N93-32152

A Decision-Theoretic Approach to the Display of Information for Time-Critical Decisions: The Vista Project

Eric Horvitz* **Corinne Ruokangas** **Sampath Srinivas**
Palo Alto Laboratory
Rockwell International Science Center
444 High Street
Palo Alto, California 94301
and
Matthew Barry
Propulsion Systems Section
NASA Johnson Space Center
Houston, Texas

Abstract

We describe a collaborative research and development effort between the Palo Alto Laboratory of the Rockwell Science Center, Rockwell Space Operations Company, and the Propulsion Systems Section of NASA Johnson Space Center to design computational tools that can manage the complexity of information displayed to human operators in high-stakes, time-critical decision contexts. We shall review an application from NASA Mission Control and describe how we integrated a probabilistic diagnostic model, and a time-dependent utility model, with techniques for managing the complexity of computer displays. The, we shall describe the behavior of VPROP, a system constructed to demonstrate promising display-management techniques. Finally, we shall describe our current research directions on the Vista II follow-on project.

1. Introduction

The Vista project was established to develop computational techniques for reducing the cognitive load of human operators that are responsible for monitoring complex systems. Decision makers under pressure to take swift action cannot afford to sift through large quantities of potentially relevant information before making a decision. Fundamental limitations in the abilities of people to process information explain why the quality of a system operator's decisions can degrade as the complexity of relevant data increases, and as the time available for a response decreases. Cognitive psychologists have found that humans cannot retain more than five to nine distinct concepts or "chunks" of information simultaneously (Miller 1956). This surprising cognitive limitation was demonstrated in a classic study by Miller, and has been confirmed repeatedly by experimental psychologists. The capacity of decision makers to consider important influences on a decision may be reduced even further if fast action is demanded in frenetic crisis situations; one cognitive-psychology study showed that people cannot retain and reason simultaneously about more than two concepts in environments filled with distractions (Waugh and Norman 1965).

There has been previous investigation of methods for managing the complexity of computational results to decrease the cognitive burden of computer users. Key concepts employed in the Vista project are similar to techniques developed in related work on the

*Also at Knowledge Systems Laboratory, MSOB X215, Departments of Computer Science and Medicine, Stanford University, Stanford, CA 94301 control of computation and decision making under bounded resources (Horvitz, et al. 1989, Horvitz 1990, Horvitz and Rutledge 1991). In the related work on the decision-theoretic control of computation, investigators have examined decisions about the tradeoff between precision and timeliness of generating a computational result (such as a recommendation for action in the world). The analogous techniques for controlling displays have centered on decisions about the tradeoff between the completeness and the complexity of information displays and computer-based explanations (Horvitz, et al. 1986; Horvitz, et al. 1989). In some of this work, multiattribute utility has been employed to control the complexity of displays (Horvitz 1987; Mclaughlin 1987) and for queuing and prioritizing the results of diagnostic reasoning (Breese, et al. 1991).

We have worked to integrate a set of flexible display strategies with a probabilistic reasoning module that performs diagnosis under uncertainty of propulsion systems of the Space Shuttle. The diagnostic-reasoning component of the approach continues to reason about the likelihood of alternative disorders in a system based on a continuous stream of sensor information. The Vista team consists of a close collaboration between members of the Palo Alto Laboratory group (Eric Horvitz (PI), Corinne Ruokangas, and Sampath Srinivas) and a propulsion systems specialist at Houston Mission Control (Mathew Barry). In our approach, we employ model-based diagnostic reasoning and flexible displays for reconfiguring the information displayed so that the most relevant information is presented in different context.

We tested our ideas about the model-based control of displays by building and validating a display manager named VPROP. The techniques explored in VPROP have application in many domains where decision makers with limited time must evaluate large amounts of information under time pressure to identify critical data.

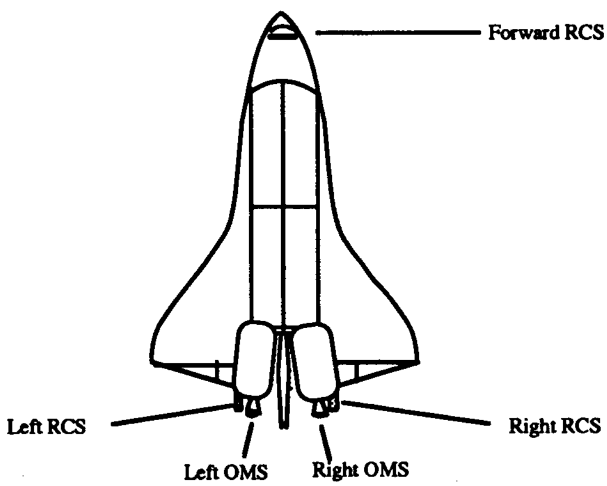

Figure 1. The position of the two orbital maneuvering system (OMS) engines and the three sets of reaction control system (RCS) thrusters on the Space Shuttle.

2. Application Area

We examined problems with the complexity of information displays used by ground controllers at Space Shuttle Mission Control in Houston. In particular, we worked with experienced flight controllers who manage Space Shuttle propulsion systems. The Propulsion Team is responsible for monitoring two different space-based thruster systems: the Orbital Maneuvering System (OMS) and the Reaction Control System (RCS). The large right and left OMS engines are fired for such critical maneuvers as orbital insertion and orbit circularization. The smaller suites of RCS thrusters are used for translation in space, for such tasks as maneuvering near another space object, as well as for the continual

The status quo computer displays for monitoring the propulsion systems consist of a cluttered, static main display of the two OMS engines, and the three banks of RCS engines. This primary propulsion-systems display is pictured in Figure 2. When ground controllers are concerned about one of the systems, alternate screens are typically requested which relay trend information about engine consumables by displaying a large table of numbers. An example of such a secondary display is pictured in Figure 3. The large amount and complex display of information that must be scanned in crisis situations can become burdensome in situations that demand quick decision making, especially for new people on the propulsion team.

3. Construction of a Probabilistic Model

The first task in developing a model-based display manager was the construction of a probabilistic diagnostic model for the shuttle propulsion-system engines. To build the diagnostic reasoner of the display manager, our team worked with propulsion experts at Mission Control to build a causal model of the shuttle's propulsion subsystems. Figure 4 displays a basic schematic of the OMS engine. In the engine, a tank of helium puts pressure on tanks of oxidizer and fuel. To fire an OMS engine, a bipropellant valve is opened which allows the fuel and oxidizer to mix and combust to provide thrust. In addition to the basic flows, ground controllers must also consider the status of a set of values between various tanks, and crossover lines that allow fuel to be shared by different engine systems. There are suites of temperature and pressure sensors at critical locations in the system. Shuttle telemetry includes information from these sensors.

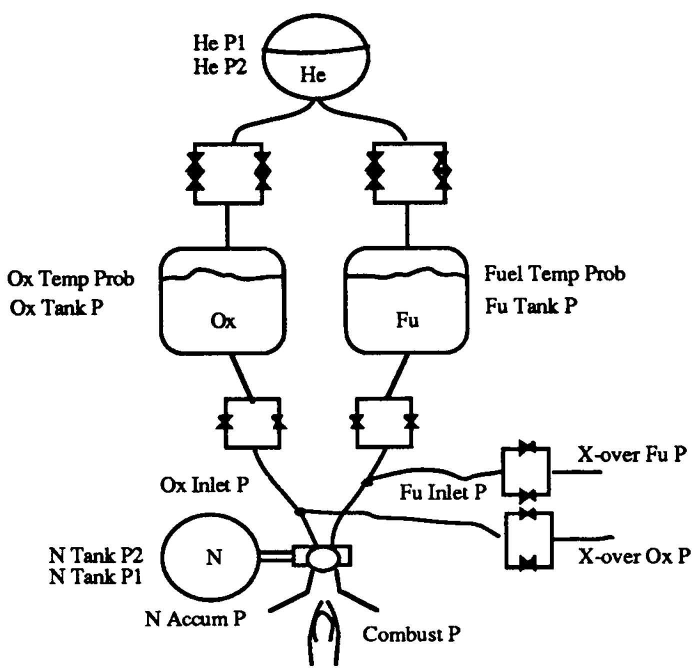

Figure 4. A schematic of an OMS engine.

We constructed a probabilistic causal network, called a belief network, to model the uncertain relationships among components of the system. Belief networks serve as the core of an increasing number of probabilistic reasoning systems. (See (Pearl, 1989) and (Horvitz, 1988) for reviews of probabilistic and decision-theoretic reasoning.) In practice, a domain expert, working with a computer engineer, structures and assesses the probabilistic relationships in a belief network. In this case, Matthew Barry, a propulsion specialist at NASA Mission Control served as expert. The belief network representation allows an expert to structure relationships about a system qualitatively, and, after, to quantify those relationships with conditional probabilities.

Belief networks serve as the kernel of diagnostic reasoning systems. Technically, a belief network is a directed acyclic graph (DAG) containing nodes, representing propositions (hypotheses, intermediate states, and observations), and arcs representing probabilistic dependencies among nodes. Nodes are associated with a set of mutually exclusive and exhaustive values that represent alternative possible states of a proposition (*e.g.*, true, false). Directed arcs capture knowledge that the value of a parent node can affect the probability distribution over the values of children nodes.

A variety of belief-network algorithms have been formulated. The algorithms propagate information through the network to compute the probability distributions over values of each node, given the observation of one or more values of a subset of nodes that represent sensors or observations. The belief network allows us to represent the notion that there may be multiple causes of problems, and the related notion that a sensor may fail, and thus provide erroneous information about the system being monitored.

Figure 5 represents an early version of a belief network constructed to model a Shuttle OMS engine. Following arcs through a belief network often tells us a story about uncertain relationships in a system. For example, the value of HELIUM PRESSURE affects the pressure readings (He P1, He P2) reported by the two independent pressure sensors on an OMS helium tank. However, the readings can also be affected, with uncertainty, by the errors in the sensor mechanisms themselves. A user with experience in sensor failures can encode his belief about the relative rate of failure of alternative critical sensors in a system.

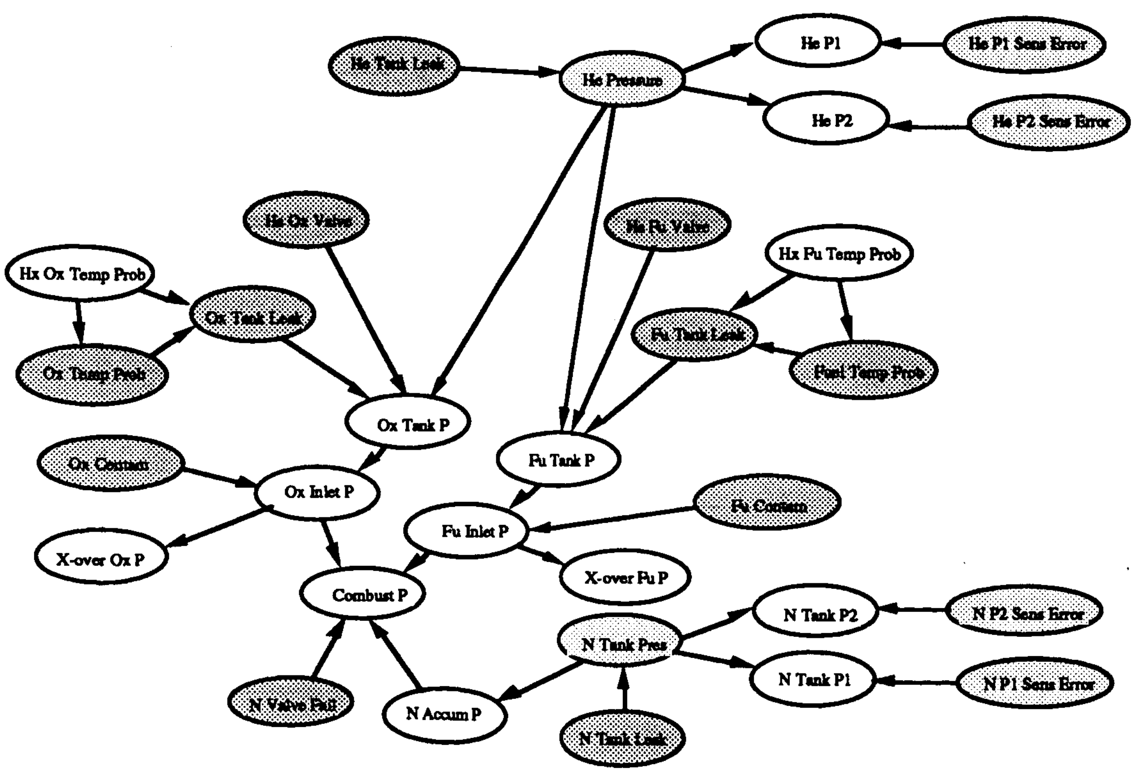

Figure 5. A causal probabilistic network for the OMS engine.

4. Display Management in VPROP

The diagnostic reasoning system was integrated with a set of display complexity-management techniques to yield a dynamic display system. Key features of VPROP's display include (1) the use of context-sensitive templates for configuring data for different phases of the mission, (2) a static screen layout of panels of information on key systems, (3) the problem-specific telescoping of system information along a simplicity--complexity vector, (4) the summarization of large quantities of information with simple status indicators, (5) the manual access of any data through an iconic information palette, and (6) the use of probabilistic and decision-theoretic reasoning to dynamically prioritize menus for accessing critical information for decision making.

4.1 Techniques for Managing Display Complexity

Figure 6 highlights these flexible display techniques. In this case, panels of information are configured for a critical OMS firing. For this context, system-specific panels containing information on the left, forward, and right RCS systems are reduced into brief summaries and information panels on the left and right OMS are expanded. By moving the cursor and clicking on any system-specific panel, a menu can be accessed that allows the detail, and concomitant area occupied by information about a particular system, to be increased or decreased beyond the standard configuration for a context.

We can simplify information about a Space Shuttle system as we shrink a panel by modulating the abstraction or the completeness of information (Horvitz, 1987). For VPROP, the expert defined dimensions of complexity which described how the details about a propulsion system should change, as a panel is allowed to occupy more screen area. That is, an expert decision maker defined how the content and granularity of information in each panel changes as that panel is enlarged or diminished.

A *display palette* appears at the lower left hand corner of the display. The display palette is laid out in a configuration that iconically reflects the main screen layout of the system-specific information panels. By clicking on soft buttons in the palette for each system, the main panels of information are incremented in detail and then reduced, as if alternate levels of complexity were stored in a circular queue. The palette doubles as a systems-status overview and summary. If telemetry about a system is within normal bounds, the corresponding display-palette soft button appears green. If there is a problem with a system, the corresponding display-palette button flashes red.

The usefulness of employing abstraction to simplify the system-specific panels and to alert an operator about system status is supported in part by findings of cognitive psychologists that people employ abstraction of items into classes to manage the complexity of reasoning about complex problems (Mesarovic 1970, Simon 1973).

4.2 Integrating Display-Complexity Tools with Diagnostic Reasoning

A belief network, that we constructed and assessed with the assistance of an experienced ground controller, is employed in VPROP to continue to monitor Space Shuttle telemetry on propulsion systems. Figure 6 highlights VPROP's reaction when a problem with the left OMS has been noted. In this case, the probability that the left OMS is normal has dropped below a threshold value. The left-OMS button in the display palette turns red and the system-specific information panel for the the left OMS automatically expands in size and complexity.

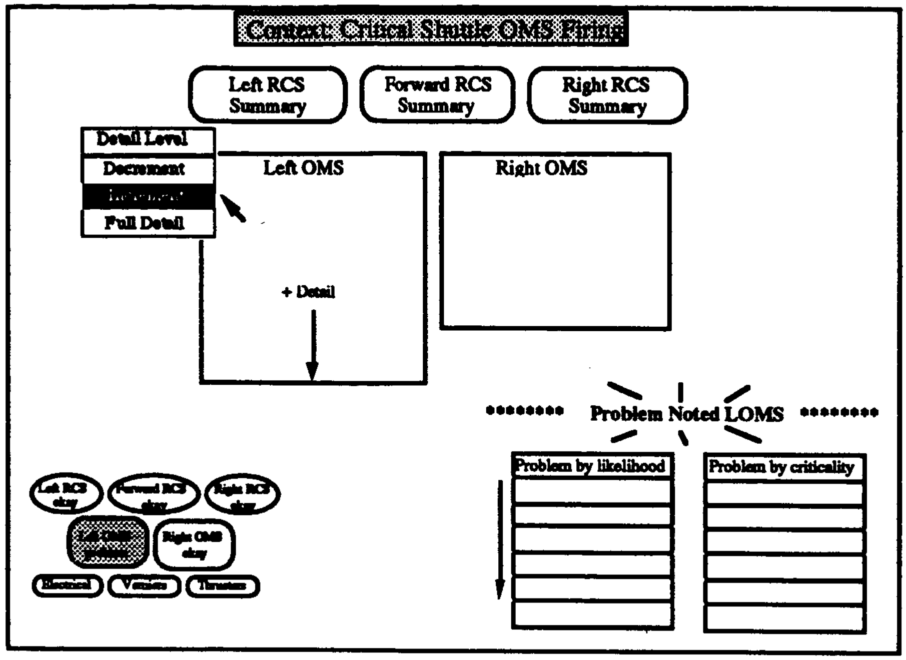

Figure 6. Key components of VPROP's flexible display.

Now, the diagnostics on the left-OMS are displayed, as portrayed by the two lists, displayed at the lower right-hand-corner of Figure 6. A summary of the problem is stated and a list of possible disorders computed by the belief network are listed by likelihood. In another column, the disorders are reordered into a set of priorities, in terms of the *expected time-criticality* of the problems. The expected time-criticality is computed by weighting a relative cost of delay for each possible disorder, assessed from the ground controller, by the likelihood computed for that disorder.

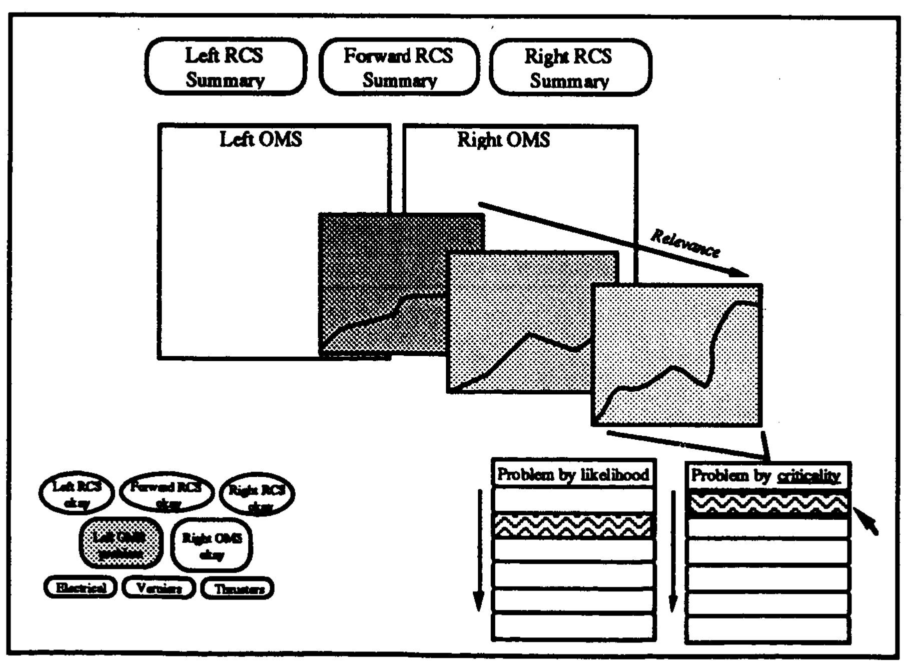

Figure 7. Accessing trend information from a prioritized list of possible faults.

Information about trends is often used by Shuttle-propulsion ground controllers to confirm disorders with the propulsion systems. VPROP allows ground controllers to access trend graphs, by positioning a cursor and clicking a mouse button on disorders appearing on the likelihood or criticality lists. Figure 7 highlights how VPROP displays the trend information in a predefined order of relevance.

The integrated Vista methodology for managing the complexity of displays, while providing diagnostic information about the likelihood of alternative faults and recommendations about time-dependent priorities, is highlighted in Figure 8. As indicated in this figure, a goal of the Vista project is to field systems that analyze real-time shuttle telemetry with probablistic causal models and that consider the most valuable configuration of the display in terms of the likelihood and time-dependent losses associated with different faults.

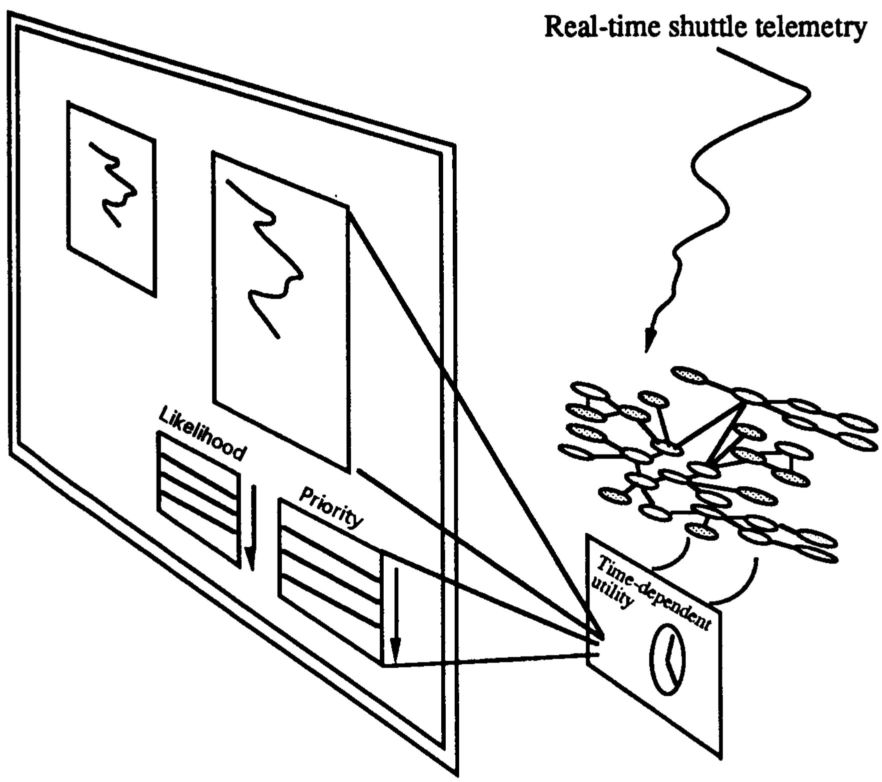

Figure 8. Key components of the Vista approach. A belief network is used to consider the diagnostic relevance of shuttle telemetry. A model of time-dependent utility is employed to prioritize alternative possible faults, given the likelihood of the faults, and to control the size and amount of detail dedicated to a panel describing the system in which the possible faults may be occuring.

A bitmap of the current VPROP screen is displayed in Figure 9. In this case, troubling telemetry has been detected by the belief network. A diagnostic screen appears, showing that, given the current sensor information, there is a high probability of engine failure. Relevant trend information has been accessed.

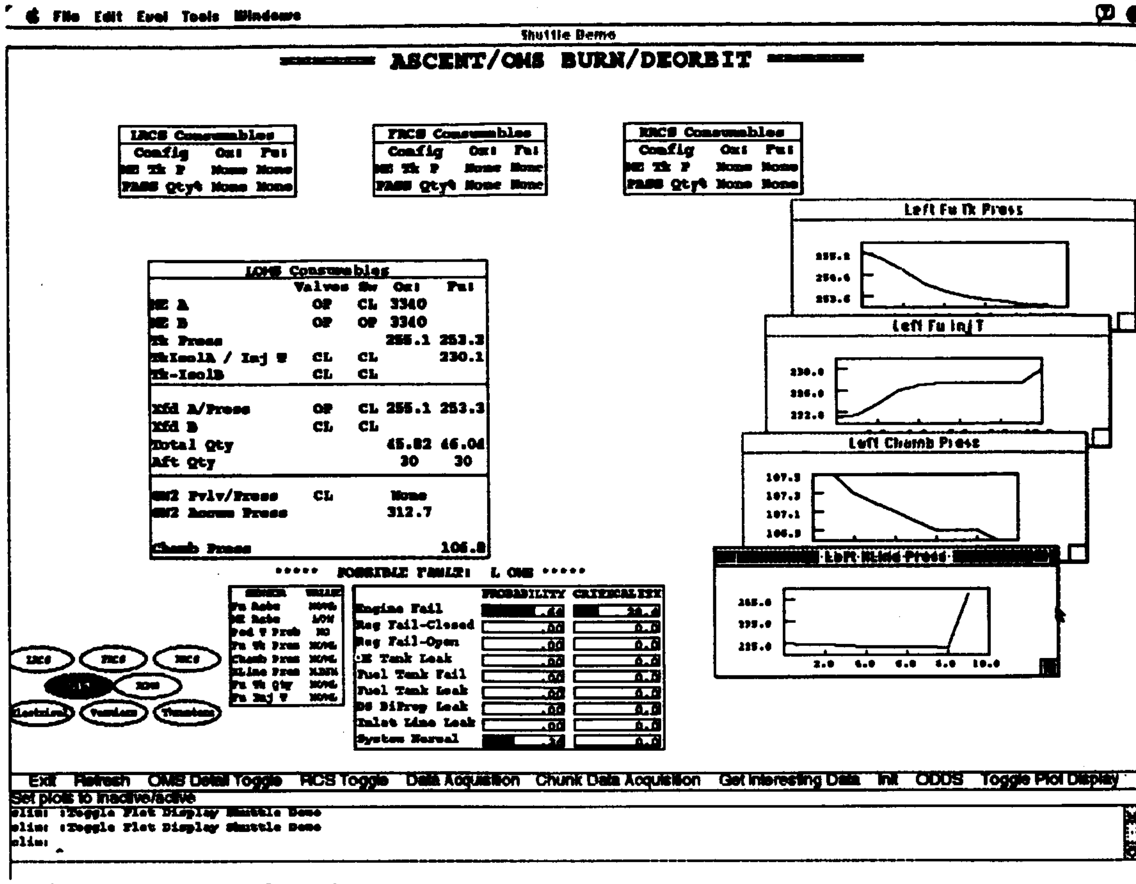

Figure 9. A screen from the VPROP system displaying how a problem with the left OMS is handled. In this case, there is a high probability of an engine failure. Relevant trend information is accessed.

5. Conclusions and Summary

The VPROP system has been reviewed by the Propulsion Team at NASA Mission Control in Houston. The system reviews were met with enthusiasm about the approach and future extensions to the methodology. Discussions and research among the Palo Alto Laboratory, the Rockwell Space Operations Company, and NASA Johnson Space Center are continuing on future projects to extend and apply Vista technology for monitoring and controlling other Space Shuttle systems as well as systems on the planned Space Station. The ability to consider alternative hypotheses under uncertainty promises to be even more crucial for monitoring systems on the space station; in comparison to the Shuttle, current plans call for the Space Station to have relatively few sensors on complex systems. This situation will make it more difficult to narrow a problem down to a single fault, given telemetry information.

Recently, we initiated the Vista II project, a Vista follow-on effort to develop a version of VPROP to be deployed on Unix workstations. The project centers on building an effective integrated display-management and diagnostic reasoning architecture, as well as the development of a new display management language for increasing the efficiency with which we can construct and custom-tailor reasoning systems for different ground-control ans space-based monitoring functions.

The integration of flexible display techniques and model-based systems with the ability to reason under uncertainty promises to provide a valuable framework for managing the cognitive load of human operators. We hope that the combination of intelligent diagnostic and display systems will allow us to keep pace with the growing complexity of machines and tasks. We foresee that intelligent reasoning systems will be important in helping to maintain--or even to reduce--the cognitive burden on human operators charged with monitoring or managing systems of increasing complexity. Indeed, the design of computational tools for managing the complexity of information about complex systems may one day be viewed as an intrinsic component of the design of the systems.

Acknowledgments

We thank John Breese, Max Henrion, Jim Martin, Michael Shaff, Philip Stubblefield, and Cynthia Wells for providing valuable feedback and assistance on Project Vista.

References

(Breese, et. al 1989)
J.S. Breese, S. Srinivas, and R.A. Fiske. A Utility-Based Pilot Vehicle Interface. *Proceedings of DARPA Symposium on Associate Technology,* pages 259-266. George Mason University, Fairfax, VA, June, 1991.

(Horvitz 1987)
E.J. Horvitz. *A Multiattribute Utility Approach to Inference Understandability and Explanation.* Technical Report KSL-87-28, Knowledge Systems Laboratory, Departments of Computer Science and Medicine, Stanford University. March, 1987.

(Horvitz 1990)
E.J. Horvitz. *Computation and Action Under Bounded Resources.* Ph.D. Dissertation, Departments of Computer Science and Medicine, Stanford University, December, 1990.

(Horvitz, et al. 1986)
E.J. Horvitz, D.E. Heckerman, B.N. Nathwani, and L.M. Fagan. The use of a heuristic problem-solving hierarchy to facilitate the explanation of hypothesis-directed reasoning. In *Proceedings of Medinfo,* Washington, DC, pages 27-31. North-Holland, New York, October 1986.

(Horvitz, et al. 1988)
E.J. Horvitz, J. Breese, and M. Henrion. Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning, Special Issue on Uncertain Reasoning in Artificial Intelligence.* 2:247-302.

(Horvitz, et al. 1989)
E.J. Horvitz, D.E. Heckerman, K. Ng, and B.N. Nathwani. Heuristic abstraction in the decision-theoretic Pathfinder system. In *Proceedings of the Thirteenth Symposium on Computer Applications in Medical Care,* Washington, DC, pages 178-182. IEEE Computer Society Press, Silver Spring, MD, 1989.

(Horvitz, et al. 1989)
E.J. Horvitz, G.F. Cooper, and D.E. Heckerman. Reflection and action under scarce resources: Theoretical principles and empirical study. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence,* Detroit, MI, pages 1121-1127. Morgan Kaufmann, San Mateo, CA, August 1989.

(Horvitz and Rutledge 1991)
E.J. Horvitz and G. Rutledge. Time-dependent utility and action under uncertainty. In *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence,* Los Angeles, CA, pages 151-158. Morgan Kaufmann, San Mateo, CA, July 1991.

(Mclaughlin 1987)
J. Mclaughlin. *The Utility-Directed Presentation of Graphical Simulation*. Technical Report KSL-87-59, Knowledge Systems Laboratory, Departments of Computer Science and Medicine, Stanford University. December, 1987.

(Mesarovic 1970)
M.D. Mesarovic, D. Macko, and Y. Takahara. *Theory of Hierarchical, Multilevel Systems*. New York: Academic Press, 1970.

(Miller 1956)
G.A. Miller. *The Magical Number Seven, Plus or Minus Two*. Psychological Review, 63: 81-97.

(Simon 1973)
H.A. Simon. *The Organization of Complex Systems* in Hierarchy Theory, H.H. Pattee, ed., pages 3-27. G.Brazilier, 1973

(Waugh and Norman 1965)
N.C. Waugh and D.A. Norman. *Primary Memory,* Psychological Review. 72:89-104.