Early AI Diagnostics at Westinghouse

Author:
James Bellows
James Bellows and Associates
Jcbellows2@aol.com

VALUE

When an unplanned (forced) outage occurs at a power plant, its impact is millions of dollars. Detecting an imminent forced outage and converting it to a scheduled outage can potentially reduce the impact significantly, even if the outage is moved from a weekday with higher peak usage to a weekend with lower peak usage by only a few days. Replacement power is typically cheaper on weekends because total demand for power is lower than during the weekdays. Even if the outage cannot be moved, detection of the equipment condition can allow preparation for the outage by having parts and tools available. A predictive diagnostic system could assist in recognizing impending outage and potentially mitigate its consequences.

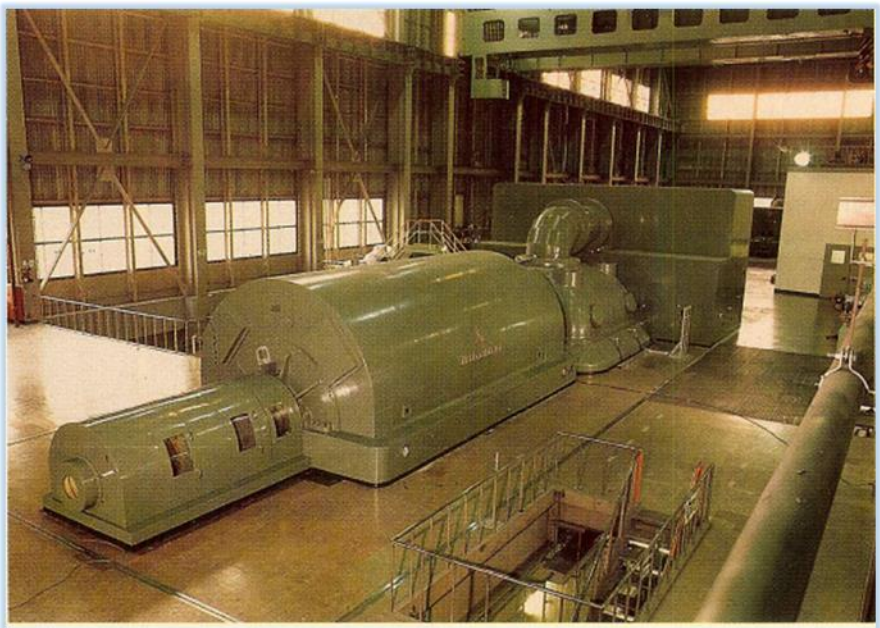

Figure 1: A steam Turbine Generator, ca. 1985. The largest component is the generator, which would be about 10 to 15 feet high. Behind the generator is the low-pressure turbine.[1]

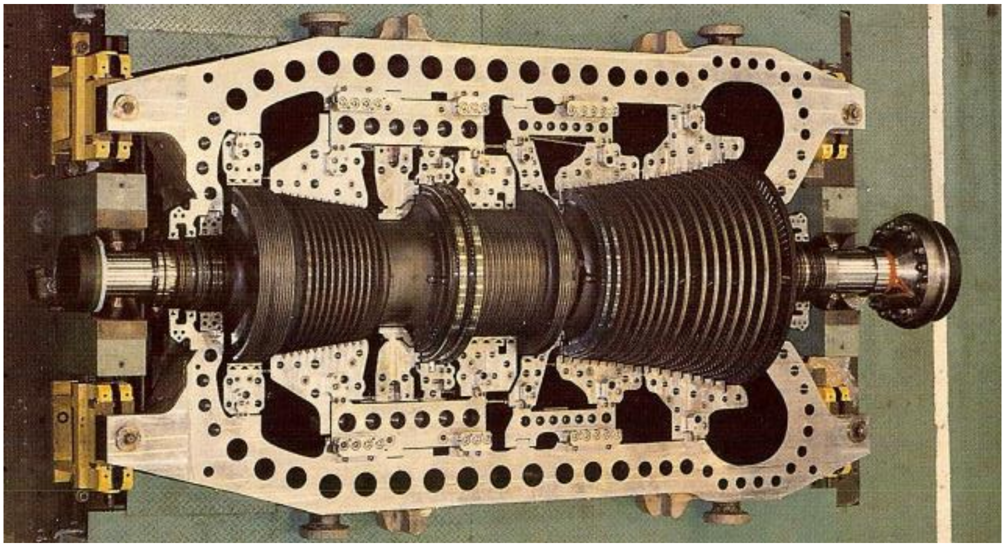

Figure 2: Combined High-Pressure and Intermediate-Pressure steam turbines, ca 1985. As the steam expands, the turbine blades become longer. At the end of the Low-Pressure turbine, they were typically 31 inches long.[1]

THE BEGINNINGS

In the mid 1970's, Robert Osborne, Manager of Controls Development in the Westinghouse Steam Turbine-Generator Division, was convinced that predictive diagnosis was the future of controls development. He supported the development of a predictive diagnostic algorithm based on conditional probabilities.[1] Lacking accurate values for the conditional probabilities, the prototype developers had control experts estimate the probabilities. Thus, as implemented, it was a precursor expert system. It became clear that the conditional probability matrix expanded as the square of the number of conditions diagnosed. This fact precluded very large diagnostic systems based on this method.

One of the problems with developing a diagnostic program was that the incidents to be diagnosed were infrequent. With reasonable maintenance, the equipment was extremely reliable. Forced outage rates (time out of service due to an unplanned outage) were below 1% with 0.05% believed to be attainable.[2] In the entire fleet, a frequent incident might occur several times per year. There were not many examples and even fewer examples with extensive data. The paucity of examples made neural

[1] Osborne, Robert L., Paul H. Haley, and Stephen J. Jennings, "Method and Apparatus for the Automatic Diagnosis of System Malfunctions" US patent 4,402,054, 8/30/1983.

[2] Meador, John T., Steam Turbines No. ANL/CES/TE-78-7. ANL (Argonne National Laboratory (ANL), Argonne, IL (United States)).

nets, which were already available at the time, a poor choice for building a diagnostic system. Neural nets also had the disadvantage that the reason for a conclusion was not explainable. Additionally, at the time, data were largely on strip chart recorders. To convert the strip chart recorder to digital data was too laborious to be practical.

THE DIAGNOSTIC SYSTEM

The solution was an expert system based on the principles of MYCIN.[3] The result was an expert system shell called Process Diagnosis System (PDS). PDS was originally written in LISP[4] by Mark Fox, then of Carnegie Mellon University.[5] It was implemented on Digital Equipment VAX computers. PDS had an interface for the knowledge engineer and inference engine to execute the knowledge base. PDS is described in reference.[6] The design consisted of nodes (ideas) and rules that connected them. The nodes were classified as sensors, hypotheses (intermediate ideas) and malfunctions (final diagnoses). Later a recommendation node type would be added. Sensors were the input, and the data could come from online sensors (through the data center at the power plant) or from manual entry of off-line data. Rules took the confidence in the input node and transferred it to an output node. See Figure 4. The rules included AND and OR functions of several types.

[3] Shortliffe. E.H., MYCIN: A Rule-Based Computer Program for Advising Physicians Regarding Microbial Therapy Selection. PhD dissertation, Stanford University (1974).

[4] Winston, Patric Henry and Berthold Klaus Paul Horn, LISP, Addison Wesley (1981).

[5] E.D. Thompson, E. Frolich, J.C. Bellows, B.E. Basford, E.I. Skiko, and M.S. Fox, “Process Diagnosis System (PDS) — A 30 Year History,” Proc. 27th Conf. on Innovative Applications of Artificial Intelligence, 3928-3933, AAAI Publications (2015)

[6] Bellows, James C., "An Artificial Intelligence Chemistry Diagnostic System," Proc. 45th International Water Conf. 15-23 (1984).

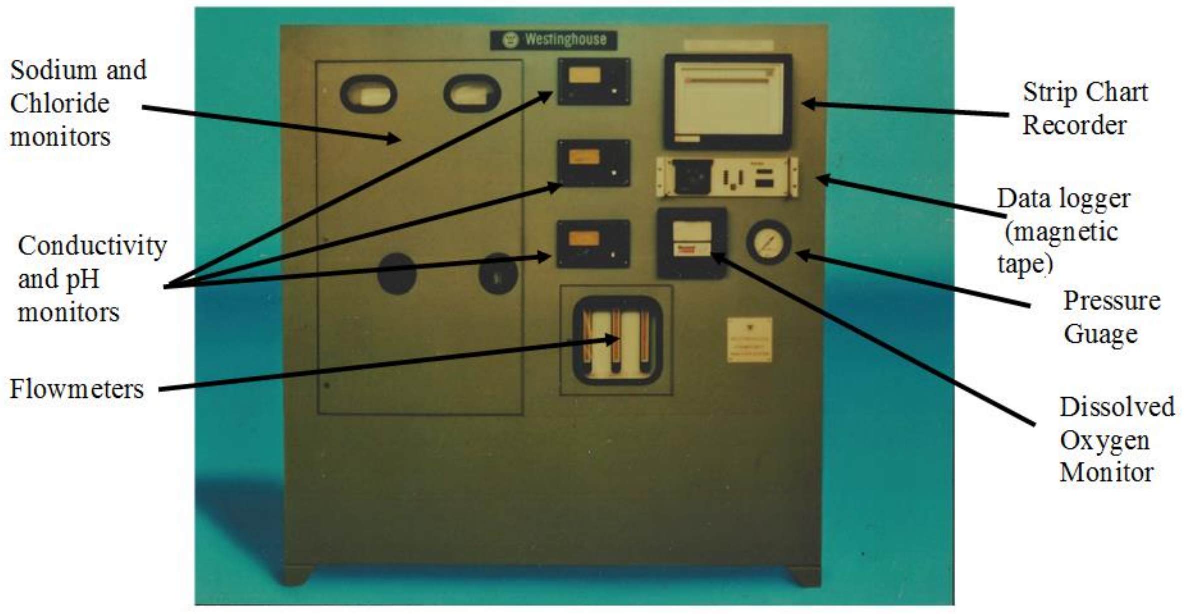

Figure 3: Westinghouse Steam Purity Analyzer System, ca. 1981. The box was about 5 feet high.

When PDS was unveiled in 1982, the Steam Chemistry Engineering section, then managed by the author, had a program of monitoring and was using data loggers in the monitor packages. A typical steam purity analyzer system is shown in Figure 3. It was placed on a sample line drawing steam from the steam generator. The sample was condensed and cooled before it reached the analyzer panel. The analyzer system included chemical monitors, a strip chart recorder and a data logger that was capable of recording a week of data on a cassette tape.[7]

Thus, the chemists had a significant amount of online data in digital form and the testing of PDS started with chemistry. Initially, the sensors were treated as correct and a small system was created.

The next iteration started to diagnose the condition and accuracy of the sensors, which were generally less reliable than the equipment they monitored because they were operating near the detection limits. Chemical sensors also had the possibility of failing with a reasonable reading. A sodium sensor, for example, might stop responding while displaying 1 ppb (µg/kg), which was a reading within the normal range. The

[7] In the 1980s, most power plants did not have steam samples sent to the central sampling system, so a separate Steam Purity Analyzer System was required. The usual samples to the central panel included makeup water, condensate, polisher effluent (if a condensate polisher were present), final feedwater (economizer inlet) and boiler drum water (if the unit had a steam drum).

diagnostic system was required to remove the sensor from consideration. This was done by creating a set of hypotheses that were validated conclusions about the sensor value. Figure 4 illustrates this process for the sodium sensor on the Condensate sample. There was a condensate-sodium-high (Cond-Na-high) hypothesis, based simply on the value of the sodium sensor on the condensate sample, and a validated-condensate-sodium-high hypothesis. Connecting them was a rule that normally transferred all the confidence from the first to the second hypothesis. However, there was a diagnosis of the sodium sensor based on other sodium sensors and other sensors on the condensate sample. If the condensate sodium sensor was deemed degraded, a Parametric Alteration Rule (Paralt Rule) would reduce the transferred confidence. Paralt rules could modify the parameters in another rule. In Figure 4, this is denoted by the valve symbol in the middle of the rule that transfers the confidence from the unvalidated Cond-Na-high hypothesis to the validated-Cond-Na-high hypothesis. This process is explained in detail in the author's paper that first revealed the chemistry diagnostic project to the power plant chemistry community in 1984.[1]

The chemistry monitoring system that supported the diagnostics had multiple sensors of the same type at various locations within the power plant. PDS was augmented to allow writing of a subsystem for each sensor type and instantiating the system multiple times with node and rule labels modified during the instantiation process. The chemistry system, which would eventually become ChemAID® for Chemistry Artificial Intelligence Diagnostics, had four iterations completely starting over before it was implemented at a power plant with over 4000 rules in 1988.[8]

One of the important goals of the diagnostic system was to be able to explain why a diagnosis occurred. The author used relatively long, descriptive names for the hypotheses and developed each diagnosis by solving the problem step-by-step. This approach makes the rulebase clear. One can trace the source of confidence in a diagnosis. These two techniques have the added value that the rulebase remains understandable even after years without maintenance.

[8] Bellows, James C., Karen L. Weaver & James Gallatin, "On-Line Steam Cycle Chemistry Diagnosis," Proc. 2nd Fossil Cycle Chemistry Conf., Seattle, Sept 1988, Electric Power Research Institute Report GS-6166.

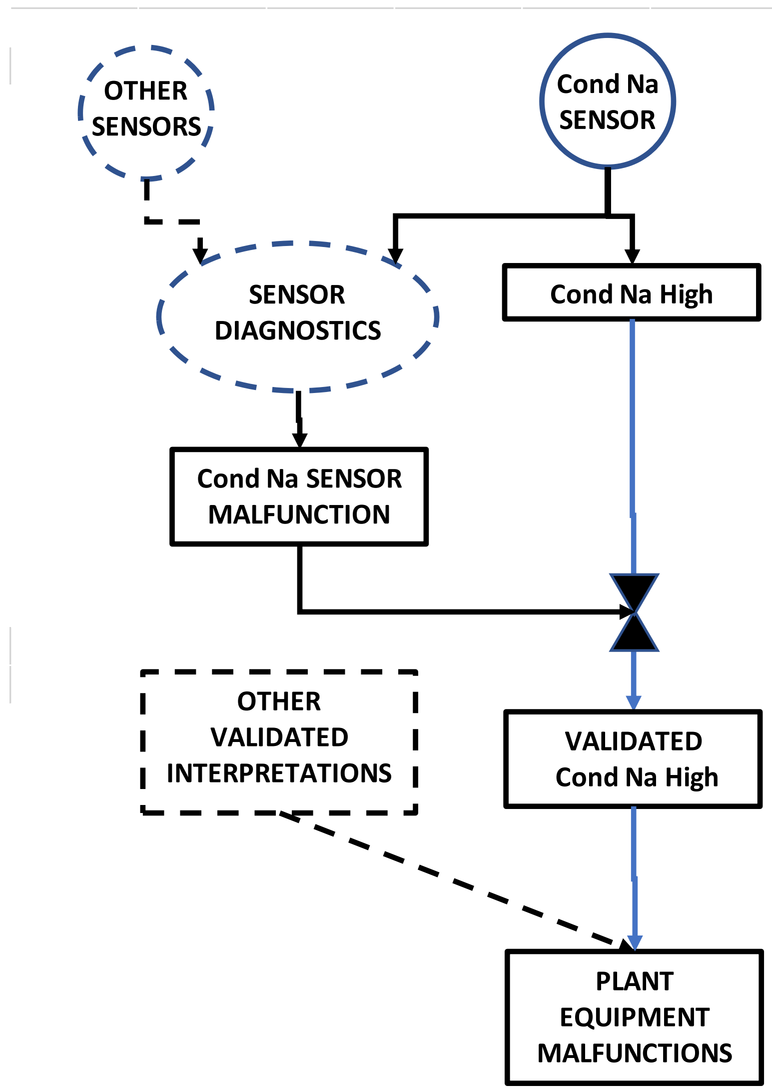

Figure 4: Schematic of diagnosis using the Condensate sodium sensor, showing sensor diagnosis and the method of using a Paralt rule to reduce confidence in the validated Condensate sodium high. Dotted shapes and lines are used to represent many hypotheses and rules.

As the Westinghouse engineers worked with the system, it became clear that the inference engine would be too slow to be practical. It was decided to write an inference engine in C to speed the process. The original inference engine placed a rule on the stack whenever one of its input nodes was updated. Since the diagnostic systems had several layers of hypotheses, this often caused a rule to be fired (executed) before all of its input nodes were updated, and then the rule would be fired multiple times. It was later decided to organize the firing of rules so that all the input to the rule had to have been updated before the rule was fired.

COMMERCIAL OPPORTUNITY

A commercial opportunity came when a customer decided to run some generators at the design limit because a new unit was late. The result was GenAID®, which was implemented on power plants in 1985. Since the chemistry diagnostic development effort had been developing the capabilities of PDS, substantially all the features needed for the generator diagnostic system were already present. GenAID® would ultimately be applied to 7 generators. It would also be the system that was converted to PCPDS (which ran on a personal computer) and became independent of the Westinghouse Diagnostic Center.

REMOTE DATA WITH CENTRALIZED DIAGNOSIS

The diagnosis was performed centrally to facilitate expert modifications of the rule bases and promote the transfer of lessons learned on one plant to all similar plants. Centralized diagnosis was initially required because the power plant control systems could not run rule-based systems. Rule-based systems were needed to allow rapid modification of the diagnostic system as lessons were learned. Data were collected in a data center in the plant and transmitted to the operations center only if significant changes were detected, e.g. the deadband was exceeded. When a datum was recorded, the limits for a significant change were calculated, and no data were recorded unless the datum was outside the limits – then new limits would be calculated. Data were transmitted from the field by leased telephone lines. Centralized diagnosis also provided the ability for Westinghouse to track the diagnoses around the clock and consult their internal experts if required. All the experts were on-call at any time. The knowledge engineers had telephone connections (1200 baud) to the central system from their homes and could work from home almost as easily as from the office. When a diagnosis appeared with significant confidence, the plant was called.

The next step was the development of the TurbinAID® system for the steam turbine,[9] [10] which was implemented at one plant in 1988 and included both mechanical and performance aspects. The performance aspect changed the value from simple outage control to daily performance improvement. TurbinAID® development was still in process when all AI development was terminated for financial reasons. Diagnostics then went through some lean years in the 1990s, but it revived with the advent of gas turbines in combined cycle plants in which the hot exhaust from a gas turbine is used to make steam for a steam turbine.[6] A boom in combined cycle plant construction occurred in the late 1990s and early 2000s. The gas

[9] Fridsma, D.E. and M.A. Moradian, "TurbineAID: Thermal Performance Component", October 21-24, 1990 WFEO/FMOI, Energy Conference, Madrid, Spain. (1990).

[10] Moradian, A., Jenkins, M.A., Chow, M. and Osborne, R.L., "Turbine AID ... Turbine Artificial Intelligence Diagnostics", ISA Tampa, FL. (1991).

turbine diagnostic system, GTAID, is used to follow over 1200 gas turbines.

The development group, at its peak about 1989, included two programmers who maintained the expert system shell, two chemists working on ChemAID, four mechanical engineers working as knowledge engineers on GenAID, four engineers working on TurbinAID and a physicist who acted as an instrument inventor so that when a measurement was needed, there was a way to get it. In addition, specialists from the entire engineering staff of Westinghouse made contributions to individual rules.

LESSONS LEARNED

One important lesson to be found in this narrative is that the audience for the expert system is important to its potential commercial success. GenAID® was written for the operator, who might know that something was wrong, but was usually not able to diagnose the problem. ChemAID® was written for the chemist, who could generally diagnose the chemistry problems. ChemAID® did not survive the lean years of the 1990s.

A second important lesson was the integration of the development team. Including the system shell programmers in the team meant that a request for a new feature could be accommodated in a few days. Because roadblocks due to limitations of PDS could be removed quickly, the knowledge engineers could continue working, almost without interruption.

A third important lesson for the construction of such systems was the step by step approach. This approach promotes transparency, and transparency is useful in explaining a diagnosis. It is also valuable when a section of the rulebase must be modified. Modification of complex rules invites unintended consequences. Modification of simple, one-step rules is much less likely to have unintended consequences.

The author would like to acknowledge helpful conversations with Eric Harper, one of the PDS programmers, who first suggested writing this paper. The author would also like to thank all the people who contributed to the development of the Westinghouse AI Diagnostics systems for a very rewarding and stimulating decade during the 1980s.

➢ Return to IIC Journal of Innovation landing page for more articles and past editions