ACOUSTIC SOURCE SPECTRAL ESTIMATION
FOR A SHALLOW WATER CHANNEL

A. A. (Louis) Beex and John F. Tilki

Digital Signal Processing Research Laboratory
The Bradley Department of Electrical Engineering
VIRGINIA TECH
Blacksburg, VA 24061-0111

INTRODUCTION

We aim to find an unbiased estimate of the spectral content of an acoustic source in a shallow water channel, from measurements on an array of acoustic sensors. The acoustic source spectrum is assumed to be a mixture of discrete and continuous spectral components. Due to multipath arrival at the sensors the spectral distribution at any individual sensor is no longer simply a scaled version of the source spectrum. Consequently, spectra associated with individual sensor signals produce biased estimates of the source spectrum.

Information about the spectral content of the source is present in the sensor array signals. The signals arriving at the sensors have, however, been modified by the channel characteristics from source to sensor. We aim to equalize the channel characteristics that produced each sensor signal, to generate multiple estimates of the source spectrum. The "volatility" of the channel characteristics for a particular sensor in the array is taken into account when fusing the source spectral estimates produced by the individual sensors.

An example integrated throughout the paper shows the efficacy of this procedure, predicated on being able to obtain accurate source-sensor transfer characteristics.

PROBLEM SCENARIO

Let us focus on the problem first. Assume that we have a shallow water channel with a depth of 104 m. Assume a vertical sensor array of 34 sensors, spaced 3 m apart, and thus spanning the depth of the channel. We emphasize that in our procedure there is no need for a vertical array, equal spacing, or spanning the depth of the channel; indeed, the procedure applies for a completely arbitrary array configuration. The acoustic source is assumed to be located at a range of 1000 m and a depth of 40 m.

A Matlab version of KRAKEN[1] is used to model the littoral channel, implementing the normal mode approach. For a given sound velocity profile, a realistic one measured in the Gulf of Mexico, bottom absorption, attenuation due to surface and bottom scatter, and the given channel depth, array sensor locations, and source position the transfer functions from the source to the individual sensors can be evaluated numerically with this water - sediment - rock layer model. This is done for a range of frequencies of interest, we used 15-150 Hz, and the results stored. Given the transfer function information we then generate the appropriate complex exponential received at a given sensor, and sample this signal at an appropriate rate (our sampling frequency was 300 Hz). Thus we use the normal mode approach to generate a sensor data array for a particular discrete frequency. Assuming the channel to be linear, the generation process can be repeated for several frequencies (we used the set $\underline{f}$=[75, 100, 125] (in Hz), with corresponding amplitudes $\underline{A}$=[1, 1/3, 1/5] and phase angles $\underline{f}$=[0, 0, 0], to produce a periodic signal as might be expected from rotating machinery), and combined into the discrete sensor data array Xd. A sensor data array Xc, due to a continuous source spectral component, was generated using the KRAKEN source-sensor impulse responses with an autoregressive input; AR(2) with poles at 90 and 115 Hz, and pole radii at 0.95. With Xn representing the wideband sensor noise, the 100x34 sensor array data matrix X is thus given by

$$X = Xd + Xc + Xn \tag{1}$$

The discrete-signal-to-wideband-sensor-noise ratio is set to approximately 30 dB, while the continuous source component is visible above the wideband noise floor at the sensors. We used 100 snapshots of sensor array signals.

SOURCE SPECTRAL ESTIMATION PROCEDURE

To find out which discrete frequencies are present in the source, we perform column-wise spectral estimation on the data matrix X. Several methods are available to do this, ranging from the classical periodogram (DFT-based) to the modern, model-based approaches. Assuming that the relevant frequency components are spaced apart (here, 25 Hz) by much more than the Rayleigh limit of resolution (300/100 = 3 Hz here), the periodogram-based methods are nearly optimal frequency estimators for the white noise case.[2] To derive accurate frequency estimates one can use lots of zero-padding. With a quadratic least squares fit in the neighborhood of spectral peaks, some immunity to the presence of noise results.

Discrete Source Component

Note that the periodogram exhibits processing gain, thereby increasing the emphasis on discrete spectral components over wideband noise as more data (before zero-padding) is used. Each of the column-wise spectral density (or power) estimates reflects the discrete source components, as modified by the source-sensor transfer functions. Consequently, the relative amplitudes and phase angles of the discrete spectral components have changed. Indeed, due to a null, at a particular frequency, in a particular source-sensor transfer function, one or more of the discrete spectral components may not be present at all. However, it is unlikely that such a null occurs for all of the sensors in the array. Information from all column-wise spectral estimates must be fused, to detect that set of discrete frequencies most likely to have occurred at all sensors. To illustrate, we simply use the sum of all the column-wise spectral estimates and then find the peak location frequencies. This results for our example, using a Blackman window and zero-padding to 8192 points, in the frequency estimates $\underline{f}$=[75.008, 99.977, 124.88]. Note that the estimation error is generally larger for the weaker discrete frequency components.

Source Localization

Knowing a discrete frequency estimate, we can use KRAKEN to generate a database of complex (amplitude/phase) sensor information, *D(d,r)*, for varying source depth and range, given the sound velocity profile and sensor array configuration. Narrowband filtering of the array data at the individual discrete frequency estimates, generates a number of filtered array data matrices (one for each discrete frequency) or, after Fourier transformation or spectral modeling, the equivalent row-wise spectral sensor measurements, *Y(f)*.

A multitude of matched field processors (MFP) is available to compare the measurement *Y(f)* with the database *D(d,r)* to yield source depth and range estimates which correspond best with the received sensor data[3]. The results for each of the discrete frequencies can be fused, to produce an improved estimate.

Our present aim is to show that bias-free estimation of source spectra is feasible. Therefore, we assume that the source can be localized accurately via MFP, meaning availability of a source depth estimate *d*=40 m, and a source range estimate *r*=1000 m.

Discrete Source Amplitude/Phase Estimation

If we let *s* represent the sensor array geometry, we can now use KRAKEN to generate the array *H(f,s,d,r)* of sensor-measured responses to a unit amplitude complex exponential source. The latter allows us to apply its inverse to the spectral sensor measurements *Y(f)*, thereby equalizing the effects of the source-sensor transfer functions. This yields, for each of the sensors, an estimated source amplitude and phase angle for each of the discrete frequency components, as shown in Figure 1. The variability of these estimates comes from 3 sources: sensor measurement noise, differences in source-sensor transfer characteristics, and the presence of the continuous spectral component in the source. For the sake of simplicity, we fuse all estimates into their weighted average, where the weighting is proportional to the absolute value of the source-sensor transfer function value at the given frequency. The latter is aimed at countering the noise amplification effects from inverting small transfer function values. This fusing process results in source amplitude and phase angle estimates of *A*=[1.024, 0.3391, 0.1952] and *f*=[-0.006499, -0.0006275, -0.001556] (in ρ radians) respectively.
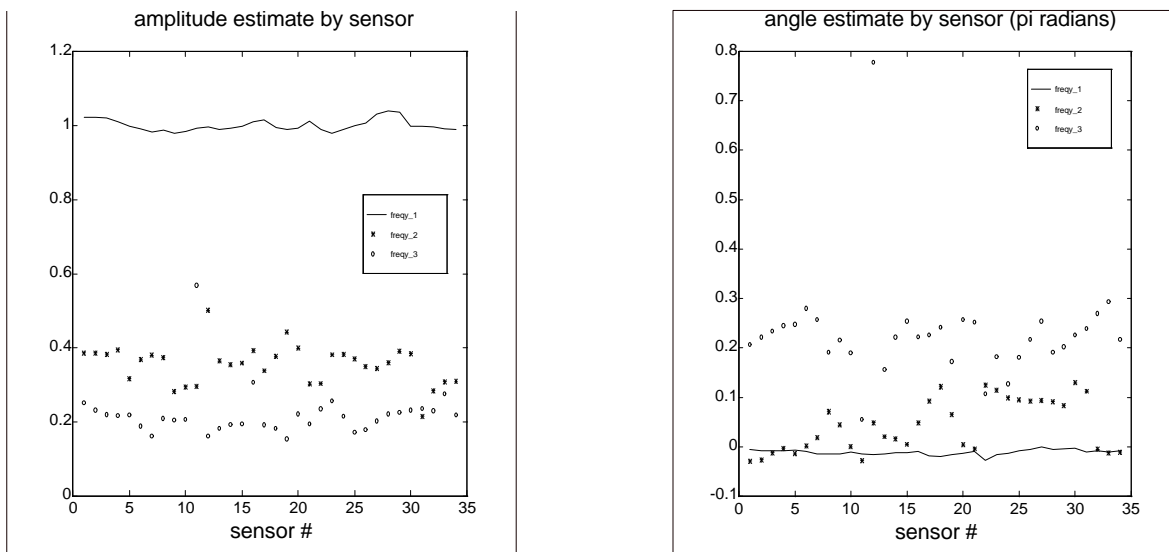

Figure 1 Discrete Amplitude Estimates (L) and Discrete Phase Angle Estimates (R).

Source Inference and Quality Control

On the basis of the source amplitude and phase angle estimates we now reconstruct or estimate the discrete source component. Inference on the source can then be obtained, for example by template matching, audio identification, or other pattern recognition procedures.
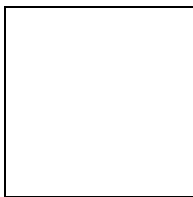
The source estimates, together with the earlier obtained source-sensor transfer function information $H(f,\underline{s},d,r)$, can be used to reconstruct or estimate Xd. Now $X\text{-}Xd$ $(=Xc+Xn)$, the continuous sensor residual, can provide an idea of how well the discrete source component was captured. When localization is in error, or when discrete frequency estimates are off, be it in frequency, amplitude, or phase angle, the continuous sensor residual is less in tune with sensor noise statistics. This quality control process can also be done spectrally, taking advantage of the processing gain for discrete components. In any case, reconstruction helps to ascertain the integrity of the many processing steps.

Continuous Source Component

So far we have seen that accurate results can be obtained for discrete source components, given conducive processing and environment. We therefore continue on the basis of the continuous sensor residual estimate.

As samples from an arbitrary signal can be represented as a linear combination of discrete frequency components, via the DFT frequencies, one might be tempted to apply the above procedure for discrete frequencies to the continuous spectral component. However, this method does not readily extend to the case of a continuous spectral component. This is seen when we examine what happens when we have a single discrete frequency component which cannot be captured by a single DFT frequency. As a result of spectral leakage the power in this single frequency component will be distributed over all, or most, of the uniformly spaced DFT frequencies. As the source-sensor transfer functions are different for each of the DFT frequencies, the DFT frequency components are multiplied by different complex constants corresponding to the transfer function values at those frequencies. The actual single frequency component, however, would undergo a change by a single complex constant, corresponding to its source-sensor transfer function value. The decomposition corresponding to the representation components that were subject to the different transfer functions no longer corresponds to the representation for the original discrete component that was subject to a single transfer function value! Consequently, a decomposition type representation can not be used here.

Assuming that the sensor signals are bandpass filtered, so that no DC component is measured and aliasing is small enough after sampling, the continuous sensor signal satisfies:

$$\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad\quad (2)$$

where $H(w)$ represents the source-sensor transfer function, and $X(w)$ represents the Fourier transform of the source signal. While in its Riemann sum approximation the integral can be thought of as a linear combination of discrete sinusoids, this will only approximate the signal $y(t)$ well if the discretization in the frequency variable w is fine enough. The latter means that the product $H(w)X(w)$ is nearly constant over the discretization interval. While $X(w)$ is fairly smooth over frequency, $H(w)$ for the underwater acoustic problem is not. This can be seen by finding $H(w)$ over the sensor bandpass filter's frequency range of interest. Figure 2 shows the

source-sensor magnitude response for sensor #17, as evaluated by KRAKEN at the 2048 DFT frequency samples. Note the resonant volatility of $H(w)$, which makes it more difficult than in the discrete source frequency case to apply the corresponding inverse, or equalizer, to the sampled measured sensor signals. In principle, $H(w)$, inclusive of our bandpass filter, corresponds to an impulse response. The latter can be approximated by discretization, resulting in the unit pulse response $h(n)$. The sampled version of (2) then expresses the received sensor signal samples $y(n)$ as the convolution of the unit pulse response and the source signal samples $x(n)$. An inverse DFT of the 2048 DFT frequency samples for $H(w)$ yields the approximate (aliased) impulse response shown in Figure 2.
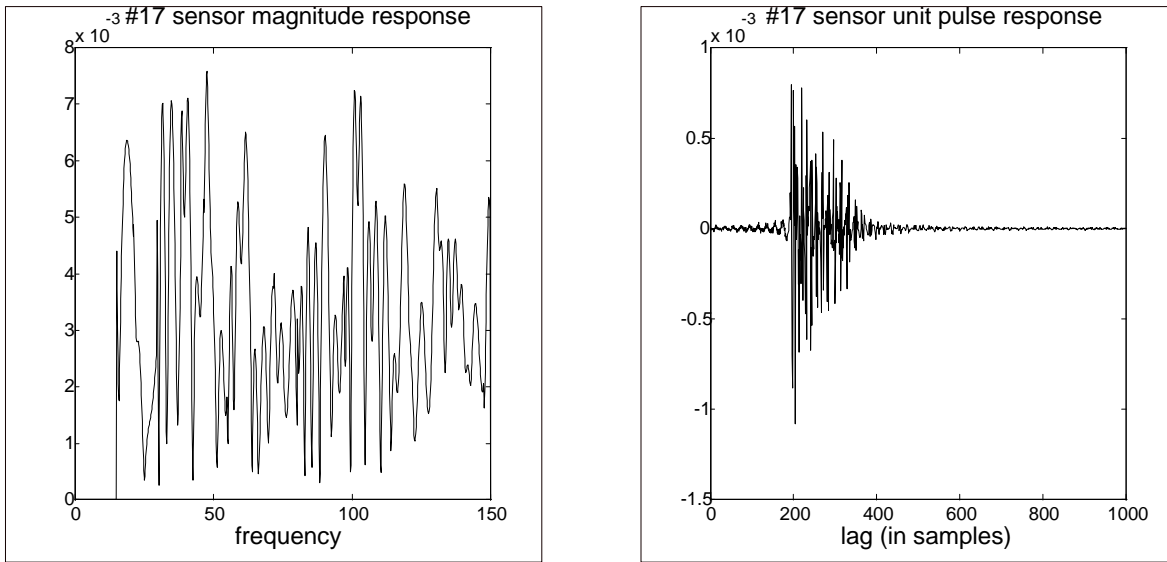

Figure 2 Source-Sensor#17 Characteristics: Magnitude Response (L) and Unit Pulse Response (R).

For a sound velocity of 1500 m/s, it takes 2/3 of a second for sound to travel 1000 meters. At a sampling frequency of 300 Hz it takes 200 samples to travel 1000 meters. The source-sensor unit pulse response must reflect this 200 sample interval, as seen in Figure 2. If too few DFT frequency samples are taken severe aliasing results, reflected in the lack of such necessary physical attributes. In the sequel 1024 point DFT results were used.

As before, the application of spectral estimation techniques to the sensor signals does not reflect the source spectrum directly, because the source signal has been subject to the source-sensor transfer functions. Knowing the accuracy of the unit pulse response, the corresponding DFT can simply be "inverted" on a frequency sample-by-sample basis. As the measured sensor signals are of a bandpass type, the DFT values in the stopbands are considered zero, and so are their "inverses." An IDFT of the "inverse" DFT frequency samples then yields a filter that, when cascaded with the unit pulse response filter, yields an equalized response on the order of 1 over the range of passband frequencies.

Improvements in the inverse filter design process, over the frequency sampling design approach taken here, are certainly possible. However, while the result is by no means perfect on a frequency-by-frequency basis, it goes a long way towards equalizing the response for a stochastic process that is broadband relative to the local perturbations in such an equalizer.

The modified covariance method of linear prediction[2] is used on each of the equalized sensor signals to yield a corresponding AR power spectral density estimate. The gain of the AR model is normalized by matching its autocorrelation at lag zero to the variance of the observed sequence. The order of the AR model used was 3, i.e. one higher than the underlying AR component, so that the extra pole could model the (equalized) wideband noise, leaving the

other two poles to model the spectral energy at 90 and 115 Hz. A weighted average, according to the energy in the corresponding unit pulse response, fuses the AR power spectral densities, the AR parameters, or the reflection coefficients from all sensors. Figure 3 shows the source power spectral density (heavy line), differently averaged AR estimators (thin smooth lines), and the averaged periodogram (noisy line), all applied to the equalized estimated continuous sensor residual. Comparing the actual AR source spectral density and the continuous residual spectral estimates indicates that our procedure estimates the major continuous spectral features reasonably well in the present noisy spectrally mixed scenario, while the effects of the "equalized" measurement noise on the continuous residual spectral estimates are not generally negligible.

continuous PSD

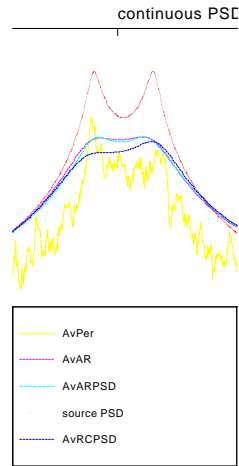| AvPer |
| AvAR |
| AvARPSD |
| source PSD |
| AvRCPSD |

Figure 3 Continuous Sensor Residual Spectral Estimates (100 snapshots).

Without measurement noise, the overall AR spectral density estimate improves noticeably around the spectral peaks. When the number of snapshots is increased the spectral estimates behave as expected, and move closer to the actual source spectral density. Therefore, no systematic bias is apparent.

CONCLUSION

A sensor-equalization approach was demonstrated, based on arbitrary geometry sensor array measurements, for estimating the discrete and continuous spectral components assumed to exist in a source spectrum. The accuracy of the estimates depends on the sensitivity of the frequency functions being sampled, over which control can be exercised, the number of snapshots of available data, the relative strengths of the continuous and discrete source components, and the strength of the measurement noise. Fusing of the results associated with individual sensors reduces estimation variance, because each sensor signal reflects the same source spectrum, while being corrupted by different sensor noise.

REFERENCES

1. J. Ianniello, A Matlab version of the KRAKEN normal mode code, TM 94-1096, NUWC/NL, 10/3/94.
2. S. M. Kay, Modern Spectral Estimation: theory & application, Prentice-Hall, 1988.
3. A. Tolstoy, Matched Field Processing for Underwater Acoustics, World Scientific, 1993.