Diagnosis Using the Formal Theory of a Signal-Processing System

HAMID NAWAB, VICTOR LESSER AND EVANGELOS MILIOS

***Abstract*—Many signal-processing systems can be viewed as transforming their input signals into a representation of the real-world scenario from which the signals originated. Such systems usually have parameters whose settings are selected on the basis of the class of expected input scenarios. Finding the appropriate parameter settings for a class of input scenarios usually involves testing the system against typical and/or important input scenarios from that class. Whenever the system output does not match the input scenario, the parameter settings responsible for the fault are identified. The system user can then adjust the system parameters to ensure correct system behavior for such scenarios. The diagnostic process of identifying the parameters responsible for system faults is generally difficult because the signal-processing system carries out a complicated mathematical transformation involving a multistage algorithm that generates an enormous amount of intermediate data. A new approach to the diagnosis of such systems is developed. The approach is based on the availability of an abstract and possibly qualitative description of the input scenario and the use of an alternative system model derived from the underlying mathematical theory that explicitly represents the phenomena responsible for any incorrect processing. This approach to diagnosis models a system as a combination of processes that transform the user-specified abstract description of the input scenario into the system output. Whenever the correct answer is obtained at the system output, each process reduces to an identity transformation at the level of abstraction of the system output. Thus system faults are viewed as being caused by one or more of the processes becoming nonidentity transformations. These processes have the advantage of theoretically relating each particular phenomenon causing the incorrect processing to just those system parameters that affect the phenomenon. Diagnosis involves finding this set of nonidentity processes through a general problem solver (GPS) type of search that operates at different levels of abstraction. An implementation of this diagnosis approach is presented for an acoustic signal-processing application.**

I. Introduction

An ACOUSTIC signal-processing system [3] for determining the directions of low-flying aircraft is an example of a problem-solving system which carries out transformations based on Fourier theory. The system receives its inputs from a microphone array on the ground that detects the sounds of nearby low-flying aircraft. The system maps these time-domain input signals into frequency-domain representations which are interpreted as indicating the number of aircraft in the input scenario and their individual directions.

In designing such a real-time signal-processing system, a trade-off is made between the scope of the input scenarios that can be correctly identified and the computational resources required. One way of gaining efficiency without limiting generality is to use *a priori* knowledge about the class of input scenarios to focus the system resources for efficient recognition. This is accomplished by introducing parameters into the signal-processing system for adjusting its behavior to a specific class of inputs (e.g., scenarios involving only helicopters). Changing to a different class of input scenarios is accomplished by changing the values of some of the system parameters.

Parameter settings are usually selected by testing the system against typical scenarios from the class of interest. Abstract and possibly qualitative descriptions of these scenarios (usually provided by a human observer) are compared to the system output. If inconsistencies exist between the system output and an input scenario description, an explanation is required for why the signal-processing system failed in that particular instance. Assuming that system faults do not arise because of errors in system implementation (hardware of software), it is advantageous if *the cause of failure is described in a way that clearly identifies the system parameter settings responsible for the faults in the system output*. The system operator can consequently adjust the system parameters to change the transformation applied by the signal-processing system appropriately.

Our research has addressed the issue of *automating the diagnosis of system faults that arise from incorrect parameter settings*. We have focused on fault diagnosis for systems whose designs are based on underlying mathematical theories. In particular, we have developed and implemented a diagnosis system for an acoustic signal-processing application with an underlying Fourier theory. We show that in such cases knowledge of the underlying theory of the system is useful for the diagnostic process. This use of the underlying theory for diagnosis comes from our observation of how experts perform diagnosis in this domain.

Diagnosis strategies that do not exploit an underlying theory can be impractical for data-processing systems (such as the acoustic signal-processing system) that generate enormous amounts of complex intermediate data. Based on our observations, experts with knowledge of the un-

Manuscript received January 15, 1986; revised August 1, 1986. This work was supported in part by the Defense Advanced Research Projects Agency. The views expressed are those of the author(s) and do not reflect the official policy or position of the US Government.

H. Nawab was with the Lincoln Laboratory, Massachusetts Institute of Technology, Cambridge, MA. He is now with Electrical, Computer, and Systems Engineering, Boston University, Boston, MA 02215.
V. Lesser is with the Department of Computer and Information Science, University of Massachusetts, Amherst, MA 01003.
E. Milios is with the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139.

IEEE Log Number 8714384.

0018-9472/87/0500-0369$01.00 ©1987 IEEE

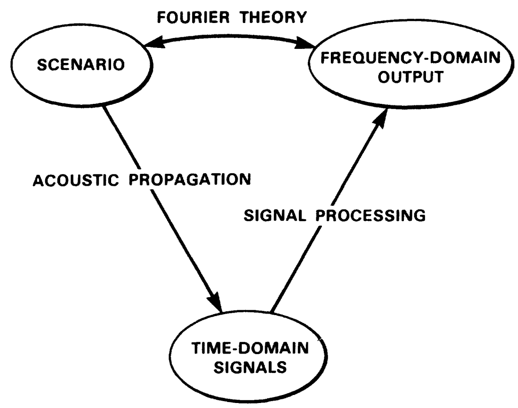

Fig. 1. Fourier theory provides alternative relationship between scenario and frequency-domain output of signal-processing system.

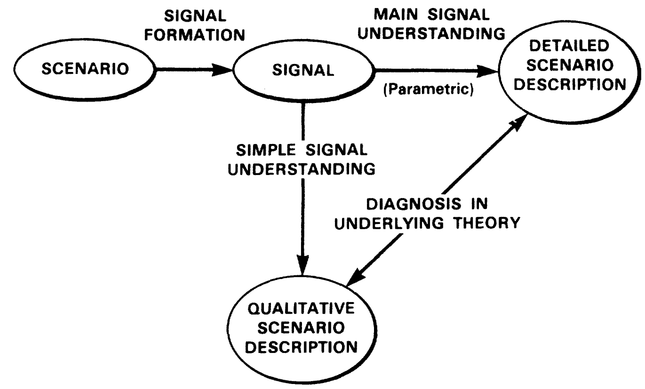

Fig. 2. Diagnosis framework for signal understanding.

derlying system theory as a rule do not use intermediate data for routine diagnosis of the signal-processing system. Without the underlying theory the only known relationship between the input and the output of a system is the transformation carried out by the system itself. In such cases, causal analysis techniques are based on the analysis of intermediate data states that have been either saved during system operation (as in [2]) or have been regenerated by simulation (as in [1]). In either case the amount of complex intermediate data that has to be analyzed can become prohibitive. In addition, these data cannot be suitably abstracted for those techniques.

However, if a system has an underlying theory and there is access to the description of the input and output of the system, the analysis of intermediate data can be avoided for the purpose of diagnosis. For example, in the case of the acoustic signal-processing system, the underlying Fourier theory of the system provides an alternative description of the relationship between the input scenario (the aircraft and their motion) and the information contained in the output of the signal-processing system. The Fourier theory also describes how changes in the various parameters of the signal-processing system affect the relationship between the scenario information and the frequency-domain output of the system. Since Fourier theory adequately describes the effects of scenario characteristics and system parameter settings on the information in the signal-processing output, our diagnosis strategy does not require inspection of the intermediate data states (including the signals received at the microphones). Thus the system is viewed not as a transformation from the time-domain microphone signals to the frequency-domain output, but rather as a transformation from a frequency-domain description of the input scenario to the frequency-domain output (see Fig. 1). Note that this view permits us to take into account any changes in the directional information that take place as the sound waves propagate from the aircraft to the microphone array.

The changes due to propagation effects as well as those due to the specific parameter settings of the signal processing can be modeled in the Fourier domain as a set of processes that transform the user-specified description of the input scenario to the system output. Each process is described in terms of how it transforms its input information. This transformation depends upon the values of the parameters associated with a process. Further, the parameters of the processes have direct correlates in the parameter set of the actual system because they are derived from the same underlying theory. When the system is operating correctly (the system parameters are appropriately set), each of the processes performs an identity transformation from its input to its output. System faults are thus modeled as the result of nonidentity transformations in one or more of the processes. Diagnosis from this perspective involves finding through a general problem solver (GPS)-type search a minimal set of nonidentity processes that explain the discrepancy between the user-specified information and the system output.

The approach presented in this paper for the *diagnosis of system faults that arise from incorrect parameter settings* has implications for signal-understanding systems in artificial intelligence. It can be used for the diagnosis of any signal-understanding system that has a theory (not necessarily Fourier theory) that provides alternative descriptions of how the desired information about the input scenario is changed by distortion phenomena to produce the information represented by the system output. Each of the distortion models must have a parametric representation in the underlying theory, and these parameters should have direct correlates in the actual system parameters. The approach requires that abstract and/or qualitative descriptions of the input scenario be available. These may be provided either by a human observer or by an alternative signal-understanding system that is simpler and less precise in the description of the input scenario while being simultaneously less likely to produce errors. The diagnosis system then identifies the parameters responsible for the errors in the output of the main signal-understanding system. The system parameters may then be adjusted to correct the behavior of the main signal-understanding system (see Fig. 2). Our current research, not discussed in this paper, focuses on strategies for the automatic adjustment of parameters of signal-processing systems with an underlying theory, once the fault parameters have been identified by the diagnosis system.

In our diagnosis strategy for the acoustic signal-processing system, we provide a GPS framework with states and operators that are at *various levels of abstraction*

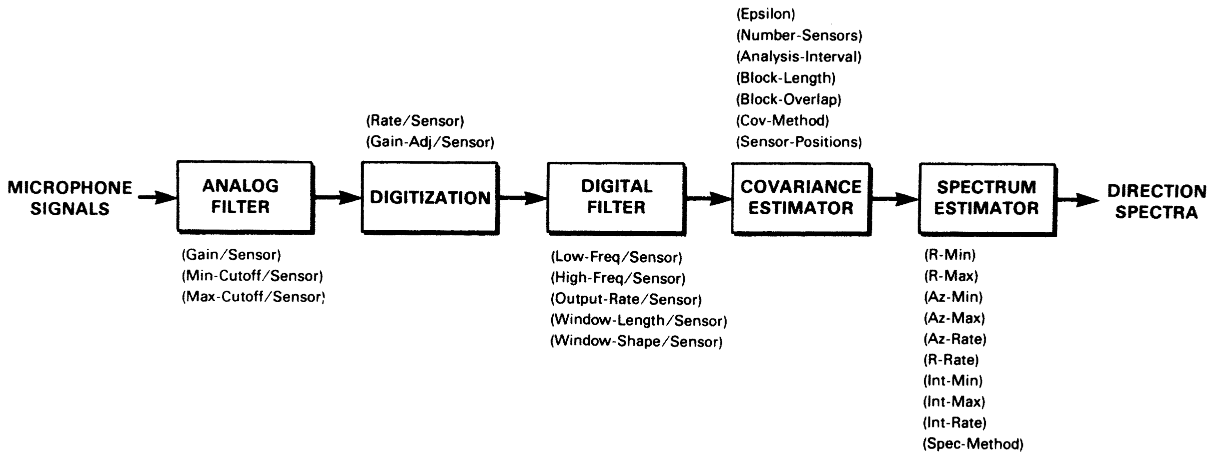

Fig. 3. Stages of signal-processing algorithm and their major parameters.

and whose descriptions are *qualitative*. In particular, the abstractions considered are not restricted to the notion of *loss of detail* but also include *changes of representation*. We believe this to be an important use of the GPS paradigm because in many real-world applications the search space is described using different representations at different levels as well as using qualitative descriptions at each of those levels.

In this paper we use the acoustic signal-processing example to illustrate our diagnostic approach and its implementation as a working system on a Symbolics 3600 Lisp machine. In Section II, we will describe the signal-processing transformation and examine briefly the Fourier theory concepts which are useful for our diagnosis system. In Section III, we describe how our diagnosis approach makes use of Fourier theory. This is followed in Section IV by an example illustrating our diagnosis approach.

II. The Acoustic Signal-Processing System

The acoustic signal-processing system determines the directions of low-flying aircraft from recordings made on several microphones that are spatially distributed on the ground. The system implements a multistage algorithm with over 20 major adjustable parameters (see Fig. 3). The signal analysis is done repeatedly every few seconds to keep the direction estimates updated. Each time the analysis is done, the output is a data structure known as a *direction spectrum*. This data structure explicitly represents the directions of each of the aircraft detected by the microphones. In Fig. 4, we have illustrated an abstraction of the direction spectrum, indicating how the information about a particular aircraft is represented.

The signal-processing system can be described in terms of frequency-domain processes that map the frequency-domain description of the input scenario into the frequency-domain output of the system. There are numerous frequency-domain processes involved in the signal-processing system. In Table I we have listed short descriptions of some of these processes. To illustrate the nature of these processes, let us consider the *equal-resolution* process. It succinctly describes the causality of events behind a phenomenon that occurs when two aircraft are very close

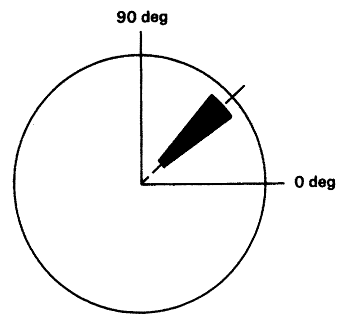

Fig. 4. Abstraction of two-dimensional direction spectrum for aircraft whose direction is between 35 and 50°. Shaded sector represents signal corresponding to aircraft. Angles subtended by sector represent range of possible directions for corresponding aircraft. Radial length of sector denotes maximum possible bandwidth for signal. In this case, minimum possible frequency for signal is 50 Hz and maximum possible frequency is 180 Hz.

in direction with respect to the microphone array and the mappings of their frequency spectra onto the direction spectrum have identical ranges (Fig. 5(a)). Depending upon system parameters, the signal-processing system may produce the direction spectrum of Fig. 5(b) which contains just one signal that is a combination of the two signals in Fig. 5(a). Such a merger is a resolution phenomenon. Since in this particular case the two original signals have equal frequency ranges, we refer to it as an equal-resolution phenomenon. The system parameters that determine whether equal-resolution has any effect in a particular situation are spread among the various algorithmic stages shown in Fig. 3. That is, the resolution process is not local to any particular stage of the signal processing chain; it is affected by each of the stages in some way. By representing the signal processing in the Fourier domain, we are thus able to isolate the system parameters that are involved in the equal-resolution phenomenon.

Some processes are more localized with respect to the actual signal-processing stages. In Table II we show the correspondence between the algorithmic stages and the various processes involved in the signal processing. For example, the *antialiasing-filtering* process is localized to the first three stages in Fig. 3. In the Fourier domain this process is modeled as the elimination of that portion of the spectrum that lies on frequencies greater than the filter

TABLE I
SUMMARY LIST OF SOME FOURIER DOMAIN PROCESSES

Equal-Resolution Operator: Merges two signals into one while retaining portions of the original signals. This occurs when the two input signals are close in direction and their frequency distributions are the same. The effects of this operator are dependent upon array aperture and the epsilon factor for the covariance matrix.

Partial-Resolution Operator: Merges two signals into one while retaining portions of the original signals. This occurs when the two input signals are close in direction but their frequency distributions are different and/or resolution does not occur over all the signal frequencies. The effects of this operator are dependent upon array aperture and the epsilon factor for the covariance matrix.

Contained-Resolution Operator: Merges two signals into one while retaining portions of the original signals. This occurs when the two input signals are close in direction and the frequency distribution of one signal is contained within the frequency distribution of the other. The effects of this operator are dependent upon array aperture and the epsilon factor for the covariance matrix.

Fast-Velocity Operator: Shifts the direction of a signal because of fast aircraft velocity. The degree of direction shift depends upon the analysis-interval parameter of the signal processing system.

Elevation-Compression Operator: Shifts the frequency downward and the signal amplitude correspondingly upward. This occurs in accordance with the elevation of the aircraft with respect to the sensor array.

Discrete-Radius Operator: Causes the frequencies and the amplitude of a signal to change. It occurs because the wavenumber spectrum is computed only at discrete points and over a finite interval over a radius in the wavenumber plane. The effects depend upon the sampling interval along a radius as well as the radial range over the computations are performed.

Discrete-Azimuthal Operator: Causes the frequencies and the amplitude of a signal to change. It occurs because the wavenumber spectrum is computed only at discrete angles. The effects depend upon the array aperture and the sampling interval.

Car-Filter Operator: Causes frequency changes in the signal because of the filter performing the computation for complex analytic representations of the signals.

Anti-aliasing Filter Operator: Causes frequency changes in the signal because of the analog filter meant for avoiding temporal aliasing in the digitization process.

Car-Ghosting Operator: Causes a ghost signal to appear at the direction which is 180° away from the original signal. The phenomenon occurs if the original signal has considered power in the lower frequencies. The effects of the operator depends upon the block-length parameter of the signal-processing system.

Spatial-Aliasing Operator: Causes the signal to wrap around at different directions due to the sparseness of the array. The effects of this operator depends upon the minimum sensor separation in the acoustic array.

Temporal-Aliasing Operator: Changes the frequencies and the amplitude of a signal because of the undersampling in the digitization process.

Peak-Picking Operator: Causes signals to disappear because of the power thresholding applied to all the signals in the wavenumber spectrum. The effect obviously depends upon the threshold parameter.

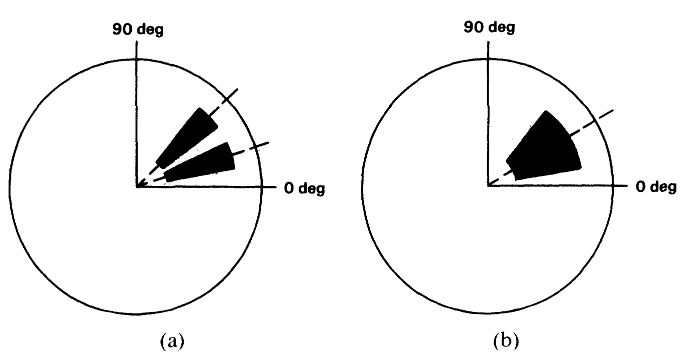

Fig. 5. (a) Direction spectrum before resolution for situation with two aircraft that are close in direction and have identical frequencies. (b) Direction spectrum after equal-resolution phenomenon has been taken into account.

TABLE II
CORRESPONDENCES BETWEEN PROCESSES AND ALGORITHMIC STAGES OF FIG. 3

Equal-Resolution: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Partial-Resolution Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Contained-Resolution Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Fast-Velocity Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Pre-CPA Doppler Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Post-CPS Doppler Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Covariance-Instability Operator: Analog-filter, digitization, digital filter, covariance estimator.

Elevation-Compression Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Discrete-Radius Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Discrete-Azimuth Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Car-Filter Operator: Analog-filter, digitization, digital filter.

Anti-Aliasing Filter Operator: Analog-filter.

Car-Ghosting Operator: Analog-filter, digitization, digital filter.

Spatial-Aliasing Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Temporal-Aliasing Operator: Analog-filter, digitization.

Covariance Size Operator: Analog-filter, digitization, digital filter, covariance estimator.

Peak-Picking Operator Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Radial-Integration Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

Range-Scaling Operator: Analog-filter, digitization, digital filter, covariance estimator, spectrum estimator.

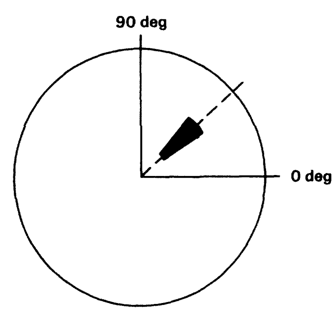

Fig. 6. Direction spectrum of Fig. 4 after its transformation by anti-aliasing-filtering process.

cut-off. For example, if the cut-off frequency is 120 Hz, the antialiasing-filtering process transforms the direction spectrum of Fig. 4 into the direction spectrum of Fig. 6. Now suppose that the cut-off frequency were 50 Hz instead of 120 Hz. Clearly, in that case the signal in Fig. 4 would be completely eliminated, and thus our signal-processing system would not be able to detect that aircraft. When such a situation occurs, it is the task of our diagnosis system to identify the offending process (antialiasing-filtering) and its associated parameters (cut-off frequency). Consequently, the cut-off frequency may be increased to ensure detection of the aircraft in that situation.

Our diagnosis strategy, which we describe in the next section, makes use of multiple levels of abstraction for the

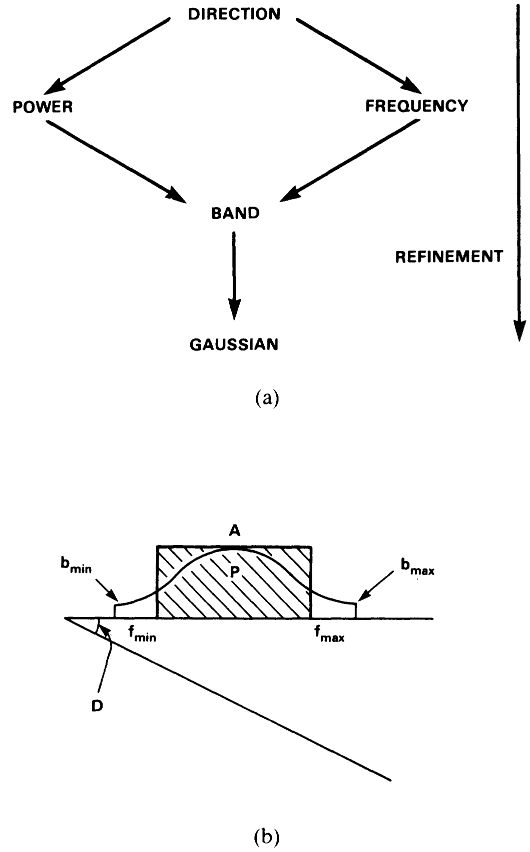

Fig. 7. (a) Refinement hierarchy of signal abstraction levels. (b) Illustration of abstraction levels. Angle D represents direction level. Area P of rectangle is added for power level, while f min and f max are added for frequency level. Band level combines power and frequency levels as well as specifying amplitude A of rectangle. Gaussian level is represented by curve and its limits b min and b max.

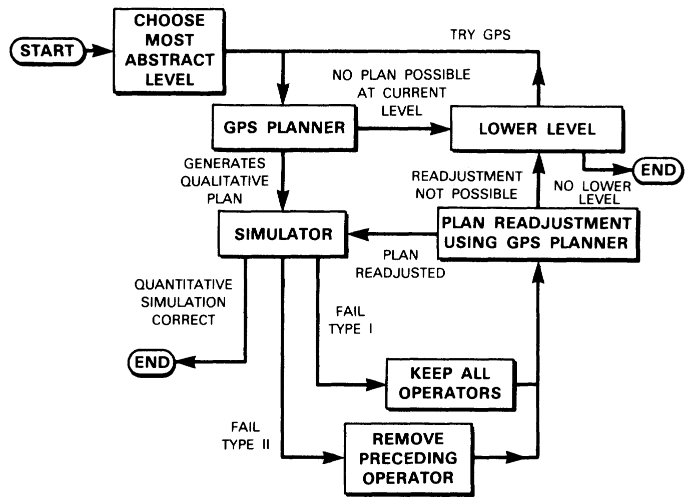

Fig. 8. Plan-and-verify strategy.

Fourier domain representation of the direction spectrum. We use five distinct levels of abstraction: direction, power, frequency, band, and Gaussian levels. Their refinement hierarchy is illustrated in Fig. 7. At the direction level, each signal is associated with just one characteristic—its direction in the direction spectrum. Other characteristics of the signals are hidden at this level. At the band level we may either choose to describe the maximum and minimum frequencies of the signal (frequency level) or just the net power in the signal (power level). The next level combines the frequency and power-level descriptions. Each signal is represented as a rectangle with a minimum and a maximum frequency and whose area represents the net power in the signal. The height of the rectangle represents the average amplitude of the signal. At the lowest level of abstraction the rectangle representation is replaced by a Gaussian representation such that the maximum amplitude of the Gaussian equals the height of the corresponding rectangle. The width of the rectangle is related to the standard deviation of the Gaussian. These are standard abstractions used in the Fourier domain descriptions of signals [4]. Note that in contrast to systems like ABSTRIPS [5], the abstraction used in our approach involves changes of representation in addition to the usual suppression of detail.

III. Diagnosis Based on Fourier Domain Models

We now describe our general strategy for diagnosis in terms of how it has been applied in the case of the acoustic signal-processing system of Section II. We assume that the user provides the diagnosis system with a qualitative and abstract description of the input scenario (in the form of a direction spectrum) for the situation under diagnosis. A *plan-and-verify paradigm* is used by the diagnosis system to identify a sequence of frequency-domain processes that are responsible for the discrepancy between the system output and the correct answer.

Given the frequency-domain representation of the acoustic signal-processing system, the problem of determining the nonidentity processes involved in a particular situation can be viewed as a state–space search; the states (to be called signal states for our application) are directional spectra, and the operators are the various processes. The initial state is derived through a straightforward computation from the user-specified description of the correct answer. The goal state is a description of the directional spectrum at the output of the signal-processing system.

We now phrase our diagnosis task as the following search problem. *Identify a sequence of processes such that* 1) *the sequence maps the initial signal state into the goal signal state and* 2) *no proper subsequence maps the initial signal state into the goal state.*

We talk about a *sequence* of processes since, as we shall see later in this section, the order in which the processes are applied to the signal states is important. The requirement that no proper subsequence exist that can map the initial state into the goal state ensures that we exclude any process not necessary for explaining the discrepancy between the correct answer and the system output. Also note that more than one sequence of processes may satisfy our search criterion. This happens when more than one explanation can justify the differences between the system output and the correct answer. Unless intermediate data states of the actual system are available for inspection,

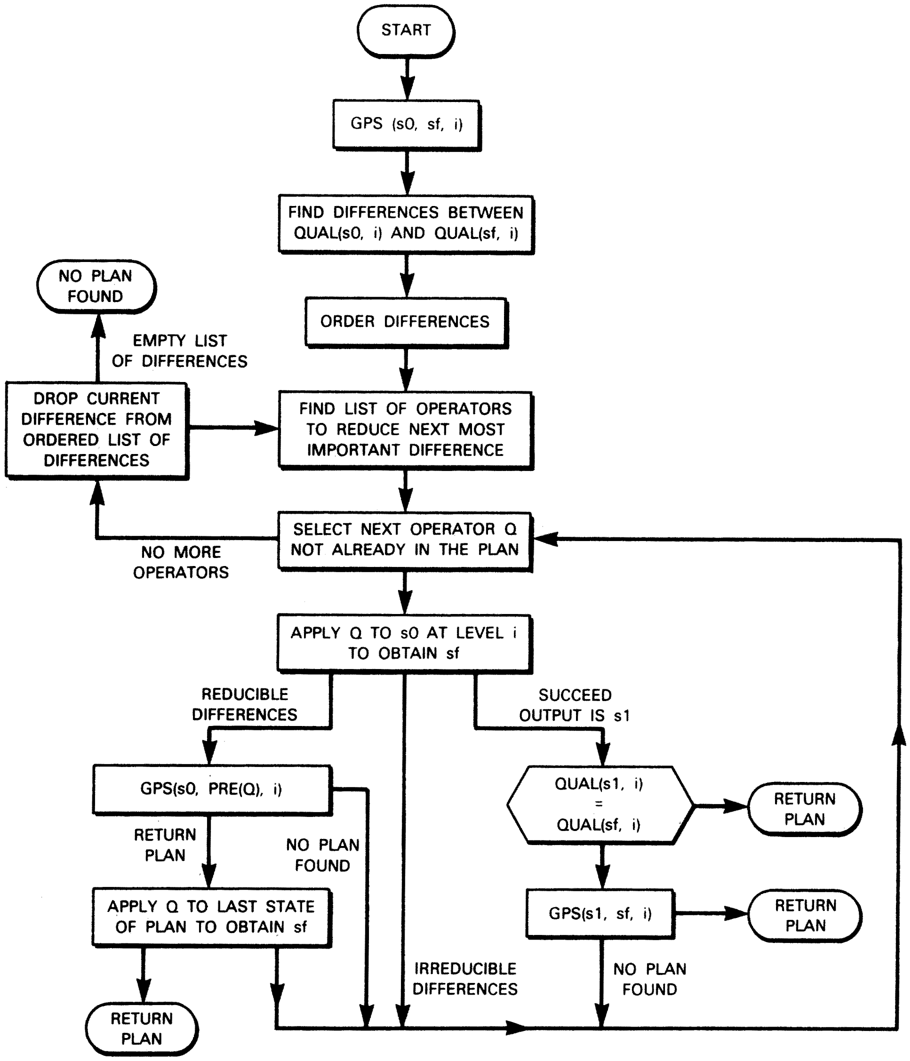

Fig. 9. GPS planner at abstraction level i.

such multiple explanations cannot be disambiguated. For our signal-processing system the intermediate data states are not available for inspection. Consequently, we have designed our search strategy so that it generally finds the sequence with the smallest number of processes. We were guided in this selection by the heuristic that *the simplest explanation is the most likely explanation for the cause of a system fault*.

Our basic search strategy is to first construct an abstract *plan* hypothesizing the sequence of processes and then to *verify* the plan at the lowest level of abstraction. A detailed description of the plan-and-verify strategy for our system is represented in the flow diagram of Fig. 8. We start by selecting the initial abstraction level of the planning phase to be the highest level of abstraction, i.e., the direction level. This is done because by ignoring as much detail as possible, we are able to generate plans with as few operators as possible. However, if the planner cannot find such a plan, it is allowed to drop the abstraction level for portions of the plan where it may be having difficulty. If such local adjustments also fail, the level of abstraction for the entire plan may be lowered. Of course, if the lowest level of abstraction has already been reached, the search ends without having obtained any answer that could explain the discrepancies.

The planning phase of our plan-and-verify strategy uses the generic means–ends analysis technique of GPS at the various levels of abstraction. A block diagram of this planner for any particular abstraction level is shown in Fig. 9. Note that the planner remains at the same abstraction level and thus follows the traditional GPS paradigm. The changes in abstraction levels are introduced by the plan-and-verify strategy of Fig. 8. Our planner for the signal-processing system classifies differences between signal states into seven categories—missing signal, unassociated signal, direction shifting, amplitude scaling, frequency shifting, resolution, and ghosting. It selects the most important difference and an operator that is likely to reduce the difference in the current situation. Control of the GPS search is accomplished through two important mechanisms.

First, no operator is allowed to appear more than once in a particular plan. This follows from the fact that each operator represents a single process in the signal-processing system. This is in contrast to GPS search strategies in problems such as algebraic simplification where the same operator can be used many times over. In those situations the GPS strategy requires "depth heuristics" to avoid fruitless searches involving many instances of the same operator. In our problem, however, the constraint that an operator may appear only once in a plan removes the need for depth heuristics.

A second mechanism for controlling the GPS search in our system is the use of an *ordering relationship* among classes of signal states. In particular, five classes of states exist. *Propagation domain* states represent plane-wave signals propagating through the atmosphere, and they must precede all other states. *Continuous-temporal* domain states come next, and they represent one-dimensional analog signals. Next, *discrete-temporal* domain states represent one-dimensional digital signals. *Continuous-spatial* domain states represent two-dimensional analog wavenumber spectra, and finally, *discrete-spatial* domain states represent digitized wavenumber spectra. Each operator specifies the allowable classes of input and output signal states. Thus our strategy does not permit plans in which the operators violate the domain requirements for their input and output states. This helps to reduce the search space considerably. Note, however, that operators whose input and output states belong to the same single domain can appear in any order with respect to each other.

An example of a GPS-operator description for the *equal-resolution* process discussed in Section II is shown in Table III. The description is qualitative because the pre- and postconditions are specified in terms of ranges rather than specific values. We employ a range-intersection criterion for testing the preconditions. That is, if a precondition requires that a parameter be within a certain range, this condition is considered to be satisfied by any state for which that parameter has a range that intersects with the range specified in the condition. This is sufficient for our purpose because we are seeking any plan that may be possible within the specified uncertainty (expressed as numerical ranges) for the initial state. Also note from Fig. 10 that an operator may have *scenario preconditions*. These are descriptions of the scenario (e.g., aircraft velocity) which are not captured by the direction spectrum representation.

When a candidate plan is generated by the GPS planner, a *verification* of the plan is attempted as the next step. The abstraction level of this verification is selected to be the lowest one at which a description of the input signal state is known. Our verification procedure makes use of the operator and state representation mechanisms used in the GPS planner. In particular, the verification procedure can be viewed as a degenerate case of the GPS planner at the detailed abstraction level. The difference is that in this case the planner does not have to perform any search for operator selection. Instead, it selects the operators in accordance with the plan to be verified. If the plan is successful in reaching the goal state, our plan-and-verify strategy ends with the executed plan representing the desired diagnosis. On the other hand, if the verification fails at some point, further diagnosis is guided by the nature of the failure as discussed next.

There are two basic types of failures. In one case, the preconditions of an operator in the plan are not satisfied

TABLE III
EQUAL-RESOLUTION OPERATOR

Input signal type	propagation, continuous-temporal discrete-temporal, continuous-spatial
Output signal type	continuous-spatial
Differences reduced	resolution
Operator parameters	
Direction	array-aperture
Power	array-aperture
Frequency	array-aperture, epsilon
Band	array-aperture, epsilon
Gaussian	array-aperture, epsilon
State preconditions	per pair of input signals
Direction	direction difference intersects [0,100/array-aperture]
Power	direction level preconditions power in [0,inf]
Frequency	minimum frequencies intersect Maximum frequencies intersect direction difference intersects [0,(1000 * epsilon)/(array-aperture * 0.0001 * maximum-frequency)]
Band	power level preconditions frequency level preconditions amps in [0,inf]
Gaussian	Frequency level preconditions with Gaussian model
Scenario preconditions	none
State postconditions	per pair of input signals
Direction	delete input signals create signal whose direction is the cover of the two input directions
Power	direction level postconditions power of output signal in [0, sum of maximum powers in signals]
Frequency	direction level postconditions minimum frequency of output same as input Maximum frequency of output same as input
Band	frequency level postconditions power level postconditions amps of output signal in [0, sum of maximum amps in signals]
Gaussian	band level postconditions with Gaussian model

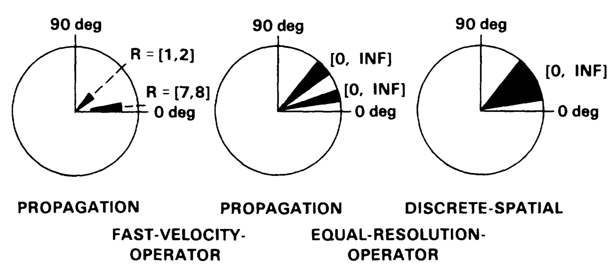

Fig. 10. Computer-generated output from diagnosis system at direction level. Circles denote signal states. Each shaded sector represents individual signal. Angles subtended by sector represent range of directions for corresponding aircraft. Radial extent of each sector denotes possible bandwidth range. Amplitude range of each signal is specified as interval label next to corresponding sector. In plan of figure, we have initial state, which is transformed by fast-velocity operator which in turn is processed by equal-resolution operator. Outputs of two operators have maximum uncertainty in bandwidths and amplitudes because operators were applied only at direction level.

by the state preceding the operator. For such a situation a plan readjustment is attempted by finding a plan for linking the state and the preconditions of the failed operator. In the second type of plan failure the state at the output of an operator does not match the qualitative description anticipated for it in the original plan. For such a situation a plan readjustment is attempted by eliminating the failed operator from the plan and devising a different plan to replace its position in the original plan. Note that the adjusted portions of a plan can be at lower levels of abstraction than the current level of abstraction chosen for the entire planning process. If local readjustment of a plan is not possible, the basic level of abstraction is lowered and a new planning process is started.

Our planning and verification strategy has much in common with the ABSTRIPS [5] problem-solving system, as well as a number of important differences. Both systems use GPS for candidate generation and use a verification environment to test the plans. They differ in that ABSTRIPS uses resolution theorem proving as the basis of the test in the verification environment; our system uses the successive application of plan operators to test whether the initial state is transformed into the goal state. More importantly, however, they differ in how they exploit the abstraction space.

Both systems use multiple levels of abstraction for their operator descriptions. In our system we carry out verification even for plans which are constructed using the highest levels of abstraction in the operator descriptions. In contrast, the ABSTRIPS system does not attempt to execute the plans formed at the higher levels of abstraction. Instead, it uses them in a process of *refinement* to produce plans at lower levels of abstraction. Thus no portion of a plan is considered ready to be executed in the verification environment until it has been refined to the lowest level of abstraction. Thus in some sense the ABSTRIPS system does a depth-first search along degrees of abstraction for plan portions to be executed in the verification environment. On the other hand, our system searches breadth first across a single level of abstraction in search of complete plans to be executed in the verification environment. Our system drops to a lower level of abstraction only when it cannot form a complete plan at the higher levels of abstraction. Furthermore, in our system changing the level of abstraction can involve change of representations, while in ABSTRIPS it involves only suppression of detail.

IV. Example

In this section we use an example to illustrate our plan-and-verify problem-solving strategy. The example was correctly analyzed by a program that implements our diagnosis strategy for the acoustical signal-processing system. In our description of the example, particular emphasis is placed on illustrating how the GPS planning, the different levels of abstraction, and the verification interact with each other. The situation considered in this example is that of two aircraft flying close to a microphone array in such a manner that their directions with respect to the node are close to each other. A resolution phenomenon often occurs in such a situation: the system detects just one acoustic source whose direction is between the two actual directions. However, as we shall see in the following example, the process of forming an explanation that involves the resolution phenomenon can be quite complex.

Let the two aircraft in our example be AIRCRAFT-1 and AIRCRAFT-2, respectively. The initial state describes these aircrafts in accordance with the user-specified information about the scenario. As detailed in the introduction, the user provides a qualitative description of the actual scenario as a result of his own observations. We will denote this initial state by $S0$ and its *qualitative* description at an abstraction level i by Qual($S0, i$). With this notation in mind, we consider the direction-level qualitative description of the initial state in our example:

Qual($S0$, direction):

AIRCRAFT-1 at direction $D1$
AIRCRAFT-2 at direction $D2$
$D1 = [0,10]$ degrees
$D2 = [35,50]$ degrees.

Let us denote the goal state for our problem by Sf and suppose that its qualitative description at the direction level is given by:

Qual(Sf, direction):

AIRCRAFT-3 at direction $D3$
$D3 = [20,20]$ degrees.

Note that since Sf represents the measured direction, it does not have any uncertainty. In accordance with Fig. 8, our GPS planning strategy for transforming the initial state into the goal state starts at the "direction" level of abstraction. Our goal is to transform Qual($S0$, direction) to Qual(Sf, direction). A set of rules determines the differences that exist between $S0$ and Sf. Two of the differences detected are resolution and directional shift. Resolution differences are characterized by having two directions in the input state while the output state has a direction between those two directions as well as possibly having the original two directions. A direction shift is a difference characterized by a direction in the input state shifting to a different direction in the output state.

Another set of empirical rules also assign priorities to the selected differences. In our example, let us assume that the resolution difference is given the highest priority. The next step is to *select* an operator Q that might reduce the highest priority difference. In our example the selected operator would be one called "equal-resolution" (see Table III). This operator acts upon two signals whose directions are close to each other and have the same minimum and maximum frequencies.

The result is a single signal with the same frequencies but a direction enclosed by the original two directions. While the equal-resolution preconditions, qual-pre(equal-resolution, direction), require that the two aircraft be closer

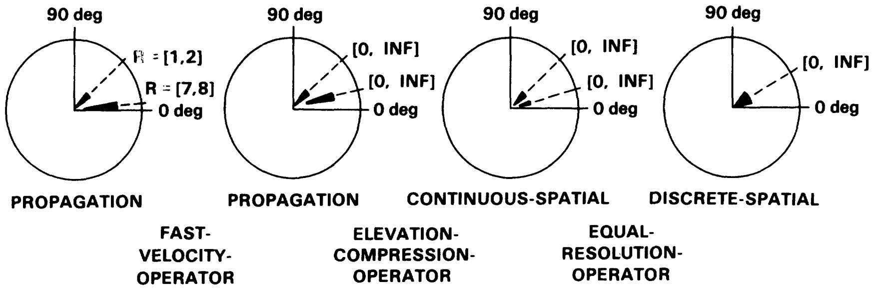

Fig. 11. Plan in example after dropping to frequency level.

than 20° apart, directions in Qual($S0$, direction) are specified to be between 25 and 50° apart. A difference of direction shift thus exists between Qual($S0$, direction) and the qual-pre (equal-resolution, direction). The planner then selects the fast-velocity operator to reduce this difference. The fast-velocity operator represents the effects of fast aircraft velocities on the direction measurements. The qual-pre (fast-velocity, direction) conditions require that an aircraft have a velocity greater than (200/analysis interval) m/s, where analysis interval is one of the system parameters. Suppose that in our example the analysis-interval parameter is set at 4 s and AIRCRAFT-1 has a velocity of 150 m/s with increasing angle. The qual-post(fast-velocity, direction) conditions specify that in such a case that fast-velocity operator will create a new state, say $S1$, in which the direction of AIRCRAFT-1 changes by 6°, making it fall in the interval [12,22]. The direction difference between [12,22] and the direction [35,50] of AIRCRAFT-2 falls in the interval [13,38]. Since this intersects with [0,20], the planner concludes that Qual($S1$, direction) matches the qual-pre(equal-resolution, direction) conditions. Furthermore, the output of equal resolution will be a signal whose direction falls in the interval [12,50]. Thus qual-post (equal-resolution, direction) conditions match the Qual(Sf, direction) conditions. Thus a complete plan at the direction level has been formulated to connect $S0$ and Sf. The plan consists of the fast-velocity operator followed by the equal-resolution operator. A graphical representation of this plan is shown in Fig. 10.

The plan just generated is then passed to the verify stage which tries the plan at the Gaussian level of abstraction. It is found that the output state $S1$ produced by the fast-velocity operator does not match the preconditions of the equal-resolution operator due to the fact that the two aircraft do not have the same minimum and maximum frequencies. (Note that since these frequencies are represented as ranges of values, two frequencies are considered to be the same only if their ranges have a finite intersection.) This is the failure labeled type I in Fig. 8. In accordance with Fig. 8 we move to the plan adjustment phase. Our goal now is to reduce the frequency-shift difference between $S1$ and the preconditions of the equal-resolution operator. The frequency-shift difference cannot be dealt with at the direction level because frequency ranges cannot be described at that level. The planner, therefore, drops to the next abstraction level known as the frequency level. The frequency ranges of the signals in Qual($S1$, frequency) are nonoverlapping while qual-pre (equal-resolution, frequency) requires them to be overlapping. The elevation-compression operator is selected for reducing this difference. The elevation-compression operator represents the phenomenon of the compression of frequency ranges for aircraft that are at high elevations with respect to an array. The qual-pre(elevation-compression, frequency) conditions require that an aircraft form a positive elevation angle with respect to the sensor array. Suppose that the elevation angle of AIRCRAFT-1 is specified to be in the interval [30,45]. Thus elevation compression applies to $S1$ and produces another state $S2$. This has the effect of reducing the minimum and maximum frequencies of AIRCRAFT-1. It turns out that the Qual($S2$, frequency) conditions and the qual-pre(equal-resolution, frequency) conditions match each other. We, therefore, have an adjusted candidate plan consisting of three operations: fast velocity, elevation compression, and equal resolution. A graphical representation of this plan is shown in Fig. 11.

The adjusted plan is then passed to the verify stage at the Gaussian level of abstraction. It is found that at the Gaussian level the output state $S2$ of elevation-compression operator does not match the preconditions of the equal-resolution operator. A frequency-shifting problem still exists. Thus the post-conditions of the elevation-compression operator did not achieve their intended goal. This is the kind of plan failure labeled type II in Fig. 8. In accordance with the strategy of Fig. 8, we drop the elevation-compression operator and seek a plan to replace it. No operator that can change the frequencies appropriately is found. The abstraction level is dropped to the band level. To reduce the frequency difference, the planner selects the CAR-filter operator. The complex analytic representation (CAR) filter changes the amplitudes of the frequency components of a signal, making some of the components zero, reducing the amplitudes of others. The CAR filter followed by the elevation-compression operator succeeds in removing the frequency difference. Thus the equal-resolution operator also becomes applicable. The final plan consists of fast velocity, CAR filter, elevation compression, and equal resolution. The graphical representation of this plan, which for our example also succeeds in the verification stage, as shown in Fig. 12. The parameters of these processes then become candidates for change.

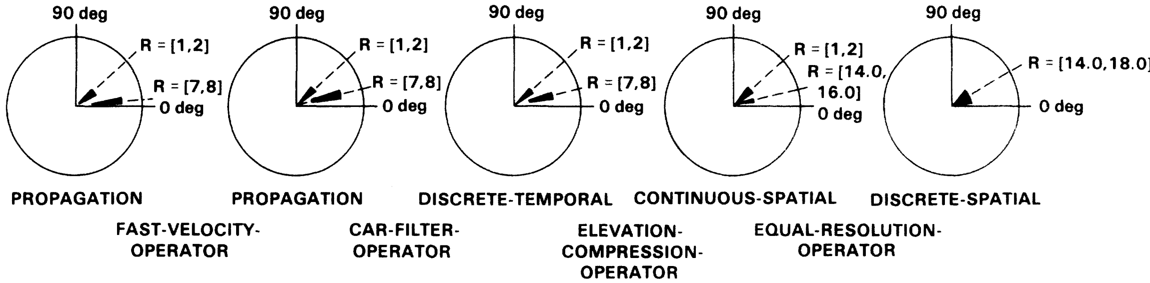

Fig. 12. Plan in example after dropping to band level.

V. Summary

We have presented a plan-and-verify strategy for the diagnosis of systems that can be modeled as carrying out mathematical transformations based on the underlying theory of an application domain. The strategy, deduced from our observation of an expert diagnostician, was described in the context of an automated diagnosis system we are in the process of developing for a signal-processing application. The diagnosis approach was motivated by the following observations about the signal processing application.

- If we have some *a priori* knowledge regarding the correct answer to the problem being solved, diagnosis can be carried out using *only* Fourier domain representations for describing the signal processing. This is in contrast to the actual signal-processing algorithms, many of which are designed to carry out operations on time-domain representations of the signals. Such a change in representation is desirable because it allows us to use conceptual abstractions readily available in the underlying Fourier theory of the signal-processing system.
- We can use abstract models of the underlying Fourier theory [3] for the signal processing to form a conceptual view of the system as a collection of interacting *processes* that can potentially modify a Fourier representation of the correct answer.
- The correctly functioning system can be viewed as having each process as an identity transformation from its input to its outputs at a certain level of abstraction. The incorrectly functioning signal-processing system can then be viewed as having one or more of its processes acting as *nonidentity* transformations. Diagnosis involves finding such nonidentity processes.
- The number of conceptual processes and the associated input and output states required to represent a particular situation in terms of the underlying Fourier theory is much smaller than the number of intermediate data states in the actual signal-processing system.

From these observations, it is clear that a major component of a diagnosis system has to be a problem-solving strategy for determining the set of nonidentity processes that explain the discrepancy between the user-specified information and the output of the signal-processing system. This involves choosing an appropriate problem-solving paradigm as well as an associated representation formalism. In this paper we described a plan-and-verify approach for this purpose. The key characteristics of this approach are as follows.

- An *abstract planning phase* uses means–ends analysis of the type used in GPS to hypothesize a plan (ordered set of processes) for transforming the correct answer (initial state) to the system derived answer (goal state) at one of five levels of abstraction. The processes are represented as operators in the GPS framework. The parameters of the operators represent the current parameter settings of the signal-processing system.
- A *qualitative representation* describes states in terms of component descriptions whose values are specified as numerical ranges and operators which act upon such qualitative descriptions. This leads to a mechanism for dealing with uncertain or approximate information; a qualitative matching criterion considers the preconditions of an operator to be matched by an input state if the ranges in the precondition descriptions of the operator intersect with the ranges in the state description. That is, if enough uncertainty exists about the value of a component of a state description, then an operator is allowed to assume that the component falls within the requirements of its preconditions. For example, a completely unknown state component would satisfy any precondition for it whatsoever.
- A *verification phase* executes the GPS-produced plan. The purpose is to check whether the plan succeeds in transforming the GPS initial state into the GPS goal state. This verification is carried out at the lowest level of abstraction.
- A *simplest plan heuristic* prefers plans that are the simplest in the sense that they involve the fewest number of operators.
- An *abstraction strategy* considers all plans at a higher abstraction level before considering plans at a lower abstraction level. This is done because plans generated at a higher abstraction level are generally simpler in the sense that they involve fewer operators and fewer parameter settings.
- A *recovery-from-failure strategy* forms the interface between the planner and the verify stage. If at any point the verify stage fails to validate the expectation of the planner, the nature of the failure is reported to the planner where GPS is invoked to modify the plan in the local vicinity of the point of failure; the level of abstraction may be lowered in such situations to generate a more detailed subplan for the adjustment to the original plan.

Although we have focused our attention upon the diagnosis of classical signal-processing systems which have Fourier theory underlying their design, our approach is more generally applicable to systems that have a formal theory underlying their design. This includes other types of signal-processing systems as well as systems that are completely outside the signal-processing area. This may include among many other mechanical, biological, and economic systems. Sufficient *a priori* information about the correct answer should be available to allow us to represent the system entirely in terms of the conceptual models. Note that the conceptual models are not just less detailed descriptions of the actual algorithms; they are instead descriptions based on a representation mechanism quite different from that of the actual system under diagnosis.

REFERENCES

[1] R. Davis, "Diagnosis via causal reasoning: Paths of interaction and the locality principle," in *Proc. AAAI-83*, Aug. 1983, pp. 88–94.
[2] E. Hudlicka and V. Lesser, "Meta-level control through fault detection and diagnosis," in *Proc. AAAI-84*, Aug. 1984, p. 153–161.
[3] H. Nawab, F. Dowla, and R. Lacoss, "Direction determination of wideband signals," *IEEE Trans. Acoust., Speech, Signal Processing*, vol. ASSP-34, pp. 1114–1122, Oct. 1985.
[4] A. Oppenheim and A. Willsky, *Signals and Systems.* Englewood Cliffs, NJ: Prentice-Hall, 1983.
[5] E. Sacerdoti, "Planning in a hierarchy of abstraction spaces," *Artificial Intelligence*, vol. 5, pp. 115–135, 1974.

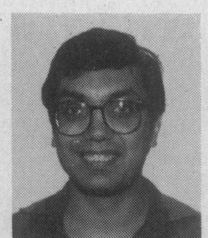

S. Hamid Nawab was born on January 13, 1955 in Swindon, England. He received the S.B., S.M., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1977, 1979, and 1982, respectively.

During his graduate studies he was a research assistant at the MIT Research Laboratory of Electronics (1977–1979) and at the MIT Lincoln Laboratory, Lexington, MA (1979–1982). His graduate research included work on fast DFT algorithms, image enhancement, and signal reconstruction from partial transform information. In June 1982 he was appointed Staff Member at MIT Lincoln Laboratory, where he conducted research on sensor array processing, acoustic signal analysis, and the application of AI techniques to signal processing/interpretation problems. Since September 1985 he has been Assistant Professor of Electrical, Computer, and Systems Engineering at Boston University, Boston, MA. His current research focuses on time–frequency signal representations and applications of AI to signal processing/interpretation problems.

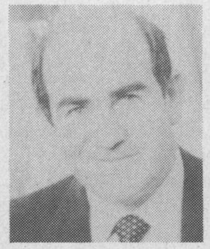

Victor R. Lesser received the B.A. degree in mathematics from Cornell University, Ithaca, NY, in 1966, and the M.S. and Ph.D. degrees in computer science from Stanford University, Stanford, CA, in 1969 and 1972, respectively.

He was a Research Computer Scientist in the Computer Science Department of Carnegie-Mellon University (CMU), Pittsburgh, PA, from 1972 to 1977. At CMU he was a principal in the Hearsay-II Speech Understanding System where he was responsible for the system architecture. While there he also was on the design team of CM* and made contributions in the areas of network operating systems and parallelism in artificial intelligence. Since arriving at the University of Massachusetts, Amherst, in 1977, where he is a Professor of Computer Science, he has lead research groups in the areas of distributed artificial intelligence and intelligent user interfaces. He is also a principal investigator of the recently awarded Coordinated Experimental Research Grant from the National Science Foundation for the study of issues in Distributed and Parallel Problem Solving, and the University Research Initiative Grant from the Office of Naval Research for research in AI.

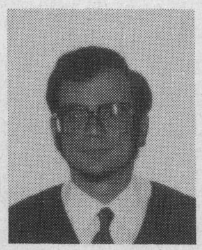

Evangelos E. Milios (S'78–M'86) received a degree in electrical engineering from the National Technical University of Athens, Athens, Greece, in 1980 and the S.M. E.E. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA in 1983 and 1986, respectively.

From 1981 to 1986, he was a member of the MIT Digital Signal Processing Group. From 1982 to 1986, he was a Research Assistant affiliated with the Machine Intelligence Technology Group at the MIT Lincoln Laboratory performing research on the distributed sensor networks project. Since May of 1986 he has been employed as a Research Scientist at the Artificial Intelligence Laboratory, Department of Computer Science, University of Toronto, Toronto, ON, Canada. His current research interests include applications of machine learning to signal classification, qualitative reasoning about signals and signal processing systems, and software engineering of knowledge-based systems.