AWE is the trading name of AWE plc
Registered office: Aldermaston Reading Berkshire RG7 4PR
Registered number 3664571

The Science & Technology Journal of AWE • Issue 22 • October 2011

Discovery 22 • The Science Engineering and Technology Journal of AWE

Science Engineering Technology at AWE

Discovery22

October 2011

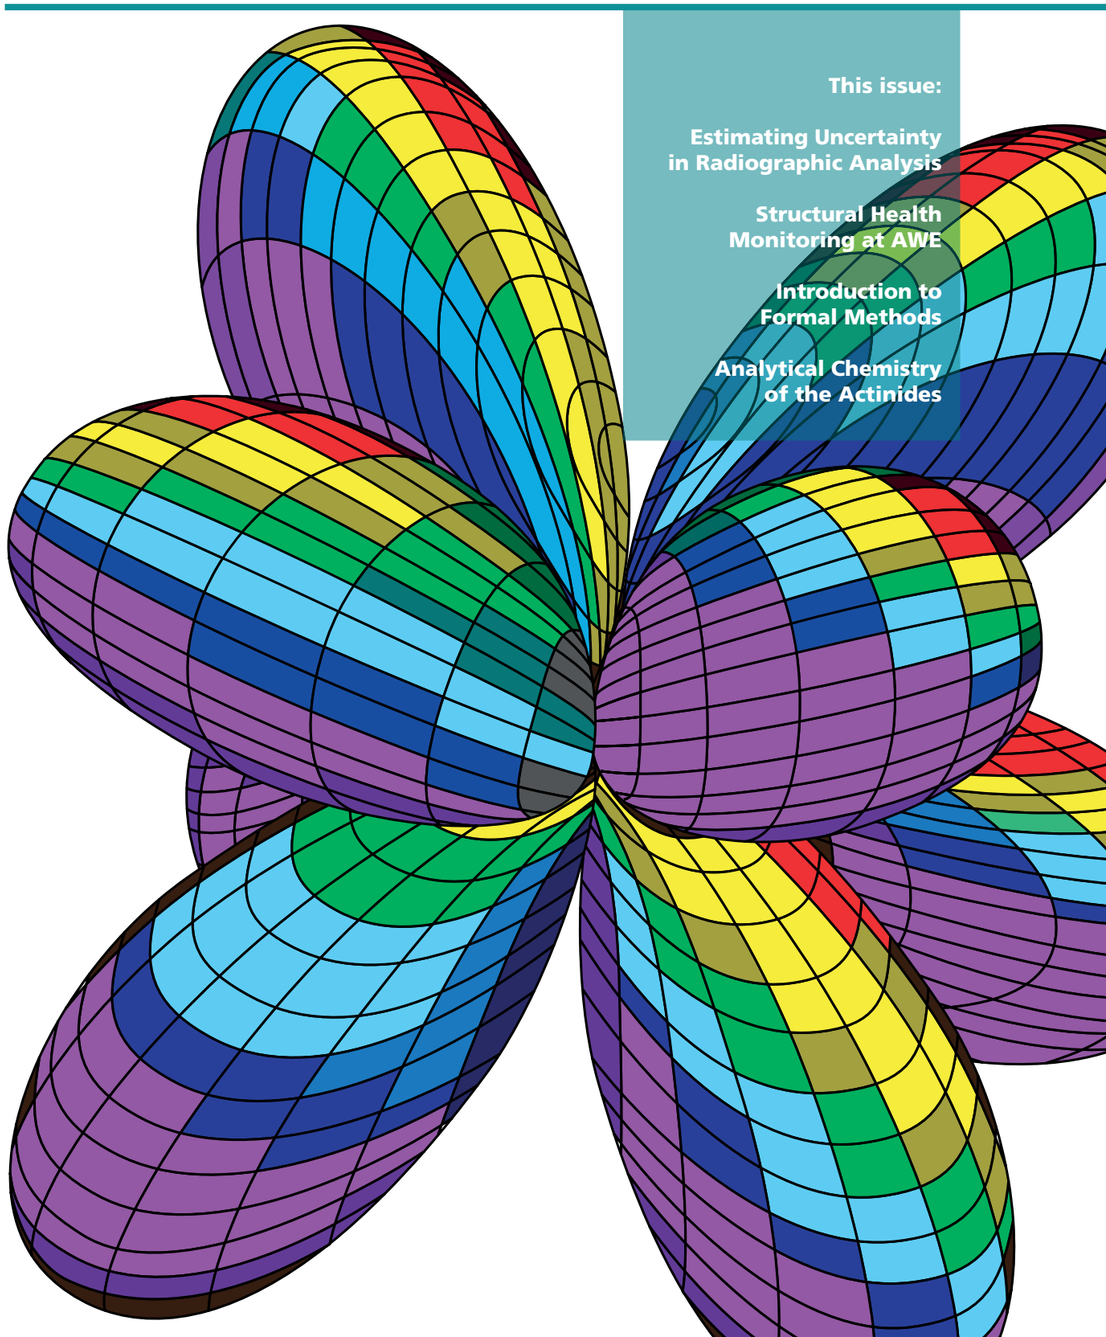

This issue:

Discovery 22

Contents

cover image
Electron orbitals

MIX
Paper from responsible sources
FSC® C020704
www.fsc.org

WORLD LAND TRUST™
www.carbonbalancedpaper.com
CBP0009310810012710

22

I am delighted to be invited to write the foreword in this publication, the 22nd issue of Discovery, which portrays a fascinating insight into the scientific and technological research undertaken at AWE. I consider this work to be vital in ensuring that the UK remains at the forefront of nuclear deterrence and national security.

In my dual roles, as Professor of Computer Science in the School of Electronics and Computer Science at the University of Southampton and as a Chief Scientific Adviser to the UK Government, I am consulted on a number of issues in the areas of complex adaptive systems and intelligent systems. It is clear that in an increasingly complex and sophisticated computer age, there is an inextricable link between security and the scientific challenges that AWE faces. The synergies between computer intelligence and AWE's research programmes are obvious.

At the University of Southampton, I am leading major research projects (ALADDIN and ORCHID) that aim to develop true partnerships between people, smart infrastructure and intelligent computing components. For me, this symbiotic relationship is vital in many aspects of science and in its translation to real-world applicability.

To this end, this issue of Discovery continues to focus on the range of scientific endeavours at AWE. It features an article on the actinide series of elements; a Bayesian approach to estimating uncertainty in the analysis of hydrodynamics experiments; an introduction to formal methods and safety critical systems and a piece on structural health monitoring which is a relatively new field in engineering.

It is also good to see the Group of Experts in Mitigation Systems colloquium featured in the Outreach section. As guest speaker at this forum, I spoke about my appreciation of AWE as an investor in science and technology in today's difficult economic environment. I was particularly impressed in the quality and breadth of work presented in this forum, and also in the way that academia, industry and government are working so closely together. This partnership is vital to further progress in this area.

Professor Nick Jennings
FREng

In closing, I hope you enjoy reading this issue and discovering more about AWE's contribution to science, engineering and technology.

Estimating Uncertainty in Radiographic Analysis

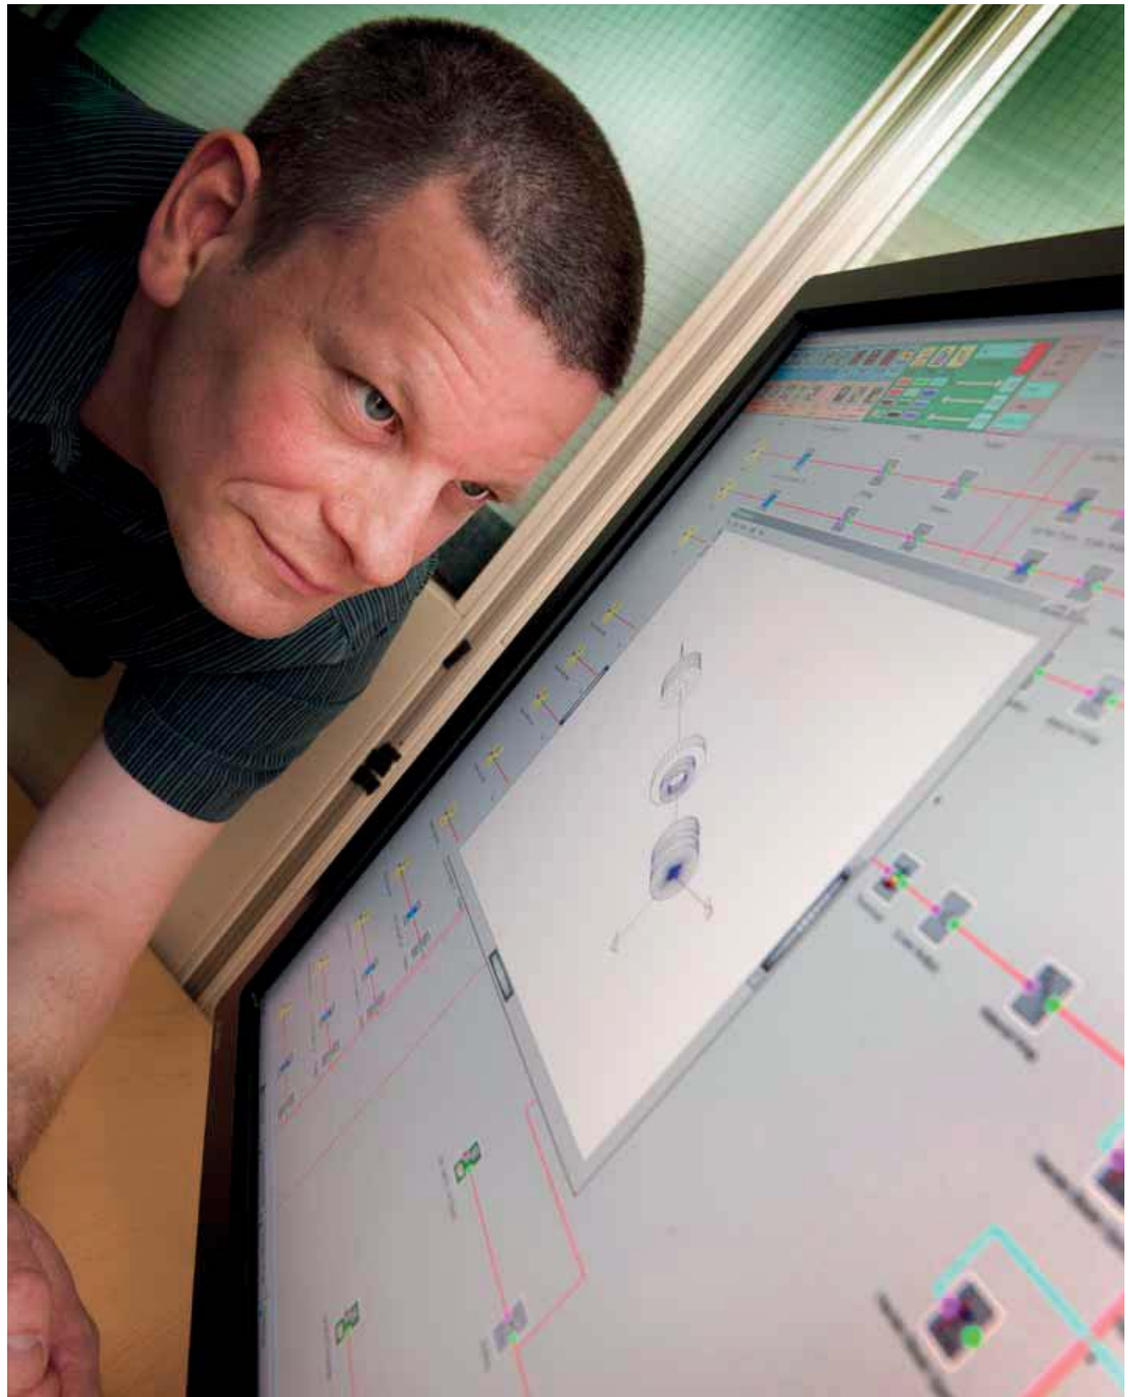

Integrated hydrodynamic experiments use inert materials to reproduce the explosively driven behaviour of nuclear warheads as closely as practicable. Analysis of the X-ray transmission radiographs from these experiments provides unique data used to validate calculational models.

Given the high cost of these experiments, it is not viable to carry out multiple experiments to assess the uncertainty of the results. A calculational methodology has therefore been introduced, Randomized Maximum Likelihood (RML), to estimate uncertainty by performing analyses on a large number of virtual experiments.

This article will present an overview of these experiments, the RML technique and its validation.

In a Core Punch (CP) experiment a very short pulse of high energy X-rays is used to look inside a representative nuclear warhead during its explosive compression. This makes it possible to observe how shock waves and interfaces develop, and to measure material densities. Such experiments have an important role in providing data to assess nuclear weapon safety and performance.

It is necessary to have specialised facilities designed to safely contain explosives as well as radiographic diagnostics capable of performing a few centimetres away from explosive devices. Figure 1 shows a specialist facility before and after an experiment.

In order to freeze motions of several kms^{-1}, the X-rays must be delivered in a single pulse of less than 100 ns. The most powerful of the flash X-ray machines at AWE, Mogul E^{+}, is 25 m long and weighs 280 tonnes. Such machines are required as the object's attenuation means that less than 1 in 100,000 of the photons leaving the source will be measured at the detector.

The pulse of X-rays is achieved by charging a metal cathode to 10 million volts, which accelerates 40,000 amps of electrons onto a small region (approximately 5 mm) of a target anode. The deceleration of the electrons in the target liberates Bremsstrahlung X-rays. This generated power of 400 GW is considerably greater than the entire UK power generation capacity of 56 GW, albeit for only 70 ns. The two Mogul machines can take two images separated by a few hundred billionths of a second. Figure 2 shows the two Mogul machines surrounding a test facility.

FIGURE 1

A chamber before and then after firing.

FIGURE 2

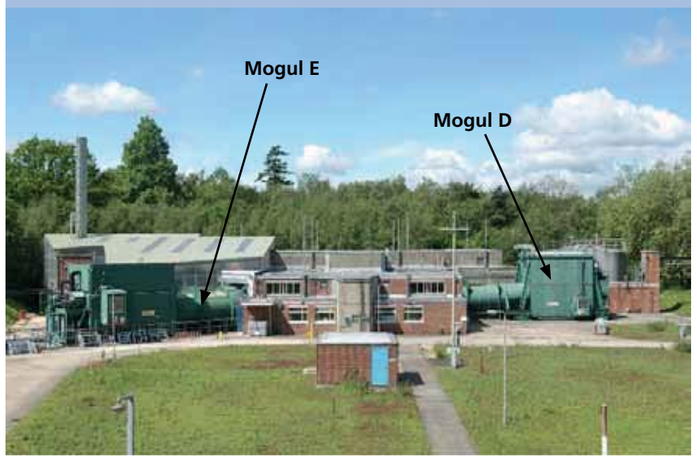

Mogul D and E X-ray sources.

Sources of Uncertainty in Radiographic Analysis

The intensity distribution (I) on a radiograph is related to the mass of attenuating material in the X-ray beam by

$$I(x,y) = I_0 e^{-\sum \mu m(x,y)} + S$$

It is possible to use the known attenuation properties of the materials (μ) to infer the mass distribution of the object (m), given the initial intensity (I_0). Using the assumption that the object is symmetric, it is possible to determine a density distribution from this mass. Scatter (S) arises from the interaction of X-rays with matter and needs to be subtracted as these photons do not provide useful information. The sum is carried out for all materials in the beam over all energies in the X-ray spectrum.

Given the challenging nature of CP radiography, there are limitations in the quality of the radiographs obtained. For the low photon flux that occurs with a highly attenuating object, quantum statistical noise results from the random statistical variation in the photons measured by the detector. Another significant issue is blur, which results from the finite size of the X-ray source spot and the energy spread within the X-ray detector.

All of the experimentally measured parameters will have uncertainties. Scatter is generally considered to introduce the greatest source of uncertainty. Deviation from the assumption of symmetry will also introduce an error into the derived density.

In the analysis, a model of the radiographic processes is created to allow a prediction to be made. The uncertain parameters in the model are then optimised so that the prediction agrees with the experimental radiograph. There are effectively a wide range of possible solutions that fit the radiographic data within the experimental noise.

An analysis that aimed to exactly fit the radiograph would yield density results in which the noise on the image could be translated to an unphysical fluctuation in the density solution.

A Bayesian approach is used to reject those solutions that are unphysical, based on our experience and understanding of the physics involved. Examples of prior knowledge typically used are mass conservation, non-negative density and density smoothness. Box 1 provides further information on the Bayesian approach.

The software used to carry out this analysis, called the Bayes' Inference Engine[2] (BIE), was originally written at Los Alamos National Laboratory (LANL)

> "A Bayesian approach is used to reject those solutions that are unphysical, based on our experience and understanding of the physics involved."

FIGURE 3

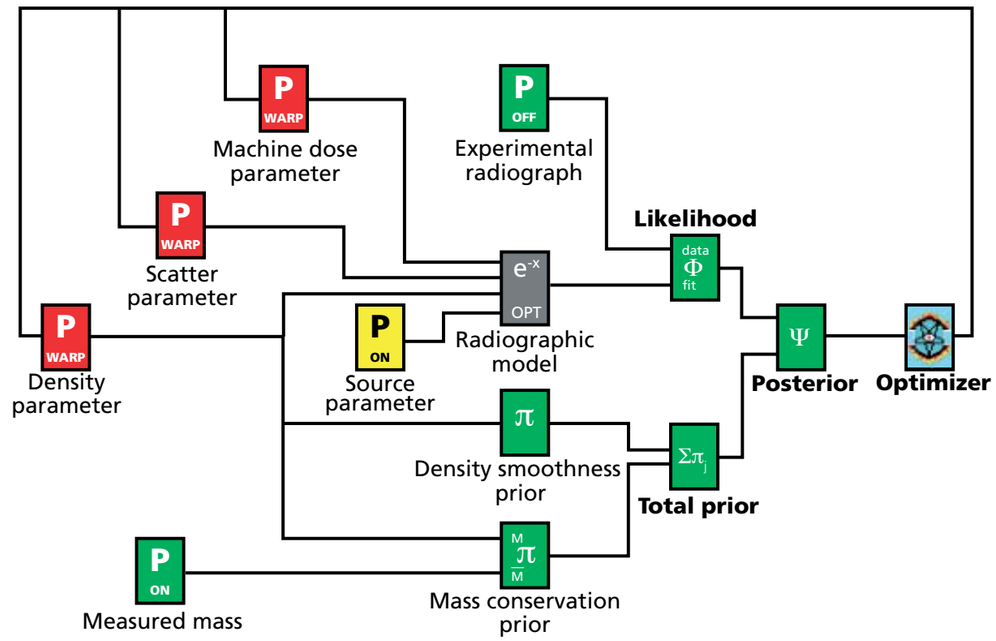

Simplified BIE network.

and since 1996 has been jointly developed between LANL and AWE. Using the BIE an analysis is carried out by representing the problem as a collection of pre-written modules in a data flow network, this is shown in Figure 3.

RML Uncertainty Estimation

Uncertainty is a measure of the confidence in a calculated or measured value. A frequency approach to uncertainty estimation would take a large number of experimental measurements and determine uncertainty from the spread of results obtained. This approach is not applicable in the case of CP experiments which are costly and liable to be performed only once.

A common technique for assessing the uncertainty in a Bayesian manner is to explore the parameter space using the Markov Chain Monte Carlo[3] (MCMC) technique. MCMC is not practical in this case due to the calculational time required resulting from the large number of parameters. The RML technique is an approximation to MCMC that was introduced to AWE by Professor M Christie of Heriot Watt University. The implementation of RML in the BIE was undertaken as part of an AWE Summer Fellowship.

From the outset it was determined that a methodology should meet the following criteria:

- Follow the Bayesian analysis techniques used in actual analysis to explore the range of possible solutions
- Consider all of the important sources of uncertainty
- Be validated both though calculation and against experimental data
- Run in a timescale that is not so long so as to prohibit its use

RML was developed by Oliver, He and Reynolds and has been

FIGURE 4

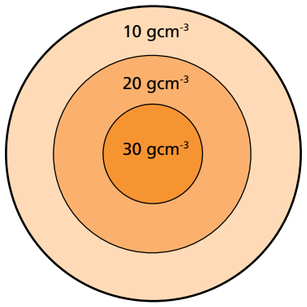

Calculation test object.

FIGURE 5

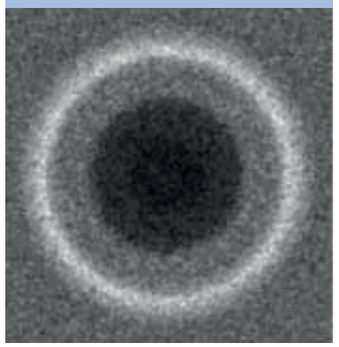

Calculational test radiograph.

successfully applied to problems in the oil industry.[4] RML aims to explore the range of possible solutions by carrying out multiple minimisation optimisations on randomly perturbed datasets within the generated expected experimental uncertainties. The model parameters are similarly given random starting values to avoid the optimisation falling into a local minimum.

RML Validation

RML has been tested on a number of synthetic cases to give confidence in its use for the analysis of CP radiographs. A simplified test object consisting of a dense central ball surrounded by less dense spherical shells, shown in Figure 4, was used to create a synthetic radiograph, shown in Figure 5. A high level of noise is included to test RML under extreme circumstances. The parameter uncertainties to be determined are the three density levels, the radiographic machine X-ray exposure, a multiplicative factor, and a uniform scatter model, an additive factor.

The data were then analysed 100 times; each time different random noise was added to the radiograph and the parameters were started at random values within their estimated uncertainties. The total mass of the test object was used as a prior, see Box 1 for further information.

Figure 6 summarises the results showing the true parameter value overlaid with the estimated parameter mean and its uncertainty bounds. One standard deviation was chosen which equates to approximately 68 % confidence. The true value of each parameter falls either within or close to the uncertainty bounds for all the RML calculations, demonstrating that the RML technique works well for this problem.

The RML technique has also been tested by applying it to a real radiograph of a test object, shown in Figure 7. The radiograph was obtained using AWE's Miniature

FIGURE 6

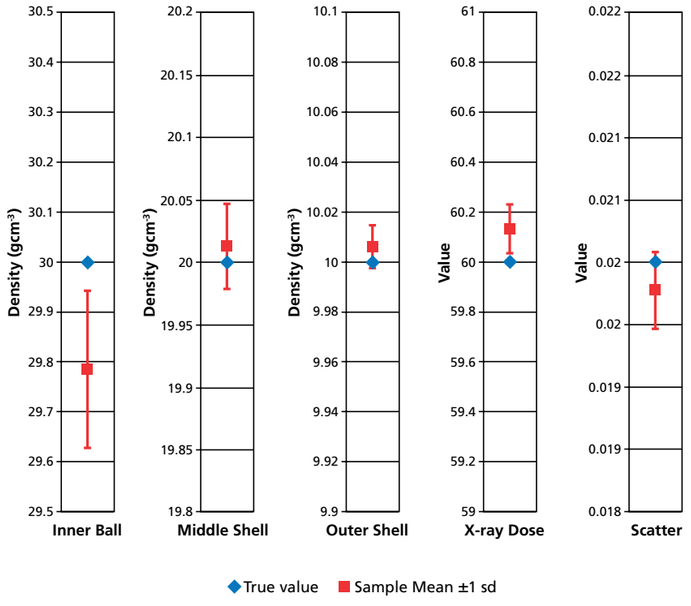

Parameter uncertainties using the RML technique.

FIGURE 7

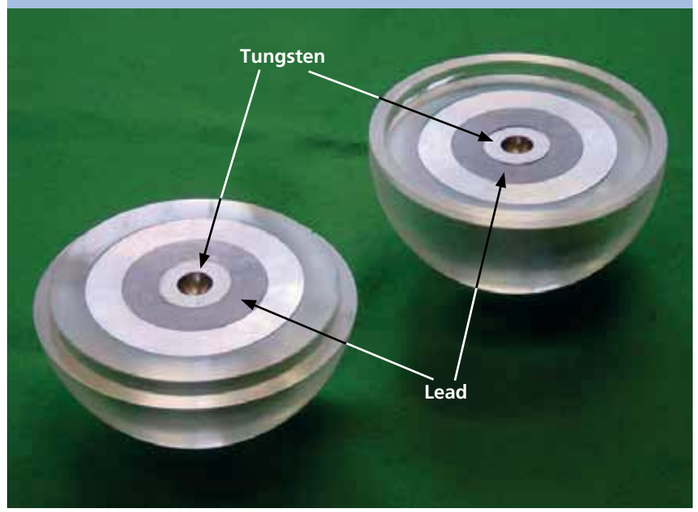

Static model of RML test sphere.

Linear Accelerator (MINAC) and is shown in Figure 8.

The priors used were density smoothness, mass conservation and non-negative density. The parameters optimised during each RML sampling were the density grid, the machine output exposure and the scatter. An estimate of the noise on the radiograph and the uncertainties of all experimentally measured parameters were required. Results are shown in Figure 9 with and without priors.

The radial averaged density profiles of the sample mean were compared with the test object and are shown in Figure 10. The true value is seen to generally fall within the bound of the mean standard deviation, again indicating that the RML technique worked successfully.

When RML was applied to the analysis of a CP radiograph the uncertainties in the result were estimated to be less than 0.25 mm on interface positions and 5 % on density. This knowledge can be used to help refine the calculational model.

FIGURE 8

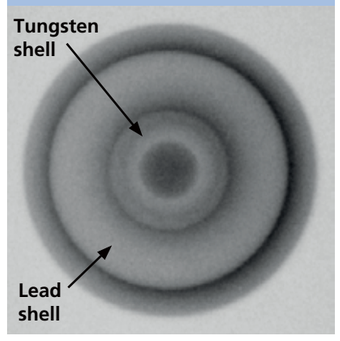

MINAC radiograph of RML test sphere.

FIGURE 9

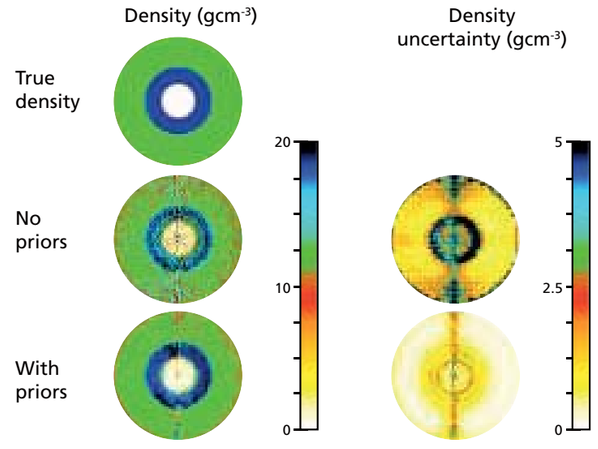

Sample Mean Density and uncertainty of the test sphere.

FIGURE 10

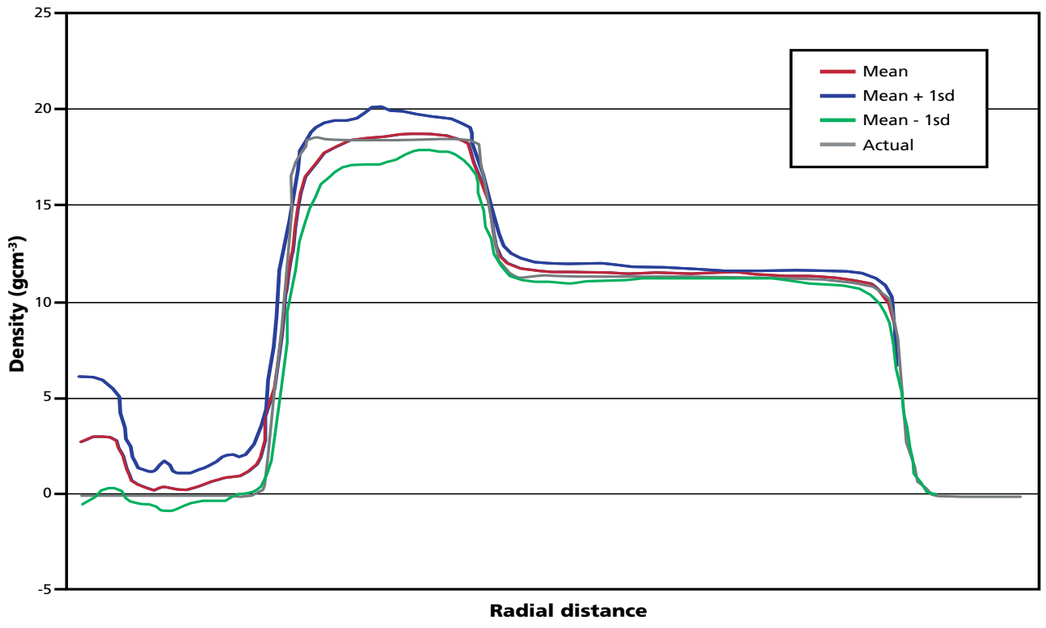

Mean radial profile.

BOX 1

Bayesian Analysis

Bayesian Inference[6] is a statistical approach for inferring the probability of an outcome by considering all the available evidence, both from the experimental data and from previous beliefs derived from other sources. It is named after the relatively obscure 18th century clergyman Reverend Thomas Bayes (1702-1761) and is now an important tool in scientific analysis. Bayes' Law[7] can be expressed mathematically as:

$$p(model \mid data) \propto p(data \mid model) \times p(model)$$

or

$$\text{Posterior} \propto \text{Likelihood} \times \text{Prior}$$

The posterior is the conditional probability of the model being correct given the data. The likelihood is a measure of how well the data predicted by the model fit the actual data. The prior is any applicable information that is not expressed in the data; this may come from sources such as previous experiments, experience, or the laws of physics. The objective of Bayesian Inference is to obtain the most probable solution by achieving a balance between the best fit to the data and the best fit to the prior knowledge.

The Future

It is important to obtain data with the lowest possible uncertainty as such, AWE is continually improving its diagnostic and analysis tools and its understanding of radiography. Experiments of known test objects are used to further refine radiographic modelling.

Work is underway to develop the next generation of experimental facilities, which will ensure AWE's world leading status is maintained in this important field of science.

Further collaborative work is underway with Heriot-Watt University to investigate the possibility of using other uncertainty estimation techniques such as the computationally efficient alternative to MCMC, the Hybrid Monte Carlo[5] (HMC) technique.

Acknowledgments

The BIE analysis tool has been developed with colleagues both in the Hydrodynamics Department at AWE and at LANL. The uncertainty estimation work contained in this report was carried out with T F Pang and utilised data from experiments carried out by S Whitwell. In addition Professor M Christie of Heriot-Watt University is thanked for his help in implementing the RML uncertainty estimation technique in the BIE.

References

1. M. Sinclair, Current Radiographic Pulsed Power Machines at AWE, Pulsed Power Conference (2005)
2. K. M. Hanson and G. S. Cunningham, The Bayes Inference Engine, Maximum Entropy and Bayesian Methods, Kluwer Academic, Dordrecht (1996)
3. N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller and E. Teller, Equations of state calculations by fast computing machine, J. Chem. Phys 21 1087 (1953)
4. D. S. Oliver, N. He and A. C. Reynolds, Conditioning Permeability Fields to Pressure Data, Proceedings of the 5th European Conference on Mathematics of Oil Recovery, (1996)
5. K. M. Hanson, Markov Chain Monte Carlo posterior sampling with the Hamiltonian method, Medical Imaging: Image Processing (2001)
6. D. A. Sivia, Data Analysis: A Bayesian Tutorial, Oxford Science Publications (1996)
7. T. Bayes, Essay Towards Solving a Problem in the Doctrine of Chances, Philo. Trans. Of the Royal Society of London 53:370-418 (1774)

AUTHOR PROFILE

Steve Calverley

Steve graduated from Loughborough University with a BSc (1st class honours) degree in Physics. In 2003 Steve joined the Hydrodynamics Department at AWE initially working on the simulation of hydrodynamics experiments and subsequently on the radiography and analysis of these experiments.

Steve has played a key role in developing AWE's radiography and image analysis capability and is recognised as a technical authority on radiography.

Structural Health

Monitoring at AWE

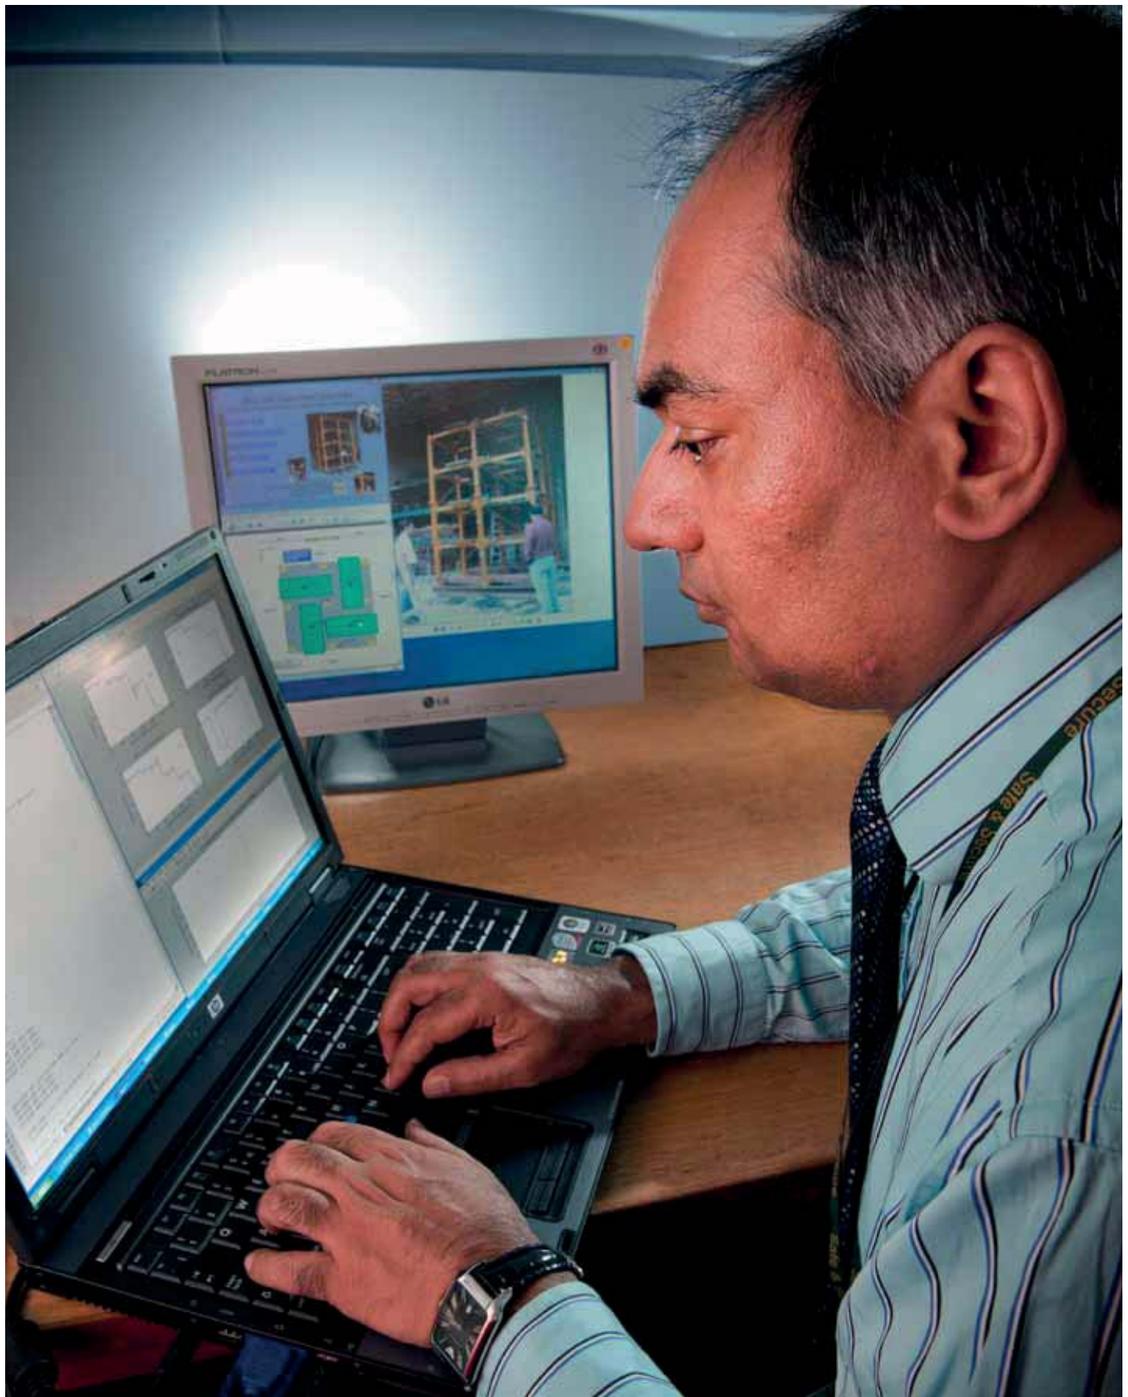

No structure is perfect; they all contain flaws or defects in varying degrees. These defects evolve over time into damage either due to operational conditions, the ageing process or extreme events such as collisions or earthquakes. The dilemma faced by engineering professionals is to then assess the integrity of these structures in order to make decisions regarding their suitability to be left in service or withdrawn for repair.

Structural Health Monitoring (SHM) is a relatively new field of engineering which makes use of sensors, actuators and mathematical software algorithms to detect the presence of damage, locate its position, assess the severity and predict the future condition of aerospace, civil and mechanical structures. The aim is to assist engineers in making decisions about life cycle management.

This article will present an overview of SHM and show the results of research conducted at AWE on a benchmark problem.

SHM evolved as a discipline over the last 30 years. The main aim for this field is to detect damage in aerospace, civil and mechanical structures.

All systems contain defects which can develop into damage and if not attended to, ultimately turn into faults. To cope with issues concerning defect, damage and fault engineers have traditionally relied on technologies like Statistical Process Control (SPC) which monitor industrial processes to alert users about the onset of faults.

Non Destructive Evaluation (NDE) encompasses techniques to assess the extent of damage in materials, components and systems once it has been located. NDE is a technology that deals with damage characterization and severity check.

Condition Monitoring (CM), which has existed a few decades before SHM, is another technology which is primarily concerned with detecting and characterizing damage in rotating and reciprocating machinery. It has had many industrial successes notably in the helicopter and wind turbine industries where it can detect the onset of damage in bearings and gears in the drive train.

SHM seeks to extend the remit of the above-mentioned techniques by offering the possibility of detecting and monitoring damage in structures. SHM shares many common techniques with CM but the two fields have evolved separately. As technology evolved, it became possible to monitor the state of health of structures using sensors, actuators and software algorithms based on pattern recognition techniques to extract damage information upon which better decisions can be made regarding life cycle management.

Properties of those structures that are sensitive to damage need to be identified and algorithms that can differentiate the undamaged and the damaged states can then be developed in order to notify the user as to the condition of the structure.

There is no universal solution to the SHM problem but rather a set of techniques is available for engineers to mould and extend in order to solve the problem at hand. One such technique is called the Statistical Pattern Recognition (SPR) paradigm and it deals with all aspects of an SHM system from design through to delivery into service. The SPR paradigm enunciates a four step process:

1. Operational Evaluation
2. Data acquisition, normalization and cleansing
3. Feature selection and information condensation
4. Statistical model development for feature discrimination

Operational evaluation is concerned with defining which damage needs to be monitored. When the damage of concern has been properly identified and defined, the task of the SHM system is to detect its onset and notify the user. Diagnosis of damage can be defined in a hierarchy of four levels; Detection, Localisation, Assessment and Prediction.[1]

Data acquisition is concerned with:

- Selecting the excitation method
- Determining sensor types, number and locations

- Determining the bandwidth of signals, sampling rates and anti-aliasing filter bandwidths

Data acquired will change due to damage, operational variations, environmental variations and noise. Noise can be removed by various types of filters and the process can be termed data cleansing. In order to detect changes due to damage alone, the effects of operational and environmental variations on the acquired data need to be removed to the extent possible by data normalization.

The main reason why SHM technology has been slow in transitioning from the laboratory to real applications is because the damage sensitive features for structures are also sensitive to changes in the operational and environmental conditions.[2] Unless the effects of damage can be isolated from the effects emanating from other sources, any detection system will suffer from unacceptably high false-positives.

Sensors cannot themselves directly measure damage, they can only measure the system responses to operational and environmental stimuli.

FIGURE 1

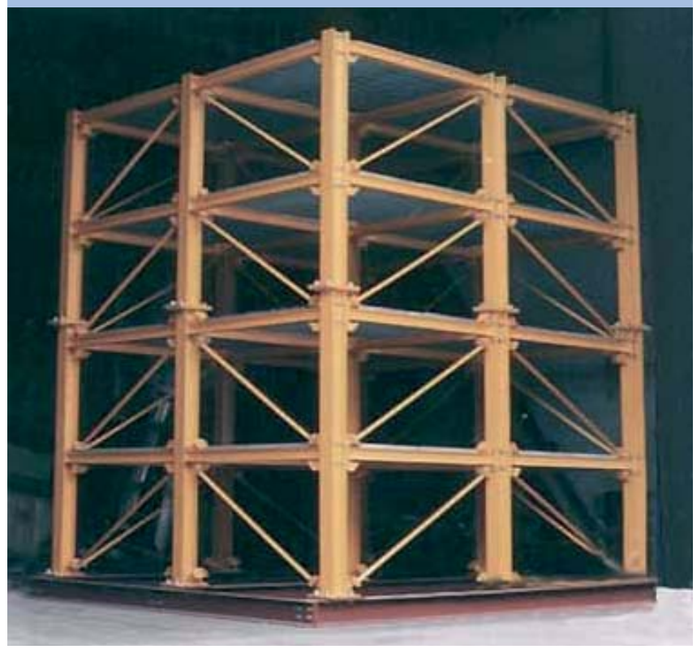

The four storey steel frame benchmark structure.

The data that are sensed are a function of the damage that may be present. Feature selection is used to identify the function that relates the data to the damage. This function cannot generally be determined based on basic physics but must instead be learned from the data using techniques from machine learning.

Three common methods of extracting damage sensitive features from data emanating from the real system or a Finite Element model of it are:[1]

1. Based on first hand observation of the degrading system, quantities can be sought that correlates well with the observed damage
2. Engineered flaws can be introduced into a system under study and features sensitive to damage isolated
3. Finite Element models can be used where damage can be introduced via computer simulations

"All systems contain defects which can develop into damage and if not attended to, ultimately turn into faults."

FIGURE 2

Case	Configuration
Config1	Fully braced configuration
Config2	All east side braces removed
Config3	Removed braces on all floors in one bay on southeast corner
Config4	Removed braces on 1st and 4th floors in one bay on southeast corner
Config5	Removed braces on 1st floor in one bay on southeast corner
Config6	Removed braces on all floors one east face, and 2nd floor braces on north face
Config7	All braces removed on all faces
Config8	Configuration 7 + loosened bolts on all floors at both ends of beams on east face, north side
Config9	Configuration 7 + loosened bolts on floors 1 and 2 at both ends of beams on east face, north side

Table of different configurations of four story benchmark structure.

Having acquired data from sensors that have been cleansed and normalized, features that contain damage information are extracted allowing engineers or computers to make decisions as to whether a newly received feature belongs to the undamaged or damaged system.

This information has to be reliable enough in that the SHM system should not give an alarm when the system being monitored is undamaged (false-positive) or neglect to give an alarm when the system is actually damaged (false-negative).

Statistical models on the other hand build a relationship between data obtained from the real or physical model. Assumptions are made about the structure of the model and then mathematical tools are used to estimate its parameters. If the assumed model structure is incorrect then even if the parameters can be fitted exactly, the model would not be a true representation of the real system and any conclusion drawn from such models would be suspect.

FIGURE 3

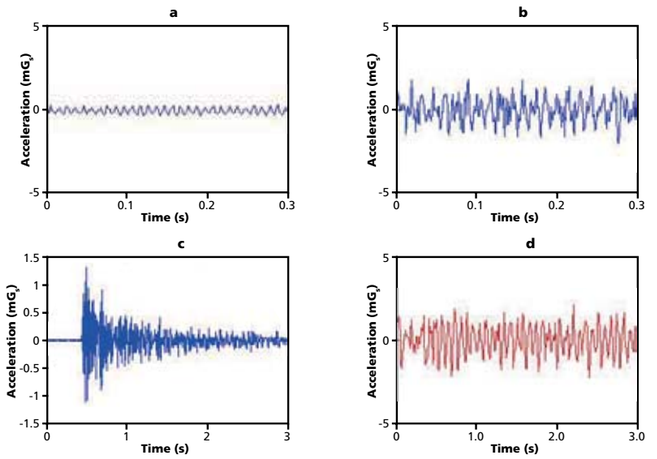

Sensor acceleration outputs for the test structure.

There are few examples where the SHM technology has transitioned from the research laboratory to practice. SHM is a significantly complex and diverse problem that may not be solved in the immediate future but progress will come in small increments over long periods of time.

SHM Research at AWE

The American Society of Civil Engineers (ASCE) have built a scale model of a four storey building, shown in Figure 1, and made available the acceleration data for the undamaged structure and for eight different configurations of damage, as shown in Figure 2. This is used to

FIGURE 4

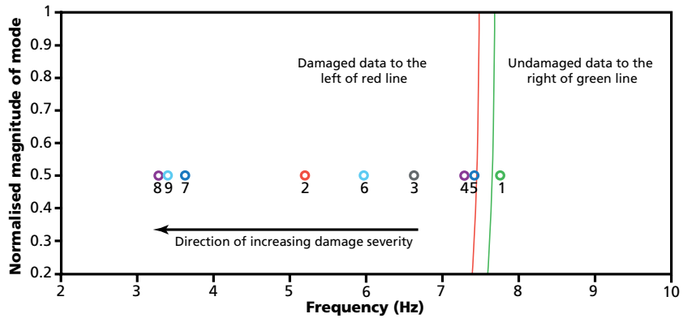

Damage classification from ambient data collected for test configurations 1 to 9.

assess techniques for damage detection, damage severity assessment, damage localisation and identification of damage type. Data representing three categories of vibration as input excitation (ambient, shaker and hammer) are also available. The structure is instrumented with sixteen accelerometers.

Figure 3 shows typical plots of sensor outputs. Figure 3a shows the acceleration data for the ambient input. Figure 3b shows the acceleration for the shaker input. Figure 3c shows the acceleration for the hammer input. These three plots are for the undamaged structure. Figure 3d shows the acceleration for the damaged structure when the input is the shaker.

Two types of damage are introduced to the structure, one is the removal of the braces that are shown as rod connections on the diagonal of each bay and the second is the loosening of the bolts on the horizontal beam. The AWE method for SHM successfully detected the damage configurations for the three input excitations.

A commonly used approach to damage detection is to use shifts in the resonant frequencies; these approaches do not work well in practice.[3-5] The AWE method uses principal component analysis to construct a new data set using a linear combination of some of the sensors.

The data set is decomposed using wavelet multi-resolution analysis which highlights the frequency window where the resonant frequency of the undamaged structure is known to reside. The resonant frequencies are extracted and classified using the Support Vector Machine (SVM) algorithm. It was found that the lower

FIGURE 5

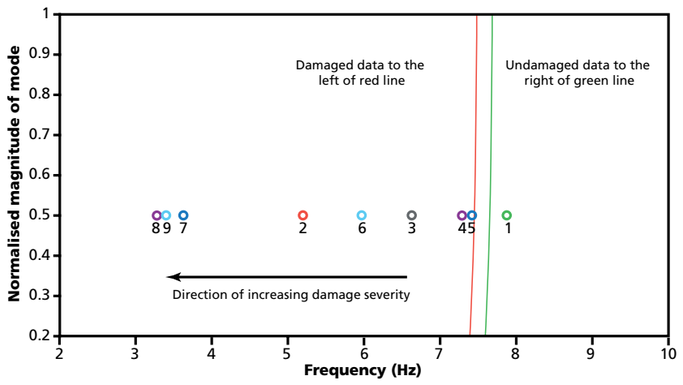

Damage classification from shaker data collected for test configurations 1 to 9.

"Sensors cannot themselves directly measure damage, they can only measure the system responses to operational and environmental stimuli."

"A distinctive feature of the AWE methodology is that it does not rely on the availability of a model of the structure"

FIGURE 6

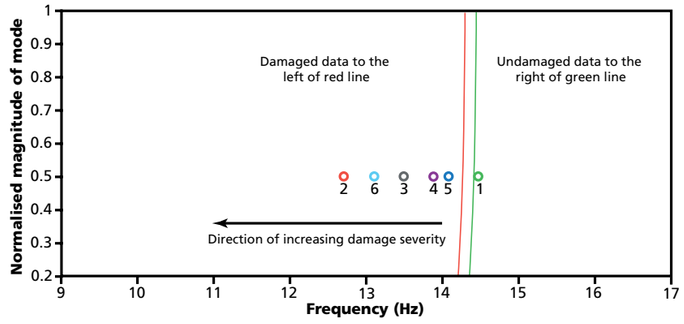

Damage classification from hammer data collected for test configurations 1 to 6.

resonant frequencies behave in a different manner to the higher ones when the damage types are different and that the amount of frequency shift is correlated with the level of damage and can be used to assess damage severity.

A distinctive feature of the AWE methodology is that it does not rely on the availability of a model of the structure, the prior availability of data for the structure in a damaged state or availability of the input excitation. On the example structure the AWE method showed that at the highest input energies, the hammer, the existence of any form of damage can be detected with almost any combination of sensors. At the lowest input energy, ambient, the choice of sensor has a significant impact on the ability to detect damage. Figures 4-7 show plots of the resonant frequencies for all the test cases with the numbering of the resonant frequencies, shown as circles, referring to the corresponding test case number as defined in Figure 2.

FIGURE 7

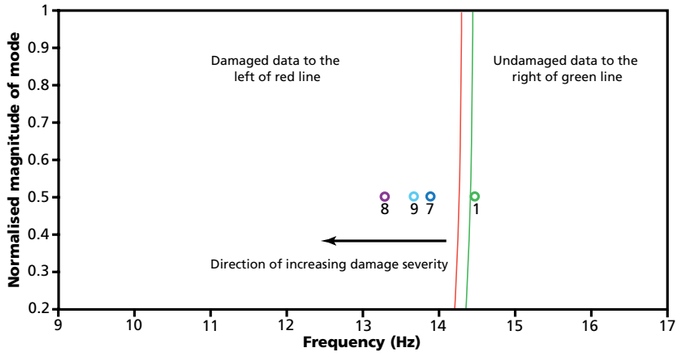

Damage classification from hammer data collected for test configurations 7 to 9.

The magnitude of the resonant frequencies are normalised to 0.5 to place them in the middle of the graph. The green and red lines, separating the undamaged and damaged cases, were generated artificially using the known location of the resonant frequency for the undamaged configuration and the knowledge that damage causes a reduction in stiffness which in turn causes a decrease in the resonant frequency.

The red and green lines were generated using the Support Vector Machine (SVM) technique which can be used for

FIGURE 8

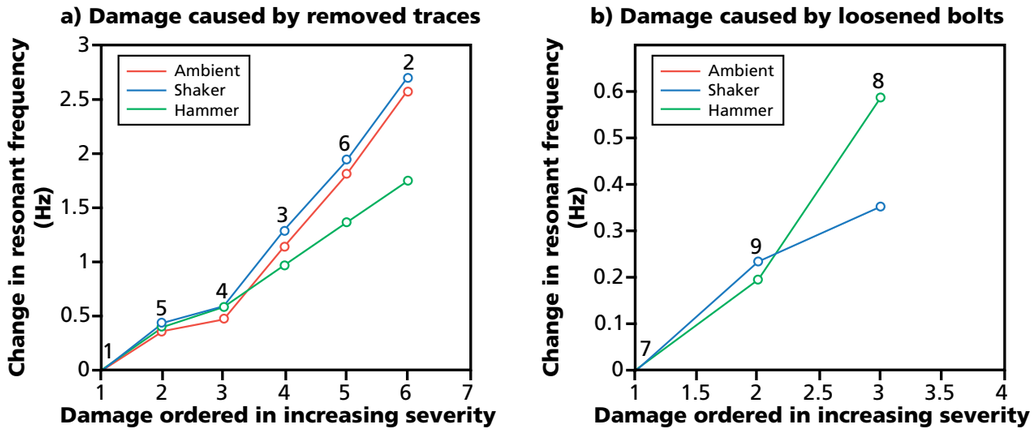

Change in resonant frequency as a measure of damage severity.

classification. This was done in order to put AWE's methodology into a SPR paradigm framework and allow sensitivity analysis and the possibility for generalisation in future work.

The SVM classifier was arbitrarily placed to the left of the resonant frequency of the undamaged case knowing that damage will cause the resonant frequency to move to the left.

Figures 6 and 7 show the test cases for the hammer input. Two plots are used for ease of visualisation because the higher resonant frequency behaves differently to the lower one for the different damage types. The choice to use the higher resonant frequency for the hammer input was made based on the energy distribution of the spectrum.

Figure 8a shows the change in resonant frequencies between damaged configurations 1 to 6 and uses the undamaged configuration 1 as a reference. It was observed that as the damage severity increases, the change in resonant frequencies becomes more pronounced and can be used to assess damage severity.

Figure 8b shows the plots for damage configurations 7 to 9 with damage configuration 7 used as a reference. The ambient and shaker plots overlap in this figure. Increasing the damage causes an increase in the change of the resonant frequencies allowing damage severity assessment of the loosening of bolts type damage as well.

> "SHM is a significantly complex and diverse problem that may not be solved in the immediate future but progress will come in small increments over long periods of time."

Summary

The AWE SHM method was able to detect a range of different damaged states, and order the states in terms of the severity of the damage. This was achieved for several different input excitations and without the need for either a numerical model of the structure, examples of the structure response for damaged states or availability of the input excitation.

This represents a useful improvement over previously published analysis on the benchmark structure. Future research will focus on issues concerning sensitivity and generalisation in order to adapt the AWE technique to tackle a wider class of SHM problems.

Acknowledgments

The author would like to thank the Condition Monitoring team at AWE for their support.

References

1. C. R. Farrar and K. Worden, An introduction to structural health monitoring, Phil. Trans. R. Soc. A 365, 303–315 (2007)

2. H. Sohn, Effects of environmental and operational variability on structural health monitoring, Phil. Trans. R. Soc. A 365, 539–560 (2007)

3. S. W. Doebling, C. R. Farrar, M. B. Prime and D. Shevitz, Damage identification and health monitoring of structural and mechanical systems from changes in their vibration characteristics: a literature review, Los Alamos National Laboratory report LA-13070-MS (1996)

4. C. R. Farrar, S. W. Doebling, and D. A. Nix, Vibration based structural damage identification, Phil. Trans. R. Soc. A 359, 131–149 (2001)

5. O. S. Salawu, Detection of structural damage through changes in frequency: a review, Engineering Structures 19 (9), 718-723 (1997)

AUTHOR PROFILE

Fazal Jamal Ahmad

Fazal is a Chartered Engineer having graduated with a Masters Degree in Control Systems from Imperial College, London.

Before joining AWE in 2007, Fazal spent 11 years designing and implementing control systems using robust control techniques in the process, power generation, automotive and aerospace industries.

Fazal is currently completing a Doctorate in Nuclear Engineering at Imperial College, London under an AWE sponsorship programme.

Introduction to Formal Methods

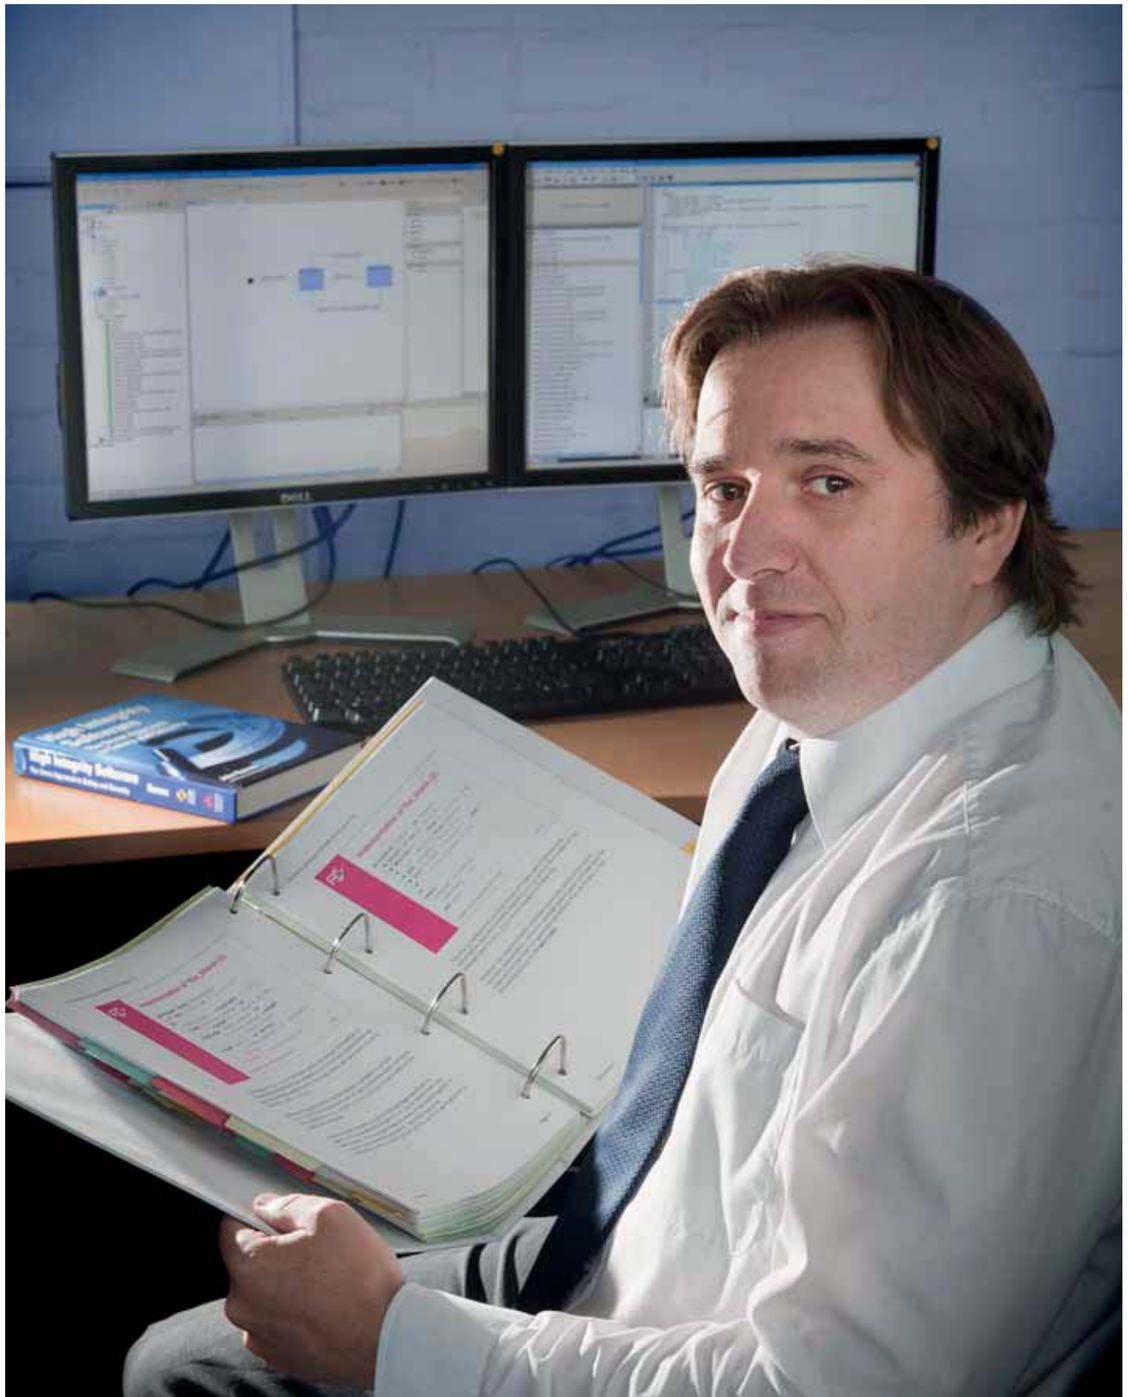

Safety and security critical systems characterise those systems whose failure would have catastrophic consequences. Whilst a lot of risks are mitigated by skilled personnel who manage operational aspects of a system, a certain amount of automation exists that is beyond the control of human operators.

Computers (the physical hardware and software) provide the automation for most modern systems. As a consequence, the burden of responsibility for the correct operation of the system lies increasingly with the suppliers of the computer hardware and software. Formal methods aim to mitigate the risks associated with the automatic parts of a system by ensuring that they operate correctly under all circumstances.

The stages involved in the development of computer based systems are similar to those of any engineering discipline. Ideally, engineers will develop a specification of the system from the customer's requirements. Then the system is designed from the specification. Finally, the system is implemented according to the design.

The engineers should provide sufficient evidence to demonstrate that the specification satisfies the requirements, the design is concordant with the specification, and the implementation is faithful to the design. It is then possible to conclude that the implementation fulfils the customer's original requirements.

Formal methods do not supersede other approaches, but must coexist and integrate with other non-formal methods and add mathematical rigour to the process.

Since the delivery of the final product is the ultimate goal of a systems engineering project there is a strong temptation to save time and money by side stepping the early stages of a development to tackle the implementation immediately. However, a common cause of project failure is lack of understanding. The benefits of incremental approaches to systems development, and formal methods in particular, can be summarised as:

- Specifications force the engineer to ask questions about the system being specified
- They can identify problems early in the development stage
- Many implementation issues can be deferred until the design has been completed
- Formal methods provide an abstract but unambiguous description of the problem to be solved
- They enable analysis phases to check that a specification meets its requirements and prove that subsequent design steps preserve correctness
- Costs are reduced because errors discovered later in the development are exponentially more expensive to rectify than errors discovered earlier
- Formal methods provide the evidence that is necessary to meet the highest safety standards

In recent years there has been an increased use of formal methods in industry. AWE has been at the forefront of the use of formal methods, with about 20 years of experience and good results. This article is aimed as a gentle introduction to formal tools and techniques that are being used at AWE and gives a short demonstration of their applicability.

Overview

Formal methods began to emerge during the 1970s, but their mathematical origins are much older. The mathematical notations adopted by formal methods are those that were used to investigate the foundations of mathematics itself. Specifically, the notations of set theory and the first order predicate calculus are ubiquitous in formal methods.

Formal methods use first order predicate logic and set theory to specify the properties of a system formally. The term 'first order' constrains the quantification of values by allowing variables to range over the elements of a set, but forbids them from ranging over higher order objects such as predicates. First order predicate logic is sufficient for the needs of formal methods.

The Z notation, for example, was first developed by the Programming Research Group at the Oxford University Computing Laboratory in the 1970s as a succinct way of formally specifying software. Z was used by AWE (in collaboration with Logica), to specify and build a compiler for AWE's microprocessor, ASP. This Demonstrably Correct Compiler

(DeCCo) was developed formally to ensure that high level source code could be faithfully translated to low level assembly code to run on the ASP microprocessor. This eliminated the possibility of errors being introduced during the compilation phase.

Most contemporary approaches to systems engineering involve the use of Computer Aided Design (CAD) tools. Tools help to minimise human error by preventing their users from making silly mistakes. Another formal method called the B Method, which evolved from Z, was developed with tool support in mind. The B Toolkit was developed by B-Core UK to support the specification and design of software systems using the B Method. The B Toolkit allows top level specifications to be constructed and, through a series of refinement steps, allows lower level designs to be developed incrementally. When a design has accumulated a sufficient amount of detail, the tool is capable of generating the software automatically. This eliminates the possibility of errors being introduced by writing the software by hand.

Collaboration between AWE and Ib Sorensen at B-Core UK resulted in the addition of hardware component libraries and a hardware descriptive language (VHDL) code generator to the B Toolkit. This enabled the development of digital hardware using the same approach.

FIGURE 1

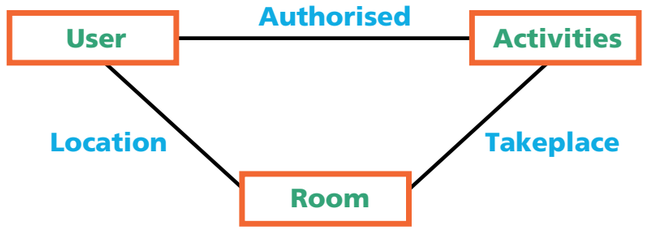

Access control system diagram.

A necessary feature of the B Toolkit is its proof obligation generator. Each design step must be shown to preserve the correctness of the earlier stages in the development. A design step cannot introduce behaviour that was not part of the original specification. The tool generates proof obligations that must be proven or, in B Method parlance, discharged to demonstrate preservation of correctness.

> "Formal methods do not supersede other approaches, but must coexist and integrate with other non-formal methods and add mathematical rigour to the process."

In many cases the tool itself is capable of discharging the generated proof obligations. In all other cases, interaction with the tool is necessary to construct the proof. Of course, it might be the case that certain proof obligations cannot be discharged. These are caused by inconsistencies in the design which must be rectified for the development to proceed. In this way, the tool helps to expose problems that might otherwise go unnoticed.

Tools are being developed to support more of the systems development lifecycle. The latest incarnation of the B Method is called Event-B. Instead of focusing on software development in particular, Event-B has evolved to model and develop systems in general. The notation used by Event-B is similar to the B Method and introduces a notion of refinement whose purpose is to manage the complexity associated with modern day systems.

Given a high level view of the system, refinement is the way lower level design details can be added incrementally and in such a way that correctness can be preserved. In Event-B, refinement takes two forms:

- Horizontal refinement, in which new features of the system can be introduced
- Vertical refinement, in which design detail can be added to existing features

Both forms of refinement are necessary to manage the inherent complexity of modern systems. By delaying the introduction of low level design/implementation details, it is possible to begin the development process with a more abstract, system wide specification that is easier to understand.

The tool support for Event-B is called the Rodin Platform. The word 'platform' indicates that it is not a single tool, but comprises a suite of tools. There is a 'core' tool which handles Event-B models and provides proof obligation generation, but there are also 'plug-in' tools that provide many other ways to analyse an Event-B model.

The ProB plug-ins, for example, include model checking support to detect whether it is possible violate a proof obligation. Using an Event-B representation of a system ProB can explore the possible operational states of the system automatically to detect whether they violate any properties that should always be true. Other plug-ins, such as the UML-B plug-in, provide different ways for the engineer to interact with the Rodin platform.

FIGURE 2

Invariant

$authorised \in User \leftrightarrow Activities$

$takeplace \in Room \leftrightarrow Activities$

$location \in User \nrightarrow Room$

$\forall u,r \cdot u \in dom(location) \wedge location(u) = r \Rightarrow takeplace[\{r\}] \subseteq authorised[\{u\}]$

Invariant declaration of access control system in Event-B.

Demonstration

Consider an access control system with the following requirements:

- Users are authorised to engage in activities
- Activities take place in rooms
- Users can access a room provided they have authority to engage in the room's activities

The final requirement tells us that access to the rooms must be controlled, but it does not tell us how it is to be achieved. This allows us to specify the problem before designing a solution. Figure 1 shows the relationships between objects in the access control system. The lines, representing relationships, connect the boxes which represent objects. The close correspondence between the requirements and this diagrammatic representation gives us some confidence that we understand the problem, even though we have done nothing formal yet.

Assuming that users, rooms and activities are fixed then the objects User, Room and Activities can be represented formally as sets in Event-B. The relationships between the objects are subject to change whilst the access control system is in operation, so the relationships authorised,

FIGURE 3

```
event Enter
  any u r
  where
    u ∈ User
    u ∉ dom(location)
    r ∈ Room
    takeplace[{r}] ⊆ authorised[{u}]
  then
    location:= location ∪ {u ↦ r}
end
```

```
event Leave
  any u r
  where
    u ↦ r ∈ location
  then
    location:= location \ {u ↦ r}
end
```

Declaration of Enter and Leave events.

takeplace and location can be represented formally as variables in Event-B.

Variables must be given a type in Event-B. Such information is specified as an invariant in the Event-B model; the invariant can also be used to specify other useful information. An invariant specifies those properties that must remain true throughout the operation of the system being specified. Figure 2 shows the declarations of the invariant for the access control system.

In addition to the type information, the fourth statement of Figure 2 captures an important property of the requirements. It says that for all users u and rooms r, such that u is in room r, the activities that can take place in r are a subset of the activities in which u is authorised to engage. Since it is an invariant, this property must always be true.

Next, events are defined using the Event-B notation to specify how the system will evolve during its operation. There is an **Enter** event, which models a user entering a room, and a **Leave** event in which a user vacates a room. Figure 3 shows how these events are declared in Event-B.

FIGURE 4

```
event AddAuth
  any u a
  where
    u ∈ User
    a ∈ Activities
  then
    authorised:= authorised ∪ {u ↦ a}
end
```

```
event RemAuth
  any u a
  where
    u ↦ a ∈ authorised
  then
    authorised:= authorised \ {u ↦ a}
end
```

Declaration of AddAuth and RemAuth events.

An event can change the values assigned to variables. However, it can do so only when its guard is enabled. The guard, delimited by the where and then keywords, specifies the circumstances under which the event can occur. Even if an event is enabled then it is not obliged to happen, but it can happen.

Upon inputting these definitions into the Rodin platform, the necessary proof obligations are generated. For the **Enter** and **Leave** events all proof obligation are proven automatically by the tool. No extra effort is needed to show that the invariant is preserved by the events.

In order to demonstrate a horizontal refinement in which a new feature is introduced, we add new events, **AddAuth** and **RemAuth**, that change the authorisation relation for a given user and activity. Figure 4 shows the definition of these events in the Event-B tool. The action of **AddAuth** simply associates an activity with a user, whereas **RemAuth** removes an existing association from the authorised relation.

The ProB model checker repeatedly executes enabled events and checks whether any erroneous states are reached. In our demonstration the invariant is violated. The tool captures the sequence of events in the violation.

A user called **User1** is authorised to engage in **Activity1** and **Activity2**. This entitles the user to enter **Room2**. However, whilst in **Room2**, the user's authorisation to engage in **Activity2** is removed. Hence the invariant is no longer true because **User1** is still in **Room2** where **Activity2** takes place.

FIGURE 5

```
event RemAuth
  any u a
  where
    u ↦ a ∈ authorised
    ∀u.u ↦ r ∈ location ⇒ r ↦ a ∉ takeplace
  then
    authorised:= authorised \ {u ↦ a}
end
```

The corrected RemAuth event.

BOX 1

Mathematics of Formal Methods

Set theory is attributed to Georg Cantor (1845 – 1918) who used it to investigate notions of infinity. More generally, a set is simply a mathematical representation of a collection of related objects. A Venn diagram is a familiar way to represent sets pictorially. An example is shown in the figure below. The rectangle represents the universal set **Σ**. The circles labelled **A** and **B** within the rectangle are said to be subsets of **Σ**; this is written formally as **A ⊆ Σ** and **B ⊆ Σ**. The overlapping region of **A** and **B** represents the set of objects common to both **A** and **B**; this is referred to as the intersection of **A** and **B**, and written formally as **A ∩ B**. The entire shaded area represents all elements in **A** or **B**; this is referred to as the union of **A** and **B**, and is written formally as **A ∪ B**. If an element **e** belongs to the set **A** then **e** is said to be a member of **A**, which is written formally as **e ∈ A**.

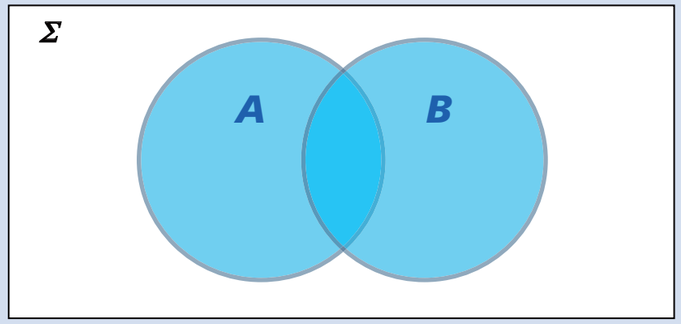

The general notion of a mathematical function can be formalised in set theory too. A function relates (or maps) the elements of one set to the elements of another. For example, the 'square' function (which multiplies a number by itself) maps the set of real numbers to the subset of non-negative real numbers. In general, functions can be categorised according to the sets they relate, and can themselves be considered elements of a set. The set of functions that map set **A** to set **B** is written formally as **A ↦ B**. Hence, **f ∈ A ↦ B** indicates that **f** is a function belonging to this set.

Formal methods use sets as a way of modelling collections of objects, and functions to relate them. For example, names, addresses and national insurance numbers could be modelled formally as individual sets, and the relationship between, say, national insurance numbers and names could be modelled as a function.

The earliest study of logic and reasoning can be dated back to Aristotle and Euclid around 300 BC. Subsequently, the nature of mathematical reasoning has been studied by many famous philosophers and mathematicians, including Descartes, Leibniz, Boole, Russell, Hilbert, Gödel and Turing. Boolean Logic, for example, is a calculus for propositions (true/false statements). In it, propositions are constructed using (among others) the logical connectives: and (**∧**), or (**∨**) and implies (**⇒**). The meaning of these connectives is demonstrated in the following truth table.

p	q	p ∧ q	p ∨ q	p ⇒ q
true	true	true	true	true
true	false	false	true	false
false	true	false	true	true
false	false	false	false	true

BOX 1

The first two columns give all possible true/false permutations for two propositions **p** and **q**. The remaining columns give the corresponding truth values when the propositions are combined using the operators. The proposition **p ∧ q** is true only when **p** and **q** are true, and **p ∨ q** is true when **p** is true or **q** is true. The truth values for **p ⇒ q** are slightly less intuitive. This proposition says that if **p** is true then **q** is true; it says nothing about **q** when **p** is false. Hence, when **p** is true and **q** is false then the proposition **p ⇒ q** is false and, regardless of the value of **q**, the proposition **p ⇒ q** is true when **p** is false.

Predicates, like propositions, make true/false statements. However, predicates are more general than propositions because their truth can depend on some variable quantity. For example, the statement '3 > 2' is a true proposition. If x is a variable then 'x > 2' is a predicate whose truth depends on the value of x. Predicate logic uses the same logical connectives as Boolean logic but, in addition, introduces the notion of quantification. A quantifier binds variables in a predicate. Given a predicate P(x), which depends on the variable x, the two forms of quantification are:

- Universal quantification: ∀x.P(x) says that for every possible value of x, P(x) is true
- Existential quantification: ∃x.P(x) says that there exists a value of x in which P(x) is true

If P(x) is the predicate 'x > 2' then ∀x.P(x) is obviously false because P(x) is false when x has the value 0, but ∃x.P(x) is obviously true because P(x) is true when x has the value 3. By introducing variables and quantifiers, predicate logic is considerably more expressive than propositional logic. The first precise development of predicate logic is attributed to Gottlob Frege in 1893. His aim was to show that all mathematical analysis and arithmetic can be founded on logic and set theory.

One solution to this problem is to strengthen the guard of **RemAuth** so that the authority for activities needed to enter a room is not removed whilst a user is in the room. Figure 5 shows this updated definition of the **RemAuth** event. If u is currently in a room then the guard checks whether authorisation for an activity a is a prerequisite to being in the room. If not the event is enabled, otherwise the event is blocked. This is sufficient to obtain a complete set of discharged proof obligations.

Through an iterated sequence of refinements (both horizontal and vertical) it is possible to introduce complexity in a controlled manner. Other features of the Rodin platform also help to manage complexity. The decomposition tool, for example, allows large specifications to be split into smaller, more manageable parts.

This simple demonstration highlights the use of formal methods as a rigorous mathematical proof tool and system engineering technique.

Summary

This article gives an introduction to the issues concerning safety critical systems, and the formal methods that aim to prove correctness of such systems. The mathematical foundations are presented, together with a demonstration of the tool support: the Rodin platform. This evolved from earlier tools that were tailored to software development. The Rodin platform, however, is tailored to system wide development. The tools contained within the Rodin platform cover the full range of system engineering activities from requirements capture to code generation.

AWE is currently using formal methods technologies on two projects. The first is the Co-design Architecture (CODA) project, in collaboration with the University of Southampton, which is working on a methodology based on Event-B for developing combined software/digital hardware systems. The second is an AWE project which is using formal methods to verify software for National Nuclear Security applications.

Recommended Reading

A. Hall, Realising the Benefits of Formal Methods, Formal Methods and Software Engineering, LNCS 3785, Springer (2005)

A. Hall, Using Formal Methods to Develop an Air Traffic Control (ATC) Information System, IEEE (1996)

Event-B/RODIN web site, www.event-b.org

S. King, J. Hammond, R. Chapman and A. Pryor, Is Proof More Cost-Effective than Testing?, IEEE (2000)

S. Schneider, The B Method – an Introduction, Palgrave (2001)

J. Spivey, An Introduction to Z and Formal Specifications, Software Engineering Journal (1989)

The Standish Group, The Chaos Report (1995)

J. Wing, A Specifier's Introduction to Formal Methods, IEEE (1990)

Acknowledgements

The author is grateful to Professor Michael Butler of the Dependable Systems and Software Engineering Group at the University of Southampton for providing the example.

AUTHOR PROFILE

Dr Neil Evans

Neil graduated from the University College of Swansea with a BSc in Computer Science. He went on to do an MSc in the Foundation of Advanced Information Technology at Imperial College and then a PhD at Royal Holloway, Univeristy of London. His PhD invesitigated information security using formal methods.

Neil has been a research academic at Royal Holloway, University of London and a reserach fellow at the University of Southampton, where he was part of the EU Rodin project. He is currently leading the Formal Methods team within AWE's Technology Division.

Analytical Chemistry of the Actinides

For the majority of analytical chemists the actinide or 5f series of elements represents uncharted territory. Not so at AWE where this region of the periodic table contains elements of fundamental importance to our technical mission. For as long as the UK has been involved in nuclear weapon design and manufacture, chemical analysis of actinide materials has played a vital part.

The chemical, metallurgical and nuclear properties of warhead components are very dependent upon the composition and purity of the actinide material from which they are fabricated. As a consequence specifications set out very stringent limits for adventitious impurity levels, isotopic composition, minimum actinide assay values etc. Chemical analysis is used to underwrite the quality requirement relating to component production, to support process development activities, pyrochemistry recycle operations, decommissioning programs and forensic analysis in support of national security. This article will provide some insight into the actinide series and some of the analysis undertaken at AWE.

Nuclear Engineering for Actinide Analytical Chemistry

AWE has a modern, purpose built actinides analysis facility with suites of interlinked gloveboxes for the safe handling of radioactive and toxic actinide materials, shown in Figure 1. Special features have been incorporated into this facility in order to cope with the corrosive aqueous environment resulting from the extensive use of concentrated mineral acids. For example, gloveboxes and extract systems are fully coated with an acid resistant fluoropolymer material and all filters and gaskets (used to maintain containment) are also suitably acid resistant.

Analytical instruments have to be extensively modified for interfacing to gloveboxes in order to meet the multiple challenges of safe containment, lack of space, ergonomic considerations and maintainability. Instruments are separated so that only the essential parts are inside glovebox containment, communicating using specialised electronic and optical breakthroughs with the control systems outside. Reagents are introduced to the glovebox through check valve reagent lines and gases through High Efficiency Particulate Air (HEPA) filtered gas lines. The main challenge when modifying instrumentation in this way is to maintain the high levels of accuracy, precision and sensitivity of the analytical techniques.

FIGURE 1

Gloveboxes at AWE actinide facility.

FIGURE 2

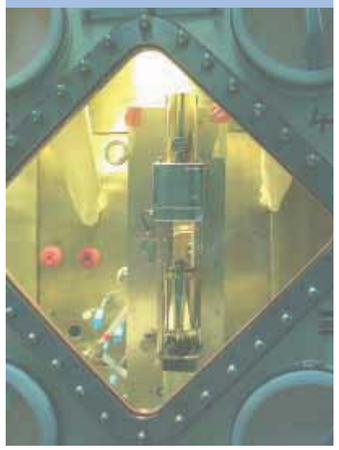

Interstitial analyser adapted for a glovebox enclosure.

FIGURE 3

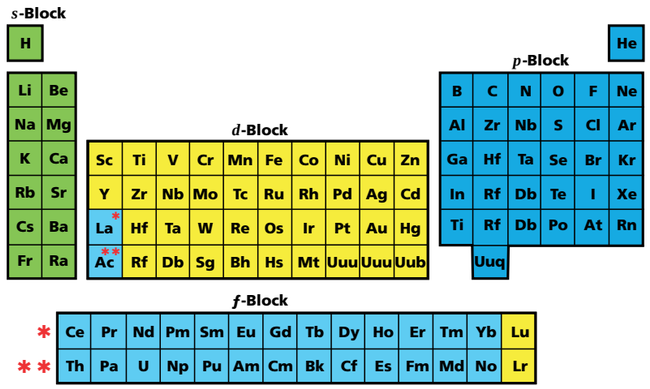

Periodic table.

Figure 2 shows a commercially available interstitial analyser adapted for a glovebox enclosure. The furnace portion of the instrument is separated from the analyser body and services to the furnace (cooling water, compressed air, gases, electricity) are supplied by specialised breakthroughs. The gases generated from the sample are passed out of the glovebox through HEPA filters, into the analyser and then discharged back into the glovebox.

The Actinide Elements

The modern periodic table is most commonly depicted with two rows set aside from the other elements, as shown in Figure 3. These rows are collectively known as the f block and are divided into the upper row, lanthanide elements, gradually filling up the 4f orbitals and the lower row, actinide elements, filling up the 5f orbitals. Box 1 provides information on orbitals and their importance with regard to the actinide elements.

The 4f orbitals of the lanthanide elements are lower in energy than the 5d orbitals. They are buried deep inside the atom, below the 5s and 5p orbitals, and so do not participate in bonding. The majority of lanthanide chemistry is therefore based around the 3+ oxidation state.

Compared as a series, the actinide elements resemble and share many characteristics with the lanthanide elements. However the light actinides, up to americium, which include the elements of most interest at AWE, have some rather interesting properties. Plutonium (Pu) is of particular interest as it occupies a transitional position in the actinide series.

In the lighter actinides the 5f orbitals are very similar in energy to the 6d orbitals and extend into space beyond the 6s and 6p orbitals, enabling them to participate in bonding and therefore exhibit variable oxidation states. After Pu, the heavier actinides become more like lanthanides as the actinide contraction takes effect. The 5f electrons become more tightly bound at higher energy levels and hence are less available for bonding. The different oxidation states of the actinide elements are shown in Figure 4. The most stable oxidation state in aqueous solution is in bold.

The oxidation state of the actinide elements vary depending on the coordinating ligand present, resulting in a variety of highly coloured complexes. Figure 5 shows a selection of colours that Pu forms in aqueous solution.

In the same way that transition metals are coloured due to d-d transitions, internal transitions

FIGURE 4

Oxidation state	Th	Pa	U	Np	Pu	Am	Cm	Bk	Cf	Es	Fm	Md	No	Lr
											2	2	**2**	
	3		3	3	3	**3**	**3**	**3**	**3**	**3**	**3**	**3**	3	**3**
	4	4	4	4	**4**	4	4	4						
		5	5	**5**	5	5								
			6	6	6	6								
				7	7									

Table of oxidation states for the different actinide elements.

FIGURE 5

Aqueous solutions of Pu in varying oxidation states.

between f orbitals in the actinides occur in the visible and near infrared region of the electromagnetic spectrum. However because these f-f transitions are forbidden by the selection rules (spin allowed but Laporte forbidden), the colours are not very intense. These transitions are characteristic for each oxidation state and so can be used in the identification and quantification of ions in solution.

The intense colours observed in actinide solutions are mostly due to ligand to metal charge transfer interactions which are allowed by the selection rules. These transitions involve electron transfer between molecular orbitals that are predominantly ligand like in character to those that are metal like in character. Solution colours are highly dependent on the nature of the actinide complex in solution i.e. by varying the ligand or oxidation state, the colour of the solution will change.

The unique properties of the light actinide elements are good news for the analytical chemist because, aside from the issue of radioactivity enabling the use of radiometric techniques, they offer a range options for their analysis. Some of the analytical techniques used are discussed.

Actinide Assay

The multiplicity of valence states achievable with Pu and uranium (U) in aqueous solution, and the relative ease with which one state can be reduced or oxidised to another, makes the redox titration a valuable means by which to quantify the actinide species. The change in solution potential as one actinide valence state oxidises to another can be monitored, using electrodes, as a function of the volume of oxidizing agent added.

This generates a typical s-shaped titration curve and the end point volume is calculated from the first derivative of this curve. Computer controlled autotitrimetry allows very precise metering of standard redox reagents at small incremental volumes, highly reproducible conditions for stirring, reagent addition and timing of sequences plus automatic temperature correction and data processing.

Special care is needed with the solution chemistry and many factors need to be evaluated. In low valence states actinides in solution are prone to oxidation

"The chemical, metallurgical and nuclear properties of warhead components are very dependent upon the composition and purity of the actinide material from which they are fabricated."

from air or dissolved oxygen, the tetravalent and pentavalent states tend to disproportionate and alpha radiation produces species that can either reduce or oxidise the actinides in solution.

A modified version of the Davies and Gray titration is used for the precise assay of U materials and is based on the U(IV)/U(VI) reaction.

Samples are dissolved in nitric acid then "fumed" with concentrated sulphuric acid to remove excess nitric acid which would otherwise interfere as an oxidizing agent. Concentrated phosphoric acid is added, which effectively provides the thermodynamic driving force to enable U in the hexavalent uranyl state to be readily reduced to the tetravalent state using the ferrous ion.

Normally the reverse reaction is favoured but in the presence of high concentrations of phosphate ion the reduction potentials of U(IV)/U(VI) and Fe(II)/Fe(III) are shifted significantly. This is due to the fact that the U^{4+} ion has a much larger charge to radius ratio than the $(UO_2)^{2+}$ ion and hence a much stronger tendency towards complex formation with phosphate ligands.

With U quantitatively converted to the U^{4+} state the complexing effect of the phosphate is negated by dilution with water. Under these conditions the reverse reaction occurs and a quantity of ferrous ion exactly equivalent to the U is produced. This ferrous ion is then oxidized titrimetrically to a potentiometric end point using potassium dichromate, catalysed by vanadyl ions.

A modified version of the "Corpel" titration is used for Pu assay and is based on the Pu(III)/Pu(IV) reaction. Total conversion of dissolved Pu to the Pu(III) state is achieved through initial addition of Ti(III). The Pu(III) is oxidized titrimetrically using Ce(IV) and taken just beyond the end point. The excess Ce(IV) is then backtitrated with Fe(II). This gives a sharper end point than trying to achieve a Ce(IV)/Pu(III) finish because the forward reaction can be completed quickly without risk of air oxidation of Pu(III).

Interstitial Analysis

Carbon, sulphur, oxygen, nitrogen and hydrogen are all measured. To determine the carbon content of an actinide sample, the sample is subjected to an induction current and allowed to combust in a stream of oxygen.

"Actinides have such a rich and diverse chemistry that they may interfere with analytical measurements"

FIGURE 6

Ion exchange chromatography.

The carbon monoxide (CO) and carbon dioxide (CO_2) generated are carried by the oxygen flow to the analyser where the CO is converted into CO_2. The total CO_2 content is then detected and quantified using infrared spectroscopy. In addition, the same combustion process converts the sulphur content of the actinide sample into sulphur dioxide (SO_2) which is again quantified using infrared spectroscopy.

Oxygen is determined by melting the sample, releasing the oxygen which reacts with the graphite crucible to form CO and CO_2. These gases are carried away from the sample to the analyser in an inert gas stream (either helium or nitrogen) where the CO is converted into CO_2 and again detected using infrared spectroscopy.

Melting of the sample also releases any nitrogen or hydrogen present, which are then carried in a suitable gas stream to the analyser and quantified using thermal conductivity detectors.

Separation Chemistry

Ion exchange chromatography is used extensively to separate elements from one another whilst retaining both elements in solution. There are two main reasons this separation is needed for the chemical analysis of actinides:

1. Safety: analytical instruments are often located outside of a glove box; radionuclides must be separated from the element of interest and left behind in the glove box
2. Analytical: actinides have such a rich and diverse chemistry that they may interfere with analytical measurements e.g. by absorbing or emitting at the same wavelength as the element that is being measured

Ion exchange chromatography involves passing the solution of interest through a column filled with resin, as shown in Figure 6. The job of the resin is to selectively adsorb some ions but not others, allowing separation to take place. Depending on factors such as the oxidation state of the ions in solution, the chemical form that they are in and sometimes even the size of the ions or complexes, precise conditions can be selected to achieve a successful separation.

Trace Elements

Trace elements are measured using glovebox adapted Inductively Coupled Plasma (ICP) spectrometers which offer the advantage of multi-element determination from the same sample. Two distinct types of ICP are used: Inductively Coupled Plasma Mass spectrometry (ICPMS) and Inductively Coupled Plasma Atomic Emission spectrometry (ICPAES).

ICPs act as very efficient ionisation sources. Argon gas passes through a quartz torch in an induction coil; an alternating current oscillates in the field generating electrical and magnetic fields at the top of the torch. A spark is then applied via a Tesla coil and some electrons are stripped off. These electrons are trapped in the magnetic field and accelerate in circles. They then collide with more argon atoms stripping off further electrons and a chain reaction occurs generating high temperatures and an annular plasma fireball at temperatures between 6000 and 10000 K.

In ICPAES large amounts of energy is transferred to the atoms and ions, promoting the excitation of their electrons to higher energy levels. When these excited atoms and ions return to their ground state or to lower excitation states they emit electromagnetic radiation in the ultraviolet/visible range of the spectrum.

Each excited element emits specific wavelengths and the intensity of the radiation is proportional to the element concentration. A spectrometer is used to separate the specific wavelengths of interest. Because atomic emission lines are very narrow lines, a high resolution detector is used which allows simultaneous detection of all elements of interest.

Due to the abundance of electrons in actinide elements and therefore the abundance of possible excited states, actinide emission spectra are extremely line rich. This can make it impossible to distinguish the emission lines for the impurity elements from the actinide background. Samples to be measured by ICPAES must undergo separation chemistry to remove the actinide prior to measurement.

However, since the majority of elements of interest are very different in mass from the actinide elements it is not necessary to separate the actinide matrix prior to measurement by ICPMS. ICPMS is an extremely sensitive technique with detection limits of the order of parts per trillion.

In ICPMS the ions produced in the plasma are drawn into the instrument by an extraction voltage of about 2000 V through a twin cone interface. The interface pump in the region between cones causes a pressure step where the cloud of gas, atoms and ions expand rapidly and speed up to supersonic speeds. This supersonic beam is then sampled by the second cone. An extraction lens 'pulls' the ions into the body of the spectrometer. Kinetic

"Solution colours are highly dependent on the nature of the actinide complex in solution"

FIGURE 7

ICPMS spectrometer.

energy is such that they overshoot this lens and then go through the focussing lenses.

The ions are accelerated up to 10 keV and focussed by the lenses into a beam of uniform kinetic energy which passes through the entrance slit into a magnetic field. The ions take a circular path according to mass/charge ratio. Following the mass separation the ions go through the electrostatic analyser and are focussed by energy onto an electron multiplier detector. Figure 7 shows a glovebox adapted ICPMS instrument.

Alloying Elements

X-Ray Fluorescence (XRF) is applied to the determination of major constituents and minor impurities typically at 50 ppm or above. This technique does not have sufficient sensitivity for trace element measurement, but has the advantage of superior accuracy and precision over techniques like ICPAES.

Incident X-rays of sufficient energy cause emission of photoelectrons from the inner shells of the atoms being analysed. Internal electronic rearrangement then follows with the hole being filled with an electron from a higher level shell. These electron transitions result in either the liberation of an Auger electron or a fluorescent X-ray of energy equivalent to the energy difference between quantum shells. It is the latter which is utilised in XRF analysis.

Due to the fact that these fluorescent X-rays result from inner shell electron transitions, the technique is largely independent of the chemical state of the element being analysed and can be applied to samples in solid or solution form.

The fluorescent X-ray photons emitted by this process can be characterised by virtue of their energy (Energy Dispersive XRF, EDXRF) or their wavelength (Wavelength Dispersive XRF, WDXRF). Both techniques are used in actinides analysis.

In EDXRF, a silicon semiconductor detector is used to collect all emitted X-rays simultaneously and energy discrimination is performed on the basis of pulse height analysis. The technique allows simultaneous analysis of all elements and is more amenable to analysis of large solid samples. Resolution is comparatively poor and interference effects are more predominant.

In WDXRF, a polychromatic fluorescent X-ray beam is collimated onto a crystal of well known lattice spacing, where the beam is diffracted into its constituent X-ray wavelengths according to the Bragg relationship. A goniometer is used to alter the angular relationship between detector and crystal and each X-ray is measured in turn by scanning with the goniometer. The analytical performance of WDXRF is usually superior in terms of resolution, accuracy and precision to EDXRF.

Speciation and Surface Analysis

Fourier Transform infrared spectroscopy can be utilised to investigate species on the surfaces of actinide samples. One version of this is attenuated total reflectance infrared spectroscopy, in which the sample is pressed against one side of a diamond crystal whilst an infrared beam is directed at the opposite side. The infrared beam penetrates the diamond crystal and impacts the surface of the actinide sample. When compared to a reference infrared beam, the infrared absorption by the sample can be calculated and an infrared spectrum of the surface species can be generated. The characteristic frequencies of these infrared vibrations may then be used to identify the surface species present.

Additional information about the vibration spectra of the surface species can be provided by Raman spectroscopy. This entails directing a monochromatic laser beam at the sample and collecting the light scattered by the sample of a different wavelength to the incident beam. This light is resolved into its component wavelengths in a spectrograph and converted into a Raman spectrum.

A technique used for looking at bulk material and surface species is X-ray diffraction. This analytical method involves directing an essentially monochromatic X-ray beam at the actinide sample and detecting the X-rays diffracted from the sample over a range of angles. The angles at which the X-ray beam is preferentially diffracted and relative intensities of these diffracted X-ray beams can be related back to the geometry of the structure of the actinide sample under investigation. In addition, the structure of surface films on the actinide sample may also be detected and resolved. The structures combined with the unit cell dimensions are characteristic of the respective actinide species.

> "Trace elements are measured using glovebox adapted Inductively Coupled Plasma (ICP) spectrometers which offer the advantage of multi element determination from the same sample."

Future Developments and Challenges

AWE maintains a comprehensive actinides analysis capability with the flexibility to respond to demand of all AWE programmes, whilst addressing important radioactive safety and environmental issues. The hazards of working with Pu solutions and the problems of processing aqueous Pu residues are well understood. These important factors demand that methods be developed to reduce analysis contact times and the volume of intractable residue produced, together with less overall reliance on aqueous chemistry.

Techniques that are amenable to automation or have the high sensitivity needed for small scale analyses offer significant advantages. Multi-element techniques are also of great benefit, particularly where interference effects can be minimised.

Where the need for precursory chemical separation is unavoidable, methods can be developed to improve selectivity and specificity of the separation. Methodologies may also be adapted and developed to produce analysis residues that are comparatively simple to process and result in waste forms that are compliant with radioactive waste disposal regulations.

Methods that eliminate the need for wet chemistry altogether are preferred, although it is difficult to envisage at this time how the current level of analytical capability could be sustained without it. Non destructive analysis methods and solid sampling techniques involving lasers and glow discharge sources are being considered.

BOX 1

Orbitals

Electron orbitals can be described as a 3 dimensional representation of where an electron is most likely to be found in space. The shapes of the f orbitals, are highly complex and are based on probabilities derived from wave functions.

The 4f orbitals are deeply bound inside the atom and do not extend far into space, hence they do not contribute to bonding. The 5f orbitals, however, are larger and more diffuse and extend far enough to overlap with orbitals from other atoms. In light actinides this leads to the 5f orbitals participating in bonding.

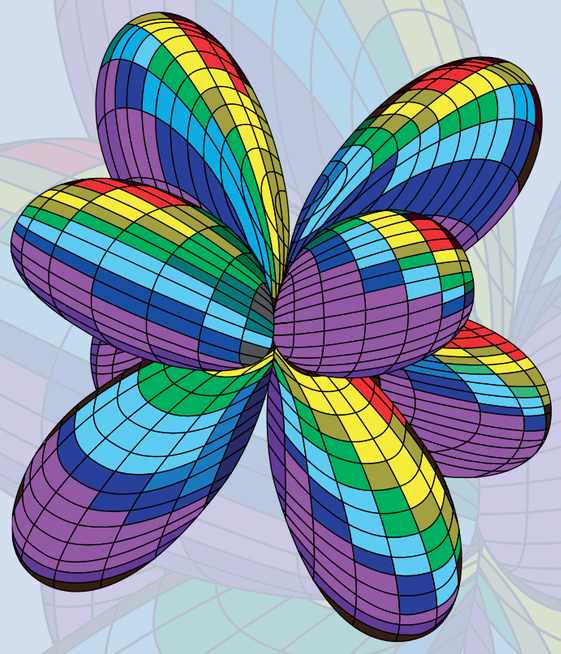

In light actinides the energy difference between the 6d and 5f orbitals is relatively low. As the number of electrons increases, however, the stability of the 5f orbital in relation to the 6d orbital increases and as a consequence the electrons preferentially occupy the 5f orbitals. The increasing atomic number and nuclear charge causes the effective nuclear charge experienced by a 5f electron to increase, leading to shrinkage in the atomic radius across the series, known as the actinide contraction. Due to this, in light actinides, it is more common to see actinide electronic configurations containing 6d electrons. Indeed the energy differences between the orbitals are in the range of chemical binding energies, and so the electronic configuration of an ion in a given oxidation state may differ according to the compound or the ligand it is associated with. Due to the actinide contraction, heavier actinides are more lanthanide like in behaviour. Plutonium occupies a very interesting position in the periodic table, sitting on a knife edge between bonding and localised behaviour, consequently its chemistry is extremely rich and varied.

Acknowledgements

The authors would like to acknowledge C Puxley and T Piper for their assistance with this article.

Figure 5 is copyright of Los Alamos National Security, LLC (2011). The authors wish to thank Los Alamos National Security for the use of Figure 5.

AUTHOR PROFILE

Pamela Thompson
Pamela joined AWE in 1987 working on the radiochemical analysis of underground test debris. During this time Pamela studied part time graduating with GRSC (Honours) from the Royal Society of Chemistry in 1992.

Pamela joined the Actinide Analysis team in 1997 carrying out the development of analytical chemistry techniques for the analysis of actinide materials. Pamela completed a company sponsored MSc in Radiometrics from the University of Liverpool in 2005. Pamela currently manages the Actinide Analysis Development team at AWE.

AUTHOR PROFILE

Katherine Hodkinson
Katherine started her career at AWE in 2005 as a Year in Industry placement student. Having graduated with a MChem (1st Class Honours) from Durham University, Katherine joined the Actinide Analysis team in 2006. Katherine is involved in developing analytical chemistry techniques for the analysis of actinide materials.

AWE's Outreach,

Major Events and Collaborative Activities

A few key events in which AWE participated during the later part of 2010 and the early part of 2011 are presented.

Plutonium Futures – The Science

The series of conferences entitled 'Plutonium Futures – The Science' convenes every two years and provides an international forum for the discussion of current research on the physical and chemical properties of plutonium and other actinide elements.

The most recent of these conferences was hosted in Keystone, Colorado, USA, on 19-23 September 2010 and was co-sponsored by Los Alamos National Laboratory, Lawrence Livermore National Laboratory and the American Nuclear Society.

Presentations covered the breadth of condensed matter physics, detection and analysis, materials science, nuclear fuel cycle and environmental behaviour.

Doctor Mark Read (AWE Team Leader Computational Chemistry), presented in the Surface Science and Corrosion session on the computer modelling of plutonium dioxide. Further oral presentations were given in the Detection and Analysis sessions by Terry Piper (AWE Group Leader Materials Characterisation) on the characterisation of uranium ore concentrates and by Pam Thompson (AWE Team Leader Actinide Development) who discussed plutonium isotopic analysis without radiochemical separation.

Doctor David Geeson, AWE's representative on the International Advisory Committee, presented a bid for AWE to co-host the next conference at the University of Cambridge in July 2012. Other co-hosts include the Commissariat à l'Energie Atomique et aux Energies Alternatives, France; Institute for Transuranic Elements, JRC/EU; and the UK National Nuclear Laboratory. The bid was subsequently accepted by the committee.

For information about Plutonium Futures 2012, please visit www.PuFutures2012.co.uk.

Condensed Matter Materials Physics (CMMP) 2010

The CMMP was held on 14-16 December 2010 at the University of Warwick, at which there was a strong AWE science and technology representation.

Doctor Graham Ball, AWE Group Leader for Material Modelling, chaired the session on 'Matter Under Extreme Conditions' for the second year.

Doctor Ball said: "Our session, featured invited talks by Doctor Mark Read (AWE Team Leader Computational Chemistry) on actinide oxides and Dan Eakins (Institute of Shock Physics) on shock-driven reactions in metal powders, plus contributed talks from Edinburgh and Warwick Universities."

The event featured an AWE exhibition showcasing work in materials modelling, Technical Outreach, Orion, and the Institute of Shock Physics.

Steering Committee members from 'Plutonium Futures – The Science'

CMMP continues to be an excelent forum for specialists to share and present the latest advances in their fields. AWE is working towards CMMP11 at which Professor Yogi Gutpta, Washington State University, has agreed to be the plenary speaker for the Matter Under Extreme Conditions session.

Group of Experts in Mitigation Systems (GEMS) 2011

The GEMS colloquium was successfully held on 5-7 January 2011 at Imperial College London. Led by the Centre for the Protection of National Infrastructure with AWE, Imperial College London and the Institute of Shock Physics co-sponsoring.

The three day forum attracted over 100 government and non-government specialists in blast mitigation. Presentations covered blast effects on structures, personnel and vehicles, mitigation of blast and fragments threats, modelling, experimental trials, and international terrorism.

Guest speaker, Professor Nick Jennings, delivering the after-dinner address at the GEMS dinner

Doctor Andrew Jupp (AWE Managing Director) said: “The work of GEMS is critical to our capability in support of the UK’s National Nuclear Security programme. Through GEMS, AWE is able to build vital collaborations across national and international fronts an example of which are our strong Strategic Alliances with Bristol, Cambridge, Cranfield, Heriot-Watt and of course Imperial.”

Guest speaker, Professor Nick Jennings, Head of Computer Science at the University of Southampton and Chief Scientific Adviser to the UK government on computer intelligence and cyber terrorism, said: “I am very pleased to be invited here today to embrace the excellent work of GEMS and am keen to emphasise the importance of this group in debating some of the UK’s complex security challenges."

Other notable figures who attended included representatives from the Defence Academy, Foreign and Commonwealth Office, Health Protection Agency, Metropolitan Police, Home Office, and MOD.

High Performance Computing conference endorses AWE’s collaborations with academia

The fourth Many-Core and Reconfigurable Supercomputing Conference (MRSC) took place on 11-13 April 2011 at the University of Bristol, an AWE Strategic Alliance partner.

Co-sponsored by AWE, MRSC is the premier event in Europe focussed on high performance computing (HPC) using advanced technologies. This is an important area of research at AWE, in support of the calculations required in weapon physics, materials science and engineering.

David Turland (AWE Technical Lead Advanced Technologies) said: “We are, once again, entering a necessary disruptive change in both the processor architectures used in HPC systems and techniques required to develop applications to target these systems. The MRSC provides an important meeting point for researchers and industry to discuss exactly these issues.”

Professor Geoff Allen (AWE William Penney Fellow) an expert in materials science and Head of the Interface Analysis Centre at the University of Bristol, attended the event to show his support to MRSC.

Professor Allen said: “I am delighted to support AWE and its association with MRSC. The conference provides an excellent knowledge forum for researchers and vendors actively involved in the development of methods, tools and applications for emerging computing platforms.”

As part of the event, the University of Bristol hosted a tour of its £7 million ‘BlueCrystal’ supercomputer, which opened in 2008. The facility revolutionises research in areas such as climate change, drug design and aerospace engineering.

If you are involved in an AWE technical event that you would like the editorial team to consider featuring in future editions of Discovery, please contact:

Paul Sagoo
Events and Communications Manager
Email: paul.sagoo@awe.co.uk

Discovery

Find out more about AWE at our website:

www.awe.co.uk

Comments and suggestions regarding this journal, please email:

Discovery@awe.co.uk

For further copies of this journal and details of other AWE publications, please write to:

Corporate Communications Office
Building F161.2
AWE Aldermaston
Reading
Berkshire
RG7 4PR