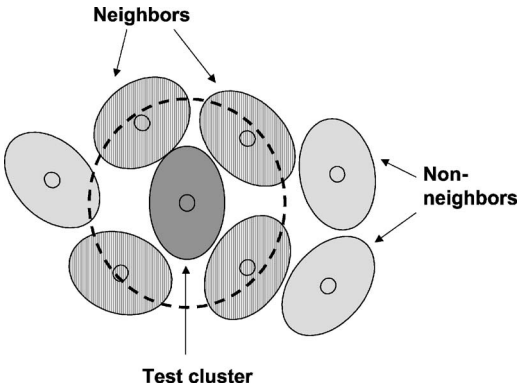

FIG. 5. Cartoon showing the proposed peak detection stage. Each cluster is compared to its "neighbors," defined as other clusters that are highly correlated with it. If the test cluster is a local peak or close to the local peak, its output is retained; if not, the cluster power is set to the estimated noise floor. The full set of clusters is then collapsed to bearings only.

of how different regions in the search space are related. A relatively simple method for using this information is described here.

One way to separate signals from noise is to require that a signal should give the peak output power in its local neighborhood. If a test cluster has output power P_{tst}, this cluster may be considered signal-like if

$$P_{\text{tst}} > P_j \quad \forall\, j \in \mathcal{N}_{\text{NB}}, \tag{14}$$

where j is the set of cluster centers in the neighborhood $\mathcal{N}_{\text{NB}}$. While noise outputs will occasionally give local peaks in output power, the majority of noise outputs will not. Noise-generated peaks will be spaced randomly in bearing and will not generate apparent source tracks over time. Sidelobe energy will also be suppressed using this test.

An approach to defining the local neighborhood $\mathcal{N}_{\text{NB}}$ is shown in Fig. 5. The neighborhood is defined as the set of all clusters whose centers are mismatched from the test cluster's center by less than some threshold

$$\text{dist}(\mathbf{v}_{\text{cen}}^{\text{tst}}, \mathbf{v}_{\text{cen}}^{j}) <= \eta_{\text{NB}}, \tag{15}$$

where distance is as defined in Eq. (5). Good results on data have been obtained by setting η_{NB} in the range of 2–3 dB, for cluster radii ranging between 0.5 and 1.5 dB. Typically the cluster neighbors identified as above extend over several adjacent bearings, and include some but not all of the clusters for each bearing.

In practice, a modification to this approach has been found to be helpful. Neighboring clusters that are not locally maximal [as defined in Eq. (14)] but have output power within 1 dB of the local maximum are retained. This modification helps in the situation where a source has energy that is split between two clusters.

The results of this detection stage are collected in a vector δ, whose entries are "1" for clusters that are found to be signal-like and "0" for clusters that are not. The full set of clusters is collapsed to bearing as

$$P_{\text{SFP-DET}}(\theta) = \max_{i \in \mathcal{I}}[P_i(\theta) \cdot \delta_i] + P_N(\theta)[1 - \max_{i \in \mathcal{I}}(\delta_i)], \tag{16}$$

where $\mathcal{I}$ is the set of clustered outputs whose cluster centers lie along θ. Thus, if there are one or more signal-like clusters along the bearing θ, the power output is set to the largest of them. If no signal-like clusters are identified, the power is set to the estimated noise floor.

A summary of the proposed processing method is as follows. Note that the first two steps below are also part of generating the ORed output described in Eq. (13) (SFP-OR).

- The full set of MFP replicas (steering vectors) is clustered, and the cluster centers are used as sampled field replicas.
- The replicas are used in beamforming. Cluster output powers and replica vectors are stored.
- An estimate of the noise floor versus bearing is made using the full set of cluster outputs.
- Each cluster's output power is examined in turn. If this power is higher than the power on any nearby cluster, the cluster is declared to be signal-like. Any neighbors with output powers that are within 1 dB of the peak are also labeled as signal-like.
- A maximization is used to collapse all signal-like clusters on each bearing to a bearings-only display. If no clusters along an azimuth have been labeled as signal-like, the output is set equal to the estimated noise floor.

Data results generated using this approach will be labeled "SFP-DET."

V. RESULTS USING EXPERIMENTAL DATA

Planewave, MFP, and sampled field approaches were tested using data from the SWellEx-96 ocean acoustics experiment (www.mpl.ucsd.edu/swellex96). Data were processed for event S5, in which the towship R/V Sproul towed both deep and shallow sources. These sources each projected a sequence of tones of varying levels. Results are shown for the bottom-mounted HLA North array, an irregularly spaced array of roughly 240 m overall length (see Fig. 6).

The sparse sampling provided by HLA North resulted in increased azimuthal sidelobes. Fig. 6(b) shows the calculated planewave response of HLA North to an endfire source at 148 Hz, one of the tones transmitted by the Sproul. As compared to a fully populated array, HLA North displays significantly elevated sidelobe levels.

The data from event S5 were beamformed using single-path (planewave and range-focused) replicas, MFP replicas, and cluster centers. Identical beamforming parameters were used for all replica types. The data were processed using 2 s FFT windows with 50% overlap. No spatial taper was applied to the array. Beams were linearly spaced in wave number, with 89 beams ranging from −70° to 70°.

Adaptive beamforming (ABF) was implemented using the DMR algorithm described above. As noted, the white noise gain parameter was set to 3 dB to give tolerance against mismatch. In the beamforming, 15 FFT snapshots were averaged in time to estimate the eigenvectors and ei-

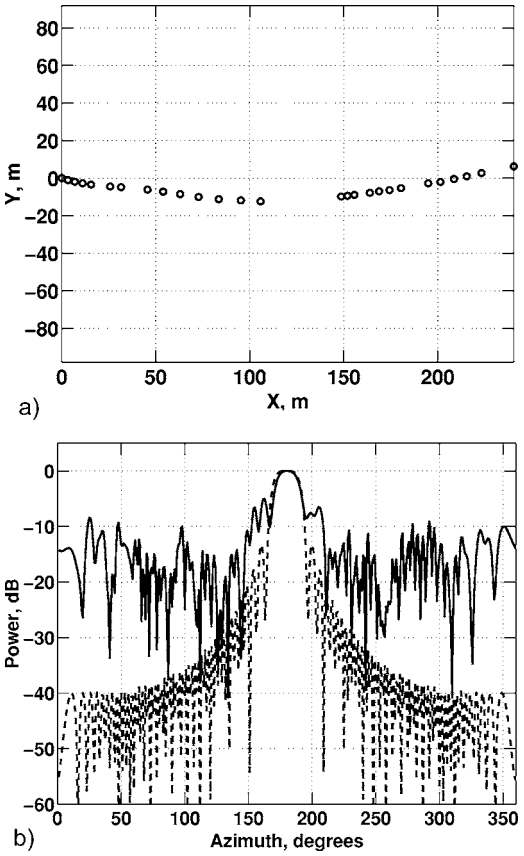

FIG. 6. Array geometry and beam patterns for bottom-mounted HLA North array from SWellEx-96. The sparse sampling causes elevates sidelobes (solid line) as compared to those expected for a fully populated array of the same length (dashed line). (a) Hydrophone locations and (b) beam patterns at 148 GHz.

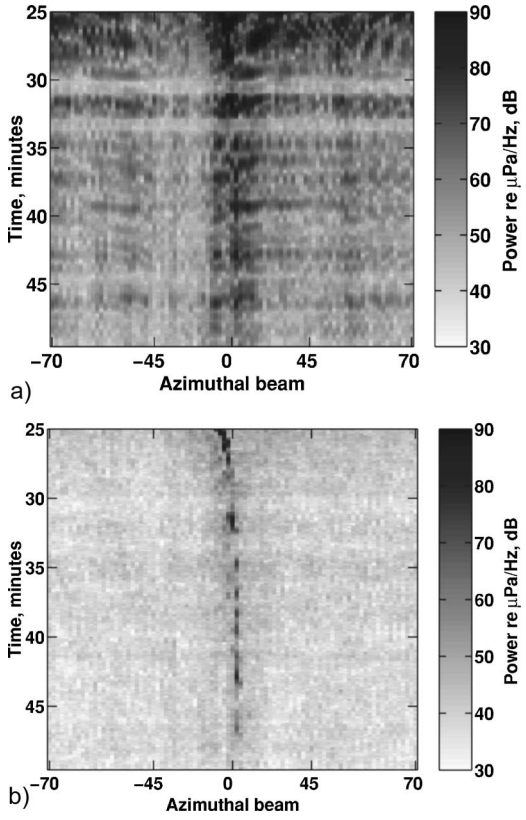

FIG. 7. SWellEx-96 bearing-time records (BTRs), 235 Hz projected tone, for event S5. Cosine-spaced planewave beams are used for beamforming. The R/V Sproul was at endfire to the array during the times shown and transmitted tones continuously throughout the event. Note that adaptive beamforming suppresses sidelobe energy at the cost of signal self-nulling. (a) Conventional planewave and (b) adaptive planewave.

genvalues corresponding to the $D=5$ strongest sources in the data. The noise floor for the adaptive beamformer was found by averaging the remaining eigenvalues ($D+1$ through N, where N is the number of array elements).

A. Example planewave and ORed full-field results

This section gives example results for planewave and ORed full-field (MFP and clustered) beamformers. The examples help to illustrate the differences between the beamformers and serve as a preface to the more detailed data analysis in Secs. V B–V D.

Figure 7 shows planewave output versus bearing and time for the 234.5 Hz frequency bin, which contains most of the energy for the 235 Hz tone projected by the deep source. The towed source is at endfire, or 0°. The planewave conventional beamforming result, seen in Fig. 7(a), shows a large amount of energy in grating lobes caused by the array sparsity. Planewave adaptive beamforming, shown in Fig. 7(b), largely rejects this grating lobe energy.

The adaptive beamformer also gives much lower energy on the target track. Since both conventional and adaptive beamformers are normalized to give unit gain on target, one would expect equal power outputs if no mismatch were present. The reduction in along-track target power therefore represents adaptive self-nulling caused by mismatch between the received signal and the assumed (planewave) signal model.

Figure 8 shows results of applying ORed full-field replicas to the same data. Figure 8(a) shows ORed MFP output [collapsed to bearing using Eq. (12)] while Fig. 8(b) shows SFP-OR output [collapsed to bearing using Eq. (13)]. A cluster radius of 1 dB was used for SFP-OR results. The ORed MFP and clustered beamformer results are very similar. The clustered output shows slightly more target self-nulling due to the fact that clustering uses a coarser sampling of the information contained in the MFP replica field.

A comparison of the towed source track (at endfire, or 0°) in Figs. 7 and 8 shows that use of full-field replicas gives significantly reduced adaptive self-nulling as compared to direct-path beamforming. Figure 8 also shows an increase in the displayed noise floor, or background, as compared with planewave output. This ORing loss reduces the overall signal-to-background ratio and motivates the use of alternate collapsing strategies.

In Figs. 7 and 8 a second apparent track is seen, initially to the right of the main track but moving quickly to the left. The source of this energy is unclear, and the source is not seen in other frequency bins. The main point of interest in Figs. 7 and 8 is the reduced adaptive self-nulling seen on the actual endfire target when full-field replicas are used.

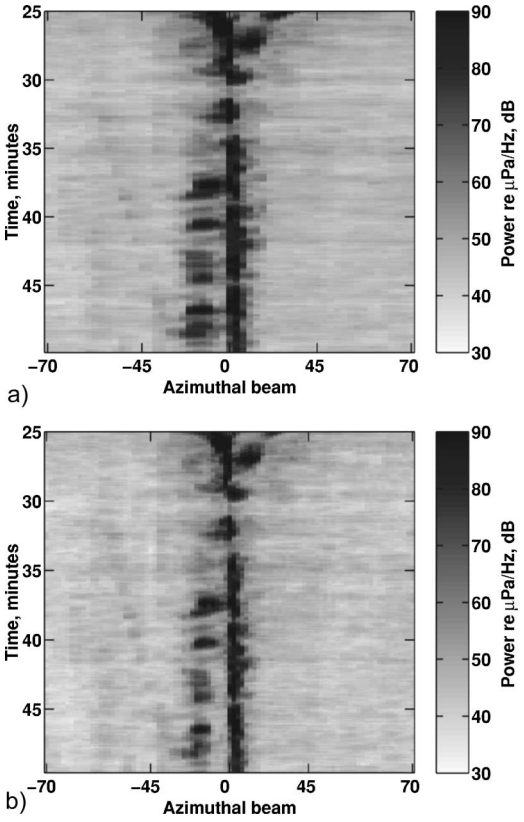

FIG. 8. ORed adaptive MFP and SFP-OR results for the 235 Hz projected tone, collapsed to bearing using Eqs. (12) and (13). Note the increases in both signal and noise levels relative to adaptive planewave results. Little difference exists between sampled field and MFP outputs when this collapsing strategy is used. (a) Adaptive MFP, ORed to bearing and (b) adaptive SFP-OR.

B. Measured mismatch reduction

This section quantifies the reduction in mismatch obtainable using clustered full-field replicas. Mismatch reduction was measured experimentally by using high-SNR tones broadcast by the R/V Sproul as calibration signals. By comparing the maximum beamformer output to the power of the signal received on the array, beamformer mismatch or signal gain degradation (SGD) can be measured as

$$\mathrm{SGD} = \max_i \frac{\|\mathbf{v}_i^H \cdot \mathbf{x}\|^2}{\|\mathbf{v}_i\|^2 \|\mathbf{x}\|^2}, \tag{17}$$

where $\mathbf{v}_i$ are the set of replica vectors and $\mathbf{x}$ is the data, assumed to come from a single source. To verify that a single source dominated the data (and thus the mismatch could be accurately computed), a singular value decomposition was calculated versus time for a covariance matrix estimated by averaging four snapshots. The ratio of the first to second eigenvalues was computed, and SGD estimates were judged to be valid if this ratio exceeded 6 dB.

Figure 9 shows computed SGD versus time for the 234.5 Hz bin. SGD values are shown for planewave replicas and for clustered beamformer outputs computed using different cluster radii. A noticeable mismatch reduction is seen when moving from planewave to clustered replicas. This mismatch reduction is equal to the gain in output power expected for conventional beamforming.

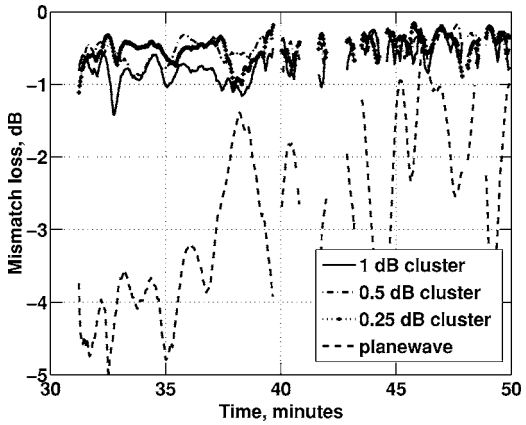

FIG. 9. Measured signal gain degradation vs time for 231.5 Hz frequency bin. Results are shown for planewave replicas and clustered full-field replicas. Results for cluster sizes of 1, 0.5, and 0.25 dB are shown. The clustered full-field replicas provide significant mismatch reductions as compared to planewave replicas.

Some mismatch reduction is observed when changing from a 1 dB cluster radius to a 0.5 dB cluster radius, while little additional reduction is seen when changing from 0.5 to 0.25 dB radii. This indicates that the inputs to the beamformer do not support a more detailed sampling of the pressure field. The remaining uncertainty is likely a combination of errors in estimated array positions and errors in the environmental modeling. This is an example of how detection gains are limited by environmental uncertainty.

Signal gain degradation values were calculated for four strong tones. Normalized histograms of the valid SGD values were generated for planewave and clustered replicas, using a 0.5 dB cluster radius. The resulting estimated probability density functions are shown in Fig. 10. The mismatch grows with frequency, as expected, but is noticeably lower for the clustered replicas. For the lowest three tones, the clustered replicas nearly always have 1.5 dB or less mismatch, while the planewave replicas can experience 2–4 dB mismatch. This dB difference can be greatly amplified in adaptive processing. Reducing mismatch by even a few decibels can dra-

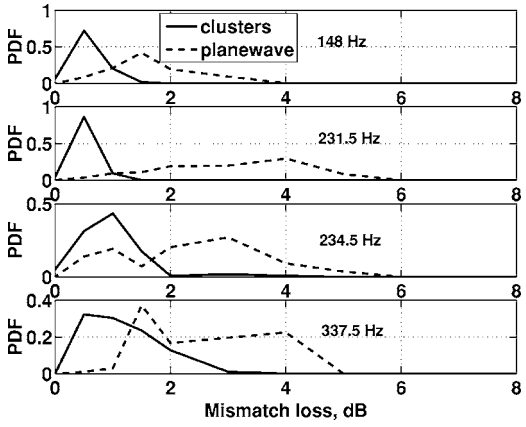

FIG. 10. Probability density of measured signal gain degradation as a function of frequency. Mismatch losses using MFP-derived clusters are always less than 3 dB, while planewave mismatch can be considerably higher.

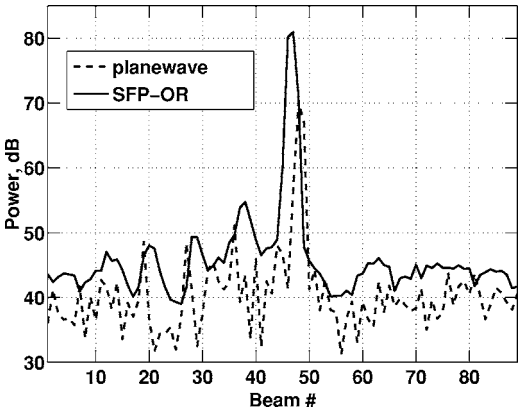

FIG. 11. Beamformer outputs at a single snapshot, comparing planewave adaptive beamfomer output to adaptive sampled field processing with simple ORing used to collapse to bearings only (SFP-OR). Mismatch reduction leads to less self-nulling and more signal power in the SFP-OR output, but ORing increases the displayed background level.

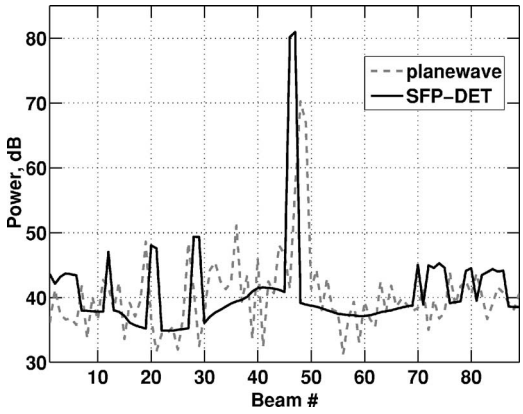

FIG. 13. Sampled field outputs using cluster-based detection (SFP-DET) compared to planewave ABF. For bearings that do not contain an identified peak in sampled field outputs, the SFP-DET output is set to the estimated noise floor for all clusters. The resulting display maintains signal gains due to mismatch reduction but minimizes the increase in the noise floor.

matically reduce adaptive self-nulling for strong sources, with smaller but still significant effects seen for weaker sources.

C. Comparison of collapsing options in data

As discussed above, collapsing MFP or clustered outputs to bearing using a simple maximization (ORing) leads to an increase in the apparent noise floor. Figure 11 illustrates this problem in more detail. Outputs of the planewave and ORed sampled field beamformers are shown versus bearing at a single snapshot in time. A cluster radius of 1 dB was used for this and all subsequent SFP results. The mismatch reduction obtained by SFP roughly 15 dB signal gain in this snapshot due to reduced target self-nulling (note that a weaker source would suffer less self-nulling loss). However, a noticeable increase in the quiet beams (away from the source) is also seen. This increase is referred to as *ORing loss*.[24]

Figure 12 compares noise floors estimated from planewave outputs, from the ORed sampled field outputs (SFP-OR), and from the *full* set of cluster outputs (SFP-DET noise floor, found using Eq. 11). This last noise estimate is used in Eq. (16) to generate the SFP-DET outputs. Visually, the SFP-DET noise floor estimated from the full set of clusters appears close to the planewave noise level. The average noise floor for the planewave curve is 39.7 dB, while the average noise floor for SFP-DET is 38.1 dB; the small difference is probably due to spillover of signal energy into the planewave noise floor estimate. In contrast, the average noise estimate for the SFP-OR outputs is 45.2 or 5.5 dB higher than the planewave output. This increase is due to the ORing operation of Eq. (13). Figure 12 suggests that the SFP-DET algorithm has the potential to produce a similar noise floor to that of planewave beamforming.

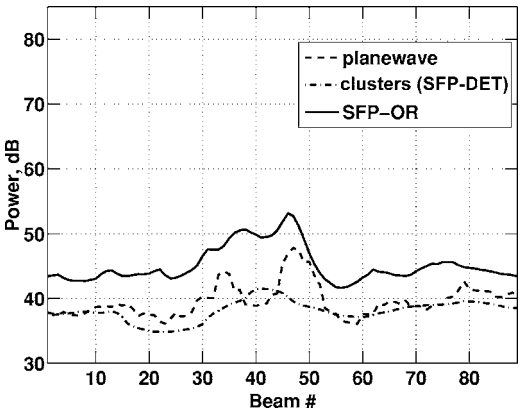

FIG. 12. Estimated noise floors for planewave adaptive beamforming, all clusters (SFP-DET noise floor), and SFP-OR results. The SFP-DET noise floor is similar to that of planewave beamforming, but ORing leads to a 5.5 dB increase in the average noise floor of SFP-OR.

Figure 13 shows the result of applying the SFP-DET approach. The planewave ABF result is shown as a dotted line, while the solid line shows the SFP-DET cluster result. For bearings where no cluster peak was found, the SFP-DET curve shows the underlying noise floor estimate plotted in Fig. 12. The SFP-DET result retains gains on the signal (which is at 0°) while avoiding a large increase in the noise floor. A number of small peaks are seen in the SFP-DET output, which may correspond to either local noise peaks or weak signals. The local SNR for these detections is of the same order of magnitude as similar noiselike peaks seen in the planewave output. An important note is that peaks generated by noise fluctuations will vary randomly, and are not expected to generate consistent tracks over time.

D. Single-tone bearing-time records

Planewave and sampled field beamformers can be usefully compared by forming bearing versus time output plots (BTR's). Results are presented for here both high-SNR and lower-SNR projected tones. Adaptive beamforming (ABF) is used in all cases. Because adaptive self-nulling increases with source level, the gains from mismatch reduction when using SFP-DET are most striking in the loud tone data. However, significant improvements are also seen for the quieter tones.

Figure 14 shows single-frequency beamformer output for the 231 Hz shallow source projected tone. This tone and the 235 Hz result shown above were among the loudest tones broadcast, with source levels of 158 dB re μPa. Planewave

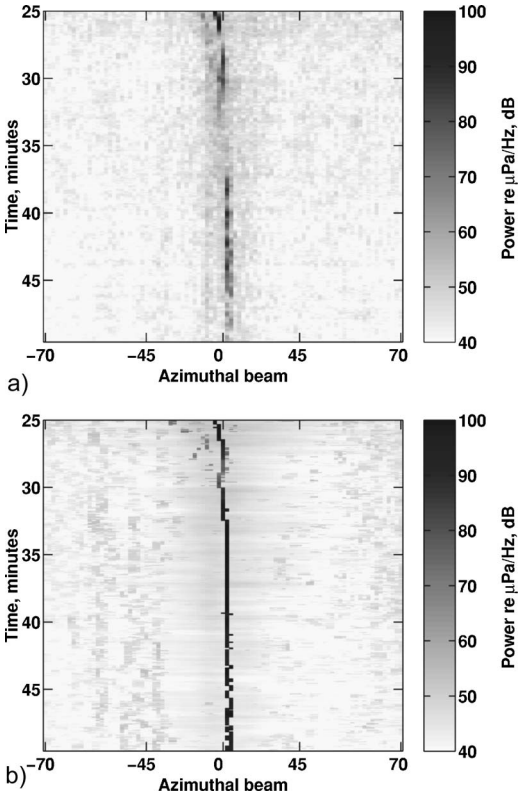

FIG. 14. SWellEx-96 bearing-time records (BTRs), 231 Hz projected tone, for time period with R/V Sproul near aft endfire (center of display). This tone was projected at 158 dB from the shallow source, which transited from roughly 2–4.5 km over the time period shown. Planewave beamforming is compared to SFP-DET sampled field results. Best results are seen with SFP-DET. (a) adaptive planewave and (b) adaptive SFP-DET.

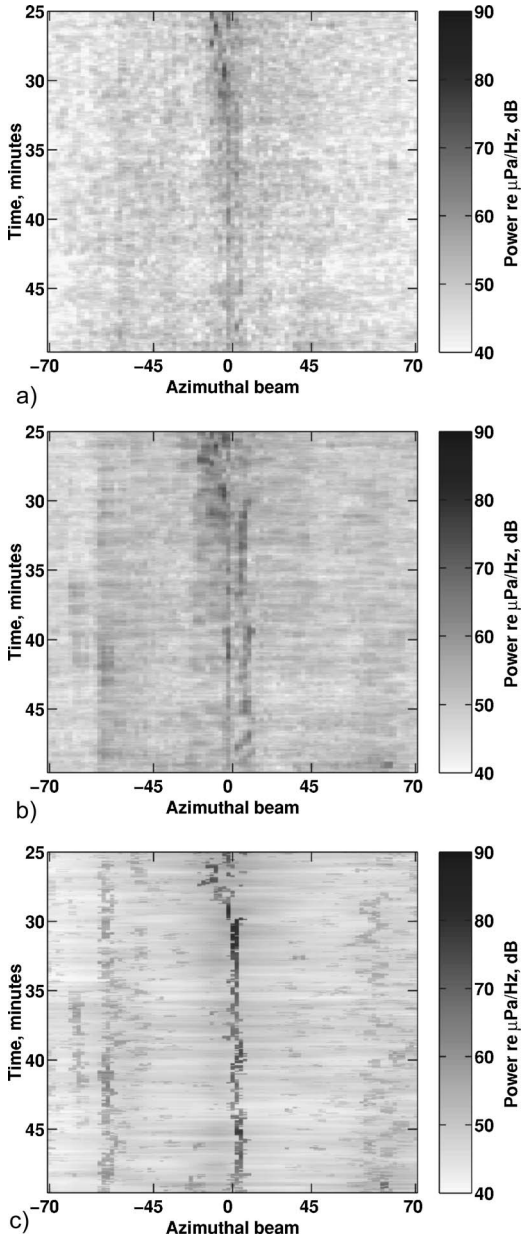

FIG. 15. Bearing-time records (BTRs) at 140 Hz. This tone appears to be an naturally radiated tone from the R/V Sproul. The SFP-DET sampled field result shows signal gains but maintains a low noise floor. The SFP-DET result performs better than either adaptive planewave or adaptive range-focused beamformers. (a) adaptive planewave, (b) adaptive range-focused, ORed to bearing, and (c) adaptive SFP-DET.

results are shown in Fig. 14(a), while SFP-DET results [generated using Eq. (16)] are shown in Fig. 14(b). The high SNR towed source is visible near 0° in both displays, but the power output along the target track is much lower for the planewave beamformer than for the SFP-DET beamformer during several time periods. Since both beamformers are normalized to give unit gain on target, the difference in power output is attributable to adaptive self-nulling in the planewave result. The SFP-DET approach is successful in reducing the signal model mismatch so that the adaptive beamformer suppresses sidelobe energy without suppressing the signal of interest. In addition, the SFP-DET algorithm has a noise floor comparable to the planewave display, resulting in a clear improvement in signal-to-background ratio (SBR).

A similar result is seen in Fig. 15. This figure plots data from the 140 Hz FFT bin, which apparently contained a naturally radiated tone from the R/V Sproul (again near 0°). Note that these results are plotted on a 50 dB dynamic range, while plots above required larger dynamic range to accomodate the extremely high SNR. While the source level for this tone is unknown, it appears to lie in between the "loud" (158 dB) and "quiet" (132 dB) projected tones. Because this tone is not as strong as the 231 Hz projected tone, the reduction in self-nulling obtained by sampled field processing is not as extreme, though an improvement is clearly noticeable.

In Fig. 15, an additional subfigure (b) is added that shows a bearing-time record formed by maximizing over range-focused direct path beamformer output, using five focus ranges uniformly spaced between 2 and 10 km. This covers the same ranges as included in the full-field replicas. This plot was generated to demonstrate that the signal gains observed in SFP-DET outputs are due to the use of full-field replicas, rather the use of range focused replicas. Wavefront curvature is not present for endfire sources, where the towed source is located, so accounting for source range is not expected to affect gain on the towed source (although it does affect off-endfire sources).

Figure 16 shows results for a quieter (132 dB re μPa) projected tone. A number of off-endfire contacts are seen in the ORed range-focused and SFP results that are not seen in

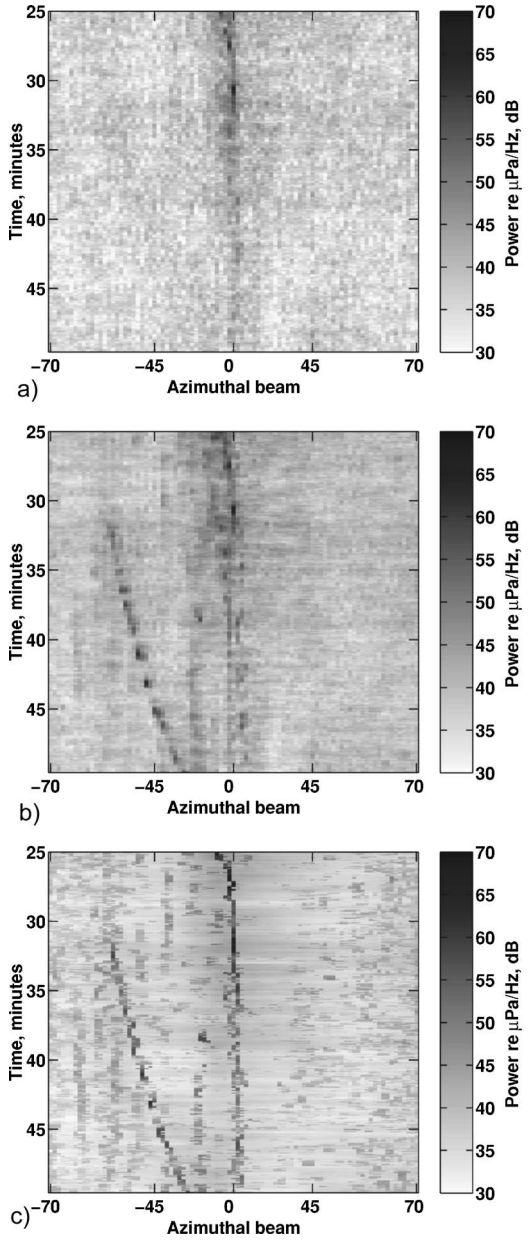

FIG. 16. Bearing-time records (BTRs) for the 237 Hz projected tone. This tone was projected from the source at 132 dB (one of the quieter projected tones). Several additional contacts are seen in the SFP-DET and range-focused outputs that are not seen in the planewave result, and are presumably due to closer-range sources. Again, the SFP-DET shows improved signal-to-background ratio on both the projected tones and observed contacts. (a) adaptive planewave, (b) adaptive range-focused, ORed to bearing, and (c) adaptive SFP-DET.

the planewave results. These are likely closer-range sources whose energy is protected when closer focus ranges are added. Results for the projected tone (near 0°) again show a noticeable signal-to-background gain for SFP-DET as compared to planewave or range-focused results.

This gain can be quantified in terms of the measured along-track signal-to-background ratio (SBR). Figure 17 shows estimated SBR for the ABF outputs from Fig. 16. The background level for each bearing of interest was estimated by averaging the output for three beams on either side, with an offset of two beams (thus the background for beam 30 is the mean of beams 25–27 and 33–35). SBR estimates were

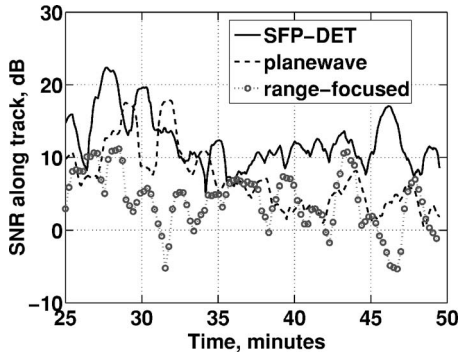

FIG. 17. Estimated along-track signal-to-background ratios (SBR), 237 Hz projected tone. Estimates are shown for adaptive SFP-DET, planewave, and ORed range-focused outputs. The median SFP-DET along-track SBR exceeds the median planewave SBR by 5.4 dB.

extracted along the observed signal track and smoothed using a 60 s moving average. The ORed range-focused result has the worst along-track SBR, presumably because the ORing operation causes an increase in background level. The SFP-DET approach has improved SBR as compared to planewave beamforming for almost all time intervals, giving an increase in median along-track signal-to-background ratio of roughly 5.4 dB. Higher gains would be expected if along-track SBR were calculated from Figs. 14 and 15, as the effects of adaptive self-nulling on planewave outputs are greater for these louder tones.

VI. SUMMARY

This paper has presented sampled field processing (SFP) as an alternative to matched field processing for exploiting the information about multipath structure provided by acoustic modeling. Unlike MFP, which attempts to simultaneously provide both detection gains and three-dimensional localization, SFP only seeks to provide detection gains. SFP uses cluster analysis to sample the predicted acoustic fields, resulting in a set of beamformer replicas that characterize typical multipath interference structures expected along each bearing. Because horizontal arrays can be used to estimate source bearing more robustly than range or depth, sensitivity to environmental uncertainty can be reduced by collapsing beamformer outputs to a bearings-only display. The information provided by cluster analysis is used in a detection scheme for collapsing beamformer outputs to bearing while preserving signal-to-background gains.

Gains resulting from application of this detection scheme (termed SFP-DET) were demonstrated using SWellEx-96 bottom-mounted horizontal line array data for an endfire source. Compared to adaptive planewave and adaptive range-focused beamformers, SFP-DET achieved better signal gain, because of its use of a full-field signal model. Compared to beamformers that collapse to bearings only by maximization (such as ORed MFP), the cluster-based SFP-DET approach achieved lower background levels. Overall, SFP-DET produced the best signal-to-background ratios of any of the beamformers considered.

The results shown here are for narrow-band (single-frequency) processing. Additional gains can be expected by

broadband processing, which combines signal energy across frequency. A simple form of broadband processing could be accomplished by incoherently averaging narrow-band BTRs such as those shown above. However, better performance would likely be achieved if the SFP-DET algorithm were modified to explicitly use broadband information. For example, a broadband version of Eq. (14) could be created which requires that a test cluster be labeled "signal-like" only if similar signal-like clusters are seen on the same bearing at other frequencies. Broadband processing using cluster analysis remains a topic for future work.

The work presented here has used cluster analysis as a preprocessing stage before beamforming is carried out. Alternatively, one could beamform using all MFP replicas, cluster the replicas corresponding to spatial peaks in the output to identify closely related peaks, and use a detection strategy to collapse these outputs to bearing. In effect, this would constitute a "beamform-cluster-detect" strategy rather than the "cluster-beamform-detect" strategy explored here. Thus, while the SFP approach examined here provided noticeable gains on experimental data, several avenues exist for further development.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Samuel Earp and the reviewers for their comments on the manuscript. This work was sponsored by DARPA-ATO under Air Force Contract No. FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government. Approved for public release, distribution unlimited.

APPENDIX: EFFICIENT CLUSTERING ALGORITHMS

The computational problem for clustering replica vectors can be posed as follows: "Find the fewest number of cluster center vectors $\tilde{\mathbf{v}}^{\text{cen}}$ such that the maximum distance between each vector $\mathbf{v}_i$ and the closest center vector $\tilde{\mathbf{v}}_j^{\text{cen}}$ is less than η," where η is the cluster radius. This can be written

$$\max_i \min_j \text{dist}(\mathbf{v}_i, \tilde{\mathbf{v}}_j^{\text{cen}}) \leqslant \eta. \tag{A1}$$

The exact solution to this problem is computationally prohibitive to calculate,[26] but approximate solutions can be found. The computational load of clustering full-field replicas was found to be greatly reduced for algorithms with the following characteristics: (1) the clustering method is "greedy," meaning that it will seek a suboptimal solution; and (2) a divide-and-conquer (denoted DnC) approach is used, in which the algorithm operates on subsets of replica vectors to build up the solution. For the cases studied, using a greedy algorithm rather than a more exact method (such as hierarchical clustering) gave a speedup of roughly 1000×. DnC gave an additional speedup of 10–20×.

A useful greedy clustering method has been proposed by Gonzalez.[26] This algorithm iteratively adds cluster centers by finding the vector that has the greatest distance from the previous set of centers. A simpler inner product clustering algorithm was also developed. This algorithm first computes the $n \times n$ pairwise distance or mismatch matrix between every pair of vectors. The vector with the largest number of close neighbors, with mismatch less than the threshold, is picked as the cluster center. All replica vectors closer than the threshold are assigned to this center. The cluster center and cluster members are removed from the vectors to be clustered, and the process is repeated until all vectors are clustered. A final step re-assigns all non-centers to the closest center. The inner product method was found to be slightly faster than the Gonzales method and was used for the results shown above.

A divide-and-conquer (DnC) approach, in which the final clustering solution is constructed from the clustering solutions found for each subgroup of replicas, provides additional gains. DnC is implemented as follows: First, a subset of MFP replicas is extracted from the overall set. For the results shown above, every tenth replica was selected. A set of clusters is formed from this small set. A second subset of MFP replicas is extracted, and the projection of these new replicas onto the previous clusters is found. Any of the new replicas that fall within previously identified clusters are assigned to them; the remaining replicas are then clustered themselves, and the new clusters are added to the overall set of clusters. This process repeats until all MFP replicas are assigned to a cluster. Typically, the run time using this approach is dominated by the clustering of the first subset. Several different methods for breaking the MFP replicas up into subsets were examined. The subdivision method used did not have a strong effect on clustering performance.

[1] D. Manolakis, V. Ingle, and S. Kogon, *Statistical and Adaptive Signal Processing* (McGraw-Hill, New York, 2000).
[2] B. Tracey, "Statistical description of matched field processing ambiguity surfaces," J. Acoust. Soc. Am. **118**(3), 1372–1380 (2005).
[3] A. Baggeroer, W. Kuperman, and P. Mikhalevsky, "An overview of matched field methods of ocean acoustics," IEEE J. Ocean. Eng. **18**(4), 401–424 (1993).
[4] M. Siderius, M. Snellen, D. Simons, and R. Onken, "An environmental assessment in the Strait of Sicily: Measurement and analysis techniques for determining bottom and oceanographic properties," IEEE J. Ocean. Eng. **25**(3), 364–386 (2000).
[5] D. Knobles, R. Koch, L. Thompson, K. Focke, and P. Eisman, "Broadband sound propagation in shallow water and geoacoustic inversion," J. Acoust. Soc. Am. **113**(1), 205–222 (2003).
[6] M. Collins and W. Kuperman, "Focalization: Environmental focusing and source localization," J. Acoust. Soc. Am. **90**(3), 1410–1422 (1991).
[7] J. Krolik, "Matched-field minimum variance beamforming in a random ocean channel," J. Acoust. Soc. Am. **92**(3), 1408–1419 (1992).
[8] J. Tabrikian, J. Krolik, and H. Messer, "Robust maximum-likelihood source localization in an uncertain shallow-water waveguide," J. Acoust. Soc. Am. **101**(1), 241–249 (1997).
[9] H. Schmidt, A. Baggeroer, W. Kuperman, and E. Scheer, "Environmentally tolerant beamforming for high resolution matched field processing: Deterministic mismatch," J. Acoust. Soc. Am. **88**(4), 1851–1862 (1990).
[10] P. Daly, "Stochastic matched field processing for localization and nulling of acoustic sources," Ph.D. thesis, Massachusetts Institute of Technology, 2000.
[11] N. Lee and C. Richmond, "Threshold region performance prediction for adaptive matched field processing localization," *Proceedings, 12th Annual Adaptive Sensor Array Processing Workshop* (MIT Lincoln Laboratory, Lexington, MA, 2004) (www.ll.mit.edu/asap).
[12] B. S. Everitt, *Cluster Analysis* (Wiley, New York, 1993).
[13] B. Tracey, N. Lee, and S. Turaga, "Statistical clustering applied to adaptive matched field processing," *Proceedings, 12th Annual Adaptive Sensor*

Array Processing Workshop (MIT Lincoln Laboratory, Lexington, MA, 2004) (www.ll.mit.edu/asap).
[14]D. Rabideau, "Clutter and jammer multipath cancellation in airborne adaptive radar," IEEE Trans. Aerosp. Electron. Syst. **36**(2), 565–583 (2000).
[15]T. C. Yang and T. Yates, "Matched-beam processing: Application to a horizontal line array in shallow water," J. Acoust. Soc. Am. **104**(3), 1316–1330 (1998).
[16]F. Jensen, W. Kuperman, and M. Porter, and H. Schmidt, *Computational Ocean Acoustics* (American Institute of Physics, New York, 1993).
[17]H. Cox, "Resolving power and sensitivity to mismatch of optimum processors," J. Acoust. Soc. Am. **54**(3), 771–785 (1973).
[18]L. Zurk, N. Lee, and J. Ward, "Source motion mitigation for adaptive matched field processing," J. Acoust. Soc. Am. **113**(5), 2719–2731 (2003).
[19]H. Cox and R. Pitre, "Robust DMR and multi-rate adaptive beamforming," *Proceedings of 31st Asilomar Conference* (1997), pp. 920–924.
[20]H. Cox, "Robust adaptive beamforming," IEEE Trans. Acoust., Speech, Signal Process. **35**(1), 1365–1376 (1987).
[21]*Mismatch* is defined here as a difference between the assumed array response of a signal (i.e., its steering vector) and the actual array response in the data.
[22]N. Lee and N. Pulsone, "Performance of sample-covariance-based adaptive sonar detectors," *Proceedings of 34th Asilomar Conference* (2000), pp. 668-672.
[23]H. Cox and D. Pace, "A fast normalizer," *Proceedings of 30th Asilomar Conference* (1996), pp. 459–463.
[24]W. Struzinski, "ORing loss data for square law detectors followed by an ORing device and an accumulator," J. Acoust. Soc. Am. **105**(4), 2170–2180 (1999).
[25]G. Bottomley, "The effects of cross-correlated noise and multi-channel signal on ORing loss," *Proceedings, IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Vol. **2** (1987), pp. 26.6.1–26.6.4.
[26]T. F. Gonzalez, "Clustering to minimize the maximum inter-cluster distance," Theor. Comput. Sci. **38**, 293–306 (1985).