SIGNAL PROCESSING IN ACOUSTICS: SCIENCE OR SCIENCE FICTION?

James V. Candy
Lawrence Livermore National Laboratory and
University of California, Santa Barbara
Livermore, California 94551

> *"...to extract the desired information and reject the extraneous..."*

Signal processing in acoustics is based on one fundamental concept—extracting critical information from noisy, uncertain measurement data.[1-5] Acoustical processing problems can lead to some complex and intricate paradigms to perform this extraction especially from noisy, sometimes inadequate, measurements. Whether the data are created using a seismic geophone sensor from a monitoring network or an array of hydrophone transducers located on the hull of an ocean-going vessel, the basic processing problem remains the same—extract the useful information. Techniques in signal processing (e.g., filtering, Fourier transforms, time-frequency and wavelet transforms) are effective; however, as the underlying acoustical process generating the measurements becomes more complex, the resulting processor may require more and more information about the process phenomenology to extract the desired information. The challenge is to formulate a meaningful strategy that is aimed at performing the processing required, even in the face of these high uncertainties. This strategy can be as simple as a transformation of the measured data to another domain for analysis or as complex as embedding a full-scale propagation model into the processor.[5-7] For example, think of trying to extract a set of resonances (damped sinusoids) from accelerometer time series. It is nearly impossible to calculate zero-crossings from the time series but it is a simple matter to transform the data to the spectral domain using the Fourier transform and then applying the property that sinusoids are impulsive-like in Fourier space facilitating their extraction through peak detection. Finding a sinusoidal source propagating in the ocean is another matter that is quite complex due to the attenuation and dispersion characteristics of this harsh environment. Here, a complex propagation model must be developed and applied to "unravel" the highly distorted data to reveal the source—a simple Fourier transform will not work. The aims of both approaches are the same—to extract the desired information and reject the extraneous, and therefore, develop a signal processing scheme to achieve this goal. In this article, we briefly discuss this underlying signal processing philosophy from a "bottoms-up" perspective enabling the *problem* to dictate the *solution* rather than vise-versa. Once accomplished, we ask ourselves the final and telling question, "Did it work or are we kidding ourselves?" Are the results science or are they science fiction?

More specifically, signal processing (Note that throughout this article we will use the term "signal processing" to encompass *all* of the techniques used to extract the useful information from data. Such techniques as image processing, tomography, array processing, spectral processing, model-based processing, etc. are implied by this terminology) forms the basic nucleus of many acoustical applications. It is a specialty area that many acousticians apply in their daily technical regimen with great success such as the simplicity in Fourier analysis of resonance data or in the complexity of analyzing the time-frequency response of dolphin sounds. Acoustical applications abound with unique signal processing approaches offering solutions to the underlying problem. For instance, the localization of a target in the hostile underwater ocean acoustic environment not only challenges the acoustician, but also taxes the core of signal processing basics thereby requiring that more sophistication and *a priori* knowledge be incorporated into the processor. This particular application has led to many advances both in underwater signal processing as well as in the development of a wide variety of so-called model-based or physics-based processors. A prime example of this technology is the advent of the model-based matched-field processor[7-9] that has led not only to a solution of the target localization problem, but also to many applications in other acoustical areas such as nondestructive evaluation and biomedical imaging. So the conclusion is the same, signal processing is a necessary ingredient as a working *tool* that must be mastered by the acoustician to extract the useful information from uncertain measurements.

Acoustics

Let us look at acoustical signal processing from a slightly different perspective. Acoustical data can be used to extract useful information about signal sources, the surrounding environment and background noise much the same as any other modality (e.g., electromagnetics: radio frequency (RF), infrared (IR), optics, etc.). The information is clearly different but can also be used effectively. The uniqueness afforded is determined by how the acoustic signals propagate within the particular environment. The information available or carried by the acoustic signal (wave) depends heavily on the source characteristics and the environment supporting the propagation in which the wave interacts or causes these signals to bounce, bend and spread in a multitude of directions distorting both their shape and arrival times at sensor locations. Localization (incoherent) of the source is performed by estimating the arrival times (time delays) and using geometric relations (triangularization). The source characteristics also determine the underlying acoustical operational frequency band and therefore the sensors used to measure the propagated data with wavelengths inversely proportional to the source frequency. Clearly, compared to RF transmissions, acoustic waves propagate at long wavelengths and are affected by materials much differently.[1-4] Scattering of waves is determined by the object size at a wavelength or less. The point is that the advantage or disadvantage of acoustical waves over other modalities is determined by the acoustical properties of the source as well as the materials composing the propagation medium or environment. For instance, the audible acoustic range is the frequency band from 20 Hz to 20 kHz and is measured by microphone sensors, while seismic signals reside in the 0 to 10 Hz band measured by networks of geophones. Inaudible ultrasonic signals are typically in the 20 kHz to 20 MHz regime measured by piezoelectric crystal sensors.[1,3,4]

Acoustics can be used to perform the usual tasks of detection, classification and localization much the same as RF with sonar replacing radar in the active problem. For instance, acoustic sources can be localized by triangularization techniques much like finding the epicenter of an earthquake by a worldwide network of seismometers using their known location and the arrival times of the seismic event. Because of the large wavelengths in acoustic signals, arrays can be designed for *coherent* (phase) processing, that is, an array of acoustic sensors can be used to localize an acoustic source and passively scan an environment through beam steering while using beamforming techniques to search for sources. The human ear is a perfect example of passive (coherent) source localization in the audible range. In more complex environments, physics (model) based techniques can be used to enhance further the measured signal and perform the localization.[5-7]

The processing of acoustic data is necessary to extract the desired information from the noisy measurements. For example, if a voice recording was available, then the microphone data could be processed to separate the voice signature from the environmental background and noise thereby decomposing the acoustic information of each. Kennedy assassination data was thoroughly analyzed from microphone data acquired from available audio and video/audio recordings. The number and location of shots fired were extracted from the recordings.[2] The frequency characteristics of the sounds were analyzed even further to extract information about the environment (time-frequency estimation). Active acoustics (sonar) is typical in many applications ranging from the underwater tracking of a moving target, to locating tumors in tissue for biomedical applications, to imaging materials using an acoustic microscope. Vibrational acoustic signals provide critical information about structures and their condition in structural integrity studies. Thus, acoustical data much like RF or radar data uniquely provides information about the source, background environment, and noise that can be processed to extract useful information depending on the source characteristics and the supporting propagation medium as well as objects populating that particular environment (e.g., urban environment with buildings).

Signal processing approach

Signal processing relies on any *a priori* knowledge of the phenomenology generating the underlying measurements. Characterizing this phenomenology and propagation physics along with the accompanying measurement instrumentation and noise are the preliminaries that all acousticians must tackle to solve such a processing problem. In many cases this is much easier said than done. The first step is to determine what the desired information is and typically this is not the task of the signal processor, but that of the acoustician performing the study. In our case, we assume that the investigation is to extract information stemming from acoustic signals either emanating from a source whether it be an autonomous unmanned vehicle (AUV) passively operating in the deep ocean or a vibrating structure responding to ground motion. Acoustic applications can be very complex especially in the case of ultrasound propagating through complex media such as tissue in biomedical applications or through heterogeneous materials of critical parts in nondestructive evaluation (NDE) investigations.[1,3,4] In any case the processing usually involves manipulating the measured data to extract the desired information, such as, location and tracking of the AUV, to failure detection for the structure, or tumor/flaw detection and localization in both biomedical and NDE.[6,7]

Another view of the same problem is to decompose it into a set of steps that capture the strategic essence of the processing scheme. Inherently, we believe that the more *a priori* knowledge about the measurement and its underlying phenomenology we can incorporate into the processor, the better we can expect it to perform—as long as the information that is included is correct. One strategy, called the "model-based approach," provides the essence of model-based signal processing.[6,7] Some believe that all of the signal processing schemes can be cast into this generic framework. Simply, the model-based approach is "incorporating mathematical models of both physical phenomenology and the measurement process (including noise) into the processor to extract the desired information." This approach provides a mechanism to incorporate knowledge of the underlying physics or dynamics in the form of mathematical propagation models along with measurement system models and accompanying uncertainties such as instrumentation noise or ambient noise as well as model uncertainties directly into the resulting processor. In this way the model-based processor enables the interpretation of results directly in terms of the problem physics. The model-based processor is really an acoustic modeler's tool enabling the incorporation of any *a priori* information about the particular application problem to extract the desired information. As depicted in Fig. 1, the fidelity of the incorporated model determines the complexity of the processor. These models can range from simple implied non-physical representations of the measurement data such as the Fourier or wavelet transforms to parametric black-box models used for data prediction, to lumped mathematical physical representations characterized by ordinary differential equations, and to full physical partial differential equation models capturing the critical details of the acoustic wave propagation in a complex medium. The dominating

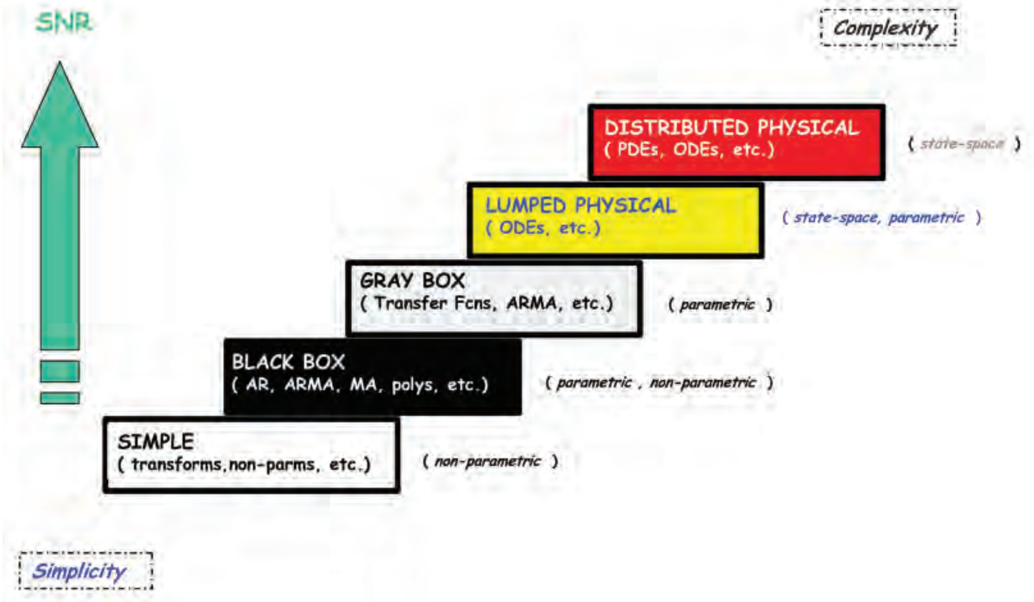

Fig. 1. Signal processing approach: the model-based "staircase." [Step 1] "Simple" non-parametric implicit models. [Step 2] "Black-box" models (transfer function, autoregressive, moving average, polynomial, etc.). [Step 3] "Gray-box" models (transfer function, autoregressive moving average, etc.). [Step 4] "Lumped" physical ordinary differential equation models (state-space, parametric, etc.). [Step 5] "Distributed" physical partial differential equation models (state-space, etc.).

factor of which model is the most appropriate is usually determined by how severe the measurements are contaminated with noise and the underlying uncertainties encompassing the philosophy of "letting the problem dictate the approach." If the signal-to-noise ratio (SNR) of the measurements is high, then simple non-physical techniques can be used to extract the desired information. This approach of selecting the appropriate model is depicted in the *signal processing staircase* of Fig. 1 where we note that as we progress up the "modeling" *steps* to increase the SNR, the complexity of the model increases to achieve the desired results. In the subsequent sections of this article, we will use the model-based framework to explain the various classes of acoustical signal processing problems and attempt to show—even at a simple level—how these schemes can evolve within this framework. This is our roadmap.[5-7]

Signal processing steps

We start with the simple first step and show how we can progress up the signal processing staircase to analyze a problem. Suppose we have a noisy acoustical measurement (Fig. 2a) of a single oscillation frequency in random noise (SNR = 0 dB, i.e., equal values of signal to noise) and we would like to extract the desired information (the single oscillation frequency). Our first "simple" approach to analyze the measurement data would be to take its Fourier transform and investigate the various frequency bands for resonant peaks. The result is shown in Fig. 2b, where we basically observe a noisy spectrum and a set of potential resonances—but nothing really conclusive. Next we apply a broadband power spectral estimator with the resulting spectrum shown in Fig. 2c. Here we note that the resonances have clearly been enhanced and appear in well-defined bands while the noise is attenuated by the processor, but there still remains a significant amount of uncertainty in the spectrum due to all of the resulting spectral peaks. Upon seeing these resonances in the power spectrum, we might proceed next to a gray-box model to enhance the resonances even further by using our *a priori* knowledge that there is essentially one dominant resonance we seek. The results of applying this processor are shown in Fig. 2d.

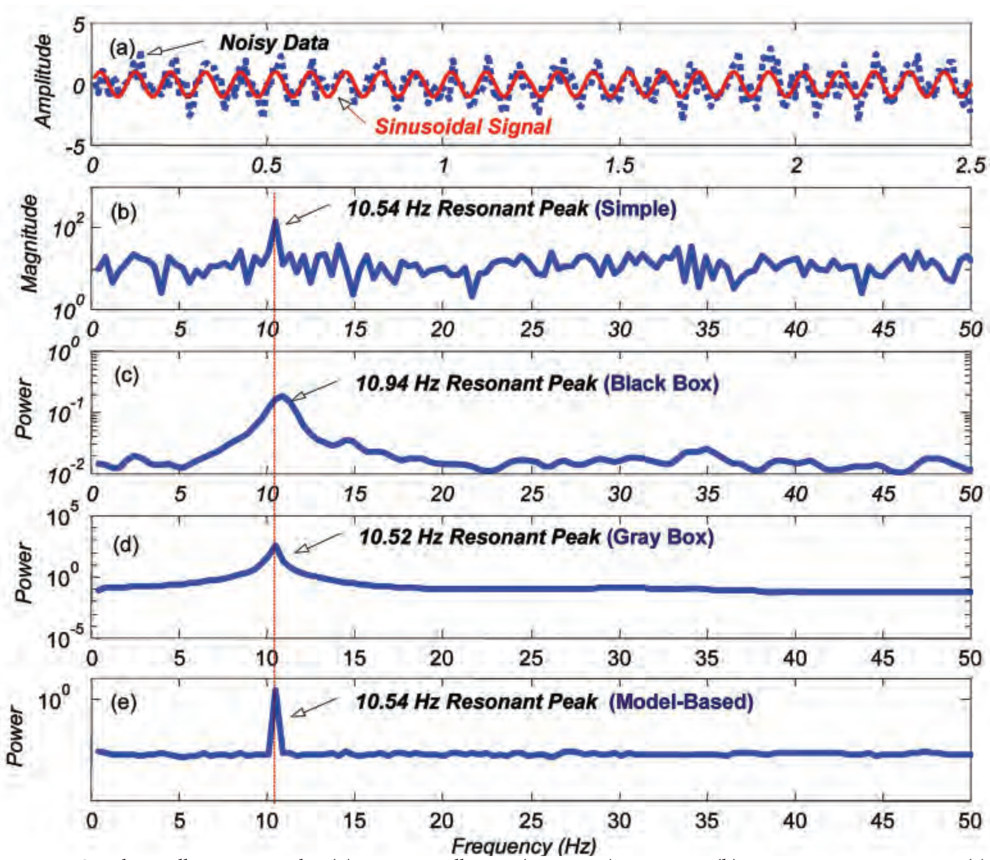

Fig. 2. Simple oscillation example. (a) Noisy oscillation (10.54 Hz) in noise. (b) Raw Fourier spectrum. (c) Nonparametric spectrum (black-box). (d) Parametric spectrum (gray-box). (e) Model-based spectrum (ordinary differential equation model).

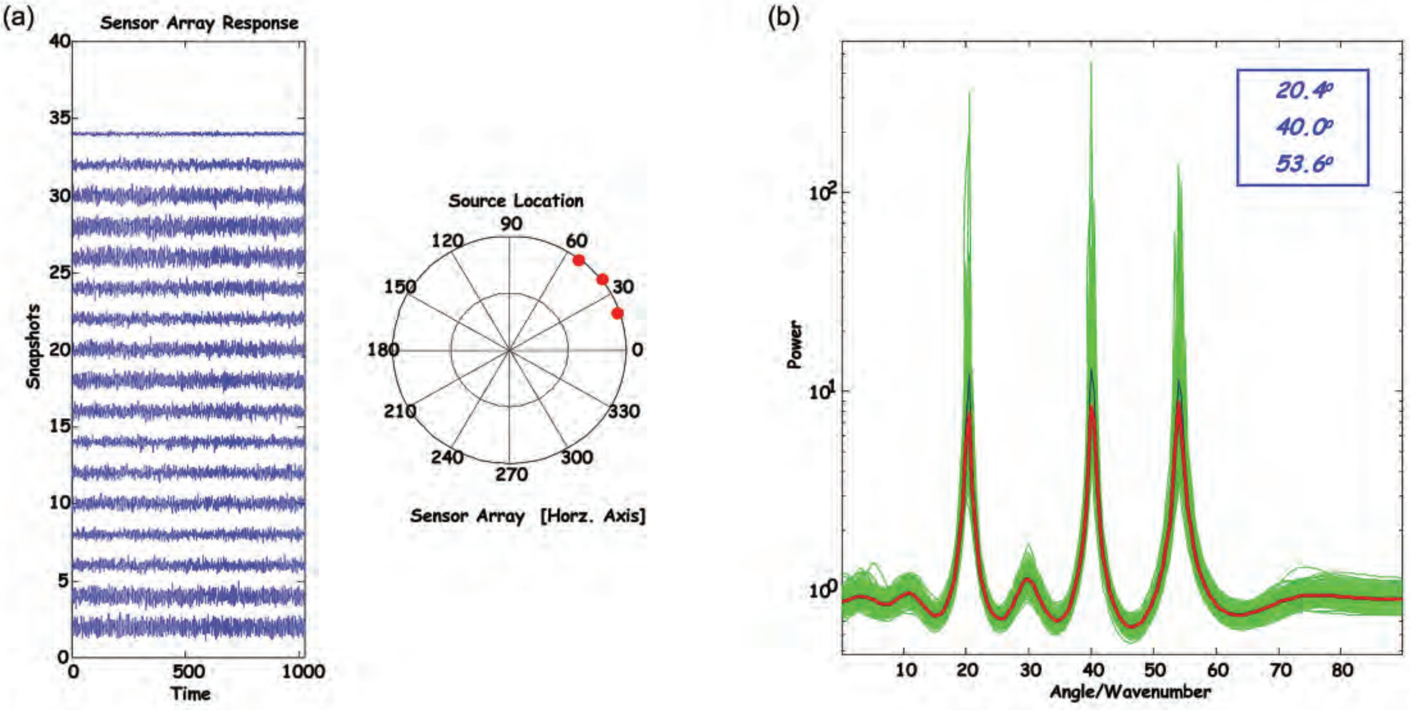

Fig. 3. Modern spectral (spatial) estimation using the maximum entropy method parametric estimator for direction-of-arrival estimation of three plane waves at 20°, 40°, 55° arrival angles; (a) Synthesized 16-element array temporal measurements (signal-to-noise = 0 dB) with true source locations; (b) Maximum-entropy-method spatial spectral estimation results with estimated arrivals at 20.4°, 40.0°, 53.6° generated from an ensemble of 256 realizations.

Finally, we use this extracted model to develop an explicit model-based processor (MBP) by developing a set of harmonic equations (ODE) for a sinusoid in noise and construct the MBP based on these relations. The results are shown in Fig. 2e. So we see that once we have defined the acoustical problem, and assessed the *a priori* information including the underlying phenomenology, then we can proceed up the staircase and exit any time we are satisfied with the result. This is the "bottoms-up" approach.[7]

Next we choose to investigate this approach in more detail by selecting some processing areas with an accompanying set of acoustical applications. Here we illustrate not only the fundamental approach to problem solving, but also to observe some of the popular processing paradigms available to the acoustician for analysis and information extraction. Our aim is to define a specific problem that represents a class of problems and then show some of the potential signal processing solutions available, demonstrating them through simulation or experiments.

Step 1: Simple spectral estimation techniques

Classical spectral analysis is a very powerful example of a set of tools that have evolved in signal processing especially in acoustics. Here a raw measurement is "transformed" to the spectral or Fourier domain for analysis. Modern techniques of spectral estimation can be considered both "black or gray-box" processors and even physics-based processors depending on the underlying application.[5] We call the black-box/gray-box methods "parametric processors," since they employ a variety of underlying model sets to achieve their enhancement and improved spectral estimation. Thus, the parametric spectral estimator consists of a processor employed to estimate the parameters of the underlying model set and then perform a power spectrum estimate using this model.

Modern spectral analysis techniques easily extrapolate to the space-time domain as long as we assume that the incoming wave front is separable in space (array) and time (or frequency). In this context, we can consider a measurement array as a spatial sampler of the arriving wave front. If we further assume that the temporal portion of the wave is restricted to a narrow frequency band, then we collapse its temporal response to a single frequency line (in the Fourier space) that can also be considered a parameter. So we see that estimating the arrival angle in the case of a planar wave front or the source location in the case of a spherical wave front can be considered a problem of "spatial" spectral estimation and all of the usual modern techniques (with some restrictions) apply to the array signal processing problem as well.[5,7]

In acoustics, a large set of problems reduce to array processing or spatial spectral estimation in this context. Such problems as ocean acoustic (sonar) signal processing for target direction-of-arrival (DOA) estimation or localization fall into this category along with ultrasonic NDE and biomedical processing searching for flaws or abnormalities.[1-4] Clearly, seismic array processing, of which most of these ideas evolve, is a root application of arrays for epicenter location and velocity estimation. With this information at hand, let us consider a simple example of a plane wave impinging on a sensor array to convey these ideas using a concrete example.

Spectral (spatial) estimation for direction-of-arrival

Suppose a 16-element linear array with acoustic sensors uniformly spaced at 2.5m is impinged upon by a set of plane waves generated from three (3) sources emanating from 20°, 40°, and 55° incidence angles. The temporal frequency is

300 Hz with corresponding propagation speed of 1500 m/sec. The sensor data are generated at 0 dB SNR. The results of applying a modern parametric estimator (maximum entropy method (MEM))[5] developed for harmonic (sinusoidal) waves in noise are shown in Fig. 3, where we observe the ensemble results of the 16-sensor channel spectral estimates. The results demonstrate that the algorithms are capable of providing reasonable DOA estimates in such a noisy environment.

Steps 2 & 3: Parametric signal processing (black/gray-box approach)

Perhaps even a more reasonable application of modern signal processing follows directly from the black/gray-box or parametric approach. In this realm of processing the acoustician can choose from a set of models that reasonably approximate the underlying phenomenology and essentially "fit" the model to the data through a variety of estimation algorithms. This type of processing evolves from the signal processing literature, from applications in speech (e.g., linear prediction, coding, recognition, etc.),[10] and controls called system identification (e.g., adaptive control, noise cancelling, etc.).[11] In this domain the parametric approach is to: (1) select a representative model set (e.g., transfer function, autoregressive (all-pole), moving average (all-zero), autoregressive moving average (pole-zero), state-space, etc.); (2) estimate the model parameters from the data; and (3) construct the signal estimate (e.g., spectrum, impulse response, etc.) from these parameters.[7] Again this can represent the "black-box" step if the parameters have *no* physical interpretation or the "gray-box" step if they do have physical relations.

Parametric signal processing for prosthetic heart valve classification application

As a parametric processing application, consider the problem of estimating flaws or cracks in prosthetic heart valves (see Fig. 4a). By placing microphones on a patient's chest and listening to the sounds radiated by the valve, its condition can be determined. From the structure of the prosthetic value and its interacting components, it is possible to isolate the sounds associated with each component and classify potential problems. Since these sounds are essentially vibrational resonances, an all-pole (autoregressive) model is selected to perform parametric signal processing and investigate the condition of the valve under test through a variety of statistical tests. Note how the acoustical problem is dictating the processing approach and potential solution. The processing is illustrated in Fig. 4. The approach selected in this application is to construct a classifier to determine in which class (failure or normal) the valve belongs (Fig. 4b). The heart

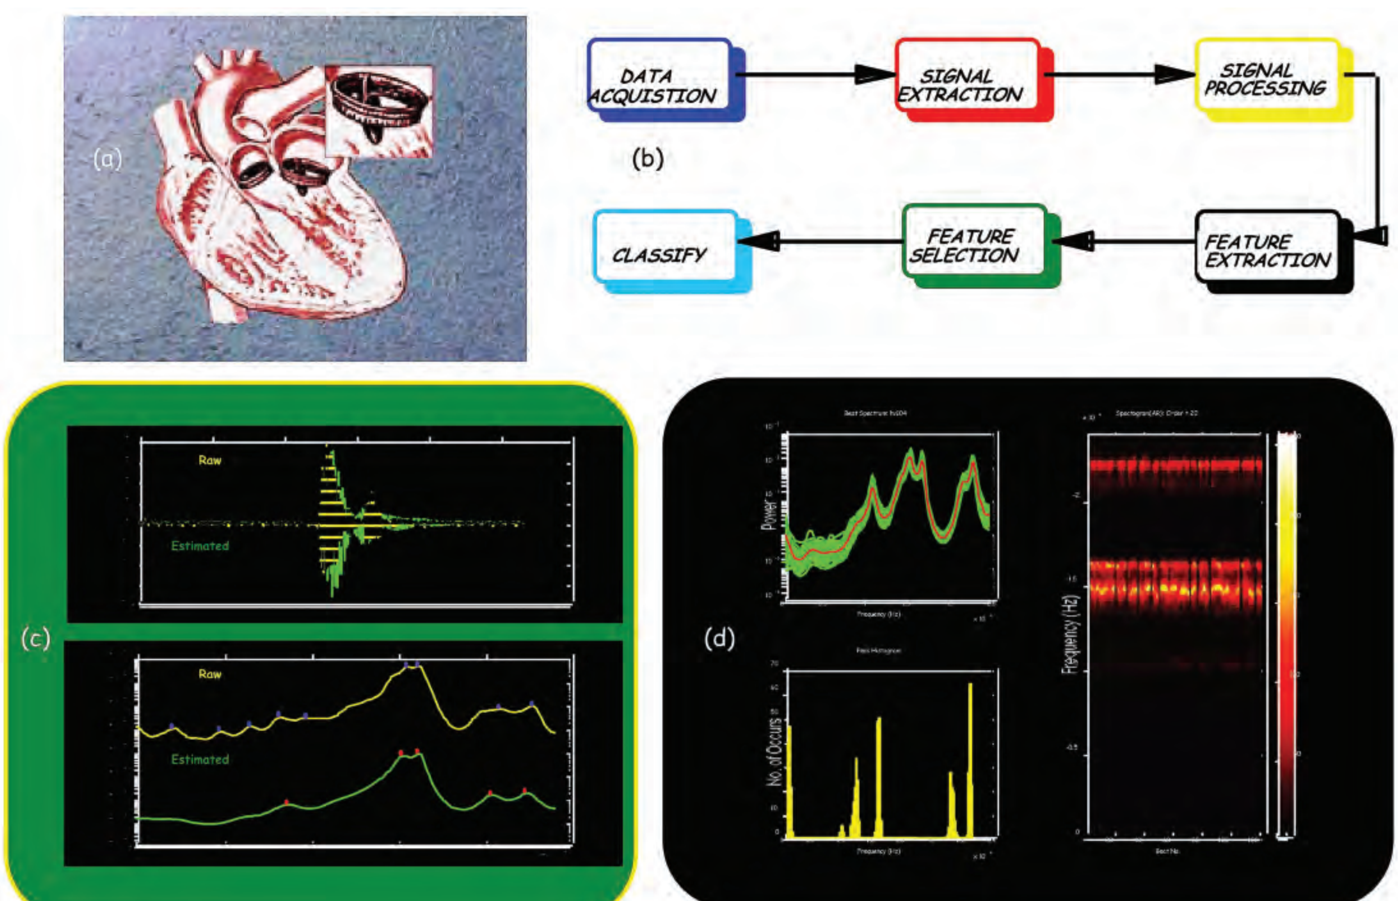

Fig. 4. Prosthetic heart valve acoustic analysis and condition detection using parametric signal processing (black/gray-box) techniques: (a) prosthetic heart valve under test; (b) overall processing paradigm to detect and classify condition; (c) parametric (all-pole) model and estimated signal/spectrum; and (d) ensemble spectrum, instantaneous spectrogram (power versus time versus frequency) image, and peak frequency probability distribution or histogram (feature) for condition classification.

valve radiates a sound with each beat that is measured by the microphone(s) positioned on the patient's chest. The noisy measurement data is enhanced and initially analyzed (Fig. 4c) and then processed further. Here a spectrum and instantaneous spectrogram are estimated from the model parameters and displayed in Fig. 4d. The frequency peaks in the spectrogram are estimated, and a peak resonant frequency probability distribution histogram is constructed to provide a feature vector that can be used to perform the classification determining the valve's condition.[12]

Step 4: Model-based processing (lumped physical approach)

Model-based signal processing[7] is the next step in the signal processing approach. It has distinct advantage over other approaches because the processor is developed directly in the acoustician's frame-of-reference, that is, in his phenomenological space. Not only does it require physics-based models of the underlying phenomenology, but it also requires knowledge of the measurement instrumentation and noise processes to construct a good processor. Here the acoustician is "thinking" directly in terms of the acoustics and not inferring results from a variety of analysis tools (e.g., Fourier and wavelet transforms) that are not directly related to the underlying propagation physics. The model-based approach is simply incorporating mathematical models of the underlying phenomenology including measurements and noise into the processing scheme—this is exactly how the acoustician becomes an integral part of the processing by providing the acoustic models that are embedded into the model-based processor (MBP). Not only is there direct benefits of thinking in the same physical coordinates of the acoustic problem, but also gaining a deeper understanding of the instrumentation performance and noise/uncertainty processes that contaminate the problem. This is the good news. Of course, if the model is inaccurate and does not represent the phenomenology adequately, then the results can be erroneous and sometimes very misleading (science or science fiction?). Fortunately, many of these processors provide "self-checking" validation tools such as residuals (difference between the measurements and model predictions) that can be statistically tested for validity and used by the processor to assure accurate performance.[7] With this in mind, we introduce two of the most popular model-based approaches in the literature: model-based matched-field processing[8,9] (imaging) and model-based recursive (in-time or in-space or both) processors (e.g., MBP or Kalman filters).[7]

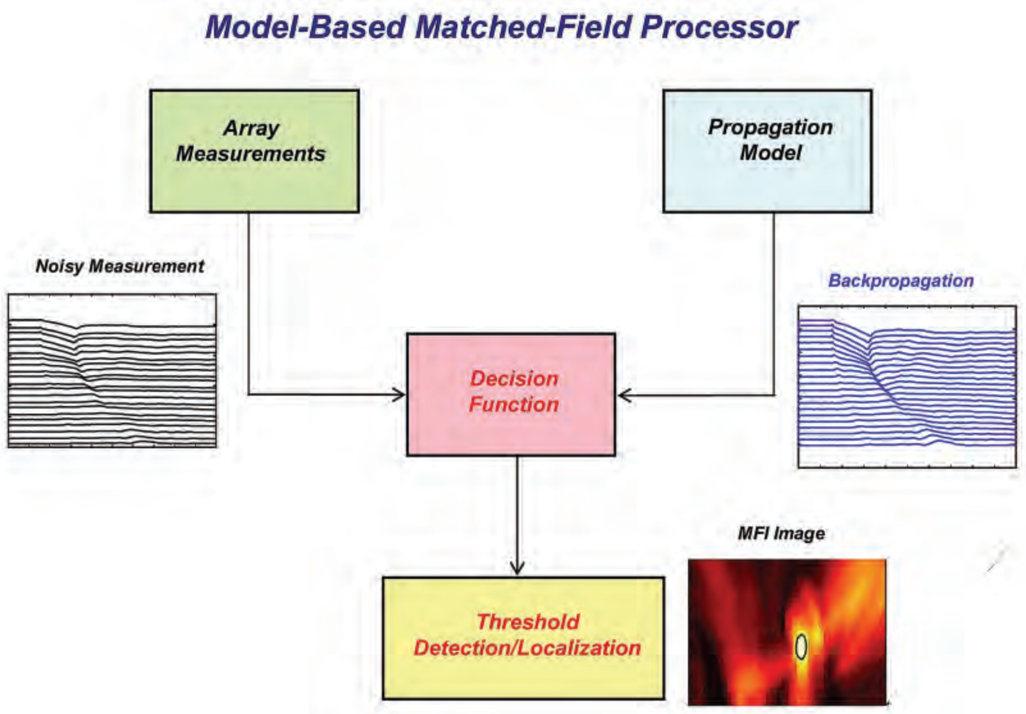

Fig. 5. Model-based matched-field processor for acoustical imaging applications. Raw array field measurements and backpropagated model measurements are compared (decision function) creating a power image that is thresholded for detection and localization.

Model-based matched-field imaging for nondestructive evaluation

Typical image processing techniques in acoustics consist of pre-processing the raw data to provide enhanced signals as input to the image formation algorithm as well as post-processing of the two-dimensional image to enhance, extract, and classify certain features of high interest. In this article, we concentrate primarily on the same theme that we have used throughout, the development of processors that incorporate more and more *a priori* information about the acoustics generating the data and its incorporation into a model-based imaging algorithm.

We saw in the previous example of a plane wave impinging on an array, how modern spatial spectral estimators (beamformers) can be used to estimate the wave's spatial and temporal spectral features. The model-based approach uses all of the *a priori* information about the plane wave propagation and noise measurements to extract the parameters directly solving the problem. The same idea can be extrapolated to the imaging problem. We assume that we have an array of sensors either physical or synthetically created, and we have developed a sequence of measurements resulting by exciting the medium under investigation. For instance, it can be an active sonar system in the ocean or an ultrasonic scanner in biomedical or nondestructive evaluation (NDE), or a passive array listening to a surveillance volume for passing airborne targets.

Here we consider the acoustic application of a laser ultrasound experiment for the NDE of an aluminum part. Our first approach is to apply the synthetic aperture focus technique (SAFT) to image the part under investigation.[13,14] We assume that the flaws can be characterized by acoustic point scatterers in the near-field. Therefore, spherical wave fronts impinge on our measurement array emanating from these flaws. The SAFT approach *creates* an image by assuming that the flaw location is at the given pixel, *calculates* the associated propagation delays and attenuations assuming a homogeneous medium, *beamforms* the measured data based on these assumptions and *estimates* the power in the beam at the assumed location (pixel or point scatterer). This procedure is repeated for each pixel until the observed power image is formed. Of course, this methodology can also be applied to image the results of time-reversal focusing and decomposition of the operator for localization purposes.[15]

Another more acoustics-oriented approach for imaging is based on replacing the beamformer with a propagation model. The same scheme (as above) applies, but the propagation model generates (backpropagation) the equivalent signal at the array and a criterion is created to "decide" whether a flaw is at a given test location (pixel or point scatterer).[7] The propagation model can be as sophisticated as deemed necessary incorporating features such as both compressional and shear waves, multipath, dispersion, and noise. This model-based technique is called *matched-field imaging* (MFI) and enables the acoustician to use the *a priori* information available in a formal procedure to create the image.[16,17] MFI is illustrated in Fig. 5 where we observe the measurements and propagation model generating the components of a decision function used to "detect" or "localize" the position of a flaw (or target).[7] Note that in the figure the plots are actual laser ultrasound array measurements. The object is a hole (flaw) drilled into an aluminum plate and measured using a synthetic aperture technique created for processing and detection.[14]

Consider a typical laser ultrasonic application where a NDE is performed on a rectangular aluminum part with two flaws. The SAFT and MFI images are shown in Fig. 6. We note that the MFI approach incorporates both compressional and shear wave fronts as well as the multipath caused by the part boundaries. The results of estimating the power at each pixel is shown where we see the high resolution and accurate results of the MFI compared to those of the SAFT processor.

Recursive (in-time) model-based processing for a towed acoustic array

Consider the following application taken from ocean acoustics to motivate the modern recursive model-based approach. Suppose we have a plane wave signal characterizing an acoustic source measured by a horizontally towed array. The plane wave is at a 50 Hz temporal frequency and a bearing of 45° impinging on a 2-element towed array at a 10 dB SNR with a tow speed of 5 m/sec. We would like to solve the problem of extracting the source bearing and temporal frequency parameters—the critical information we seek. From the model-based perspective, the bearing/frequency estimation or equivalently, localization problem can be considered a problem of estimating a set of parameters from noisy hydrophone measurements.[18,19]

The *classical* approach to this problem is to assume that the signal is separable in space and time and select a single sensor channel to perform spectral analysis/peak detection on the time series to estimate the temporal frequency parameter. The bearing can then be estimated independently by performing spatial spectral estimation (as before) or beamforming on the array data. A beamformer can be considered a spatial spectral estimator that is scanned over bearing angles indicating the true source location at the bearing of maximum power. The result of applying this approach to our problem is shown in Fig. 7a. The figure illustrates the classical outputs of both spectral estimators peaking at the correct frequency and bearing angle parameters.

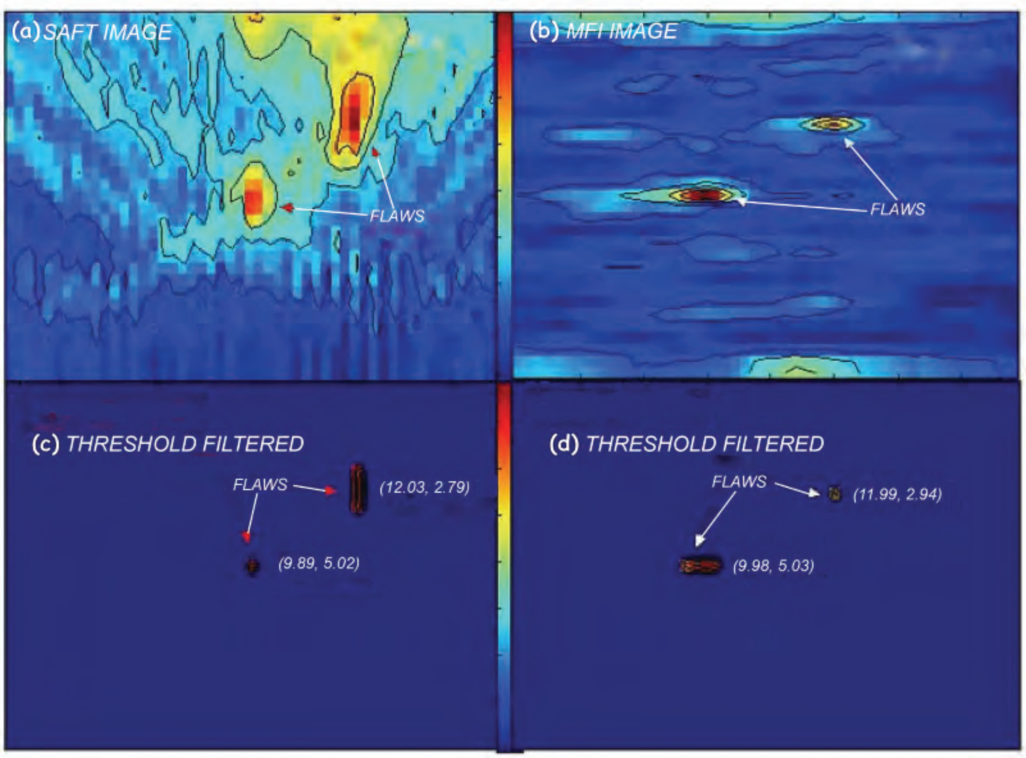

Fig. 6. Acoustical imaging of laser ultrasound flaw detection for NDE of aluminum part: (a) SAFT imaging; (b) model-based matched-field imaging; (c) thresholded SAFT image for flaw localization [(12.03 mm, 2.79 mm),(9.89 mm, 5.02 mm)]; (d) thresholded model-based matched-field image for flaw localization [(11.99 mm, 2.94 mm), (9.98 mm, 5.03 mm)].

The MBP is implemented by incorporating the plane wave propagation, hydrophone array, and noise models. However, the temporal frequency and bearing angle parameters are unknown and must be estimated. The solution to this problem is obtained by augmenting the unknown parameters into a MBP state-space structure and solving the so-called joint estimation problem.[7] This is the parametrically adaptive form of the MBP used in most ocean acoustic applications.[18] Here the problem becomes nonlinear due to both the measurement model (plane wave) and the augmentation. It is

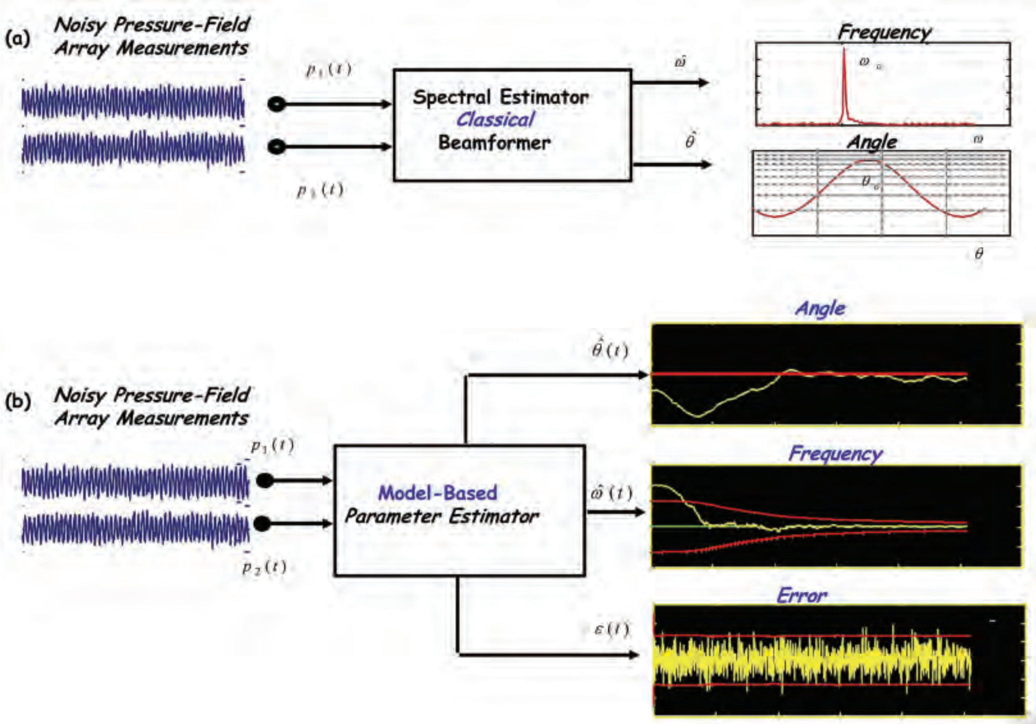

Fig. 7. Plane wave impinging on a 2-element hydrophone array: Frequency and bearing estimation problem: (a) classical spectral (temporal and spatial) estimation approach; (b) Model-based approach using parametrically adaptive (nonlinear) processor to estimate bearing angle, temporal frequency and the corresponding residual or innovations sequence.[18,19]

we have discussed or time/space varying (e.g., ocean) where we construct the processors using recursive-in-time or recursive-in-space techniques or both to capture the ever-changing medium or motion of the acoustic problem that must be solved.

Model-based tomographic imaging for biomedical imaging

Tomographic imaging is governed by partial differential propagation equations evolving from a full-field wave model to generate its solutions.[20,21] It can be thought of as a methodology of solving an inverse problem. Consider the development of a tomographic image of breast tissue to determine if a cancerous tumor is or is not present. This technology is based on Fourier imaging techniques and diffraction tomographic reconstruction of the tissue using multiple scans through the breast (object). Fourier imaging essentially obtains projections in object space using Fourier transforms (1 dimensional) to "fill in" regions in the 2 dimensional (2D) Fourier space and is based on the Fourier diffraction theorem equating a projection (plane wave propagation) in object space to a region (arc) in 2D-Fourier space. Once the space is filled, a 2D inversion is performed using the model to perform a backpropagation similar to a numerical time-reversal to reconstruct the object and create the image. This is similar to computer-aided tomographic (CAT) reconstruction employing x-rays using the Fourier slice theorem along a line, but here the acoustic propagation is modeled by the wave equation.[20,21]

The application of this model-based (distributed) approach to the breast tissue is shown in Fig. 8.[22] Here we observe the usual ultrasonic reflection image (not tomogra-

more computationally intensive than the spectral estimators used in the classical approach; however, the results are appealing as shown in Fig. 7b. We see the bearing angle and temporal frequency estimates as a function of time converging to the true values of 50 Hz and 45° bearing angle. The MBP also produces the "residual or innovations" sequence, (shown in the figure) that is used in determining its overall performance for validation. In this case the sequence must be statistically zero-mean and white (uncorrelated) for optimal performance.[7]

Thus, the classical approach simply performs spectral estimation temporally and spatially (beamforming) to extract the parameters from noisy data, while the model-based approach embeds the unknown parameters into its propagation, measurement, and noise models enabling a solution to the joint estimation problem. The performance of the MBP is validated by analyzing the statistics of its innovations sequence. This completes the application.

Step 5: Model-based processing (distributed physical approach)

This step can be the most complex depending on the embedded model representation. Typically, the phenomenology is represented by a set of partial differential equations characterizing the propagation medium and a set of noisy measurement equations that still can be characterized by a state-space representation in some form or another. Whether the application is based on time-invariant statistics like most of the applications

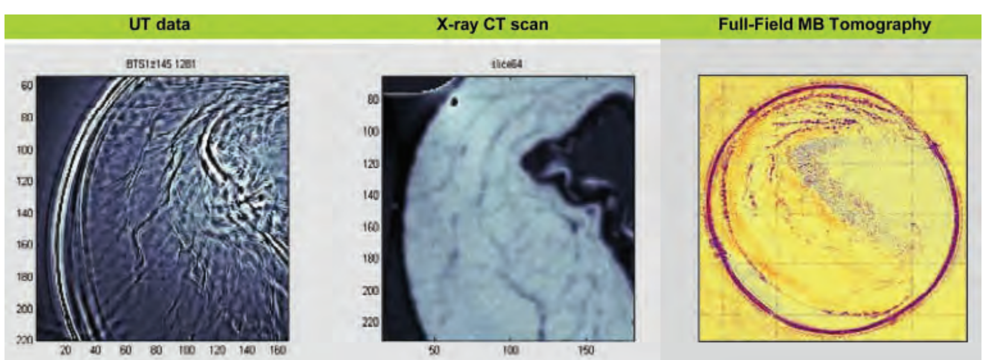

Fig. 8. Ultrasonic diffraction tomography is a model-based method based on the wave equation propagation model and backpropagation techniques to construct a meaningful reconstruction. The model-based tomographic reconstruction shown on the far right image compares quite well to the usual ultrasonic scanner shown on the far left (reflection ultrasound) and the x-ray CAT scan (middle).

phy) from a hospital scanner compared to an x-ray CAT image and finally the reconstructed image of the model-based processor using the embedded wave equation propagation model. The results demonstrate the advantage of such a sophisticated approach—especially for this application.

Summary

We expressed the basic notion that signal processing is concerned with the extraction of useful information from uncertain measurement data. From this perspective, processing is an essential ingredient that the acoustician cannot ignore and therefore must be included in a daily regimen for problem solving. To explain the place that signal processing occupies, we developed a conceptual signal processing framework to demonstrate how signal processing techniques "fit" into this plan. Using a "bottoms-up" perspective we illustrated conceptually the progression up the so-called signal processing staircase (Fig. 1) to illustrate a variety of acoustical processing problems. We have discussed some of the modern techniques in acoustical signal processing employing the philosophy that this approach incorporates more and more of the *a priori* acoustical information available into the processing scheme that typically takes the form of a mathematical model. The incorporation of these models into the processor leads to the model-based approach or equivalently, the physics-based approach to signal processing. We started with a simple representation of the staircase showing that as the models get more complex so does the processor using some simple examples for motivation. We demonstrated some acoustic applications in sonar and nondestructive evaluation and compared these results to the more classical approaches. We concluded the discussion with a tomographic imaging technique demonstrating the evolution of model-based approaches to complex acoustical problems. The answer to the question of "science or science fiction" therefore lies in the hands of the acoustician who must be able to sort through the processed results with the aid of statistical testing to assure the validity of the findings. **AT**

Acknowledgments

This work was performed under the auspices of the U. S. Department of Energy by the Lawrence Livermore National Laboratory under Contract No. DE-AC52-07Na27344.

References

1 M. J. Crocker (Ed.), *Encyclopedia of Acoustics* (John Wiley, New York, 1997).

2 W. M. Hartmann, *Signals, Sound, and Sensation* (Springer Science, New York, 1997).

3 D. R. Raichel, *The Science and Applications of Acoustics* (Springer, New York, 2006).

4 T. D. Rossing, F. Dunn, W. M. Hartmann, D. M. Campbell, and N. H. Fletcher (Eds.), *Handbook of Acoustics* (Springer, New York, 2007).

5 M. Hayes, *Statistical Digital Signal Processing and Modeling* (John Wiley, New York, 1996).

6 J. V. Candy, "Modern techniques in acoustical signal and image processing," *Proc. 9th Intl. Congress Sound & Vibration* (2002).

7 J. V. Candy, *Model-Based Signal Processing* (John Wiley/IEEE Press, New Jersey, 2006).

8 H. P. Bucker, "Use of calculated sound fields and matched-field detection to locate sound in shallow water," J. Acoust. Soc. Am. **59**, 329–337 (1976).

9 A. Tolstoy, *Matched Field Processing for Ocean Acoustics* (World Scientific, New Jersey, 1993).

10 J. Markel and A. Gray, *Linear Prediction of Speech* (Springer-Verlag, New York, 1976).

11 L. Ljung, *System Identification: Theory for the User* (Prentice-Hall, New Jersey, 1987).

12 J. V. Candy and H. Jones, "Processing of prosthetic heart valve sounds for single leg separation classification," J. Acoust. Soc. Am. **97**(6), 3663–3673 (1995).

13 K. J. Langenberg, M. Berger, T. Kreutter, K. Mayer, and V. Schmitz. "Synthetic aperture focusing technique in signal processing." NDT International **19**(3), 177–189 (1989).

14 J. V. Candy, R. D. Huber, D. H. Chambers and G. H. Thomas, "Matched-field imaging of laser generated ultrasound for non-destructive evaluation," LLNL Report, UCRL-JC-133397 (1999).

15 M. Fink, "Time reversed acoustics," Physics Today **50**, 34–40 (1997).

16 A. Baggeroer, W. Kuperman and H. Schmidt, "Matched-field processing: Source localization in correlated noise as an optimum parameter estimation problem," J. Acoust. Soc. Am. **83**, 571–587 (1988).

17 E. J. Sullivan and D. Middleton, "Estimation and detection issues in matched-field processing," IEEE Trans. Ocean. Eng. **18**, 156–167 (1993).

18 J. V. Candy and E. J. Sullivan, "Ocean acoustic signal processing: a model-based approach." J. Acoust. Soc. Am., **92**(12), 3185–3201 (1992).

19 E. J. Sullivan and J. V. Candy, "Space-time processing of a moving array: A model-based approach." J. Acoust. Soc. Am. **102**(5), 754–765 (1998).

20 A. J. Devaney, "A filtered backpropagation algorithm for diffraction tomography," Ultrasonic Imaging **4**, 336–350 (1982).

21 A. C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging* (IEEE Press, New York, 1988).

22 J. V. Candy, "Ultrasonic breast imaging program," LLNL Report, UCRL-JC-124026 (1998).

Jim and Pat Candy

James V. Candy is the Chief Scientist for Engineering and former Director of the Center for Advanced Signal & Image Sciences at the University of California, Lawrence Livermore National Laboratory. Dr. Candy received his B.S.E.E. degree from the University of Cincinnati and his M.S.E. and Ph.D. degrees in Electrical Engineering from the University of Florida, Gainesville. He is an Adjunct Full-Professor at the University of California, Santa Barbara and a registered Control System Engineer in the state of California. Dr. Candy is a Fellow of the Institute of Electrical and Electronics Engineers (IEEE) and the Acoustical Society of America (ASA) and was recently elected as a Life Member (Fellow) at the University of Cambridge (Clare Hall College). He is a member of Eta Kappa Nu and Phi Kappa Phi honorary societies and was elected as a Distinguished Alumnus by the University of Cincinnati. Dr. Candy received the IEEE Distinguished Technical Achievement Award for the "development of model-based signal processing in ocean acoustics." Dr. Candy was selected as an IEEE Distinguished Lecturer for oceanic signal processing. He was recently awarded the prestigious Helmholtz-Rayleigh Interdisciplinary Silver Medal by the ASA "for his contributions to signal processing and ocean acoustics." He has published over 200 journal articles, book chapters, and technical reports as well as written three texts in signal processing with a fourth in press. He is currently the IEEE Chair of the Technical Committee on Sonar Signal and Image Processing and served as Chair of the ASA Technical Committee on Signal Processing in Acoustics and as an Associate Editor for Signal Processing of *JASA Express Letters*.