Towards an Online Prognostic System for Predicting the Axial Shrinkage of AGR Cores

Graeme M. West[1], Christopher J. Wallace[2], and Stephen D. J. McArthur[3]

[1,2,3]*Institute for Energy and Environment, Dept. EEE, University of Strathclyde, Glasgow, G1 1XW, UK*

graeme.west@strath.ac.uk
christopher.wallace@strath.ac.uk
s.mcarthur@strath.ac.uk

ABSTRACT

In the UK, there is the desire to extend the operation of the Advanced Gas-cooled Reactor (AGR) power plants beyond their initial design lifetimes of 35 years. As part of the justification of extended operation, an increased understanding of the current and future health of the graphite reactor cores is required. One measure of the health of the AGR power plants is the axial height of the graphite core, which can be determined through measurements undertaken during statutory outages. These measurements are currently used to manually make predictions about the future height of the core, through identifying the relevant data sources, extracting the relevant parameters and generating the predictions is time-consuming and onerous. This paper explores an online prognostic approach to support these manual predictions, which provides benefits in terms of rapid, updated predictions as soon as new data becomes available. Though the approach is described with reference to a case study of the UK's AGR design of power plant, similar challenges of predicting passive structure health also exist in other designs of power plant with planned license extensions.

1. INTRODUCTION

The condition of the graphite core of an AGR power plant is a major life-limiting factor. The graphite core provides moderation of the nuclear fission process and provides the structure in which to house the uranium fuel and provides pathways for both the gas coolant and for control rods. Three underlying processes govern the degradation of the graphite. Firstly, the neutron moderation process causes dimensional change to the graphite when fission neutrons collide with the nuclei of the carbon atoms. In parallel, the strength of the graphite is altered by both the neutron irradiation, which increases the strength and radiolytic oxidation, which reduces it (Shennan, 1993). The dimensional change in the graphite provides a useful measure of the degradation of the core, and is easier to measure in-situ than the strength of the graphite. Unlike most other components of nuclear power plants (NPP) (for example PWR Reactor Vessel heads and internals and the core pressure tubes in CANDU reactors (IAEA technical report, 2008) the graphite bricks, which comprise the core of an AGR, are irreplaceable. Therefore, predictions of Remaining Useful Life (RUL) for the graphite bricks also dictate to a large extent the remaining useful life of the power plant.

Making predictions about the dimensional change of graphite in AGR stations is not new. For example, Shennan (1983) describes predictions relating to dimensional change made about the graphite in Hinkley Point B power plant in 1979 (Hinkley Point B began generating electricity in February 1976). When the power plants were built, predictions were made based on theoretical and experimental data (as no operational data was available). A lifetime of operation has produced vast quantities of data, and data capture and storage equipment has advanced significantly, first with electronic loggers replacing pen and plotter devices, and then these first generation electronic loggers in turn being replaced by newer models. However, the original plant components being monitored, such as boilers, graphite bricks and pressure vessels have not been replaced. Managing and interpreting this increased volume of data necessitates the introduction of automated analysis techniques to support a limited number of experts in the field, as traditional manual approaches do not scale well. The remainder of the paper is organized as follows – the next section discusses the drivers for prognostics in nuclear power plants, followed by a section discussing the approach adopted within AGR power plants, including the manual method of prediction currently adopted. To investigate the possible use of an online prognostic system to replace this manual approach, three case studies are then presented. The results of these case studies are discussed along with a proposed system development for such an online system, as the next stage of this work.

2. DRIVERS FOR NUCLEAR PROGNOSTICS

In general, the nuclear industry is cautious in the adoption of new technologies for reasons of safety, preferring to utilize existing proven techniques, unless there is no viable alternative. For example, the use of wireless technologies to transmit data is used in many applications, but has yet to find significant use within the nuclear domain (though progress is being made towards this, as reported in Hashemian, 2011). The field of prognostics faces similar challenges, where there are very few examples of deployed prognostics systems in use within nuclear power plants. Varde and Pecht (2012) provide a tabulated summary of the current state of prognostics, grouped by component/system type, and provides information for both online and offline prognostics. The paper indicates that rotating plant prognostics is the most mature area, and that online prognostics is, at best, still in the R&D stages for NPP components. For general reviews of prognostics applied to rotating plant, but not constrained to nuclear applications Jardine, Lin and Banjevic (2006), and Heng et al (2009) both provide useful summaries. In other non-nuclear domains there has been significant work undertaken. For example, Goh (2006) and Kothamasu, Huang and VerDuin (2006) both provide reviews of the application of prognostics to manufacturing. Prognostics in electronic systems is another area with significant progress in prognostics, with Goebel et al (2008) describing an application to battery health and Pecht (2008) providing wider coverage of the domain. Though lagging other industries, there is the recognition that prognostics could provide significant benefits to the nuclear industry, illustrated by projects such as the US light-water reactor sustainability program (US Department of Energy, 2013). Prognostics could be applied to both passive components, such as pipework, pressure vessels and graphite moderator bricks and active components, such as turbine generators and reactor coolant pumps. The first structural degradation prognostic system for nuclear power plants is mentioned in Bond, et al (2011), but this is still at the demonstration, rather than deployment stage. Coble, Humberstone and Hines (2010) have investigated prognostics for simulated data for the new IRIS power plant, and again this is at the technology demonstration stage.

3. APPROACH

The usefulness of prognostic systems depends on the predictive accuracy of the model used and the availability of data to generate RUL estimates from the model. For large fleets of similar assets, models can be based largely or entirely on statistical behavior, if the assumption is made that the future components will degrade in the same manner and the operating conditions are fixed (Coble and Hines, 2008). This is termed Type 1 prognostics. If information relating to the operating environment, which influences the degradation process, then Type II may be applicable (e.g. proportional hazards models, shock models). However where there are a limited number of assets, there is also often a limited quantity of data available across the lifetime of the asset. Commonly where prognostics is applied to high criticality assets, such as a component of a nuclear reactor, it is impractical from either a safety or financial perspective to allow a component to run to failure, and therefore gather sufficient data to create these statistical models.

For this reason, prognostic models for such components often include a physical model of the state of the component, based on material, chemical or some other understanding of the physical processes causing degradation in the asset. These models typically describe a macroscopic behavior of a material (structural strength, dimensional change, etc) based on experimental data or on a microscopic model of the physical processes. Since it is not practical to model microscopic behaviors of the entire asset, the most common approach to overcoming these limitations is to use the physical model as the basis for a prognostic model, estimating macro scale behavior of the asset based on some dominant characteristic of the physical model. The prognostic estimates based purely on the use of such a model are likely to be of limited usefulness, given that the material of which the asset is constructed and the environment is located are unlikely to be identical to those in the experiments used to derive the model.

1.1 Combining Model Based and Data Driven Prognostics

An alternative approach is to use a physical model as the basis of prognostic analysis, but to fuse this with a data driven approach (Pecht and Jaai, 2010) as it becomes available. It is argued that where a general model is applicable, there may be unknown parameters within a particular asset that cause the measured data to diverge from the general model, as in Figure 1. An automated, iterative approach, whereby a revised estimate of the model is generated every time new data is available allows a prognostic model to evolve as the knowledge of the asset increases. The rapid assessment of new data may revise a prognostic model such that a pre-defined limitation will be reached sooner than anticipated based on existing models. It can be seen therefore that such an approach has clear safety benefits.

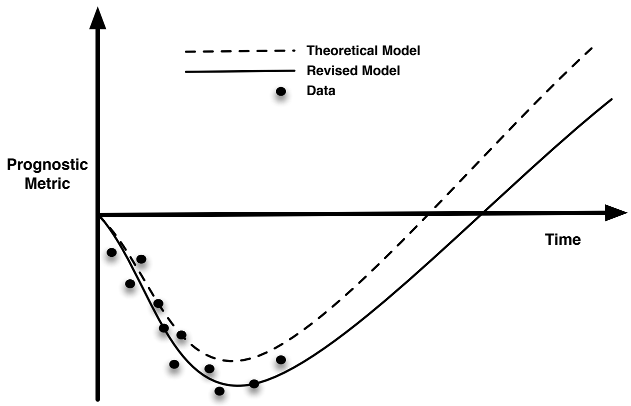

Figure 1. Generalized model of graphite dimensional change

For the graphite core used in the AGR, a model of dimensional change reported in Brocklehurst and Kelly (1993) has been adopted to estimate the evolution of dimensional change in the graphite bricks as they degrade with exposure to radiation. This model is generalized in Figure 1 and is based on previous experimental work and can provide at best a general model of the evolution of graphite under comparable parameters. A large volume of data is available describing various features of AGR cores, from inspection and monitoring activities and it is proposed that this data can supplement existing models of core evolution. Previous work by West, et. Al (2010) explored the possibility of using fuel grab load trace data, a set of monitoring measurements gathered during refuellling operations, to trend the shrinkage of the core. Though this work is related, this paper focuses solely on the use of inspection data gathered during routine outages.

2.2 An evolving prognostic system for the AGR

Based on existing data collection and condition monitoring analyses, it is possible to trend certain characteristics of the AGR core, such as brick heights and channel bore diameters, in order to calculate geometric changes in the structure of the components within the core. Using a statistical sample of different components within the reactor, a core-wide estimate of shrinkage can be derived and compared to existing predictions. Should the shrinkage be significantly smaller or larger than forecast, the prognostic model can be revised, using the most conservative data available so as to maximize safety margins. The estimation of shrinkage is used to support strategic decisions about the health of the core, and whether it can still fulfill its role as fast neutron moderator, and maintain the necessary physical structure to permit the un-impeded movement of fuel and control rods and to allow the passage of coolant gas.

3.1. Existing analysis

During statutory outages of the AGR nuclear power plants, a few fuel channels (typically 20-30) are assessed to determine the current health of the graphite bricks in these fuel channels, which in turn is extrapolated to provide an indication of the state of the whole core. This provides a short-term view of the health of the core and assessment of the data gathered during inspection needs to be assessed before the plant is returned to service. Formal documentation is produced which provides a clear statement of the current condition of each channel inspected. The raw inspection data is retained, though not in a database, but as analysis files containing both the raw data and results of analysis performed on the data.

For a longer-term view of the core, some of this data is collated and trended to provide an indication of the overall degradation in the core and to support statements that on given target dates the core is predicted to be in good health. However, gathering and filtering this data is a labor-intensive task, as suitable channels need to be identified, the associated data files located and the relevant information extracted.

There are two sources of data used to make the predictions. The first set of data is accurate measures of the internal bore of the graphite bricks throughout the height of the fuel channel. This data is gathered through specialist inspection equipment and is made available to the engineers as a raw data file. From this data three parameters are extracted which provide a representation of the shrinkage in the core:

1. A direct measurement of height
2. An average measure of brick shrinkage taken from the mid-points of each brick
3. An average measure of the full channel diameter shrinkage.

The second source of data required is to make the prediction is a measure of the cumulative irradiation the bricks have been subject to. Time is not a suitable measure, as the graphite only degrades while the plant is in use, so when predictions of RUL are made, the duty cycle of the station needs to be factored in. To ensure maximum conservatism, it can be assumed that the station is run at a constant rate, thus ensuring an under-prediction in the RUL. This cumulative irradiation measure is obtained from operational plant data, and is not recorded directly during inspection.

Values for the three measures of shrinkage are plotted against cumulative irradiation. For each set of data both a linear and a second order polynomial line are fitted to the data to provide predictions of future shrinkage. An example of such a manually produce prediction is shown in figure 2.

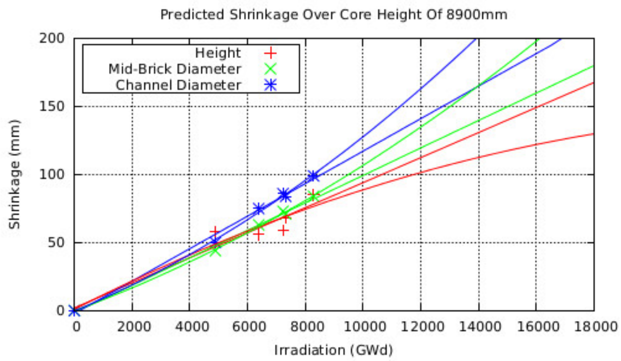

Figure 2. An example of an existing manual prediction of core shrinkage

From all the predictions, the most conservative estimate is used to inform other activities that require the use of these shrinkage predictions. These predictions are necessarily conservative and though first and second order polynomial models are simplified representations, they do provide suitably useful results.

Currently, the existing data management strategy for this inspection data creates a significant barrier for deploying this as an online system, as the primary use for the inspection data is to provide confidence that the channels inspected are in suitable condition to allow the power plant to be returned to service following an outage, and the management of the data is tailored to optimally fulfill this role. However, there is a project currently underway to improve accessibility to the raw, verified data for use by other functions, such as possible online predictions of core degradation. In order to test proposed prognostic approaches before the live station data is made available, simulated data can be used. This has the advantage that the input data and model of core degradation can be tightly controlled and the state of the core simulated at any given level of irradiation, unlike the operational data which can only provide data up to the current level of irradiation. Using simulated data can provide confidence that robust predictions of future state can be made, provided of course that the degradation model is suitably representative of the actual core degradation.

A series of case studies have been developed to demonstrate the possible online prognostic capability that could be achieved if validated inspection data were made available. It is recognized that several assumptions and simplifications have been made in these case studies, and that the resulting predictions of core shrinkage should not be taken literally, but instead as a means of demonstrating the techniques. Furthermore, the approach has been adopted to mimic the existing manual approach in the first instance, though recognizing that future work could see other prognostic approaches applied.

4. CASE STUDIES

In order to demonstrate the framework, a case study is presented which uses simulated brick degradation data to explore the application of online prognostics. Though the most common output of a prognostic algorithm is a measure of RUL, the objective of these predictions is to estimate the shrinkage of the core at certain levels of future cumulative irradiation. MATLAB R2012b was used to implement these case studies.

4.1. Approach

Based on the manual analysis mentioned in the previous section a set of simulated data was created which would provide an approximation of the behavior of graphite degradation. It should be noted that the purpose of this data was to generate a suitable volume of data to explore the prognostic algorithms rather than being an accurate representation of the degradation of graphite. Using a suitable 2nd order polynomial, shown in Eq. (1), and approximated from the manual analysis data, a set of graphite shrinkage values was generated for a series of data ranging from no irradiation to 16,000GWd accumulative core irradiation, a reasonable estimate for a lifetime of operation of an AGR core (based on existing analysis).

$$s = -0.00000022559i^2 + 0.01110778967i \qquad (1)$$

A component of noise was added to this signal to simulate errors in measurement. This component was based on a normal distribution and a new noisy component was generated and applied each time the program was run.

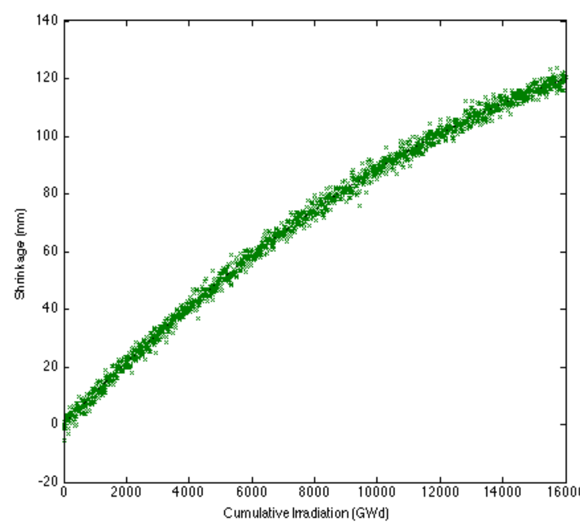

Figure 3. Simulated core shrinkage data including a noise component

Figure 3 shows the output from one such run. A second order polynomial was fitted to the noisy data and was compared to the original simulated data to provide a measure of the quality of fit that would be obtained if all the data were used to make the prediction. In this case a least squares error value of 0.99987 was obtained, reflecting a good fit to the mode from the data. In particular, the value of shrinkage at 16,000GWd was 119.9736mm from the original simulated data and 119.7853mm from the estimate derived from the noisy data, for this particular instance. In general, the difference is in the region of ±0.2% accuracy, which provides an indication of the maximum possible accuracy that could be expected of the predictions.

4.2. Case 1: Train and predict

Using the noisy data set, the first 250 data points were used to fit a 2nd order polynomial and this was used to make a prediction of the shrinkage at 16,000GWd. Figure 4 shows the raw training data used, the prediction and the actual response using an example noisy data set.

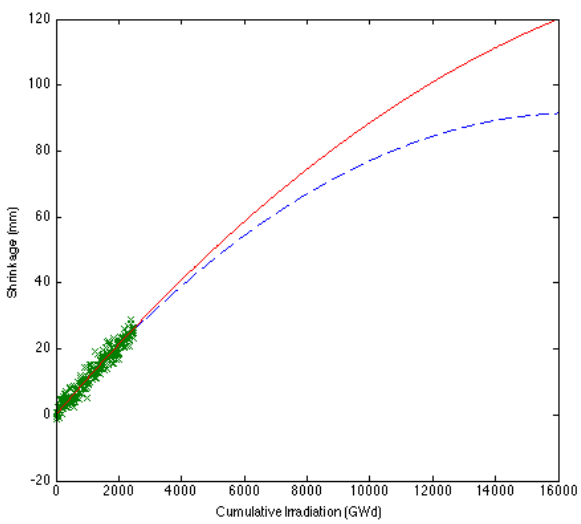

Figure 4. Predicted shrinkage response based on the first 250 measured data points shown in the dotted line. The idealized response is shown as the solid line.

The choice of the initial training set size will have an effect on the early predictions, but with each new measurement beyond the initial training set, the predictive model (2nd order polynomial) could be revised using all available data. For example, Figure 5 shows the progress of the prediction of shrinkage at 16,000GWd as more of the raw noisy data is available.

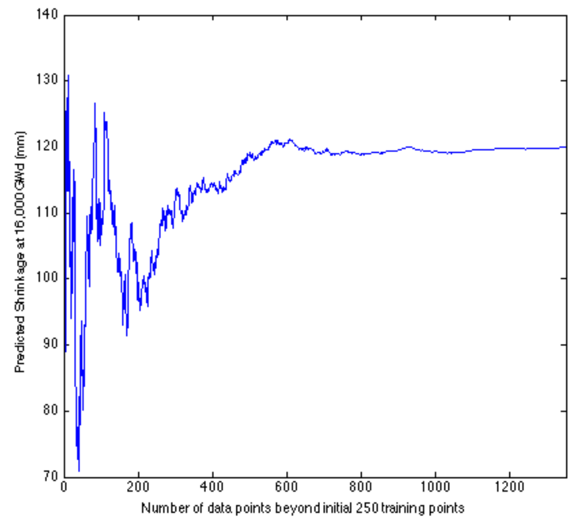

Figure 5. Progress of predicted shrinkage as more points are considered, demonstrating convergence on 120mm

In this instance, after about 600 additional data points above the initial 250 training points (850 data points total), the prediction has converged on the actual final shrinkage of 120mm. The rate of convergence will depend on the initial data set used, and figure 6 illustrates the case where 20 example noisy responses have been generated from the same underlying source signal.

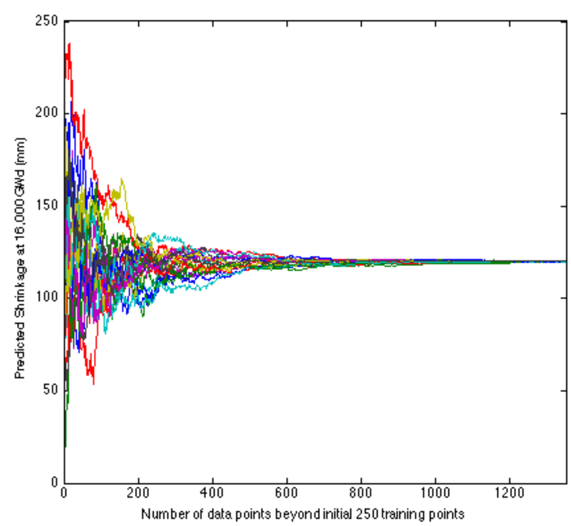

Figure 6. Illustration of responses from 20 different sets of input data

4.3. Case 2: Response to an artificial outlier

The second case investigated was to introduce an artificial outlier to investigate its response to the overall predictions made. Using the same raw data an error of +30mm was introduced to measurement point 900 and this is shown in figure 7.

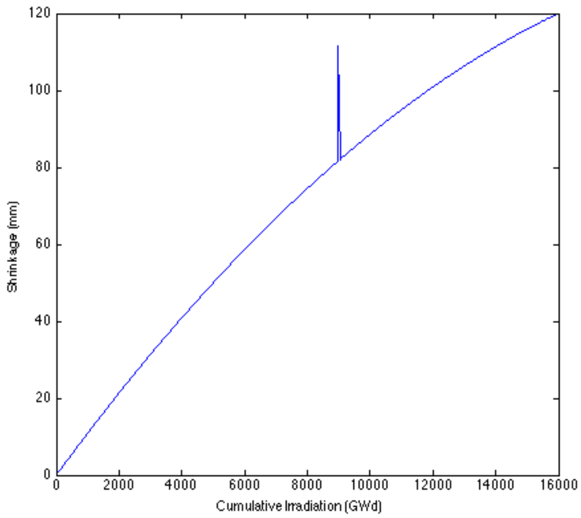

Figure 7. Raw data with an artificial outlier inserted

The corresponding response of how this affected the predicted final shrinkage at 16,000GWd was generated and as expected this did not have a significant effect on the prediction. Figure 8 shows a small perturbation caused by the introduction of the outlier at point 650 on the x-axis, but the influence of this single outlier on the overall prediction is quickly dominated by the contribution of the other data points.

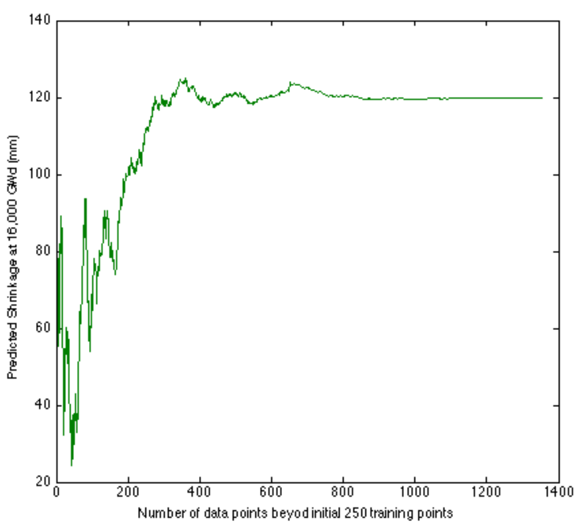

Figure 8. Progress of predicted shrinkage

4.4. Case 3: Simulation of change in underlying model

There are a number of possible reasons that the underlying degradation model may change. For example, a change in the operating temperature of the reactor may affect the underlying shrinkage rate, as might other factors such as the planned injection of methane and carbon monoxide to inhibit the rate of radiolytic oxidation. In other prognostic applications, maintenance actions could also result in a step change in the underlying degradation model.

To simulate a change in underlying model a set of data was generated based on one polynomial model for the first half of the data and switching to a second polynomial model for the second half of the data. Figure 9 shows a plot of the ideal data. The solid line represents the initial polynomial model and is extended through to 16,000GWd. The dotted line simulates a change in the underlying degradation response at 8,000 GWd and shows the new degradation path.

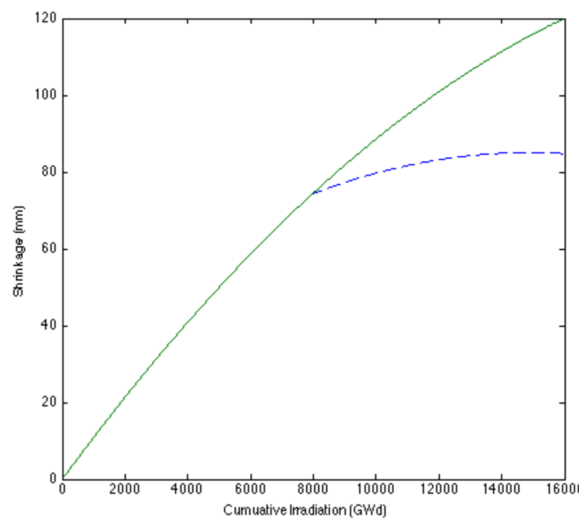

Figure 9. Input data containing a change of underlying model at 8000GWd

As with the previous cases, random noise was applied to the ideal signal, the first 250 points used as training data and a set of predictions calculated based on the available data. The results are shown in Figure 10. As before, the predictions converge towards the 120mm value from the first polynomial model then following the change of model converge towards the second value of approximately 80mm.

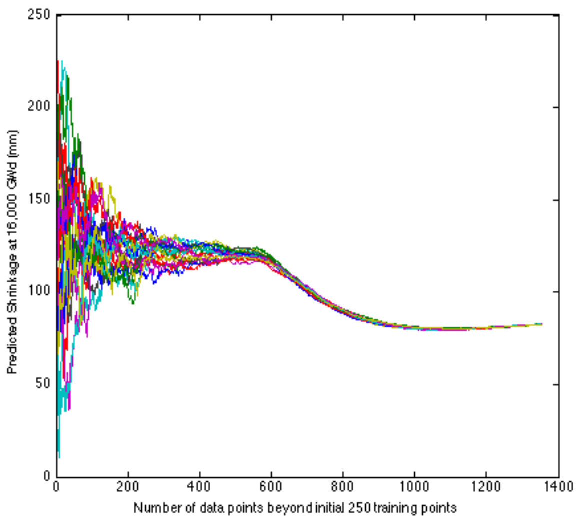

Figure 10. Progress of prediction in response to change of underlying model from 20 sample cases

5. Discussion

The case studies have demonstrated the applicability of prognostics to a continuous feed of measured shrinkage data. A concern with an online prognostic system which continuously updates its prediction as new data is made available is the concern that an erroneous input measurement might significantly alter the prediction. The case study which examined this contained an extreme outlier measurement as an input and though it did affect the final prediction slightly, its affect was minimized by the volume of other measurement data. When there was a change in the underlying model, there was a lag in dealing with the change in response. As an alternative, if the point at which the underlying model changed was known, then a new predictive model could be built from that point forward. However, this new model would require a period of time to build up a suitable number of training points to be able to predict forward accurately. Another option might be to run the prediction based on the full set of available data in tandem with the data from the change point onwards then switch across to the second model once enough seed data has been gathered. It should be noted that for the graphite core, these underlying changes would be very infrequent so this change would not need to be undertaken very often.

5.1. System Development

In order to create an on-line prognostic system the following elements will be required by the system:

- Data management strategy: The prognostics system needs to be able to access the appropriate input data streams as they become available. In particular, the proper QA grade of the input needs to be assured.
- Suitable predictive algorithm
- Appropriate feedback mechanisms: as new data is fed to the system, this data can be used to either increase confidence in the predictions made, can be used to enhance the predictive model or can be used to identify an error in the data gathering. Ensuring the system is robust and able to deal with the possibility that the input data is erroneous is also important.

5.1.1. Data Management Strategy

This is critical to the successful deployment of an online prognostic system. Often data comes from existing monitoring systems, but the prognostic system may require input from a number of different monitoring systems, which were not designed to exchange data. In our case study, two sources of information are required, the cumulative damage and the measures of bore diameter. Neither of these are available as an existing consolidated data source, so this needs to be addressed.

A related issue with developing an online predictive system is the need to deal with historical data. If an online prognostic system were being implemented in a new power plant, then this could be built into the initial design. However, with the AGR power plants (and equally with a large number of other designs of plant) they were built many years ago and the data capture and storage technologies were very different. Including this legacy data is important as it does provide a good baseline reference and useful input to the predictive models.

5.1.2. Choice of predictive algorithm

There is an assumption that the chosen prognostic model is representative of the graphite degradation. In order to be deployed as an online system then some of these assumptions will need to be addressed. Li, Marsden & Fok (2004) describe in detail the relationship between the bore profile (the input to the predictive model) and irradiation induced dimensional change, and should be used as the basis for building a more complete predictive model.

5.1.3. Feedback mechanism

Both the choice of predictive algorithm and also the measure of confidence are important for the successful deployment of an online predictive system. The predictive algorithm should be capable of predicting remaining useful life as the final output of the system, but should also be able to predict interim locations along the degradation curve to allow new operational data to be compared to predictions to assess their accuracy. This provides the opportunity to update the prognostic model if the operational data and predicted data are significantly different. It also provides the opportunity to attribute a measure of confidence to the predictions that are made.

6. CONCLUSION

The condition of core of the AGR nuclear power plant is a major life-limiting factor and being able to make predictions about its future health is important for continues and extended plant operation. Currently these predictions are undertaken manually. This paper has explored the potential for implementing an online predictive system by providing updated predictions of core shrinkage, a measure of core health, as and when new inspection data becomes available. In the first instance, the prognostic approach closely follows the manual approach to predicting shrinkage as the authors feel that this will ease the transition towards acceptance of an online prognostic approach. However, it is recognized that there are other state of the art techniques which could have application here, and these should be investigated in future. Case studies of an erroneous measurement and a change in the underlying degradation model have been explored through the use of simulated data and a discussion provided as to how these results could be incorporated into a system which would allow online prognostics to be performed.

ACKNOWLEDGEMENT

This work was funded by EDF Energy. The views presented by the authors do not represent the views of EDF Energy

REFERENCES

Brocklehurst, J. E. & Kelly, B. T. (1993), Analysis of the dimensional changes and structural changes in polycrystalline graphite under fast neutron irradiation, *Carbon*, Vol 31 Issue 1, pp. 155-178.

Coble J. B., and Hines, J. W. (2008), Prognostic algorithm categorization with PHM Challenge application, *International Prognostic and Health Management Conference 2008,* p1-11

Coble, J. B., Humberstone, M and Hines, J. W, (2010) Adaptive monitoring, fault detection and diagnostics, and prognostics system for the IRIS nuclear plant, *International Prognostic and Health Management Conference 2010*

Goebel, K., Saha B., Saxena, A., Celaya, J. and Christophersen, J. (2008) “Prognostics in battery health man- agement,” IEEE Instrumentation & Measurement Magazine, vol. 11, no. 4, pp. 33–40

Goh, K.M. (2006), A review of Research in Manufacturing Prognostics, *IEEE Conference on Industrial Informatics*, 16-18 August 2006, Singapore, pp417-422

Hashemian, H. M., Kiger, C. J., Morton, G. W. & Shumaker, B. D. (2011), Wireless Sensor Applications in Nuclear Power Plants, *Nuclear Technology,* Vol. 173, No. 1, pp 8-16

Heng, A., Zhang, S., Tan, A. C. C and Mathew, J (2009) Rotating machinery prognostics: State of the art, challenges and opportunities, *Mechanical Systems and Signal Processing*, Volume 23, Issue 3, p. 724–739

IAEA (2008), Heavy Component Replacement in Nuclear Power Plants: Experience and Guidelines, IAEA Nuclear Energy Series No.NP-T-3.2, Vienna (2008)

Jardine, A.K.S., Lin, D and Banjevic, D. (2006), A review on machinery diagnostics and prognostics implementing condition-based maintenance, *Mechanical Systems and Signal Processing*, Volume 20, Issue 7, p. 1483-1510.

Kothamasu, R., S.H. Huang, and W.H. VerDuin (2006), “System Health Monitoring and Prognostics – A Review of Current Paradigms and Practices,” International Journal of Advanced Manufacturing Technology 28: 1012 – 1024.

Li, H., Marsden, B. J. & Fok, S. L. (2004), Relationship between nuclear graphite moderator brick bore profile measurement and irradiation-induced dimensional change, *Nuclear Engineering and Design,* No. 232 pp.237-247

Pecht, M. (2008), Prognostics and Health Management of Electronics, Wiley, Interscience, New York, NY, 2008.

Pecht, M. and Jaai, R. (2010), A prognostics and health management roadmap for information and electronics-rich systems, *Microelectronics Reliability*, Volume 50, Issue 3, p. 317–323

Shennan, J.V. (1983), Graphite R&D reveals long life for AGRS. *ATOM,* No. 323, pp.188-191

Bond, L.J., Ramuhalli, P., Tawfik, M.S. & Lybeck, N.J. (2011), Prognostics and life beyond 60 years for nuclear power plants. *IEEE Conference on Prognostics and Health Management (PHM), 2011,* pp.1,7, June 20-23, doi: 10.1109/ICPHM.2011.6024316

U.S. Department of Energy (2013) Light Water Reactor Sustainability Program: Integrated Program Plan, INL/EXT-11-23452, Revision 1

Varde, P. V. and Pecht, M. G., (2012) Role of Prognostics in Support of Integrated Risk-based Engineering in Nuclear Power Plant Safety, *International Journal of Prognostics and Health Management*, ISSN 2153-2648, 2012 008

West, G.M., Wallace, C.J., Jahn, G. J., McArthur, S. D.J., & Towle, D. (2010), Predicting the ageing of advanced gas-cooled reactor (AGR) graphite bricks. *Seventh American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation, Control and Human-Machine Interface Technologies,* Las Vegas Nov. 2010.

BIOGRAPHIES

Graeme M. West is a Research Fellow at the Institute for Energy and Environment in the Department of Electronic and Electrical Engineering at the University of Strathclyde. He received a BEng (Hons) degree in Electrical and Mechanical Engineering and a PhD in Electrical Engineering in 1998 and in 2002 respectively, both from the University of Strathclyde.

His current research interests include intelligent system applications and data mining within power engineering and in particular applying the techniques to applications within the nuclear industry.

Christopher J. Wallace is a Research Assistant at the University of Strathclyde researching the application of Multi Agent Systems for data mining and data fusion of condition monitoring data from nuclear power stations. His research interests include the application of AI techniques to condition monitoring and engineering decision support and web technologies for human interface.

Stephen D. J. McArthur received his B.Eng. (Hons) and PhD degrees from the University of Strathclyde in 1992 and 1996 respectively. He is a Professor in the Institute for Energy and Environment, within the Department of Electronic and Electrical Engineering. He is Director of the EDF Energy Advanced Diagnostics Centre. He chaired the IEEE PES Working Group on Multi-agent Systems and the IEEE Intelligent Systems subcommittee. His research interests include intelligent system applications in power, condition monitoring, fault diagnosis, smart grids and multi-agent systems.