

*Automated Software Analysis
of Nuclear Core Discharge Data*

*Ted W. Larson
James K. Halbig
Jo Ann Howell
George W. Eccleston
Shirley F. Klosterbuer*



Se

CONTENTS

ABSTRACT	1
INTRODUCTION	1
AUTOMATED SOFTWARE ANALYSIS	5
IDENTIFICATION OF AREAS OF INTEREST AND FUEL DISCHARGE EVENTS	6
CORRELATION OF EVENTS	7
MONITORING REACTOR POWER LEVEL.....	7
STATISTICAL PHENOMENA OF CDM DATA	8
NEURAL NETWORKS FOR SOLVING THE FUEL GEOMETRY PROBLEM.....	14
NEURAL NETWORKS FOR FUEL BURNUP COMPUTATION	18
RESULTS OF NEURAL NETWORK MODELS FOR SOLVING GEOMETRY AND BURNUP.....	21
CONCLUSIONS.....	22
REFERENCES.....	22

AUTOMATED SOFTWARE ANALYSIS OF NUCLEAR CORE DISCHARGE DATA

by

**Ted W. Larson, James K. Halbig, Jo Ann Howell,
George W. Eccleston, and Shirley F. Klosterbuer**

ABSTRACT

Monitoring the fueling process of an on-load nuclear reactor is a full-time job for nuclear safeguarding agencies. Nuclear core discharge monitors (CDMs) can provide continuous, unattended recording of the reactor's fueling activity for later, qualitative review by a safeguards inspector. A quantitative analysis of this collected data could prove to be a great asset to inspectors because more information can be extracted from the data and the analysis time can be reduced considerably. This paper presents a prototype for an automated software analysis system capable of identifying when fuel bundle pushes occurred and monitoring the power level of the reactor. Neural network models were developed for calculating the region on the reactor face from which the fuel was discharged and predicting the burnup. These models were created and tested using actual data collected from a CDM system at an on-load reactor facility. Collectively, these automated quantitative analysis programs could help safeguarding agencies to gain a better perspective on the complete picture of the fueling activity of an on-load nuclear reactor. This type of system can provide a cost-effective solution for automated monitoring of on-load reactors significantly reducing time and effort.

INTRODUCTION

Nuclear power stations in the United States contain reactor cores, which can only be accessed from one end, usually the top; fuel can only be accessed when the reactor is shut down. One safeguard advantage to this type of reactor is that it is relatively easy for a nuclear safeguarding agency to monitor the fueling process. An inspector from a safeguarding agency can be sent to the site to oversee the fueling procedure. On-load nuclear reactors are different from United States reactors in that operators may obtain access to the core from both ends, and they can be continuously fueled without shutting them down. Such an operation offers an interesting challenge from a safeguards perspective because it provides a greater opportunity for the diversion of nuclear material. On-load reactors are well-suited for producing plutonium from their standard fuel bundles. Safeguarding an on-load reactor requires keeping track of fuel as it is pushed through the reactor. When a fresh fuel bundle is pushed in one side of the reactor, a spent fuel bundle is simultaneously discharged into a collection mechanism on the other side. Using this fueling scheme, a typical on-load reactor will

discharge 55 to 65 fuel bundles per week. Figure 1 shows a conceptual diagram of this fueling cycle. Because this is an ongoing process, it is not simple for a safeguarding agency to continuously monitor fueling; it is rather labor intensive to have a safeguards inspector on-site all of the time.

There are several approaches to the problem of monitoring the constant fueling process of on-load reactors. Obviously, some type of a measurement system for providing continuous, unattended monitoring of fueling is the most attractive alternative because it is the least labor-intensive. Core discharge monitors (CDMs)¹ provide a convenient system for safeguarding on-load reactors because they can be installed as an independent, tamper-resistant package. A CDM uses radiation detectors to monitor the movement of fuel between the reactor core and the fuel storage area. Core discharge monitoring can be performed by an electronics package called GRAND (gamma ray and neutron detector) and associated detectors developed at the Los Alamos National Laboratory (Los Alamos). A typical measurement station consists of a GRAND, which is the data acquisition electronics, and four detectors mounted in one enclosure. Each GRAND continuously collects data from its associated detector at discrete time intervals and transmits the data to an MS-DOS compatible computer for recording. The four detectors in each enclosure include two neutron detectors (fission chambers) and two gamma-ray detectors (ion chambers), one shielded and one unshielded. Because of low-enriched fuel and low exposure of the spent fuel, relatively few neutrons are emitted by the spent fuel. The neutron detectors are encased in a container of heavy water. In the ideal situation, the neutron detectors are sensitive to neutrons created by (γ,n) reactions in the deuterium surrounding the detectors. To produce a neutron in the (γ,n) reaction requires a gamma-ray energy threshold of 2.2 MeV. To monitor core discharges, we mounted four GRANDs around the nuclear core, two on each reactor face. For purposes of discussion, Fig. 1 shows an example of how a typical on-load

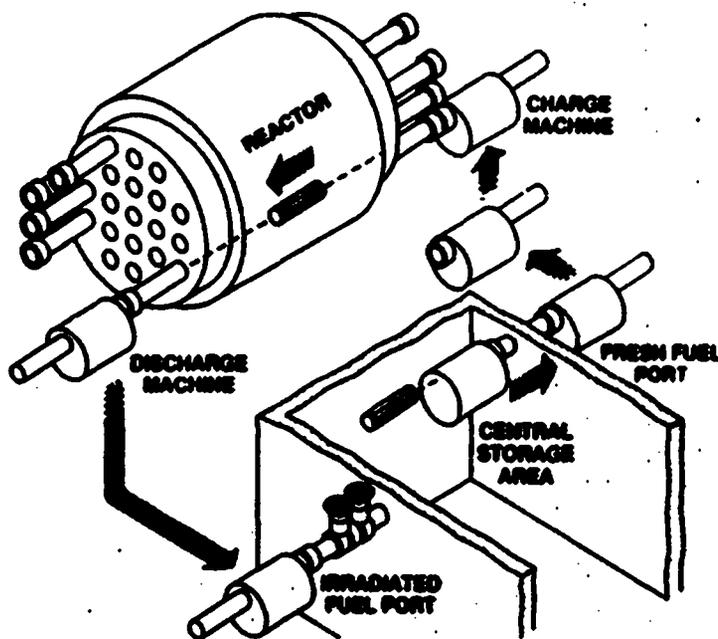


Fig. 1. Conceptual diagram of fueling cycle.

reactor might be laid out with the detectors and GRANDs installed. In Fig. 2(a), the core is mounted in the building such that the core faces are on the east and west side of the reactor. So, fueling takes place from east to west, or west to east. Because the core is mounted in this fashion, each measurement station is designated by its location in relationship to the core, either the southeast (SE), northeast (NE), southwest (SW), or northwest (NW) corner.

Each GRAND collects nuclear radiation data from the detector enclosure, adds 10 time stamps to it, and temporarily stores it. The data are then fed to the collection computer upon request for more permanent storage. At a later time, data can be off-loaded from the collection computer for off-line review. The detector data fed from the GRAND consist of five channels of information. The channels are labeled as follows: fission chamber A, fission chamber B, fission chamber C, ion chamber 1, and ion chamber 2. Fission chamber A corresponds to the first neutron detector in the detector enclosure. Fission chamber B is another view of the first neutron detector, which can be used for possible tamper detection. The second neutron detector in each detector enclosure is labeled as fission chamber C. This neutron detector is not wired to its corresponding GRAND, but rather to the GRAND on the opposing face. For example, the NE fission chamber C is wired into the NW GRAND, and the NW fission chamber C is wired into the NE GRAND. This provides the overall system with a backup, in case the GRAND for one of the detectors fails. This cross wiring is shown in Fig. 2(a) as the splice box between the two GRANDs on each side of the reactor core. Finally, the two gamma-ray detectors correspond to the ion chamber 1 and 2 channels, respectively. Figure 2(b) shows the layout of a detector enclosure. An in-depth discussion of the detector assemblies and the GRAND electronics package can be found in "The Design and Installation of a Core Discharge Monitor for CANDU-type Reactors," by J. K. Halbig, et al.²

The data collection computers sample the GRANDs at a pre-determined interval, usually every 10 or 11 seconds. The total number of data points collected by one detector on one GRAND is limited to 7 855 points per day. For all 20 detectors, this works out to be about 157 100 data points per day or 4 713 000 per month. To store all the data points for one month requires around 150 816 000 bytes of storage. It is impractical to analyze this amount of data. Because of these data storage considerations, the GRANDs employ a data compression technique to reduce the amount of data. During periods of low activity, the GRAND sends only a representative 2% of the incoming data to the collection computer. This reduces the data set considerably. Because the data are time stamped, these gaps do not present a problem in analysis or review of the data. Shown in Fig. 3 are graphs of data from two detectors during one particular day.

A qualitative analysis of the data by a safeguards inspector can yield a considerable amount of information on the fueling activity at the reactor. Each large spike on the graph corresponds to a pair of fuel bundles being discharged from the reactor. Smaller spikes or decay curves or both on the graph may correspond to other activities such as the rotation of the fueling machine, or the radioactive decay, called cooling, of the spent fuel fission products being held in the fueling machine during a shuffling operation. The reactor power level can also be determined from the data because the amount of background the detectors are sensing corresponds to the current power level of the reactor. The background in this context is considered to be the amount of radiation the reactor emits during normal operation when no fuel is present outside of the core. Currently, a safeguards inspector can make qualitative judgments about reactor activity by visually examining the graphs of detector activity on both faces of the reactor. For example, an inspector can count the number of spikes on the graph and determine the total number of fuel pushes the reactor made in a particular day. The total number of fuel pushes counted could then be compared to facility declarations for safeguards verification. Unfortunately, reviewing all the information collected by all the detectors is a

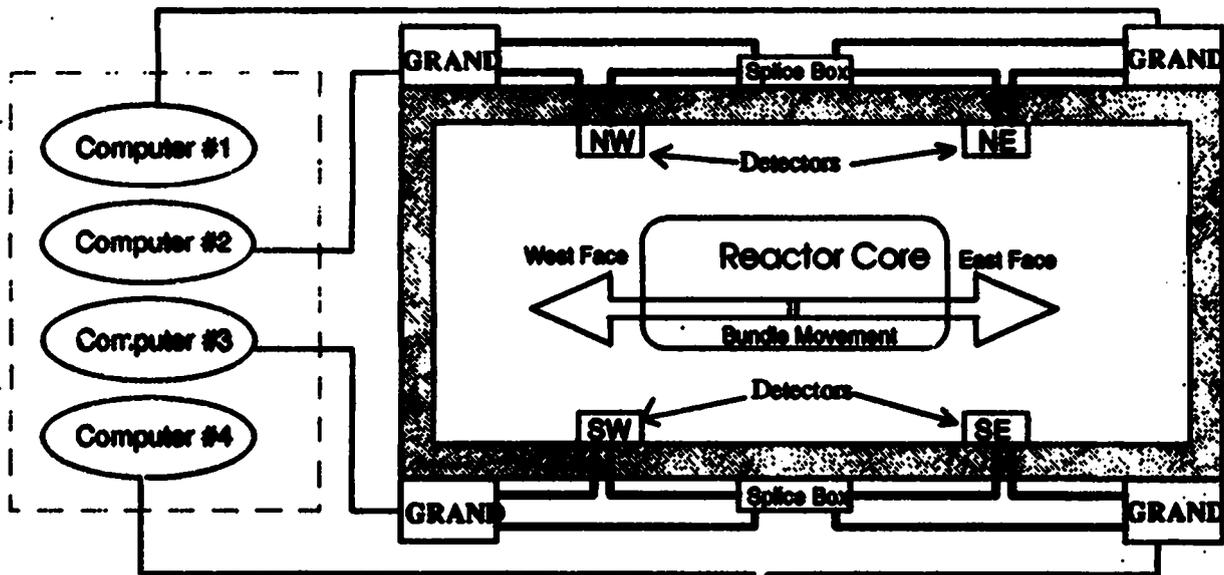


Fig. 2(a). Sample layout of a typical on-load reactor.

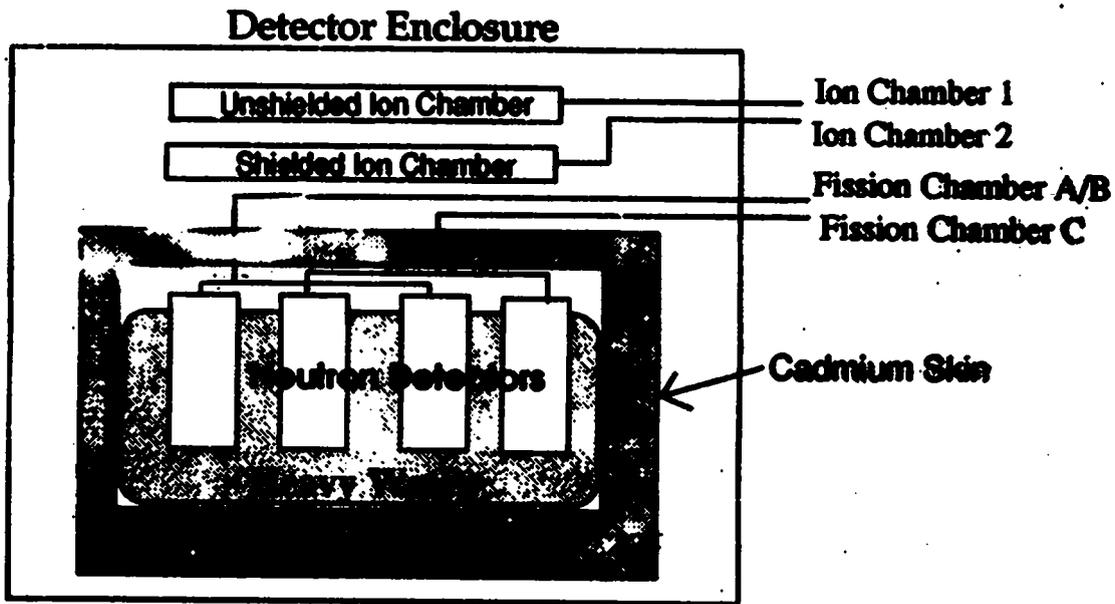


Fig. 2(b). A typical detector enclosure.

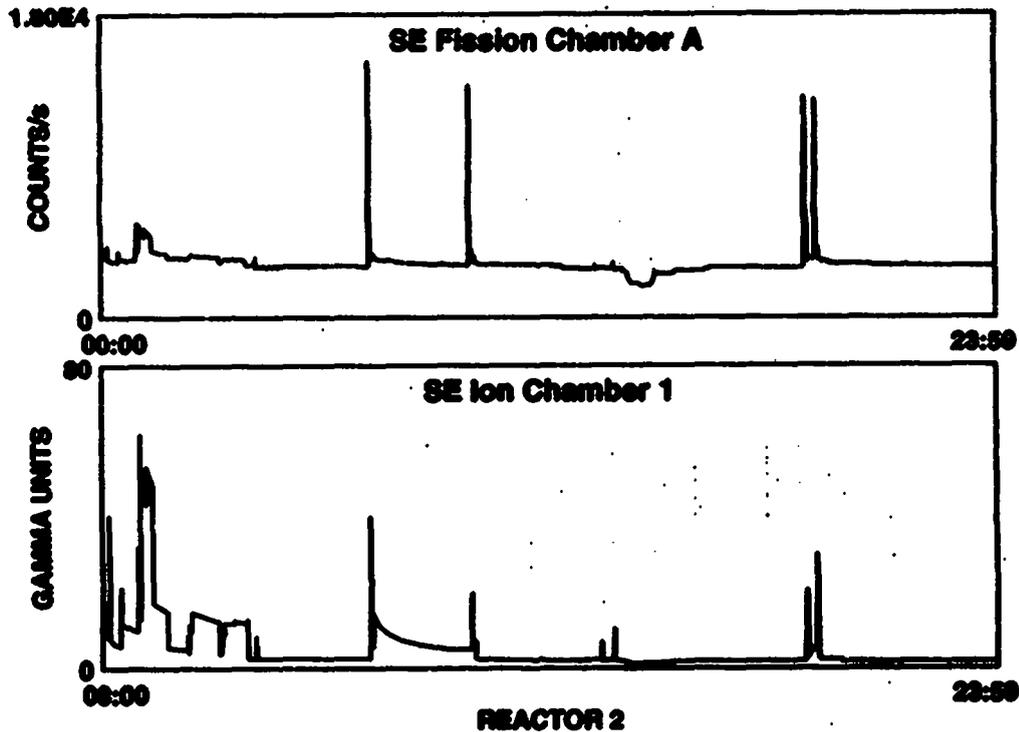


Fig. 3. Sample CDM data from an on-load reactor.

time-consuming task for a safeguards investigator. An automated process could considerably reduce the review time and make it easier for a safeguards investigator to wade through the volumes of data available from an operational on-load reactor. This paper examines the feasibility of developing an automated software system capable of quantitatively analyzing all available data to help a safeguards investigator review reactor defueling activities.

AUTOMATED SOFTWARE ANALYSIS

Analysis of fueling data from an on-load reactor is a complex task. As a result, we designed the automated analysis system to investigate the feasibility of several objectives, each successively more difficult to implement. These objectives are as follows:

1. Identify areas of interest in the data for a safeguards investigator to examine in greater detail.
2. Count fuel bundle pushes and identify when they occurred.
3. Determine reactor power level as a percentage of full power.

4. Correlate an event for one detector channel with all other detector channels. This could be used to detect possible tampering and to ensure that all the channels are operating correctly.
5. Identify the fueling channel from which the spent fuel was discharged.
6. Compute the burnup of the spent fuel bundles.

Each of these objectives extracts features from the entire set of collected data. These features can then help a safeguarding inspector assess all the data collected for a particular on-load reactor. To facilitate these objectives, a prototype analysis tool was developed called CDM analysis. The CDM analysis package accomplishes objectives one through four reasonably well. We developed this tool to explore the possibility of developing a production-grade analysis tool. Objectives five and six were attempted using a neural network modeling paradigm described later. To develop and test CDM analysis as completely as possible, we needed a considerable amount of data. Unfortunately, only about 30 days of data was available for analysis. We concluded that although the total amount of data used for testing the analysis software was sparse, the software still performed very well.

IDENTIFICATION OF AREAS OF INTEREST AND FUEL DISCHARGE EVENTS

A significant problem with an automated analysis of CDM data is identifying areas in the data that a safeguards inspector would be interested in examining. Identifying areas of high activity can considerably reduce the amount of time the inspector needs to peruse the data. At such an early stage in program development it is important that inspectors do not use this system as a substitute for visual examination of all the data, but rather as an aid in the review process. When the reactor is running at a constant power level, a baseline can be established as an average of the background noise collected by the detectors. Important data regions can be identified as large changes in the slope of the average signal, above or below the baseline. This is one technique we used to identify areas of interest.

CDM analysis makes two passes over the CDM data during its search for areas of interest. In the first pass, it slides an average along the signal looking for significant changes. When the slope of the signal jumps above or below the sliding average by more than 10%, the data points are flagged as something to look into later. In the first pass, a huge amount of data points may be flagged as interesting. To reduce the clutter, a second pass is made over just the areas that were flagged. Areas near each other in the time series are clustered together with the maximum data point being marked as the middle of the event. From the resulting list of areas of interest, a report can be generated to alert the safeguards investigator to specific areas of the data. During this development, CDM analysis was not expanded to explain to the investigator what is actually occurring in the underlying system. It currently tells the investigator to check out certain areas for activity. In the future, a detailed analysis of an actual reactor in operation along with the collected data could be used to develop a system that could generate such a report. Finding just the refueling spikes is as easy as finding all other events; CDM analysis just needs to be more discriminating. If the threshold is set very high, at 50%, this two-pass technique will yield as output all the fueling spikes for the given data set. These fueling-event spikes will be marked by the center, or high point, of the events. Coincidentally, this solves objective number two. It is important to realize that this analysis is rather subjective because it assumes the CDM data were recorded during *normal* reactor operation.

CDM analysis can generate graphs of collected data. Figure 4 is an example of an output graph from CDM analysis. Two channels are graphed, which are typical of all events occurring on both faces of the reactor. The vertical lines represent marked fuel discharge events or power level excursions of the reactor or both. The fuel discharge events and the power level excursions are all marked in different colors on the computer screen to make it easy to differentiate between them. The threshold value mentioned earlier is marked by a horizontal line. In this case the threshold is high to locate only fuel discharge and power level events.

CORRELATION OF EVENTS

Once all the fuel discharge events are identified and clustered, each event is then located on the data from all the other channels. To accomplish this correlation, we had to overcome the problem that the clocks on the GRANDs are not synchronized. Correlation is difficult because an event marked with one time stamp on one GRAND may not be found at the same time on another GRAND. Finding events on GRANDs that are on the same reactor face is not too difficult. CDM analysis simply finds the spike height maximum nearest to the time at which the event was recorded. The problem of finding events that occurred on the opposite face in one particular GRAND's data, can be resolved by using the cross wiring of the C fission chambers. CDM analysis finds the event on fission chamber C and uses the spike to perform a pseudo-synchronization of the GRAND's clock. This allows CDM analysis to find where the peak should be on a detector viewing an event on an opposing face.

MONITORING REACTOR POWER LEVEL

Once the areas of interest are identified, power level monitoring is rather simple. When no events are occurring, background radiation is sensed by the detectors. The average of the background can be

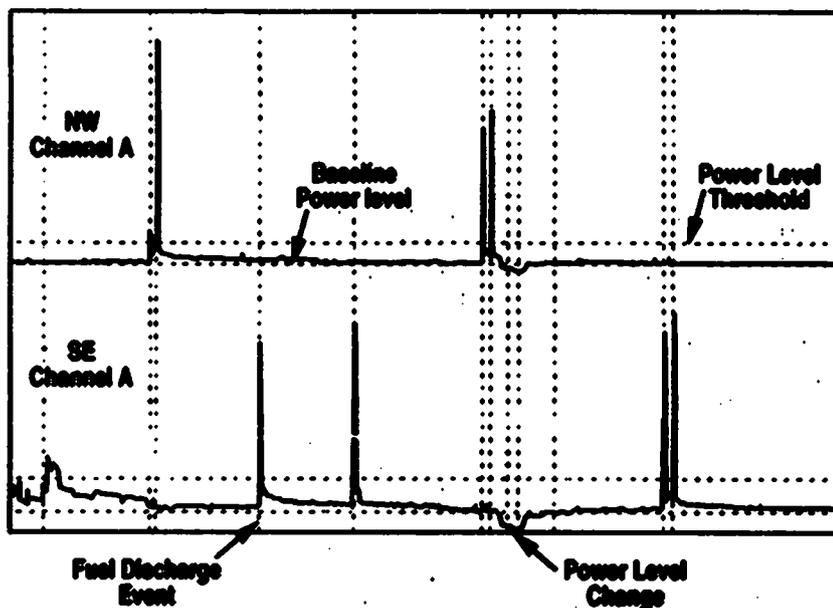


Fig. 4. Sample output from the CDM analysis program.

used to compute the power level by establishing a baseline reading of what is considered to be full power. This baseline is computed by examining data from a reactor that is operating at a fixed power without fuel outside the core. The average value recorded by each detector is used as the baseline. This baseline is marked on the graph in Fig. 4 by a horizontal line. If the average value of the background moves from this baseline, then the power level is changing. The data has shown that most power changes occurred in a stepwise fashion. CDM analysis evaluates the power excursions in the following manner. If the reactor power is raised or lowered, the slope of the average background starts to become very steep. This is marked as the beginning of a power excursion. When this slope flattens out again, the end of the power excursion is marked. The new value at which the average background comes to rest is considered the new power level of the reactor. The average background as a percentage of the pre-defined baseline is the percentage of full power at which the reactor is running. Currently, CDM analysis does not examine more than one channel on one detector when making its power level computations. In a production-level analysis package, this percentage should be an average of all the percentages computed from all channels on all detectors. By taking power level measurements from all sides of the reactor core and averaging them, a more accurate power level reading could be obtained. Even though examining just one channel gives a fairly accurate reading, within 5%, examining all channels is a much better strategy because it provides a redundancy check. Figure 5 is an example of the power level of a reactor being raised from micro-power to 100%. Notice that the excursion occurs in multiple steps. CDM analysis is also capable of printing a report that details each step of the power level change and what power level the reactor moved to. An example of a report for the data graphed in Fig. 5 is shown in Fig. 6.

STATISTICAL PHENOMENA OF CDM DATA

To better understand the reactor fueling process, we statistically analyzed the available data. The height of the radiation spikes representing fuel discharge events can be correlated to the fuel burnup and the location of the fuel channel from which the fuel is being discharged. The burnup contribution to the spike height is a result of fission product buildup in the spent fuel bundles. If one knows the spike heights from all 20 channels for a particular event, one should be able to compute the location and possibly the burnup. Determining burnup is a difficult, multivariate problem that will be discussed in greater detail later. Do all 20 channels make a valid contribution to computing the location of the fuel bundles? The most important question is whether or not detectors on one face of the reactor see events occurring on the opposite face significantly enough so that the error in measurement is small. Figure 7 demonstrates that detectors on one face do not see events on the opposite face. When an event is occurring on one face of the reactor, the spike appearing on the detectors on the opposing face is insignificant. This means that one should treat each face of the reactor separately when attempting to determine the location of fuel bundles being discharged. A correlation does exist between the two spike heights on opposing detectors on the same reactor face. For a given burnup of the fuel bundles being discharged, and as the source is closer to one detector and farther from another, the signal will be larger on the nearer detector and smaller on the farther detector. The graph of the signal strength from two opposing neutron detectors on the same face is shown in Fig. 8.

The relationship is cone shaped because the error in data from a channel becomes greater as a function of the distance of the discharged fuel from the detector. The variance depicted in Fig. 8 is not large enough to prohibit developing a model for locating fuel bundles on the reactor face. Each point on the plot in Fig. 8 corresponds to a position on the reactor face as a function of two opposing

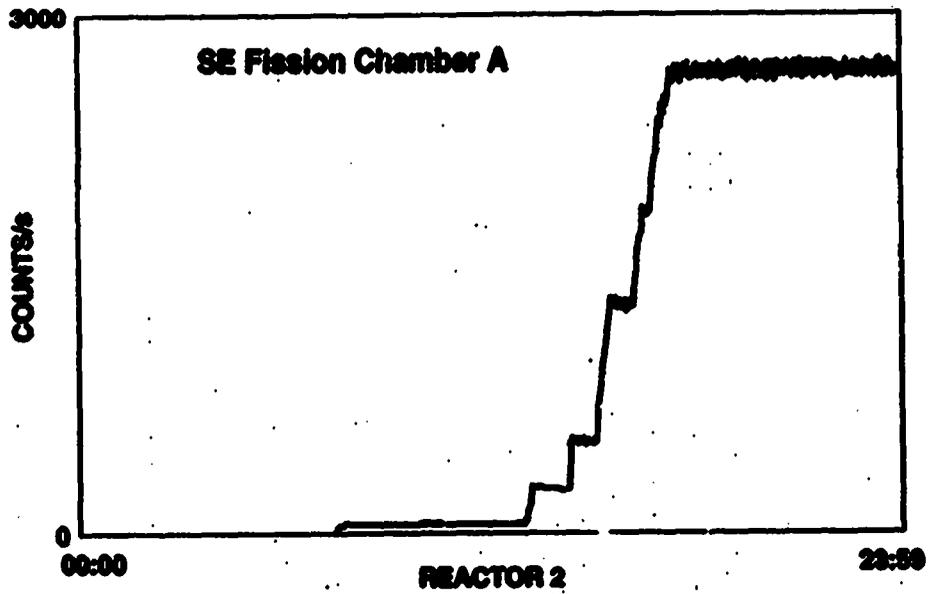


Fig. 5. The multiple steps of a power level change.

Power Excursion Report				
Date	Time		Percent of Full Power	
	Start	End	From	To
90.11.02	13:00:05	14:12:25	1.0	9.8
90.11.02	14:12:48	14:29:15	9.8	21.3
90.11.02	15:03:28	15:31:47	21.3	33.1
90.11.02	15:32:34	15:52:18	33.1	50.9
90.11.02	16:06:14	16:56:34	50.9	100.0

Fig. 6. A sample power excursion report.

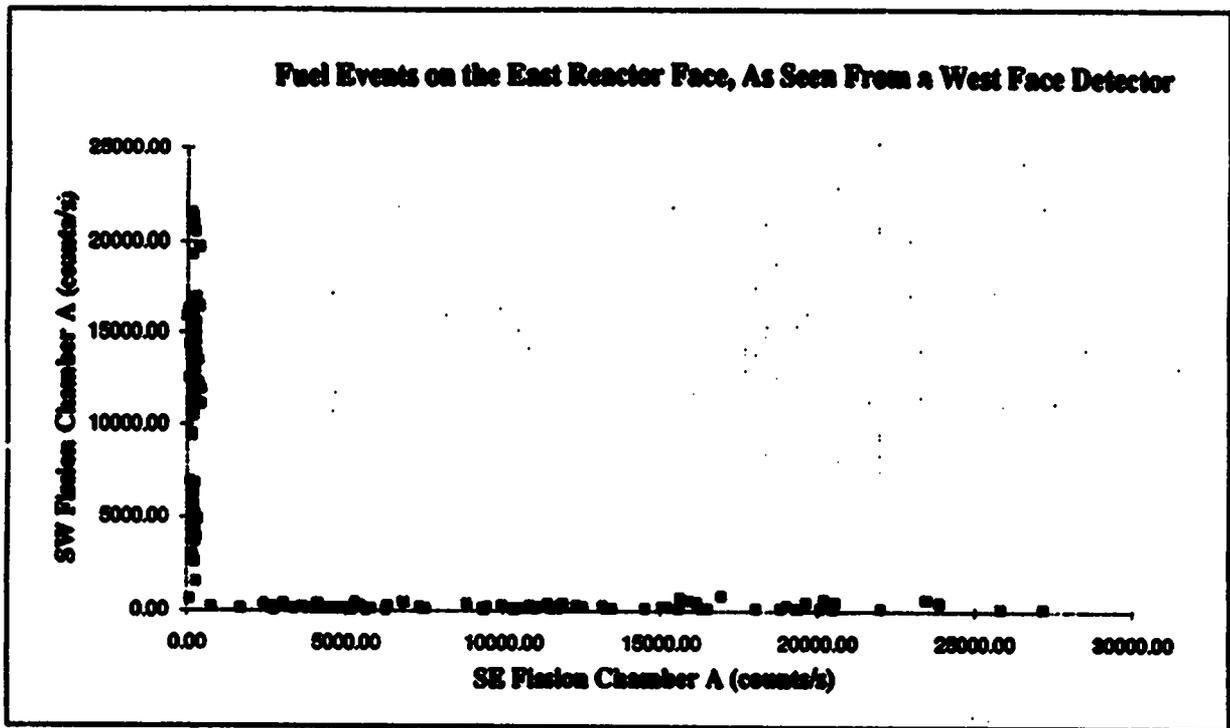


Fig. 7. Fuel events as seen on opposite reactor faces.

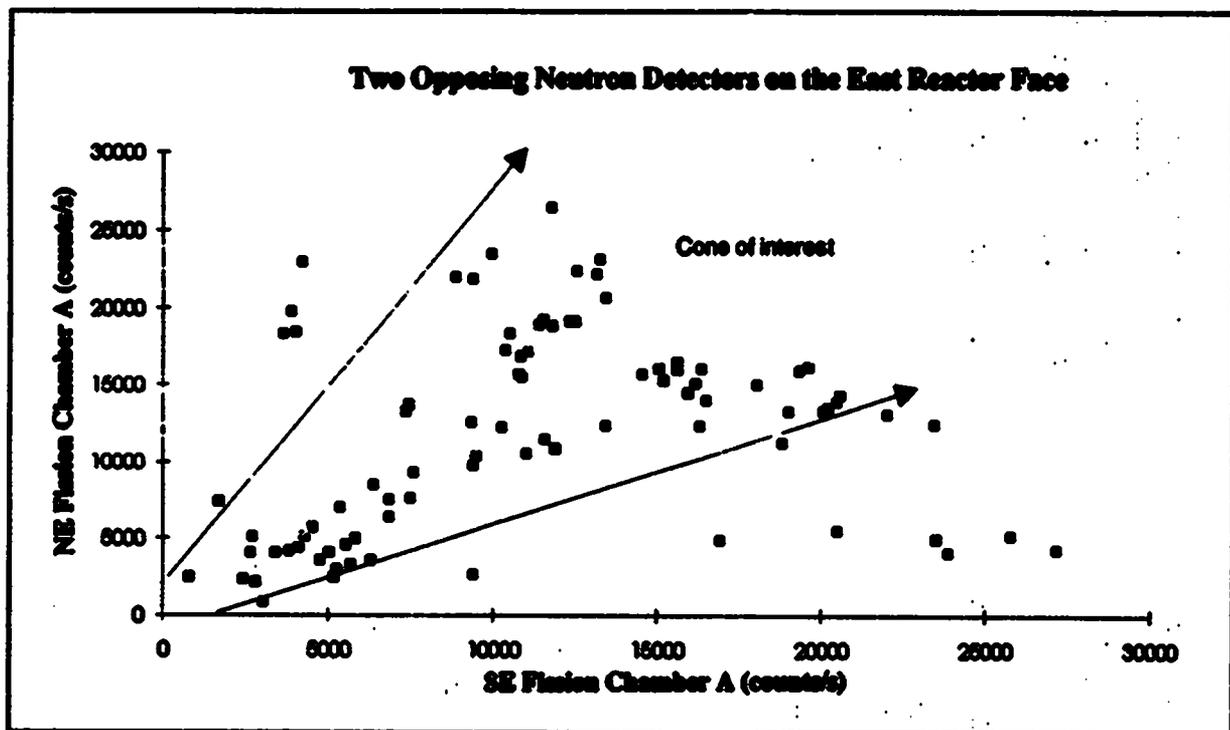


Fig. 8. Relationship between two opposing neutron detectors.

detectors. The neutron channels appear to be the least error prone. Figure 9 shows the same view of the east reactor face, except from the viewpoint of the unshielded ion chamber channel.

Unfortunately, the variance in the ion chamber data is very pronounced. There are a considerable number of outliers in the data set because of the operation of the data acquisition system. The fission chamber data are integrated over the entire sampling time, whereas the ion chamber is a point sample over about 50 ms. Because the ion chamber data are not integrated over the entire data acquisition time, the resulting data are not as representative of the radiation fueling event. Integration causes an averaging effect to occur over the entire sample time yielding much more representative data. This causes less variance to occur in the neutron detector measurement than the ion chamber measurement. We attempted to use the area of the entire fueling event to minimize variance during the analysis, but it did not yield better results. This occurred because the fuel spikes do not have a well-defined shape on either the neutron or the gamma-ray channels. As a result, it was difficult to determine the time interval to integrate over. Often, there are decay curves following fueling events, which corrupt the integral value. A fixed interval analysis was tried, but it did not reduce the variance in the data enough to be significant. Changing the data collection to a higher sampling rate would create a better data set for analysis like this in the future.

Geometrically speaking, the detectors from which these data were taken are set up on each reactor face with the south-side detectors being higher than the north. If the distances from each fuel channel to each of the two detectors are computed, we can derive an equation from physics principles that correlates the ratios of these distances to the ratios of detector spike height. Figure 10 shows how the detectors are set up and gives sample variables for deriving a functional form of the distances as a function of the spike height.

An equation can be derived for the spike height or signal strength ($Sig1$ or $Sig2$), as a function of the source term S_i and the distance to the fuel channel, $R1$ or $R2$. Using a point source functional form for each detector, for which each fuel bundle pair is a point source, the equations are as follows:

$$Sig1 = S_i \left(\frac{\alpha_{21}}{R1^2} + \frac{\alpha_{11}}{R1} \right)$$

$$Sig2 = S_i \left(\frac{\alpha_{22}}{R2^2} + \frac{\alpha_{12}}{R2} \right)$$

In these equations, the α values are coefficients of fit. When the two equations are combined as a ratio, the source term drops out of the equation thus, yielding the following final form.

$$\frac{Sig1}{Sig2} = \frac{\frac{\alpha_{21}}{R1^2} + \frac{\alpha_{11}}{R1}}{\frac{\alpha_{22}}{R2^2} + \frac{\alpha_{12}}{R2}}$$

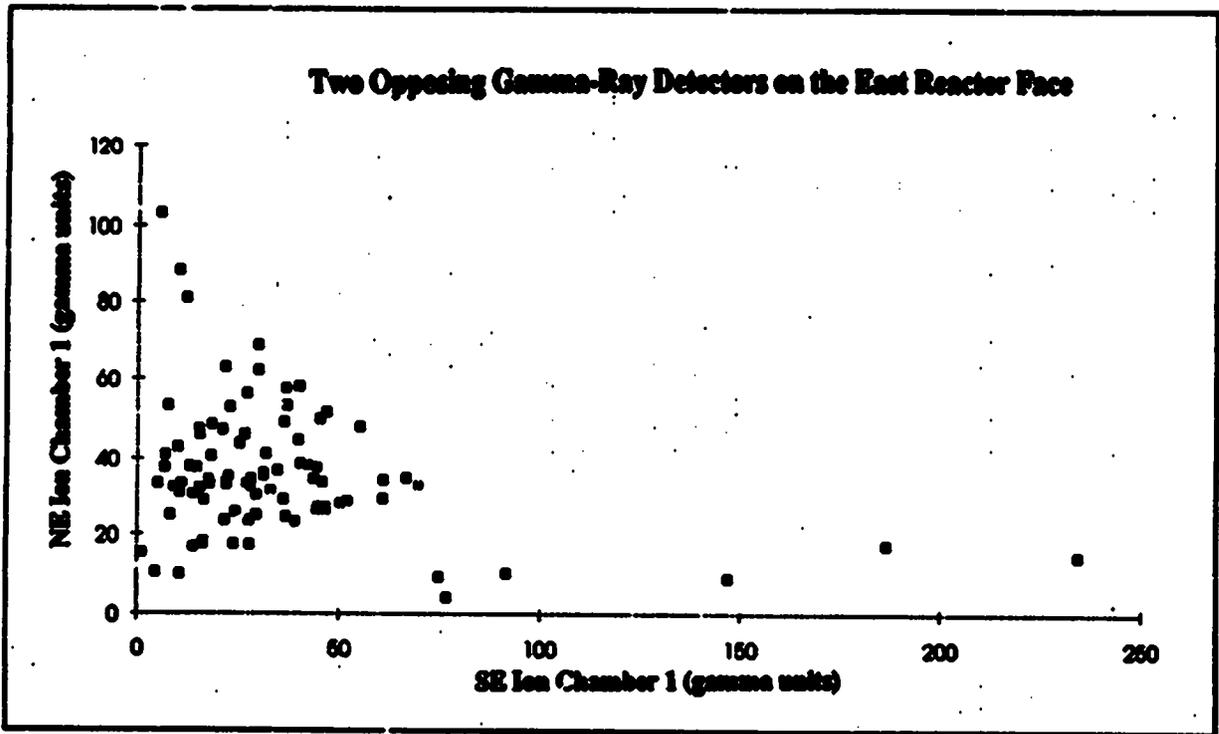


Fig. 9. Relationship between two opposing gamma-ray detectors.

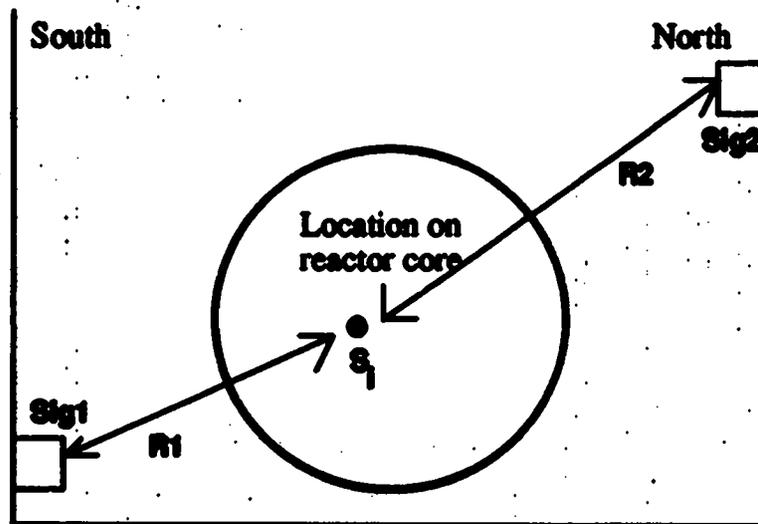


Fig. 10. Typical detector setup.

A nonlinear optimization was performed to obtain the values for the equation. These values were determined to be the following:

$$\alpha_{21} = 33.82$$

$$\alpha_{11} = 0.01$$

$$\alpha_{22} = 13.19$$

$$\alpha_{12} = 0.00$$

It is interesting to notice that the $1/R$ component makes a minimal contribution to the equation after this fit. The graph in Fig. 11 shows the relationship between the function and the actual signal values for fission chamber A. If the fit were perfect, one would expect the plot to be linear. Unfortunately, there is variance in the data from this exact, functional form. This means that the functional form cannot be used for determining exact locations on the reactor face of fuel discharge events. Some type of error-tolerant model is needed to make this determination. Neural network models are perfectly suited for such applications. The variance shown in Fig. 11 can be compensated for by using a representative sample of the data, which conforms to the $1/R^2$ relationship, to make a mathematical model of the distance-to-spike-height system. If the model is fairly robust, the other fuel channels can be extrapolated from only a few examples out of the total data set.

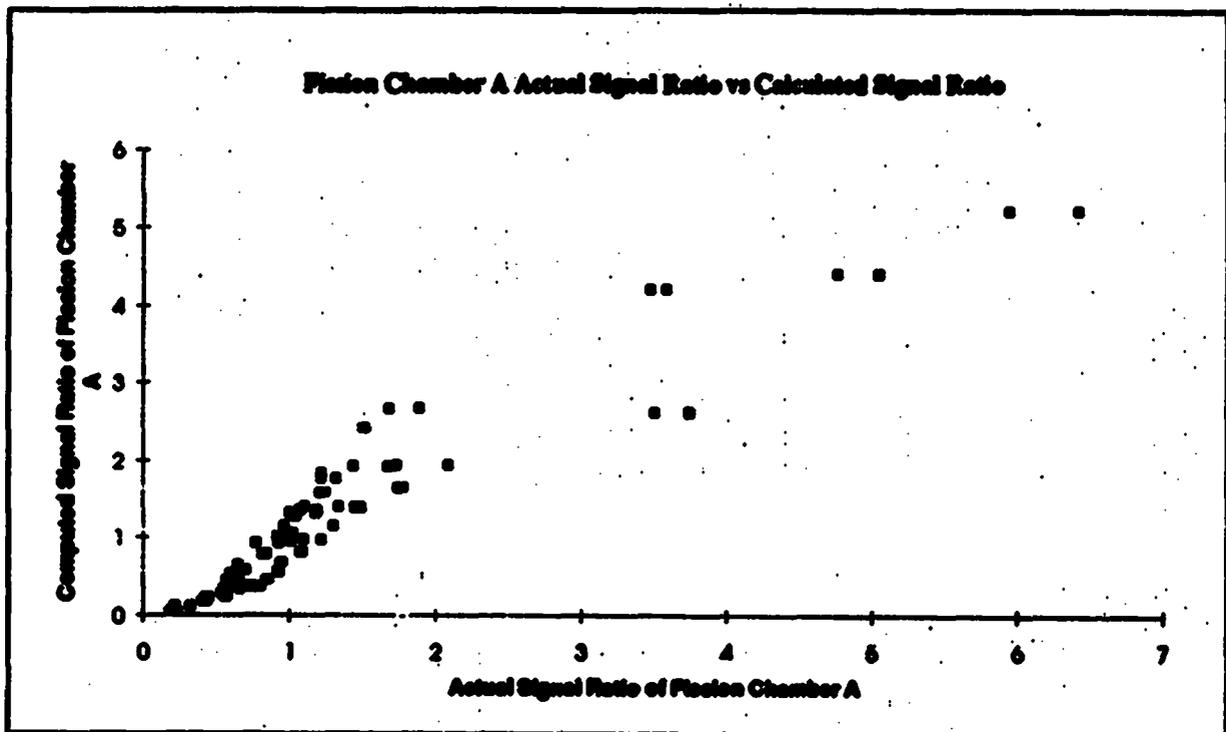


Fig. 11. Actual signal values compared to fit.

NEURAL NETWORKS FOR SOLVING THE FUEL GEOMETRY PROBLEM

Neural network models are mathematically derived from nerve cell biology. They are an abstraction of a collection of fundamental units organized in a hierarchical fashion. In biological systems each unit would represent a nerve cell. Learning, and thus reacting to stimuli, is accomplished in a *black box* scheme. The black box is molded when patterns of stimuli and their corresponding responses are presented in a repetitive fashion. Given the input and the expected output, the black box *learns* the mapping from input to output. The learning inside the black box takes place by updating a series of weighted mathematical functions after each pattern presentation. The goal is to minimize the global error over the entire training set. After training, the model is presented with new stimuli. If the model has been trained to a minimal amount of error, it should generalize and predict the appropriate output for an arbitrary input pattern. Further details regarding the internal, technical aspects of the neural network modeling paradigm may be found in *Parallel Distributed Processing*, by David E. Rumelhart and James L. McClelland.³ The neural networks used in this proof-of-principle were created using NeuralWorks Professional II/Plus,⁴ a commercial neural network development tool manufactured by NeuralWare, Inc. A Sun SPARCstation computer was used to perform the analysis and develop the neural network.

Neural networks seem well-suited for automation of continuous processes because of their capability to generalize given only a representative sample of data. For example, Nippon Steel has utilized neural networks to achieve greater reliability in their continuous casting process (Hidetaka 1991).⁵ Neural networks are also moving into the nuclear power arena because of their predictive capability. Roh and Cheon have developed several neural network models to attempt to predict thermal load requirements of a nuclear power station and have achieved very promising results.^{6,7} Bartlett and Uhrig have also applied neural networks to nuclear power by creating a model to automate status diagnostics.⁸ The particular neural network applications described in this paper are considerably different from previous applications to nuclear power in that we are performing nuclear safeguards rather than optimizing the power generation capabilities of the reactor.

The 30 days of available data yielded around 170 examples of fuel discharge events. Unfortunately, these events only used 90 out of the 460 available fuel channels in the reactor core. As well, all the 170 events occurred as a result of fuel shuffling operations during initial reactor startup. This provides a challenge because data from shuffling operations are not necessarily indicative of normal reactor fueling activity. Regardless, we thought that an appropriately trained neural network model could classify fuel events into different regions on the reactor face. This type of classification could be a potential benefit to safeguards because an investigator could then verify the event against facility declarations. Regions on the reactor face must be used because of the symmetry of the geometry in the problem. It is possible that there are fuel channels for which the proportion of the distances from each detector is the same. Figure 12 shows the proportions of the distances for all the events on the east face of the reactor: 88 fuel discharge events.

The overlapping points represent examples of the geometric symmetry problem. This problem is that the ratio of the distances from the two opposing detectors to a given fuel channel could be the same for more than one channel on the reactor face. This problem can be solved either by using regions of the reactor face, which include symmetric points in the same region, or by eliminating the symmetric points during creation of the model and allowing the neural network to infer the region in which they lie. The latter method was used in the models presented in this paper, although we think that using symmetric regions would be a good solution to the problem as well. The potential capability to extrapolate the fuel discharge location, based on only a few examples, is one of the key

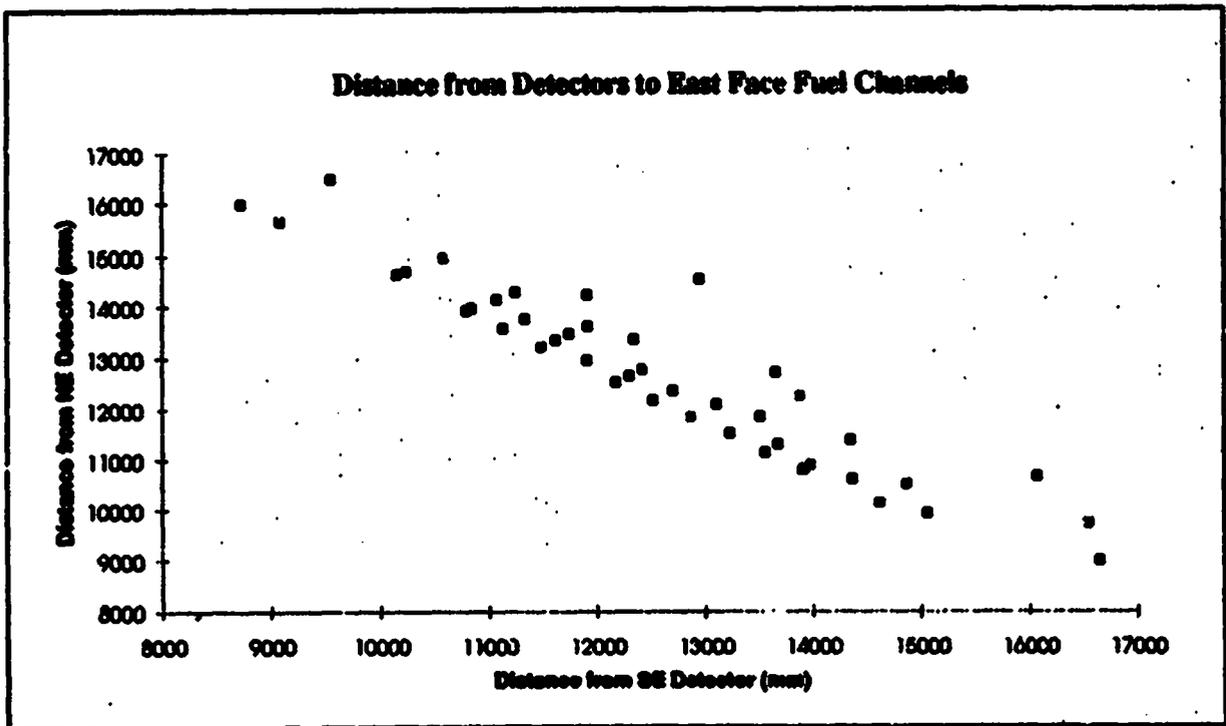


Fig. 12. Distances of fuel events from two opposing detector enclosures.

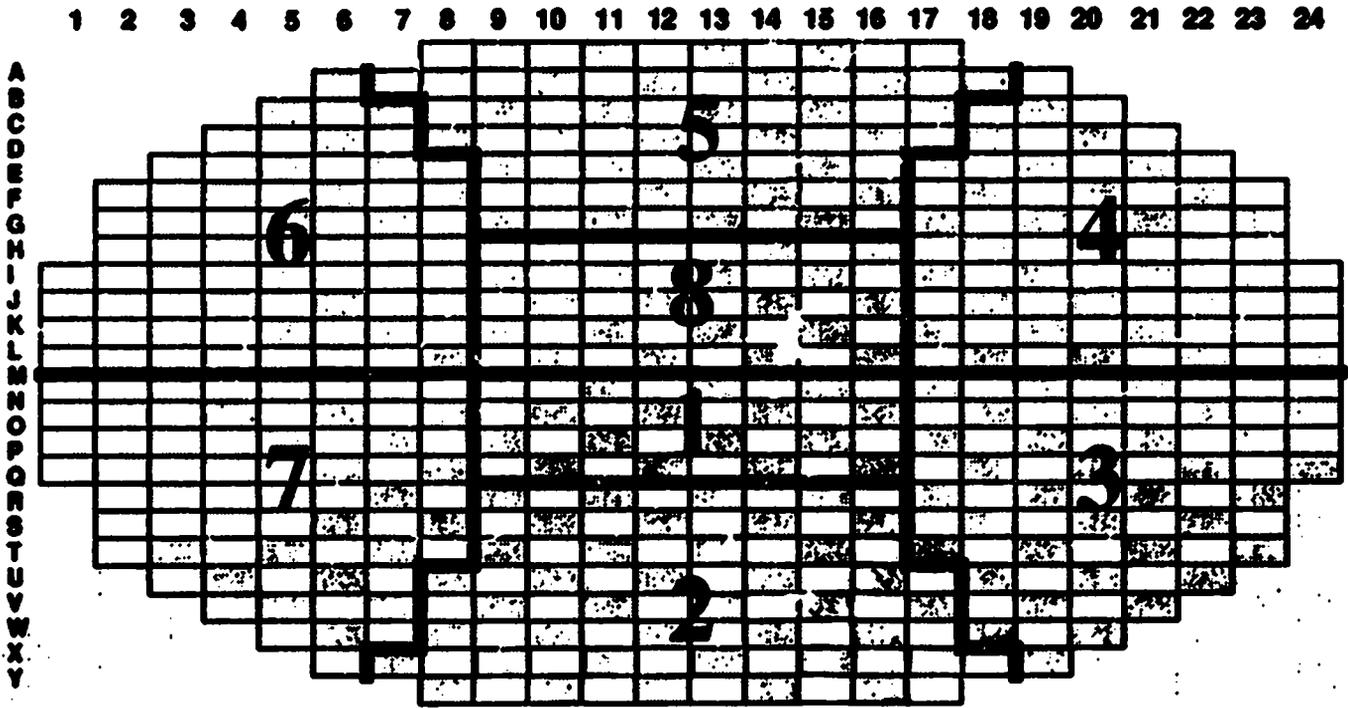
reasons we chose a neural network model to solve this problem. We estimate that by training the network on only a few examples of each region, the other events in that region could be inferred. This would allow the model to use events that have occurred to classify events that will occur.

The first neural network model divided the channel map into eight regions. This channel map and the eight regions are shown in Fig. 13. Almost all the regions were chosen because of the distribution of the points in the available data.

Because detectors on one face do not reliably see events on the opposing face, only 10 channels from the same face out of the 20 total channels were used in the neural network model. Although the ion chamber channels are virtually useless in the computation, they are needed to bring the neural network to convergence. The ion chambers act as noise during the training process to help separate the input vectors into appropriate categories. Back-propagation was chosen as the modeling paradigm because of its ability to use real-valued inputs.⁹ The resulting neural network model was composed of 10 inputs, 2 hidden layers each with 9 nodes, and 3 outputs. The three outputs were used to perform a binary mapping of the eight possible regions. Figure 14 shows a graphical representation of the eight-region model.

Because the peak heights of the correlated channels are used as inputs to the neural network, there is a scaling problem. The numeric values for the heights can be very large numbers to feed to a neural network. To overcome this problem, the inputs are normalized to within the range of -1.0 to 1.0. The result of this normalization is a very tight clustering of the information, which leads to a training problem for the network. Normally, one would use the normalization range of 0.0 to 1.0 because that is the dynamic range of a standard sigmoid transfer function. In this case, a hyperbolic

East Reactor Face - Looking West



- West fueling channel.
 - East fueling channel.

Fig. 13. Eight-region map of reactor face.

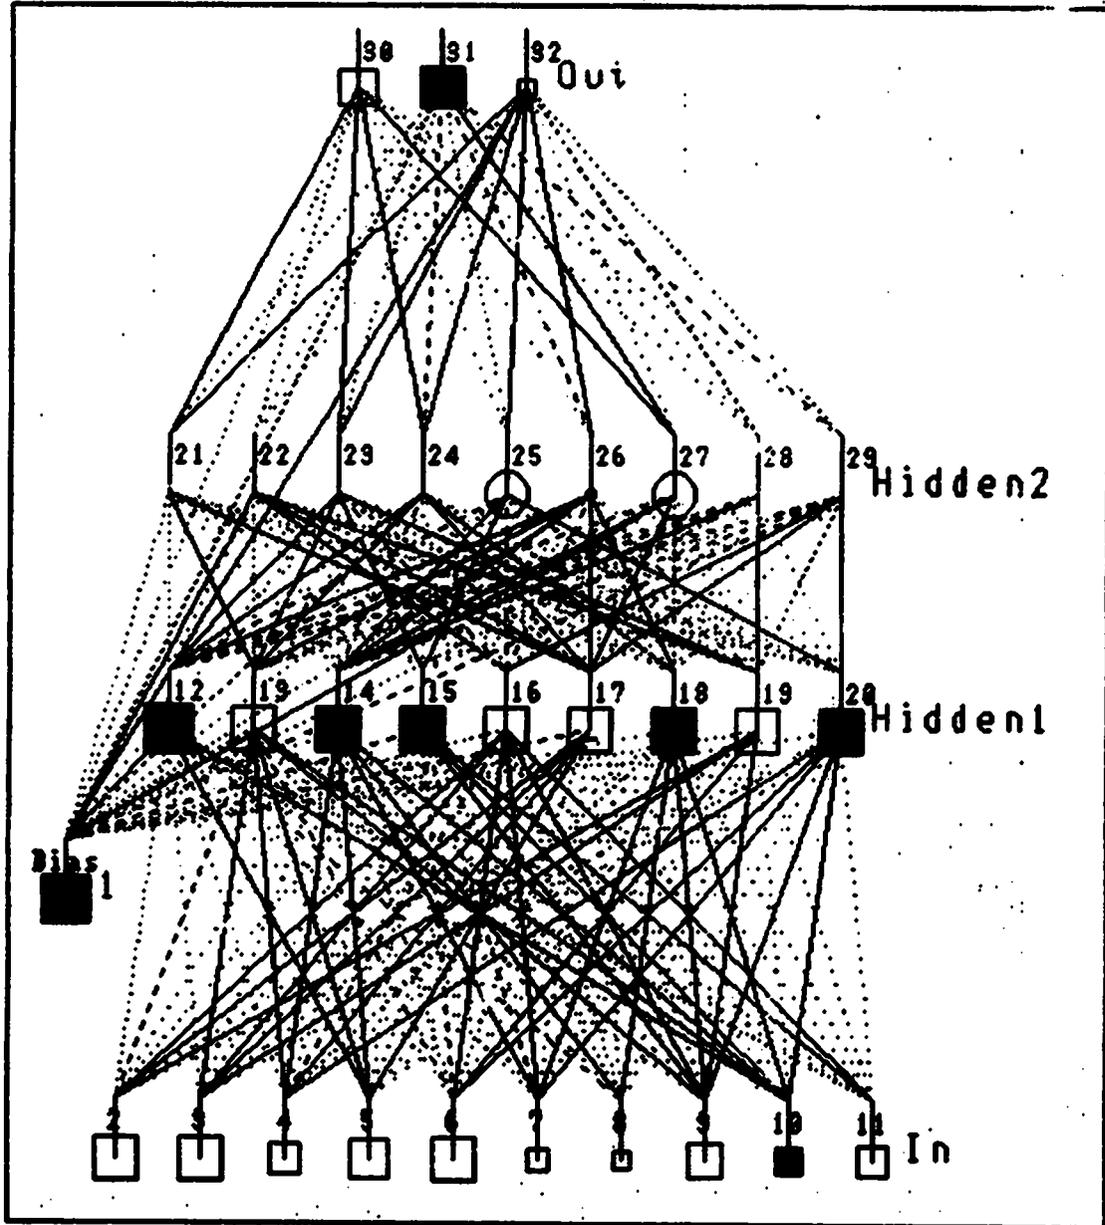


Fig. 14. Graphical representation of eight-region neural network.

tangent function was used as a transfer function because its wider bandwidth stretches the attenuation. The hyperbolic tangent equation used for the transfer function is¹⁰

$$f(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$$

As a result, the error equation for the local error for a particular node at level n as a function of the errors at level $n+1$ is¹⁰

$$\epsilon_j^{[n]} = (x_j^{[n]} + 1) \cdot (1 - x_j^{[n]}) \cdot \sum_{k=1} (\epsilon_k^{[n+1]} \cdot w_{kj}^{[n+1]}) ,$$

where $\epsilon_j^{[n]}$ is the error of node j , $x_j^{[n]}$ is the output of node j , and the $w_{kj}^{[n+1]}$ are the weights on layer $n+1$.

To get the neural network to converge, it is necessary to use slight variations of the standard back-propagation algorithm's delta-rule. Instead of the standard delta-rule, the cumulative, generalized delta-rule was used.¹⁰ Instead of updating the weights after each pattern presentation, a cumulative weight update was tallied and then applied at the end of the training epoch. This helps speed up the training of the network and overcome any problems that may exist with order in the training data. The model also used a bias to assist in speeding up the convergence of the network.

A four-region neural network was also created with 10 inputs, 1 hidden layer containing 15 nodes, and 2 outputs. The internal structure, such as the transfer function and the learning rule, are the same as for the eight-region model. The two outputs of the four-region model also performed a binary mapping to the four regions on the reactor face. The regions in this case were chosen randomly to be quadrants on the face of the reactor. The four-region channel map is shown in Fig. 15.

In the case of both the eight- and four-region models, binary mappings of the region numbers were used instead of bin-sorted outputs because it seemed to improve the generalization skills of the model. Both models were able to achieve convergence in under 50 000 training epochs, with an epoch size of 12. We used two hidden layers in the eight-region model in the hope that it would improve the generalization of the network, but we found that both one- and two-layer networks seemed to generalize equally well. For this problem, the two hidden-layer networks were more difficult to train and required a greater number of training epochs. The four-region network shown in Fig. 16 was created using only one hidden layer because of this finding.

NEURAL NETWORKS FOR FUEL BURNUP COMPUTATION

Computing fuel burnup is the most difficult of the original objectives because it contains the greatest number of variables. The factors affecting fuel exposure and the radiation levels detected by the CDM are as follows:

1. The fuel channel where the bundle is located (1 of 460)
2. The position of the bundle inside each fuel channel (1-13)
3. The time at each position
4. The time after shutdown of a particular fuel push
5. The rate at which the fuel push occurs

East Reactor Face - Looking West

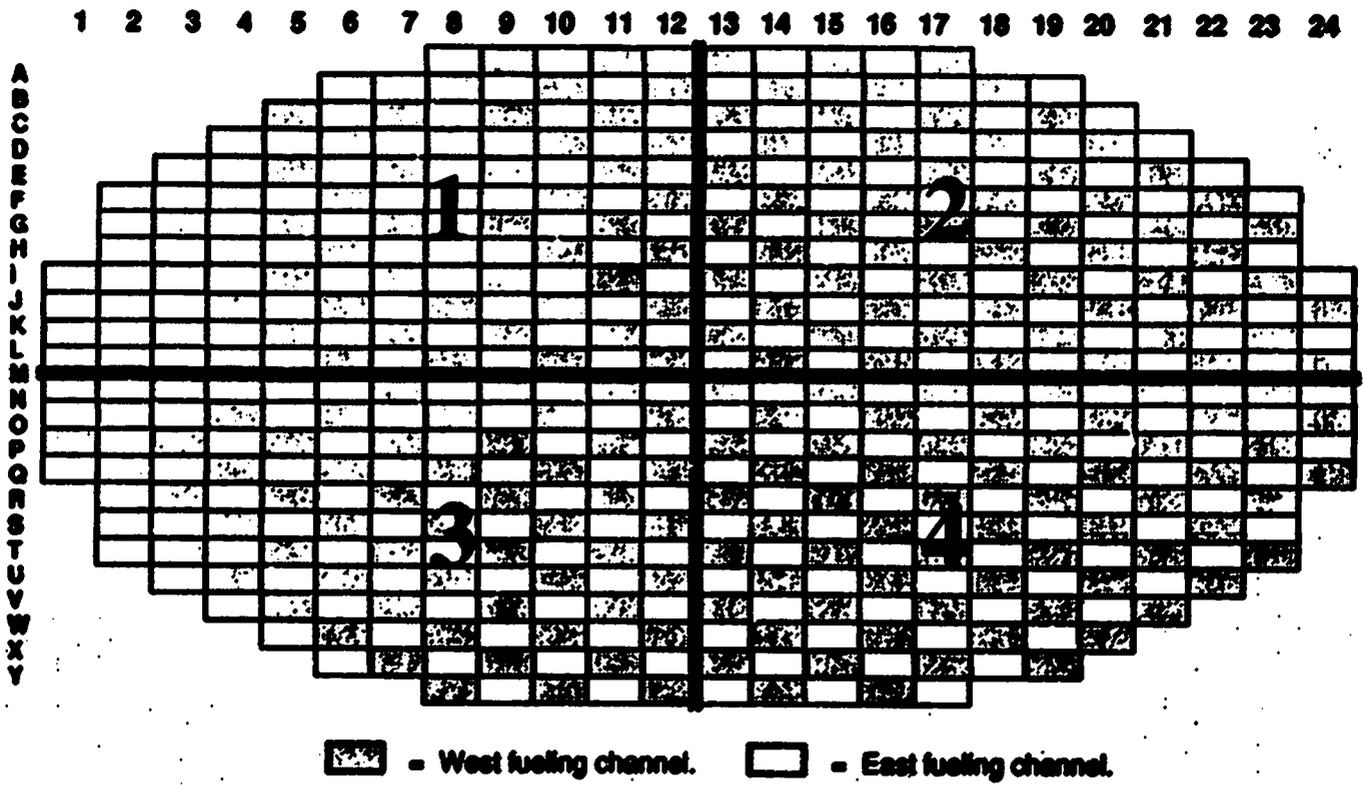


Fig. 15. Four-region map of reactor face.

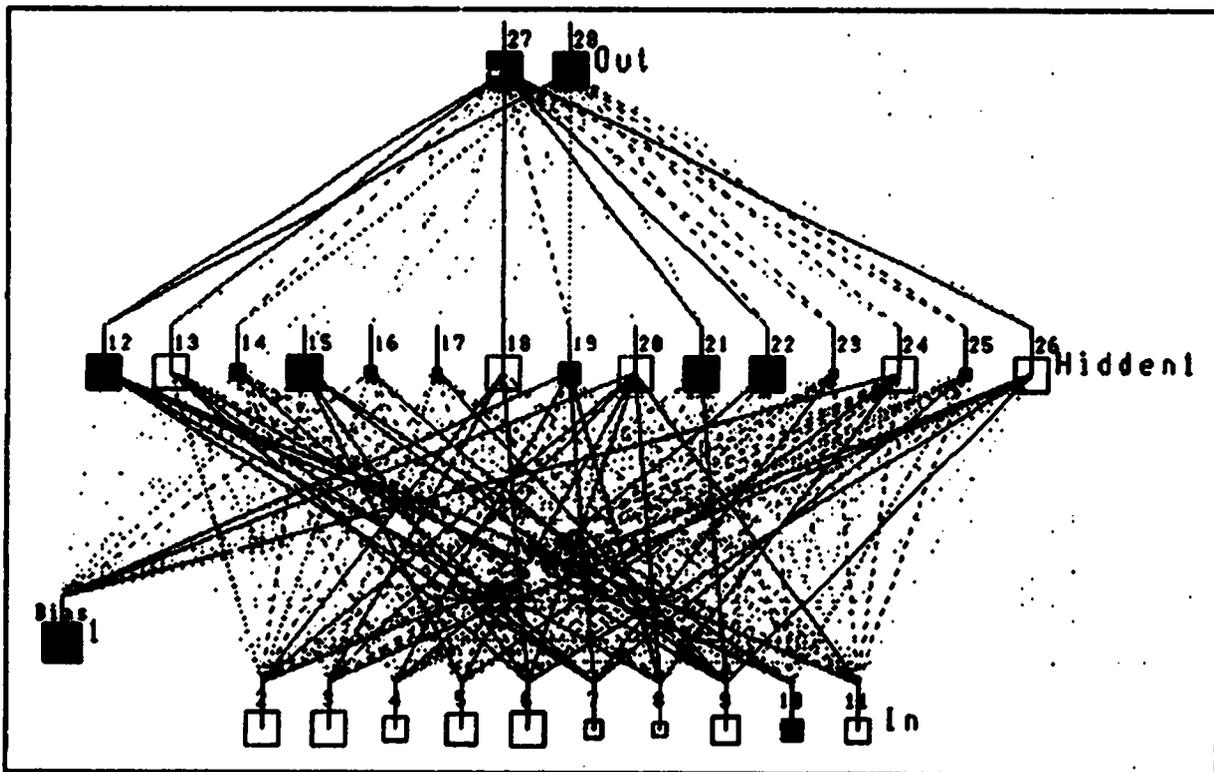


Fig. 16. Graphical representation of four-region neural network.

For safeguards, it may be important to be able to determine if the facility is discharging low-burnup fuel from the reactor. As in the fuel geometry problem, we think that the fuel burnup problem can also be solved using a neural network to classify the burnup into different categories, one being low-burnup. It is difficult to actually compute numbers for the burnup of each individual fuel bundle because the spike recorded by the CDM is an additive value of n bundles being discharged simultaneously. The measured activity from which the burnup computation is made also depends on the recent irradiation history of the fuel bundles being discharged. For normal reactor operation, this factor would need to be taken into account. Because the data used for this research only included shuffling operations, the irradiation history becomes irrelevant. Figure 17 shows the distribution of the known, additive burnups for the discharged fuel bundles on the east reactor face. It is interesting to notice that the burnups fall into one of four distinct regions.

Because the available data only contained fuel events during the initial reactor startup and shuffling operation, the fuel in the reactor only occupied the last four positions in each of the 90 fuel channels for which there were recorded events. The fuel bundles discharged from the reactor were never at different positions in the fueling channel. This means the fuel-exposure-dependent variables have a negligible effect on the burnup prediction, hence, it is possible to create a neural network model for classifying the burnup into one of the four categories based upon the CDM data.

The neural network model used for the burnup calculation contained 10 inputs, 1 hidden layer with 10 nodes, and 2 outputs. The two outputs were used as a binary representation of the four possible regions. The internal structure of the network, such as the transfer function and learning rule, was the same as those used in the geometry-problem neural networks.

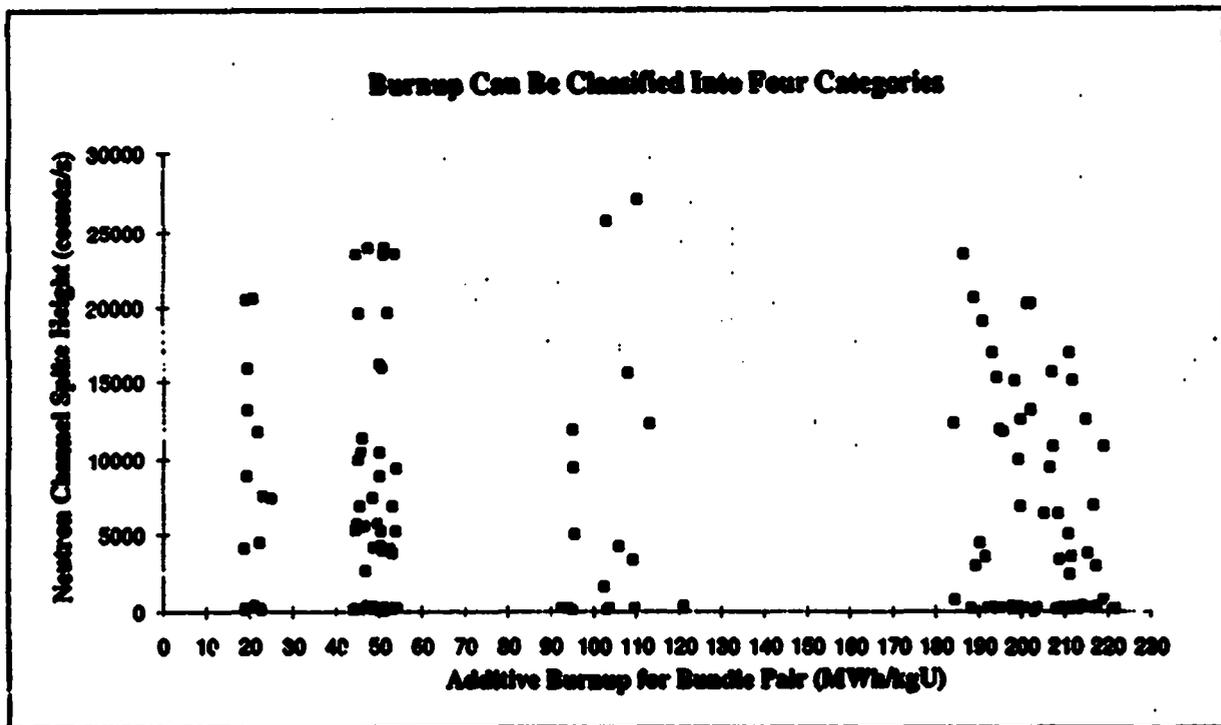


Fig. 17. Four categories of recorded burnups in sample data.

RESULTS OF NEURAL NETWORK MODELS FOR SOLVING GEOMETRY AND BURNUP

The neural networks used for solving the geometry problem were trained and tested on data from the east face of the reactor, although the west face could have been used just as well. Because of the limited amount of data, a representative number of fuel events were chosen as the training set in an attempt to extrapolate fuel events on channels the network had never seen before. For the eight-region model, 63 patterns were used in the training set out of the 88 total fuel events on the east reactor face that were used as the test set. As a result, 72 events were classified correctly into 1 of the 8 regions on the reactor face. This means the network generalized and extrapolated the correct region for nine events, and correctly classified the region 82% of the time. Because the region was correct for 9 out of 25 possible extrapolated events, the network extrapolated correctly 36% of the time.

The four-region model performed slightly better than the eight-region model, with only 32 patterns in the training set out of the total 88. The model classified 68 events correctly into 1 of the 4 regions on the east reactor face. Although this is only 77% accuracy, it generalized and extrapolated the position for 36 events out of a possible 56. This means it extrapolated correctly 68% of the time. This is considerably better than the eight-region model with only nine correct extrapolations.

The neural network for computing burnup performed especially well. It is important to reiterate that part of this performance is explained by the fact that the exposure-dependent variables were missing from the equation. The burnup neural network was trained on 48 patterns out of the 88 total in the test set. The network classified the burnup correctly into 1 of the 4 categories 81 times. This is an accuracy of 92%, with 33 patterns being extrapolated out of 40. This is an extrapolation correctness of 82%.

CONCLUSIONS

The CDM analysis tool created here will act as a prototype for studying and creating a more robust CDM data analysis tool. The potential of this tool as a bundle counter and a power-level monitor has been demonstrated. Neural network implementations for determining the area of the reactor face from which the fuel was discharged and the additive burnup of the two fuel bundles have been successful enough to warrant further research. We think that with a more complete set of representative data from an operating on-load reactor, neural network models could be created to render up to 100% accuracy in position and burnup. The data needed to achieve this capability should include fuel pushes from all 460 channels of the reactor and a complete cycle of fuel through all 13 positions in every channel. The results obtained in this research are extremely promising considering the limited amount of data available. We believe that the results can only get better with more data.

Future work should include devising a more accurate technique for determining areas of interest in the CDM data, rather than a sliding average. Power level monitoring using an average over all 20 channels will also yield a more accurate power level calculation. Deficiencies in the collection of quantitative data should be corrected. We need more samples of data per unit time and a gamma-channel reading more representative of the measurement period. In addition, different types of neural network models should be tried once a representative amount of data has been obtained. The portability of neural network models to other reactors of the same type should also be investigated. Neural network models hold great promise for future work in the area of core discharge monitoring and automated examination of large volumes of continuously collected data for better nuclear safeguards. We firmly believe that a commercial-grade tool for monitoring power and counting fuel bundles from CDM data should be developed.

REFERENCES

1. J. K. Halbig and A. C. Monticone, "Proof-of-Principle Measurements for an NDA-Based Core Discharge Monitor," *Nucl. Mater. Manage.* XIX (Proc. Issue), 847-852 (1990).
2. J. K. Halbig, A. C. Monticone, L. Ksiezak, and V. Smiltnieks, "The Design and Installation of a Core Discharge Monitor for CANDU-type Reactors," *Nucl. Mater. Manage.* XIX (Proc. Issue), 839-846 (1990).
3. David E. Rummelhart and James L. McClelland, *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, (The MIT Press, Cambridge, Mass., 1986).
4. Neural Computing: NeuralWorks Professional II/Plus Users Manual (NeuralWare Inc., Pittsburgh, PA, 1991).
5. Kominami Hidetaka, "Neural Network System for Breakout Prediction in Continuous Casting Process," Nippon Steel Technical Report No. 49 (April 1991).
6. M. S. Roh, S. W. Cheon, and S. H. Chang, "Power Prediction in Nuclear Power Plants Using a Back-Propagation Learning Neural Network," *Nucl. Technol.* 94, 270 (1991).
7. M. S. Roh, S. W. Cheon, and S. H. Chang, "Thermal Power Prediction of Nuclear Power Plant Using Neural Network and Parity Space Model," *IEEE Transactions on Nuclear Science* 38(2), 2 (1991).

8. **Eric B. Bartlett and Robert E. Uhrig, "Nuclear Power Plant Status Diagnostics Using an Artificial Neural Network," *Nucl. Technol.* 97, 272 (March 1992).**
9. **Philip D. Wasserman, *Neural Computing Theory and Practice*, (Van Nostrand Reinhold, New York City, NY, 1991).**
10. **NeuralWorks Professional II/Plus Reference Guide (NeuralWare, Inc., Pittsburgh, PA, 1991).**