

Heat Exchanger Demonstration Expert System

D. G. Bagby

Communications Systems Research Section

R. A. Cormier

Radio Frequency and Microwave Subsystems Section

This article describes a real-time expert system intended for detecting and diagnosing faults in a 20-kW microwave transmitter heat exchanger. The expert system was developed on a LISP Machine, Incorporated (LMI), Lambda Plus computer using Process Intelligent Control (PICON) software. The Heat Exchanger Expert System has been tested and debugged. Future applications and extensions of the expert system to transmitters, masers, and antenna subassemblies are discussed.

I. Introduction

An expert system is a computer program that embodies organized knowledge concerning some specific area of human expertise sufficient to perform as a skillful and cost-effective consultant [1]. The goal of the Heat Exchanger Expert System Project was to write a program that achieved a high level of performance in diagnosing and troubleshooting heat exchanger problems in the Deep Space Network. This article describes the results of that undertaking and discusses their implications for future DSN activities. It was concluded that the project was successful and that the same techniques can be applied elsewhere in the DSN.

With inadequate monitoring and control, problems can remain unidentified until serious damage has been done and can then take too long to fix. Moreover, operations become more complex as functions are added to support new deep space missions. Expert systems can enhance operator productivity by quickly identifying problems, diagnosing the causes of the problems, recommending appropriate solutions, and predicting impending problems that have not yet occurred. The heat exchanger expert system is a harbinger of things to come.

II. Heat Exchangers

The radio frequency power required to transmit navigation and command signals to distant spacecraft generates large amounts of heat which are dissipated by the transmitter heat exchanger. Just as a water pump cools an automobile engine by circulating water through the engine and through a radiator, so too does the heat exchanger cool a transmitter by passing water through it to a cooling coil (see Fig. 1).

When a heat exchanger fails to perform its cooling function properly, the entire transmitter becomes inoperable. Many hours can be lost in finding the source of the problem, and many days may be required to repair it.

In 1980, approximately 25 percent of all transmitter discrepancy reports arose from heat exchanger malfunctions. This finding led to the installation of a new generation of heat exchangers for the DSN's low-power (20-kW) transmitters, which by 1987 had reduced the share of all transmitter discrepancy reports attributable to heat exchanger failures to 15 percent. A typical example of heat exchanger failure occurred in Spain during the winter of 1986–87, when 70 mph winds and 20°F temperatures combined to freeze all three

coils. The problem took 4 months to correct. On another occasion, in 1988, a failed sensor on a California heat exchanger indicated a flow when there wasn't any. The unsuspecting operator turned the transmitter on, and the resulting heat load burned up the klystron focus magnet, thereby deactivating the antenna for many weeks.

Mishaps such as these and many others could have been prevented had a skilled engineer been available to constantly monitor all heat exchanger dials. Such an arrangement is, of course, economically infeasible. Failing this, the creation of an expert system with many of the capabilities of a skilled engineer seemed to be a worthwhile investment. This was the motivation behind the heat exchanger expert system.

III. The Expert System

The task was to test the feasibility of developing an expert system for maintaining a heat exchanger. Since no active DSN heat exchangers were available, an old heat exchanger returned from the field to the Jet Propulsion Laboratory was used for testing purposes. Sensors and a data collection device were added to the heat exchanger, and a computer was leased to process and analyze the data and to provide diagnostics.

The first task was to determine which problems the heat exchanger expert system should address. This involved a careful analysis of the problems that field personnel encountered, the amount of time they spent solving these problems, and the measures that they would consider helpful in performing their work. The principal conclusions of this study were that the expert system should:

- (1) Isolate faults to the least replaceable unit.
- (2) Provide supervisors with equipment status and maintenance reports.
- (3) Furnish operators with alternatives if equipment is not operational.
- (4) Perform continuous on-line data acquisition and analysis.

These conclusions constituted the design requirements for the expert system. Each task is discussed below.

A. Isolate Faults

The most difficult part of the job was to acquire from the expert the knowledge for solving heat exchanger problems and to translate that knowledge into a machine-readable data structure suitable for automated problem solving. For every heat exchanger component, the expert was asked to identify all possible modes of failure, the symptoms or indicators of such

failures, the sensor information necessary to isolate the failures, and the corrective actions required to remove them. The expert's responses were then coded by a knowledge engineer into a set of "if-then" rules having the general form:

IF <situation> THEN <action>

where each rule corresponded to a piece of knowledge furnished by the expert. A representative example of a rule is shown below:

IF inlet resistivity < 2 megohms and outlet resistivity > 2 megohms,

THEN conclude that the main coolant loop is being contaminated.

The final knowledge base of the Heat Exchanger Expert System contained over 300 such rules, organized as shown in Fig. 2.

B. Status Reports

In addition to isolating faults, the expert system had to provide the user with current status reports on all important heat exchanger components (e.g., whether or not the heat exchanger was on, which pumps or fans were on, if pumps and fans were cycling properly, and whether or not the heat exchanger was operating under local control). Measurements such as these were essential for improving equipment availability and reducing operational costs. Other important data for heat exchangers included equipment operating hours, momentary fault reports, past equipment failures, and repair records. These data elements were provided automatically by the expert system.

C. Furnish Alternatives

The expert system was also required to give the operator guidance about how best to proceed whenever the heat exchanger was "not operational." In that event, there were two possibilities: either the heat exchanger was "available," which meant that it could still be used but that there were things wrong with it, or, alternatively, the heat exchanger was "not available," which meant that the heat exchanger could not be used under any circumstances. The expert system had to be able to distinguish between these two situations and to issue appropriate instructions accordingly (see Fig. 3).

The failure of the heat exchanger control connection due to pin corrosion is a type of failure which need not prevent transmitter operation; the heat exchanger can be operated locally independent of the transmitter. This failure actually occurred at a DSN station, and the heat exchanger was subsequently

turned on locally. Many other failures also allow normal operation provided appropriate action is taken. In some cases, the heat exchanger may be operated in a limited capacity until repairs can be made. Typical examples include fan failure (one cooling fan is usually adequate except on the warmest days when maximum uplink power is also required) and unstable coolant temperatures (stable coolant temperature is critical only for experiments requiring an extremely stable RF carrier, such as gravity wave searches). The expert system also provides a printed copy of this type of information to the operator whenever a problem occurs, thus improving equipment availability.

D. On-Line Data Acquisition

It is not enough to know that a fault has occurred. The user would also like to be able to predict when a fault is likely to occur so that preventive maintenance may be performed. Continuous on-line data acquisition was incorporated into the expert system design to make such predictions possible. A least squares regression line was applied to selected sensor values to direct the operator's attention to impending failures. The expert system first alerts the operator to the presence of a potential problem, and then, by comparing those trends in on-line sensor values with certain prespecified parameters, the system predicts a time when those trends will become critical. This mechanism mimics an attentive engineer and permits detection of wear-out and component degradation before actual failure. Repairs can then be scheduled.

An example of a slow failure that can be detected in a timely manner with on-line data monitoring is a slow coolant leak. Such leaks often go undetected until the heat exchanger is turned off because the coolant has tripped the low-level sensor. Another example is loss of cooling efficiency. Cooling radiators collect dirt, bugs, and miscellaneous debris, all of which interfere with the cooling. This problem is detected by on-line monitoring of the core air flow pressure. Upward trends in pressure are evidence of fouling. The monitoring of ambient air temperature can also be used to protect the coils from freezing. When the temperature decreases below freezing, automatic turn-on of the circulating pumps provides protection if the cooling fans are not turned on. (If the heat exchanger is turned on with the cooling fans operating, this will freeze the coils almost immediately.)

When actual failures occur, the status of parameters just prior to the failure can provide information that reduces the time of isolating a fault because pertinent knowledge is localized, organized, and interpreted. Quite often, once a failure has occurred, it is impossible to operate the equipment. Short-circuited parts or failed interlock sensors are examples of faults that can be identified from data monitoring information of the operation just prior to failure.

IV. Implementation

A. Inference Engine

In addition to the knowledge base of rules, a mechanism was needed for manipulating these rules to make logical deductions and diagnostics. This mechanism, kept separate from the knowledge base, is essentially a mathematical theorem prover called an "inference engine." Its job was to derive conclusions about the status of the heat exchanger from current sensor values and the rules furnished by the expert system. Thus, if the knowledge base contained the rule:

IF A THEN B

and a sensor reported that A was true, the inference engine would deduce that B was also true. Then the assertion that B was true might in turn trigger other rules having the form:

IF B THEN C

and so on.

Program execution consists of a continual sequence of such cycles that persists until either no rule executes or a halt is commanded. At each cycle, all rules whose preconditions are satisfied by the contents of the working memory are determined. If more than one rule is activated, one is selected by means of some suitable "conflict resolution" strategy. All the actions associated with the selected rule are then performed, and the database is changed accordingly.

Alternatively, if B is suspected to be the cause of a problem, then, armed with a rule of the form IF A THEN B, the inference engine may work backward and ask the sensor system to provide any confirming evidence about the existence of A.

The most difficult thing about applying expert system techniques to a heat exchanger is that everything is time-dependent. Consequently, all rules and variable values must be couched in terms of a particular point in time. For example, the simple rule:

IF A > B THEN C

becomes

IF A > B at time T THEN C at time T + 1

e.g.,

IF tank level > (tank level as of 10 minutes ago),

THEN conclude tank level is rising.

The tank level problem is one of the simplest examples. Evaluating sequences of events over time with an expert system proved much more difficult. As another example, for the expert system to ascertain whether or not two pumps were cycling properly, it had to have an internal model describing the last time a particular pump was on compared with the other pump. This placed a significant demand on the inference engine, which not only had to compare events that occurred at different points in time but also, for purposes of isolating a fault that had transpired, had to recall past data or request confirming or disconfirming data from the data supplier—all in real time!

At the time of this study, there was only one commercially available inference engine possessing such capabilities: the Process Intelligent Control, or PICON, produced by LISP Machine, Incorporated (LMI), of Cambridge, Massachusetts. Accordingly, PICON was procured from LMI, together with its Lambda 2 Plus computer.

B. Sensors

Before a single rule was developed, 81 sensors were placed on the heat exchanger. For various reasons, not all of them could be used in the development of the rules. Consequently, the rules were written for a subset of 57 sensors. Of the 24 unused sensors, 3 failed, 4 did not provide any information due to equipment failure, 4 monitored parameters that were never actually used, 7 monitored transmitter heat load, which proved unusable, and 6 provided information available by other means. The sensors that had been intended to capture the heat load could not be used because of the uncontrolled heat loss in the connecting pipes running between the heat exchanger and the transmitter. This piping, exposed to rain, snow, wind, and sun, made heat balance calculations impractical. Even test loads produced large standard errors.

With the 57 sensors that were used, it was possible in some cases to identify a particular failed component, but in general fault isolation was confined to specifying the appropriate procedures and test equipment necessary to further isolate the fault to the defective part.

C. Simulations

Simulation was chosen as the method for modeling the heat exchanger. The entire expert system was divided into independent pieces and separately simulated sensor values. Once confidence was achieved that the individual pieces were operating properly, they were joined together. Finally, simulated sensor values were replaced with actual sensor values.

There were a number of shortcomings with this approach. It failed to account for sensor noise, signal timing, and unfore-

seen parameter variations. The rule base that worked so well for simulated sensor values was often inadequate when applied to real sensors. For example, the rule:

IF reservoir level is decreasing,

THEN send "Differences in tank level indicate a leak" to the engineer

worked well in a simulated mode but misfired when confronted with fluctuations in the readings of the actual reservoir sensor, even when the actual reservoir level itself was constant. The use of smoothing did not eliminate this problem, for noise was not the only factor that had to be considered (Fig. 4). The reservoir level dropped whenever the pump was turned on. This difference was traced to the expansion of the piping caused by the increased pressure of the coolant with the pump operating. The effects of temperature, pressure, time delays for equipment to turn on and off, and sensor hysteresis were other considerations that were initially overlooked for the simulated data.

D. User Interfaces

It is not enough for an expert system to simply come up with the right answer; that answer must also be presented to the operator in such a way that it can be understood. Here, the LISP Machine contained icons, graphs, and displays that made every expert system operation clear. Three different types of heat exchanger problems illustrate these capabilities: a coolant flow interruption, a fan failure, and a leak in the main reservoir tank.

1. Coolant flow interruption. One way to create a flow interlock is to partially close the klystron collector flow valve of the transmitter. Once flow falls, the expert system immediately detects a fault, and the words "Flow Interlock" appear on the console display (Fig. 5). At the same time, the set of instructions listed in Fig. 6 is sent to the operator via a printer. The sequence of rules and sensor values leading up to that message can also be explicitly displayed.

2. Fan failure. To create a fan failure, one could open the fan circuit breaker. When the circuit breaker opens, the expert system detects the source of the fault and notifies the operator with the message "Warning! Fans not operating properly." As before, the operator can call up detailed instructions and graphs to help pinpoint the source of the problem.

3. Reservoir leak. An illustration of the predictive capabilities of the expert system can be obtained by attaching the reservoir level transducer to a stochastic ramp function. The regression line fitted to these reservoir level values and a pa-

parameter specifying the minimum safe reservoir level (20 gallons) then trigger the graph and warning message shown in Fig. 7.

IV. Tests

A test procedure was developed to verify the accuracy of the rule base and to measure its ability to provide operational information useful in the maintenance of the 20-kW transmitter heat exchanger. Key features of the plan included a 72-hour soak test, an identification of various sensor failures, discovery of temperature and flow interlocks, the isolation of short-circuit faults, and the location of leaks. The test procedure was applied to the expert system on two separate occasions: in August of 1987 and in October of 1987. There was significant disagreement between the tester and the system engineer about which tests had passed and which had failed due to different interpretations of test specifications. Figure 8 presents the more conservative results, i.e., the stricter interpretation of the specifications. As the diagram illustrates, the first series of tests yielded correct responses 54 percent of the time, while the second series of tests had a correct response rate of 70 percent. The August tests were performed before the rule base had been debugged and before the heat exchanger was fully operational. By October, the rule base had been debugged but the heat exchanger was still in need of repair. All remaining errors were explainable and repairable.

The tests results indicated that an operational heat exchanger expert system could be developed.

V. Computer System Problems

The unexpected bankruptcy of LMI in April 1987 limited the amount of software and hardware support available during the entire development process. Although the LISP Machine itself is now working properly, initially it had serious problems that would have been impossible to resolve without customer support from LMI. These problems included a hard disk crash, a keyboard failure, and a problem within the terminal itself. Numerous software crashes were traced to unsuitable environmental conditions such as nearby arc welding and high-power transmitter testing. These disappeared with the introduction of a line conditioner.

VI. Lessons Learned

The principal lesson learned from this project was that real-time expert systems are feasible in DSN applications. The LMI Lambda computer and its PICON software were easy to work with and expedited the creation of the Heat Exchanger Expert System.

More time spent initially on requirements and problem definition would pay off in an implementation. Detailed test specifications would be written and agreed upon between testers and designers early in the design process. This would be done after agreement was reached about the scope of the system and before the detailed design was begun. This could eliminate discrepancies such as those found in the test results described above.

More effort defining the types of sensors that are required to diagnose a problem would be useful. Less effort spent trying to decide how best to use already available sensors is indicated. The latter strategy produced many overly complex and cumbersome rules that could have been eliminated by the addition of a single sensor.

A difficult aspect of developing a real-time expert system is establishing reliable connections with external devices and sensors. An elaborate tailor-made C computer program was written just to effect communication between R-TIME, a proprietary real-time data interface, and the 57 heat exchanger sensors.

Another problem encountered in this study was that sensors were not reliable, and additional sensors were needed to check the accuracy of other sensors. The sensor data tended to be noisy, and many sensors failed to operate as expected. Much time could have been saved had rules been added to the knowledge base to look for inconsistencies among collections of sensor values. Simulations need to model the statistical characteristics of both the signals and the sensors.

Real-time expert systems are difficult to build and to debug. The systems can be fooled or misled by the execution sequence of data requests and rule firings. This causes the expert system to generate erroneous messages, which greatly complicate the debugging process and validation of the knowledge base itself. Extra time must be allowed for the debugging process.

VII. Summary and Conclusions

This project developed and demonstrated an expert system that monitors and diagnoses transmitter heat exchanger failures in real time. Sensors and data collection devices were added to a heat exchanger returned from the field. An expert system was developed to collect and provide heat exchanger data and provide diagnostics. All of this was accomplished in 8 months with the expenditure of 2 work-years of effort.

The system will undergo further tests and refinement before it is made fully operational. Eventually, the expert system

will be transferred to Goldstone, California, for testing by DSN operators. If it is found that the expert system reduces transmitter downtime, then the same technology will be applied to other troublesome subsystems. One likely candidate

is the 70-meter antenna assembly, which now accounts for over half of all DSN downtime (54.8 percent). Other possibilities include the traveling-wave maser or the 20-kW CW microwave transmitter itself.

Reference

- [1] F. Hayes-Roth, D. A. Waterman, and D. B. Lenat, *Building Expert Systems*, Reading, Massachusetts: Addison-Wesley, 1983.

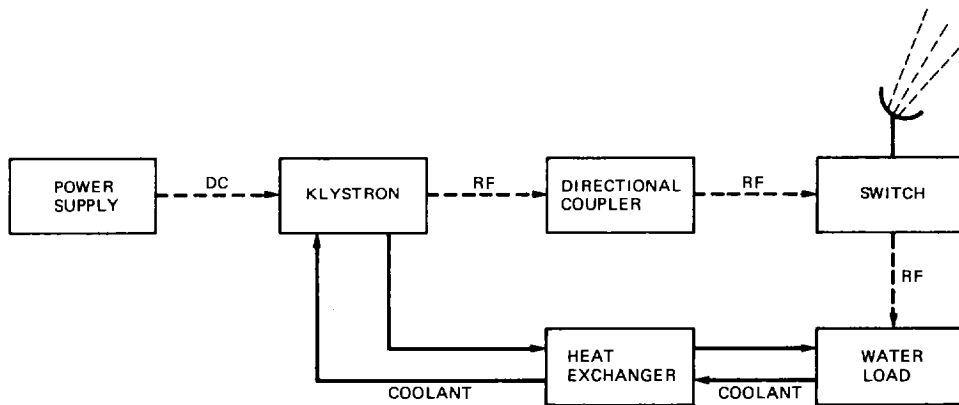


Fig. 1. Block diagram of the 20-kW transmitter

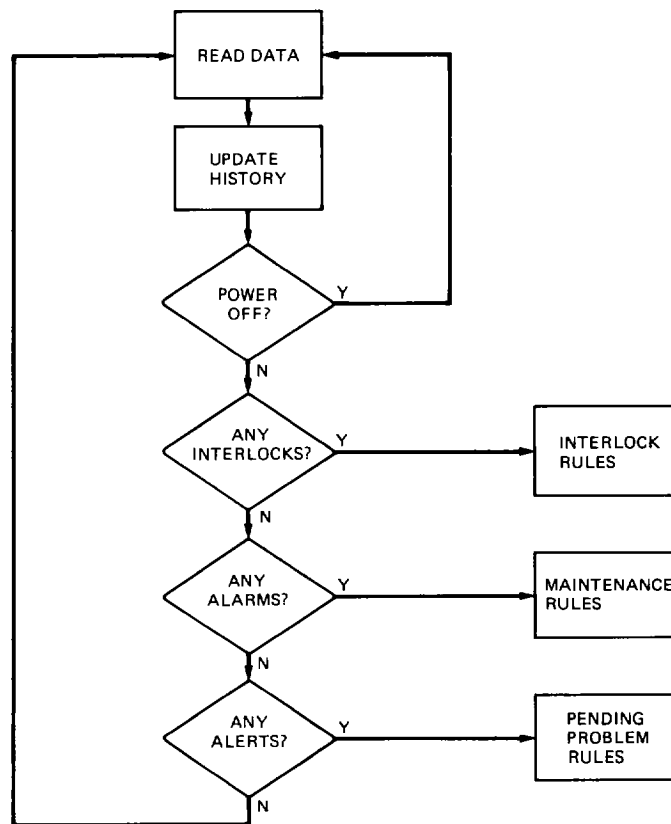


Fig. 2. Expert system rule set structure

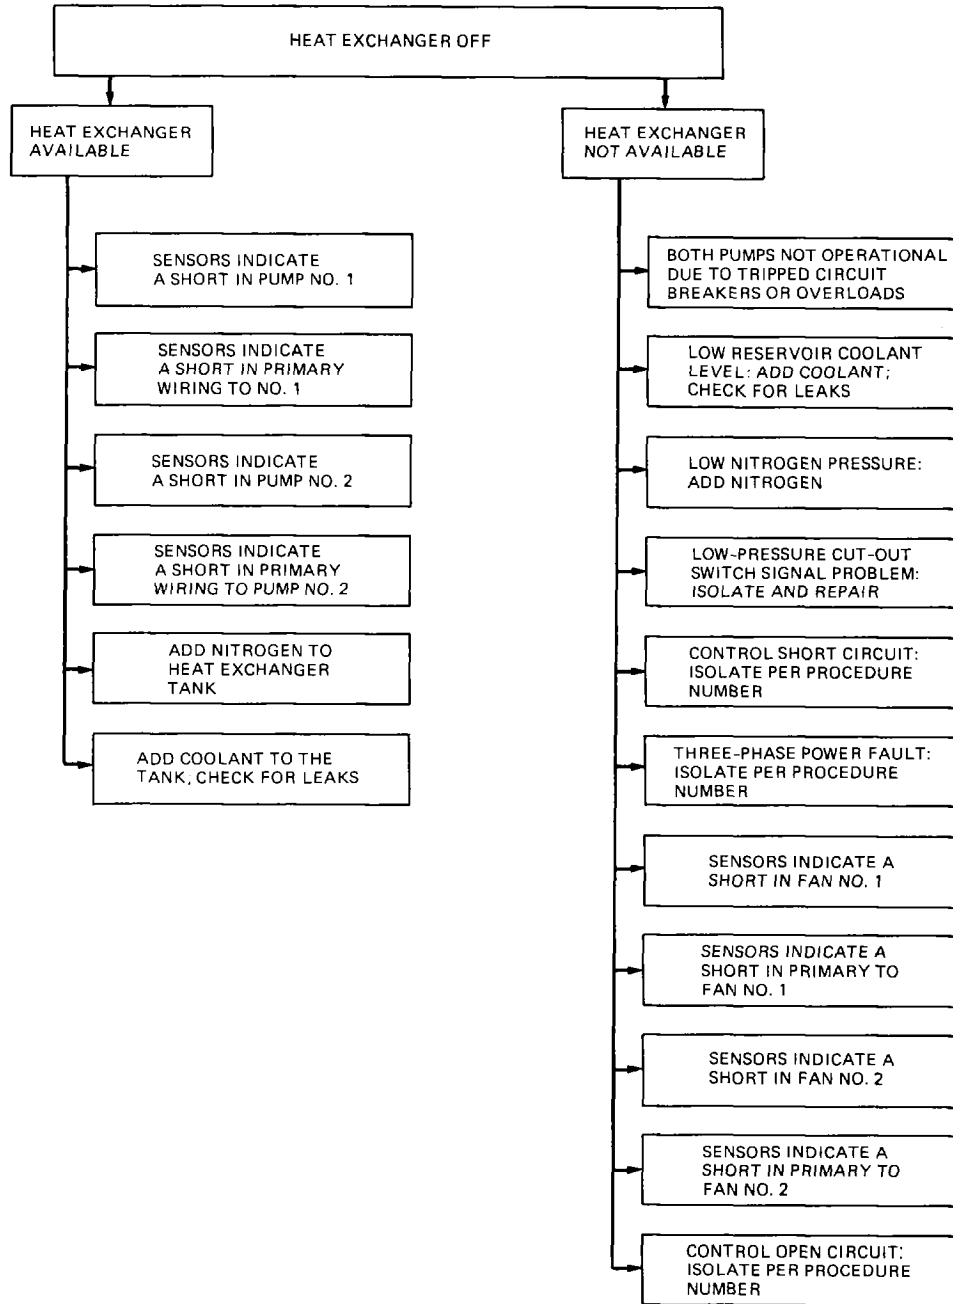


Fig. 3. Heat exchanger "off" problem categories




Fig. 4. Comparison of (a) unsmoothed and (b) smoothed data



Fig. 5. Expert system console

```
.....  
.....  
PROCEDURE 2.1.1.4  
  
A flow interlock is indicated. Go to the power amplifier and to the  
heat exchanger and determine that:  
  
a) Heat exchanger input and output shutoff valves are fully open.  
b) Power amplifier input and output shutoff valves are fully open.  
c) All coolant circuit return valves are fully open.  
d) Adjust coolant inlet throttling valves for each cooling  
circuit to required flow, per manual TM 03202.  
  
If the above procedure has cleared the interlock then the  
problem was not a heat exchanger problem.  
  
.....  
.....
```

Fig. 6. Sample expert system printout

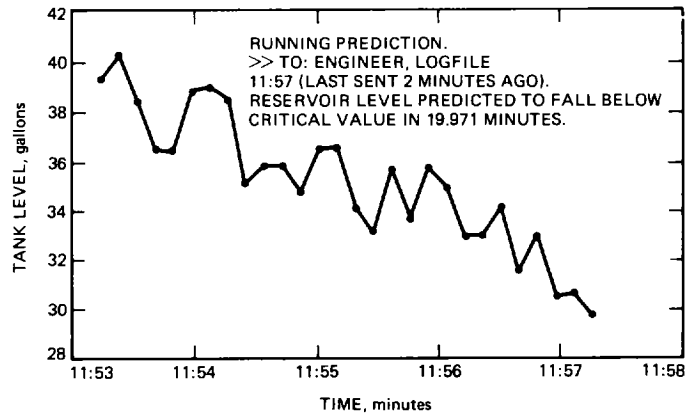


Fig. 7. Sample expert system predictive analysis (heat exchanger reservoir level)




Fig. 8. Test results of the heat exchanger expert system