CONF-8806146--2

CONF-8806146--2

DE88 012048

APPLICATION OF AI TECHNOLOGY TO NUCLEAR PLANT OPERATIONS

by

Dr. J. I. Sackett

Argonne National Laboratory

Idaho Falls, Idaho 83403-2528

ASEE Annual Conference

Portland, Oregon

June 19 - 24, 1988

MASTER

*Work supported by the U.S. Department of Energy, Civilian Reactor Development, under Contract W-31-109-Eng-38

APPLICATION OF AI TECHNOLOGY TO NUCLEAR PLANT OPERATION

Dr. J. I. Sackett

Argonne National Laboratory

P.O. Box 2528

Idaho Falls, ID 83403-2528

Abstract

In this paper, applications of Artificial Intelligence (AI) Technnology to nuclear-power plant operation are reviewed. AI Technology is advancing rapidly and in the next five years is expected to enjoy widespread application to operation, maintenance, management and safety. Near term emphasis on a sensor validation, scheduling, alarm handling, and expert systems for procedural assistance. Ultimate applications are envisioned to culminate in autonomous control such as would be necessary for a power system in space, where automatic control actions are taken based upon "reasoned" conclusions regarding plant conditions, capability and control objectives.

Introduction

The need for improvement is a common theme heard in the U.S. nuclear industry, not the least of which is the need to improve the reliability and efficiency of many of its plants. This is particularly evident when the U.S. experience is compared to that of Canada and Japan (Table 1). (The U.S. plant capacity factor in 1987 was ~59% while that for Japan was ~77%.) Close examination of this record reveals a number of areas where advanced computer technology can offer substantial help, if a way can be found to integrate it into existing plants.

In this respect, the U.S. is different from many other countries (notably Canada, France, and Japan) where computer technology is being pursued aggressively for specific application to nuclear plants, but there is not a lack of capability or interest on the part of U.S. developers. The U.S. nuclear industry suffers from the fact that no new plants are being designed and built. One can do much more with advanced controls and diagnostic systems if one is developing new plant designs within which the technology can be integrated.

The challenge in the U.S. is one of retrofit of computers to existing plants and of convincing operators of existing plants of their benefits. Fortunately, this may not be so difficult as it may first appear. The existence of a large number and variety of plants will allow a staged approach in development and application of the technology, using equipment already largely in place. It is also spawning a great many diverse projects. The technical problem to be addressed first is effective management and use of information from the plant. Effective use of information is a goal not restricted only to operation. It extends to maintenance, engineering, management and regulation. Indeed, some of the best early application of AI technology

Table I.

Power Generation and Capacity Factors Worldwide*

Nation	1987 No. of Units	1987 MW Gross	1987 Year's Total Generation	1987 Capacity Factor	1986 No. of Units	1986 MW Gross	1986 Year's Total Generation	1986 Capacity Factor
Finland	4	2,400	19,369,281	92.28%	4	2,400	18,779,692	88.79%
Hungary	4	1,760	10,895,454	88.58%	3	1,320	7,403,389	77.15%
Switzerland	5	3,079	22,957,974	84.66%	5	3,079	22,452,235	82.92%
Belgium	7	5,718	41,926,645	83.25%	7	5,718	38,634,559	76.17%
South Korea	7	5,808	37,848,202	82.36%	6	4,766	28,321,216	60.70%
Spain	8	5,812	41,222,014	79.71%	8	5,812	37,462,635	74.84%
Japan	36	28,046	186,182,064	77.40%	33	24,686	167,216,343	76.20%
Yugoslavia	1	664	4,494,654	77.27%	1	664	4,018,858	69.09%
Sweden	12	10,095	67,238,867	77.05%	12	10,095	69,949,843	80.40%
West Germany	19	19,848	130,553,685	74.87%	17	18,295	117,389,823	77.91%
Taiwan	6	5,146	33,107,923	73.34%	6	5,146	26,940,795	61.14%
Canada	18	12,786	80,635,815	71.77%	17	11,813	74,470,130	72.80%
Argentina	2	1,005	6,464,762	67.67%	2	1,005	5,699,939	67.07%
Netherlands	2	540	3,210,360	75.03%	2	540	4,179,700	83.16%
France	53	52,095	265,484,455	59.54%	49	47,170	254,216,251	66.18%
U.S.	106	97,180	478,960,126	58.85%	98	87,241	437,104,009	59.51%
Britian	38	12,940	56,180,839	53.63%	38	12,940	59,076,458	55.92%
India	6	1,330	5,482,212	46.63%	6	1,330	5,136,218	44.28%
South Africa	2	1,930	6,598,532	39.03%	2	1,930	9,322,269	55.14%
Pakistan	1	137	308,310	25.69%	1	137	527,380	43.94%
Brazil	1	657	973,302	16.91%	1	657	144,256	2.51%
Italy	3	1,330	173,628	2.45%	3	1,330	8,650,912	74.43%
Total		341,270,305	1,500,359,103			321,248,073	1,397,096,910	
Average/Unit		793	4,399,880	65.02%		773	4,352,327	66.25%

* From Nucleonics Week, Vol. 29, No. 5, February 1988.

may well be in providing information for such "off-line" use. Such information can be made available and used with minimal capital investment, if effective ties can be established to existing plant data-aquisition systems.

Background of Expert Systems Development

Artificial Intelligence (AI) has been a field of a development for approximately the last 20 years. Present AI applications are commonly referred to as Expert Systems utilizing rules and knowledge about systems and methods of problem-solving learned from experts.

Typically, an Expert System consists of two components; (1) an inferrence engine which processes the knowledge according to defined rules and strategies, and (2) a knowledge base which incorporates the necessary information about the system and its peformance. The inferrence engine is the heart of the problem-solving process which allows solutions to be developed along two basic techniques. One called forward chaining and another called backward chaining. In a forward chaining technique of solution, the inference engine begins with available information and, according to a defined reasoning process, reaches conclusions about that system. A backward chaining solution assumes conclusions about the system and works backward to check agreement with the information available to determine whether or not the conclusion can be supported.

The success of the expert system depends in great measure upon the extent and accuracy of the knowledge base upon which it operates. To develop a knowledge base requires one to gleam extensive information from experts or other sources about the system in question. This information can be gained either by interviewing individual experts or utilizing information available in procedures, maintenance books, text books and the like. The tax code is an example of a well-structured set of rules that could be represented well in a knowledge base; people could ask questions of an Expert Systems Advisor to obtain advice and information in the same way they could from a Tax Accountant. As with most situations, there are circumstances where clear answers are not available. One of the unique aspects of Expert Systems is that they can provide useful information and conclusions in the absence of complete information. It is this characteristic that gives them their advantages and involves them in a process that could be termed Automated Reasoning. It is also this characteristic that will make their introduction into a regulatory process more difficult. However, the systems do have the advantage that their reasoning utilize to reach a given conclusion can be documented specifically as a part of the process, providing a form of cross-check on the accuracy of the conclusion.

There have been a number of notable successes with Expert Systems, one of the first and best known being the program MYCIN. It can make diagnosis and prescriptions for antibiotic treatment of disease utilizing test results and general information supplied by the physician, allowing best choices to be made in the absence of complete information.

It has been shown to provide guidance better than the medical experts in the field, representing as it does a collection of experience from many physicans. Another notable success, has been with an Expert System called XCON which configures the components of VAX computers. VAX computers are quite modular in construction with many, many components all of which are not compatible. The program is used routinely to configure VAX systems to provide optimal configuration to meet individual user needs.

Especially in the last five years, there has been increased commercialization of Expert Systems. There are probably over 2000 systems in use at present. They tend to be very application-specific and are involving a large number of systems dealing with many individual and relatively small problems. In this sense the technology has been consolidating around techniques now proven to be beneficial.

Expert Systems depend upon new ways of processing information. This has required that development of specific programming languages to accommodate them. Basically, these languages are concerned with data that are largely symbolic and with relatively ill structured problems for solutions, by methods primarily concerned with knowledge representation and symbolic computation.

There are two basic languages used in expert systems. They are LISP, developed in the United States, and PROLOG, developed in Europe. Most of the existing Expert System Programs are developed in LISP, used primary because of the simplicity of its syntax and relative ease of using the language. Prolog has definite advantages for some problems, however, especially those involving a highly structured set of rules. LISP and PROLOG may both be contrasted to languages like FORTRAN and other procedural languages which require very strict structure for solving particular problems and require complete information for input. It should be noted, however, that many expert system "shells" are utilizing other languages, notably "C."

Expert Systems are gaining very rapidly in popularity and will likely soon find many applications in nuclear power plant operation. Examples are in providing procedural compliance, ensuring technical specification compliance, scheduling maintenance activities, and intepreting information from alarms or plant operating data. While there are many potential applications, it must be emphasized that the development of an expert system is an activity which should be approached very carefully and with the full understanding about what the work involves. For example, capturing an expert knowledge is still as much a craft as it is a science. It is also very difficult to articulate common sense and intuition. Expert Systems tends to break down where less and less can be specified definitely about the system. One of the more intriguing aspects is that in their development and testing, one will often find that the information available is incomplete in ways not before realized. The very process of development can provide valuable feedback to understanding and completeness of the problem being addressed.

A Process for Automation

Automation of existing plants can be described as a logical process wherein analog controllers are first replaced by digital programmable controllers, then as a next step, networked through a central computer to provide overall supervisory control. The computer systems used for plant data acquisition can increasingly be used for diagnostics. As the diagnostic tools are developed, they in turn can be integrated into plant control to provide functions of sensor validation, best estimates of sensor values, and estimates of plant state to greatly enhance the reliability of control and accident management. The data handling functions can be spun off to include information for maintenance, management, and interfacing with regulatory bodies. Sophisticated graphics become an important part of the system, providing the operator with a good mental picture of what is occurring. Such graphics emphasize iconic representations or display of patterns to describe the process. Other functions include alarm handling, procedure prompting and expert-system guidance to the operator in the course of both normal and abnormal operation of the plant. This progression is illustrated in Figure 1.

A number of demonstration projects with actual plants are underway. An example in the U.S. is the Experimental Breeder Reactor II (EBR-II) which is serving as a test bed for much of the associated technology, with the goal of ultimately demonstrating the feasibility of an automated control system that uses the best of AI technology.[1] The Advanced Test Reactor (ATR), operated by DOE at the INEL, is also being

Figure 1. Elements of AI Technology Implementation for Reactor Plant Control

Fully Automatic Control	Full Authority Computer-Based Control	Autonomous Supervisiory Control	Predictive Simulation and Display	Automatic Damage Control	Reasoned Control Response
Automated Reasoning Technology	Microprocessor Controllers/ Hierarchical Control	Optimal Supervisiory Control	Iconic/ Pattern Display	Fault-Tolerant Hardware/ Software	On-line Procedure Prompting
Expert System Technology	Programmable Digital Controllers	Supervisiory Controls/ Automation of Procedures	Integrated Graphics Display	Sensor Validation/ Smart Sensors	Alarm Filtering/ Expert Advisors
Digital Technology	Local Digital Controllers	Manual Control	Digital Display	Signal Conditioning/ Automatic Data Handling	Digital Alarms/Procedural Guidance
Analog Technology	Analog PID Controllers	Manual Control	Analog Display	Redundant Sensors	Procedural Response
	Local Control	System Control	Display	Data Handling	Control Action

utilized as a development bed for much of the technology.[2] A number of U.S. utilities are also exploring applications of AI technology in individual plants, with some showing great promise of success. The purpose of this paper to review a number of these that represent useful application to operation and maintenance.

Sensor Validation

Sensor validation is extremely important because it is quite common for signals or information to be in error, particularly during off-normal or accident conditions. Errors can result from miscalibration of instruments, from sensors giving erroneous readings because of location, or because the variable being measured is out of its intended range of operation. Also, during transients, instruments may lag the course of the transient and give, therefore, erroneous information.

There are several basic techniques being pursued for signal validation[3]. The simplest is by methods of comparison wherein several sensors reading the same value are compared and erroneous data identified. Such "voting" techniques require at least three sensors and, therefore, can be relatively expensive to implement.

The second technique involves analytic redundancy wherein the process being monitored is modeled, allowing values to be inferred even though they are not measured directly. This allows information to be obtained in the absence of direct measurement, using instrumentation from other locations to be utilized. The difficulty associated with analytic redundancy is the cost and complexity associated with the modeling and the fact that the model itself may contain errors for modes of operation where limited data exists.

One of the more intriguing techniques to have been developed recently is pattern recognition. Pattern recognition may be especially valuable in application to existing plans simply because there is a large base of data from which to develop the necessary information. A pattern recognition system which is proving to be very successful is the System State Analyzer (SSA) developed by EI at EBR-II.[4] It is reviewed as an example of the technology.

The SSA provides an estimate of actual values compared to current readings as well as an estimate of overall plant state. (These are probably the two most important functions that can be provided in support of plant operation.) A unique aspect of this system is that it does not use an analytic model of the plant or associated systems to provide the information, but uses pattern recognition based on learned states from previous operation. (Its operation is based upon techniques of matrix manipulation of large number of signals to establish "patterns" representing best fits to previous operation.) Because of this, it may be applied to existing plants easily and with little cost. The information required exists in historical records of plant operation. Particularly in the U.S., where most AI applications are retrofit to existing plants with an existing operating history, these are important qualities.

To use the system, one provides a list of the signals considered to have a relationship to the item of interest, e.g., power level, flow rates, etc. In EBR-II, the system is being used to provide an estimate of reactor power level, reactor outlet temperature, flow rates and primary pump performance, among others. Once the parameters of interest are chosen, the next step is to select periods of reactor operation considered typical for the plant conditions of interest. These periods of operation are used to "teach" the system what to expect from each of the signals involved. For example, in EBR-II 129 signals are used in the learned patterns from which reactor power level may be estimated. If any of the related signals exhibit strange behavior relative to the others, its behavior is flagged. During a recent period of operation, the sensor measuring mixed-mean reactor outlet temperature at EBR-II drifted, giving an incorrect reading of reactor outlet temperature. However, the SSA was able to provide a predicted reactor outlet temperature that was quite accurate (see Fig. 2) as well as identify the fact that the sensor was drifting. The system is capable of providing an accurate prediction of outlet temperature (or any of the other 129 parameters) in the absence of an actual sensor reading for that parameter. It will continue to be tested, with objective of further developing the system to accommodate plant transients. Of particular interest were recent plant dynamic tests which demonstrated that the system could successfully follow widely changing plant states, providing accurate representation of plant signals (Fig. 3).

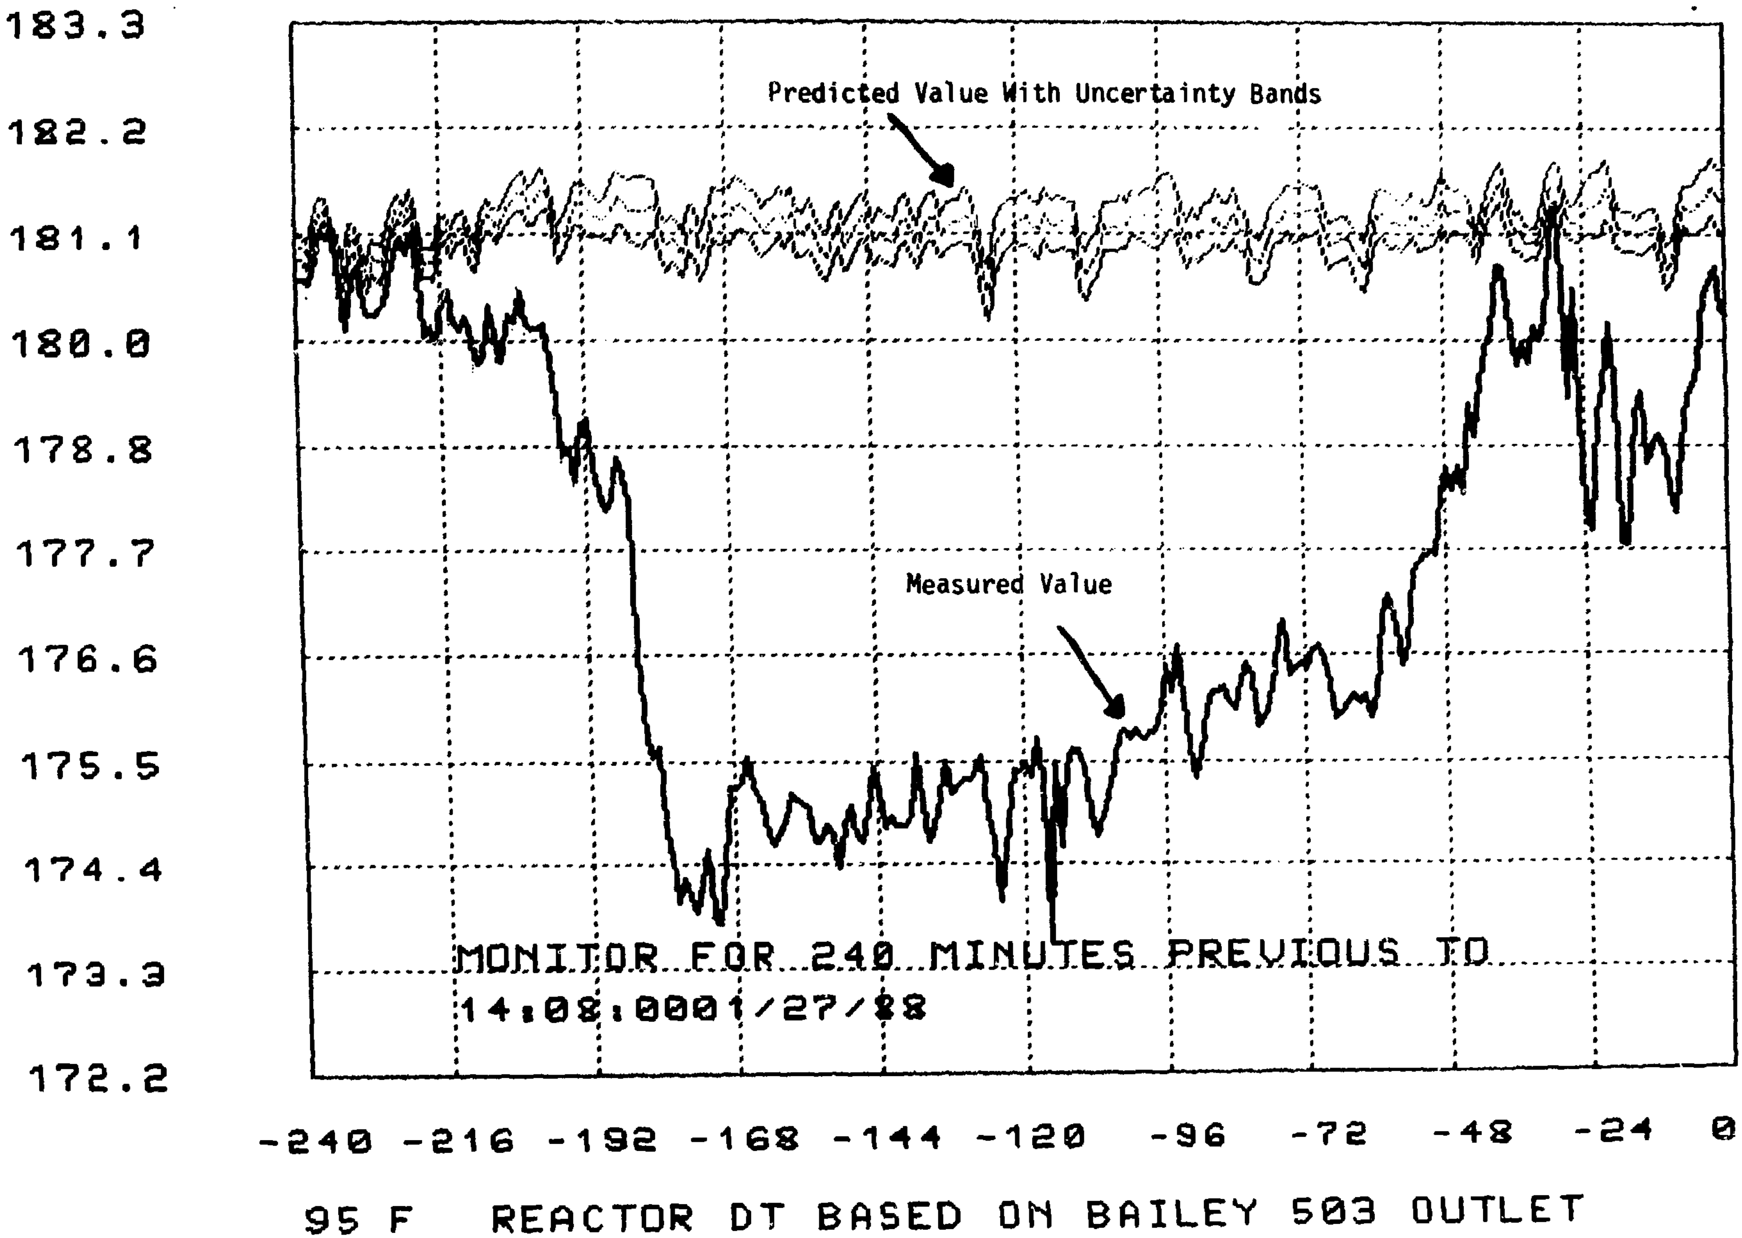

183.3
182.2
181.1
180.0
178.8
177.7
176.6
175.5
174.4
173.3
172.2
Predicted Value With Uncertainty Bands
Measured Value
MONITOR FOR 240 MINUTES PREVIOUS TO
14:08:0001/27/88
-240 -216 -192 -168 -144 -120 -96 -72 -48 -24 0
95 F REACTOR DT BASED ON BAILEY 503 OUTLET

Fig. 2

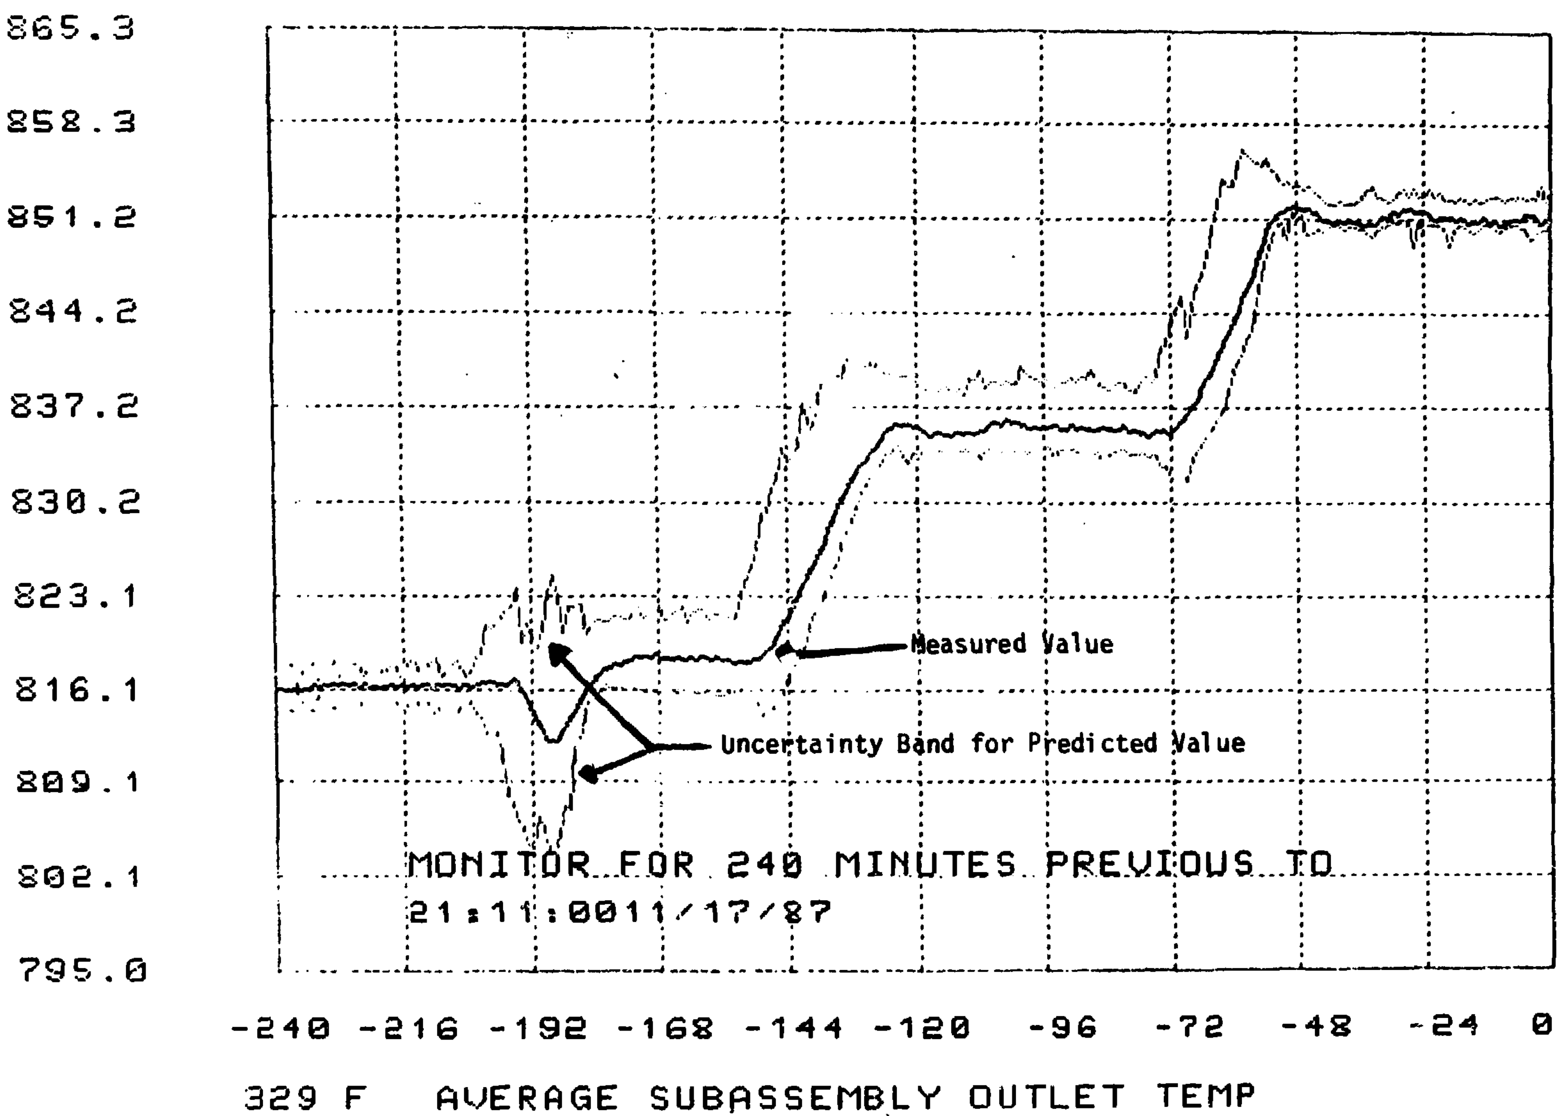
865.3
858.3
851.2
844.2
837.2
830.2
823.1
816.1
809.1
802.1
795.0
Measured Value
Uncertainty Band for Predicted Value
MONITOR FOR 240 MINUTES PREVIOUS TO
21:11:0011/17/87
-240 -216 -192 -168 -144 -120 -96 -72 -48 -24 0
329 F AVERAGE SUBASSEMBLY OUTLET TEMP

Fig. 3

Another very promising approach to sensor validation is that being developed by B. R. Upadhyaya, et al.[5] This approach utilizes a number of signal validation modules to operate on a single set of sensors. The advantage is that methods can be compared one to the other and those that are particularly well suited for a given set of conditions can be applied, whereas other techniques may operate better under other conditions. An on-line expert system can be used to discriminate between techniques and determine the best results. Such an approach may be especially attractive as computer hardware capability continues to improve; it will be further developed and tested with real plant data.

Their system consists of a number of sensor validation modules operating in parallel and incorporating the following items:

1. Generalized consistency checking and sequential probability ratio tests.
2. Process emperical modeling.
3. Multi-variant data driven modeling.
4. Jump, pulse, noise diagnostics modeling.
5. An expert system for qualitative signal validation.

Expert Systems

By far the most common application to plant operation and maintenance is the use of expert systems to guide the decision-making process. The large number of excellent expert-system shells now available has made development of such systems fairly straight-forward and fast.

The most promising applications include maintenance of complex pieces of equipment, complex plant operation (such as startup) and emergency response. Scheduling tasks are also readily adaptable to the use of expert systems, in combination with optimization schemes to maximize work accomplished at minimum cost. Expert systems require experts or well developed instructions and involve a process of putting information into useful form for complex application. A few selected applications will be reviewed.

The first is a system for scheduling of maintenance work, developed by Takayasu Kasahara, Yasuo Nishizawa, Kanji Kato and Takasi Kiguchi of Energy Research Laboratory, Hitachi, Ltd.[6] The system takes as input the list of tasks to be performed, the limitations associated with personnel available to do the work and the plant conditions required for each individual task. As much detail about each of the tasks as possible is provided, including estimated times for completion, data from previous experience for similar jobs, etc. Iterations are performed to optimize scheduling of the work while eliminating interferences between tasks, ensuring the existence of proper plant conditions, and ensuring the existence of necessary personnel to complete the tasks. It is a good example of many similar applications in scheduling which promise to improve maintenance planning. Perhaps the most important capability of these tools is the ability to respond to unexpected problems and to reschedule work as appropriate. As experience grows with these applications, it is expected that their value will become more readily apparent.

The use of expert systems in the operation of specific plant equipment where many individual options exist is another fruitful area of application. A good example is the Fuel Insert Shuffler, developed by Robert Colley and Joseph Naser of EPRI, John Gaiser of Intellicorp and Thomas Brookmire of Virginia Electric Power Company.[7] The function of the fuel insert shuffler is to plan the crane movements of the fuel handling system of a PWR. It evaluates the many combinations of movements possible in handling fuel between the core and the storage pool and develops an optimum plan for such handling. As pointed out by the authors, the system is a good example of the success of work between three important parties in its development; a utility to define the problem and to provide the basic expertise for the task, a vendor who is expert in AI technology, and a research and development organization to provide the funding and direction for the work. It should be noted in fact, that much of the R&D in AI will be successful only when close ties are established between the developer and the user. One virtue of expert systems is that they require this close cooperation, since they are a distillation of information from the source. In the case of the fuel shuffler, information was available in the form of written procedures from which the initial approach could be established. As is so often the case, much of the information regarding strategy and constraints was not written down and could only be developed through interviews with the operators. This is typical of development of such systems and is in fact, one of their virtues. Information that would ordinarily not be apparent or considered unimportant, while in fact it plays an important role, is often brought to light and put into context through the development process.

The third area of useful application of expert system technology is to operating procedures themselves, particularly for complicated operations or for response during emergencies. A good example is the OPA (Operator Advisor), for emergency guidance and monitoring during the course of off-normal reactor operation. Developed as a prototype by Luc Mampaey, TRACTEBEL, S.A., Belgium, it monitors and guides operator actions during accident conditions.[8] It addresses a common problem with written procedures, especially during emergencies. Written procedures often do not fully anticipate the sequence of events involved with emergencies and may be incomplete or even misleading. (To overcome this problem, much attention has been given to preparation of procedures which are more general and key more to observed behavior as evidence of an underlying root cause). Also, during the stress associated with an emergency situation, it is often difficult for the operator to respond well to a set of written instructions, especially when the response is quite complicated. The OPA consists of a "knowledge acquisition unit" which allows the introduction of the procedural steps and related knowledge and a run-time unit which provides advice to the operator during the course of an emergency. The demonstration of the system consists of a real-time interactive simulation of a steam generator tube rupture accident.

Another system which provides procedural help to the operators during the course of an accident is the Procedural Prompting System developed at HEDL for application to the primary pump controllers at FFTF.[9] This system reasons about the system, given a description of its components and the operating constraints which apply. It is also given a set of operating objectives for the system and from this information constructs a set of operating instructions from the conditions and operating objectives as they exist in real-time. Also included is a separate AI program that acts to oversee the logic and actions of the procedure prompter. It has proven to be very effective in demonstration tests and represents the next step from expert systems into automated reasoning.

Given the potential and the success of many of these systems, it is hard to understand why they haven't found wider application in plant operation. It is important, however, to remember that to be effectively applied to an existing power plant, they must be easily developed and applied and confidence must be gained by the user through extensive testing. Techniques which limit the amount of specific modeling and which can fully utilize existing plant operating experience would appear to be the best suited for retrofit to existing plans. Conversely, techniques which depend in larger measure on analytic models and simulation appear to be better suited to new designs.

Display Systems

Traditional approaches to display of data have not succeeded very well in establishing a good understanding by the operator of what is occurring when a large number of signals are involved. However, recent work is changing this. For example, in alarm handling much of the early work was directed toward developing an analytic description of the sequence of alarms and their implications. It was found after much effort, that because of the extensive number of combinations possible, such an attempt was an overwhelming endeavor.[10] Recent work has been considering patterns of alarms as that can be recognized by an operator, and perhaps, ultimately by an AI system for the root cause and sequence of events underway. Again, pattern recognition systems appear to offer significant benefits.

Work at the Halden Reactor Project has been looking at means of display filters utilizing logic techniques to filter and combine alarms, and to graphically display combination of alarms to better depict the part of the plant affected. In this manner it may be possible in an iconic display of the plant to provide information not only of the systems affected but the progress of the disturbance through the plant as, for example, appropriate colors change around an affected region. Such techniques begin to take advantage of the basic human trait of pattern recognition.

A similar approach has been developed by Beltracchi[11] and Lindsay[12], for depicting the course of normal and off-normal operation of EBR-II. The basic idea is to provide an iconic display of the conservation of energy in the system. Aspects which upset the balance of energy conservation may be quickly and easily seen. For example, in this system, the production of heat in the primary coolant system is shown as a rise and fall of temperature around the heat transport loop. The transport of heat to the secondary and steam systems is likewise shown, with the unique aspect being the use of the temperature-entropy diagram to fully characterize the condition of the steam system (see Fig. 4). This approach also involves a number of iconic displays that increasingly detail systems in the plant as one pages down. The idea is being extended to other concepts such as conservation of mass, and appears to have great promise for giving a clear picture of what it is that must be controlled and maintained. Again, these approaches emphasize recognition of patterns that allow operators to quickly understand overall performance. It has been found in the course of extensive testing at EBR-II that not only are patterns more easily evident, but also inconsistencies associated with errors in sensor readings. It is a means of sensor validation through visual methods.

Logic Programming

"Automated reasoning" is the next step beyond some of the early expert systems. It involves careful description of the system, its components, and the logic/performance structure that ties the components together. With the structure thus established, conclusions can be reached regarding performance or constraints of the system. Systems that appear to be especially suited for early demonstration of logic

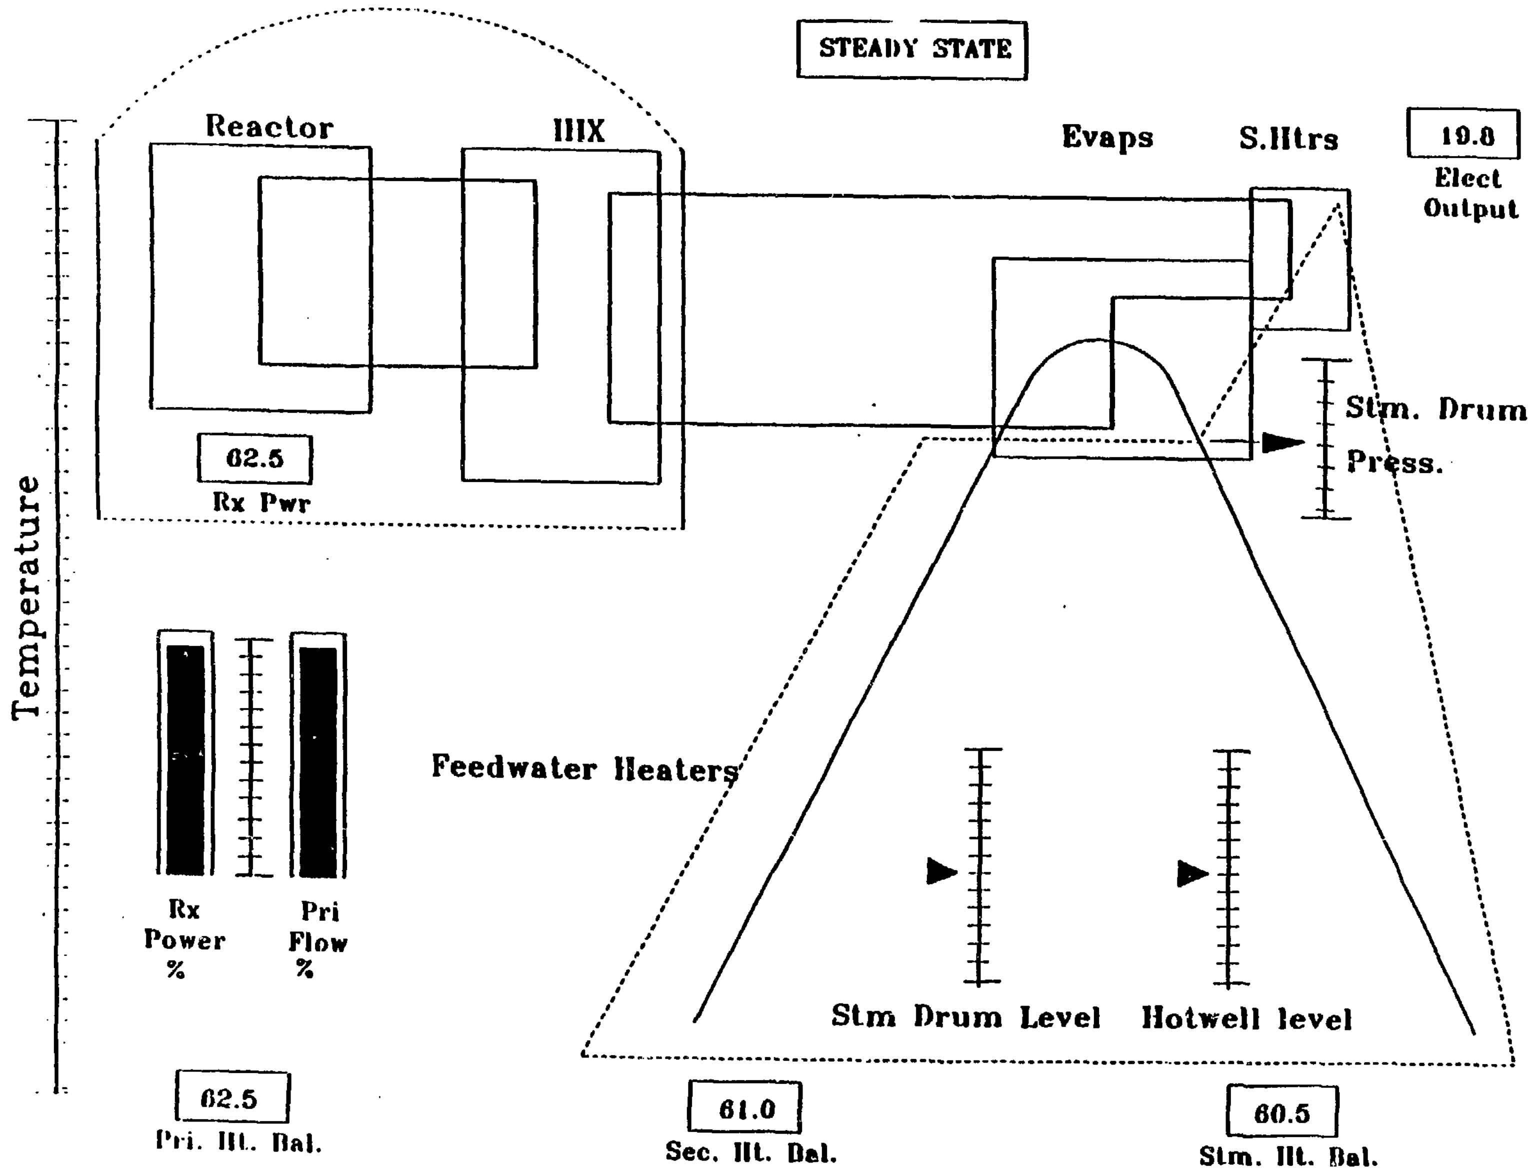

MAIN DISPLAY

Fig. 4

programming are those for ensuring compliance with technical specifications, since they represent a structure of very specific rules linked together in a formal logic structure. The first system discussed is a Prolog-based system for computerizing technical specifications developed by Larry Lidsky of MIT.[13] He points out that the structure of technical specifications, which are essentially a collection of rules which are all interrelated and reflect constraints depending upon plant conditions, are especially well suited to Prolog as a language. In fact, the translation of technical specifications to Prolog is a very easy thing to accomplish. He has also found in development of a system for a typical power reactor that there is feedback both ways. That is, the consistency of technical specficiations themselves may be more easily understood by putting them into a Prolog framework. Other efforts include those of Raghef and Abdelhai from the University of Illinois who have produced a system for tracking of technical specifications.[14] Their approach also uses symbolic programming to describe the logical structure represented by technical specification requirements.

Another similar effort has been that of Steve Epstein from Management Analysis Company in development of the program called TEST[15], which used object oriented programming to develop the package which allows the plant operator to analyze the given system configuration in light of technical specification requirements. This area seems to be especially well suited to AI applications (logic programming) and will likely be one of the first to be implemented with significant payoff.

There are a number of AI based programs directed toward assisting the operator in efficient operation of the plant. Those dealing with plant procedures are especially interesting and again represent fruitful ground for application in reactor power plants. The work at Halden Reactor Project, under way since 1985, is addressing questions such as whether it is feasible to computerize all types of procedures currently in use for power plant control and operation. The system at Halden has been implemented in both Prolog and List, and has shown good success.

Summary

Effective implementation of AI technology in U.S. nuclear power plants will require a well understood framework within which individual systems can be built and proven. The course of evolution that the technology must take for existing plants may be the single most important factor in its success. Emphasis must be given to close interaction with operators and designers of plants. Very effective tools are being developed for sensor validation, expert system guidance for operators, maintenance and accident management. As these technologies are proven, they will be increasingly incorporated into the plant automation.

It is important in the U.S. that a step-wise approach be taken that is easily implemented on existing plants and allows confidence to be gained. To be accepted by the user community, real problems must be addressed in a way that provides clear benefit from AI Technology. As evidenced by the increased number of demonstration projects now "on-line," good progress is being made.

REFERENCES

1. J. I. Sackett, et al, "Inherent Safety and Operability Testing in EBR-II," Proceedings, Topical Meeting on Safety of Next Generation Power Reactors, May 1-5, 1988, Seattle, WA.

2. Dan Corsberg and Larry Johnson, "A Nuclear Reactor Alarm Display System Utilizing AI Techniques for Alarm Filtering," Proceedings, ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

3. A. L. Sudduth, Knowledge Based System in Alarm Filtering, Signal Validation, and Fault Diagnosis," Duke Power Company, Charlotte, North Carolina, Topical Meeting on Safety of Next Generation Power Reactors, May 1-5, 1988, Seattle, WA.

4. J. E. Mott, "A Generalized System State Analyzer for Plant Surveillance," Proceedings, ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

5. B. R. Upadhyaya, et al, "An Integrated Approach for Signal Validation in Nuclear Power Plants," ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

6. T. Kasahara, et al, "Automatic Scheduling of Maintenance Work in Nuclear Power Plants," Proceedings, ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

7. R. Colley et al, "Fuel Insert Shuffler: A Case Study of Expert System Development," Proceedings, ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

8. L. Mampaey, "OPA (Operator Advisor), Emergency Guidance and Monitoring System," Proceedings, ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

9. D. E. Smith and S. E. Seeman, "A Control System Verifier Using Automated Reasoning Software, "Proceedings ANS Topical Meeting on Computer Applications for Nuclear Power Plant Operation and Control, September 1985.

10. Mike Bray, this series.

11. L. Beltracchi, "A Model-Based Display," Proceedings, ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

12. W. S. Schorzman and R. W. Lindsay, "Data Handling of EBR-II for Advanced Diagnostics and Control Waste," EPRI Seminar on Data Acquisition Control and Communication in Power Plants, February 1988, San Diego, California, February 1988.

13. L. M. Lidsky, "The Use of Prolog for Computerized Technical Specifications" Proceedings, ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

14. M. Ragheb and M. Abdelhai, "Production - Rule Analysis System for Nuclear Plant Technical Specifications Tracking," ANS Topical Meeting on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.

15. S. Epstein, "TEST: AI and Technical Specifications," Proceedings, ANS Topical Meetng on AI and Other Innovative Computer Applications in the Nuclear Industry, August 31 - September 2, Snowbird, Utah.