AI Magazine Volume 7 Number 4 (1986) (© AAAI)

Robert L. Osborne, Ph.D

Online, Artificial Intelligence-Based Turbine Generator Diagnostics

Introduction

The need for online diagnostics in the electric power-generation industry is driven by a number of significant factors. Due to the low number of new power plants being built by electric utilities, the average age of existing power plant equipment in the United States and its susceptibility to failure is increasing rapidly. Figure 1 shows the percentage of power-generation equipment over 20 years old as a function of year. Note the rapid increase of average age after 1980 and the fact that by the year 2000 fully 50 percent of all generation equipment in the United States will be over 20, the oldest average age of power plant equipment ever experienced by U.S. utilities. Thus, there is a need to know what the actual operating condition of the equipment is at all times, so that outages can be avoided by taking corrective actions at the earliest possible time and by preplanning for outages if they become necessary in order to to minimize their length.

In order to provide increased information on the actual operating condition of the equipment, the utility industry has installed additional monitoring capability utilizing power plant computer systems that measure system variables and present these variables effectively. Using data highways, these power plant computers allow the operator to display monitored variables on color CRTs in a variety of ways. Displays include showing all variables above an alarm level, sensors that are out of service, variables superimposed on diagrams of the equipment, and variables plotted as a function of time. The value of variables can be printed periodically or when required.

Unfortunately, regardless of how sophisticated the plant computer system is, knowing the value of variables such as temperature, pressure, and vibration level, that are reported by the computer does not always allow you to know what is wrong with the equipment; that is, the system does not tell the operator a bearing is failing or a conductor is broken. Today, it takes a skilled person to interpret the value of measured variables to determine what is actually wrong with the equipment.

In order to aid the equipment operator in making better operating decisions, the need was recogniized in the mid-1970s to place in a computer the capabilities of diagnosticians who know the relationships between measured variables and the condition of the equipment. At this time, an approach utilizing probabilities was pursued; a small microprocessor-based system was built and demonstrated to the utility industry at a symposium in 1980. The results showed that there was utility interest in this kind of product. It was also recognized, however, that a commercial-size system using hundreds of variables and identifying hundreds of conditions could not be obtained through this approach. Placing the knowledge in the computer was awkward and time consuming, and only one malfunction at a time could accurately be identified. Thus, we were left knowing what the product was that we wanted but not an acceptable method of implementing that product.

Abstract The development of an online turbine generator diagnostic system is described from conception to initial field verification. The system is composed of a data center located in the power plant that collects data from online measurement devices and communicates these data to a centralized diagnostic facility in Orlando, Florida, where the actual diagnosis is done. The resulting diagnosis and recommended actions are transmitted to the power plant where they are displayed to the operator by the data center. The marketplace need, initial approaches to the product, system specification generation, rule base development, and initial system field verification are described. The artificial intelligence (AI) diagnostic program has been diagnosing seven large utility generators since July 1984 and has correctly diagnosed a significant number of generator and instrumentation problems. Issues such as a centralized approach, rule base quality control, and the range of resources needed for a successful product are discussed.

Dr Robert L. Osborne is manager of Diagnostics, Monitor and AI Development at Westinghouse Electric Corporation, Service Technology Division in Orlando, Florida, 32826-2399.

Expert System Approach

A decision to pursue expert systems as the basic tool in the development of diagnostic systems was made in 1981; this decision was dictated by a number of factors. The development effort is lessened considerably with an AI approach when compared with conventional programming in a language such as Fortran. Although expert knowledge can be written into conventional software, the program input process requires not only experts but also programmers. Using AI software, nonprogrammer knowledge engineers can create expert systems through interactive input sessions with experts. Using AI software also allows easier modification to an existing expert system. Conventional fault tree analysis does not indicate the validity of the answer, but an expert system, just like a human diagnostician, can give the confidence associated with each of the diagnoses presented to the user. This confidence is extremely important because few power plant equipment diagnoses are 100 percent certain, and the user can be basing multimillion dollar decisions on the information presented by the system. An expert system can display the method it used to reach its conclusion. This clarifies the diagnosis and provides an excellent vehicle for training new people in the field. Conventional programming discourages this method of training. Thus, when applied in a cost-competitive environment, AI becomes very effective compared with standard programming approaches.

The first step in the development of the diagnostic system product was to initiate the development of the basic AI tool. The result of this development was the Process Diagnostic System (PDS). PDS is a forward-chaining, rule-based system in which sensors, hypotheses (see figure 2), malfunctions, rules, and turbine structure are represented as schemata in SRL (Fox 1979; Wright and Fox 1982). The initial implementation utilized the MYCIN (Shortliffe 1976) approach to the representation and propagation of certainty. The implementation was modified, though, because of the existence of erroneous sensor data due to sensor degradation or spurious readings. In particular, the following modifications were made:

- The fuzzy minimum for conjunctive evidence was replaced with a weighted average that reflected the degree to which the evidence should be considered useful in the decision. These weights can dynamically be altered according to sensor health. This removes the problem of MYCIN's underestimation of belief identified by Wise (1986).
- A rule's certainty factor was extended to include the specification of the necessity and sufficiency of the evidence.
- Earlier sensor readings and hypotheses were archived to allow time-series analysis as part of the reasoning.
- Logical sensors which are the composite of multiple physical sensors, were used to reduce the impact of erroneous readings.

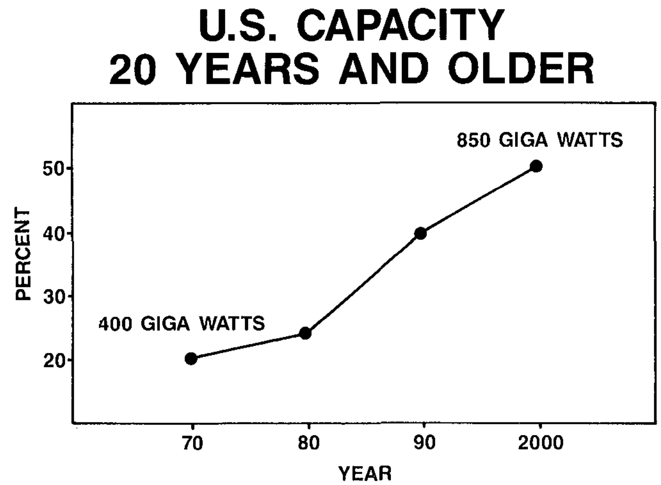

Figure 1. Percentage of Electric Power Generation over 20 Years Old.

```
{{ pds node
   MB: "level of belief in the node being true"
   MD: "level of disbelief in the node being true"
   CF: "level of certainty = mb - md"
   SUPPORTING-RULES: "rules for which this node is hypothesis"
   SUPPORTED-RULES: "rules for which this node is evidence"
   SIGNAL: "contains signal schema name(s)"
   DESCRIPTION: "English description of the node"
   HAS-IS-A: (or sensor hypothesis malfunction) }}
```

Figure 2. Generic Node in PDS.

- Metarules were used to alter diagnostic rules, weights, and certainty factors when sensor degradation was identified.

Once the forward-chaining deduction cycle is completed for the current sensor readings, the operator can enter into a mixed-initiative interaction in which the system can elicit information not available through sensors. PDS was developed using SRL and Franz Lisp. The production version is written in C. For a complete description of PDS, the reader is referred to Fox, Kleinosky, and Lowenfeld (1983).

In the initial stage of product development, one of the critical decisions was whether to place the expert system in a computer located in each power plant or to utilize a centralized computer. The centralized approach was chosen for a number of reasons. First, the knowledge of the relationships between variables that can be measured and the condition of the turbine generator resided primarily with design and service engineers who had worked with the equipment for a number of years. A centralized approach allowed a diagnostic rule base to be written that could improve over the years because of its growing knowledge of actual component problems and its ability to recognize malfunctions that had not occurred to date.

Because of the relative infrequency of equipment failures on any one turbine generator, a basic question was how the knowledge base could be improved in the future, utilizing both the designers knowledge and the users experience, without a centralized facility. By utilizing a central diagnostic center approach, any knowledge gained in one power plant to improve the rule base could immediately be made available to all users. This is in contrast to a decentralized system where every individual computer in each power plant would have to be reprogrammed each time a rule base improvement was made. In addition, by manning the center 24 hours a day, if a new malfunction appeared in any unit that was not in the rule base, a human diagnostician would have many resources immediately available to aid the operator and to update the rule base if necessary. These resources include databases that contain the outage histories of all units being diagnosed, detailed design information, and the equipment design engineers themselves.

After developing the basic tool, PDS, the next step was to convince management a commercial-sized system, that is, one with hundreds of variable inputs and hundreds of equipment conditions diagnosed, could be made a commercial reality. It was decided that a steam chemistry diagnostic rule base would be generated. Knowing the state of steam chemistry in a power plant is important in preventing damage to those parts of the power plant where the steam flows. Steam chemistry was chosen for the first diagnostic rule base primarily because more steam chemistry upsets occur per month than actual equipment problems; more upsets meant that this rule base would be exercised considerably more frequently than one for a turbine or a generator, allowing much more evaluation of the AI approach in a shorter period of time. Also, much more steam chemistry field data were stored that could be used offline to evaluate the diagnostic rule base.

In early 1982, a small 25-rule chemistry diagnostic rule base system was written using PDS. This system demonstrated that the AI approach was practical and could be expanded to large systems.

Additional work on the steam chemistry rule base (Bellows 1984) resulted in a demonstration, primarily for management, of a much larger rule base. This ability to do larger rule bases convinced division management to significantly expand the entire program, which led to establishing the diagnostic center at the power generation headquarters in Orlando, Florida.

Diagnostic Center

The diagnostic center was designed to centralize our entire diagnostic program in one geographic location. Within the power generation headquarters building, the center was located next to the service engineering department because that department would have the most interaction with the diagnostic center on a day-to-day basis.

Figure 3. Diagnostic Center Information Area..

The diagnostic center has five distinct areas: the information area, the diagnostic conference room, the operations center, a customer laboratory, and two AI laboratories. Each of these areas has its own mission as part of the diagnostic product line.

Initiating a new technology in a 70-year-old division presented an educational challenge to the employees in the division as well as to our customers. The information area, shown in figure 3, is one of several effective techniques used to inform people about the diagnostic product and the AI technology it is based on. So as not to interrupt the engineers working on the product, an IBM-XT is programmed to allow anyone to walk up and request, through the use of a "mouse," subjects such as AI, the diagnostic center, and expert systems. For each subject requested, several paragraphs of explanation are presented to the onlooker. When the person has finished reading these paragraphs they can return to the menu and choose another subject.

Another portion of this information area contains an infrared detector that senses the presence of anyone in the area, triggering a four-minute recording keyed into a selectively lighted diagram which explains the diagnostic center and how it is integrated into the service engineering department. Although this method was effective, especially for groups, the frequent replaying of the recording distracted the engineers in the area, so the recording is now only used when groups are visiting the center.

The diagnostic conference room, shown in figure 4, is an important part of the diagnostic center. This futuristic-looking area provides a place where personnel can interact in problem-solving sessions; demonstrations of the products can be made; and direct communication between division

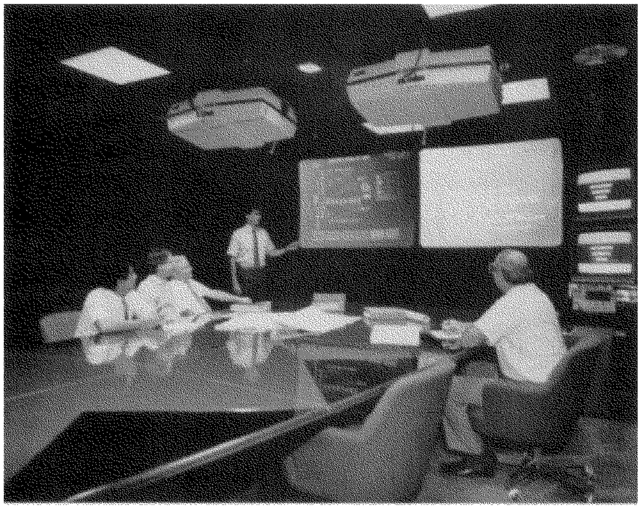

Figure 4. Diagnostic Center Conference Room.

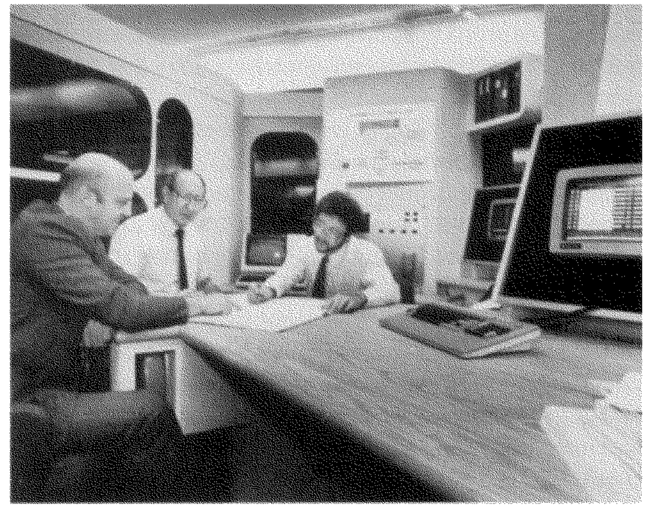

Figure 5. Artificial Intelligence Laboratory.

personnel and other locations, such as our research and development center and the field, can instantly take place. The room itself was designed around the concept of a "paperless conference room." Ideally, the only paper that needs to be brought into the room is for personal note taking. All other information is contained in databases or comes in live from the field and can be accessed by personnel in the room. To facilitate live viewing, there are two single-gun video projectors that can be seen in the upper portion of figure 4. In addition, there are two video monitors and a video tape deck to the right of the main screens at the front of the room. These machines facilitate reviewing of videotapes taken of malfunctioned equipment in the field and sent to the diagnostic center.

The operation of the entire room is controlled from a console at the back. From this console, information in various company computers can be accessed. For example, engineering information on the design of a particular turbine generator can be obtained, thus allowing personnel to concentrate on problem solving rather than searching through volumes of papers and notebooks for information. All data and diagnoses of equipment in the field are available instantaneously to everyone in the room.

The diagnostic conference room is used as an education area for both customers and company personnel. For those who have had no exposure to the product or to AI, several demonstrations are available. One demonstration is composed of an introductory portion that gives the overall goals of the project, a medical example to explain expert systems and how they are used in diagnoses, and an explanation of the diagnostic system being demonstrated.

Next, a demonstration is presented of the diagnostic system in operation when a generator is experiencing an abnormal condition. Field data stored in the computer are fed into the diagnostic program, and the results are presented on the screen. As the condition of the generator continues to deteriorate, each diagnosis is presented to the audience as if the diagnostic center were actually diagnosing the unit using live data. The audience is drawn into the action and begins to understand what the product is and how it works in a way that is hard to obtain using other methods. Over a thousand people have now seen this demonstration.

The customer laboratory is another area of the diagnostic center, used to test software and hardware systems before shipment to the field. It also displays how the data center and various field monitors are integrated into the diagnostic center to form the diagnostic system product.

The operations center is dedicated to 24-hour-per-day, seven day-a-week support of our utility customers. The center is staffed by a diagnostician who reviews all online diagnoses being done on customer equipment. A description of how the operations center works with the data centers in the field is given in The Data Center.

Two AI laboratories are also in the diagnostic center. The interior of one of the laboratories is shown in figure 5. The knowledge engineer sits in the area in the center of the picture, in front of two CRT screens. The experts sit at the table in the foreground and have CRTs that contain the same information as those in front of the knowledge engineer. The laboratory is designed to facilitate the transfer of knowledge from the experts to the computer. The knowledge engineer is responsible for rule base generation. The engineer's responsibilities include interviewing the experts in each equipment area, extracting their knowledge, and placing that knowledge in the computer. The knowledge engineer is the only person who has the authority and the responsibility for entering the rules into the computer. When experts disagree, the knowledge engineer must determine what actually is entered into the rule base. Whenever possible, two or three experts are interviewed at a time. This provides a synergy where the experts reinforce and add to each other's knowledge. It also provides checks and balances on the knowledge that goes into the rule base because several experts are likely to identify all questionable rules or confidence factors, preventing them from going into the rule base. Interviewing sessions are generally limited to a half day because the process is very tiring.

Problems of noncooperation of experts has been virtually nonexistent, which is a tribute to those involved. Also, a demonstration of the operation of the diagnostic system is given to the experts so that they have a clear understanding of the product they are contributing to before the interview begins. The interview process involves the experts and the knowledge engineer concentrating on a particular equipment area. When all the rules have been generated in a session, the experts leave. The knowledge engineer subsequently draws a diagram of the rule base and enters the rules into the computer. Once the knowledge engineer is satisfied with the operation of the new portion of the rule base, the experts are reconvened. The engineer then exercises the rule base by inputting data. The experts are asked if they agree with the resulting diagnosis. If they agree, no changes are made. If they disagree however, corrections are agreed to, the rule base is modified, and the resulting rule base is again exercised with appropriate data. The process is repeated until a satisfactory rule base is created.

Another method used to ensure a high-quality rule basis is the design review. Design reviews are conducted by a panel of engineers knowledgeable of the equipment being diagnosed. Most of these panel members are not involved in the generation of the rule base. A list of action items is generated by the panel, and the knowledge engineer is responsible for modifying the rule base to satisfy these action items. The design review is a very effective tool in rule base development.

In order to be effective, the knowledge engineer must have a good knowledge of the equipment the diagnostic rules are for. Without such knowledge, the engineer is severely handicapped in dealing with the experts and in making judgments of what finally goes into the rule base. Basically, the experts should feel that the knowledge engineer understands the details of the knowledge being transmitted. However, the knowledge engineer does not need to know Lisp. Generally, it is very difficult to find a knowledge engineer who knows the equipment and who also has a Lisp background.

Engineers both in Orlando and at the research and development center in Pittsburgh who know Lisp and can have the PDS tool modified when necessary, are available to the knowledge engineer. Virtually all of the many improvements made in the diagnostic center programs are to satisfy the needs of the knowledge engineers. We have found that there is a continual need to improve and extend the programs to suit the particular requirements of the diagnostic products being developed.

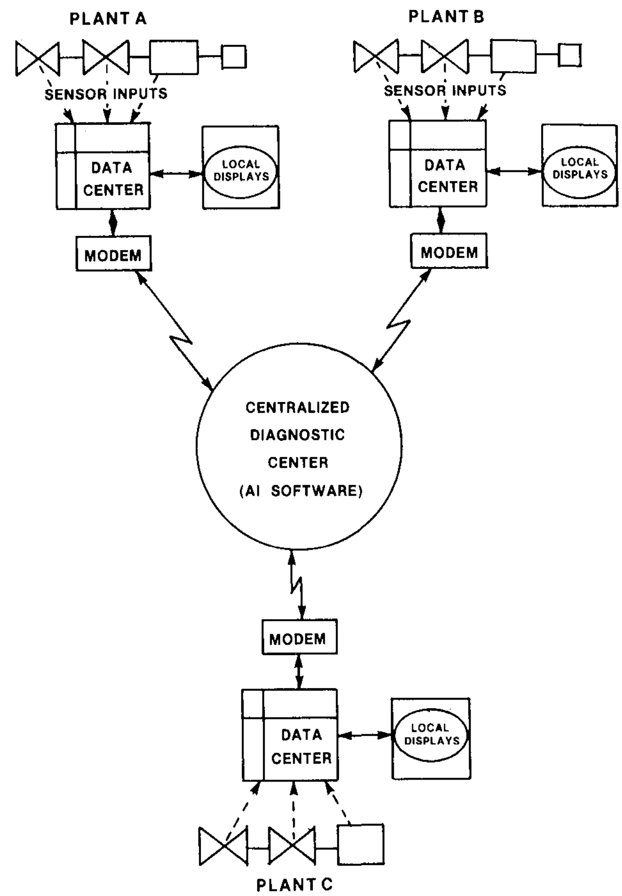

Figure 6. Diagnostic System Schematic.

The Data Center

Other extremely important parts of the diagnostic system product are the data center (Osborne, Kemper, Emery, Trosky, Logan, and Rodriquez 1985) and the monitors that are located in the power plant. Figure 6 is a diagram of the entire diagnostic system, showing the data centers feeding data to the diagnostic center and receiving the resulting diagnosis for display to the operator.

Figure 7 shows a data center. The console is designed to handle one or two generators. The operator utilizes a touch screen, color CRT for primary communication with the system as shown on the right side of the picture. On the left is the input/output cabinet. A keyboard is available for the few times that data are entered or for the electronic mail feature that is used to communicate with the diagnostic center.

The data center was designed to be both an advanced, in-plant data monitoring system and an extension of the diagnostic center that can display the diagnosis to the operator and give direct recommendations on what actions should be taken. The diagnosis is presented in the form of a list of

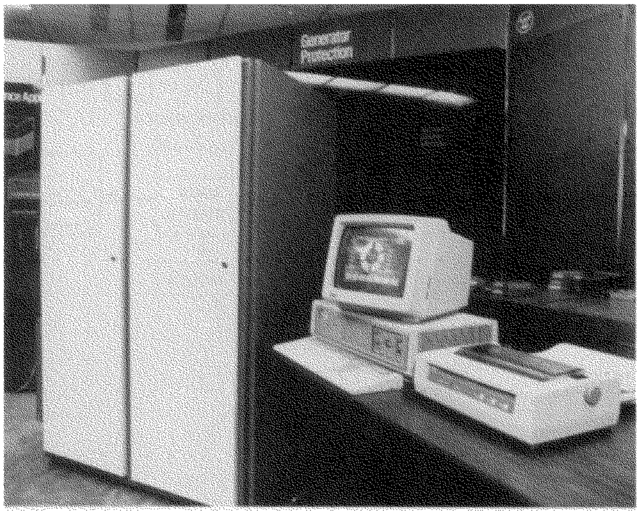

Figure 7 Data Center Console.

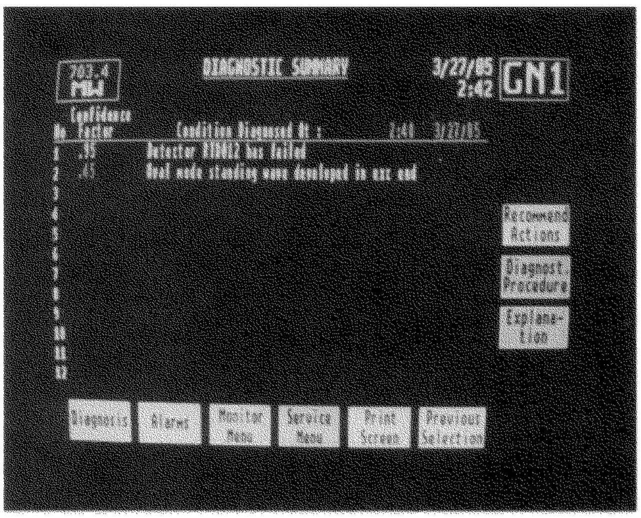

Figure 8. Typical Diagnostic Display.

''candidate conditions.'' A typical list is shown in figure 8. The condition (malfunction) with the highest confidence factor is always at the top of the list. Confidence factors have values from negative one to positive one. A confidence factor of positive one means the diagnostic system is absolutely sure that the condition actually exists and a confidence factor of negative one indicates absolute surety the condition does not exist. Confidence factors near zero indicate uncertainty. This condition typically occurs when measured variables associated with the condition deviate very little from normal. This is consistent with what a human diagnostician would tell the operator under similar circumstances.

The system is designed to give the operator the ability to ask the diagnostic center computer through the in-plant data center for a procedure that can be implemented to increase the quality of the diagnosis. For example, if two different equipment conditions have similar effects on the monitored variables, there might be an action the operator can take, such as change the load or read a pressure gage and enter the value into the data center. The diagnostic center then sends a new diagnosis based on both the online variables and any values manually input into the data center. The reason such procedures are given is that if a variable is used infrequently and primarily to increase the confidence factor, installing an online sensor can't be justified economically. The procedure, in essence, substitutes for the online sensor during the few times that data is needed.

An electronic mail capability is built into the data center to allow the power plant operator to send messages to the human diagnostician at the diagnostic center and vice versa. Because the diagnostic center is manned 24 hours a day, assistance can be obtained beyond the online diagnosis. For example, the diagnostician can access databases on equipment history and detailed engineering design information. One of the key functions the diagnostician performs is recognizing those few instances when a potential problem occurs that is not in the rule base. The diagnostician assists the operator with the immediate diagnosis and then has the rule base reviewed and updated with new rules that relate measured variables to the malfunction. Without this centralized approach, it would be difficult to improve the rule base.

Customer Involvement

An important part of the development of the diagnostic system has been customer involvement. A project to verify the diagnostic system was jointly undertaken in 1984 with the Texas Utilities Generating Company (TUGCO) (Osborne, Gonzalez, Weeks, and Martin 1986). The objective was to monitor and diagnose seven of their generators in east Texas.

Phase 1 of the project included modifying the existing generator monitoring systems at each plant to permit them to transmit reports of data every hour as well as when an alarm condition took place. These data were then translated into a diagnosis by the generator diagnostic system and displayed at the diagnostic center in Orlando. If a developing abnormal condition was diagnosed, diagnostic center personnel contacted the plant operators with the diagnosis and recommendations.

The phase 1 system was limited because the modified monitoring system did not have the capability of communicating between the Orlando diagnostic center and the plant. Any communication in this direction took place by a separate electronic mail system or by telephone line. However, the data center now has the capability of full, two-way communications, so that the diagnostic output is accessible to the plant personnel.

The data center data-acquisition system monitors approximately 110 different sensors in a gas-cooled generator and its auxiliaries. This is compared to the phase 1 system total of 22 sensors, which limited the number of identifiable conditions in the phase 1 system to 94. The phase 2 system is presently able to identify approximately 350 conditions, using over 8,500 rules.

Another major difference between the phase 1 and phase 2 systems as it applies to diagnostics is the frequency of data transmissions. The data center (phase 2) can transmit continuously through a dedicated data link, thus allowing a high degree of trending. Trending was less feasible with the phase 1 system, which called once an hour and only transmitted one set of data at a time.

Operational Experience

The phase 1 system started operation in July 1984. Since that time, a number of incidents have been correctly diagnosed. The incidents range in importance from sensor failures to broken conductors in the stator winding of the generator.

The value of the diagnostic system is best illustrated by an incident on January 25, 1985, when broken phase coil conductors in one of the monitored units caused the unit to be brought offline. The diagnostic expert system correctly diagnosed the situation 2-1/2 hours before the conditions reached the alarm level. The unit was removed from service before serious damage occurred and was on turning gear ready for synchronization four days later. The utility avoided a potentially costly event by recognizing the condition early and taking the appropriate action.

The capability of this system to discriminate false alarms from real emergencies was demonstrated during phase 1. In February 1985, correction factor inaccuracies in a temperature-normalizing algorithm caused temperature deviations to exist. These deviations could have been interpreted as broken conductors on the generator stator winding. However, the diagnostic system correctly diagnosed the situation as bad correction factors. Thus, an unnecessary unit shutdown was avoided.

In September 1984, the expert system made use of a temperature comparison scheme to identify a sensor malfunction that otherwise could not be identified through conventional range checks. In a second incident in September 1984, the diagnostic system diagnosed a conductor discontinuity at a particular location. Because the severity was not high, a recommendation was made to continue running for the remaining three weeks before a planned outage but to be ready to repair the problem at that time. Inspection of the generator three weeks later proved the diagnosis to be accurate. This incident illustrates the value of the diagnostic system as a predictive maintenance tool.

The phase 1 expert system, due to the sensor input limitations, was only able to monitor and diagnose conditions in the generator stator winding system. The phase 2 system, however, will not only be able to do what the phase 1 system did but will also be able to diagnose problems in the remainder of the generator, which includes the hydrogen auxiliary system, the seal oil system, the generator rotor, and the excitation system and its control.

Presently, the data centers have been installed in each plant, and initial testing has been done by TUGCO instrument and control engineers. As a result of both this testing and earlier discussions with TUGCO personnel, a number of changes to the data centers have been made. These changes have improved the functionality of the system and increased its user friendliness to the operators. Other improvements include the ability to utilize time histories of variables and the use of procedures displayed to plant personnel to obtain additional data.

Conclusions

Based on our field experience to date, several conclusions can be made. The use of a software tool such as PDS is very important in the development of diagnostic systems because it allows personnel who are experts in equipment diagnosis to be utilized as knowledge engineers, even though they do not know a specialized computer language such as Lisp. The use of an AI approach to online equipment diagnostics is practical. To create a commercial online diagnostic system takes a large number of highly trained, experienced engineers in a number of disciplines dedicated solely to the system's development.

References

Bellows, J. C 1984 An Artificial Intelligence Chemistry Diagnostic System In Proceedings of the Fortyfifth International Water Conference, 15-25. Pittsburgh, Pennsylvania

Fox, M S. 1979. On Inheritance in Knowledge Representation. In Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokyo, Japan

Fox, M S ; Kleinosky, P ; and Lowenfeld, S. 1983. Techniques for; Sensor-Based Diagnosis In Proceedings of the Eighth International Joint Conference on Artificial Intelligence, 158-163. Menlo Park, Calif.: American Association for Artificial Intelligence

Osborne, R. L ; Gonzalez, A J ; Weeks, C A ; and Martin, J 1986. First Year's Experience with Online Generator Diagnostics Paper presented at 1986 American Power Conference, forthcoming. Chicago, IL.

Osborne, R. L ; Kemper, C T ; Emeιy, F. T.; Trosky, W T ; Logan J R ; and Rodιiquez, R. 1985 Turbine Generator Data Center for Artificial Intelligence-Based Diagnostics Paper presented at 1985 Joint Power Generation Conference, no published proceedings Milwaukee, WI

Shortliffe, E. H 1976 *Computer-Based Medical Consultations· MYCIN.* New York: American Elsevier.

Wise, B. P 1986. An Experimental Comparison of Uncertain Inference Systems Ph D diss , Dept of Engineering and Public Policy, Carnegie-Mellon Univ

Wιight, J M., and Fox, M S 1982 SRL/1.5 User Manual, The Robotics Institute, Carnegie-Mellon Univ