

Image Cover Sheet

CLASSIFICATION

UNCLASSIFIED

SYSTEM NUMBER

513106



TITLE

Transmission of Acoustic Cross-Spectral Matrices Over Low-Bandwidth
Communications Channels

System Number:

Patron Number:

Requester:

Notes:

DSIS Use only:

Deliver to:



TRANSMISSION OF ACOUSTIC CROSS-SPECTRAL MATRICES OVER LOW-BANDWIDTH COMMUNICATIONS CHANNELS

Garry J. Heard, David J. Thomson, Gordon R. Ebbeson, and Ronald I. Verrall

Defence Research Establishment Atlantic

P.O. Box 1012 Dartmouth, Nova Scotia, CANADA B2Y 3Z7

INTRODUCTION

This paper is concerned with an initial look at how to transmit spectral information from a remote underwater acoustic array to a receiver that may be located anywhere on earth. This goal is a common desire in the fields of underwater acoustics, bio-acoustics, and oceanography.

The only viable way to succeed at this goal is to make use of satellite telemetry. Unfortunately, restrictions imposed by the necessary satellite coverage (open oceans and polar regions), the costs, and the need for relatively simple, small devices at the remote array location strongly restrict the bandwidth of the communications channel.

Previous remote systems have made use of the relatively inexpensive ARGOS satellite system, but these applications are limited to those with extremely low data rates on the order of 0.5-3.5 kBytes/day [1]. The Naval Research Laboratory (NRL) has investigated the other end of the spectrum by employing an 8.5 m-high surface buoy supporting a 1.5 m-diameter, stabilized, satellite transmitting-receiving dish antenna [2]. Neither of these systems are suitable for the majority of remote acoustic array applications: one because the data rate is far too low and the other because of the cost and sheer size of the equipment in addition to the high cost of the satellite time.

At present, and in the near future, it does not appear that there are any satellite systems that could be considered ideal for the majority of remote-array applications, at least for the marine field where the antenna can be expected to be constantly moving. There are two low-bandwidth, relatively low-cost satellite options that offer at least a partial solution: ORBCOMM and IRIDIUM [1]. These systems employ relatively simple antennae that should be adaptable for marine buoy use. Both of these systems have limitations related to availability, data packet lengths, and number of messages. Neither system would be capable of continuous operation and in both cases the inbound (remote location to fixed receiver) data rate is 2400 baud. Further investigation and experiments will be required to determine the applicability of these communication systems to marine remote-array applications.

Having established the motivation for this paper and discussed the options for actual implementation of remote satellite-linked arrays, we now turn our attention to the problem of reducing the array data to a volume manageable over a low-bandwidth communications link. Our particular interest is in transmitting information from a remote array that would allow us to locate and track a source of noise. Our preferred method of acoustic noise location in the ocean is known as matched field processing (MFP) [3].

In its simplest form, MFP consists of correlating spatially-distributed spectral information, often appearing as a cross-spectral matrix, with field replicas generated from an acoustic propagation model. The replicas are generated for hypothetical source locations (e.g., different ranges and depths) spread over a search grid. The correlations produce an ambiguity surface whose maximum value is interpreted as the true location of the acoustic source. Usually,

the problem is complicated by the fact that the acoustic environment is not well known and global optimization over a multi-dimensional state-space is required to produce useful results. Such processing is quite demanding and not currently feasible to locate in a small remote array; hence, it is necessary to return array data to a location where it can be processed.

In the next section we consider the transmission rates necessary to return raw acoustic data, spectral data, and finally selected cross-spectral matrices to a laboratory location. In the third section we illustrate two methods for reducing cross-spectral matrices based on observed properties of such matrices and show the effect of the compression on the MFP results for synthetic test data. We do not consider the use of established compression algorithms which would be applied as a matter of course.

DATA RATES

In order to gain a perspective on the scope of the problem, consider a hypothetical array with 20 hydrophones, each sampled at 2500 Hz in order to provide a useful array bandwidth of 1000 Hz. Consider also that each sample is composed of 16 bits or two bytes. This example array would result in a raw data rate of 105 Bytes/sec. In addition to the raw acoustic data it is also necessary to establish a data protocol to allow for error detection and subsidiary information such as array identification, frequency, sample time, control commands, and averaging time. An additional physical layer protocol may be imposed by the communication system and this protocol may or may not be included in the published data transmission rates of the satellite system. Both protocols essentially act as an overhead on the raw acoustic data. Each data packet could easily require 24-48 Bytes of non-acoustic information. If data messages of 1 kByte are possible, then the overhead would be approximately 5%; however, it is not clear that packet lengths of 1 kByte will be possible and overhead could exceed 40%.

Recall that the existing satellite systems have inbound data rates of 2400 baud corresponding to approximately 240 Bytes/sec. Transmitting raw data from an array is clearly impossible, as together with protocol overhead, a compression ratio of about 700 would be required.

Another possibility is to transmit beam spectra from the array. With 1-Hz resolution, the 1000-Hz acoustic bandwidth would require 1000 numbers for each spectrum. Eight-bit (1 Byte) representation would be sufficient, implying that at least 5 seconds would be required to transmit each spectrum. Since a 20-element array would be easily capable of providing 10 independent beams it would require about 1 minute to transmit a snapshot. This situation would have limited use as spectra would be required more frequently for many applications. The high degree of order and similarity in the spectra would allow for reasonable compression ratios using any number of standard techniques. It is conceivable that data compression techniques would allow enough spectra to be transmitted for the majority of applications.

Cross-spectral matrices comprise other data types that could be

12

100

100

100

100

100

100

100

100

100

transmitted and have the advantage of allowing MFP and other advanced signal processing techniques to be carried out at the receiver location.

In the case of our hypothetical 20-element array, a cross-spectral matrix (CSM) is composed of 400 complex numbers. This array would require 3200 Bytes with single-precision representation. Transmitting an entire matrix would take about 16 seconds. Tracking and localization applications would require numerous CSM's at different times and frequencies. Since applications would employ averaging times of between 0.5 to 25 seconds, it is clear that we would soon fall behind and the result would be snapshot acoustic information with relatively low duty-cycles.

Fortunately, CSM's are Hermitian and therefore we can immediately save channel bandwidth by sending only the diagonal and upper-triangular part of the matrix. Although this symmetry reduces the data to be transmitted by almost half, even this data reduction is not sufficient and further compression is desired.

CSM COMPRESSION

Two relatively simple schemes for reducing the amount of data in a CSM are discussed in this section. The first scheme consists of a threshold operation on the CSM and a subsequent sparse matrix representation of the thresholded matrix. The second method consists of an eigenvalue/eigenvector decomposition of the CSM followed by elimination of eigenvalue/eigenvector pairs based on the magnitude of the eigenvalues. The CSM matrices are reconstructed from the reduced data sets and MFP is applied. The effects of successive increases in the level of compression are observed on the peak ambiguity value and the determined location (range, depth) of the acoustic source.

Two synthetic data sets are included in the current work. These data sets are COLNOISE A (40 dB SNR, source at 9.1 km, 66 m depth) and B (-5 dB SNR, source at 9.7 km, 58 m depth) taken from the matched field processing benchmark problems [4]. The COLNOISE data include a coloured spectral background due to breaking waves at the ocean surface. The test environment consists of a shallow water region of 100-m uniform depth. A weakly downward refracting sound-speed profile exists in the water column and the bottom is consistent with a clay sediment overlying a homogeneous half-space.

CSM's generally contain a few dominant values and many smaller values. It was expected that the small values could be eliminated. Indeed, this seems to be true, as the MFP results (Fig. 1) show that for both high and low SNR test cases a significant portion of the CSM could be set to zero without affecting the localization result. For the high SNR case, we were able to lower the threshold to 65% of the modulus of the largest element without disturbing the localization result. For the low SNR case, we were able to lower the threshold to about 30% of the maximum modulus value. In both cases the sparse matrix representation resulted in over 80% compression without including the additional 50% compression available due to the symmetry of the CSM. Thresholding provides a data compression rate that may be independent of the SNR, clearly many applications could obtain a potential 10:1 reduction in the data rate.

Eigenvector based MFP has been successfully carried out by others [5] in the past. It has long been known that a few eigenvalues are generally much larger than the rest and that MFP can be successful with just one or a few of the largest eigenvalue/eigenvector pairs.

The use of eigenvalue/eigenvector decomposition appears attractive for compressing the CSM data since each CSM can be reduced to $RE(N+1)$ Bytes, where R is the number of bytes in the floating point representation, N is the number of hydrophones, and E is the number of eigenvalues retained. For our example array with single-precision representation this works out to 168 Bytes/eigenvalue, or only 5.25% of the original matrix size. In both the COLNOISE A and B test cases, using only the largest eigenvalue/eigenvector pair proved to be sufficient, providing an almost 20:1 reduction in the data rate.

SUMMARY

This paper has discussed the potential for transmitting spectral data from a remote array via existing low-bandwidth satellite communication channels. Two methods of reducing the necessary data have been presented and illustrated with synthetic test cases. It is apparent that the existing satellite channels are insufficient for continuous operation, but should meet the requirements of some applications where periodic snapshots of data will suffice.

REFERENCES

- 1 Mah, K. Satellite data telemetry. Seimac Ltd. Dartmouth, NS, TR-421-99-002, 1999.
- 2 Rupa, M.A., Goldstein, J.A., Krout, T.L. Data telemetry and acquisition system for acoustic signal processing investigations. NRL Report NRL/MR/5550-96-7820, Naval Research Laboratory, Washington, DC, 1996.
- 3 Tolstoy, A. Matched field processing for underwater acoustics. World Scientific, Singapore, 1993.
- 4 Porter, M.B., Tolstoy, A. The matched field processing benchmark problems. J. Computat. Acoust. 2(3), 161-185, September 1994.
- 5 Ozard, J.M. Matched field processing in shallow water for range, depth, and bearing estimation: results of experiment and simulation. J. Acoust. Soc. Am., 86(2), 744-753, August 1989.




Figure 1. Matched field localization results using CSM with elements smaller than threshold set to zero. (a) sparse matrix compression percentage exclusive of additional reductions available from symmetry, (b) depth estimate, (c) range estimate.

#513106