25.0 Digital Signal Processing

Academic and Research Staff

Prof. A.V. Oppenheim, Prof. J.S. Lim, Prof. B.R. Musicus, Prof. A.B. Baggeroer, Dr. M. Feder, G. Aliberti

Graduate Students

M. Bace, J. Bondaryk, D. Cobra, M. Covell, D. Griffin, D. Harasty, J. Hardwick, S. Isabelle, J. Jachner, T. Joo, T. Pappas, G. Wornell, R. Wright, A. Zakhor

Support Staff

P. Eiro, D. Gage, C. LeBlanc

Part-Time Assistants/Special Projects

A. Donato, M. Dove, H. Gharavy

25.1 Introduction

The Digital Signal Processing Group is carrying out research in the general area of signal processing. In addition to specific projects handled on campus, there is close interaction with Lincoln Laboratory and the Woods Hole Oceanographic Institution. While a major part of our activities focuses on the development of new algorithms, there is a strong conviction that theoretical developments must be closely tied to applications. We are involved with the application areas of speech, image, video, and geophysical signal processing. We also believe that algorithm development should be closely tied to issues of implementation because the efficiency of an algorithm depends not only on how many operations it requires, but also on how suitable it is for the computer architecture it runs on. Also strongly affecting our research directions is the sense that while, historically, signal processing has principally emphasized numerical techniques, it will increasingly exploit a combination of numerical and symbolic processing, a direction that we refer to as knowledge-based signal processing.

In the area of knowledge-based signal processing, there are currently two research projects. One involves the concept of symbolic correlation, which is concerned with the problem of signal matching using multiple levels of description. This idea is being investigated in the context of vector coding of speech signals. Symbolic correlation will entail the use of both symbolic and numeric information to efficiently match a speech signal with stored code vectors. The second project in this area deals with the representation and manipulation of knowledge and expressions in the context of signal processing. This work examines issues such as the representation of knowledge, derivation of new knowledge from that which is given, and strategies for controlling the use of this knowledge.

In the area of speech processing, we have, over the past several years, worked on the development of systems for bandwidth compression of speech, parametric speech modeling, time-scale modification of speech, and enhancement of degraded speech. Recently, a new model-based speech analysis/synthesis system was developed. This system is capable of high-quality speech production, and it is being used in several low- and mid-rate speech coding systems. Our newest speech coding system has achieved a bit rate of 4.8 kbps while maintaining high speech quality. Research continues on adaptive noise cancellation techniques in a multiple microphone environment. An approach based on maximum likelihood estimation has shown substantial improvements over previous techniques.

In image processing, several restoration and enhancement projects were completed recently. One project involved the estimation of coronary artery boundaries in angiograms. This research produced a more robust model of the coronary angiograms which, consequently, improved the estimates of the arterial dimensions. A second image processing project studied the removal of ghosts from television signals. This form of degradation is caused by multi-path channels and can be removed by the use of an appropriate inverse filter. The stable filter which results is in general non-causal and, therefore, some form of time reversal must be used to implement the filter. Other research included motion compensation for moving pictures, and magnitude only reconstruction of images.

In the area of geophysical signal processing, our research is focused on the transformation of side scan sonar data. In practice, this data is corrupted by a number of factors related to the underwater environment. Our goal is to explore digital signal processing techniques for extracting the topographic information from the actual sonographs. Concepts under study include the removal of distortions caused by towfish instability and reconstruction based on multiple sonographs taken from different angles.

We are pursuing a number of projects which are directed toward the development of new algorithms with broad potential applications. For some time, we have had a considerable interest in the broad question of signal reconstruction from partial information, such as, Fourier transform phase or magnitude. We have shown theoretically how, under very mild conditions, signals can be reconstructed from Fourier transform phase information alone. This work has been extended to the reconstruction of multidimensional signals from one bit of phase and, exploiting duality, zero-crossing and threshold crossing information. Current research includes reconstruction from distorted zero-crossings. In addition, the reconstruction from multiple threshold crossings is being studied. This problem has been shown to be better conditioned than reconstruction using only a single crossing. Also, we are examining the problem of narrowband signal detection in wideband noise. This project intends to compare several different techniques under a number of computational constraints. Research continues on relationships between information theory and stochastic estimation. We are exploring applications to statistical problems, iterative signal reconstruction, short-time analysis/synthesis, and parameter estimation.

With the advent of VLSI technology, it is possible to build customized computer systems of astonishing complexity for very low cost. However, exploiting this capability requires the design of algorithms which use few operations but have a high degree of regularity and parallelism, or can be pipelined easily. We are exploring directions which include systematic methods for designing multi-processor arrays for signal processing, isolating signal processing primitives for hardware implementation, and searching for algorithms for multidimensional processing that exhibit a high degree of parallelism. We are also investigating highly parallel computer architectures for signal understanding, in which a mixture of intensive computation and symbolic reasoning must be executed in an integrated environment.

25.2 Motion Compensation for Undersea Cameras

National Science Foundation Fellowship
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Jae S. Lim, Matthew M. Bace

Undersea cameras have been used for many years in undersea exploration, reconnaissance, and salvage operations to go where divers cannot. While these cameras have proven to be very useful, they are still somewhat limited by their motion stability.

Typically, an undersea camera is towed behind a boat at a depth of several feet above the ocean floor. The video from the camera is displayed on monitors in the boat. Because the camera is not rigidly attached to anything, motion instabilities arise due to several major sources. Waves on the ocean surface cause the boat and therefore the towed camera to move slowly up and down. Variations in the currents near the ocean floor lead to unpredictable changes in the pitch and roll of the camera. The combined motion produces very disturbing effects in the resulting video sequence, even to the point of inducing sea-sickness in the technicians viewing the display monitors. The variations in camera depth and angle are also a source of difficulty in producing large "mosaic" pictures from sequences of pictures taken from horizontally adjacent positions. The ideal camera would cruise at constant depth with a constant angle.

While it may be impossible to mechanically stabilize an undersea camera, recent work in the field of image processing indicates that it may be possible to process the video from the camera in such a way so that it appears as if the camera is stable. The goal of this research is to develop an algorithm for accomplishing this motion compensation. First, an estimate of the motion of the undersea camera will be obtained from the input video sequence. Then, the motion estimate will be used to compensate for the vertical and rotational components of the camera's motion so that only the horizontal component of the motion is present in the output video sequence.

25.3 Reconstruction Of Nonlinearly Distorted Images From Zero Crossings

National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Joseph E. Bondaryk

It has been shown theoretically that bandlimited, multidimensional signals can be specified to within a constant factor by the information contained in the locations of their zero crossings. It has been shown experimentally that two-dimensional, bandlimited signals can be reconstructed to within a constant factor from zero crossing information alone. The two-dimensional signals used were derived from images and their zero crossings corresponded to the threshold crossings of the images.

In this research, the problem considered is that of two-dimensional, bandlimited signals which have been affected by memoryless, nonlinear distortions. It is shown that such distortions retain the information required by the above theory, if they contain a monotonic region. Therefore, reconstruction to within a scale factor of an original, bandlimited image from the threshold crossing information of a distorted image is made possible. The zero crossing coordinates of the two-dimensional signal derived from a distorted image are substituted into a Fourier Series representation of the original signal to form a set of homogeneous, linear equations with the Fourier coefficients of the original signal as unknowns. The least squares solution to this set of equations is used to find the Fourier coefficients of the original signal, which are inverse Fourier transformed to recover the original image. By comparison of the distorted and reconstructed images, the nature of the distortion can be described. Some of the numerical problems associated with the reconstruction algorithm are also considered.

The reconstruction process is particularly stable for images which have bandedge Fourier components of high magnitude. It is shown that reconstruction from the zero crossing information of these images is similar to reconstruction from halftones, binary-valued images used to represent images which contain a continuous range of tones. This theory is finally extended to include reconstruction of bandlimited images from the zero crossings of distorted halftones.

25.4 Digital Processing of Side Scan Sonographs

National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Daniel T. Cobra

Since its introduction in the early sixties, side scan sonar has proved to be a very important tool for underwater exploration and,in particular, for marine geology. Its applications include surveying the sea floor, the search and location of objects on the bottom of the sea, and the prospection of mineral deposits.

The information contained in reflected sound waves is used by side scan sonar to produce a graphic record, called a sonograph, which constitutes a composite representation of the topographic features and the relative reflectivity of the various materials on the sea bed. Due to several factors, however, sonographs do not provide a precise depiction of the topology. Geometric distortions can be caused by motion instability of the towfish on which the transducers are mounted. This can be due to a number of factors, including variable ship speed and sea currents. The record can also suffer from interferences such as those caused by dense particle suspension in the water, shoals of fish, or by ultrasonic waves generated by passing ships. As a result, the interpretation of sonographs often requires extensive practice and can be a tedious and time-consuming task.

Our general goal is to explore the application of digital signal processing techniques to side scan sonar data, with the purpose of producing enhanced sonographs. At present, we are studying the specific problem of estimating and correcting the distortions caused by towfish instability. This project is being conducted under MIT's joint program with the Woods Hole Oceanographic Institution, with the cooperation of the U.S. Geological Survey.

25.5 Representation and Manipulation of Signal Processing Knowledge and Expressions

National Science Foundation Fellowship
National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Michele M. Covell

The phrase "signal processing" is used to refer to both "symbolic" and "numeric" manipulation of signals. "Symbolic" signal processing manipulates the signal description as opposed to the signal values with which "numeric" signal processing is primarily concerned. Efforts have been made to create computer environments for both types of signal processing.[1,2] Some issues that arise as a result of this work concern uniform representation of knowledge, derivation of new knowledge from that which is given, and strategies for controlling the use of this knowledge. This research will be concerned with these areas and how they apply to digital signal processing.

Representations that have been used in symbolic signal processing[1,2,3] have been largely distinct from those used in numeric signal processing.[1,4] The types of representations used are further separated by the control structures that the numeric and symbolic information commonly assume, the distinction essentially being the same as the distinction between Algol-like languages and logic programming languages. This dichotomy results from the differing amounts of available knowledge about appropriate approaches to the problems being addressed. By separating the control structure from application knowledge, this dichotomy can be avoided.

Strategies for controlling when knowledge about a signal is used should be provided and new strategies should be definable, since these control structures provide additional information about the problem space, namely, approaches that are expected to be profitable. Control strategies can also be used to outline new approaches to a problem, approaches that would not be considered by simple trigger-activated reasoning.

Finally, the ability to derive new knowledge from that which is given is desirable. This ability would allow the amount of information initially provided by the user to be minimized. The environment could increase its data base with new conclusions and their sufficient pre-conditions. Two immediate advantages of providing the environment with this ability are the reduction in the programming requirements and the possible "personalization" of the data-base. A reduction in programming requirements is available since information that is derivable from given information need not be explicitly encoded. Commonly, this type of information is provided to improve the performance of the derivation process. Secondly, since the environment would add information to the data set according to conclusions prompted by the user's queries, the data set would expand in those areas which the user had actively explored.

References

[1] C. Myers, *Signal Representation for Symbolic and Numeric Processing,* Ph.D. diss., MIT, 1986.

[2] W. Dove, *Knowledge-based Pitch Detection,* Ph.D. diss., MIT, 1986.

[3] E. Milios, *Signal Processing and Interpretation using Multilevel Signal Abstractions,* Ph.D. diss., MIT, 1986.

[4] G. Kopec, *The Representation of Discrete-Time Signals and Systems in Programs,* Ph.D. diss., MIT, 1980.

25.6 Iterative Algorithms for Parameter Estimation from Incomplete Data and their Applications to Signal Processing

National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Meir Feder

Many signal processing problems may be posed as statistical parameter estimation problems. A desired solution for the statistical problem is obtained by maximizing the Likelihood (ML), the A-Posteriori probability (MAP) or some other criterion, depending on the a-priori knowledge. However, in many practical situations the original signal processing problem may generate a complicated optimization problem, e.g., when the observed signals are noisy and "incomplete."

An iterative framework for maximizing the likelihood, the EM algorithm, is widely used in statistics. In the EM algorithm, the observations are considered "incomplete" and the algorithm iterates between estimating the sufficient statistics of the "complete data" given the observations and a current estimate of the parameters (the E step), and maximizing the likelihood of the complete data, using the estimated sufficient statistics (the M step). When this algorithm is applied to signal processing problems, it yields, in many cases, an intuitively appealing processing scheme.

In the first part of this research, we investigate and extend the EM framework. By changing the "complete data" in each step of the algorithm, we can achieve algorithms with better convergence properties. In addition, we suggest EM type algorithms to optimize other (non ML) criteria. We also develop sequential and adaptive versions of the EM algorithm.

In the second part of this research we examine some applications of this extended framework of algorithms. In particular we consider:

1. Parameter estimation of composite signals, i.e., signals that can be represented as a decomposition of simpler signals. Examples include:
   - Multiple source location (or bearing) estimation
   - Multipath or multi-echo time delay estimation
2. Noise cancellation in a multiple microphone environment (speech enhancement)
3. Signal reconstruction from partial information (e.g., Fourier transform magnitude).

The EM-type algorithms suggested for solving the above "real" problems provide new and promising procedures, and they thus establish the EM framework as an important tool to be used by a signal processing algorithm designer.

Portions of this work were supported in part by the MIT-Woods Hole Oceanographic Institution Joint Program.

25.7 Multi-Band Excitation Vocoder

National Science Foundation (Grant ECS 84-07285)
Sanders Associates, Inc.
U.S. Air Force - Office of Scientific Research (Contract F19628-85-K-0028)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Jae S. Lim, Daniel W. Griffin

The problem of analyzing and synthesizing speech has a large number of applications, and as a result has received considerable attention in the literature. One class of speech analysis/synthesis systems (vocoders) which have been extensively studied and used in practice are based on an underlying model of speech. For this class of vocoders, speech is analyzed by first segmenting speech using a window such as a Hamming window. Then, for each segment of speech, the excitation parameters and system parameters are determined. The excitation parameters consist of the voiced/unvoiced decision and the pitch period. The system parameters consist of the spectral envelope or the impulse response of the system. In order to synthesize speech, the excitation parameters are used to synthesize an excitation signal consisting of a periodic impulse train in voiced regions or random noise in unvoiced regions. This excitation signal is then filtered using the estimated system parameters.

Even though vocoders based on this underlying speech model have been quite successful in synthesizing intelligible speech, they have not been successful in synthesizing high-quality speech. For clean speech, the synthesized speech often exhibits a "buzzy" quality. For noisy speech, severe "buzziness" and other degradations often occur resulting in a large drop in intelligibility scores. The poor quality of the synthesized speech is, in part, due to the excitation models and the parameter estimation methods used in existing vocoders.

The Multi-Band Excitation Vocoder contains a speech model which allows the band around each harmonic of the fundamental frequency to be declared voiced or unvoiced. Accurate and robust estimation methods for the parameters of the new speech model were developed, as well as methods for synthesizing speech from the model parameters and methods for coding the speech model parameters. An 8 kbps vocoder was developed as well.

This 8 kbps Multi-Band Excitation (MBE) Vocoder was compared with a more conventional Single Band Excitation (SBE) Vocoder (1 V/UV bit per frame) in terms of quality and intelligibility. Informal listening indicated that the "buzzy" quality of the SBE Vocoder was eliminated by the MBE Vocoder with the improvement being most dramatic in noisy speech. Intelligibility tests (Diagnostic Rhyme Tests) for speech corrupted by additive white noise (approximately 5 dB SNR) produced an average score of 58.0 points for the MBE Vocoder, 12 points better than the average score of 46.0 for the SBE Vocoder. In addition, the average score for the MBE Vocoder was only about 5 points below the average DRT score of 63.1 for the uncoded noisy speech. This represents a much smaller intelligibility decrease in noise experienced by most vocoders.

This project was completed in March 1987.

25.8 Television Signal Deghosting by Noncausal Recursive Filtering

National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Daniel J. Harasty

A ghosting channel is typically modeled by a finite impulse response filter. Such a filter has an all-zero system function. A deghosting system, the inverse of the ghosting channel, mitigates the effects of the ghosting. For a nonminimum phase channel, that is, for a filter which has a discrete-time system function with zeros outside the unit circle, the deghosting system has poles outside the unit circle. Such a deghosting system has a stable, noncausal infinite impulse response, which can be decomposed into components which are strictly causal and strictly anticausal. Our research considered the implementation of a deghosting system.

A recursive filter can be used to implement the time-flipped impulse response of the anticausal component of the deghosting system. However, the input and output signals must also be flipped in time. Rather than flipping the signals in their entirety, piecewise time reversal is used: This algorithm partitions the signal into overlapping blocks which are independently time reversed in flip buffers. The overlap accommodates the transient response of the recursive filter implementing the time-flipped anticausal impulse response. A consequence of piecewise time reversal is that the anticausal impulse response is no longer infinite, but the flip buffer and overlap lengths can be chosen to provide an arbitrarily close approximation.

A discrete-time ghosting channel model, deghosting system decomposition, and piecewise time reversal were studied. The advantages of the direct form recursive filter over other filter structures were studied, as well as the simulation program and experimental results.

This research was made possible by the cooperation of the David Sarnoff Research Center with MIT through the VI-A Internship Program. This project was completed in August 1987.

25.9 A 4.8 Kbps Multi-Band Excitation Speech Coder

National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Jae S. Lim, John C. Hardwick

Recently completed research has led to the development of a new speech model. This model, referred to as the Multi-Band Excitation Speech Model, has been shown to be capable of producing speech without the artifacts common to model-based speech systems.[1] This ability makes the model particularly applicable to speech coding systems requiring high-quality reproduction at a low bit rate. In reference 2, a 9.6 kbps speech coder based on this model was first described. Later work resulted in an 8.0 kbps speech coding system. Both of these systems have been shown to be capable of high-quality speech reproduction in both low and high SNR conditions.

The purpose of this research is to explore methods of using the new speech model at lower bit rates. Results indicate that a substantial amount of redundancy exists between the model parameters. Recent research efforts have focused on methods of exploiting this redundancy in order to more efficiently quantize the model parameters. A 4.8 Kbps speech coder has been developed as part of this research. Preliminary tests have shown that this bit rate has been achieved with the same quality level as found in the previously developed coding system. Extensive evaluations of this new 4.8 kbps speech coding system are currently being conducted.

References

[1] D.W. Griffin and J.S. Lim, "A New Model-Based Speech Analysis/Synthesis System," IEEE International Conference on Acoustics, Speech, and Signal Processing, Tampa, Florida, March 26-29, 1985, 513-516.

[2] D.W. Griffin and J.S. Lim, "A High Quality 9.6 kbps Speech Coding System," IEEE International Conference on Acoustics, Speech, and Signal Processing, Tokyo, Japan, April 8-11, 1986.

25.10 Image Interpolation Using Edge Information

AT&T Bell Laboratories Doctoral Support Program
National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Jae S. Lim, Steven H. Isabelle

One application of interpolation in the area of digital image processing is in the increase of scale or resolution of still images. In this application, unknown sample values of a continuous 2-D image function must be approximated based on the known sample values of a digital image. This interpolated image is an approximation to the ideal digital image created by sampling the original continuous image on a more dense sampling grid. This approximated image should be in some sense close to the ideal digital image. Because pictures are meant to be viewed by humans, the idea of closeness of the interpolated image to the ideal image is related to the response of the human visual system. For human viewers, edges convey much of the information in images; therefore, an interpolation system should place emphasis on recreating edge information accurately.

This research is aimed at developing an image interpolation system which incorporates edge information. This edge orientation information is computed using a simple edge model along with local characteristics of the image. Preliminary results indicate that artifacts introduced by conventional approaches such as pixel averaging can be substantially reduced using this method.

25.11 Multi-Level Signal Matching for Vector Quantization

Canada, Bell Northern Research Scholarship
Canada, Fonds pour la Formation de Chercheurs et l'Aide a la Recherche Postgraduate Fellowship
Canada, Natural Science and Engineering Research Council Postgraduate Fellowship
National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Jacek Jachner

Our research investigates the use of multi-level signal representations, or hierarchies of signal representations with increasing levels of abstraction of detail, to perform efficient signal matching for Vector Quantization in a speech coder. The signal representations range from the numeric sequence that completely characterizes a signal, to high-level representations that abstract detail and group signals into broad classes.

The use of abstraction of detail in signal matching has been proposed in the context of a helicopter sound signature detection problem.[1] The current work focuses on low-bit-rate speech coding, for which recently proposed approaches, such as the Code Excited Linear Predictor (CELP) structure,[2] rely on a Vector Quantizer to code the residual signal after linear prediction stages. The residual is matched with one codeword signal out of a codebook of such signals, such that an error criterion between input and codeword signals is minimized. In the CELP coder, the error criterion is a time-varying linearly weighted mean-square error. The signal matching operations required for Vector Quantization (VQ) represent the major source of complexity in the CELP structure, and limit the practical size, hence performance, of VQ codebooks.

Our work seeks to reduce matching complexity by using multi-level signal representations simplified by abstraction of detail. The computational benefit of such signal representations is twofold. First, the matching operations on simplified or partial signal representations, termed partial errors, are substantially simpler to perform than applying the error criterion to the complete signal representation. Next, the simplified signal representations partition the codebook into equivalence classes of code words that share a partial signal representation value. The partial error is the same for all code words in a class and needs to be computed only once. The partial errors are useful when they provide sufficient information to eliminate some code words from further consideration. By matching using multiple signal representations with progressively greater representation detail, a branch-and-bound procedure[3] is used to prune the set of eligible code words. The computational advantage of this approach stems from the substitution of simpler partial error evaluations for the more complex application of the error criterion to the complete representation.

The formulation and evaluation of suitable multi-level representations and partial error functions for VQ in the CELP coder are the main directions of our current work. Performance simulations are conducted using the SPLICE[4] signal processing environment on LISP machines.

References

[1] E. Milios, *Signal Processing and Interpretation using Multi-level Signal Abstraction,* Ph.D. diss., MIT, 1986.

[2] M. Schroeder and B. Atal, "Code-Excited Linear Prediction (CELP): High-Quality Speech at Very Low Bit Rates," IEEE International Conference on Acoustics, Speech, and Signal Processing, 1985.

[3] E. Lawler and D. Wood, "Branch-and-Bound Methods: a Survey," Operations Research 14(4):699-719, 1966.

[4] C. Myers, *Signal Representations for Symbolic and Numerical Processing,* Ph.D. diss., MIT, 1986.

25.12 Detection of Narrowband Signal in Wideband Noise

Amoco Foundation Fellowship
National Science Foundation (Grant ECS 84-07285)
Sanders Associates, Inc.
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Tae H. Joo

The search for radio signals transmitted by extraterrestrials is a complex, multidimensional search problem because little is known about the transmitted signal. Current searches for extraterrestrial intelligence (SETI) collect data from a predetermined range of signals. These data are then processed to detect all synthetic components. (Synthetic components of signals are those which do not originate naturally. This assumes that the synthetic component is generated by extraterrestrials.) The assumption that the transmitted signal is a continuous wave (CW) at certain frequencies is commonly used in determining the range. Existing SETI systems use a frequency of 1450 MHz, atomic hydrogen line.

Due to uncertainties in the transmitter location, the relative velocities and the receiver antenna beamwidth, the frequency of the CW signal is unknown but is within 200 KHz of 1420 MHz. The propagation experiences multi-path which spreads the CW signal to a bandwidth of about 0.05 Hz. Therefore, SETI systems must search a wide frequency band (approximately 400 KHz) to detect a very narrowband (0.05 Hz) signal in poor signal-to-noise ratio (SNR) conditions.

Current SETI systems use FFT's to compute the spectrum. Each spectrum is then compared to a threshold to detect a peak. Because the SNR is low, the frequency bin size of the FFT is matched to the bandwidth of the narrowband signal. Therefore, a 2^{23}, or approximately 400 KHz/0.05 Hz, length FFT is required. In an existing system known as mega-channel extraterrestrial array (META)[1], this FFT is computed in two steps. First, the signal is filtered by 128 band-pass filters. Second, each band-pass filtered signal is transformed by a 64K length FFT. These computations are made using fixed point arithmetic. There are alternative implementations of this DFT-based method. The performance of different implementations, within constraints of the finite register length and other computational limitations, will be examined.

If the received signal is modeled as a sinusoid in white noise, modern spectrum estimators (e.g., the maximum entropy method) or frequency estimators (e.g., Pisarenko's method) can be employed. The performance and applicability of these algorithms, within constraints of computational limitations, will be examined.

Reference

[1] P. Horowitz, J. Forster, and I. Linscott, "The 8-Million Channel Narrowband Analyzer," In *The Search for Extraterrestrial Life: Recent Developments,* ed. M.D. Papagiannis, 361, 371. Hingham, Massachusetts: Kluwer Academic Publishers, 1985.

25.13 Estimation of Coronary Artery Dimensions from Angiograms

National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Jae S. Lim, Thrasyvoulos N. Pappas

A new approach was developed for the measurement of the severity of coronary obstructions from coronary angiograms. An angiogram is an x-ray picture of arteries in which a contrast agent has been injected. Existing techniques are heuristic and their performance is not satisfactory. Our approach exploited the characteristics of the signals involved. A model of the film density of the coronary angiograms was developed and used to estimate the diameter and cross-sectional area at each point along the vessel. Our model accounts for the structure of the vessel and background, as well as the distortions introduced by the imaging system. Both a one-dimensional and a two-dimensional model of the angiograms were studied. The algorithms were tested on synthetic data, on x-rays of contrast-medium-filled cylindrical phantoms, and on real coronary angiograms. Both algorithms were shown to have better performance than current methods. Moreover, the two-dimensional algorithm was shown to be better than the one-dimensional algorithm.

This research was completed in April 1987.

25.14 Chaotic Dynamics in Digital Signal Processing

U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Gregory W. Wornell

This research is aimed at exploiting the rapidly emerging theory of chaotic dynamical systems in problems of interest to the digital signal processing community. Chaotic dynamical systems are deterministic systems capable of exhibiting highly stochastic behavior. Due to their exponentially sensitive dependence on initial conditions, the state evolution of these systems becomes increasingly unpredictable. It is well-known that chaotic phenomena are not restricted to high-order systems; very complex, chaotic behavior can be observed in simple non-linear dynamical systems of arbitrarily low-order.

Present work is aimed at identifying and characterizing chaotic dynamics at work in some digital signal processing scenarios. For example, the effect of quantization in digital filtering is clearly a deterministic, non-linear phenomena typically modeled as noise-like. It is conceivable that a chaotic model of quantization effects (rather than an additive-noise model) may lead to improved techniques for reducing these effects.

Other phenomena under investigation as chaotics systems include speech generation, non-linear signal distortion, and signal modulation.

25.15 Image Texture Modeling

U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Jae S. Lim, Rosalind H. Wright

Textured regions are the nemesis of many image processing algorithms such as those used for segmentation or compression. Such algorithms rely on assumptions of stationarity or on a high degree of correlation of the image gray-level data. In textured regions, where correlation may exist between spatial patterns instead of between gray-levels, these assumptions are inadequate. It is desirable to develop a model which captures the texture's inherent pattern in a concise way.

One desirable property of a texture model is an ability to synthesize syntactically regular "structural" textures (e.g., a tiled wall) as well as the more random "stochastic" textures (e.g., shrubbery). Current models usually assume only one of these cases. This research has so far focused on the appropriateness of causal and non-causal 2-D Markov random fields as a model for image texture. The Gibbs distribution of statistical mechanics, which is equivalent to the Markov model, offers a milieu which is better suited to studying the regular neighborhood interactions which give textures their "texture" quality. The use of the (stochastic) Gibbs model in modeling and randomizing structural textures will be explored. The computation required to analyze and synthesize 2-D non-causal texture models is high, and encourages the examination of two more issues. The first issue involves modeling texture after reducing the number of gray-levels. The second issue involves determining the spatial resolution necessary to characterize the texture.

The goal of this research is to develop a model which is capable of synthesizing a continuum of structural and stochastic textures. Although analysis is a major part of the synthesis effort, the focus is on synthesis to see if one can characterize the texture "completely" as opposed to just extracting features which discriminate it from another texture. Such a model can later be used to justify features for texture classification and discrimination, as well as for controlled experimentation with human perception of visual pattern.

25.16 Reconstruction of Multidimensional Signals from Multiple Level Threshold Crossings

Fannie and John Hertz Foundation Fellowship
National Science Foundation (Grant ECS 84-07285)
U.S. Navy - Office of Naval Research (Contract N00014-81-K-0742)

Alan V. Oppenheim, Avideh Zakhor

It has been shown theoretically that under mild conditions multidimensional signals can be recovered from one-level crossings (e.g., zero crossings). However, the accuracy with which locations of the one-level crossings need to be specified is large enough to limit the applicability of such a method in many practical situations. In this research, we have found two major sampling strategies for reconstruction of signals from multiple-level crossings.

In our first approach, we extend new theoretical results in multivariate polynomial interpolation theory, in order to define a variety of semi-implicit sampling strategies. These strategies, which provide sufficient conditions for recovery of multidimensional signals from non-uniform samples on lines of rational slope, are ultimately applied to the problem of reconstruction from multiple-level crossings. Although these semi-implicit results are general enough to be used for recovery from signal crossings with arbitrary functions, they do not provide conditions for reconstruction of signals from an arbitrarily small number of thresholds. In order to circumvent this difficulty, we are taking a second approach which is implicit, and uses algebraic geometric concepts to find conditions under which a signal is almost always reconstructible from its multilevel threshold crossings.

A problem distinct from that of uniquely specifying signals with level crossings is that of developing specific algorithms for recovering them from level crossing information, once it is known that the signals satisfy the appropriate constraints. We have developed a variety of reconstruction algorithms for each of our two approaches, and demonstrated results for several images. Preliminary investigation of their quantization characteristics seems to indicate that the dynamic range and bandwidth requirements for representation of signals via multiple level threshold crossings lie in between those of Nyquist and zero crossing representation. Moreover, under certain circumstances, our semi-implicit and implicit sampling strategies become identical to Nyquist sampling. This bridges the gap between explicit, semi-implicit, and implicit sampling strategies, unifies seemingly unrelated sampling schemes, and provides a spectrum of sampling schemes for multidimensional signals.

This research was completed in October 1987.

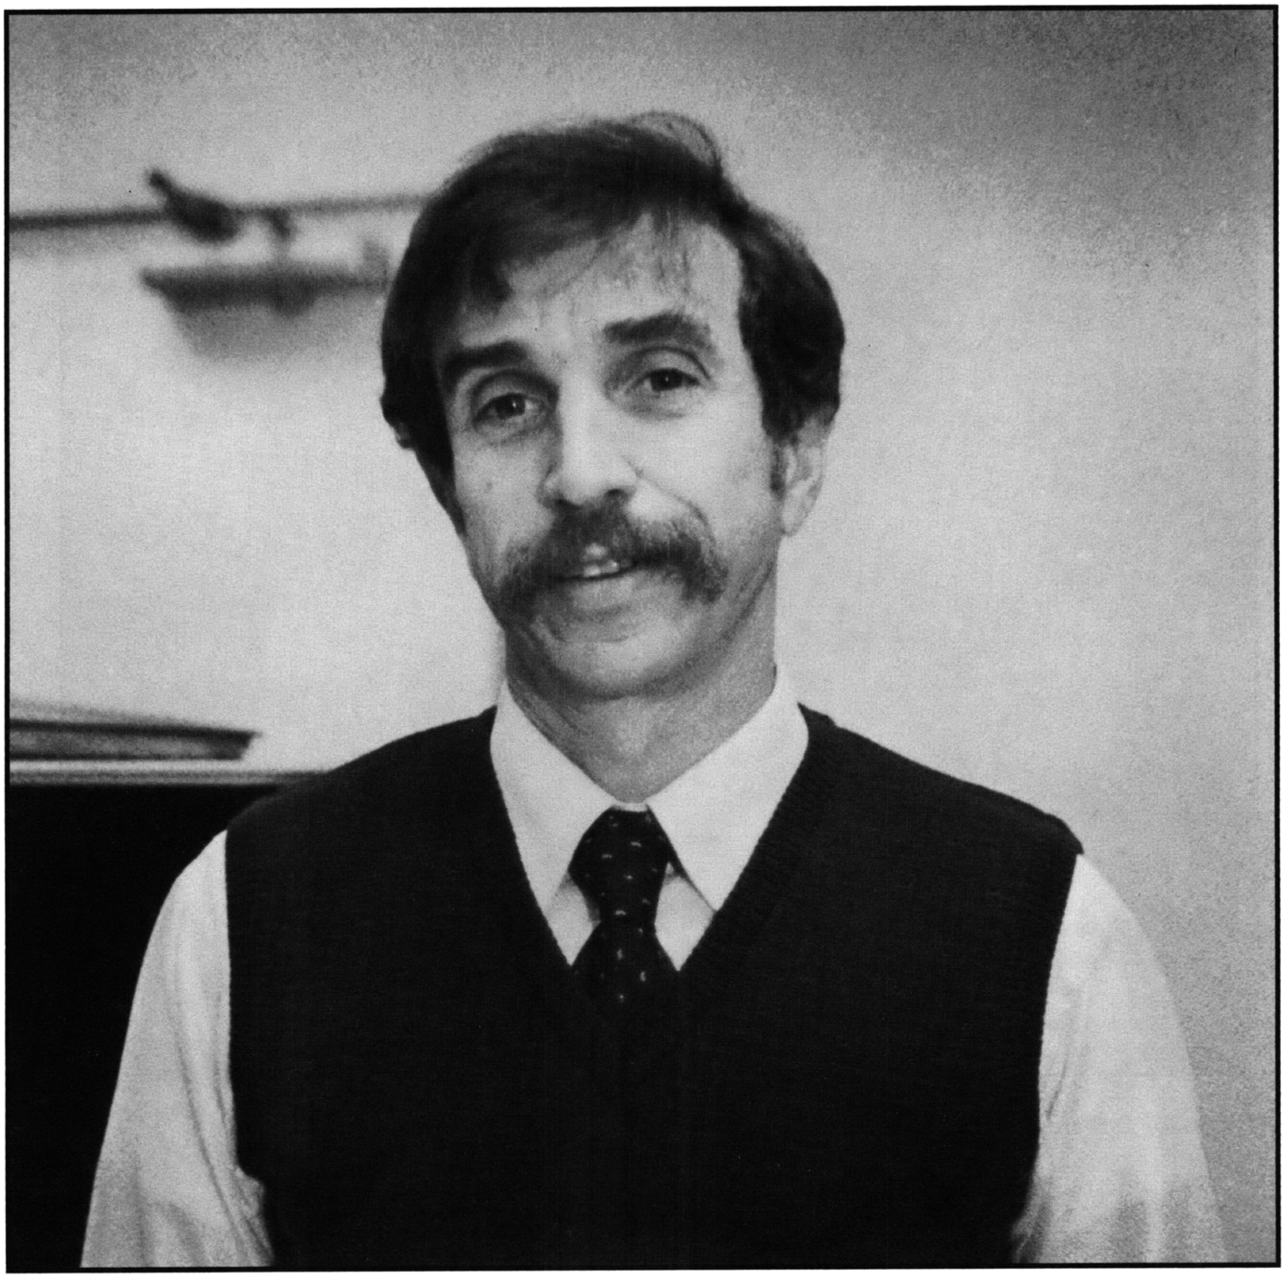

Professor Alan V. Oppenheim

Professor Donald E. Troxel