Transactions, SMiRT-23
Manchester, United Kingdom - August 10-14, 2015
Division III, Paper ID 544

HOW AGR CORE MODELLING SUPPORTS REACTOR SAFETY CASES

Tom Bush[1]

[1] Senior Structural Analyst, AMEC Foster Wheeler, UK

ABSTRACT

EDF Energy can only operate an Advanced Gas-cooled Reactor (AGR) if the safety of operation is demonstrated by a safety case. The graphite core of the AGR design has the potential to be the life-limiting component. The knowledge, understanding and confidence that core modelling provides is a vital component in producing safety cases; therefore it is unlikely that it would be possible to justify operating the AGRs without the advantage of core modelling. This paper explores the part that core modelling plays in expanding knowledge and reducing uncertainties, illustrated using some examples of how it has been used effectively to underwrite safe reactor operation.

While core inspections and monitoring are employed to test the theoretical understanding of the condition of the AGR cores, such empirical methods cannot be used to predict future core states. Core modelling, both computational and physical, provides one of the few means of anticipating the possible future states of the AGRs. However, simplifying assumptions and hypotheses are inherent in any computational model, and if they are poorly understood then quantifying the uncertainties associated with a model can be difficult. This paper describes some strategies for reducing the uncertainties associated with computational models.

INTRODUCTION

This paper contains a discussion of the role that core modelling, both computational and physical, plays in supporting AGR core safety cases. The paper should provide an overview of certain aspects of the graphite core project, which is intended to highlight some of the benefits that core modelling provides. Although the scope of the examples used is limited, it is hoped that they serve to make some points that are more generally applicable.

AGR GRAPHITE CORE SAFETY CASES

For EDF Energy to be able to operate an AGR it must produce a safety case, which demonstrates that it can be operated safely. Safety cases are submitted to the UK Office for Nuclear Regulation (ONR) for consent to operate. Typically a safety case covers a specific period of time, in which case it must be updated periodically.

Principally, an AGR graphite core must be able to meet the Fundamental Nuclear Safety Requirements (FNSR), which are that it:

- Allows unimpeded movement of control rods and fuel assemblies, Bradford (2007).
- Directs gas flows to ensure adequate cooling of the fuel and core, Bradford (2007).
- Provides neutron moderation and thermal inertia.

To produce a robust graphite core safety case, EDF Energy has adopted the framework of a ‘six-leg’ safety case. The six legs are:

- The Core Component Condition Assessment (CCCA) leg, i.e. the state of the aged components.
- The Core Damage Tolerance Assessment (DTA) leg, i.e. the state of the core with aged components.
- The Core Inspection leg.
- The Core Monitoring leg.
- The Consequences Assessment leg.
- The ALARP Assessment of Plant Modifications leg.

Each leg of the safety case focuses on different ways of understanding core behaviour and may include the use of different methods, for example making computational predictions or acquiring empirical evidence. The evidence produced in each leg complements, and can be used to reduce potential uncertainties in, the remaining legs. This allows a robust safety case to be constructed.

CORE DESIGN

Although the AGR core designs are not identical, they do share certain features. The assembly is typically a 16-sided structure with approximately 12 layers of graphite bricks; each layer approximately one metre high. There are two main types of brick: fuel bricks and interstitial bricks. The bricks are stacked vertically in columns and linked horizontally by a keying system. The fuel brick columns house fuel assemblies, while the interstitial columns house control rods. A typical view of part of a core assembly is shown in Figure 1. In the plan view shown, the inner 308 channels are the fuelled channels.

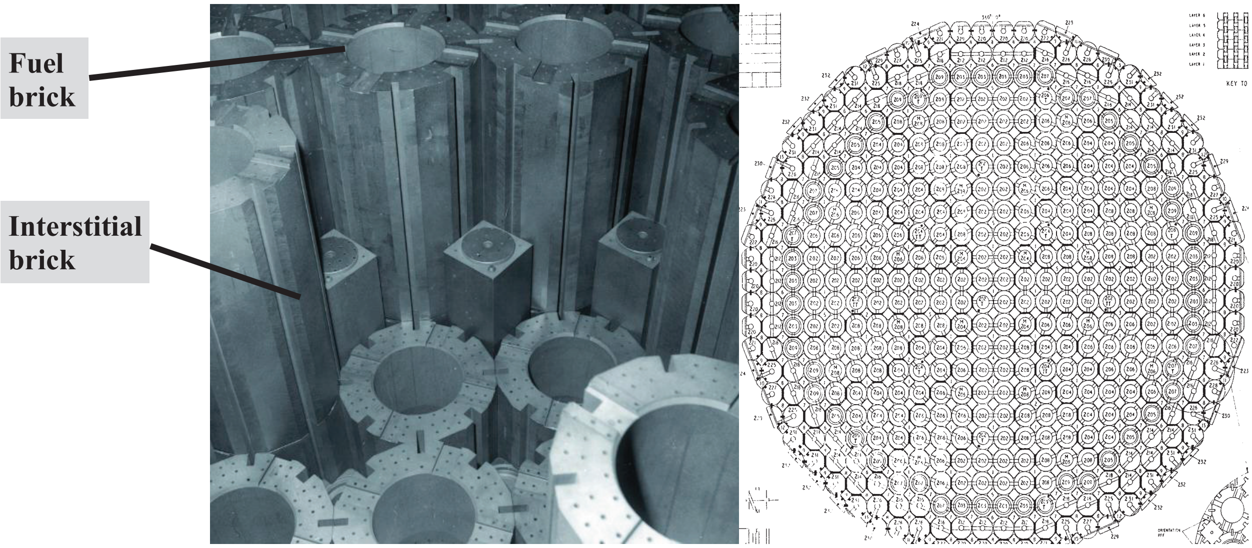

Figure 1. Views of an AGR graphite core assembly. Isometric (left) and plan (right) views.

Understanding the ageing of the graphite core is a difficult task because not only is the design an intricate one, with a complicated set of key-keyway interactions, but the component properties are also changing with time, due to the effects of irradiation and oxidation, see Figure 2. A further complication is that the rate-of-change of a component's properties varies with time depending on its location in the core. For example, the rate-of-change is greatest for components in the most highly-dosed layers of the core, or those closest to the fuel elements.

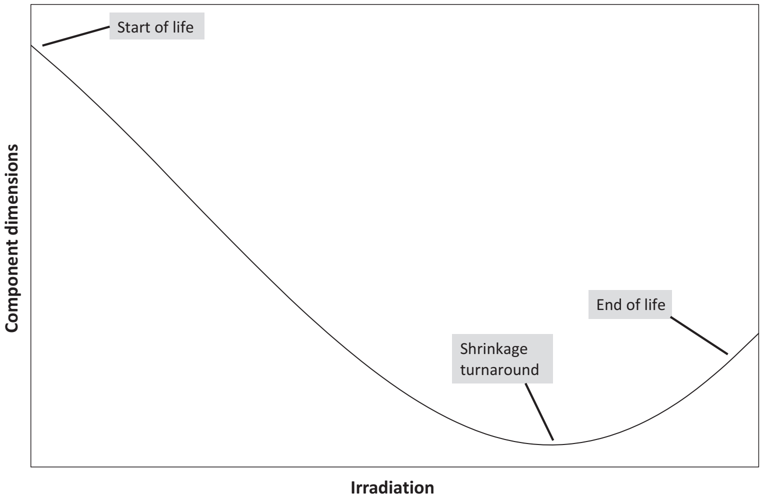

Figure 2. Graphite dimensional change with irradiation.

CORE MODELLING

A theoretical understanding of the behaviour of ageing graphite in AGRs and a good understanding of the design of each AGR core means that engineers and scientists have an expectation of how the AGR cores will behave in the future. To an extent these expectations can be tested by regular comparison with empirical evidence, but observations are of only limited use for predicting the future. Nevertheless, making predictions about the future behaviour of the cores - one of the primary benefits that core modelling provides - is important in making safety cases because it is necessary to demonstrate that the AGRs are safe to operate for some time into the future. Core modelling allows engineers and scientists to explore the operational limits of the cores, beyond what is observable under normal operating conditions.

The ageing effects of nuclear graphite manifest themselves in a variety of ways, such as changes in material properties, weight loss, dimensional changes, brick cracking and changes in core geometry. These changes mean that it is a significant challenge for the cores to continue to meet the FNSR; a challenge which will become ever greater.

SAFETY SPACE

Numerous parameters are important in determining the status of an AGR core. In a sense the state of the cores as a function of the operational parameters can be regarded as a multidimensional landscape or ‘safety space’, comprising regions that are safe or potentially unsafe with respect to the FNSR. But, while there could hypothetically be a function of all the parameters, that could be used to calculate the status of the core, such a function cannot be defined due to inadequate knowledge and information.

One of the ways in which the graphite core project addresses the problem is by dividing core behaviour studies into sub-disciplines which can be understood more easily in isolation. Examples of sub-disciplines include the different legs of a safety case. But within an individual leg, such as the DTA leg, there is further specialisation. The component models are said to be decoupled, which has certain advantages, see Crawford (2014). An advantage of a decoupled model is that it can often be made simpler than a holistic model because, for example, it has fewer parameters. This allows analyses to focus on tractable problems. Drawing together the results of the decoupled models it is then possible to derive conclusions about the system as a whole. Figure 3 simplifies things further, for the purpose of illustration. It shows a hypothetical two-dimensional safety space and a curve within which the core is ‘demonstrably’ safe. In

Figure 3 the parameters could be, for instance, the number of cracked bricks and the number of failed end-face keys. The hypothetical function suggests that increasing either of the parameters ultimately leads to a potentially unsafe state. Of course, in reality there are far more than two parameters.

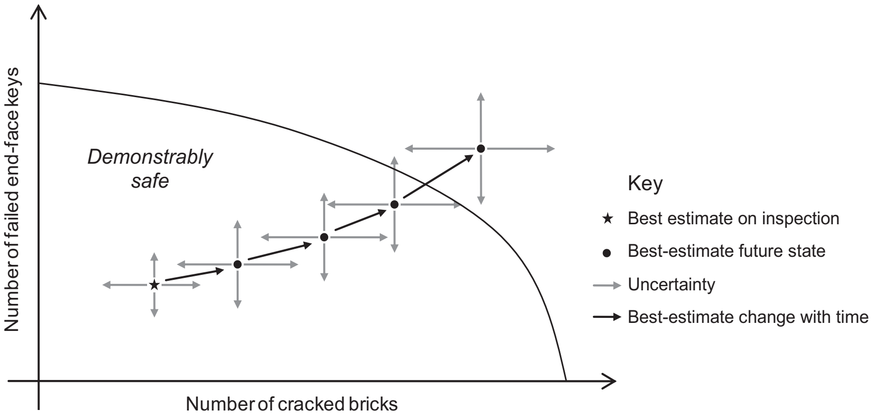

Figure 3. Simplified hypothetical two-dimensional safety space (based on Crawford (2014)) showing the 'demonstrably safe' region as a function of the number of cracked bricks and the number of failed end-face keys. The hypothetical function suggests that increasing either of the parameters ultimately leads to a potentially unsafe state.

Whole-core models typically have many input parameters, all with associated uncertainties. If models are applied deterministically then, potentially, a large number of load cases would need to be run to fully explore the parameter-space; a process which is likely to be prohibitively costly and time-consuming. Therefore, to get the most benefit from a model it is essential to identify parameters which do not significantly affect the results and assign them a fixed sensible value, while focusing attention on varying the parameters that have a greater influence on the results.

One of the key challenges facing EDF Energy and its partners on the graphite core project, whose work contributes to AGR core safety cases, is to gain a better understanding of the location of the boundary between the safe and potentially unsafe regions of the safety space. Another is to determine the current and future locations of the cores within safety space, as illustrated in Figure 3. Therefore, it is possible to successfully extend the lifetime of a reactor by increasing the range of the demonstrably safe region, refining the positions of the observed or predicted core states, or reducing the uncertainties associated with those core states.

Lack-of-cliff-edge studies

An important type of assessment that has formed part of graphite core safety cases is a 'lack-of-cliff-edge' assessment. The value of such an assessment is described by Crawford (2014): "one needs to be aware of the possibility of a cliff edge: a sudden reduction of factors of safety, over a short distance in the safety space or a short time". Core models, both physical and computational, are ideally suited to these types of analyses since it is feasible to use hypothetical parameter values, which extend beyond the time period covered by the safety case. This is illustrated in Figure 4: the uncertainties for these analyses may be greater since the region of application is now more remote from any real measurements. Nevertheless, when combined with analyses that focus on the region that is intended to be demonstrably safe, lack-of-cliff-edge analyses reduce the overall uncertainty about the safety of the core because they are bounding.

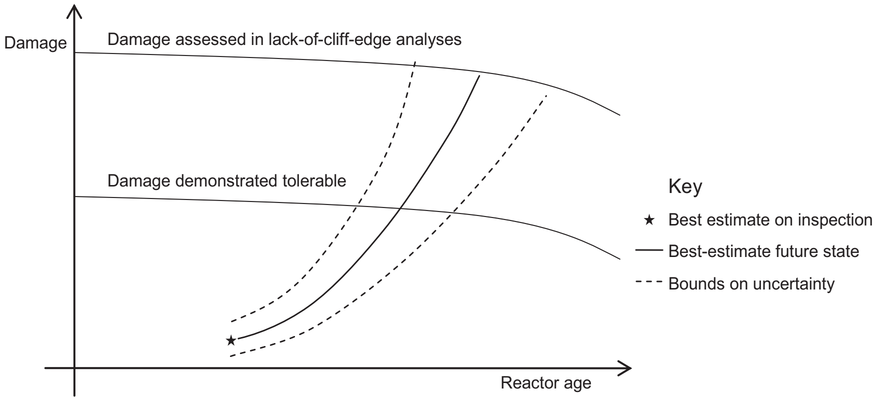

Figure 4. Lack-of-cliff-edge assessments have greater uncertainty because the region of application is more remote from any real measurements. But when combined with analyses that focus on the region which is intended to be demonstrably safe, lack-of-cliff-edge analyses reduce the overall uncertainty about the safety of the core.

FROM CONSERVATIVE TO BEST-ESTIMATE

A report by the International Atomic Energy Agency (2009) discusses the merits of replacing traditional conservative analyses with a more refined “best estimate plus uncertainties” approach. It identifies two main deficiencies of conservative analyses: firstly, a conservative assessment may mask important safety issues; and secondly, it may not show margins to acceptance criteria which, in reality, could be used to obtain greater operational flexibility. This shift in methodology has gradually been adopted as part of the graphite core project over recent years in the work carried out to support AGR core safety cases.

Conservative methods

To enable continued operation of the AGRs in the event of graphite ageing effects such as brick cracking, EDF Energy produced a ‘Revised AGR Core Safety Case’ (RACSC). It was to meet the requirements of the RACSC that whole-core models, such as the AGRIGID code described by Shaw et al. (2007) were instigated. The same paper noted that prior to the development of the RACSC models it was not possible to undertake whole-core modelling since there were neither suitable computational codes nor sufficient computing power to perform such analyses within practical timescales and costs. A similar balance presently affects the graphite core project.

In spite of the fact that AGRIGID was successfully used to support a number of ‘RACSC’ safety cases, the way in which the model was applied, described by McLachlan et al. (2007), was highly conservative. The method of application was conservative because a bounding loading was applied to the model, effectively simulating the core being tilted on its side to envelop realistic distortions. This is not representative of the loads that a core experiences during normal operation, but more onerous. Shaw et al. (2010) confirmed that representative loading of the core gives rise to less distortion than the bounding loading. Other inputs used for the early AGRIGID studies also tended to be rather coarse. For instance, in the absence of any observations or representative predictions of the number of cracked bricks in the cores, analyses were based on bounding assumptions. Current assessments make use of predictions of bore cracking based on fitting statistical models to historical observations, for example Robinson and Maul (2013), and representative predictions of keyway root cracking.

Ultimately, a point will be reached at which a conservative model will predict that there is no margin on the tolerability of the core to the predicted core distortion for a given postulated core damage; while in reality there may still be available margin and hence scope for continued safe AGR operation.

Best-estimate methods

The philosophy behind best-estimate methods is that the computer codes used should not be inherently conservative and the input data should be realistic. Therefore a best-estimate result can be determined and the uncertainties in a result can be better quantified by taking into account the uncertainties in the input data and the models themselves.

The whole-core modelling methodology developed in Shaw et al. (2007) (referred to as a 'stick-spring' model), based on the use of rigid beams and non-linear springs to represent the structure of bricks and key/keyway interactions respectively, was intended to be an equally valid but less computationally expensive alternative to Finite Element (FE) solid modelling. When the simplified model was developed, the conservative assumptions which distinguished the stick-spring method from an ideally realistic model were acceptable because the model was fit for purpose; it was simple to use and able to successfully generate evidence to support safety cases to justify graphite core lifetime extension.

Currently more realistic models are being developed as part of EDF Energy's Plant Lifetime Extension (PLEX) programme. The purpose of the new models is to augment and ultimately supersede those developed for the RACSC safety cases. Modelling methodologies such as SALCOR, and SOLFEC (see Brasier and Rogers (2013)), eliminate the sources of some of the most obvious conservatisms in the stick-spring models.

SALCOR is a whole-core FE solid modelling methodology that has been developed as part of the PLEX programme. It has the potential to overcome many of the conservatisms and limitations associated with stick-spring models. Some of the advantages, listed below, make possible a reduction in the uncertainties associated with stick-spring models, which place ostensible limits on AGR core safety margins:

- More realistic representation of component geometry.
- Deformable rather than rigid bodies.
- Use of realistic material models rather than springs with fixed stiffnesses.
- Realistic simulation of core component damage mechanisms and consequences.
- Realistic simulation of key/keyway disengagement, a constraint on the region of validity for stick-spring models.
- Potential for more refined output.

A comparison of the representation of the core in SALCOR, and its stick-spring predecessor, AGRIGID, is shown in Figure 5. To illustrate the kind of benefit that SALCOR may offer: AGRIGID typically predicts that as the level of brick cracking increases key/keyway disengagement becomes more likely, while other aspects of the damage tolerance assessment for normal operation still show healthy margins. This places potential limits on the range of AGRIGID studies because eventually their validity comes into question. However, SALCOR may be able to show that disengagement does not occur until cracking is more widespread, making it possible to demonstrate tolerability of greater levels of core damage. This is equivalent to raising the level of the 'damage demonstrated tolerable' boundary in Figure 4, thereby helping to justify graphite core lifetime extension.

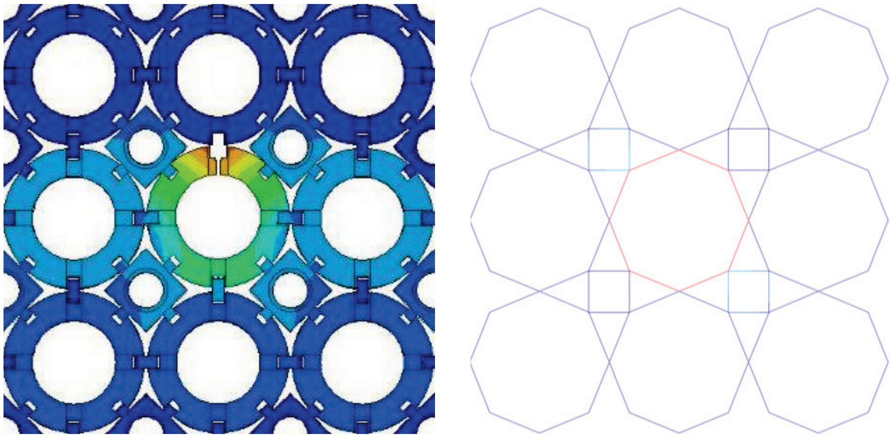

Figure 5. Comparison of a typical representation of core components in SALCOR (left), which uses finite element solid modelling, and AGRIGID (right), which is based on rigid beams and springs.

DETERMINISTIC AND STATISTICAL APPLICATIONS

Another challenge facing whole-core modelling strategies arises from the limitations of using deterministic methods. As suggested previously, a pragmatic approach can be adopted, whereby parameters to which models are insensitive remain fixed, minimising the required number of runs. However, this leads to a reliance on a degree of engineering judgement to identify which combinations of parameter values should be chosen in any set of analyses and there is a risk that onerous combinations of parameters may be overlooked.

There may also be some parameters whose values not only have uncertainty associated with them, but also inherent variability. Variability in AGR graphite cores, as discussed by Maul (2013), is an irreducible characteristic of certain phenomena, such as the precise details of crack distributions. To investigate parameters such as the variability of crack distributions more effectively, it is beneficial to adopt more of a quasi-probabilistic or statistical approach. To do this would involve running far larger numbers of load cases than are typically undertaken in current studies. The inherent challenges with this strategy are:

- To be able to reduce runtimes for individual load cases or run greater numbers in parallel. Computing and software licensing resources typically place constraints on this.
- Streamlining the post-processing of large numbers of load cases to give easily comparable measures of tolerability.

An example of streamlining the measurement of tolerability is the channel functionality assessment code LEWIS described by Crawford (2013), which assesses predicted channel shapes with respect to geometrical constraints and calculates a relatively small number of measures of the margin on channel shape tolerability. Thus, it reduces a potentially complicated set of results to something far more manageable and easily comparable.

VALIDATION AND CROSS-COMPARISON

Where it is practicable, computational models are validated by physical models. However, there are instances in which it is not possible or economically viable to replicate a model with a physical test. In such cases cross-comparison of independent computational models may be carried out in place of, or to supplement, validation.

Validation

The use of a model to provide evidence to support safety cases, and its acceptance by the engineering and scientific community, may sometimes depend on adequate validation by comparison of predicted and measured results. Models can easily become abstruse, so building confidence in them is important.

Nevertheless, validation is not the only purpose for carrying out physical testing. It is suggested by Roscow et al. (2010) that models, "where sufficiently representative, can be used to obtain information of direct applicability to the cores". Indeed, there may be instances in which it is preferable to build a physical test model rather than a computational model. An example is the case of work carried out in 2010, in which physical tests with quarter-scale aluminium fuel bricks were used to demonstrate that key/keyway disengagement was not a concern for bricks with complex crack morphologies. The work was used to support a safety case for Hartlepool and Heysham 1 power stations.

Cross-comparison

In his discussion of uncertainty with relation to statistical models, Maul (2013) makes the point that,

> "Wherever possible multiple models that are consistent with the available data should be considered in order that conceptual model uncertainty can be investigated."

This statement may have been made in the context of statistical modelling but it is equally applicable to other types of core model, including whole-core models. As an example, a cross-comparison of independent computational models has been carried out, in which results from the static whole-core modelling code, AGRIGID, were compared with those from the seismic code, GCORE (the development of which is described by Kralj et al. (2007)). The outcome of the study led to an improved understanding of the sensitivity of both models to various mutual parameters (such as the spring stiffnesses used to simulate different types of contact) and, hence, the significance of those parameters.

CONCLUSIONS

Producing safety cases to justify extending the lifetimes of AGR graphite cores relies on information from a range of sources:

- Empirical data from core inspections and monitoring, to provide accurate measurements of the current states of the cores.
- A theoretical understanding of the behaviour of ageing nuclear graphite at micro and macro scales.
- State-of-the-art models, built on the previous points, to provide predictions of the future states of the cores.

AGR core modelling helps to justify safe reactor operation in the following ways:

- It is directly related to the core FNSR for shutdown, cooling and refuelling.
- It is an independent leg to the core safety case.
- It has a predictive power, once validated.
- It enables hypothetical scenarios, such as faults and seismic hazards, to be explored.
- It allows uncertainty and 'lack of cliff-edge' to be studied.
- It may be used on both a best-estimate and a conservative basis.
- It is amenable to both deterministic and statistical applications.

Thus, AGR lifetime extension may be maximised through a gradual evolution of the state-of-the-art technology, assessment methods and best-estimate input data. A likely constraint is that the tools used to support safety cases at any given time must be economically practicable.

ACKNOWLEDGEMENTS

The author would like to thank EDF Energy for permission to present this paper. The views expressed in this paper are those of the author and do not necessarily represent those of EDF Energy. Special thanks to Richard Crawford, whose work has had a particular influence on this paper.

REFERENCES

Bradford, M.R. (2007). “An Overview of British Energy’s Graphite Core Assessment Methodology”, *Proceedings of the conference on Management of Ageing in Graphite Reactor Cores*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

Brasier, S. and Rogers, S. (2013). “Solid-body Simulation of the Dynamic Response of an Array of Graphite Bricks”, *Proceedings of the conference on Modelling and Measuring Reactor Core Graphite Properties and Performance*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

Crawford, R.J. (2013). “The LEWIS Method for Channel Functionality Assessment”, *Proceedings of the conference on Modelling and Measuring Reactor Core Graphite Properties and Performance*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

Crawford, R. (2014). “Lessons from the graphite core project”, *Nuclear Future*, Volume 10, Issue 6.

International Atomic Energy Agency (2009). “Deterministic Safety Analysis for Nuclear Power Plants”, *Specific Safety Guide SSG-2*.

Kralj, B., Humphreys, S.J. and Duncan, B.G.J. (2007). “Seismic Modelling of an AGR Nuclear Reactor Core”, *Proceedings of the conference on Management of Ageing in Graphite Reactor Cores*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

Maul, P.R. (2013). “Statistical Modelling of the Ageing of Graphite Cores”, *Proceedings of the conference on Modelling and Measuring Reactor Core Graphite Properties and Performance*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

McLachlan, N., Shaw, D.J., and Salih, M.H.S. (2007). “Application of Whole Core Modelling Methodology to Life Extension of AGR Reactor Graphite Cores”, *Proceedings of the conference on Management of Ageing in Graphite Reactor Cores*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

Robinson, P.C. and Maul, P.R. (2013). “Statistical Modelling of Graphite Properties”, *Proceedings of the conference on Modelling and Measuring Reactor Core Graphite Properties and Performance*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

Roscow, A.R., Skelton, J. and McLachlan, N. (2010). “Continuing Development of the Quarter Scale Core Model to Support the Safe Performance of AGR Cores”, *Proceedings of the conference on Securing the Safe Performance of Graphite Reactor Cores*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

Shaw, D.J., Salih, M.H.S. and McLachlan, N. (2007). "Development and Validation of Whole Core Modelling Methodology for AGR Graphite Cores", *Proceedings of the conference on Management of Ageing in Graphite Reactor Cores*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.

Shaw, D.J., McLachlan, N. and Powrie, C. (2010). "Application of Whole Core Modelling Methodology with Representative Core Loading to Life Extension of AGR Graphite Cores", *Proceedings of the conference on Securing the Safe Performance of Graphite Reactor Cores*. G.B. Neighbour (Ed.), The Royal Society of Chemistry, Cambridge, UK.