Report 83-44
Stanford -- KSL

Scientific DataLink

Signal-to-Symbol Transformation:
Reasoning in the HASP/SIAP Program.
H. Penny Nii,
Dec 1983

card 1 of 1

Stanford Heuristic Programming Project **December, 1983**
Memo HPP-83-44

Signal-to-Symbol Transformation: Reasoning in the HASP/SIAP Program

by

H. Penny Nii

Computer Science Department
Stanford University
Stanford, California 94305

Reprinted, with permission, from IEEE Acoustic, Speech and Signal Processing, Spring, 1984.

Signal-to-Symbol Transformation: Reasoning in the HASP/SIAP Program*

H. Penny Nii

Heuristic Programming Project
Computer Science Department
Stanford University
Stanford, Ca. 94305

ABSTRACT

In the past fifteen years, artificial intelligence scientists have built several signal interpretation, or understanding, programs. These programs have combined "low" level signal processing algorithms with knowledge representation and reasoning techniques used in knowledge-based, or expert, systems. [4] They have shown how the use of task domain knowledge combined with symbolic manipulation techniques can be of use in making signal understanding systems more effective and efficient. HASP/SIAP is one such program that tries to interpret the meaning of passively collected sonar data. In this paper we explore some of the AI techniques that contribute in the "understanding" process. We also describe the organization of HASP/SIAP system as an example of a programming framework that show promise for applications in a class of similar problems.[1]

THE PROBLEM

HASP/SIAP sleuths in the deep ocean. Using data from concealed hydrophone arrays, it must detect, localize, and ascertain the type of each ocean vessel within range. The presence and movements of submarines are of most interest, but there are strategic and tactical motives for monitoring all vessel types.

The program starts with digitized data from hydrophone arrays that monitor an ocean region from its periphery. The arrays have some directional resolution. Ideally each look direction produces a data channel with sound energy only from vessels near its axis, a spatial partition resembling spoke gaps on a bicycle wheel. In

---

[1] Much of the content of this paper appeared in more detail in an *AI Magazine* article. [8]

*© IEEE Acoustics, Speech and Signal Processing: Spring, 1984.

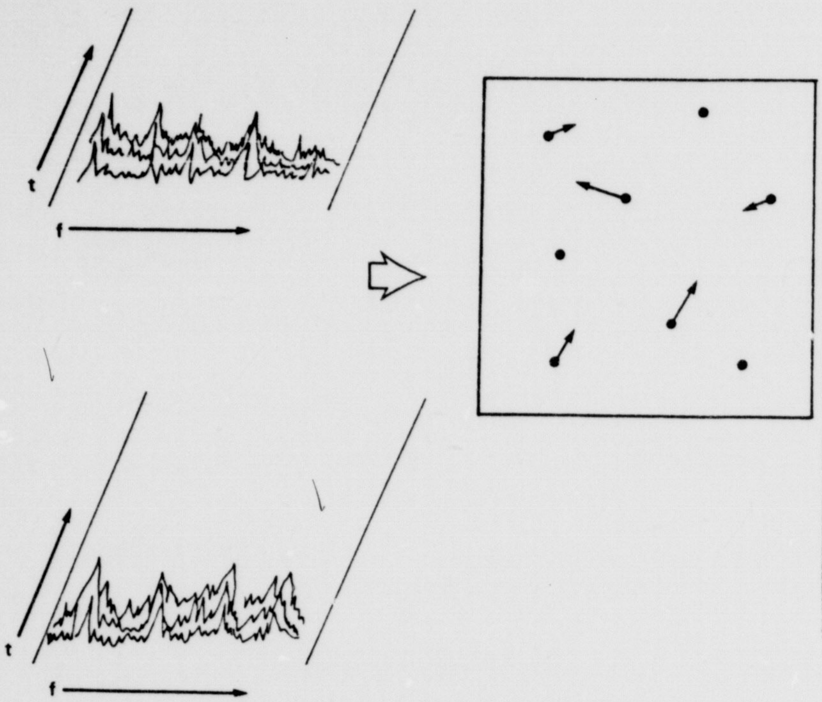

Figure 1: The Problem

practice, radiation from a single vessel may spread across several gaps, and many vessels may be located in any one gap, or in adjacent gaps, a situation that can produce a kaleidoscope of sound.

Rotating shafts and propellers, and reciprocating machinery on board a ship are major sources of the intercepted acoustic radiation. The **signature**, or sound spectrum, of a ship under steady operation contains persistent fundamental narrowband frequencies and certain of their harmonics. Imagine the ship's propeller saying "ahhhhh" on its way across the ocean. On an analyst's sonogram, this sound would appear as a collection of dark vertical lines against a fuzzy gray background.

Sonogram displays used by a sonar analyst are analog histories of the spectrum of received sound energy. New data on a channel are portrayed by varying the intensity of pixels on the display. Greater concentrations of energy at a given frequency are translated into higher intensities at corresponding horizontal positions. Synchronous horizontal sweeps, therefore, leave vertical lines on a display where persistent frequencies are present. (See Figure 2 for a simulated sonogram.) Starting, stopping, frequency shifting, and even subtle traces, are discernible to a trained eye.

Sonar analysts have been trained to recognize the sound signature traits of ships on their sonogram displays, and to classify

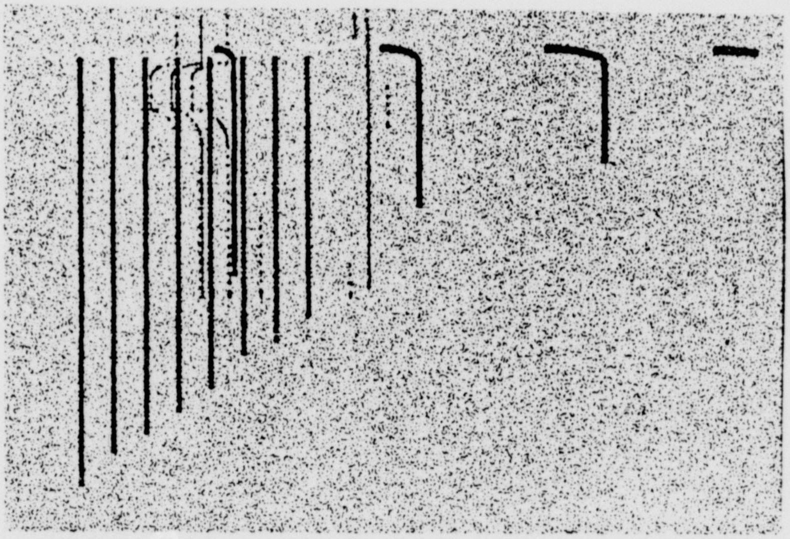

time/frequency

Figure 2: A Simulated Sonogram (with noise suppressed)

a signature into one of several important classes. If only one ship is present on a channel, the problem is essentially to match the measured signature (it may be imperfectly measured or partially present) to a collection of stored references for the best fit. Most computer-aided ship classification schemes have achieved some measure of success at this level.

When several ships radiate into the same array channels, the problem becomes more demanding. Highly skilled analysts use a collection of tricks in sonogram interpretation to disentangle the signatures for classification. Their procedures are not strictly separable for these tasks. That is, they do not disentangle signatures first without regard to classification information.[2] HASP/SIAP is unique among current machine-aided classifiers in imitating this non-separable approach.

There are other difficulties. The program has to overcome the problems of non-cooperative subjects in a noisy, complex medium. Ocean-going vessels typically move within fixed sea lanes, but storms and currents cause shifts in routes and operations. The background noise from distant ships is mixed with storm-induced and biological noises. Sound paths to the arrays vary with diurnal and seasonal cycles. Arrival of sound energy over several paths may suddenly shift to no arrivals at all,

---

[2] "Disentangle" means to separate correctly signature components of different vessels. This process is aided by contextual information about the plausible behavior of the sources on ships.

or arrivals only of portions of vessel radiation. Sound from one source can appear to arrive from many directions at once. Characteristics of the receivers can aiso cause sound from different bearings to mix, appearing to come from a single location.

ROLE OF REASONING

HASP/SIAP experience has shown that a combination of artificial intelligence with signal processing techniques not only work but can result in a powerful tool for applications in many areas. To date, a programming framework used in HASP/SIAP, called the blackboard framework, have been used successfully in problems of speech understanding [3], interpretation of protein crystallographic data [2], photographic image understanding [7], intelligence analysis [6], planning [5], battlefield situation understanding, and other information fusion problems.

All of these problems are characterized by the need to work with noisy and/or incomplete data. Ordinary statistical classification methods can organize signal data along a variety of dimensions, but "understanding" requires the use of information not present in the signals themselves, knowledge about the objects represented in the signal, and the semantic relationships among these objects. Too often the knowledge required for understanding the situation are diverse, uncertain, and incomplete. The primary problem solving strategy used under such circumstances is to either generate credible hypotheses, or to generate partial solutions and combine them using diverse knowledge available about the task. The credibility of the hypothesized solutions are then evaluated using other knowledge. To accomplish such a complex task, a program must be knowledgeable about and be able to reason about several things:

- assigning meaning to data consistent with, and derived from, situational context;
- discriminating important or critical pieces of information from irrelevant information;
- exploiting redundancy in the data;
- allocating signal processing resources wisely and providing data and task feedbacks to these resources; and
- detecting promising avenues that lead toward a solution, and recognizing when a plausible solution has been found.

In addition, the programs must be able to represent the partial solutions and uncertain hypotheses so that they can be used and modified by many different sources of knowledge.

In the following sections we describe the basic concepts of the programming framework used, followed by a short description of HASP/SIAP. These sections will show how HASP/SIAP addresses and implements many of the reasoning tasks described above.

TERMS AND CONCEPTS IN A BLACKBOARD FRAMEWORK

The *understanding* of sonograms often requires using information not present in the signals themselves. Major sources of information are reports from other arrays and intelligence reports. More general knowledge, like the characteristics of ships and common sea-lanes, also contributes significantly. Each such **source of knowledge** may at any time provide an inference which serves as a basis for another knowledge source to make yet another inference, and so on, until all relevant information has been used and appropriate inferences have been drawn.

Essential to the operation of the program is its **model** of the ocean scene. The model is a symbol-structure that is built and maintained by the program and contains what is known about the unfolding situation. The model thus provides a context for the ongoing analysis. More commonly known as the situation board to the analysts, the model is used as a reference for the interpretation of new information, assimilation of new events, and generation of expectations concerning future events. It is the program's *cognitive flywheel*.

The task of understanding the situation from the sonogram and other data is accomplished at various **levels of analysis**. These levels are exhibited in Figure 3. The most integrated, or the highest, level represents the situation board describing all the ships hypothesized with some confidence. The lowest level, that is, the level closest to the data, consists of connected line segments containing features derived from the signal data.

At each level, the **units of analysis** are the **hypothesis elements**. These are symbol-structures that describe what the available evidence indicate in terms that are meaningful at that particular level. Thus, on the Vessel level, in Figure 3, the descriptive properties that each Vessel element can have are: Vessel Class, Location, Current speed, Course, Destination. Each of the values of the properties has associated with it **weights**, an informal measure of confidence in the hypothesis. The example below shows a part of a hypothesis element on the Source level with different expressions of confidence.

SOURCE-1

TYPE	(Engine .5) (Shaft .3) (Propeller -.3)
LOCATION	((lat 34.2) (long 126.5) (error 9))
SUPPORT	(HARMONICS-4 HARMONICS-7)

[*SUPPORT attribute points to other hypothesis elements.*]

Links between the levels of analysis are built from **sources of knowledge**. A knowledge source (KS) is capable of putting forth the **inference** that some hypothesis elements present at its "input" level imply some particular hypothesis elements(s) at its "output" level. A source of knowledge contains not only the knowledge necessary for making its own specialized inferences, but also the knowledge necessary for checking the inferences made by other sources of knowledge. The inferences which draw together hypothesis elements at one level into a hypothesis element at a higher level (or which operate in the other direction) are represented symbolically as links between levels. The resulting network, rooted in the input data and integrated at the highest level into a descriptive model of the situation is called the **current best hypothesis** (CBH), or **the hypothesis** for short.

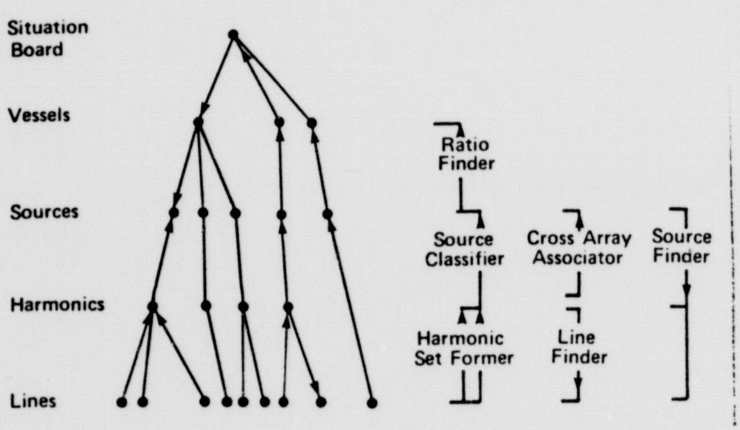

Figure 3: Levels of Analysis and Some Knowledge Sources

Each source of knowledge holds a considerable body of specialized information that an analyst would generally consider "ordinary". Sometimes this is relatively "hard" knowledge, or "textbook" knowledge. Also represented are the **heuristics**, that is, "rules of good guessing" an analyst develops through long experience. These judgmental rules are generally accompanied by estimates from human experts concerning the **weight** that each rule should carry in the analysis.

Each KS is composed of "pieces" of knowledge. By a piece of knowledge we mean a **production rule**, that is, an IF-THEN type of implication formula. [1] The "IF" side, also called the **situation** side, specifies a set of conditions or patterns for the rule's applicability. The "THEN" side, also called the **action** side, symbolizes the implications to be drawn (or various processing events to be caused) if the "IF" conditions are met. Following is a heuristic represented as a production rule:

If Source was lost due to fade-out in the near-past, and
Similar source started up in another frequency, and
Locations of the two sources are relatively close,

Then They are the same Source with confidence of .3.

[**Source** *refers to some noise producing objects, such as propellers and shafts on ships.*]

Hypothesis formation is an "opportunistic" process. Both **data-driven** and **model-driven** hypothesis formation techniques are used within the general **hypothesize-and-test** paradigm. The *knowledge of how to perform*, that is, how to use the available knowledge, is another kind of knowledge that the analysts possess. This type of knowledge is represented in the form of **control rules** to promote flexibility in specifying and modifying analysis strategies.

The **unit of processing activity** is the **event**. Events symbolize such things as "what inference was made", "what symbol-structure was modified", "what event is expected in the future", and so on. The basic control loop for these **event-driven programs** is one in which lists of events and a set of control rules are periodically scanned to determine the "next thing to do".

HASP/SIAP ORGANIZATION

The task of HASP/SIAP is to interpret continuous streams of data and to maintain a current situation model. The primary input data consists of 5-minute segments describing, in effect, a summary of observed signals at various frequencies (from each channel, for each array). The secondary inputs are information available from other arrays and a variety of reports routinely available to the analysts. The output of the program is a hierarchic data structure, the Current Best Hypothesis, containing the program's best explanation of the current input data considered in conjunction with previous analyses of earlier-received data (see Figure 6).

The data structure used in CBH is isomorphic to the human analysts' decomposition of the problem domain. The advantage of this isomorphic representation [9] is that the constraints and relations of the objects in the task domain are represented in the structure itself and need not be made explicit. Thus, for example, the fact that sound producing components are parts of a vessel is represented by the data structure itself. The same data structure is used as a medium through which integration of many diverse data and knowledge source is accomplished. The interpretation is accomplished by bidirectional, step-wise transformations between signals and symbolic description of objects at various levels of abstraction. Thus, for each level in the hierarchy, there must be at least one KS that can transform information on one level into information that is meaningful on some other level. For example, the following rule transforms a segment into a line by attaching it to a previously identified line:

If Characteristics of a new segment "match" an earlier line,
Source associated with the line is not currently heard, and
Source had disappeared less than 30 minutes ago,

then Source is being picked up again with confidence .5, and
Segment is assigned to the old line.

[*Note that if the "match" algorithm or heuristic is not very good, the weight of the confidence can be reduced.*]

The word "transformation" is used loosely to mean a shift from one representation of an acoustic source(signal segments, for example) to another (propeller, for example) using any formal or informal knowledge. The major pieces of the HASP/SIAP program is shown in Figure 4.

Kinds of Knowledge Represented

There are several kinds of knowledge used in HASP/SIAP, each represented in a form that seems the most appropriate.

Knowledge about the environment: The program must know about common shipping lanes, locations of arrays and their relation to map coordinates and known maneuver areas. These knowledge are represented in procedures that compute the

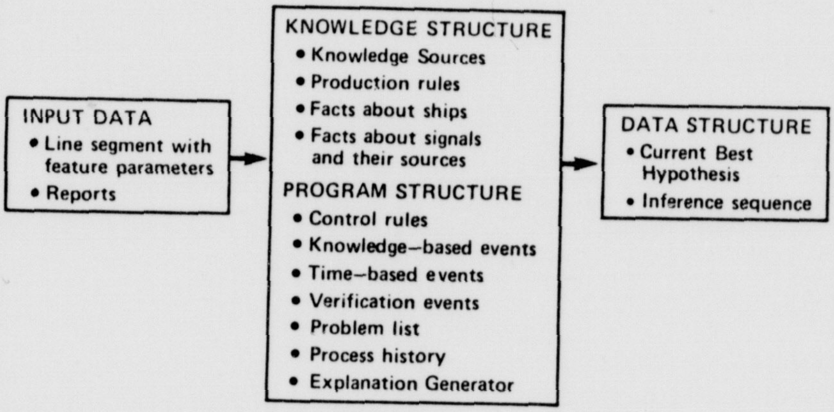

Figure 4: A General Program Structure

information.

Knowledge about vessels: All the known characteristics about vessel types, component parts and their acoustic signatures, range of speed, home base, etc., are stored in frame-like data objects. [10] Each object is a symbolic model of some aspect of the vessels the program knows about. This knowledge base is used by rules whenever a specific class of vessels is being analyzed. When a vessel class is hypothesized from a piece of data, the knowledge base also provides information on other pieces of data the program should look for to support the hypothesis. The use of this model-driven hypothesis formation technique reduces the amount of computation by directing other KSs to look for specific data.

Interpretation knowledge: All knowledge about transforming information on one level of the CBH to another level (i.e. assign meaning to data on one level in terms of another level) is represented as sets of production rules. The rules in a KS usually generate inferences between adjacent levels. However, some of the most powerful KSs generate inferences spanning several levels. For example, a line with particular characteristics may immediately suggest a vessel class, bypassing analysis for harmonic set formation and source identification. This type of knowledge is very situation-specific. Many of these specialized, detailed knowledge used by the analysts were elicited from human experts and were put in the program.

Knowledge about how to use knowledge: Since the problem solving strategy is opportunistic, the program must know when an opportunity for further interpretation has arisen and how best to capitalize on the situation. In HASP/SIAP this type of knowledge is also explicitly represented and used.

How the Knowledge is Organized and Used

How well an expert system performs depends both on the competence of the KSs and on the appropriate use of these KSs. When and how a KS is used depends on its level of expertise and its relevancy. Relevancy in turn depends on the state of the CBH. The control mechanism for KS selection needs to be sensitive to, and be able to adjust to, the numerous possible solution states which arise during interpretation. Given this viewpoint, what is commonly called a "control monitor" can be viewed as another type of domain-dependent knowledge, albeit a high level one.

In a hierarchically organized control structure problem solving activities decompose into a hierarchy of knowledge needed to solve problems. On the lowest level is a set of knowledge sources whose charter is to hypothesize plausible interpretations. We refer to KSs on this level as **Specialists.** At the next level there are knowledge sources that know when to use the various Specialists. The highest level knowledge source analyzes the current solution state to determine what information to analyze next. This activity is called *focussing-of-attention*, and it plays a major role in the efficient use of resources in the hypothesis formation process.

The KS hierarchy should be clearly distinguished from the hierarchy of analysis levels. The hypothesis hierarchy represents an *a priori* plan for the solution determined by a *natural* decomposition of the task domain. The KS hierarchy, on the other hand, represents a plan for organizing the problem-solving activities needed to fill in information on the hypothesis hierarchy. Compare Figures 3 and 5.

KSs on the Specialist Level: Each Specialist has the task of creating or modifying hypotheses, evaluating hypotheses generated by other Specialists, and cataloging missing evidences that are essential for further analysis. Although a KS has access to the entire hypothesis structure, it normally "understands" only the descriptions contained in two levels, its input level and its output level. Some examples Specialist rules shown below.

Hypothesis-Generation:

If A Harmonic set "match" another set on another channel,
then Both sets are from the same source with .6 confidence.

Inference-Evaluation:

If Source belongs to Vessels of class Cherry or Iris, and
Harmonics of Source have been stable for a while,
then Increase the confidence of Cherry and Iris by .3.

Problem-Cataloging or Expectation generation:

If Report exists for a vessel class Rose to be in the area, and
A Source likely to be from Rose has been detected,
then Expect to find other Sources associated with Rose class.

KSs at the KS-Activation Level: A KS on this level has the task of invoking the appropriate Specialists given the processing strategy being employed. For example, a KS charged with calling the appropriate KSs within a model-driven strategy has a different task than one charged with a data-driven strategy. KS-activators can also implement different priorities; for example, the fastest-first-KS or the most-accurate-KS first. HASP/SIAP has two KS-activators, the Event-driver and the Expectation-driver. The Event-driver chooses a specialist-KS with the highest degree of specialization, and assumed accuracy, if there is more than one Specialist available to process a focussed event. The Expectation-driver processes items on the Problem-list on the basis of how critical the needed evidence is to the emerging situation hypothesis.

KS on the Strategy Level: The Strategy KS reflects a human analyst's problem-solving strategy. Its expertise consists of determining how accurate the current hypothesis is and in deciding what activity will have the highest impact on the CBH. It has a variety of possible activities to choose from. "Is it time to check for specific data?" "Has anything happened since the last processing cycle that might invalidate the hypothesis?" "Has an expected event occurred?" "What is the best region in the CBH to work on next; i.e., what is the region of minimal ambiguity?"

In HASP/SIAP there are no formal mechanisms to measure the differences between the current best hypothesis and the "right answer". For one thing, the system presupposes a heuristic that "Ships can't just disappear; let's wait and see if the next set of data can shed some light". A strategy level KS detects when the solution hypothesis is "on the right track" by a periodic use of heuristic criteria. For example, a consistent inability to verify expectation-based hypothesis elements may signal an error in the hypothesis. A more general indication of ineffective hypothesis formation appears as a consistent generation of hypotheses whose confidence values are below a threshold value; and which therefore indicates that the analysis is "bogged down".

The Basic Execution Cycle

The KS execution filters down from the highest level KS to the lowest level specialist KSs, and back to the top. The cycle consists of (1) focussing attention on either pending timedependent activities, on verification of a hypothesis, or on one of the recently modified hypothesis elements (Strategy KS); (2) selecting the appropriate KSs for the attended event (KS Activation); and (3) executing the selected KSs (Specialist KSs). The Specialist KSs will generate new events -- a new CBH or expectations. The cycle is then repeated. Figure 5 shows the basic flow of information and control.

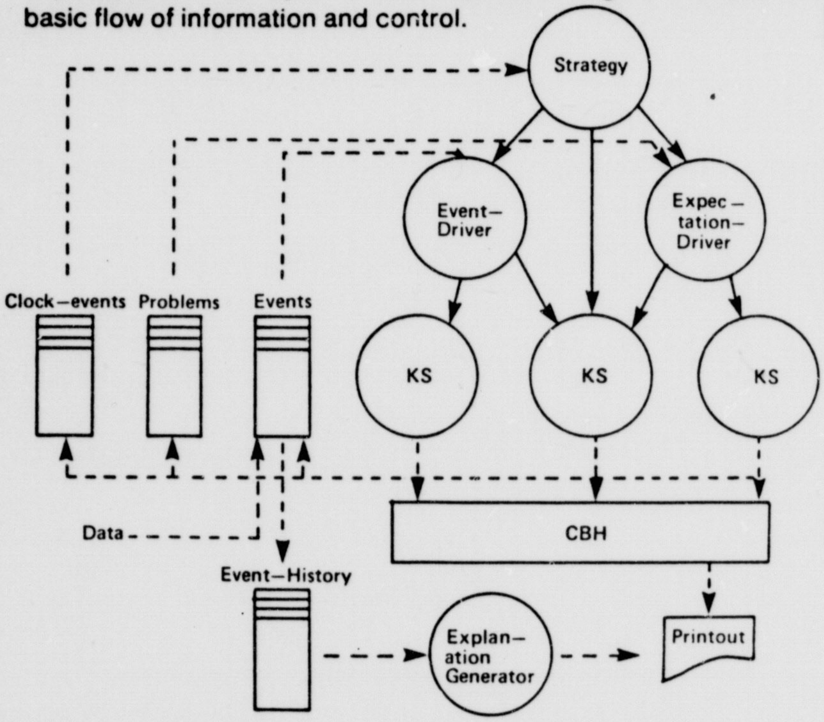

Figure 5: Information and Control Flow

The Current Best Hypothesis

As KSs are applied to a stream of input data, a solution hypothesis emerges. As mentioned earlier, the hypothesis structure represents the up-to-date "best hypothesis" at any given time for the data available up to that time. The structure of the CBH is a linked network of hypothesis elements, each representing a meaningful aggregation of lower level hypothesis elements. A link between any two hypothesis elements represents a result of some action by a KS. A link has associated with it directional properties: (1) A link that goes from a more abstract to a less abstract level of the hypothesis is referred to as an **expectation-link**. The node at the end of an expectation-link is a model-based hypothesis element, and the link represents *support from above* (i.e. the reason for proposing the hypothesis element is to be found at the higher level). (2) A link which goes in the opposite direction, from lower levels of abstraction to higher, is referred to as a **reduction-link**. The node at the end of a reduction-link is a data-based hypothesis element, and the link represents *support from* *below* (i.e. the reason for proposing the hypothesis element is to be found at a lower level).

THE CURRENT BEST HYPOTHESIS AT TIME 20455

Vessel-1

Class (OR (Cherry 8.4) (Iris 6.9) Tulip 6.2) (Poppy 4.8) 20455 . . .)
Location ((Lat 37.3) (Long 123.1) (Error 37))
Speed 15.7
Course 135.9
Sources (AND Source-1 Source-5)

Source-1

Type (OR (Cherry Propeller 5.5) (Poppy Shaft 2.5)
(Poppy Propeller 2.0) (Cherry Shaft 2.5) 20455 . . .)
Dependency Unknown
Regain (20230)
Harmonics (Harmonic-1)

Harmonic-1

Fundamental (224.5 20520)
Evolution (fade-in 20230 fade-out 20210 . . .)
Lines (AND Line-1 Line-2 Line-6 Line-12)

Source-5

Type (OR (Cherry Shaft 6.0) (Poppy Shaft 4.0)
(Iris Propeller 5.0) (Tulip Propeller 2.0) 20455)
Dependency 6
Harmonics (Harmonic-5)

Harmonic-5

Fundamental (162.4 20455)
Evolution (fade-in 20455)
Lines (AND Line-25)

ASSIMILATION (RATIO Source-1 Source-5 .5) 20455)

Problems List

(EXPECT Vessel-1 (SUPPORT Cherry) (Dependency Propeller 5))
(EXPECT Vessel-1 (PRED.LOC (Lat 37.2) (Long 123.) (Error 41.3))
(REPORT REPORT-GEN Rose (Signature (Engine 30 166.7) . . . .))

Figure 6: A Part of a Current Best Hypothesis

Representation of Continuous Interpretation: As mentioned earlier, the primary input data is a series of 5-minute segments of continuous data. These continuous segments are reduced by signal-processing programs to a symbolic segment description containing its frequency, bandwidth, intensity. These segments must then be integrated into Line level hypothesis elements. In addition, lines must be integrated into harmonic sets; harmonic sets must be associated with acoustic sources; sources must be attributed to vessels. The CBH serves as the historical context by which the integration can occur. For this purpose, the CBH must keep track of the times when specific events occurred in addition to maintaining a network of relationships between Lines, Sources, and Vessels. Time markers are placed with hypothesized values;

new time markers are added only when a change is made to the values. In the example below, there were no change to the hypothesized Source type for two hours even though its weight may have changed during that time.

SOURCE-1
Type [(Cherry Propeller 7.5) 10650 (OR (Cherry Propeller 6.0)
(Poppy Propeller 2.5)) 10450]

Explanation of the Line-of-Reasoning:HASP/SIAP has the capability of producing an explanation based on the line-of-reasoning (i.e. processing history) it followed in hypothesizing the situation. However, the CBH contains enough information about evidential support and evolution of the various elements of the hypothesis that it serves as a useful explanation by itself.

CONCLUSION

In signal processing applications, involving large amounts of data with poor signal-to-noise ratio, it is possible to reduce computation costs by several orders-of-magnitude by the use of knowledge-based reasoning rather than brute-force statistical methods. *It makes little sense to use enormous amounts of expensive computation to tease a little signal out of much noise, when most of the understanding can be readily inferred from the symbolic knowledge surrounding the situation.*

Of the reasoning requirements mentioned in the early part of this paper, HASP/SIAP did not address the problem of feedback to the signal processing front-end (only to symbolic processing parts), and the allocation of signal collection and processing resources. Some possible approaches to these problems are discussed here.

There are two kinds of feedback that the CBH can provide signal processors. First, special purpose detection and parameter measurement algorithms depend on higher level information -- the CBH has the necessary information. Second, threshold values for the front-end need to be adjusted according to the current needs and expectations. In both cases, the processing of information over a period of time leads to modifications of parameters in the front-end for more accurate subsequent processing.

In the undersea surveillance problem, a variety of signal processing techniques can be brought to bear in the front-end. Some of these fall in the category of "standard" algorithms that are more or less always used. But others are specialized algorithms [illegible]t cannot be used for all channels at all times because [illegible]eir cost. This resource allocation problem is knowing *when* to invoke the special signal processors. The approach to its solution lies in rules that can recognize when the context will permit the special processor to resolve important ambiguities. In HASP/SIAP only a rudimentary capability for this process was used, but its value was conclusively demonstrated.

There is an additional cost saving possible. Sensor bandwidth and sensitivity is expensive. From a symbolic model it is possible to generate a set of signal expectations whose emergence in the data would make a difference to the verification of the ongoing model. Sensor parameters can then be "tuned" to the expected signals and signal directions; not every signal in every direction needs to be searched for.

The HASP/SIAP experience indicates that some good AI can cover many signal processing inadequacies, as well as direct the employment of appropriate signal processing algorithms. The intelligent combination of AI and signal processing views the signal processing component as another knowledge source, with rules on how best to employ algorithms and how to assign meaning to their output.

References

[1] Davis, R. and J. King.
An Overview of Production Systems.
In E.W. Elcock and D. Michie (editor), *Machine Intelligence 8: Machine Representation of Knowledge*, . John Wiley, New York, 1977.

[2] Engelmore, R. S. and A. Terry.
Structure and Function of the CRYSALIS System.
Proc. of IJCAI 6 :251 - 256, 1979.

[3] Erman, L.D., F. Hayes-Roth, V. R. Lesser, D. Raj Reddy.
The Hearsay-II Speech Understanding System: Integrating Knowledge to Resolve Uncertainty.
ACM Computing Survey 12:213 - 253, June, 1980.

[4] Feigenbaum, Edward A.
The Art of Artificial Intelligence: I. Themes and Case Studies of Knowledge Engineering.
Proc. of IJCAI 5 :1014 - 1029, 1977.

[5] Hayes-Roth B. and F. Hayes-Roth.
Modelling Planning as an Incremental, Opportunistic Process.
Proc. of IJCAI 6 :375-383, 1979.

[6] Lenat, D. B., A. Clarkson and G. Kiremidjian.
An Expert System for Indicators & Warning Analysis.
Proc. of IJCAI 8 :259 - 262, 1983.

[7] Nagao, M., T. Matsuyama, and H. Mori.
Structured Analysis of Complex Photographs.
Proc. of IJCAI 6 :610 - 616, 1979.

[8] Nii, H. P., E. A. Feignebaum, J. J. Anton, and A. J. Rockmore.
Signal-to-Symbol Transformation: HASP/SIAP Case Study.
AI Magazine, Vol. 3, No. 2 :23 - 35, 1982.

[9] Pentland, Alex P. and Martin A. Fischler.
A More Rational Viwe of Logic.
AI Magazine, Vol. 4, No. 4 , 1982.

[10] Stefik, Mark.
An Examination of a Frame-Structured Representation System.
Proc. of IJCAI-6 2:845 - 852, 1979.