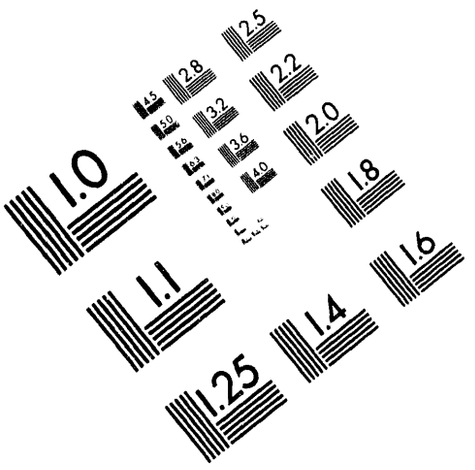

Association for Information and Image Management

1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910

301/587-8202

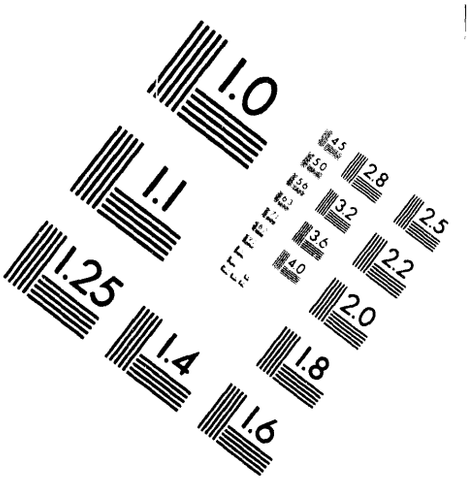

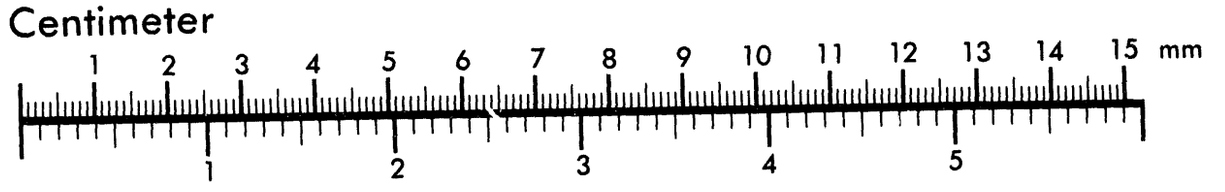

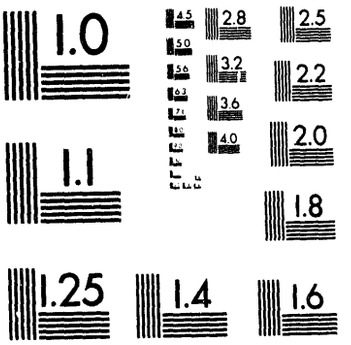

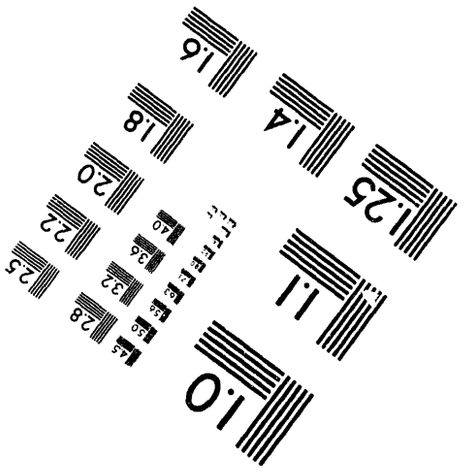

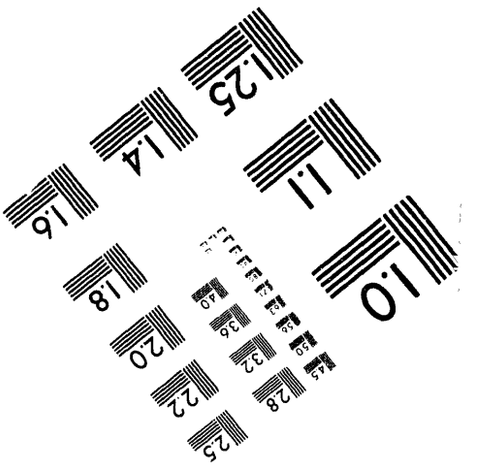

1 of 1

CONF-940449--15

Instrumentation and Controls Division

Wavelet-based acoustic recognition of aircraft*

W. B. Dress and S. W. Kercel

Invited Paper for the

SPIE International Symposium on Optical Engineering and Photonics in Aerospace Sensing

Conference on Wavelet Applications

Orlando, Florida

April 5 - 8, 1994

---

*Research sponsored by the U. S. Department of Energy, under contract DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc.

Wavelet-based acoustic recognition of aircraft

W. B. Dress and S. W. Kercel

Instrumentation and Controls Division
Oak Ridge National Laboratory
P. O. Box 2008
Oak Ridge, Tennessee 37831

ABSTRACT

We describe a wavelet-based technique for identifying aircraft from acoustic emissions during take-off and landing. Tests show that the sensor can be a single, inexpensive hearing-aid microphone placed close to the ground. The paper describes data collection, analysis by various techniques, methods of event classification, and extraction of certain physical parameters from wavelet subspace projections. The primary goal of this paper is to show that wavelet analysis can be used as a divide-and-conquer first step in signal processing, providing both simplification and noise filtering. The idea is to project the original signal onto the orthogonal wavelet subspaces, both details and approximations. Subsequent analysis, such as system identification, nonlinear systems analysis, and feature extraction, is then carried out on the various signal subspaces.

1. INTRODUCTION

Identifying aircraft types from acoustic emissions has been a long-standing problem. Earlier methods, based on the fast-Fourier transform and a distance measure between correlation matrices, had only limited success and did not lead to deployable instrumentation. System identification by autoregressive, moving average (ARMA) models was likewise unsuccessful due to the non-stationary nature of the acoustic event. Most of the published work in aircraft acoustic emission has been concerned with classifying the sound levels and types for purposes of noise abatement rather than aircraft identification.[1,2] Few, if any, of the previous attempts at classification have been published in the open literature.

As a time-frequency approach seemed most likely to succeed, we examined several time-frequency methods for extracting identifying information from acoustic signatures. The method of wavelet analysis was found to be the most promising from a practical standpoint due to the ease and effectiveness of implementation as a finite-impulse response (FIR) filter bank.

In addition to using wavelets to extract relevant features of the acoustic event, we explored several methods of feature classification. Fuzzy set classifiers and neural networks are among the most promising for this particular problem. Thus, we have overcome earlier problems by a combination of wavelet analysis and pattern recognition. The resulting algorithm is suitable for use in a small instrumentation package due to its efficiency, speed, and accuracy.

Section 2 describes the data-collection effort, including the recording equipment and physical configuration. Data analysis and the particular use of wavelets is presented in Section 3, while Section 4 is a discussion of classification techniques that gave satisfactory results. Section 5 continues the wavelet analysis in more detail, showing how other signal-processing methods may be combined with wavelets to extract physical parameters of an event. The results are summarized and conclusions are presented in Section 6.

2. DATA COLLECTION

After exploring the availability and quality of existing aircraft acoustic-data sets, we decided that first-hand measurements of acoustic signatures would be necessary for a comprehensive and accurate analysis. The Metropolitan Knoxville Airport Authority (Knoxville, Tennessee) very graciously cooperated with our effort, allowing us to set up our recording instrumentation at the end of several of their major runways.

On four occasions, we were able to obtain acoustic-signature data of aircraft operations on three different runways at McGhee-Tyson Airport, near Knoxville. Data were recorded for a total of 162 aircraft events, spanning a wide variety of aircraft types and actions. In total, approximately 30 aircraft types were recorded; they fall in three main classes: large jets, twin-engine turbo props, and single or double piston-engine aircraft. We attempted to record statistically significant samples of both take-off and landing events for each aircraft type, but due to the restricted sampling methods imposed by time constraints, this goal was not always achieved. Additional events occasionally recorded included preflight throttle tests, wherein the sound source was stationary but intense preceding a takeoff, and echoes from the large jets reflected off nearby buildings, ariving some tens of seconds after the aircraft passed by the sensors.

The primary sensor was a laboratory quality, B&K 4133 condenser microphone having a frequency range from 17 Hz to 39 kHz, flat to within ±2 dB. The acoustic events were recorded by a Honeywell Model 101 FM multichannel magnetic tape recorder. The recordings were made (on different occasions) at 7.5 ips and 15 ips; the recorder bandwidth at the former speed is 25 kHz, +1, -2 dB. Each event was voice-annotated by the operator; this annotation was recorded on another channel of the same one-inch wide tape. For the airport recordings, the microphone was mounted on a foam-padded plywood square, 18 inches on a side, and approximately two inches above the foam. The microphone was pointed upward at about 30 degrees. The axis of the microphone's directional pattern was normal to the runway. The microphone itself was located about 300 feet from the take-off end of the runway, some 12 feet from the paved edge. This restricted the recordings to the last segments of take-offs and the first segments of landings, as the aircraft was typically airborne as it passed the microphone.

Events were logged in a notebook as well as on the tape itself. Airport personnel identified each event as to aircraft type, take-off, landing, or flyby prior to the actual event. This gave the recording crew ample notification as to when to start recording the event and provided an additional safety factor as well. There were approximately 10 events per hour on the selected runways. Tail numbers of most aircraft were noted. A small selection of events is shown in Table 1.

Table 1. Selected Aircraft Acoustic Events. The event time-of-day is given along with the aircraft type, tail ID number, and the perceived action . The column labeled "Index" is an identifier for the location on the magnetic tape. The "Comments" column contains other information about the event or the aircraft.

Index	Time	Aircraft Type	Tail ID	Comments	Action
1.9	08:31	Piper PA32			landing
2.1	08:54	Piper PA32		1 eng turbo	landing
2.2	09:03	Jet Stream	N163PC	2 eng	landing
2.3	09:16	Jet Stream	N861AE	Am Eagle	rev up
2.3a	90:16	" "	"	" "	take off
2.4	09:10	MD88			landing
2.5	09:34	Lear Jet	N952M	mixed w/cessna?	landing
2.7	09:47	Single Engine	N336EJ	1 eng	landing
2.9	09:53	Mosquito	N2FF	1 eng piston	take off
2.12	10:21	Saab	N344CA	2 eng turbo	landing
2.14	10:39	MD80	N983DL	cargo jet	takeoff
3.1	10:56	DC 9		passenger jet	landing
3.2	11:04	Saab	N275UE	2 eng turbo	landing
3.3	11:06	Cessna	N2RG	1 eng piston	rev up
3.6	11:21	737-300	N917UA	passenger jet	take off
3.8	11:36	King Air	N60SC	2 eng turbo	take off
3.9	12:06	Fokker	N494US	like DC9,MD80	landing
3.10	12:15	Cessna 152		1 eng piston	take off
3.11	12:30	King Air	Army	2 eng turbo	landing

2.1 Acoustic Interference

A common problem in data collection of this type is interference from ground reflections. At certain frequencies, depending on the geometry of the measuring system, acoustic energy reflecting off the pavement or ground will interfere (both constructively and destructively) with the main signal traveling directly from the aircraft. This effect is most evident in a sonogram display of the event, wherein the time-frequency plot takes on a layered or "onion skin" appearance. The interference effect is also described in the literature.[3] As such interference is not indicative of the behavior of the source, we minimized the effect by appropriate placement and orientation of the microphone. In addition, by suitably averaging the various wavelet transform levels, and not extracting phase information, the acoustic interference effect does not contribute any confusion to the feature vectors used in the classification schemes.

2.2 Other Sources

One of the problem areas appearing in previous work was the inability of certain methods to discriminate between aircraft events and other likely occurrences such as motorized ground traffic and thunderstorms. In addition, the echoes noted above

might well appear to be separate aircraft events to a system having low discrimination capability. To investigate the possibility of discriminating against such "uninteresting" events, we collected data from both ground vehicles and acquired thunderstorms signatures.

For motorized vehicular events, data was taken with the same recording instrumentation, but located at an overpass on Interstate 40, near Knoxville. 55 highway events (trucks and automobiles) were recorded. Preliminary Fourier analysis of a jet take-off and a highway event of a large tractor-trailer passing by clearly demonstrates the type of problem encountered. Figure 1 shows a comparison between a binned, Fourier power spectrum of a jet take-off and a truck drive-by.

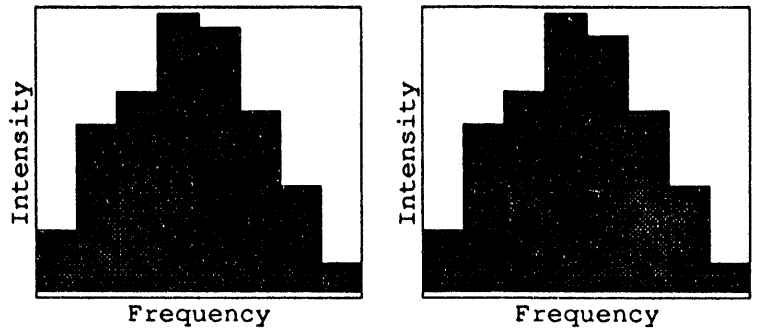

Figure 1. Aircraft events can be similar to ground-traffic events. The sound intensity shown as a function of frequency accumulated into third-octave bins spans the range from about 100 to 633 Hz. The aircraft signature, from an MD80 "touch-and-go" event, is on the left; the truck highway event is on the right. Both histograms have been normalized to unity at the most intense bin.

Each of the events shown above are recorded for their entire duration of the take-off or drive-by, defined somewhat arbitrarily as that period when the average sound intensity was greater than twice the background sound. This period for the aircraft take-off was about 22 seconds, while the truck drive-by lasted about 3 seconds. The third-octave histograms show very little difference between the two events. One simple resolution to this dilemma is presented in the next section.

To complete the set of over 200 recordings, 11 "natural" events of recorded thunderstorms were extracted from a commercial cassette recording. A commonly found feature in acoustic events stemming from natural causes is a 1/f behavior in their power spectral density (PSD). Below, we show that some aircraft events exhibit only a slight departure from the 1/f behavior noticed in the thunderstorm data.

3. DATA ANALYSIS

Acoustic signatures (a temporal sequence of pressure waves characteristic of certain events) contain information in their temporal structure as well as the structure of their power spectra. Wavelet analysis is one method of extracting such temporal behavior in a mathematically consistent and practical manner. Many other techniques have been developed over the past half century and could be likewise explored in any complete exposition on signature analysis. For example, there have been several powerful methods related to the Fourier approach (e.g., the Gabor transform). Other spectral-based methods rely on linear models and have produced the ARMA method, the related maximum-entropy spectral approximation method, and several other variations on linear-system modeling. The Wigner transform (based on the Wigner-Ville distribution) is an exceptionally popular method for analyzing chirp signals and is the basis of the ambiguity transform used in analysis of radar signatures. In addition, Markov methods provide a convenient means of analyzing structured sequences of acoustic events found in connected speech.

In keeping with a design goal of a small, inexpensive, and deployable instrument, all the signatures obtained in the data-collection stage were digitized to 8 bits at a rate of 11,000 samples per second and stored on digital magnetic media. This digitized data set served as the input to the analysis stage.

The objective of the airport study was to determine the feasibility of extracting feature vectors suitable for distinguishing between classes of aircraft. A key problem in the identification of complex events is to find a set of features that are similar for different signatures of the same class, but distinctly different for signatures of different classes. Additionally, one would like the feature-extraction process to be algorithmic for automatic operation. Since other researchers had devoted considerable effort to this problem using variations on Fourier analysis, we knew that this technique was not suitable. However, Fourier analysis remains an excellent point of departure for any signal analysis problem.

3.1 Preliminary Fourier analysis

Each of the different event classes was examined with a discrete Fourier transform (DFT) to determine if any outstanding features were present. Jet aircraft (passenger and cargo planes) exhibited few outstanding features in the frequency domain. The propeller aircraft and turbo-prop aircraft showed some spectral peaks that could easily be identified with the periodic disturbance of air displaced by the tips of the propellers. A typical spectral peak (shown below in Section 5) was at 86 Hz in the case of a four-blade propeller rotating at the rate specified for take-off.

The Fourier transform of an acoustic signal taken from a large jet is shown in Figure 2. The most noticeable characteristic is the approximate 1/f behavior of the signal power as a function of frequency. While we can certainly expect to distinguish this event from a turbo-prop take-off based on the difference in spectral features (particularly in the 100 Hz range), there is little here to suggest that Fourier analysis might produce unambiguous features for distinguishing this event from other similar, but distinct events—either landings from take-offs, or between diverse large jet aircraft.

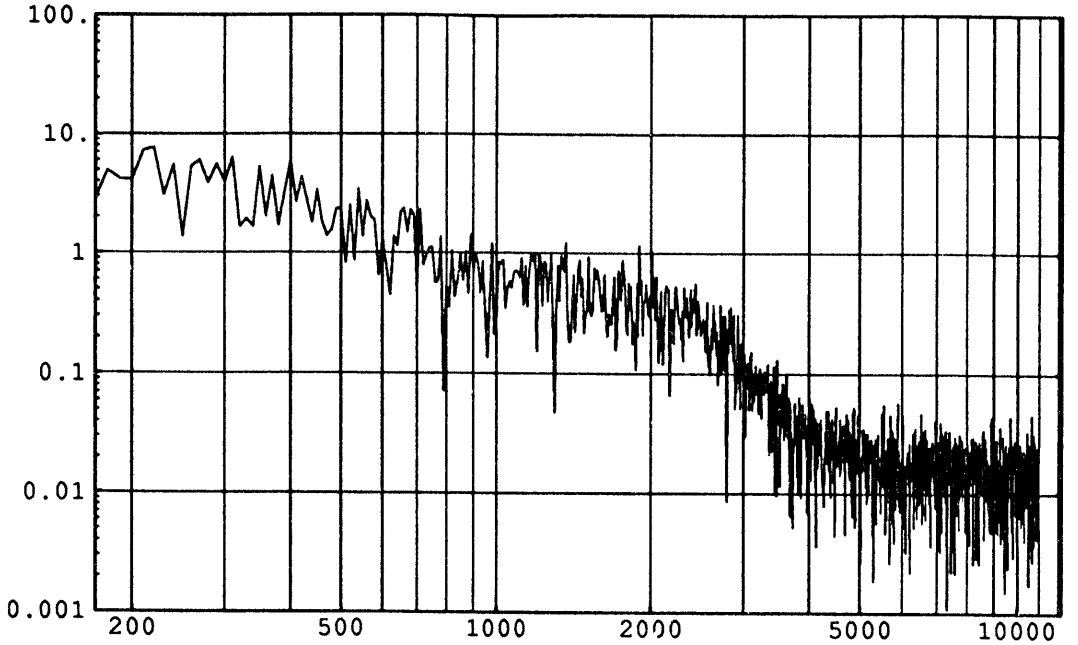

Figure 2. Fourier power spectrum of an MD80 take-off event. The power spectrum of an entire take-off event shows a departure from the 1/f characteristic in the range of 1200 to 3000 Hz; otherwise, there are few outstanding features that might be conducive to classification. The ordinate is in arbitrary units of acoustic power; the abscissa is in Hz. The recorded signal was digitized at a rate of 22 kHz for this analysis.

As most natural acoustic events (e.g., thunderstorms) tend to show a 1/f character in their PSD, the above aircraft event is barely distinguishable over its duration from such natural phenomena. There is slight departure from 1/f behavior at both extremes (below about 300 Hz and above about 7000 Hz) where the detected power seems to drop off or rise slightly. The main departure from 1/f behavior lies in the range of 1000 to 3000 Hz where we might expect to find the a contribution to distinguishing features. This region corresponds approximately to levels 3, 4, and 5 in the wavelet transform described below.

In Section 2, we saw that a third-octave description of an MD80-class jet differed little from that of typical highway noise. Here, we assert that there will be little qualitative difference between an aircraft event and certain natural phenomena likely to be present. Thus, the problem is to extract a set of features that are sensitive enough to distinguish between aircraft of various types, but not confound any aircraft classifications with highway events or thunderstorms.

3.2 Wavelet analysis

The digitized signals were projected onto discrete wavelet subspaces by the FIR wavelet method as described by Mallat.[4,5] We found that certain wavelet bands contained little information, depending on aircraft type, while others were rich in relevant distinguishing information. This simple observation provides a preliminary classification scheme, and is refined below.

The wavelet transform coefficients for the FIR filter bank were taken from Daubechies,[6] and a 12-level wavelet transform was computed for each of the events. A typical transformed event, in this case another large-jet take-off, is shown in Figure 3.

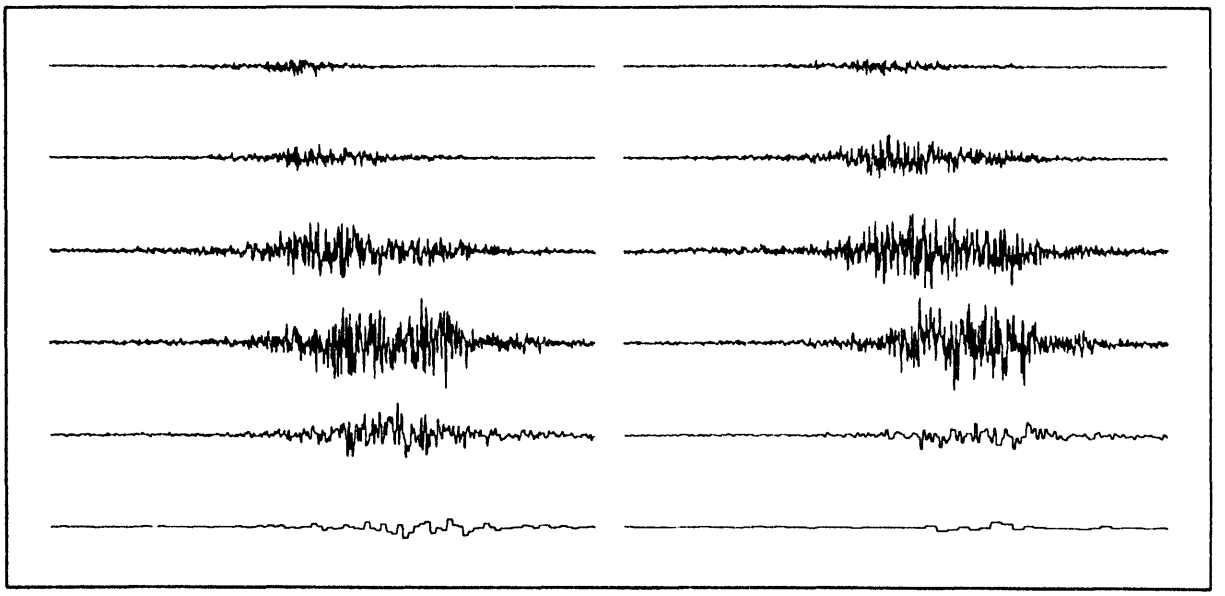

Figure 3. A 12-level wavelet transform of an MD80 take-off using the D4 discrete wavelet transform. Level 1 (smallest details) is the trace in the upper left; level 2 is shown in the upper right. The next row contains level 3 and level 4, and so on. The units are arbitrary, but chosen so that the entire event, some 20 seconds along the horizontal axis, is displayed.

The approximations to the signal at each level were discarded as we are only concerned with the details over several of the levels. Since there are of the order of a hundred thousand samples for a typical event and the transformed signals were averaged in time to obtain reasonably sized signatures, the particular wavelet transform used was relatively unimportant. For shorter signals with no averaging, the particular choice of analyzing wavelet can make substantial differences in the descriptor sets. One effect of analyzing the signal on wavelet subspaces is that the temporal character is not lost; the time axis is merely compressed by a factor of two for each transform level. If we now bin the square of the intensities of the transformed signal according to a given temporal epoch during an event, we obtain a density plot as shown in Figure 4.

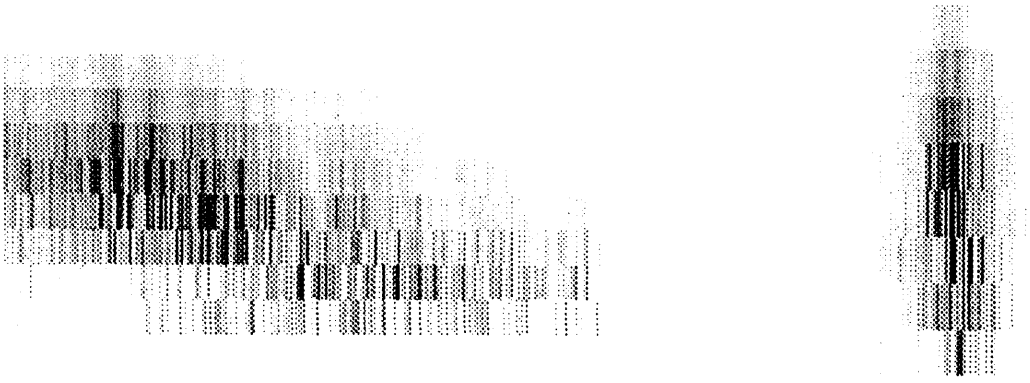

Figure 4. Wavelet transforms for the aircraft and highway events shown in Fig. 1. A temporal average (into 40 ms bins) has been carried out on each of the 8 most significant wavelet levels. Time increases to the right, and scale from top-to-bottom, with the finer scale (higher frequencies) at the bottom.

In this view (still greatly compressed compared to the original signals), there is an observably great difference between the aircraft and highway events introduced in Section 2. Even if we attempt to correlate the truck event (right) with a portion of the aircraft event, we obtain a poor match in all instances.

Fig. 4 graphically illustrates a possible method of classification. Here, we have a simple function of two dimensions (temporal bins and wavelet levels). Defining the onset of an event in relation to the most intense region or maximum of the function, we can now compare events. This high level of discrimination was found to be unnecessary for successful classification of the aircraft, but remains a possibility for future applications to higher-discrimination needs, such as intra-class discrimination.

4. EVENT CLASSIFICATION

Subsequent analysis extracted information for feature vectors from these various subspaces. Some of the methods employed were higher-order zero-crossings[7], ARMA modeling, Fourier transforms of the subspace signals, and additional wavelet analysis. In this section, we employ Fourier transforms of the subspace signals to obtain compact yet descriptive feature vectors.

4.1 Feature extraction

The wavelet transform provided us with 12 subspace signals for each aircraft event. As mentioned above, in most cases only the details of levels five through eight are sufficient to characterize the event. However for some of the "big-jet" events, a significant portion of the energy resides in levels one through four; this provides a way to distinguish these events from the turbo-prop events. A feature vector can be defined for each of the events as a normalized vector obtained by partitioning a Fourier power spectrum on each of the four wavelet subspaces into third-octave bins. This results in a vector consisting of eight components on each wavelet subspace, or 32 components in total. These vectors were normalized by dividing each component by the largest component in each vector. Thus, we have compressed the acoustic information necessary for classification by nearly four orders of magnitude from that of the digitized time-domain signal.

4.2 Fuzzy set membership

A fuzzy set[8] approach provides a simple method for creating event clusters. It has a certain robustness that allows reasonable identifications to be made on small-event sets. A fuzzy set describing a particular event is simply the normalized feature vector described above.

Once a fuzzy set is computed for each event, the standard methods of fuzzy set theory as discussed in Ref. 8, can be used to group events into classes and compute class membership possibilities for a test event belonging to an unknown class. The operations of fuzzy union, intersection, and membership degree are all the machinery required for a fuzzy-event classification system. The Figure 5 shows the results of comparing a selection of events with the fuzzy union of all the Saab events. The Saab is a twin-engine turbo-prop aircraft; in all we recorded 6 Saab events (both take-offs and landing) for 3 different models. Accordingly, the "Saab" class is sparse from statistical considerations (we would like to have about 30 events of each type). The height of each bar represents the possibility that the event is a "typical" Saab event. The bars at locations 1, 2, 26, 30, 32, and 44 show the response of the Saab cluster to the Saab events (of course, they each show 100% "Saab possibility"). The other events showing high Saab possibility are of interest; these are events 13, 20, 29, and 41. Each of these events was a aircraft of similar type to the Saab. The events with low Saab possibility, such as those labeled 10 and 40, stem from jet liners, and not small turbo-props.

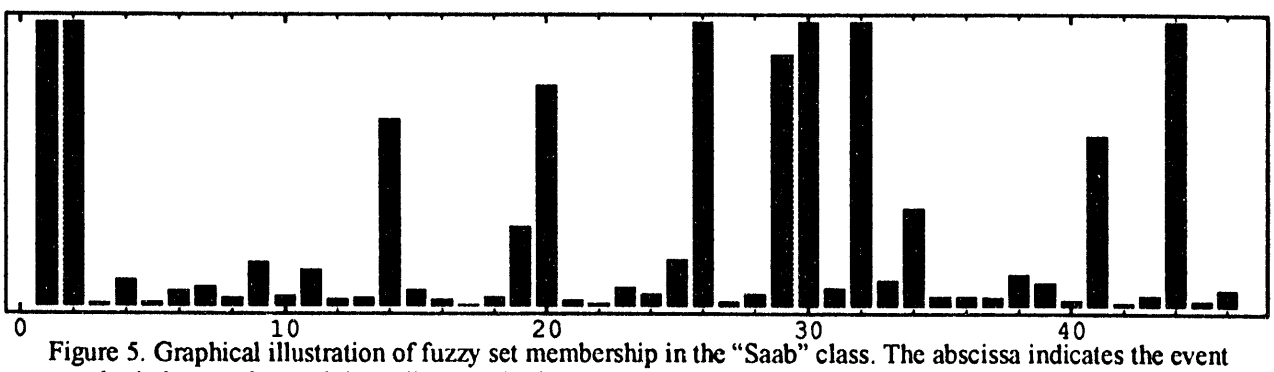

Figure 5. Graphical illustration of fuzzy set membership in the "Saab" class. The abscissa indicates the event type by index number and the ordinate is the fuzzy possibility, ranging from 0 to 100%.

The fuzzy-set approach presented here resulted in better than 90% correct identification of event type from the universe of the 200 recorded events. This confidence measure was obtained by a cross-correlation of all the pair-wise fuzzy-set membership possibilities and then selecting the higher possibilities (i. e., those above 95%) as defining class membership.

4.3 Neural network classification

For certain data sets, the method of functional approximation[9] known as artificial neural networks (ANN) is more efficient than either the usual multivariate approaches or fuzzy-set classification. The normalized feature vectors extracted above are ideal candidates for inputs to any of several ANN models—most notably the multilayer perceptron network (MPN) model and the

radial basis function (RBF) model. The primary difference between these two approaches lies in their particular basis functions and methods of parameter estimation. The MPN uses basis functions consisting of non-orthogonal "sigmoidal" functions such as the inverse hyperbolic tangent, and commonly employs a form of gradient descent to locate a local optimum in parameter space. (The model parameters are referred to as "weights" in the ANN vocabulary—this usage is to be distinguished from that in regression analysis.) On the other hand, the RBF uses a radially symmetric basis function set, again not necessarily orthogonal, usually based on, but not restricted to, the Euclidean distance between multidimensional data points and cluster (or node) centers. The method of parameter estimation is either gradient descent in weight space for the MPN or, in case of simple network architectures, a direct least-squares fit using the matrix pseudo-inverse is quite effective.

As we wished to classify each event into a predetermined set, such as "large jet landing," these *a priori* classes were identified with output nodes (or units) of the network. We selected a total of six classifications for testing the ANN: take-offs and landings for each of the three categories of big jets, twin-engine turbo-props, and single- or twin-engine piston aircraft. There were 32 input nodes, corresponding to each of the 32 components in the canonical feature vector. These nodes actually do nothing but serve as notational and conceptual convenience in the ANN vocabulary. The intermediate layer in the network, the "hidden" layer in ANN parlance, consisted typically of 20 to 60 nodes where each basis function was applied to the sum of the inputs over all 32 components of the feature vector. For the MPN, each of the input vector components is scaled by a parameter or weight that can be different for each input component and each hidden node. For the RBF case, the distance between the input vector and each node center was computed before applying the basis function; weight multiplication was done after the function application in contrast to the MPN method. In passing from the hidden layer to the classification layer (network output), each hidden node's output was scaled by another parameter before summation at the classification nodes. There were approximately 2400 degrees of freedom in the 60-node MPN network. The RBF network was considerably smaller due to the fewer hidden nodes required; 20 nodes were sufficient, resulting in 886 degrees of freedom.

Both network models gave approximately the same degree of correct classifications. At the 86% level, this was not significantly different than the fuzzy-set method discussed above. The main difference between the two classification schemes was in the number of classes able to be supported. Due to the sparseness of the data set, the ANN approach could only function with six or fewer clusters. The fuzzy-set method operates quite well with two or three times the number of clusters. Thus, we could easily detect differences between types of turbo-props using the latter method.

4.4 Non-linear analysis

We now wish to exploit the simplifying aspect of the wavelet projection operators. For orthogonal wavelet bases, the original signal is projected onto non-overlapping subspaces. Each of these spaces is somehow "simpler" than the original signal space as certain details have been removed in the case of the approximations and certain low-frequency behavior has been discarded in the case of the details. For some types of signals and certain wavelet functions, this reduction in complexity can translate to "cleaner" attractors, if not simpler ones. This is analogous to the digital filtering methods used by Lawkins, et al,[10] but goes beyond their approach in assuming possible decomposition of high-dimensional attractors to simpler geometrical structures by means of wavelet projections.

This is obviously correct in the case of the Saab 340 take-off analyzed below. The attractor on approximation levels that preserve the fundamental frequency of rotation for the propeller tips are easily seen to be limit cycles (compact ellipses in an embedding dimension of two), whereas the attractor for the original time series is not quite so easily interpreted. Additional analysis of the dynamics on the details subspaces will be carried out at a later time.

5. PHYSICAL PARAMETER ESTIMATION

The estimation of physical parameters from a one-dimensional time-domain signal is a model-dependent problem, with at least as many solutions as there are models available. In the case of a turbo-prop aircraft event, the models are restricted by known aircraft types (number of engines, number of blades on each propeller, and nominal take-off and landing engine-power levels). There is some information to be gained from a Fourier analysis of the digitized time-domain signal. When Fourier methods or other approaches are applied to the projections of the signal on wavelet subspaces (now the approximation signals are considered as well as the details), both the accuracy and the amount of extracted information increases. For the analysis presented in this section, we consider the central, intense portion of a Saab 340 during take off. The portion where a physical model is extractable covers about six seconds when the aircraft was closest to the microphone.

Figure 6 shows a 1/4-sec segment of the digitized acoustic signature during the first part of the six-second interval when the aircraft was approaching the microphone, approximately one second before opposition. Note the presence of a number of other signals at higher frequencies that are modulated by the fundamental. The intensity, as measured by peak-to-peak distances is also increasing (the aircraft is approaching the microphone).

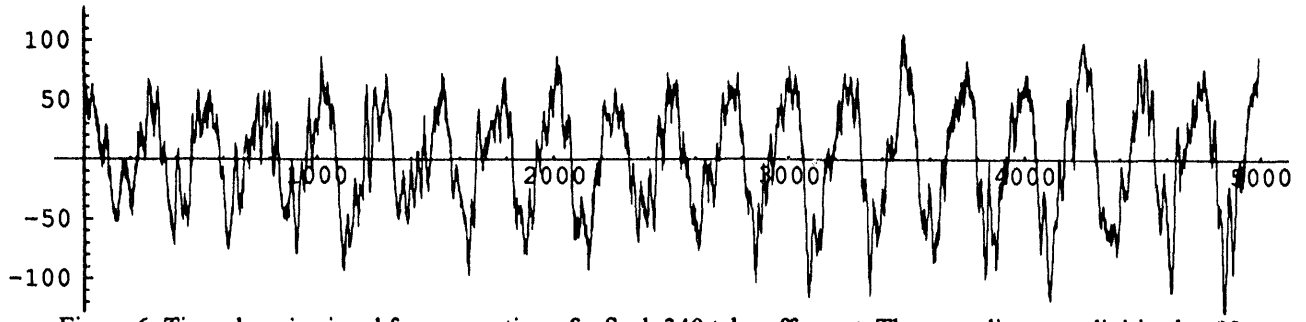

Figure 6. Time-domain signal from a portion of a Saab 340 take-off event. The recording was digitized at 22 kHz, so the length of this portion is about 250 ms. The strong pressure pulses due to the turbo-prop tip passage are clearly visible.

We now proceed to analyze this event using standard DFT methods and conclude with a wavelet analysis followed by a nonlinear fit to a Doppler model.

5.1 Fourier parameter estimation

Figure 7 shows the DFT of the above segment. The strong fundamental located at bin 20 corresponds to about 88 Hz. A second harmonic is just visible at bin 40.

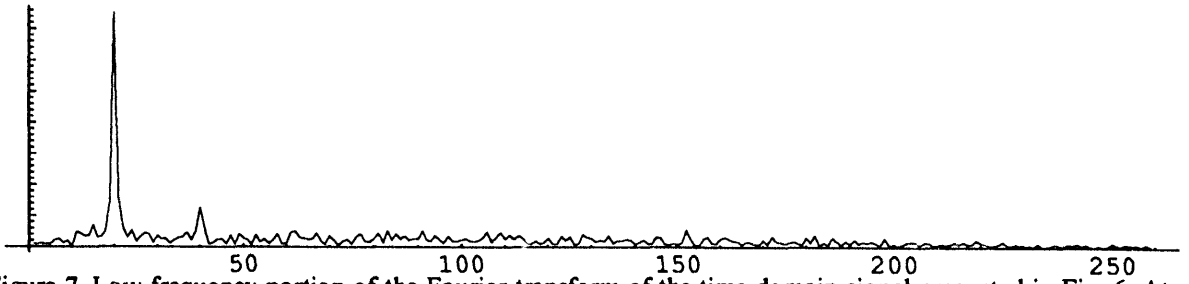

Figure 7. Low-frequency portion of the Fourier transform of the time-domain signal presented in Fig. 6. At the scale shown, there is almost no information beyond bin 250 (1100 Hz). The abscissa has units of frequency bins, where each bin spans 4.4 Hz. The ordinate is in arbitrary units of power.

The resolution of the measurement is ±1 bin or ±4.4 Hz and the weaker harmonics may, in general, be unidentifiable at this resolution due to the ratio of the power in the peak to the power in the noise over such a wide frequency range (8.8 Hz).

The standard way of increasing the resolution of a DFT is to include more samples in the transform. This has the disadvantage of increasing the time required for the calculation; there are other disadvantages as well. If accurate measurements are required (for, e.g., Doppler effect determination), including more samples hides the very information sought since the frequency is changing significantly over a time scale of 250 ms for a typical small-craft take-off. Including more samples would merely give an average over that extended time. Figure 8 shows the results of an alternate method: padding a 100 ms segment of the original signal with zeros to increase the (apparent) resolution.

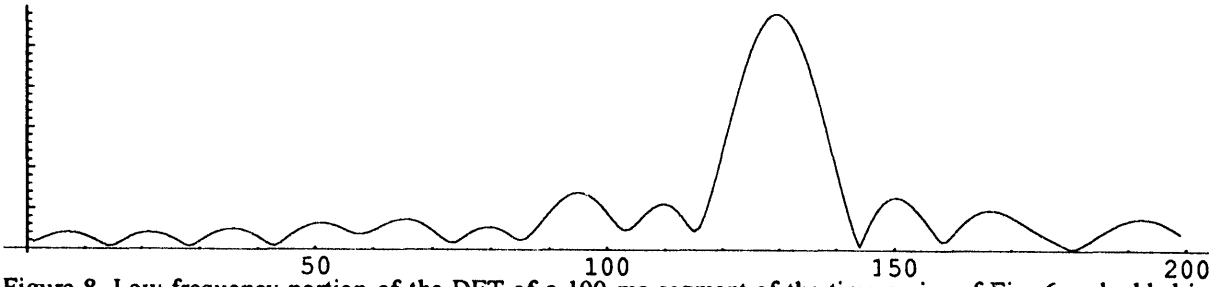

Figure 8. Low-frequency portion of the DFT of a 100 ms segment of the time series of Fig. 6 embedded in 32000 zero samples. Here, an abscissa bin corresponds to 0.69 Hz, and the peak is located at bin 130.

Again, we can locate the peak to within ±1 bin (without employing sophisticated peak-fitting methods), and obtain a central frequency of 89.4±0.7 Hz—much better resolution than before. The peak shown above owes its pleasant symmetry to the set of

non-zero samples set within a much larger span of zeros, reflecting more the appearance of a Fourier transform of a square pulse than the actual data samples from the aircraft signal.

Once the fundamental has been identified, the other peaks are usually not too difficult to locate as long as the signal-to-noise ratio is favorable. The harmonics can be found by a modified ARMA method wherein the fundamental is part of the model, or by a DFT by searching for peaks at proper locations above the fundamental.

5.1.1 Prop-Blade Harmonics

Identification of the fundamental frequency due to the propeller is a straight-forward measurement using either the DFT or ARMA methods. The effect of "noise," defined as interfering or masking signals present in the signal being analyzed, can reduce the accuracy of either method to the point of failure. These effects can be minimized by appropriate wavelet filtering prior to spectral analysis. The DFT or ARMA method is then applied to a scaled version of the original signal, resulting in a much more robust result.

Another approach is to project the signal onto a set of wavelet subspaces and identify the harmonics as the principal component on each of these subspaces. This seems to work for both the second and third harmonics, depending on the aircraft and event type.

Once the original signal has been projected onto the discrete approximation and detail subspaces, further analysis on these subspaces can enrich the understanding of the physical system producing the one-dimensional time series.

5.2 Signal processing on wavelet subspaces

One of the more valuable benefits of carrying out a wavelet analysis is the multiresolution character of the analyzed signal on a set of nested, orthogonal subspaces. The original signal is split into independent subsignals by projecting it onto these wavelet subspaces. As an example of how the combined wavelet-spectral-analysis method might be used to obtain some of the underlying physical parameters of the event from those subspaces, consider the wavelet transform of a 100 ms segment (2200 samples) of the above data set. For the subsequent analysis, we will want to examine smoothed versions of the scaled original signal, so we choose a wavelet transform that ensures smooth approximations, such as D8 with 16 coefficients (see Ref. 6, p. 195). The result of applying this particular transform is shown in Figure 9.

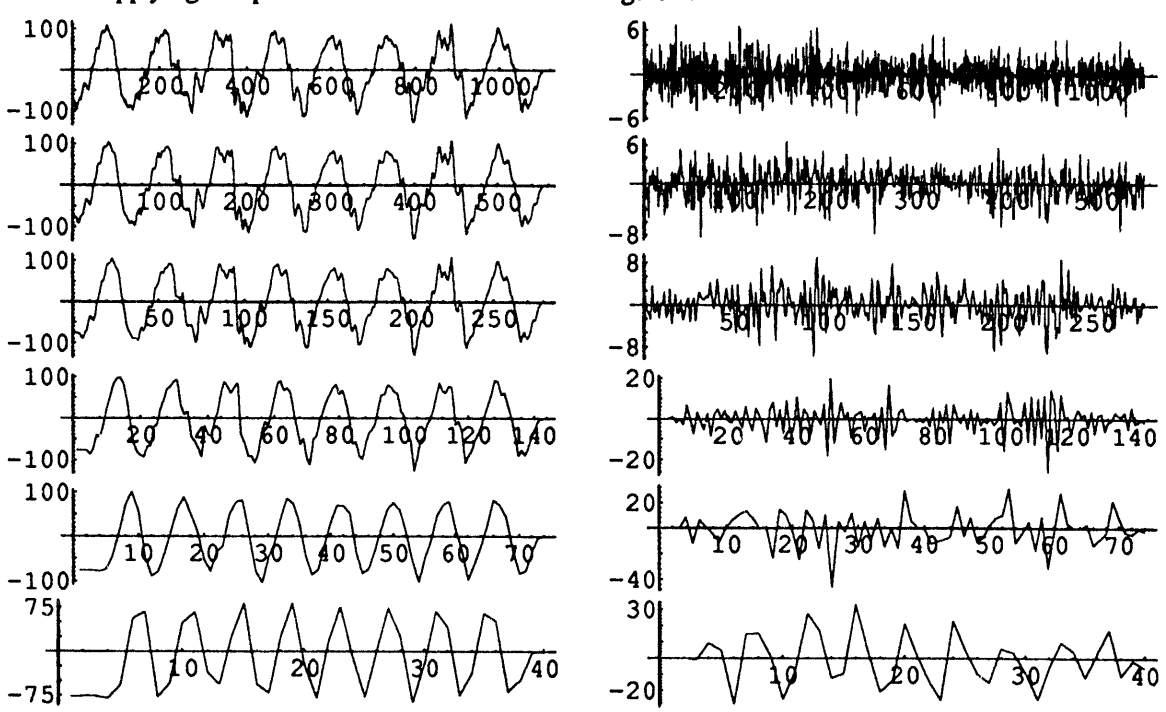

Figure 9. Wavelet transform of a portion of the time-domain signal shown in Fig. 6. The approximations are on the left and the details on the right. Level 1 results are at the top, and level 6 at the bottom. Note that the fundamental oscillatory nature of the signal is preserved at least through level 6.

If we use the approximation data of level 5 (fifth plot from the top on the left in Fig. 9) as input to the an ARMA model for spectral estimation, we find a very sharp peak at 89.32 Hz (with an error of about 0.81 Hz given by the fwhm, or ±0.4 Hz). The Fourier transform, above, gave the fundamental as 88±4.4 Hz. The spectral density obtained by plotting the frequency response of the ARMA model is shown in Figure 10, where the sharpness of the spectral peak is stems from the filtering action of the wavelet transform in producing the level 5 subspace signal.

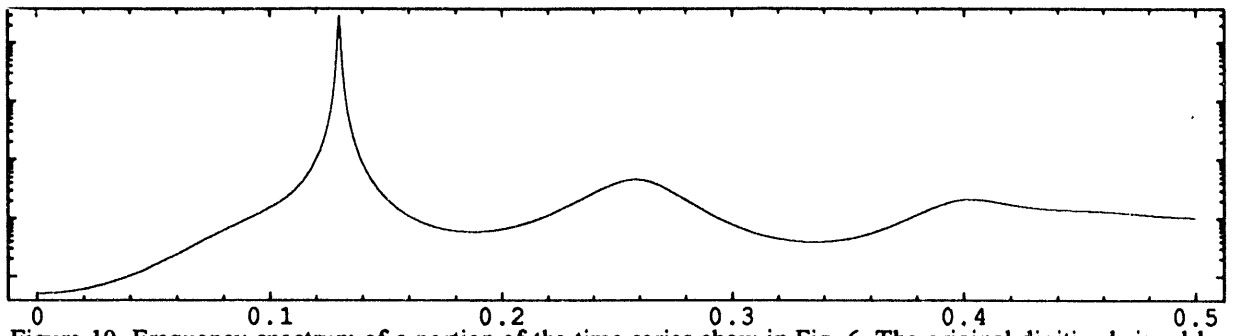

Figure 10. Frequency spectrum of a portion of the time series show in Fig. 6. The original digitized signal has first been projected onto the fifth level of the approximation subspace, a portion of which is shown in Fig. 9. The plotted function was obtained by a least-squares fit to a 12th order ARMA model. The abscissa represents normalized frequency, and the ordinate arbitrary power units.

The fundamental and two harmonics are visible in this spectral plot. The first harmonic is located at 178.73 Hz, which is twice the fundamental well within the error. Since the analysis was done on a level-five subspace, the normalized frequency corresponds to 2^{-5} x 22000 Hz, or about 686.9 Hz.

5.3 Analyzing a subspace projection

Figure 11 shows the central portion of the level 7 approximation of the wavelet transform similar to that shown in Fig. 9, but over a much longer time span. As a point of reference, the details in levels seven and eight have by far the largest amount of "energy," sharing about 50% of the total power in the signal.

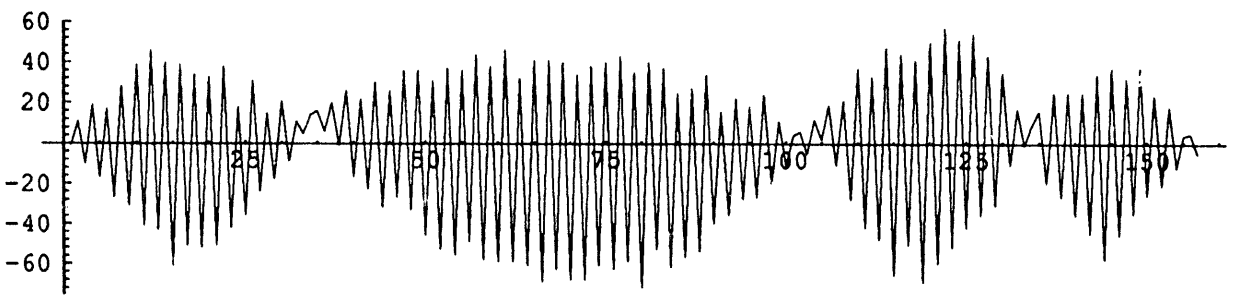

Figure 11. Central portion of the seventh approximation level of a D8 wavelet transform of the Saab 340 take-off, sampled at 11000/sec. The abscissa represents time, and spans about 1.8 seconds. The ordinate is the intensity of the subspace signal in arbitrary units.

There are two obvious features in this subspace projection: (1) the waveform is that of a high-frequency oscillation modulated by a very low-frequency wave, and (2) the high-frequency part is "at the Nyquist limit" for this discrete signal space. Since this is a level-seven projection, the Nyquist limit (the concept is still valid for a scale-space signal) corresponds to an oscillation of 2 parts in 2^7 times the Nyquist limit of the original sample (5500 Hz). Thus, the high-frequency oscillation in Fig. 11 is a direct consequence of the fundamental 86 Hz secular propeller frequency seen in the Fourier spectrum in Section 5.1 above. This is not at all surprising due to the octave scaling nature of the wavelet transform. What is a pleasant surprise, however, is that the appearance of the fundamental when projected onto this wavelet subspace is a simple (and remarkably clean) modulation of the underlying 86-Hz component.

The separation into high- and low-frequency parts becomes quite evident if we take the wavelet transform of the entire level 7 signal. This is shown in Figure 12 where the approximation is on the left and contains about 6% of the total power in level 7. The bulk of the signal is in the details (on the right), which is now clearly seen as a low-frequency modulation. The high-

frequency information is contained in the specification of the subspace itself, namely that it is a details subspace corresponding to the eighth octave of the original signal. It has been pushed "beyond" the Nyquist limit and is now present in the geometry of the nested projections rather in an oscillatory nature of a particular time series.

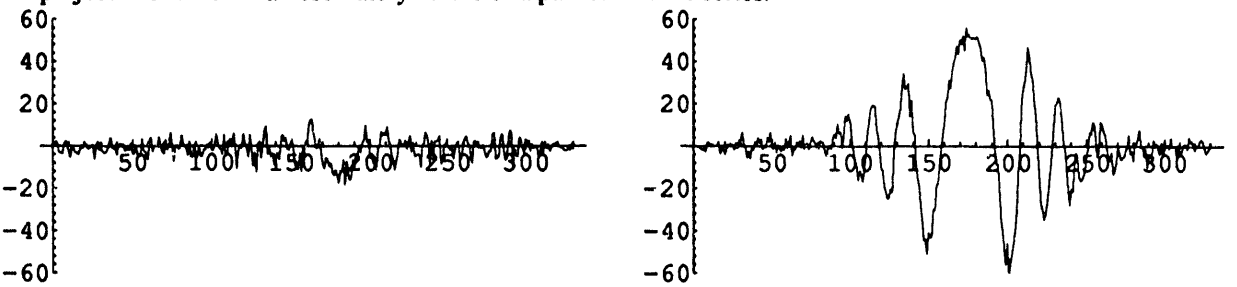

Figure 12. D4 wavelet transform of the level 7 approximation, the central portion of which is shown in Fig. 11. The approximation is on the left and the details on the right. The abscissa spans about 8 seconds, showing the entire take-off event. The ordinate units are the same as in Fig. 11 and show comparable intensities.

Further analysis is possible on the details (signal shown on the right in Fig. 11) section of this transform. It may well be possible to infer underlying mechanisms as to the physical cause of the modulation from the above signal or use its characteristics for identification purposes. For example, the period of the low-frequency beat is about 0.3 sec, and is consistent with a slight mismatch in rpm of the two engines. If such a mismatch occurs at the maximum throttle during take-off in a consistent manner for similar aircraft, the features of the signal may serve to identify the particular aircraft (at least until the engines are re-tuned). These ideas await further study before definitive statements can be made.

5.4 Frequency tracking

Matching zero-crossings or peak positions[11] across several wavelet levels can give accurate timings of certain discontinuous events—in this case, the forcing of air currents in the direction of the sensor by the passage of propeller-blade tips. The method described in Ref. 11 is particularly suited for this type of signal as well as to more complicated phenomena such as speech processing. Although the peak-matching method gave excellent results for our simpler turbo-prop data, here, we employ a variant of the ARMA method based on recursive least squares to extract the fundamental propeller-tip frequency as a function of passage time.

The signal is analyzable by this method during the period when the aircraft is less than about 100 m away from the microphone, with closest approach of about 50 m. A small amount of amplification could extend this range by perhaps twice. As we are building the ARMA model on a low-scale approximation subspace, noise introduced by amplification should not be a problem.

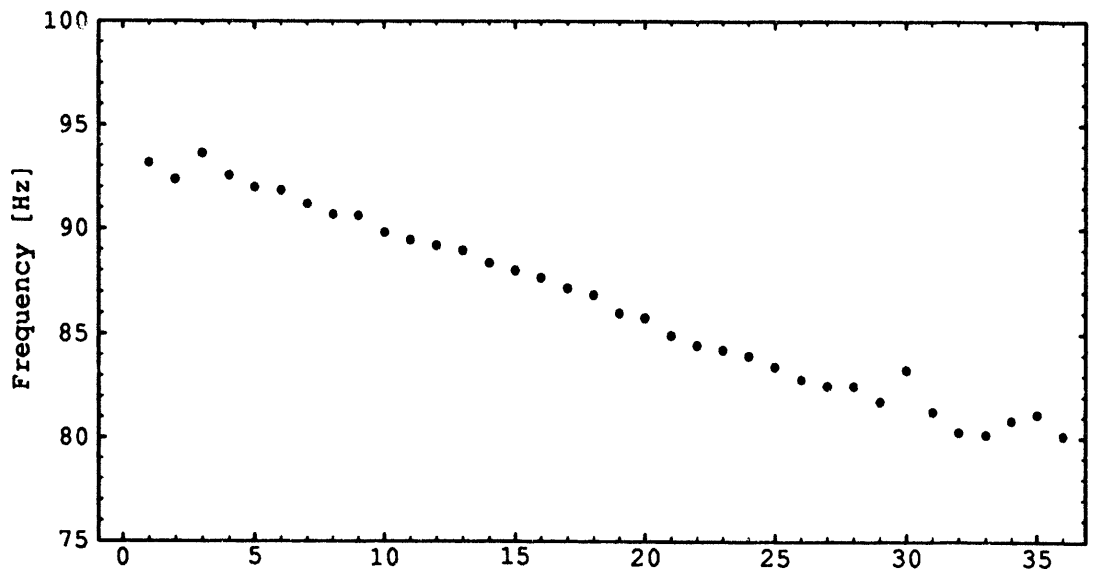

Figure 13. Frequencies obtained by a recursive, least-squares ARMA model tracking the fundamental frequency of the dynamic aircraft system. The abscissa is in units of 51 ms (see text).

Figure 13 shows that measured frequency drops off from about 93 Hz to about 80 Hz during the event. The units along the time axis are 51 ms due to the averaging method used (each frequency point was computed from a 50% overlapping, 100 ms-long moving window), so the chart spans about 1.8 sec. The dispersion of the points at each end of the region is due to the low level of the acoustic signal when the aircraft is beyond about 100 m from the sensor.

5.4.1 Doppler model: extracting the physics of the event

The frequency observed at the microphone due to a moving signal source is given by Equation 1. Here, f(t) represents the frequency at time t in the stationary (microphone) frame and f_0 is the constant frequency as measured in the moving frame (actual signal frequency). The velocity of the signal source is v_0, measured in Mach units; the distance of closest approach of the source to the microphone is d, measured in seconds per Mach unit. Linear motion is assumed.

$$f(t) = \frac{f_0}{1 + \frac{v_0^2 (t - t_0)}{\sqrt{d^2 + v_0^2 (t - t_0)^2}}} \tag{1}$$

If we now fit the above data set to a model of the Doppler shift as expressed in Eq. 1 using nonlinear methods, we obtain the results summarized in Figure 14.

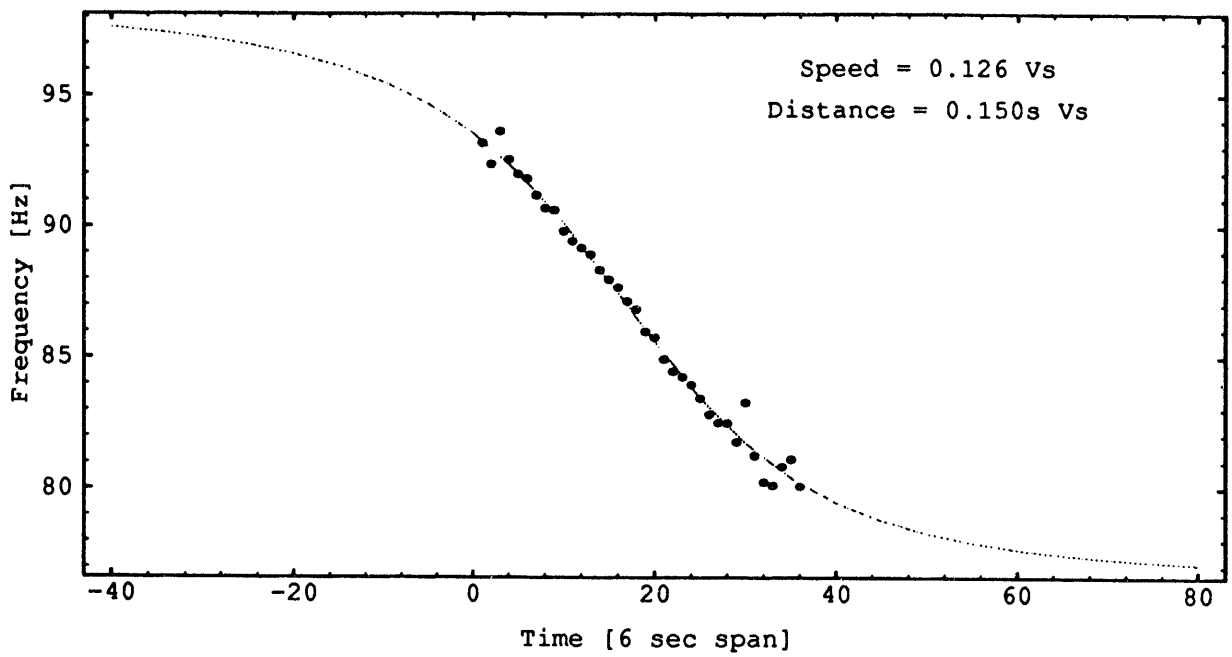

Figure 14. Fit of a Doppler model (Eq. 1) to the data shown in Fig. 13. Vs is the speed of sound at the runway. Abscissa units are the same as in Fig. 13.

The parameters of the fit are t_0 = 18.7 t-units (aircraft opposite microphone), transverse frequency or f_0 = 86.2 Hz, which is very close to the mean of the data set (86.4 Hz), d = 2.955 t-units, and v_0= 0.126 Mach units. The t-units are about 51 ms (see above) and the speed of sound (standard atmosphere at 100 m altitude) is about 340 m/s. Thus, the two essential parameters of the motion of the aircraft are (assuming constant velocity, which is borne out by the symmetry and quality of the fit):

speed of aircraft = 42.8 m/s
distance to sensor = 51.9 m.

While the analysis can be extended along similar lines to obtain additional information about the event, the parameters derived above are the easiest to obtain and should be available in real time given an adequate hardware platform and appropriate algorithms. A simple error-propagation analysis can be applied to assign confidence limits to these results and aircraft acceleration could be included in the model without much additional computation.

6. DISCUSSION

This paper presented a set of signal-processing methods based on wavelet analysis that has been successful in high-reliability identification of aircraft type from acoustic measurements. In addition, we demonstrated a method for discriminating between aircraft sounds made by other technological sources as well as those generated in the natural environment. Moreover, extraction of physical parameters related to both the aircraft type and event was shown to be a straight-forward task. While there are several variations possible for each particular methods developed—for example that of Ref. 11 is more efficient at following the Doppler shift than the recursive ARMA method discussed above, the essential idea of projecting the original signal onto multiple wavelet subspaces prior to feature extraction and system-identification allows for a more comprehensive analysis with higher quality results and less overall computation.

6.1 Results of wavelet analysis

Wavelet filter-bank analysis of the digitized acoustic aircraft signatures is a computationally inexpensive step in the general problem of signal processing for recognition and identification. We showed how to use characteristics of the projected subspace signals for extracting features suitable for event classification and frequency tracking. The latter provided the input for a Doppler model that gave accurate indication of distance to the single microphone as well as transverse aircraft speed. No doubt, additional physical parameters remain to be extracted from the data sets. Some possibilities are wing vibrations that could modulate the fundamental oscillatory excitations and, perhaps, an estimate of turbine parameters is feasible for the large jets.

6.2 Results of classification

Both neural networks and fuzzy-set clustering were explored as classification schemes. A constraint maintained throughout the project was computational efficiency since a deployable instrument package remains a goal. This eliminated several good classification schemes and reduced the choice to fuzzy-set classifiers and radial-basis function networks. The former showed that a signature can be numerically characterized by its degree of membership in a fuzzy set, while the latter indicated that certain ANN models provide efficient event classification.

Both classification methods demonstrated a reasonable degree of success in recognizing complex signals in spite of sparse data set. At this time, neither is offered as the ultimate classification scheme for the signal types explored in this work.

6.3 Conclusions

Although the data collected for this project were analyzed with a computer workstation, we have written similar code for a digital-signal processing platform where the computation takes place in real-time. In addition, the problem of using a small and inexpensive microphone has been explored by preprocessing the raw digitized data with a filter whose response function approximates one of the small and inexpensive commercially available devices. No degradation in classification capability was noted. It seems feasible to package a suitable analysis system in a small, inexpensive and portable device for both surveillance and additional data-acquisition endeavors.

The work on classification should be extended to include a more extensive analysis, using much larger data sets, both for the fuzzy-set method and the ANN approach. Additionally, invariants extracted from each of the wavelet subspaces by the techniques of nonlinear time-series analysis should provide additional discriminating features as well as serve as a point of departure for more extensive system-modeling efforts.

In spite of the preliminary nature of work described, the method shows great promise for classification of complex, time-varying signals and as an aid in system modeling and identification.

7. ACKNOWLEDGMENTS

We wish to thank R. S. Thomas and B. M. Beets for their data-gathering excellence and V. M. Baylor and S. A. McKenney for their moral and financial support.

This research was performed at Oak Ridge National Laboratory, operated by Martin Marietta Energy Systems, Inc., for the U. S. Department of Energy under Contract No. DE-AC05-84OR21400.

8. REFERENCES

1. C. M. Harris, *Handbook of Acoustical Measurements and Noise Control*, McGraw-Hill, New York, 1991.

2. M. K. Kallergis, "Flight tests on propeller noise pressure time history characteristics," *Noise Control Engineering Journal*, Vol. 29, pp. 32-37, 1987.

3. R. N. Hosier and D. A. Hilton, "Some Effects of the Atmosphere and Microphone Placement on Aircraft Flyover Noise Measurements," *NASA TM X-72791*, NASA, 1975.

4. Stephane G. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Trans. on PAMI*, Vol. 11, No. 7, pp. 674-693, July 1989.

5. Stephane G. Mallat, "Multifrequency Channel Decompositions of Images and Wavelet Models," *IEEE Trans. on ASSP*, Vol. 37, No. 12, pp. 2091-2110, December 1989.

6. Ingrid Daubechies, *Ten Lectures on Wavelets*, SIAM, Philadelphia, 1992.

7. Benjamin Kedem, *Time Series Analysis by Higher Order Crossings*, IEEE Press, Piscataway, New Jersey, 1994.

8. Bart Kosko, *Neural Networks and Fuzzy Systems*, Prentice-Hall, Englewood Cliffs, New Jersey, 1992.

9. T. Poggio and F. Girosi, "Regularization Algorithms for Learning that are Equivalent to Multilayer Networks," *Science*, Vol. 247, pp. 978-982, 1990.

10. William F. Lawkins, C. Stewart Daw, Darryl J. Downing, and Ned E. Clapp, Jr., "Role of low-pass filtering in the process of attractor reconstruction from experimental chaotic time series," *Phys. Rev. E*, Vol. 47, No. 4, pp. 2520-2535, 1993.

11. S. Kadambe and G. F. Boudreaux-Bartels "Application of the Wavelet Transform for Pitch Detection of Speech Signals," *IEEE Trans. on Inf. Th.*, Vol. 38, No. 2, pp. 917-924, 1992.

DATE
FILMED
10/12/94
END