

XXII. DIGITAL SIGNAL PROCESSING

Academic and Research Staff

Prof. A.V. Oppenheim
Prof. A.B. Baggeroer

Prof. L.J. Griffiths*
Prof. J.S. Lim
Prof. J.H. McLellan

Dr. M.R. Portnoff
Y-t. Li⁺

Graduate Students

T.E. Bordley
W.P. Dove
G.L. Duckworth
C. Esmersoy
D.B. Harris
M.W. Hauser

M.H. Hayes
S. Holtzman
A.L. Kurkjian
S.W. Lang
D.C. LeDoux
N.A. Malik

D.R. Martinez
D.R. Mook
B.R. Musicus
S.H. Nawab
T.F. Quatieri, Jr.
S.G. Sifferlen

The Digital Signal Processing Group is carrying out research in the general area of digital signal processing with applications to speech, image, and geophysical data processing. In addition to specific projects being carried out on campus, there is close interaction both with Lincoln Laboratory and with the Woods Hole Oceanographic Institution.

In the area of speech processing, over the past several years the Digital Signal Processing Group has been working on the development of systems for bandwidth compression of speech, parametric modeling of speech using pole-zero models, and enhancement of degraded speech. Our work in the speech area is currently heading toward an increasing involvement with the problem of enhancing degraded speech and a related problem, that of the development of algorithms for robust speech compression in the presence of additive noise.

In a related area the methods of speech compression using linear predictive encoding are being applied to the compression of data recorded in ocean-bottom seismometers. These methods are being tested with data provided by the Woods Hole Oceanographic Institution.

The areas of image and geophysical data processing in general both involve the processing of multidimensional signals. The theoretical projects in 2-D signal processing include filter design (e.g., 2-D all-pass design to match phase

*Professor, Department of Electrical Engineering, University of Colorado, Boulder, CO.

⁺Visiting Scientist from Tsinghua University, Peking, China.

(XXII. DIGITAL SIGNAL PROCESSING)

response), the synthesis of good 2-D filter implementations, 2-D spectrum analysis, and 2-D deconvolution. We have been pursuing a number of projects specifically related to geophysical data processing. We are applying some of the filter design results to seismic-wave migration by implementing a program on our MAP processor. Another project, which has been carried out in collaboration with the Woods Hole Oceanographic Institution, is the development of an algorithm for data processing to measure the acoustic reflection coefficient from the ocean bottom. Out of this work has come a Hankel transform algorithm which has potential applications to a number of other problems. Another problem area is that of velocity analysis on array data. The specific application that we are considering is that of velocity analysis on well logging data. We are also pursuing a number of other problems associated with the analysis of well logging data, including the development of techniques for event detection. In another application of velocity analysis, we have applied adaptive array processing to measure the reverberation of acoustic signals in the Arctic Ocean, as well as the phase velocity of the seismic paths within the seabed. Acoustic imaging from a submersible often generates an image dominated by strong highlights because of the specular reflections introduced by the relatively long wavelengths. We are working on an adaptive array processing method to suppress the deleterious effects of these highlights in the image.

There are also a number of projects related to image processing that we are currently pursuing. With regard to the problem of image restoration/enhancement, we have developed a new image restoration system which is effective in reducing various image degradations such as additive noise, multiplicative noise, blurring, etc. This image restoration system has also been applied with successful results in reducing quantization noise in PCM image coding. With regard to the problem of phase-only image reconstruction, we have several theoretical conditions under which an image can be exactly recovered within a scaling factor only from its phase function. In addition, numerical algorithms have also been developed to perform image reconstruction from its phase. These results have the potential for significant impact on various image processing problems such as image coding, image registration, etc.

In both the context of image processing and array processing, we are also exploring such topics as high resolution, multidimensional spectral estimation,

and two-dimensional short-space signal processing. The work on image processing is being carried out in collaboration with Lincoln Laboratory.

1. LINEAR PREDICTIVE ENCODING OF SEISMIC DATA

U.S. Navy – Office of Naval Research (Contract N00014-75-C-0852)

National Science Foundation Fellowship

Thomas E. Bordley, Arthur B. Baggeroer

In studying the earth beneath the ocean floor, the use of ocean-bottom sensors physically isolated from the controlling research vessel offers many advantages over the use of surface sensors or deep-sea arrays. Per se, since the sensor is in contact with the bottom, the complex effects associated with the passage of a signal across the earth-water interface and through the sea are eliminated. Furthermore, because the signals are observed in the earth, the shear and stress components of the returns can be recorded as distinct signals (through the use of horizontally and vertically polarized sonars), thus increasing the amount of information gathered. Finally, unlike free-floating deep-sea arrays, the location of an ocean-bottom sensor is secure. Thus, the difficult problem of controlling the position of an array is eliminated as well as the noise effects associated with the residual movement of the sensor.

In such work, storing data in digital form is desirable since it simplifies gain-ranging and other techniques used to ensure high quality reproduction of the observed signals. Unfortunately, at present, digital recording requires an order of magnitude more storage than analog recording. Thus, since seismic signals are of long duration and of wide dynamic range, the limited amount of storage which can be placed on the sea floor presents a severe constraint on the amount of data which can be gathered using all-digital techniques. For this reason, it is of considerable interest to develop a scheme for minimizing the amount of storage required to faithfully represent these signals. In effect, one wishes to reduce the cost associated with the digital approach by using the available storage as effectively as possible.

This research examines the effectiveness of linear predictive encoding in performing seismic data compression. This technique was selected for study originally

(XXII. DIGITAL SIGNAL PROCESSING)

because of its widespread use in speech processing and general data communication. The basic idea of this technique is to pass a signal through a linear, shift-invariant, causal and causally invertible whitening filter, and then to store the filter parameters and the whitened signal instead of the original data. The term linear predictive encoding is appropriate because whitening with this filter is equivalent to subtracting from each signal point the linear least-squares prediction of the point based on previous signal values. Thus, the whitened signal is the sequence of prediction errors made by the linear least-squares prediction filter. The rationale for using this approach to achieve data compression is that if a signal is sufficiently predictable (i.e., if the energy in the error signal is sufficiently small compared to the original signal energy), then the dynamic range of the whitened signal will be significantly less than the dynamic range of the original signal. Thus, the total number of bits required to represent the whitened signal and the filter parameters will be much less than that needed to represent the waveform directly.

This work was performed on data supplied through Dr. Graham M. Purdy of the Woods Hole Oceanographic Institution.

2. EVENT DETECTION IN SONIC WELL LOGGING

National Science Foundation (Grant ENG76-24117)

Webster P. Dove, Alan V. Oppenheim

Oil wells are analyzed by acoustically testing at many places along their depth, from which a sound velocity profile can be developed. For each test a pulse of sound is generated at the bottom of a 13 meter long test probe and received at four microphones spaced at one meter intervals at the top of the probe.

The signal received at each microphone is the sum of many overlapping dispersed pulses, each of which has travelled a different path. To find the velocities in the paths of interest the arrival time of each related pulse must be determined accurately. That requires a method of reducing the dispersion of the pulses so they become distinct separate arrivals.

To do this signal processing, we are using Recursive Least Squares prediction

(the covariance method) to cancel the filtering effect of the different paths. Then the arrival time of the first and second pulses (which are the ones of interest) should be apparent either in the output of the predictor or the behavior of the predictor coefficients.

3. ADAPTIVE ARRAY PROCESSING FOR HIGH-RESOLUTION ACOUSTIC IMAGING

U.S. Navy — Office of Naval Research (Contract N00014-77-C-0257)

National Science Foundation Fellowship

Gregory L. Duckworth, Arthur B. Baggeroer

Determination of the internal structure of a medium opaque or ill-suited to electromagnetic radiation is a problem encountered in many different applications. High-resolution visualization of underwater objects through turbid seawater is the problem currently being dealt with; however, other applications include real-time viewing of internal movements of the human body without x-ray's potentially harmful effects, nondestructive testing of metallic and low x-ray contrast objects, and determination of the earth's subsurface structure.

Because of their analogous behavior to electromagnetic radiation with respect to reflection, diffraction, and refraction, but differing attenuation and physiological properties, short wavelength acoustic pressure waves can be used to perform the above tasks, but with a new set of inherent advantages and difficulties. For example, in the context of the undersea environment, acoustic imaging has an advantage over optical imaging in that the attenuation of the acoustic-pressure waves is dependent primarily on the temporal frequency, and relatively independent of the density of suspended solids, whereas light is subject to intense backscattering from cloudy water. A result of this is that the "range-to-reverberation" limit is larger for acoustic imaging, and although absorption at wavelengths adequate for reasonable resolution is high, we can theoretically increase the illuminating power and obtain the desired range capabilities.

The problems with acoustic imaging stem from the need to keep the wavelengths long enough for adequate range, SNR, and power consumption, and short enough for good resolution and small receiver apertures. These considerations ultimately lead to systems with small numerical apertures with the diffraction field under-

(XXII. DIGITAL SIGNAL PROCESSING)

sampled in space and hence, poor resolution and aliasing problems. Resolution seems to be the most problematic issue since the large point-spread function generated by classical (Fresnel transform) processing is subject to tremendous amounts of sidelobe leakage from specular reflections. Typically a great deal of specularity is encountered since the illuminating wavelengths are large compared to the surface roughness of the objects to be imaged.

These problems have led us to the thrust of the current research — application of the "Maximum Likelihood" technique of spectral analysis to adaptive array processing of the diffraction-pattern samples. It has been found that the adaptive point-spread function of a system incorporating this technique yields better resolution for distributed objects as long as care is taken in estimation of the spatial-covariance function. The subtlety involves making the spatial covariance appear like it was formed by reflections from statistically independent incremental areas. The work also involved determination of the statistics of the estimators when an inadequate number of data vectors are used to ensure that the spatial-covariance matrix is distributed in a complex Wishart manner. Two-dimensional arrays that are optimized in some sense for good resolution and aliasing reduction with the minimum number of sensors are also examined.

This work was completed in September 1979 and is reported in a Master's thesis entitled, "Adaptive Array Processing for High Resolution Acoustic Imaging," by G.L. Duckworth. The results were presented at the 9th Symposium on Acoustic Imaging and Holography in December 1979.

4. DESIGN OF TWO-DIMENSIONAL FILTERS

National Science Foundation (Grant ENG76-24117)

David B. Harris, James H. McClellan

Two methods for the design of two-dimensional (2-D) rational digital filters are developed. Both methods utilize a reflection coefficient function (RCF) representation for the system function. The RCF characterization is an extension of the 1-D Levinson recursion. Its use substantially simplifies the stabilization of filters in the design process.

In one approach, the design problem is formulated as a frequency domain error norm minimization. The norm is a measure of the difference between a desired frequency response and the response of the approximating filter. In this context, the RCF representation is used to construct a set of barrier functions which augment the error so as to restrict the domain of approximation to the set of stable filters. A nonlinear optimization algorithm is employed to perform the minimization.

The second approach to rational filter design is based on a 2-D linear prediction algorithm. The prediction algorithm is used to obtain an all-pole approximation to a desired spectral function. The approximation is an infinite Levinson expansion in terms of reflection coefficient functions. The expansion is truncated, retaining only the most significant RCF's, in order to obtain a realizable filter. With modifications, the method is shown to be useful for all-pass filter design. A method for transforming the phase approximation problem to a spectral factorization problem is presented, in which the prediction algorithm is employed to perform the factorization. The spectral factor is used to construct the appropriate all-pass transfer function.

In addition to simplifying stability constraints, the RCF characterization leads to an efficient lattice structure realization for 2-D recursive filters. Two lattice configurations are examined in detail.

A number of design examples are presented, illustrating the effectiveness of both techniques. The examples are filters useful in the exploration seismology problem. Special attention is paid to the design of fan filters for wave equation migration processing.

5. SIGNAL RECONSTRUCTION FROM PHASE OR MAGNITUDE

U.S. Air Force (Contract F19628-80-C-0002)

Monson H. Hayes, Jae S. Lim, Alan V. Oppenheim

For both continuous-time and discrete-time signals, the magnitude and phase of the Fourier transform are, in general, independent functions, i.e., the signal cannot be recovered from knowledge of either alone. Under certain conditions,

(XXII. DIGITAL SIGNAL PROCESSING)

however, relationships exist between these components. For example, when the signal is a minimum phase or maximum phase signal both the log magnitude and phase can be obtained from the other through the Hilbert transform. This relationship has been exploited in a variety of ways in many fields including network theory, communications, and signal processing. As a result of some research into the general problem of signal reconstruction from phase or magnitude, several sets of conditions have been specified for which a discrete-time sequence is uniquely specified, to within a scale factor, by the phase of its Fourier transform, without the restriction of minimum or maximum phase. The first set of conditions considers the case in which the phase is specified at all frequencies whereas the second set considers the case in which the phase is specified only at a discrete set of frequencies. Once these conditions have been established, a dual set of conditions are presented for which a discrete-time sequence is uniquely specified to within a sign and a delay by the magnitude of its Fourier transform, again without the restriction of minimum or maximum phase.

For those sequences satisfying the appropriate set of conditions, three algorithms have been devised for reconstructing a sequence from its phase or samples of its phase. The first algorithm is conceptual in that it provides some insight into the fundamental results of phase or magnitude signal reconstruction. The remaining two algorithms are numerical and provide a practical method of performing phase-only signal reconstruction. The first numerical algorithm is an iterative technique in which the estimate of the desired sequence is improved at each iteration. The second is a closed-form solution for which the desired sequence is found by solving a set of linear equations.

Although stated only in terms of one-dimensional sequences, the results are readily extendible to multidimensional sequences. This extension is accomplished by utilizing a few well-known results on projection and slices. With this extension, a potential application for phase-only reconstruction is the blind deconvolution of images which have been blurred by a symmetric point-spread function.

Further research is expected to address such questions as the noise sensitivity of reconstruction techniques, the convergence properties of iterative solutions, the existence of more efficient algorithms that may be realistically applied to processing large amounts of data, alternate forms of the multidimensional theorems which do not depend on projection/slice ideas, and the extension

of the known results to include additional constraints on the sequence (such as positivity).

6. TIME-SCALE MODIFICATION OF SPEECH

U.S. Navy — Office of Naval Research (Contract N00014-75-C-0951)

Samuel Holtzman, Michael R. Portnoff, Jae S. Lim

We have implemented an analysis-synthesis system on our PDP-11 computer that performs uniform-rate speed transformations on speech signals. The problem of spectral degradation and introduction of noise, which usually occur in similar systems, are not found in this one. This is because the system performs the transformations in the frequency domain by means of the discrete short-time Fourier transform rather than in the time domain.

In order to achieve an even more natural-sounding result, we are at present introducing nonuniformities into the speed transformations to incorporate a dependency of the system on local features of the speech signal being transformed.

Our work has been directed toward the development of an algorithm to automatically segment the speech signal into a sequence of passages for which an expected level of degradation, caused by uniform time-scaling, can be determined. The purpose of segmenting the signal in this manner is to allow the degree of local time-scale modification to be decreased whenever the expected level of degradation is high.

The algorithm uses a statistical analysis of the speech signal to determine a level of local quasi-stationarity which, based on our model of speech production, is highly correlated with the expected level of local degradation.

7. THE ESTIMATION OF FORMATION PARAMETERS IN SONIC WELL LOGGING

U.S. Navy — Office of Naval Research (Contract N00014-75-C-0951)

Schlumberger-Doll Research Center Fellowship

Andrew L. Kurkjian, Alan V. Oppenheim

The sonic well logging problem is modeled as an ideal point source in an infinite fluid cylinder surrounded by an elastic solid. There is an ideal point

(XXII. DIGITAL SIGNAL PROCESSING)

source on the borehole axis whose waveform is assumed to be known. The borehole radius and parameters of the fluid are also assumed to be known. An ideal array of point receivers is placed along the borehole axis a known distance from the source. The problem is to estimate the parameters of the solid formation from the received signals.

Existing estimation schemes are based on overly-simplified physical models of the acoustical field. Consequently, these methods have achieved only limited success. The purpose of this research is to develop a signal processing procedure which takes into account more of the fundamental physical behavior of the problem than existing methods. This interdisciplinary approach to the problem is expected to lead to superior estimation techniques.

8. SPECTRAL ESTIMATION FOR SENSOR ARRAYS

National Science Foundation (Grant ENG76-24117)

Hertz Foundation Fellowship

Stephen W. Lang, James H. McClellan

The processing of signals from sensor arrays has applications to areas such as radar, sonar, seismology, and radio astronomy. In several important cases, far-field imaging, for example, the type of processing desired can be classified as power spectral estimation. Current areas of research interest are the investigation of techniques for estimating power spectra and the interaction between the selection of a spectral estimation algorithm and the design of arrays.

There is no one best way to estimate spectra; the circumstances and requirements of a particular application may favor one estimation technique while putting another at a disadvantage. For example, the data or input requirements of a spectral estimation technique can affect its utility. Does it require time series data from each sensor or does it begin with correlation measurements between sensors? Does it require regularly sampled data or can it handle irregular sampling? Computational considerations are also important. Does an algorithm exist to compute the spectral estimate with a reasonable amount of effort? Performance is an issue. How well does the estimation technique perform with the spectra likely to be encountered?

In recent years many new spectral estimation techniques have been introduced. However, many of them have only been developed in the context of one-dimensional spectral estimates using evenly sampled data. It would be of great interest to use ideas which led to these one-dimensional techniques in developing spectral estimation techniques and computational algorithms more suitable for use with sensor arrays: multidimensional estimates from irregularly spaced sensors.

Array design is intimately related to spectral estimation; the design of the array determines what information is available to be processed. The relationship between array design and the performance spectral estimation is therefore of interest.

9. SPEECH ENHANCEMENT

U.S. Navy — Office of Naval Research (Contract N00014-75-C-0951)

Jae S. Lim, Alan V. Oppenheim

Degraded speech occurs in a variety of contexts, and its enhancement is desirable for many practical applications. In our past research on this problem, we have developed several systems for enhancement and bandwidth compression of noisy speech by attempting to estimate the parameters of a specific underlying speech model based on the Maximum A Posteriori (MAP) estimation procedure. When the systems were implemented and applied to real speech data, they performed well as enhancement and potential bandwidth compression systems of noisy speech at various S/N ratios. As a practical application area, we have applied our speech enhancement systems with successful results in reducing quantization noise in PCM speech coding. Our future research in this area will include investigation of methods to improve our current speech enhancement systems, development of new systems, and application of speech enhancement systems to practical problems.

(XXII. DIGITAL SIGNAL PROCESSING)

10. TWO-DIMENSIONAL POWER SPECTRUM ESTIMATION

U.S. Air Force (Contract F19628-80-C-0002)

Government of Pakistan Scholarship

Naveed A. Malik, Jae S. Lim

Power spectral density estimation occupies a place of fundamental importance in such diverse fields as speech processing, seismic signal processing, image restoration, radar, sonar, radio astronomy, etc. The current research is directed toward the characterization and understanding of existing two-dimensional power spectrum estimation (PSE) methods, and the development of high-resolution techniques applicable to two-dimensional and higher signals.

Most of the conventional PSE methods extend in a straightforward manner from one to two dimensions. However, important high-resolution techniques, such as the Maximum Entropy Method (MEM), pose highly nonlinear problems when applied to two-dimensional signals. The MEM solution has been shown to have an autoregressive or all-pole form for problems of any dimensionality. Further, an important new result by Woods proves the existence and uniqueness of the MEM solution to the two-dimensional PSE problem under certain constraints.

In order to gain an understanding of the problem, the performance evaluation of various ad hoc autoregressive (all-pole) spectral estimators has been performed. These models were derived via a minimum mean-square-error algorithm, using a linear prediction type of problem formulation.

The shape of the autoregressive (AR) filter mask is arbitrary in two dimensions, and for the purposes of PSE, even causality of the filters is not an issue. Thus the AR problem formulation is not unique in two dimensions. Various filter




Fig. XXII-1. Filter mask shapes.

mask shapes have been investigated, including the 'L'-shaped filter, the symmetric filter, and single-quadrant filters (Fig. XXII-1).

It was found that the spectral estimate was 'stretched' or distorted from the true shape, the distortion being a function of the mask shape employed. Further, the mask shape also dictates the resolution properties of the filters, which can be very directional in nature. It has been conjectured that the stretching is due to the nonsymmetric shape of the region over which information about the signal is available (i.e., a nonsymmetric autocorrelation function region). As such, complementary filter sets (such as the 'backward L' and 'forward L' shapes) can be found such that the average of their spectral estimates provides a reasonably good, undistorted estimate of the power spectrum of the signal. The AR filters certainly perform better than the conventional periodogram-type estimators, and given some a priori information about the signal generating process, it may be possible to make an intelligent choice of the filter mask to obtain a good power spectrum estimate.

The current area of interest is the true MEM problem, and an iterative method of obtaining the MEM solution is being investigated. Tests of the algorithm in the one-dimensional case (where the exact MEM solution is available for comparison) have been completed, and now the extension to two dimensions is being performed. The technique seems to hold a lot of promise in cases where the signal-to-noise ratio is very low. The method is completely general, and if results indicate its viability as a PSE technique, the extension to problems of any dimensionality should be fairly straightforward.

11. EXTRACTION OF ACOUSTIC PLANE-WAVE REFLECTION COEFFICIENT FROM THE SOUND FIELD GENERATED BY A POINT SOURCE

U.S. Navy — Office of Naval Research (Contract N00014-77-C-0196)

Douglas R. Mook, Alan V. Oppenheim, George V. Frisk

[G.V. Frisk is with the Woods Hole Oceanographic Institution.]

Under the assumption of horizontal stratification, the ocean floor can be characterized by an acoustic plane wave reflection coefficient which is a function of angle. When measurements are made using a point source, this function

(XXII. DIGITAL SIGNAL PROCESSING)

can be extracted from the reflected pressure field by using the Hankel transform. The merits of this approach for practical problems were examined during the past year. Future work will concentrate on using theoretical results to derive a mode based scheme for extracting the reflection coefficient.

12. MAXIMUM LIKELIHOOD ESTIMATION WITH NOISY DATA

U.S. Navy — Office of Naval Research (Contract N00014-75-C-0951)

Bruce R. Musicus, Jae S. Lim

Maximum Likelihood (ML) Estimation is a powerful tool for estimating model parameters or signals from observed system output. Not only does it yield estimates with nice theoretical properties, but the estimates are also easily calculated for many useful signal models. Unfortunately, when both the parameters of the system as well as the system output must be estimated from observations corrupted by noise, Maximum Likelihood Estimation usually requires a difficult non-linear optimization. Three different ML approaches have been proposed for estimating the signal and parameters of a system from noisy observations. We have found iterative algorithms for solving each of the three problems, which effectively decouple the uncertainty in the parameter and signal values, thus simplifying the calculation required. When applied to a particular pole-zero model, all three algorithms iterate back and forth between linearly filtering the observations to estimate the signal and fitting parameters to the signal estimate by solving linear equations. The theoretical properties of the algorithms, their relationship to methods previously proposed by Lim, and their application to a variety of signal models have been studied. Present work on these algorithms is focusing on testing their performance with real data, and on developing adaptive real-time versions of the algorithms.

13. PROCESSING OF SATELLITE IMAGERY

U.S. Air Force (Contract F19628-80-C-0002)

Syed H. Nawab, Alan V. Oppenheim, Jae S. Lim

Investigations are being carried out to develop and apply signal processing techniques for problems associated with the processing of multispectral satellite imagery (e.g., LANDSAT). In particular, the problems of restoration and registration are being addressed. Because of the multispectral as well as multitemporal availability of the image data, the signal processing aspect is challenging and has substantial potential for innovation.

14. EVALUATION OF CIRCULARLY SYMMETRIC TWO-DIMENSIONAL FOURIER TRANSFORMS AND ITS APPLICATION TO THE MEASUREMENT OF OCEAN-BOTTOM REFLECTION COEFFICIENTS

U.S. Navy — Office of Naval Research (Contracts N00014-75-C-0951
and N00014-77-C-0196)

Alan V. Oppenheim, George V. Frisk, David R. Martinez

[G.V. Frisk is with the Woods Hole Oceanographic Institution.]

[D.R. Martinez was with the M.I.T.-W.H.O.I. Joint Program in Oceanography/Oceanographic Engineering; he is currently with ARCO Oil and Gas Company, Dallas, Texas.]

In a variety of applications the need arises for the evaluation of the two-dimensional Fourier transform of circularly symmetric functions. Because of the circular symmetry, the two-dimensional Fourier transform reduces to the Fourier-Bessel or Hankel transform. This research considers a method for evaluating this transform using the "projection-slice" theorem for multidimensional transforms. The method is applied specifically to the measurement of the plane-wave reflection coefficient of a horizontally stratified ocean bottom using the fact that, for a point source, the bottom-reflected field and the plane-wave reflection coefficient are circularly symmetric and are related through a two-dimensional Fourier transform.

(XXII. DIGITAL SIGNAL PROCESSING)

15. SHORT-TIME FOURIER ANALYSIS

Michael R. Portnoff

Short-time Fourier analysis is based on the notion of a multidimensional representation for a one-dimensional signal. Specifically, a one-dimensional time signal, $x(t)$, is represented by a two-dimensional function of time and frequency, $X(t,\omega)$, called a short-time Fourier transform (STFT). In its simplest form, the STFT $X(t,\omega)$ is defined as

$$X(t,\omega) = \int_{-\infty}^{\infty} x(\tau) w(t - \tau) e^{-j\omega\tau} d\tau,$$

where $w(t)$ is a window function that is, in some sense, narrow in time or frequency, or both. In its more general form, the STFT is defined using a window that is allowed to depend on both time and frequency.

Short-time Fourier analysis is particularly useful for studying "slowly time-varying" phenomena such as speech, music, and other acoustic signals, because rapidly varying local features appear as functions of frequency in the STFT, whereas slowly varying global features appear as functions of time. Thus, the STFT is a formal mathematical description for our notion of a "time-varying spectrum." Furthermore, short-time Fourier analysis is, in many ways, analogous to the acoustic processing performed by the human auditory system.

Our research in this area has been completed with a better understanding of this method of signal analysis, and we are in the process of reporting our research results in the literature.

16. PHASE ESTIMATION

U.S. Navy — Office of Naval Research (Contract N00014-75-C-0951)

Thomas F. Quatieri, Jr., Alan V. Oppenheim

The phase estimation problem for discrete-time sequences was first encountered in the development of a mixed-phase homomorphic vocoder, where the smooth-phase estimate of the vocal-tract impulse response led to harsh-sounding synthetic

speech. The sensitivity of the unwrapped-phase envelope to time-domain perturbations was observed to be greater than that of the log-magnitude spectrum. These observations initiated a number of questions and answers about properties, relations, and estimation procedures of the magnitude and phase of a Fourier transform.

In particular, a theoretical framework was developed for unwrapped-phase estimation from harmonic spectra (voiced speech) through smoothing real and imaginary spectral components. Short-time homomorphic analysis and a short-time harmonic model have led to pitch-adaptive duration and alignment requirements on time-domain windowing. The underlying phase envelope is consequently preserved so that cepstral windowing can be applied. The result is a mixed-phase homomorphic vocoder of somewhat higher quality than its minimum-phase counterpart.

In addition, two alternative mixed-phase vocoders were considered: the first is based on linear interpolation of complex harmonic peaks, and the second on Lim's homomorphic spectral-root deconvolution scheme.

Also, iterative procedures have been developed for recovering the magnitude from the phase function, and, likewise, the phase from the magnitude function. These algorithms are applicable to phase and magnitude estimation, and are easily extended to two dimensions. The reconstruction of phase or magnitude from the known component relies on constraints such as finite duration and causality. In some cases, convergence of the algorithm to a unique phase or magnitude can be demonstrated.

We are currently considering the application of these iterative techniques to the restoration of degraded images with emphasis on phase estimation when a magnitude estimate is known.

Finally, a new method of phase estimation has been developed which transforms the phase estimation problem to a magnitude estimation problem. The advantage of this procedure is that established magnitude estimation techniques (e.g., Wiener filtering) can be applied to phase estimation.

