Seventh American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies
NPIC&HMIT 2010, Las Vegas, Nevada, November 7-11, 2010, on CD-ROM, American Nuclear Society, LaGrange Park, IL (2010)

CONTROL ROD MONITORING OF ADVANCED GAS-COOLED REACTORS

C. J. Wallace, G. M. West, G. J. Jahn, and S. D. J. McArthur
Institute for Energy and Environment,
University of Strathclyde
cwallace@eee.strath.ac.uk, graeme.west@strath.ac.uk, gordon.jahn@eee.strath.ac.uk,
s.mcarthur@eee.strath.ac.uk

D. Towle and G. Buckley
Graphite Core Project Team, EDF Energy
Barnwood, Gloucester
dave.towle@edf-energy.com, graham.buckley@edf-energy.com

ABSTRACT

The UK's fleet of Advanced Gas Cooled Reactors (AGR) are approaching, and have in some cases exceeded, their original design lives. Continued operation is under enhanced safety cases based on monitoring, inspection and component condition assessment of the core and related systems. This paper presents an analysis of the regulating control rods of an AGR, which are used to manage the power and reactivity of the core. Current manual analyses attempt to detect possible restrictions in the motion of the rods due to degradation of the graphite core, however the development of an automated intelligent analysis of the control rod data provides a repeatable and auditable method of analyzing the data. It is shown, by means of an example data set, that despite some limitations in the scope of the recorded data, it is possible to estimate the performance of the rods and present this information to the engineer in a way that more easily indicates abnormal behavior than existing analyses. It is also noted that though this work was initially conceived as a method of detecting restrictions in the motion of the regulating control rods, the results are potentially more useful is characterizing control rod performance and has potential application in predictive maintenance.

***Key Words*: Condition Monitoring, Control Rods**

1. INTRODUCTION

The graphite cores of Advanced Gas Cooled Reactors (AGR) experience dimensional change, cracking and weight loss (due to radiolytic oxidation) under prolonged exposure to neutron and gamma irradiation in the presence of carbon dioxide. This results in distortion, cracking and weight loss of the graphite bricks that comprise the core. Each AGR core consists of around 300 fuel channels and around 75 control channels and it is of critical importance that the fuel and control assemblies can move through their respective channels without restriction. In order to ensure that the reactors are safe to continue operating, inspections of the core every two to three years measure the dimensions and distortion of selected fuel channels, however it is not practical to inspect every channel on a regular basis. For this reason, in recent years a greater emphasis has been placed on Condition Monitoring (CM) and the use of existing reactor data sources to derive more detailed knowledge about the core between outages. There are several potentially useful sources of such data, such as measurements recorded during refuelling or the continuously updated logs of reactor parameters, including various temperatures and pressures. This paper discusses the use of the height data recorded for the control rods which are used to manage the power of the core. This data is manually analysed to various extents at some stations, however a formal method of determining where the data may indicate a negative core condition or other anomalous behavior does not currently exist. This paper discusses the potential of the automated analysis of this data as a further addition to the core-monitoring regime and gives an example of an analysis of the data and shows how a physical interpretation of the results can be used to more rigorously determine the performance of the control rod systems.

2. ADVANCED GAS COOLED REACTOR

2.1 AGR Design

The AGR is an evolution of the MAGNOX design, also using graphite as the core structure and the moderator. The use of higher pressure carbon dioxide coolant and enriched uranium fuel, however, allows for a significantly higher power density, and hence a smaller core. The core is contained within a pre-stressed concrete pressure vessel within which the coolant, held at a pressure of around 40bar, transfers heat from the fuel assemblies to the heat exchangers. The AGR was designed with the intention of regular on-load refueling, however this only occurs at some stations at reduced power due to issues with coolant flow during refueling.

2.2 Graphite Core

The graphite core is composed of thousands of cylindrical interlocking graphite bricks[1], arranged in between 11-13 concentric rings over 10 to 15 layers, each of around 850-900mm in height, which form fuel channels. The interstitial bricks placed in the gaps between the fuel channels form smaller channels that are used for inserting control rods and for carrying the coolant flow.

2.3 Control Rod Design

AGR cores feature two set of rods: bulk rods and regulating rods. The bulk rods are used for long-term hold down or, reactor trips and are held above the core during normal operation. Regulating, or trimming control rods in AGRs are driven automatically by auto-controllers that continually attempt to balance the power of the core in relation to a manually set reference value. As the primary method of control and

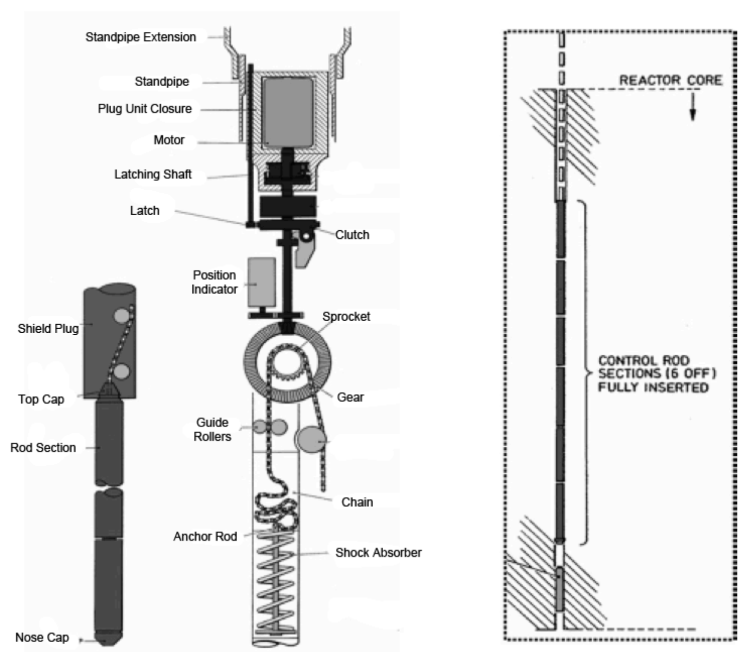

Figure 1. On the left, the design of the upper control rod drive mechanism and on the right, the articulated rod which is inserted into the core.

shutdown system of AGRs, the ability of the rods to move freely into the core is critical to safe operation of the reactor.

With the knowledge that the graphite of which the core is composed suffers damage, there exists the possibility that the motion of the control rods might be restricted. The height data for each rod is currently manually inspected for indications that the rod may be impacting the channel wall. There are, however, several different operational causes of rod motions that appear abnormal on first inspection and require significant amounts of an engineer's time to investigate.

The specific control systems for regulating rods vary somewhat between AGRs, with some stations using sector controllers to move multiple rods simultaneously, whilst others move each rod independently. The specific drive orders issued to each rod are not recorded, therefore there is not a direct method of determining whether a rod is moving as requested.

3. MONITORING

The relative infrequency of channel inspections is augmented by more regular analysis of monitoring data, which is considered at quarterly Monitoring Assessment Panel (MAP) meetings at each station. These meetings have currently only been implemented at the oldest stations, however they are scheduled to be introduced to all AGRs. Observations of various reactor monitoring parameters and information from channel inspections are contained within the Intelligent Monitoring Assessment Panel System (IMAPS) [2]. IMAPS allows an engineer to view the state of the core based on particular observation types, the most recent observations or best estimate of the state of the core on a particular date. Suitably qualified engineers enter observations of a number of reactor parameters in advance of the meeting, including refuelling data, channel power analyses and control rod movements. These observations are then validated at the Monitoring Assessment Panel (MAP) meetings to help determine the current health of the core.

3.1 Control Rod Analysis

As the primary control system in the AGR, the ability of the control rods to move without restriction is critical. The main use of control rod data at the MAP therefore is to attempt to detect instances where the motion of the control rod may have been affected by damage or distortion to the graphite channel. This is currently done manually, whereby a trained engineer attempts to visually detect instances of rod motion that appears to be restricted at a particular height. This process is clearly not optimal, as the process is neither repeatable nor quantitatively useful for other related analyses. The next section therefore describes a method of analysing the motion of the regulating rods, for the case of stations in which the rods are grouped into sectors, within which all of the rods receive the same drive orders.

4. ANALYSIS

4.1 Control Rod Motion

A simple approach to perform the task of detecting restricted rod motion would be to write a filter that detects extended periods at a particular height, however there are several problems with this approach.

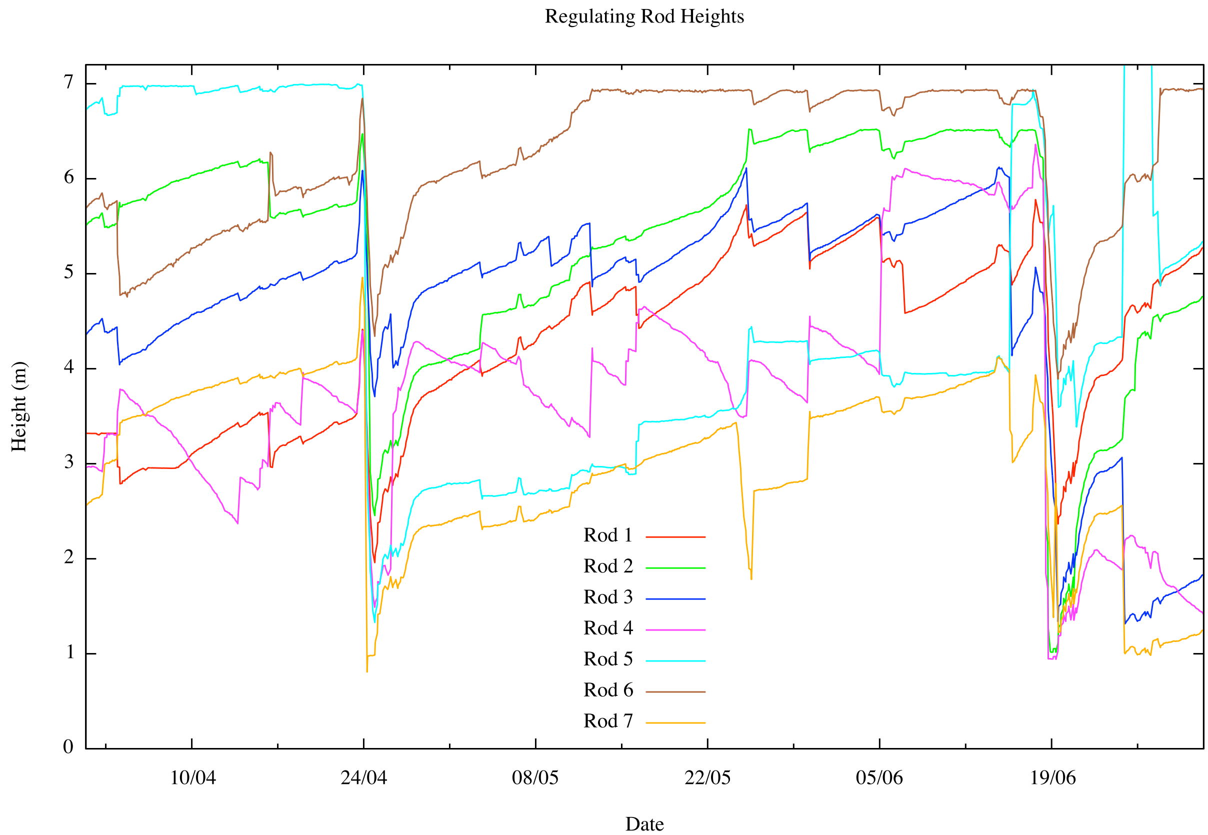

Figure 2. An example set of height data for seven rods in a sector, all of which have been issued the same drive order.

The motion of the control rods, as in Figure 2 is generally upwards after refueling events, which are indicated by the more dramatic motions, centered around 24/4 and 19/6. During refueling, the control rods are inserted further into the core as the power is reduced, and remain at a lower height after refueling compared to before refueling as the new fuel is a stronger neutron source. As the fuel burns up, the release of neutrons slows and less absorption by the control rods is required to maintain fission, and the rods are slowly withdrawn. It can be clearly seen that the motions of the rods are rarely regular. This is a result of the automatic control of the rods and their movement to balance the temperature of the core. As all of the rods have, in principle, been issued the same drive orders, a more comprehensive analysis attempts to estimate the initial drive order then compare the motion of each rod to the drive in order to detect anomalous behavior.

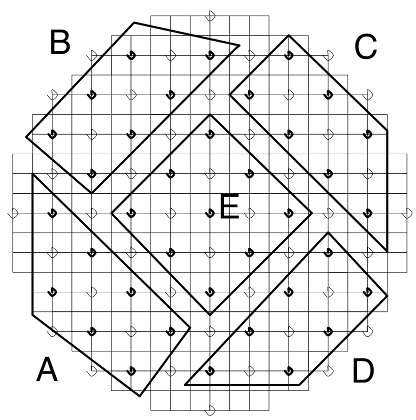

Figure 3. The notional sectors of a core. Each rod within a sector receives the same drive order.

4.2 Sector Drive

The sector drive order is the command issued to a sector of rods, of which there are notionally five per core, with four sectors containing seven rods and the other containing 9. All of the rods are designed to respond to the orders in the same way, that is: in an idealized system, all of the rods would move by the same amount at the same time. These orders are not recorded, so it is not possible to say exactly what drive orders were issued by the auto controller.

An initial estimate of the sector drive order can be made by segmenting the data into smaller windows of one hour and then comparing the net motion of each of the rods in the sector. A set of rules was applied to the segmented data which classified each hour segment as an *upward-drive, downward-drive* or *no-drive* based on the number of rods which drove upwards, drove downwards or did not change in height. The average upwards or downwards drive for each segment is then calculated, which can be considered the first estimate of the sector drive order.

4.3 Correlation

A comparison is then made between the segmented motion and this initial estimate of the sector drive order, by calculating the Spearman's Ranking Correlation (SRC) Coefficient [3], defined as

$$\rho = 1 - \frac{6\sum d_i^2}{n(n^2-1)} \tag{1}$$

Where n is the number of data values in each segment and d is the difference between the ranked values of each data set. The SRC was specifically selected because of the non-linear nature of the motion of the rods. There is significant noise and variance in the data as a result of variations in core temperature, hence the drive order, which is better suited to the SRC, than the more traditional Pearson's Correlation which is more suited to linear data. More specifically, the SRC coefficient is a calculation of how well two data sets can be described by a monotonic function.

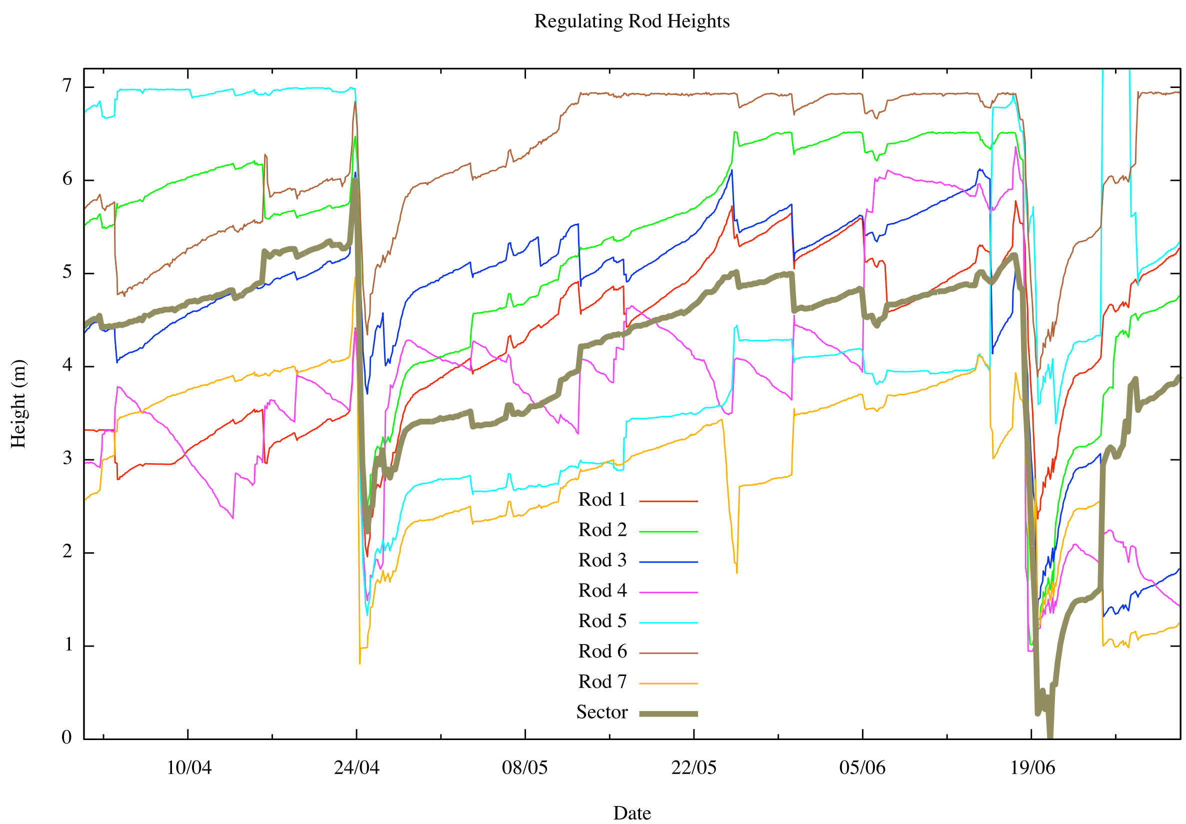

Figure 4. The estimated sector drive of a hypothetical rod with a perfect response compared to the actual motion of the rods in the sector.

4.4 Refined Sector Drive

Now that the correlation of each rod to the estimated sector drive order is known, it is possible to select only those rods with a correlation greater than a particular cut-off value and use this subset of the rods to calculate a more refined average. This cut-off can be set arbitrarily, however several simulations and consultation with previous MAP analysis conclusions and data suggest that this value is best set to between *0.5* to *0.75*, due to noise in the motion of the rods and inefficiencies in the mechanical systems which drive the rods. By recalculating the sector drive order as before, using the subset of rods with a correlation coefficient greater than the cut-off, we can compare the estimated sector drive order for a hypothetical rod with a perfect response to the rods within the sector, as in Figure 4.

4.5 Rod Correlation

With an estimate of the sector drive orders that were issued to each sector, it is now possible to calculate the correlation coefficient for each rod, relative to the sector drive, using the segmented data, in Figure 5.

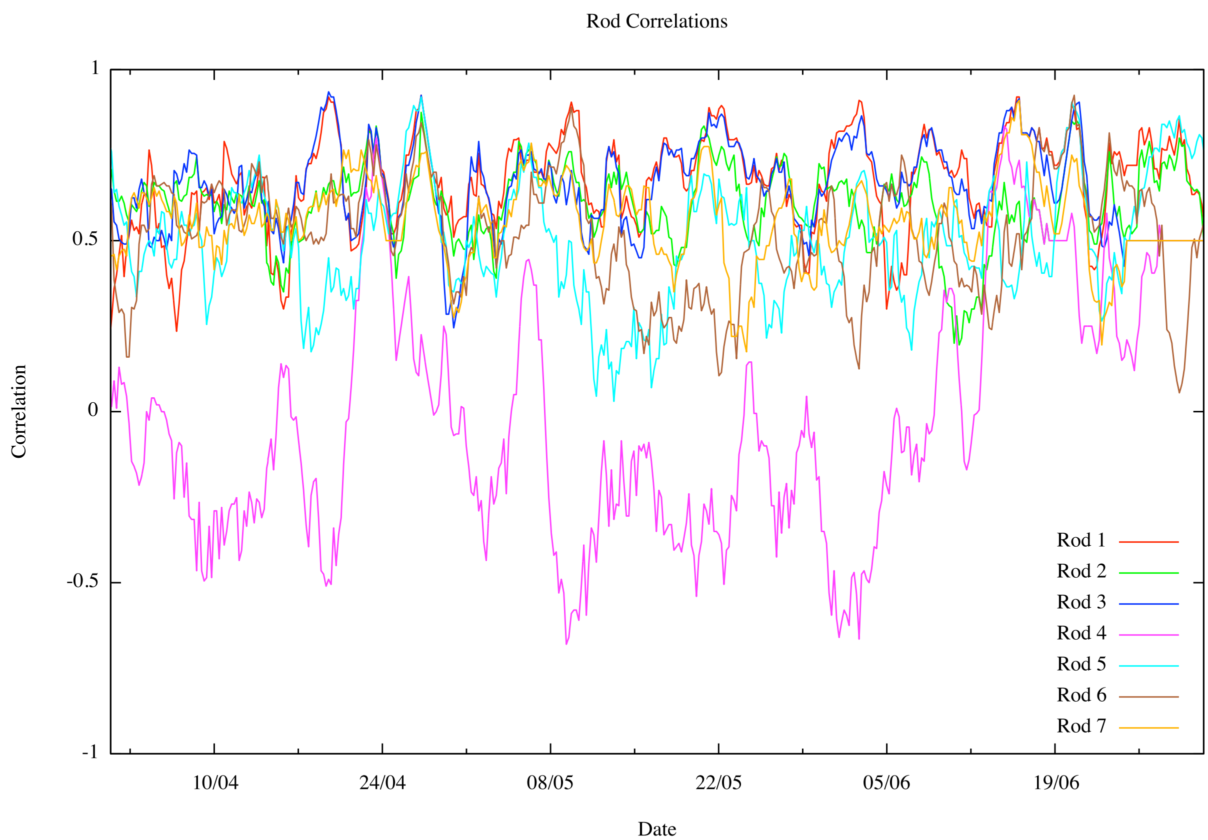

Figure 5. The calculated correlation coefficients for each rod in a sector of 7. Note that the majority of the rods are grouped around the value of *0.5*, with the exception of Rod 4 which consistently experiences negative correlations.

In this example, the correlation plots for the rods are fairly similar, with a few notable exceptions, particularly *rod 6*, which has a correlation consistently lower than the main group, and often negative.

4.6 Interpretation

Figure 5 shows what can be considered the output from the analysis, however the correlations on their own, which may give some indication of anomalous behavior, are of limited use without proper interpretation, specifically:

- A correlation approaching ***+1*** indicates that the motion of the rod matches very closely the estimated sector drive.

- A correlation greater than, but close to ***0*** indicates that the rod moved in the same direction as the estimated sector drive, but at a much slower rate.
- A negative correlation close to ***0*** indicates that the rod moved slightly in the opposite direction from the estimated sector drive.
- A highly negative correlation, approaching ***-1***, indicates that the rod was moving in the opposite direction from the estimated sector drive at close to the same rate,

It is recognized that these descriptions are relatively subjective, but for the purposes of this initial piece of work on control rod analysis, the results are intended only to indicate the behavior that is likely to exist rather than make specific quantifiable classifications.

5. RESULTS

5.1 Practical Implementation

Figure 6 presents a more practical example of how the correlations, which are calculated for each rod, might be more effectively presented to the user. This example illustrates how the correlations can be interpreted with respect to the relevant height data.

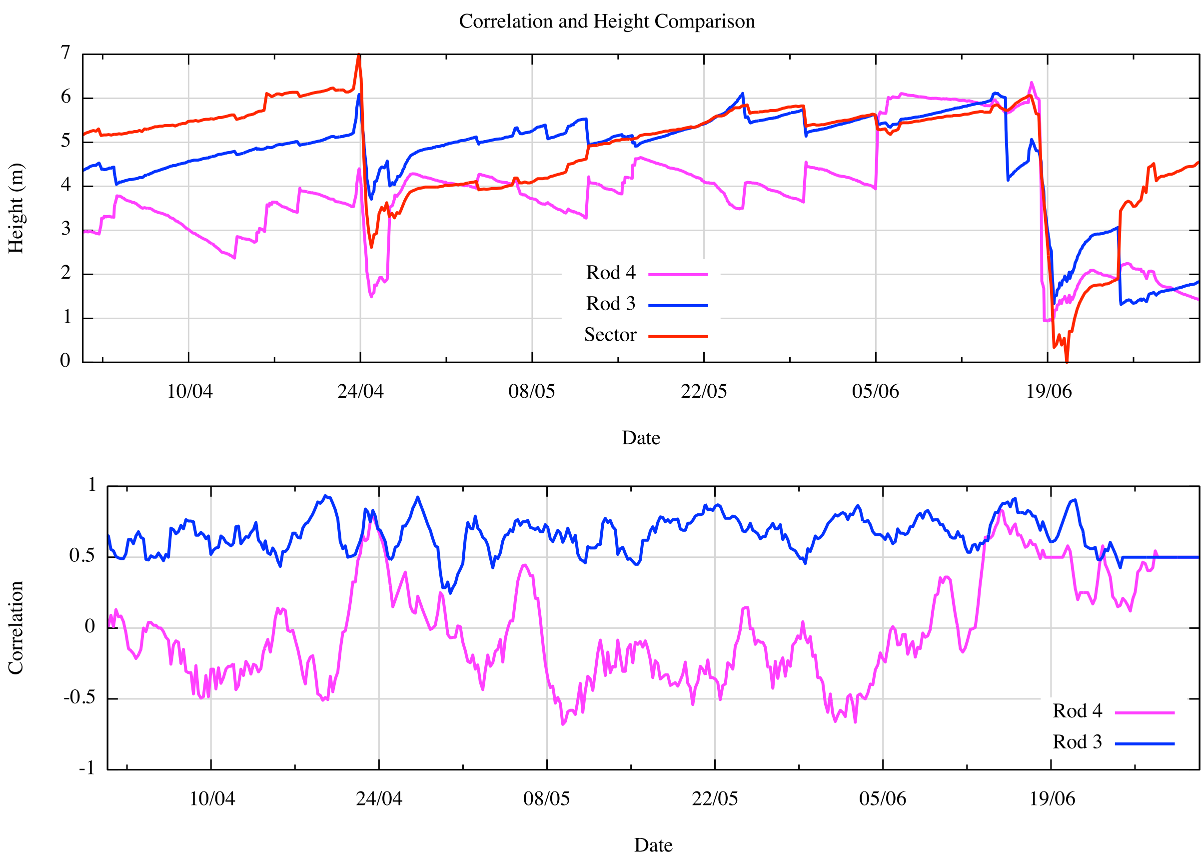

Figure 6. An example of how the calculated correlations can be used to indicate how well the motion of a rod matches the estimated sector drive.

In this example there can be considered to be two types of operation, specifically the motion of the rods during refueling and the motion between refueling. The example presented here, contains two refueling events, centered on *04/2010* and *06/2010* respectively and is indicated by the rapid change in height as the rods are inserted into the core.

In the region before the first refueling, the estimated sector drive trends upwards, as does *rod 3*, which is reflected by positive correlations of around *0.5*, however *rod 4* moves downwards over almost the entire period, with some short, rapid, upward drives that are likely manual corrections to their position by the operators. The downward motion of *rod 4,* relative to the upward motion of the sector drive, is reflected by the generally negative correlations (between *0 and -0.5*), however during the short, rapid, upward drives that also occur in the estimated sector motion, the correlation noticeably increases.

During both refueling events, the correlations of both rods are around *0.75*, suggesting that when under manual control (i.e. not the small, corrective motions of the auto-controller), the rods respond better to the drive orders.

In the larger interval between the refueling events, the responses of both rods considered in Figure 6 are similar to their respective responses prior to refueling. In the case of *rod 3*, the correlation with the estimated sector drive remains largely between *0.5* and *1*, indicating consistent response to the drive orders. *Rod 4* continues to slip into the core, with the exception of short upward drives, for which there are corresponding increases in correlation, however the constant nature of the slipping suggests that a mechanical fault in the drive mechanism is the reason for this rod's poor response.

6. DISCUSSION

6.1 Physical Considerations

In principle the regulating rods are moved automatically to balance the core power, however in practice there are often operational reasons for moving rods manually or rods not moving because they are at the limits of the core. Actions and limits such as these complicate the calculation of the correlation of the rods to the estimated sector drive because they do not represent anomalous behavior however the calculated correlations would indicate that the rods are not moving as expected. Small corrections were therefore required in order to ensure that the impression given by the correlations was as reflective of the motion of the rods relative to the drive orders as possible. It has not yet however been possible to entirely remove the behavior of operators who may force a rod to remain at a specific height for operational reasons.

6.2 Structural

Though the initial motivation for this work was to attempt to detect rod limiting, the correlation plots for the data analyzed thus far have not definitively identified any rod that appears to be axial limiting for which there has not been another explanation, such as a mis-set limit, a blown fuse on the drive assembly or some other reason. There are two main aspects to this problem: firstly, the clearances within the fuel control channels are relatively large and it is quite likely that such an event has not occurred yet and secondly, the analysis does not yet feature the ability to produce quantifiable classifications of behavior, but rather gives an indication of the 'behavior' of the rod. As a basis for further analyses however, the technique presented here provide an intital approach for automatically detecting abnormal control rod behavior and with continual degradation and ageing of the graphite core, the probability of a control rod experiencing restriction will increase.

6.3 Performance

The analysis has produced particularly clear representations of rods that are not driving as would be expected over extended periods, as opposed to a rod that remains at a fixed height. Rod slipping for example, whereby the drive mechanism allows the rod to slip slowly back into the core while under auto control, is quite evident in the examples presented here. The ability to quantify such performance could possibly lead to more predictive maintenance and long-term health trending of the control rods.

7. FURTHER WORK

7.1 Refinement

The most important requirement from the next stage of this work is to derive quantifiable classifications from the correlation calculations described here. This will involve the use of a large set of historical data for which there are known and validated manual analysis results. The system can then be trained to use rule sets to classify particular behaviors and return diagnoses to the user.

7.2 Implementation

As mention earlier, various stations use different auto control systems, often with each regulating rod moving independently. The technique developed here for estimating drive orders based on the motion of a group of rods is not directly applicable, however a more generalized description of the normal motion of the rods between refueling events coupled with some historical knowledge of each rod should allow for a similar analysis.

A desktop version of the software is being developed that can be used with extracts of control rod height data to automatically generate reports to compliment the existing trending analysis. Initially it is expected that these reports might be used for comparison at the MAP and eventually may aid the engineer in detecting anomalous motion in the rods that may go un-noticed by the current analysis.

7.3 Distributed Analysis

There is only very limited interaction between the analyses of different datasets within AGR core condition monitoring, often despite qualitatively understood physical relationships between related core systems. It has therefore been proposed that an expanded system of monitoring across other datasets, coupled with tighter integration of existing analyses and available data along with the introduction of data mining techniques could increase the quality and quantity of useful information about the core available from monitoring. In order to facilitate the implementation of such a system and given the large amount of work generally involved in installation or modification of equipment in any nuclear power station, this system will be designed to be distributed across several different networks and locations, using existing interfaces.

A solution to the lack of interoperability between various existing condition monitoring is currently being investigated by means of a Multi-Agent System (MAS). MAS have already been proposed for use in Generation IV reactors [4], however given the distributed nature of the problem of nuclear condition monitoring as a result of ageing IT infrastructures and regulatory limitations on the installation of new equipment, MAS could play a pivotal role in managing analyses and information. The interfaces of both existing systems and new analyses, such as the control rod monitoring example presented here could be contained within Intelligent Agents, that interact, share information and cooperate autonomously. These agents would be developed for the analysis of each different data set and would allow for a flexible, distributed system to which new analyses could be added as required.

Initial research at the University Of Strathclyde has already seen intelligent agents developed for communication between the IMAPS system and the British Energy Trace Analysis (BETA) system[5], that performs automated analysis on refueling data, and for performing the analysis of the control rods described earlier. Further work is underway to develop other IA that encompasses the existing MAP analyses and data retrieval and exchange.

8. CONCLUSION

This paper has described a method for detecting anomalous behaviour in the motion of AGR regulating rods by comparing the rod motion to the estimated drive order issued to the rod. A case study was presented which illustrates that the technique can be used as a basis for a system that automatically diagnoses and classifies the behaviour of the rods during operation. The work also illustrates an example of condition monitoring on a critical core system without the need for extra monitoring equipment.

9. ACKNOWLEDGMENTS

This work was funded by EDF Energy and EPSRC. The views presented in this work do not necessarily represent the views of EDF Energy.

10. REFERENCES

1. A.G. Steer, "*AGR Core Design, Operation and Safety Functions*" *Management of Ageing Processes in Graphite Reactor Cores*, University of Cardiff, United Kingdom, 28th-30th November 2005, Royal Society of Chemistry Publishing, pp11-18
2. Jahn, G. McArthur, S. Reed, J. Towle, D. "*Staged Implementation of an Agent Based Advanced Gas-Cooled Reactor Condition Monitoring System",* IEEE Power Energy Systems General Meeting, 2007, pp1-4
3. B. Ostle, L. C. Malone, "*Statistics in Research: Basic Concepts and Techniques For Research Workers".* 1988, Wiley Blackwell, pp217
4. R. Uhrig, L. Tsoukalas, "*Multi-Agent-Based Anticipatory Control For Enhancing The Safety and Performance of Generation-IV Nuclear Power Plants During Long-Term Semi-Autonomous Operation"* Progress In Nuclear Energy. 2003;43(1):113-120.
5. G. West, G. Jahn, S. McArthur et al. "*Data Mining Reactor Fuel Grab Load Trace to Support Nuclear Core Condition Monitoring"* IEEE Transactions on Nuclear Science. 2006, 53