CONF-841017--7

NOTICE
PORTIONS OF THIS REPORT ARE ILLEGIBLE.
It has been reproduced from the best available copy to permit the broadest possible availability.

CONF-841017--7

DE85 002469

By acceptance of this article, the publisher or recipient acknowledges the U.S. Government's right to retain a nonexclusive, royalty-free license in and to any copyright covering the article.

LONG-TERM AUTOMATED SURVEILLANCE OF A COMMERCIAL NUCLEAR POWER PLANT¹

C. M. Smith* and R. C. Gonzalez**

*Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831

**The University of Tennessee, Knoxville, Tennessee 37916

ABSTRACT

A pattern recognition system for monitoring nuclear reactor signals is presented. The system is based on detecting deviations from baseline signatures learned automatically during normal plant operation. The capabilities and limitations of the recognition approach were investigated during a 2 1/2-year continuous, on-line series of experiments at the Sequoyah-1 nuclear power plant.

KEYWORDS

Pattern, recognition, nuclear, reactor, surveillance, noise, analysis, monitoring.

INTRODUCTION

Interest in procedures for performing automatic monitoring functions in nuclear power plants stems from a need to provide assistance to plant operators in assimilating quickly the large number of interrelated signals that provide insight into the operational status of the plant.

The signal recognition problem in this field can be broadly divided into two categories: surveillance and diagnosis. The surveillance problem is one of classifying the status of the components being monitored into two classes: normal or abnormal. The diagnosis problem is concerned with identifying the source and degree of a detected abnormality (Gonzalez and Howington, 1977).

Unlike many pattern recognition applications where representative patterns from each class of interest are available, it is difficult to obtain patterns characterizing abnormal behavior in the components of a nuclear plant. Simulation of abnormal behavior by planned failure of components in a plant is usually not a feasible alternative for generating the desired abnormal patterns because of governmental regulations and implementation costs. The design of a pattern recognition system for surveillance is thus reduced to characterizing normal behavior and establishing limits beyond which the operation of a component is classified as abnormal. This approach is based on the assumption that the data which are labeled as normal are derived from plant components whose operation is within design specifications.

One of the principal features of a pattern-recognition system designed to aid plant operators is that the system output should contain information beyond a simple "normal" or "abnormal" message. This information can be in the form of graphs, tables, or numerical data, but, in the presence of an abnormality, it must contain reference to the characteristics and limits of normal behavior. This information aids the operator in detecting trends and also in evaluating the relative importance of an alarm.

The system described in this paper is oriented in this direction. It is based on a set of eight discriminants designed to detect specific features in power spectral density (PSD) estimates (Piety, 1977; Piety and Robinson, 1976). One of the principal features of the system is its ability to learn the limits that characterize the behavior of the signals being monitored and then verify the accuracy of these limits via an extended learning period. Another important feature of the system is its capability to monitor up to 20 reactor signals and to output a variety of summary tables and plots which have proved to be of significant value to signal analysts whose aim is to gain insight into the operational status of the plant components being monitored.

The performance of the pattern recognition system was evaluated through a 2 1/2-year series of continuous on-line experiments at the Tennessee Valley Authority's Sequoyah-1 nuclear power plant. The principal purposes of these experiments were (1) to investigate the data reduction aspects of

¹Research sponsored by the U.S. Nuclear Regulatory Commission, Office of Nuclear Regulatory Research under Interagency Agreement No. 40-551-75 and performed at Oak Ridge National Laboratory, operated by Martin Marietta Energy Systems, Inc., for the U.S. Department of Energy under Contract DE-AC05-84OR21400.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

gmk

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

the system in the sense that it stored only those signals that were flagged as being different from their corresponding baseline signatures; (2) to evaluate system behavior over a long period of operation; and (3) to determine the applicability of the discriminants to various types of signals commonly monitored in a nuclear plant. The results of these experiments indicate that significant data reduction is possible. In our case, the reduction ranged from 75 to 90%, depending on the signals monitored. The results on learning baseline behavior were encouraging in the sense that little change in the operational limits was observed after the initial learning period, thus indicating stability in these limits and lending an added degree of credibility to the signatures that were flagged by the system as being suspect.

The discriminants followed our theoretical predictions when applied to random (e.g., neutron) signals. However, considerable variation from expected behavior was observed with signals that contained strong nonstationary components, thus indicating that monitoring these signals requires selective filtering to remove these components.

STRUCTURE OF THE RECOGNITION SYSTEM

System Elements

Figure 1 shows a block diagram of the power spectral density recognition system (PSDREC) described in this paper. As indicated in the previous section, the system is capable of monitoring up to 20 signals. The data acquisition stage consists of amplifying, filtering to reduce the effect of aliasing, and sampling to produce the desired digital data. The PSD computation stage performs the required fast-Fourier transformer operations, computes the PSD of each signal, and carries out the averaging needed to produce smoothed PSD estimates.

The discriminant computation stage yields the eight discriminants used by the system to perform recognition. The learning stage establishes the limits of normal operation and tracks the parameters needed to compute the discriminants. Finally, the recognition stage determines whether a given PSD observation is within the limits of normal operation and performs all operations associated with output data formatting. The results of recognition can be displayed, stored, or transmitted to a remote site via a communication link.

Discriminants

The discriminants (features) used by the recognition system are based on ratios between a test PSD, $P(f_i)$, and an appropriate baseline PSD, $Q(f_i)$, where f_i is the i th frequency component, and i is in the range $1 < i < N$. A logarithmic transformation is applied to these PSD ratios in order to make all results compatible with PSD displays which, because of dynamic range considerations, are usually plotted on a logarithm scale. All logarithms used in this paper are $\log_{10}$, and are denoted simply as $\log$. It is also important to note that the discriminants outlined in the following discussion were defined empirically to highlight specific types of changes between a test PSD and a given baseline PSD (Piety, 1977).

The first discriminant, D_I , is sensitive to differences between the overall integral power in the test and baseline PSDs. Discriminants D_{II} and D_{III} are local measures of deviation, and discriminant D_{IV} is the average difference between these two PSDs on a $\log_{10}$ scale. Discriminant D_V is a measure of the variance of a set of log ratios. The sixth discriminant, D_{VI} , measures the number of log ratios above or below the median. Finally, discriminants D_{VII} and D_{VIII} are based on the number and length of sequences of consecutive log ratio runs above or below the median. The mathematical details of these discriminants are given by Piety (1977).

Limit Selection

The discriminants discussed in the previous section were designed to detect specific characteristics in the ratios of PSDs. Consequently, the behavior of these discriminants is examined on an individual basis by the recognition system. The approach is to place bounds on normal discriminant behavior by establishing operational limits. These limits are selected initially based on statistical assumptions about the input data and then modified as the system gains experience from monitoring the plant.

Two types of limits are used by the system: alert limits and alarm limits. Both limits are computed according to the same basic methods but use different statistical criteria. For example, a typical value used for the alert limits of D_{IV} and D_V is the percentile corresponding to five standard deviations, while the percentile corresponding to ten standard deviations is used for the alarm limits. Adjustments of the initial limits after the learning period are discussed in the following section.

LEARNING AND MONITORING METHODOLOGY

One of the principal features of PSDREC is a comprehensive mechanism for managing data obtained during the learning and monitoring phases of system operation. In this context, monitoring refers to the set of actions taken by PSDREC based on whether or not the observed discriminants are within the specified limits. Learning is the subphase of monitoring during which the system is allowed to adjust the limiting criteria based on the behavior of the discriminants over some specified time interval.

Terminology

In its present form, PSDREC can monitor a maximum of four signals simultaneously, each referred to as an "inner loop." The process of monitoring any set of up to four signals simultaneously is referred to as an "outer loop." In order to monitor more than four signals, PSDREC sequentially monitors several outer loops. A "run" denotes the sequential monitoring of several (up to 20) outer loops.

The operation of PSDREC is based on four types of PSDs computed for each member of an inner loop. These four types of PSDs, referred to as PSDZ, PSDB, PSDT, and PSDC differ only in the number of data blocks used in their computation.

PSDZ is the short-term trend spectrum. Of the four types of PSDs, the PSDZ spectrum is formed over the shortest time span (i.e., the smallest number of data blocks). Changes in consecutive PSDZ spectra, called "short-term trend fluctuations," detect rapidly occurring signal changes. The PSDZ spectrum is used as the basic building block from which the other three spectra are formed by ensemble averaging different numbers of PSDZ spectra.

PSDB is the baseline spectrum, formed over the longest time span. It is formed only once for each operational state, and it is the absolute reference spectrum used to determine how much the short- and long-term trend spectra have varied.

PSDT is the long-term reference trend spectrum. It is formed over a span several times longer than a PSDZ spectrum but usually shorter than a baseline spectrum. Like PSDB, this spectrum is a reference spectrum for the long-term trend, but it is a relative rather than an absolute reference. Whereas PSDB is absolute and thus never changes, the PSDT spectrum can be updated. This updating process allows PSDT to reflect slow changes in the signal that are deterministic over the time span of several PSDT spectra but that have less than the normal statistical fluctuation for a single PSDT spectrum. PSDC is the long-term candidate trend spectrum. It is formed over the same time interval as PSDT and is related to the updating of this spectrum.

The relationship between the PSDB, PSDT, and PSDC spectra and the number of PSDZ spectra that compose them is explained in Fig. 2, which shows the PSDZ as the average of "Z" data blocks, the PSDB as the average of "B" PSDZs, the PSDC as the average of "C" PSDZs, and the PSDT as the average of "T" PSDZs. Also note that each PSDZ is an integer number ($= Z$) of data blocks; each PSDC is an integer number ($= C$) of PSDZs; each PSDT is an integer number ($= T$) of PSDZs; C equals T; each PSDB is an integer number ($= B$) of PSDZs; and each PSDB is an integer number of PSDCs. A trend interval is the number of data blocks (and thus the number of PSDZs) between the formation of consecutive PSDCs or PSDTs. One PSDC is formed per trend interval.

Fig. 1. Diagram of surveillance system.

Fig. 2. Interrelations of PSDs in PSDREC.

As indicated earlier, all discriminants are based on PSD ratios. PSDREC computes three ratios for this purpose: $\frac{PSDZ}{PSDT}$, $\frac{PSDC}{PSDT}$, and $\frac{PSDC}{PSDB}$. The function of each of these ratios is explained below.

The $\frac{PSDZ}{PSDT}$ ratio is used to monitor short-term trend fluctuations, which are indicative of rapid changes in the monitored signal. Monitoring for short-term trend fluctuations occurs each time a new PSDZ is formed after the long-term reference trend (PSDT) has been established.

The $\frac{PSDC}{PSDT}$ ratio is used to monitor long-term trend fluctuations, which are indicative of slowly occurring signal changes. The monitoring of the $\frac{PSDC}{PSDT}$ ratio is part of the methodology by which the PSDT spectrum is kept current by updating. Monitoring for long-term trend fluctuations and long-term reference updating occurs whenever a new PSDC spectrum is formed.

The $\frac{PSDC}{PSDB}$ ratio is used to monitor the movement of the long-term trend from the absolute reference baseline. Since the long-term trend varies from the baseline just as the short-term trend varies from the long-term trend, the movement of the long-term trend from the baseline must be tracked in order to determine the total movement of all trends. Monitoring for movement of the long term trend with respect to the absolute reference baseline also occurs whenever a new PSDC is formed.

Learning

The adaptive capability of PSDREC is achieved through a learning period, during which discriminants are tested against the predicted limits in order to determine how closely the data follow the stated assumptions. The learning period normally occurs only at the beginning of the execution of a PSDREC run. The length of the learning period is specified in terms of an integer multiple of the number of blocks in the trend interval. It is assumed that the plant is in steady state and operating normally during learning.

The discriminant limits are initially established using the approach discussed in the previous section. At the end of learning, the limits are updated to reflect discrepancies between the theoretical predictions and actual data behavior. The updating procedure for the first five discriminates is based on normalizing these discriminants by the observed means and standard deviations. Let $M(D_1)$, $\sigma(D_1)$, $M(D_1)$, and $S(D_1)$ be the theoretical and computed means and standard deviations, respectively, of discriminant D_1 . After learning is completed, the upper and lower limits for D_1 are updated according to the following relations:

$$t_{IU} = [T_{IU} - m(D_1)][S(D_1)/\sigma(D_1)] + M(D_1)$$

$$t_{IL} = [T_{IL} - m(D_1)][S(D_1)/\sigma(D_1)] + M(D_1)$$

where T_{IU} and T_{LU} are the theoretical upper and lower limits. Clearly, t_{IU} and t_{IL} reduce to these limits when $S(D_1) = \sigma(D_1)$ and $M(D_1) = m(D_1)$.

The second and third discriminants are compared against single limits which are updated as above:

$$t_i = [T_i - m(D_1)][S(D_1)/\sigma(D_1)] + M(D_1)$$

for $i = II, III$.

Discriminants D_{IV} and D_V are Gaussian variables according to the central limit theorem. Accordingly, the limits for these discriminants are adjusted simply by using the computed means and variances in place of their theoretical counterparts. In other words, we compare the standardized variables $z_i = [D_i - M(D_1)]/S(D_1)$, $i = IV, V$, against the limits determined by specifying a desired percentile. Similar comments apply to D_{VI} and D_{VII} , with the exception that their means are updated relative to a proportion of the predicted median. The theoretical standard deviations for these two discriminants are retained because they are conservative estimates. Finally, discriminant D_{VIII} is updated by using an iterative procedure discussed by Plety (1977).

The learning period in PSDREC is divided in two parts: initial learning and extended learning. Initial learning specifies a period at the end of which the limits are updated, as discussed above. All computed parameters (e.g., means and standard deviations) are stored at the end of this period. Extended learning starts with these parameters and updates them based on more observations. The extended learning period is used as a check on the stability of the parameters determined during initial learning. The limits are updated in exactly the same manner as discussed above but use the extended parameters.

Monitoring

Initially, the operator defines the signal-specific information for monitoring each signal (inner loop information), the analysis parameter information for each group of signals (outer loop information), and the procedural information for analyzing the groups of signals (run information).

Specifically, the operator determines the following for each signal being analyzed:

1. How many data blocks form a baseline spectrum (PSDB).
2. How many data blocks form a reference long-term trend spectrum (PSDT).
3. How many data blocks form a candidate long-term trend spectrum (PSDC).
4. How much longer a trend interval is than a long-term candidate trend.
5. How many data blocks form a short-term trend spectrum (PSDZ).
6. How many data blocks will be analyzed during initial learning.
7. How many more data blocks will be analyzed during extended learning.
8. What the bounding conditions that determine an operational state are.

IMPLEMENTATION

The PSDREC system is implemented in a Digital Equipment Corporation (DEC) PDP 11/34A minicomputer that executes the surveillance algorithms and controls all major system components. The computer runs under RT-11 Version 4.0 and has a floating-point processor (FPP). It is equipped with 96 Kwords of semiconductor memory and two RK05 type disk drives housing three magnetic disk storage surfaces which provide 7.5 megabytes of storage capacity. A standard DEC LA36 terminal provides input and output capabilities for controlling the system, and a Tektronix 4006 terminal provides plotting capability.

Signal conditioning amplifiers and filters for the PSDREC surveillance system consists of 16 computer-controlled amplifiers (including sample/hold circuitry), 16 computer-controlled Rockland Model 816 analog low-pass programmable filters for anti-aliasing, 1 DEC DR11-K parallel input/output device to control all amplifiers and related functions, and a second DR11-K to control all filters. Signal sampling is accomplished through the use of an ADAC Model 600/11AD 12-bit analog-to-digital converter (ADC), a DEC KW11-P programmable clock, and the sample/hold circuitry located in the signal amplifiers. The ADAC ADC performs the analog-to-digital conversion at a maximum rate of 100,000 samples/s.

RESULTS

The system described in the previous section was installed at Tennessee Valley Authority's Sequoyah-1 nuclear power plant on September 1, 1980, and operated continuously through March 31, 1983. The major purposes of this experiment were

1. To investigate the data reduction aspects of the system by monitoring the number of observations flagged as suspect versus the total number of observations performed by PSDREC over multiple fuel cycles,
2. To evaluate the learning behavior of the system over a long period of operation, and
3. To evaluate the applicability of the discriminates to various types of signals.

Data Reduction

At Sequoyah-1, the surveillance system monitored the 20 reactor signals listed in Table 1 over the frequency range 0.02 to 49 Hz. During the first fuel cycle, the system flagged only a few suspect signatures, greatly reducing the number of signatures requiring examination by a noise analyst. Figure 3 summarizes the benefits in terms of the number of observed and suspect signatures for each signal. For nine of the signals, approximately 10% of the signatures observed were suspect; for six others approximately 20% were suspect; for the remaining five signals (reactor coolant system pressure, two core exit temperatures, steam generator feedwater flow, and reactor vessel head acceleration), >20% of the observed signals were suspect.

An investigation into the causes for the large percentage of suspect signatures in the last grouping has been conducted. Results from the study of the reactor coolant system pressure signal revealed that its signature is influenced by the presence of several other sensors that share the same pressure tap line. Studies of the core exit temperature signals show that one is corrupted by intermittent noise pickup, possibly electrical interference, and the other displays numerous short-term changes in its overall signature, which are yet to be explained. The steam generator feedwater flow signature has two major resonances whose frequencies vary significantly with time, thus causing the large number of suspect signatures. The sources of these resonances and the cause

of their nonstationarity has not yet been determined. The reactor vessel head accelerometer signature possesses a significant resonant structure, and the amplitudes and widths of the resonances display large fluctuations which the surveillance algorithm does not adequately handle when limits are set for one particular group of discriminants. While monitoring of these signals proved troublesome, the surveillance algorithm performed well in monitoring all neutron signatures. For the neutron signals, many of the suspect signatures represented either a temporary loss of signal (i.e., to calibrate the chamber) or a change in an identifiable noise source.

Fig. 3. Comparison of the number of observed and suspect signatures for each monitored signal at Sequoyah-1.

TABLE 1 The Reactor Signals Monitored by the Surveillance System at Sequoyah-1

Signal description	Number
Lower power range ex-core neutron chambers	4
Average power range ex-core neutron chambers	2
Intermediate range ex-core neutron chamber	1
Reactor coolant system pressure	1
Reactor coolant system flow	1
Pressurized level	1
Reactor coolant system hot leg temperature	1
Reactor coolant system cold leg temperature	1
Core exit thermocouples	3
Steam generator steam flow	1
Steam generator feedwater flow	1
Steam generator water level	1
Steam generator pressure	1
Reactor head accelerometer	1

Learning Behavior

PSDREC initially sets the monitoring limits based on statistical assumptions about each discriminant and then adjusts these limits after learning. Typical limit adjustments for these discriminants are illustrated in Fig. 4. The bounding bars indicate the alert limits used by the system during

Fig. 4. Discriminant 1 values and limits for the second fuel cycle of Sequoyah-1
(a) end of learning (b) end of extended learning.

each of the learning periods labeled in the plots. We observed that in most cases the limit values were increased for all discriminants, indicating that the theoretical limits were on the low side. The amount of limit increase was proportional to the discrepancy between the predicted and actual data behavior. It was also observed in our experiments that there tended to be little change in the limits at the end of extended learning.

Another important parameter addressed by the experiment at Sequoyah is the number of individual PSDs (denoted by I) that are averaged to form a working PSD, and the number of these PSDs (called "data blocks" and denoted by B) that are averaged to obtain PSDB, PSDT, PSDC, and PSDZ. One of the principal factors in selecting I is the large number of discriminants that can be obtained in a given time period. For example, PSDREC was able to cycle through all 20 signals in approximately 5 h with I = 20. This yielded on the order of five sets of discriminants per day for each of the monitored signals. On the other hand, I = 100 yielded approximately one set of discriminant a day per signal. We found that increasing I (and proportionally, B), tended to produce tighter limits at the end of initial learning. This is not unexpected since longer averaging clearly reduces statistical variations.

PSDREC was set up at Sequoyah in a dual operating mode where the system alternated between I = 20 and I=100 in order to test learning behavior between one and five discriminant observations per day. We found no appreciable difference in system performance other than a slight increase in the number of flagged discriminants for the larger value of I as a result of the tighter limits. We believe that this type of experimentation requires further refinement by focusing especially on fewer signals to allow faster cycle time through these signals and to reduce the massive number of resulting observations.

Table 2 lists the values of B inputs used to set up PSDREC (see section on Monitoring) at Sequoyah as a function of I. The last entry indicates the allowed percentage variation before a new operational state is declared. For example, if PSDREC starts monitoring the plant at 90% power, a range from 81 to 99% power variation is allowed before a change in operational state is declared by the system.

TABLE 2 Power Spectral Density Recognition System Setup

List item in Monitoring Section	I = 20*	I = 100*
1	500	5,000
2	100	500
3	100	500
4	100	500
5	20	100
6	1,000	10,000
7	1,000	10,000
8	±10% variation	

*Entries correspond to values of B.

CONCLUSIONS

To our knowledge, the work reported in this paper represents the most comprehensive continuous, on-line experiment carried out at a commercial nuclear power plant to date. When PSDREC was installed at Sequoyah-1 on September 1, 1980, the principal objectives were to establish a baseline signature for each signal and to evaluate the capability to automatically determine deviations from those baselines.

As indicated in the previous section, the system was able to learn and track the behavior of random, stationary signals with a high degree of success. Nonstationary signals were considerably more difficult to track, as evidenced by the proportionally higher rate of flagged PSDs associated with these signals.

Initially, PSDREC was set up to analyze each signal over the entire frequency range of its sensor. Our experience at Sequoyah-1 indicates that a considerable amount of effort is required to establish analysis parameters, such as frequency ranges and monitoring periods, for each signal prior to the onset of automatic surveillance operation.

REFERENCES

- Gonzalez, R. C. and L. C. Howington (1977). Machine recognition of abnormal behavior in nuclear reactors. IEEE Trans. Syst., Man & Cybern. (USA), SMC-7(10), 717-728.
- Piety, K. R. (1977). Statistical algorithm for automated signature analysis of power density data. Proc. Second Specialists Meeting on Reactor Noise (SMORN II), Pergamon Press, NY, 781-802.
- Piety, K. R. and J. C. Robinson (1976). An on-line reactor surveillance algorithm based on multivariate analysis of noise. Nucl. Sci. & Eng., 59, 369-380.