SAND91-1436C

EXPERT SYSTEMS TO AID IN WIND FARM OPERATIONS

L. L. Schluter
Sandia National Laboratories
Albuquerque, New Mexico

F. Nateghian
U.S. Windpower, Inc.
Livermore, California

G. F. Luger
University of New Mexico
Albuquerque, New Mexico

Abstract

An expert system is a knowledge-based program that provides solutions to problems in a specific domain by mimicking the behavior of a human expert. Expert systems can have several advantages over traditional programming methods; however, developing an expert system generally involves a considerable amount of time and money. Therefore, careful investigation must be done to ensure that a problem is suited for an expert system application. This paper examines several areas where an expert system may help wind farm operators lower their operational costs. Justifications for using expert systems rather than traditional programming methods are given. This paper also discusses some of the design decisions that were made in developing an expert system for U.S. Windpower that will aid it in diagnosing wind turbine failures.

Introduction

Department of Energy estimates indicate that current wind energy conversion systems (WECS) can produce a kilowatt-hour of electricity for approximately 7 to 10 cents. To be competitive with conventional fuels (i.e., oil, coal, etc.) the cost to produce a kilowatt-hour of electricity must be reduced to less than 4 cents. Environmental impact laws may close the cost gap in the future, but to become more cost competitive today, wind farm operators must continue to reduce their costs. In the past, expert systems have helped companies reduce their operation and maintenance costs. One example of this is Digital Equipment Corporation's use of a program known as XCON. XCON is an expert system that configures all of Digital's VAX and PDP-11 series computers. Before Digital started to use XCON, it was experiencing considerable financial expense because of configuration errors. If a computer was shipped with the wrong or missing parts, Digital had to correct the problem quickly. The use of XCON has greatly reduced this situation and thus reduced Digital's operational cost. Also, every computer that is configured by XCON is configured in the same way. This has effectively produced a standard configuration for particular computers. Now when a maintenance person works on a VAX 780 it will be configured the same as all of the other VAX 780s that he has worked on. This reduces the amount of time the maintenance person needs to acquaint himself with a customer's system and reduces the maintenance cost [1]. This is one example in which an expert system has proven to be useful.

This paper discusses a project that was done in cooperation with U.S. Windpower. U.S. Windpower is the leading manufacturer of wind turbines in the United States. The company also operates a large wind farm in the Altamont Pass located just east of San Francisco, California. This wind farm consists of approximately 4000 wind turbines. The ultimate goal of this research project is to develop an expert system that will aid U.S. Windpower to lower the operation and maintenance cost associated with its wind farm. To accomplish this, an area in which U.S. Windpower will benefit from an expert system needs to be defined. U.S. Windpower has several areas in which they feel an expert system will aid them in lowering their operational costs. These areas need to be investigated to determine if an expert system is appropriate. This paper discusses the areas investigated and gives arguments justifying an expert system as an appropriate approach. The paper then explains some of the design decisions that were made in the development of an expert system for U.S. Windpower. The expert system will aid its maintenance teams in the diagnostic procedures associated with wind turbine failures.

Expert Systems

The two main parts of an expert system are a database and an inference engine. The database holds two kinds of information: a knowledge base that holds the specific information about the domain, and a fact list that holds the current state of the system. The knowledge base contains cause-and-effect relations, in the form of IF-THEN rules, that represent how an expert would solve a problem. The facts contained in the fact list can be data input into the system by an outside source (i.e., user, read in from a file, etc.) or facts created by rules as the expert system executes. The inference engine uses both sets of information to find an answer to the problem (see Figure 1) [2]. The inference engine will fire rules (execute the THEN portion) when there are facts contained in the fact list to satisfy the rule (meet the requirements of the IF portion) until a goal is reached or until no more rules can fire.

In contrast to a traditional programming approach, which uses data and algorithms to create a program, expert systems are created using knowledge and inference. Heuristics, rules of thumb, used by humans to solve a particular problem are coded into the knowledge base. Inference is used by expert systems to derive new facts from existing facts, imitating one way human experts solve problems [3].

Expert System

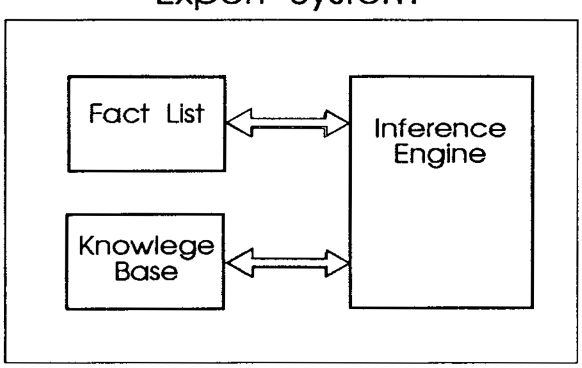

Figure 1: An Expert System.

Defining Areas for Expert System Application

The first major task for this research project is to define a process for which U.S. Windpower will benefit from an expert system. U.S. Windpower believes several of its operation and maintenance (O&M) processes could benefit from using expert systems. These processes involve

1) Using distributed meteorological sensor data to optimally start up and shut down wind

   turbines in a large wind plant sited in complex terrain.

2) Analyzing the alarm data signatures to automate decision making for alarm "clearing" and restarting of turbines.

3) Analyzing alarm and operational data to predict imminent failure.

4) Analyzing alarm data to help better prepare crews with probable causes of wind turbine failures.

The main goal associated with all of the processes is to increase the amount of time the wind turbines are producing electricity. The first process mentioned will accomplish this goal by keeping a turbine operating when wind conditions are appropriate, but shutting the turbine down when excessive fatigue would occur, thus reducing down time for the turbine due to maintenance. Normally when an alarm occurs on a wind turbine, it indicates conditions that could damage the wind turbine, so the turbine shuts down. The second process listed will increase the operating time of a turbine by automatically restarting the turbine when conditions are appropriate rather then waiting for a human operator to initiate the restarting process. Based on past experience, U.S. Windpower personnel believe that before a component fails, the wind turbine will give a series of alarms indicating the component is faulty. The third process would use an expert system to analyze alarm history to predict the component failure. The fourth process diagnoses a wind turbine failure more accurately to reduce the amount of troubleshooting that a windsmith (field maintenance person) must perform[4].

Initially all of these areas seem suitable for an expert system since they mimic the behavior or thought process of a human expert. All of the areas involve humans using heuristics and inference to make decisions. There are several guidelines that researchers have developed for determining whether a problem is appropriate for an expert system solution. Three are taken from Luger and Stubblefield [5]. One of these guidelines is: "The need for the solution justifies the cost and effort of building an expert system." U.S. Windpower did not inform us of its current O&M costs or by how much they expect to reduce their O&M costs with the aid of an expert system. The fact that the company is willing to spend its time and money in support of the project states that it believes this first guideline is satisfied in the areas listed.

Another guideline is: "Cooperative and articulate experts exist," which implies that the knowledge to build the rule base exists. For example, a large amount of experience currently exists in diagnosing wind turbines that have been taken out of service. This is the type of knowledge required for analyzing alarm data to help better prepare crews with probable causes of wind turbine failures. This area also seemed promising because of the large number of expert systems that have been developed in the diagnostic problem domain. One of the earliest expert systems designed was called MYCIN. MYCIN was designed to diagnose and recommend treatment for blood infections [6]. Many examples of expert systems used to diagnose faults in computers, automobiles, and electronic components can be found in today's literature. The initial research in this area, for U.S. Windpower, was aimed at building a rule base that would more accurately pinpoint the problem for a windsmith before he gets to the turbine. After several discussions with U.S. Windpower personnel, it was apparent that the expert system, with the information that the windsmith currently uses, could not more accurately pinpoint the problem.

The reason an expert system is not able to produce a more accurate diagnostic is that the alarms are ambiguous, and several different components may have caused a single alarm. The alarms will point the windsmith to a general subsystem in which to start troubleshooting, but he must then physically go to the turbine to pinpoint the problem. For example, if an automobile will not turn over when the key is turned, a mechanic knows that some component in the starting system is at fault (the battery may be dead, the battery cables may be loose, the starter may be at fault, etc.). To fix the automobile, a mechanic must check each component until the faulty component is found. Since the expert system cannot physically go to the turbine, it will not be able to pinpoint the problem. The consensus of the windsmiths interviewed was that a more accurate diagnosis could be done if the expert system could incorporate information such as the turbine configuration, electrical environment, condition of neighboring turbines, and snap shots. A snap shot is detailed information on the state of the turbine when an alarm occurred. The amount of this type of information is so large no windsmith currently uses it; therefore, the expertise to decipher it is not currently present. The expertise to use this information could be developed over time, but to get an expert system working in the near future, experts must be available to describe how to use the information in diagnosing a failure.

This does not exclude an expert system from being useful in this area. Referring to the automobile example, a mechanic will know from past experience that the battery is the most likely component to have caused the problem or perhaps is the easiest component to check. This implies that there is a preferred method for troubleshooting the automobile. If this preferred method is followed, the average time for troubleshooting the automobile is reduced. U.S. Windpower personnel believe that an expert system will be very useful in capturing the knowledge of the more experienced windsmiths on how to troubleshoot a wind turbine. A windsmith with more experience will know the most likely component to have caused the alarm and will start troubleshooting at that point. An expert system can give prioritized suggestions for a procedure to follow when troubleshooting a wind turbine. The default priority for the suggestions will be based on the alarm condition for which the turbine was taken out of service. Work history will be involved in changing the default priorities of the suggestions. For example, if the pitch sensor was replaced last week, the suggestion that the pitch sensor is at fault will have its priority lowered. By following the procedures used by the more experienced windsmiths, the average time spent troubleshooting will be reduced.

Another guideline in determining whether a problem is appropriate for an expert system is whether the problem may be solved using traditional computing methods. If it can, then it is not a candidate for an expert system. Because expert systems rely on heuristic approaches, it is unlikely that an expert system will out perform an algorithmic solution if one exists. In creating a suggested diagnostic procedure, an expert system will have several advantages compared to a traditional computing method. The number of problem codes that currently exist is approximately 75. A procedural program could be written to handle the number of rules required by the problem codes. However, work history must also be taken into account. If there are approximately four suggestions for each problem code, then there are 12 different priority lists for each problem code. This is based on the number of permutations that a set of four items can have [7]. This means at least 900 different results could occur. With a problem of this magnitude, it is best to start out with a program that will address a subset of the problem and then gradually increase the scope of the program. For example, the problem codes that occur most often can be addressed first. Rules can be added later to include other problem codes not addressed initially. Because of the ease with which the rule base of an expert system can be expanded, an expert system is more appropriate for this type of approach than a traditional computing method.

Another justification for using an expert system is that the rules may change. If U.S. Windpower discovers a component that has a high failure rate, it may switch manufacturers to reduce the failure rate. If the company is successful, the most likely cause of an alarm today may not be the most likely cause a year from now. Thus, the rule base will change. Since the rule base is a separate entity within the expert system, rules can be easily added, deleted, and/or modified as needed. This is not necessarily true with a traditional computing method.

A final argument in favor of an expert system involves the use of a learning system. Two goals of learning systems are to provide more accurate solutions and to cover a wider range of problems without re-programming [8]. Because of the large number of possible outcomes and because the rule base will change with time, a learning system has many advantages. More research needs to be done in learning systems to determine if such a system is appropriate. However, an expert system can evolve into a learning system more easily than a traditional computing method.

Based on these arguments, an expert system is the appropriate approach to better preparing work crews with the probable causes of wind turbine failures.

The analysis of alarm and operational data to predict imminent component failure was also studied to see if an expert system would be applicable in this area. The number of possible combinations of alarms is over 16 million. In order to predict imminent component failure, a series of alarms must be studied to find a pattern that indicates a component is going to fail. The possibilities are too vast for any traditional computing method. An expert system using heuristic search rules and a rule base that is easily expanded is better suited for this application. Even with an expert system using heuristics to search the problem space, it may take years to develop a rule base to encompass all of the possibilities. Again, based on the number of rules involved, a learning system that can cover a wider range of problems without re-programming seems to be appropriate.

The major disadvantage with implementing an expert system in this area is that there is currently very little expertise in dealing with alarm history. Thus, no heuristics currently exist in this area. This expertise and associated heuristics will be developed over time. Several U.S. Windpower employees believe, from past experience, that an expert system in this area would save significantly in the maintenance cost of wind farming. So even though the expertise does not currently exist, any expert system should be implemented in such a way that the data to predict imminent component failure is available. As this expertise is developed, the rule base can be expanded to include the new knowledge.

Developing an Expert System to aid in Diagnosing Faulty Wind Turbines

An expert system has been developed and is currently being evaluated by U.S. Windpower. The first specification to be defined was the type of inference engine that should be used. There are two main ways to construct an inference engine; forward-chaining and backward-chaining. Forward chaining uses information from the fact list to move through the IF-THEN rules to reach a result or goal. This method is referred to as data driven. Backward chaining is the reverse of forward chaining. In backward chaining the expert system starts with a result and uses the information on the fact list to confirm or deny the result. To determine whether a forward-chaining or backward-chaining inference engine was more appropriate, we considered the type of output that was required. If only one possible solution is required as the output, then a backward-chaining system is appropriate. If it is necessary to find all possible solutions, then a forward-chaining system is more efficient [9]. The expert system needs to list all possible solutions. Thus, a forward-chaining inference engine will be more efficient in solving this problem. The data required to reach a conclusion is given in the workorder and work history of a turbine. The expert system does not have to acquire any additional data to reach a conclusion. This also suggests that a forward-chaining inference engine is more appropriate than a backward-chaining inference engine [5]. A final argument in favor of a forward-chaining inference engine is that there are a large number of potential goals (approximately 900) but only a few ways to use the given information [5]. For example, the search space is significantly reduced given the alarm information.

It should be noted that for predicting imminent component failure, a backward-chaining system *may* be more appropriate. There are over 16 million possible alarm combinations. Even if the number of components that can fail on a wind turbine is 1600, the number of possible input combinations significantly exceeds the number of possible outcomes. Since there are a large number of possible alarm combinations, the system may require fewer rules to start with a conclusion to prove or disprove [5]. For example, the expert system can start with the conclusion that the brakes need to be replaced. Then from the alarm history, it must prove or disprove this conclusion. This would also allow the expert system to include within its initial rule base just the major components of a wind turbine. The rule base could gradually be expanded to include other components that are considered important. However, not enough research has been done on how to design the rule base for predicting imminent component failure to determine which type of inference engine needs to be used. Therefore, a forward-chaining inference engine will be used because it is more appropriate for diagnosing the cause of a wind turbine failure.

Another important issue is that the expert system should be able to run on a variety of computers. The expert system will be run on a workstation at U.S. Windpower, on a different type of workstation to simulate the system at Sandia National Laboratories, and U.S. Windpower will want the system to run on a personal computer in the future. It was therefore necessary to have the source code available to compile on the different machines.

The inference engine that the system uses is supplied by NASA. It is written in, and fully integrated with, the C language and is called the 'C' Language Production System (CLIPS). It is a forward-chaining inference engine that provides reasonable performance on a wide variety of computers [10]. An additional benefit is that it is supplied to government contractors at no cost.

The information the expert system needs to produce a troubleshooting procedure is the alarm condition for which the turbine was taken out of service and the work history for the turbine. This is the information that the windsmiths currently use to start the diagnostic procedure. The alarm condition is contained in a workorder that is created when a turbine is taken out of service. At U.S. Windpower the workorder and work history of a turbine are stored on the field maintenance computer. U.S. Windpower's data acquisition system records the occurrence of all the alarms. At some point in time, the total number of occurrences for each alarm is sent to the field maintenance computer. This is the type of information the expert system will need in order to predict imminent component failure. Therefore, the expert system must reside on the field maintenance computer or on a computer system that is networked to the field maintenance computer. Figure 2 illustrates the inputs and the output of the expert system.

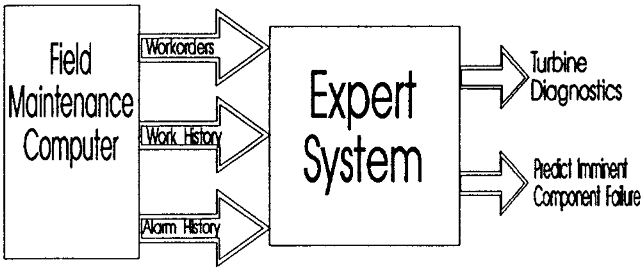

Figure 2: Expert System Inputs and Outputs

CLIPS is designed to be embedded within other programs. Calls to CLIPS are made like calls to any other subroutine. Therefore, a main routine has been created that loads the rule base, initializes the fact list, and then calls CLIPS to execute the rules (see figure 3).

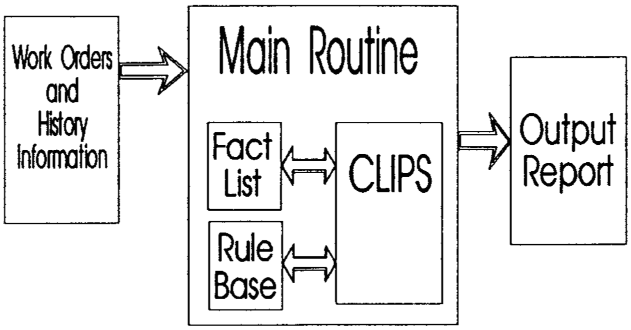

Figure 3: Final Configuration of Expert System

The initial rule base was created by obtaining a maintenance manual with several of the alarms that commonly occur, along with a troubleshooting procedure for each alarm. Data sheets were created containing the alarm that occurred, along with a list of suggestions as to what might have caused the alarm. These sheets were then distributed to several different U.S. Windpower employees as questionnaires. We not only gave them to field maintenance personnel, but also to several engineers who had considerable experience in dealing with the wind turbines. We asked them to see if they agreed with the suggestions listed. If not, we asked them to write down what their suggested approach to the problem would be. This is how the initial rule base was created.

Schluter [11] contains a detailed discussion on the implementation details of the expert system.

Example Output from the Expert System

Site Number: 1023 Problem Code: PC000

Operator Comment:
None

Problem:
The blades could not be pitched to the feather position fast enough.

Possible Causes According to Priority:
-->If the machine is new, the pitch sensor has failed.
-->The brake assembly is faulty.
-->The pitch nut lubrication assembly is not working.
-->There is excessive wear on the bolt threads.
-->The hub could be binding.

Evolving the Expert System into a Learning System

As the sets of collected data grow in our wind farm applications we intend to design a simple learning system. Machine learning will take two forms: first, the attempt to find patterns in the collected data, and second, a generalization technique for retaining information on successful interventions in the wind farm operation.

The first approach to learning has a strong inductive flavor. It will look, for example, into sets of alarm data to aid our understanding on alarm "clearing." Normally, when an alarm occurs on a wind turbine, it indicates conditions that could damage the wind turbine, so the turbine shuts down. Under certain conditions it may be appropriate to clear the alarm and restart the turbine rather than take the turbine out of service. As an analogy, consider a car that has stalled. If there is no apparent reason for the car to have stalled, the operator restarts the car. On the other hand, if the oil light was on or the temperature was excessively high prior to the car's stalling, the operator would not restart the car because the engine might be seriously damaged. We expect to find that certain data sets, when associated with a particular model turbine, create this type of alarm. This information can help in a number of ways: first, it may allow the addition of new rules to our reasoning system to automatically ignore certain alarm situations. Second, even if the new information is never "automated," it can still assist the human expert in better understanding the alarm clearing problem. Thus, the ability of the computer to find alarm patterns aids the human expert.

The second type of learning is often called "case-based learning." The idea is that when the human or computer discovers a remedy to a wind turbine problem, or for that matter any situation of alarms or maintenance, that remedy can be saved for future use in that particular "case." This saved case is made as general as possible so that it can be used to remedy any of a class of future problems. For instance, if a turbine of a particular serial number has some problem that is detected and repaired, the information can be saved for all turbines of that same model (or of some other category that is more general than that particular serial number).

Finally, the machine-learning component of programs such as ours should be seen as simple search techniques that can go over the information that is being continuously collected. The result of this search is to find new patterns in the data to help us understand and remedy problem situations. The "case based" component also allows us to collect and generalize successful interventions in the wind farm operation. These successful cases will allow us to better intervene when new situations arise. More information on models and techniques for machine learning may be found in Chapter 15 of Luger and Stubblefield [5].

Summary

Before developing an expert system, careful investigation needs to be done to ensure that the problem to be solved is suited for an expert system application. U.S. Windpower presented a list of areas for which they felt an expert system would be useful. Initially, all of the areas seemed appropriate for an expert system application because humans were using heuristics and inference to make decisions. However, this observation, by itself, does not justify the use of an expert system. Two processes were found to be appropriate applications for expert systems. They are aiding U.S. Windpower in diagnosing wind turbine faults and predicting imminent component failure. Expert systems have two advantages over traditional programming techniques that these areas can take advantage of. One is that the rule base can be easily modified and expanded. This allows the system to start with a small subset of the problem and gradually increase the rule base to cover more of the problem domain. This reduces the overall complexity of the problem. The second advantage is an expert system can evolve into a learning system. Because there is a vast amount of information that must be studied and since the heuristics may change over time, a learning-based system that can produce more accurate results and cover a wider range of the problem domain without re-programming is desirable.

An initial expert system has been developed and implemented in cooperation with U.S. Windpower. The goal of this expert system is to aid U.S. Windpower in diagnosing wind turbines. The output of this expert system is currently being evaluated. Based on the type of information involved in diagnosing a wind turbine, a forward-chaining inference engine was found to be more appropriate than a backward-chaining inference engine. The inference engine that is being used is supplied by NASA and is called CLIPS. An initial rule base has been created, which contains the problem codes that occur frequently.

References

[1] McDermont, J., "A Rule Based Configurer of Computer Systems," *Artificial Intelligence* 19, pp. 38-59, 1982.

[2] Schildt, Herbert, Artificial Intelligence Using C, McGraw-Hill, Berkeley, California 94710, 1987.

[3] Yazdani, M., Artificial Intelligence: Principles and Applications, Chapman and Hall Ltd., New York, New York, 1986.

[4] Private communication with the following U.S. Windpower personnel:
Geoff Brown, Jason Creighton, Joe Fleck, John Gregory, Tom Harris, Bill Holley, Aaron Howard, Lou Manfredi, Faramarz Nateghian, Rick Owen, Joe Pasquariello, Joel Ryther Paul Smith, John Spikula, and Bob Szymanski.

[5] Luger, George F. and Stubblefield, William A., Artificial Intelligence and the Design of Expert Systems, The Benjamin/Cummings Publishing Company, Inc., Fort Collins, Colorado, 1989.

[6] Buchanan, B., and Shortliffe, E., Rule Based Expert Systems: The MYCIN Experiments of the Stranford Heuristic Programming Project, Addison-Welsey Publishing Company, Menlo Park, California, 1984.

[7] Hillman, A., Alexanderson, G., Grassl, R., Discrete and Combinatorial Mathematics, Dellen Publishing Company, San Francisco, California, 1987.

[8] Forsyth, R. and Rada, R., Machine Learning: Applications in Expert Systems and Information Retrieval, John Wiley and Sons, New York, New York, 1986.

[9] Rich, E., Artificial Intelligence, McGraw-Hill Book Company, San Francisco, California, 1983.

[10] Giarratano, J., CLIPS User's Guide: Volume 1, NASA Lyndon B. Johnson Space Center, Information Systems Directorate, Software Technology Branch, Houston, Texas, 1991.

[11] Schluter, Larry L., An Expert System to Aid Wind Farm Operators, Electrical Engineering Computer Engineering Department, University of New Mexico, Albuquerque, New Mexico, 1991.