The Science and Technology Journal of AWE

Discovery 21

December 2010

Discovery

Contents

cover image
Adhesive columns formed during peel test.

21

I am delighted to write the foreword to Discovery 21 in my role as Director Strategic Technologies within MOD where my responsibility is to ensure strategic capabilities are developed and maintained to meet HMG requirements, many of which are AWE-centric.

AWE has undergone a significant investment and enhancement of capabilities in recent years. The depth and breadth of current AWE science and technology is epitomised by Discovery. It is particularly satisfying to see ample evidence of the quality of staff from articles appearing in this edition.

One article is on High Performance Computing, an area needing continuous development to model the science underpinning nuclear warheads in an era of no underground tests – a real technical challenge for stockpile certification. Another key component of certification is high energy density physics to simulate nuclear weapon environments, and it is pleasing to see two articles on plasma physics, an area where significant investment is also ongoing in the new Orion laser. The manufacturing and material science associated with organic materials is equally complex and requires strong scientific understanding to support the stockpile as it ages. The adhesives article also illustrates the role of outreach to support AWE programmes.

I do hope you enjoy this edition of Discovery.

Peter Sankey OBE
Director Strategic Technologies, MOD

Photonic Displacement Interferometer

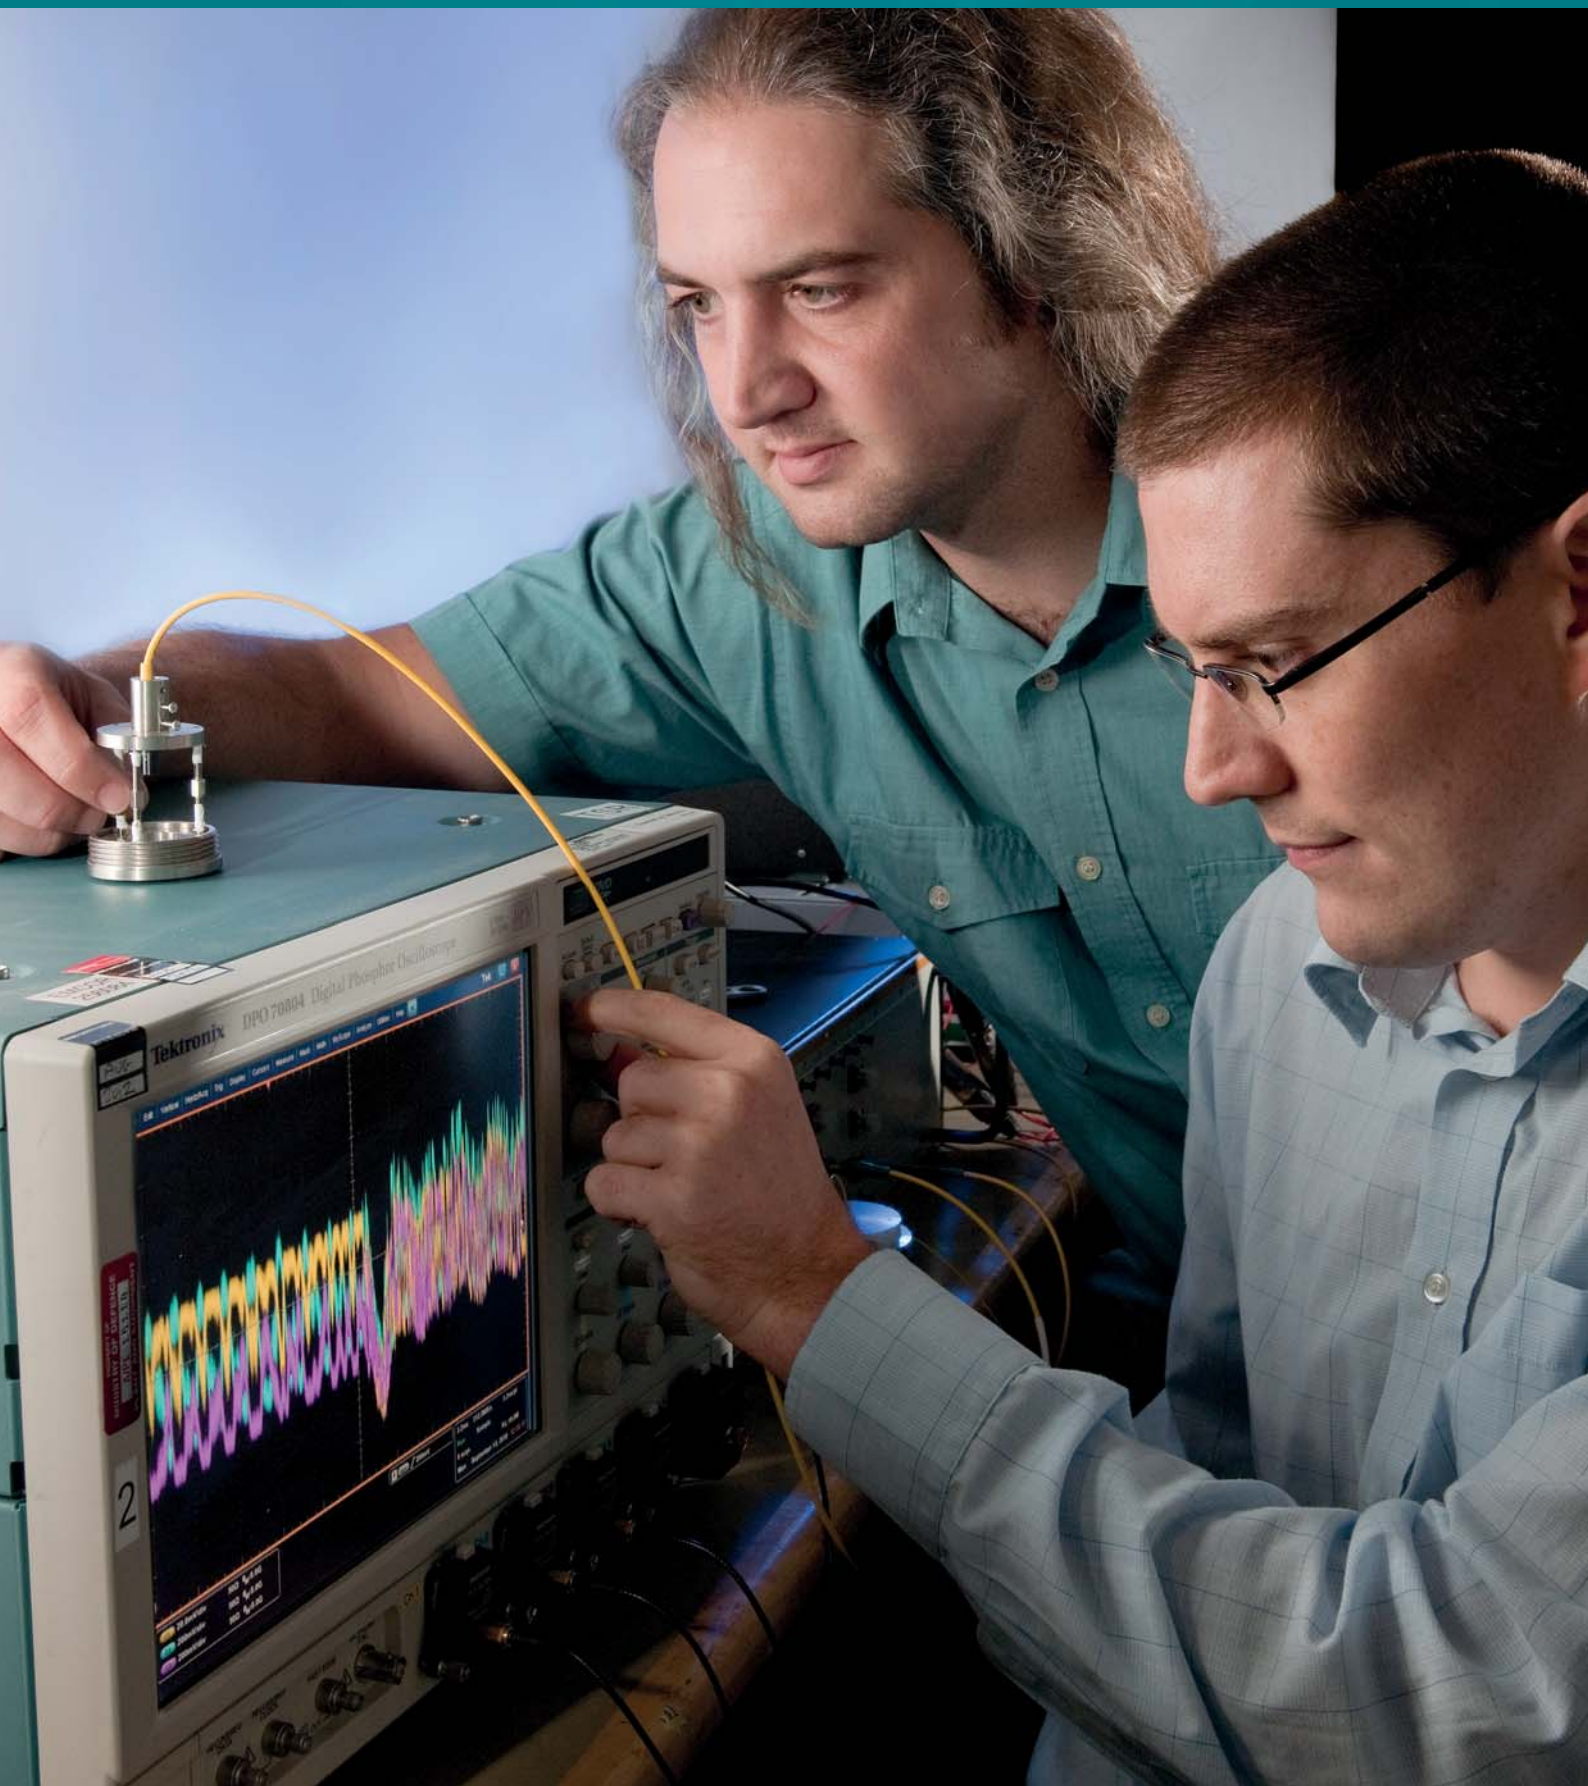

An interferometer has been developed for use in the investigation of impulse and thermo-mechanical effects arising from the deposition of a several nanosecond pulse of either keV X-rays or MeV electrons into materials. Energy deposition rates are slow compared with the Debye frequency and so the energy can be assumed to be immediately coupled to the lattice phonons thus increasing the pressure.

The resulting pressure wave causes back surface displacements of the material in the order of a micron, with peak velocities of tens of metres per second. In this regime conventional interferometers used in shock physics are unsuitable.

Impulse tests have been conducted on the Saturn X-ray pulsed power machine at Sandia National Laboratories (SNL) to demonstrate the capability of this interferometry technique.

Photonic Displacement Interferometer

An interferometer required to diagnose surface motion resulting from thermo-mechanical effects must be able to measure very small displacements, velocities, and also distinguish reversals in direction. A normal VISAR (velocity interferometer system for any reflector) is inappropriate for this task because gaining a suitable fringe count from a long delay leg would cause the time resolution to be too large.

An interferometer set-up as a heterodyne velocimeter (Het-V), [1] as used in Hydrodynamics Division at AWE, would fail as too few wave cycles would be seen to perform an adequate sliding-FFT. A displacement measurement where the change in phase is proportional to the distance moved by the measurement surface is the only viable option.

The displacement interferometer developed, the photonic displacement interferometer (PDI), is of the conventional Michelson-type and shares many of its components with the all-fibre Het-V system. All of the components are from telecommunications stock and are therefore based upon a 1550 nm laser and single mode fibres. The polarisation state of the laser propagating along the single mode fibre can change and needs to be controlled with a fibre polarisation controller.

The key components used are the optical circulators and the 3x3 coupler; which act in an analogous way to their microwave counterparts. In the circulator light passed into one port is directed to the next port in a clockwise direction. The 3x3 coupler consists of three bare fibres twisted and fused together. The evanescent electric field, outside of the fibre core, of the light transmitted through one of the three fibres is coupled into the other two fibres. There is a coupling ratio of approximately a third between each output fibre.

As shown in Figure 1, light from the laser is split into two beams for the two arms of the interferometer, the reference and measurement arms. In each arm light travels from the splitter through the circulator and lens to a reflective surface, it is then collected again by the lens, travels back along the same fibre to the circulator. The circulator then directs the light to the coupler.

FIGURE 1

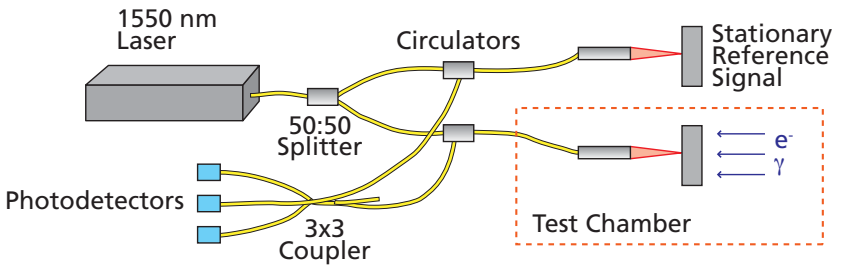

PDI schematic.

The two beams are coupled within the 3x3 coupler using two of the inputs, leaving the third blank. A third of the reflected measurement and reference beams are collected by each fibre where they interfere. The modulation depth of the interference fringes (measured by the photo detectors) is proportional to the cosine of the distance moved. The fringes are given by

$$I_1 = [a^2E_1^2 + d^2E_2^2 + 2adE_1E_2cos(kx+\phi_1)]$$

$$I_2 = [b^2E_1^2 + e^2E_2^2 + 2beE_1E_2cos(kx+\phi_2)]$$

$$I_3 = [c^2E_1^2 + f^2E_2^2 + 2cfE_1E_2cos(kx+\phi_3)]$$

where I_i is the intensity, x is the position of the measurement surface, k is the wave number of the laser, E_1 and E_2 are the amplitudes of each of the beams and ϕ_i is a constant phase and is a function of the coupling ratios, a^2 through to f^2. The coupling ratios are nominally a third in an ideal case, although will vary due to the manufacturing tolerances. It can be shown for a 3x3 coupler that is symmetric and lossless that [2,3]

$$a^2 + b^2 + c^2 = 1$$

$$d^2 + e^2 + f^2 = 1.$$

Since the laser intensity must be conserved between the light entering and exiting the 3x3 coupler it can be shown that the extra phase elements, ϕ_i, must all differ by 120° such that the sum of the three fringes will equal the two input intensities, E_1^2 and E_2^2. With phase elements differing by 120° the cosines will sum to zero.

"The low photon energy results in a short penetration depth into the test sample leading to the material experiencing a large energy density."

The phase is measured from the change in intensity of the fringes and reversals in direction are distinguished by using all three fringes similar to conventional quadrature, in this case it has been named triature.

Considering the ideal case where the intensity of each interferometer leg is equal and the 3x3 coupling ratios are all one third causing each cosine to differ in phase by 120°, then the phase angle, kx, can be calculated from all three fringes as [4]

$$\tan(kx) = \sqrt{3}\,\frac{I_3 - I_2}{2\,I_1 - I_2 - I_3}$$

FIGURE 2

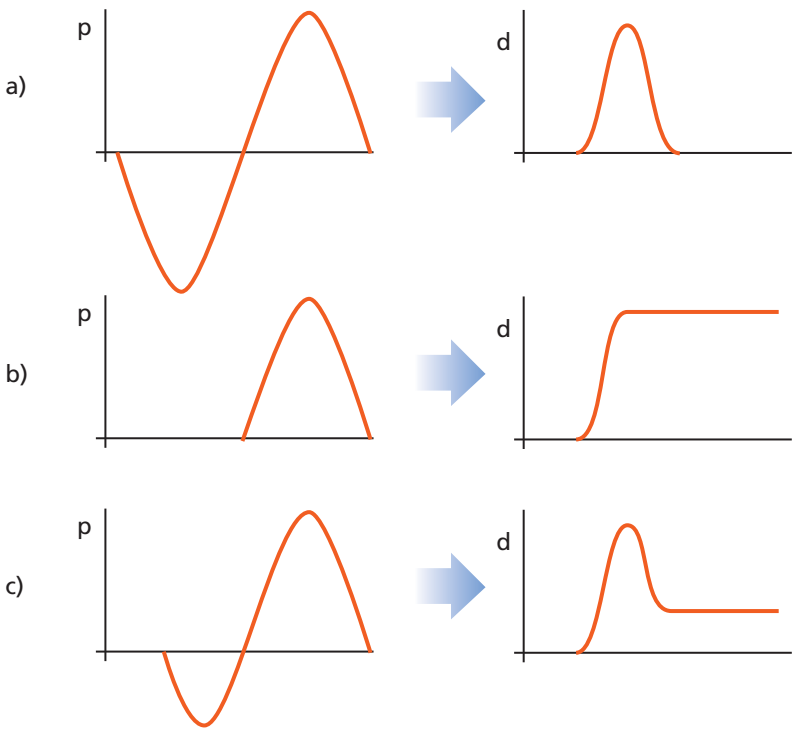

The left side shows pressure waves generated by (a) no phase change, (b) vapour dominated impulse and (c) melt dominated impulse generation. The resulting particle displacement is shown on the right.

PDI measurements at Saturn

The PDI system was used during a series of impulse experiments at the Saturn facility. Saturn provides a source of X-rays by discharging a very high current through a medium. This medium forms a plasma and transitions between electron states produce emission lines which are characteristic to the medium used.

Normally the medium is a metal in the form of a series of thin wires. The shots used a puff of Argon instead which emitted a spectrum located about 3.1 keV, in the Argon k-shell, with a yield of 10-20 kJ.

The low photon energy results in a short penetration depth into the test sample leading to the material experiencing a large energy density. High temperature rises are generated creating high pressures and possible phase changes close to the sample's surface. In particular vaporisation generates a large impulse which propagates through the sample material.

FIGURE 3

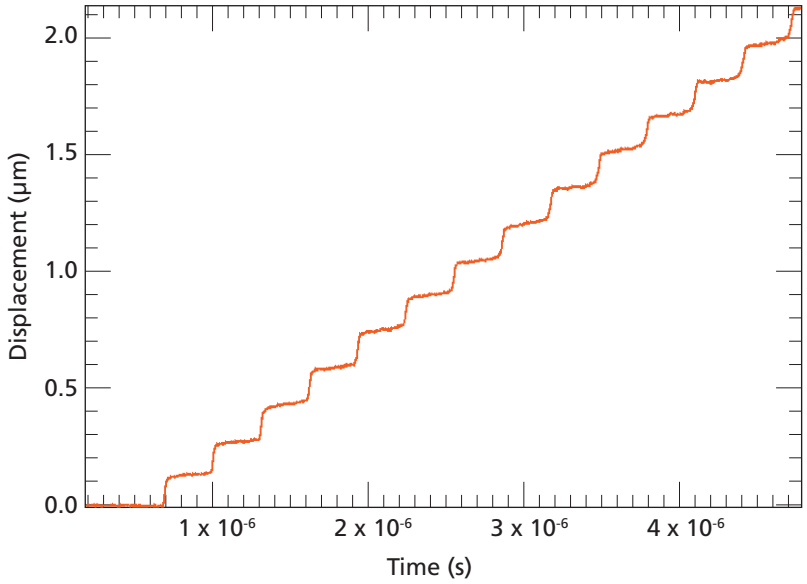

Displacement data for a 0.5 mm thick gold sample. The impulse for this sample was 4.9 Pas.

The effect of pressure waves entering the front surface for the three cases of no phase change, melt and vaporisation is illustrated in Figure 2. The rear surface displacement is also shown and assumes the material exhibits an elastic response where the displacement at the rear is proportional to the time integral of the pressure at the front surface.

X-ray impulse experiments require knowledge of both the total and time dependence of the impulse entering the sample. Previous experiments have relied on measuring the shape of the propagating stress wave using carbon gauges and measuring the total impulse using different techniques such as TRIM (time resolved impulse) gauges. PDI with triature however allows the simultaneous measurement of both quantities using the same technique and is completely non-invasive.

"The key to measuring the total impulse relies on the ability of the PDI system to measure displacements for a long period of time regardless of the number of direction reversals."

"High temperature rises are generated creating high pressures and possible phase changes close to the sample's surface."

The key to measuring the total impulse relies on the ability of the PDI system to measure displacements for a long period of time regardless of the number of direction reversals. This capability means we can determine the average velocity of the sample after the impulse and from conservation of momentum determine the total impulse received.

The velocity of the centre of mass of the sample can be calculated from the average velocity of the rear surface. Where the response is elastic it can be performed over a single displacement step but where inelastic response occurs the average gradient must be found over an integer number of steps.

The measurements at Saturn were very successful. Several different materials were investigated providing numerous examples of melt and vaporisation generated impulse, which will be used for future model development and validation. Each sample was chosen with a high aspect ratio such that a one dimensional response was expected for several stress wave transits, before edge effects become dominant. Figure 3 shows displacement data acquired for a 0.5 mm thick gold sample, where vaporisation has occurred. Characteristic displacement steps associated with the impulse are easily resolved with typical resolutions of 10-20 nm. The average velocity was measured to be 0.5 ms^{-1} giving an impulse of 4.9 Pas.

Summary

A displacement interferometer for suitable X-ray induced impulse measurements was successfully developed and has been tested at the Saturn facility.

Work is ongoing at the Z facility. The Z facility generates much higher yields of cold X-rays allowing a greater range of materials to be studied under X-ray impulse effects. It is also interesting from a modelling perspective as a result of having a significantly shorter pulse time than Saturn.

"The Z facility generates much higher yields of cold X-rays allowing a greater range of materials to be studied under X-ray impulse effects. It is also interesting from a modelling perspective as a result of having a significantly shorter pulse time than Saturn."

Acknowledgements

The authors would like to thank S. C. Jones and G. Chantler of Sandia National Laboratories for their support in the development and testing of the interferometer, and also N. Barnes and R. Burrell for their contribution towards the Saturn experiment.

References

1. O. T. Strand et al, Rev. Sci. Instrum., 77, 083108 (2006)
2. R. G. Priest, IEEE Transactions of Microwave Theory and Techniques, Vol. MTT-30, No. 10, October 1982
3. J. Pietzch, J. Lightwave Tech., 7, No. 2, 303 (1989)
4. D. H. Dolan, S. C. Jones, Rev. Sci. Instrum., 78, 076102 (2007)

AUTHOR PROFILE

Andy Sibley can be contacted on e-mail: andrew.sibley@awe.co.uk

Andy Sibley

Andy graduated in physics from Birmingham University in 1999 where he then obtained a PhD in condensed matter physics in 2002 specialising in low temperature superconductivity phenomena. He then spent three years working on the application of novel technologies to the field of optical metrology at the National Physical Laboratory. In 2005 he joined AWE working on thermal mechanical effects in materials.

AUTHOR PROFILE

Adrian Hughes can be contacted on e-mail: adrian.hughes@awe.co.uk

Adrian Hughes

Adrian graduated from the University of Manchester with a degree in Physics with Theoretical Physics in 2004. Adrian joined AWE after graduation to work on radiation effects in materials. Adrian is currently developing the computer simulation capability of the thermal mechanical effects department.

Imaging using Bragg Reflection Optics

This article describes progress in high resolution imaging using spherical Bragg reflectors. Spatial resolution of 2 μm has been demonstrated in HELEN experiments and it has also been demonstrated that multilayer optics are effective in the environment of short pulse laser experiments by their application to imaging wire targets heated by protons.

Proton heating is also currently an active research project with the aim of developing an experimental platform for producing warm dense matter under controlled conditions. These conditions can be diagnosed by X-ray imaging of the heated and expanding sample.

The use of geometrical optics to image X-ray sources is severely limited by the refractive indices of materials at short photon wavelengths. Lenses have only weak focussing effects[1] even when they do not heavily attenuate the radiation.

Mirrors have to be used at grazing incidence, and are useful for moderate (~10 μm) resolution imaging of soft (<5 keV) X-rays. The use of curved Bragg reflectors, either single crystals or multilayers, has produced images having ~2 μm spatial resolution. This is consistently the best of any microscopy system fielded on large laser systems. Also, because of their narrow bandwidth, object self-emission and high energy background can be greatly reduced.

Bragg's law states that

$$2d\sin(\theta) = n\lambda$$

where d is the spacing of atomic or coated planes in the chosen direction, θ is the angle of incidence the radiation makes with the planes, n is an integer known as the 'order' and λ is the wavelength.

Crystal reflectors are fabricated using a thin crystal wafer forced to conform to a curved substrate, whereas alternating coatings are applied to a substrate to form a multilayer reflector. At near normal incidence crystals can reflect from ~0.5-6 keV (n=1, first order), based on the limited range of lattice spacing available, although this is no guarantee that a sufficiently perfect, flexible, stable crystal is available.

Multilayer coatings can be as closely spaced as $2d$=40 Å, giving a maximum photon energy of ~0.3 keV, which because of their fabrication method can be more tightly curved.

Crystals generally have significantly lower bandwidth than multilayer reflectors, and therefore lower reflectivity, but have the advantage of a greater choice of detectors and filter materials.

Experimental arrangement

Experiments were carried out on the HELEN laser to demonstrate a Bragg reflection imaging system to provide high resolution data on the expansion of isochorically heated foils, heated by proton deposition.

The experiment required high spatial resolution (<5 μm) imaging coupled to a high time resolution (10 ps) detector to record the motion of a sample. The experimental arrangement is shown in Figure 1. Where the sub-picosecond beam for heating enters from the upper left. The

FIGURE 1

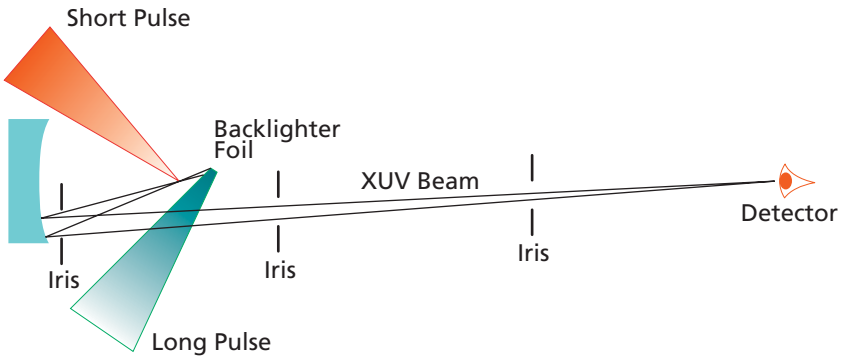

Schematic diagram of experimental layout.

nanosecond beam for backlighting enters from the lower left. The short pulse laser target was placed near one focal length from the Bragg reflector, with the backlighter target far from focus behind it.

Spherically curved multilayer ($2d$~44 Å) reflectors were positioned at a distance from the object s_o just greater than the focal length f to produce high magnification images of test objects and heated targets. The optical parameters were f=50 mm, s_o=51 mm, to give s_i~2500 mm (image distance) and M~50 (magnification) as per the lens equations

$$\frac{1}{f} = \frac{1}{s_i} + \frac{1}{s_o}$$

$$M = \frac{s_i}{s_o}$$

A long-pulse beam provided a source of quasi-thermal X-rays in the bandwidth of the Bragg reflector (~280 ± 2 eV). Laser pulses of ~100 J in 1–2 ns at 527 nm produced sufficient X-rays, with the laser spot size being adjusted to vary the image brightness. This source was ~5 mm from the object. The source, object and Bragg reflector were on a common axis.

The optical configuration was chosen to overcome the limiting spatial resolution of the time-resolved detector, an X-ray streak camera (XRSC), which had a contrast transfer function (CTF) equivalent to a Gaussian point spread function of 3 μm full-width at half maximum (the spatial resolution) in the object plane.

As shown in Figure 1, the system employed an off-axis aperture to reduce spherical aberration, exposure to debris and damage, and to ensure that the imaging rays passed to one side of the object. An aperture of 2.5 mm diameter was used 2.5 mm off the axis. With this arrangement, the contribution to the system resolution due to the optical aberrations should have been 1 μm[2] and the diffraction contribution 0.1 μm.

Additional apertures were used to decrease the amount of scattered unabsorbed laser light and expanded plasma emission. Thin aluminium filters of 0.16 μm thickness supported on a nickel mesh were used to stop unwanted emission reaching the camera.

The exposed area of the Bragg reflector was heavily damaged by X-rays, ablation, and debris. After each shot the reflector was rotated about its axis of symmetry, so a new section was presented. As the optical axis was not completely coaxial with the rotation axis, this necessitated re-alignment of the system.

In order to assess the spatial resolution of the optical system, a time-integrating CCD detector was used in place of the XRSC to record static two-dimensional images of a backlit object. The CCD had a pixel pitch of 13.3 μm, equivalent to 0.27 μm in the target plane. The test objects were grids of 1000–2000 lines per inch (LPI) used exclusively with the CCD detector.

Results

Figure 2a shows a CCD camera image of a backlit 1000 LPI grid and a corresponding lineout. The fine diagonal grid is a contact radiograph of the nickel mesh support on the 0.16 μm filter. The area of low signal in the bottom left is a repair to the filter. The grid was mounted over a pinhole to aid alignment. By convolving a Gaussian point-spread-function (PSF) of 2.5 μm FWHM with a synthetic grad pattern, the modulation pattern in Figure 2b was reproduced. This was typical of a number of shots which showed resolutions of 2-3 μm.

A streaked image of a backlit, undriven 20 μm wire is shown in Figure 3; where the horizontal axis is time, vertical axis is space and signal intensity increases from black to white. The spatial resolution from the wire edge is 5.5 μm. Given the CTF-equivalent

"This was typical of a number of shots which showed resolutions of 2-3 μm."

FIGURE 2

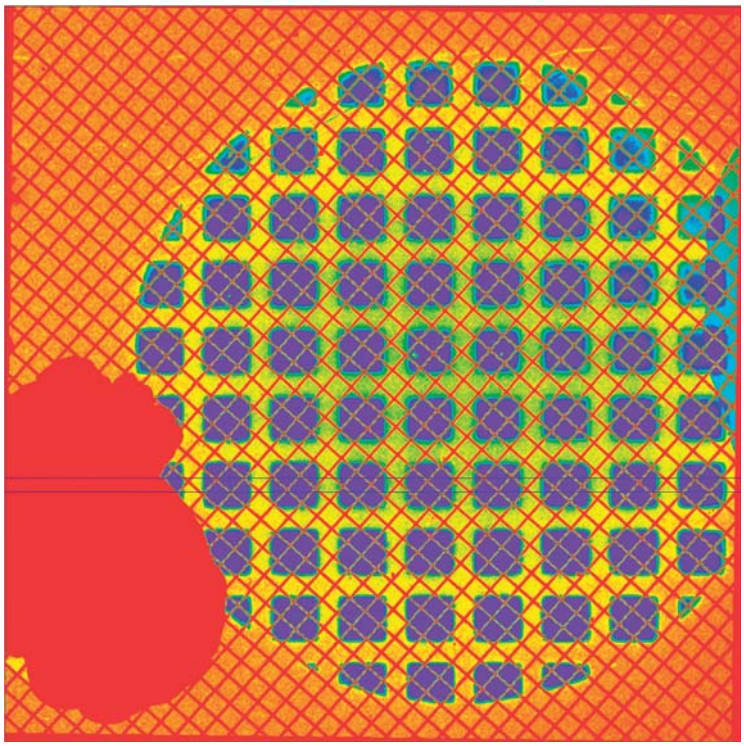

(a) X-ray CCD record of backlit object with associated 1000 LPI grid.

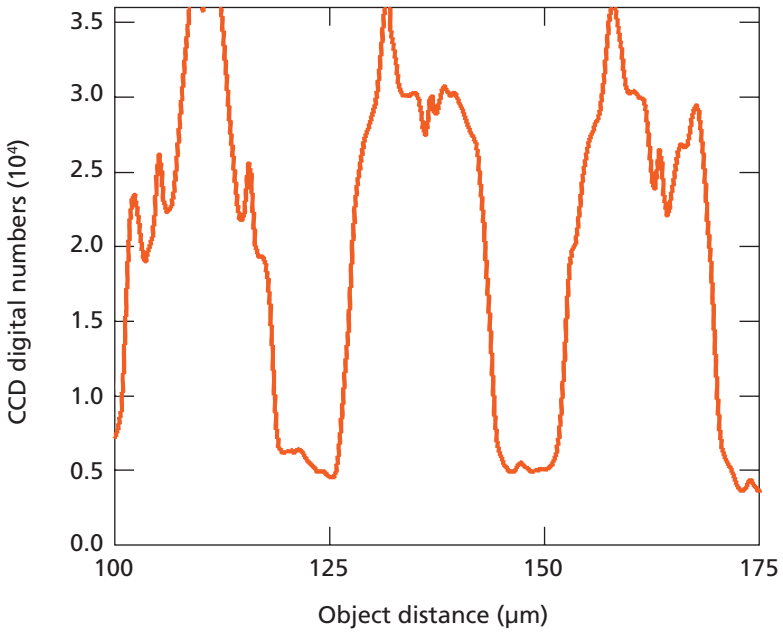

(b) Gaussian PSF modulation pattern.

of the streak camera (3 μm) and the spatial resolution of the time-integrated imaging system (2–3 μm), this is only slightly worse than the expected resolution of the streaked imaging system of ~4 μm. There is no evidence of expansion of the wire due to the backlighter radiation field evident in Figure 3.

A foil disc (15 μm thickness, 175 μm diameter) is imaged in the same way, shown in Figure 4. The short-pulse laser beam was timed to arrive 400 ps after the rising edge of the long pulse beam, and was focused on the centre of the disc. Self-emission is evident from the disc around the time of the short-pulse interaction, and the signal is comparable with that from the backlighter. The interface where backlighter absorption increases sharply is seen to expand rapidly (10^5 ms^{-1}).

Proton heating experiments

The short pulse laser is used to generate a beam of high energy protons. In order to demonstrate isochoric proton heating the target configuration, shown in Figure 5, is used. The section of the 10 μm wire which intercepts the proton beam was placed at the focus of the imaging system, and relayed onto the XRSC photocathode via careful optical alignment.

A result is shown in Figure 6. The rate of expansion is less than in Figure 4 due to the coupling inefficiencies of this method, but is still substantial. The temperature of the heated wire is estimated to be ~20 eV.

FIGURE 3

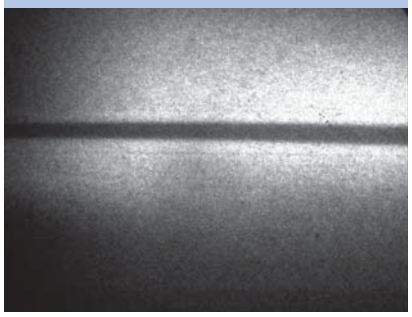

Streak camera record showing an undriven wire, backlit by a 1.5 ns pulse onto a orthogonal streak camera slit of 200 µm width.

FIGURE 4

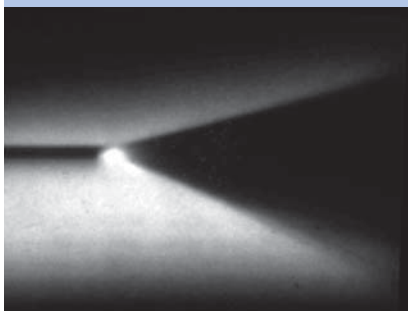

Streak camera record showing a gold disc, undriven for ~400 ps, then illuminated by a prompt laser pulse. Expansion occurs at a rate of ~10^5 ms^{-1}, backlighter energy was 75 J in 1.5 ns.

FIGURE 6

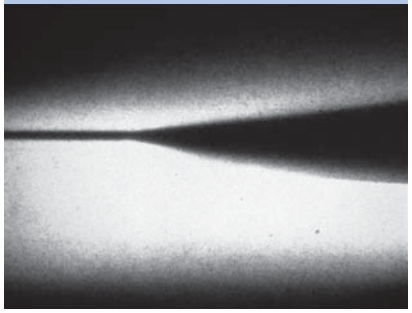

Streak camera record of proton heated experiment. Short pulse laser output was 57 J in 0.5 ps, I~5x10^{19} Wcm^{-2} at 1053 nm.

FIGURE 5

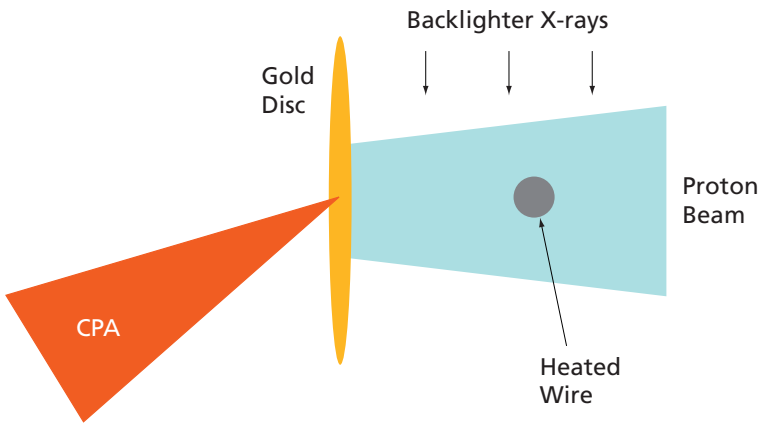

Layout of the proton heating experiments.

Conclusion

Although the spatial resolution achieved was slightly poorer than that predicted for an optimal system, it is comparable or better than has been achieved with equivalent systems.[3] The most likely reasons for this discrepancy concern the Bragg reflector in its non-ideal surface roughness and non-optimal position.

System alignment was performed using a telescope to centre the optical axis of the Bragg reflector to the target-detector axis, followed by adjustment of the reflector orientation to locate an image on the detector. Small errors in the positioning of the reflector would be corrected by adjustments to the orientation, with a less than optimal CTF.

The surface roughness of the reflector substrates was specified to be <1 Å r.m.s., with the coating surface finish not thought to degrade this specification significantly.[4] Exposure to high X-ray fluxes and debris will degrade the reflector surface finish. Without a detailed study of the amplitude, frequency and distribution of surface non-conformity, it is difficult to assess the effect on resolution.

Perturbations of lower frequency cause degradation of the reflector figure and hence the spatial resolution of images. The most likely cause of such perturbations is the deposition of plasma or molten debris onto the reflector. The use of a filter to protect the exposed reflector surface is impractical due to the proximity and fragility of the materials which can reasonably be used at these photon energies.

Acknowledgements:

The authors would like to thank the target fabrication and laser operations teams for their efforts during this experiment, David Lavender for designing key mechanical parts and Dr David Hoarty for his advice and assistance.

References

1 C. C. Smith; unpublished.
2 J. A. Koch et al, Appl. Opt. 37 1784 (1998)
3 Y. Aglitskiy et al, Phys. Plasmas 3 3438 (1996)
4 D. L. Windt et al, Appl. Opt. 33 2025 (1994)

AUTHOR PROFILE

Steven James can be contacted on e-mail: steven.james@awe.co.uk

Steven James

Steven joined AWE in 2000, having completed a MPhys in Physics with Astrophysics at the University of York. Steven initially worked in the Plasma Physics diagnostics team and in 2002 moved to the Plasma Physics Material Properties team. Steven carried out experiments on the HELEN laser and is currently involved of the commissioning of the Orion laser.

AUTHOR PROFILE

Edward Gumbrell can be contacted on e-mail: edward.gumbrell@awe.co.uk

Edward Gumbrell

Edward graduated with a 1st class honours BSc in Physics from Liverpool University. Edward went on to study for an MSc at Imperial College, were he stayed to do his PhD research on high-field non-linear optics and ultra-fast energy transport processes in plasmas. Edward received a Marie Curie Post-doctoral Fellowship and then joined Carl Zeiss, where he developed patented technologies for next-generation optical metrology systems.
Edward joined AWE's Plasma Physics Department in 2003 , where he has played a key role in experiments utilising laser-driven charged particle beams. Edward is also a honorary lecturer in the Quantum and Laser Optics Group at Imperial College.

Adhesive Testing

Adhesives are an important part of any system that requires the joining of materials. Understanding an adhesive's structure and behavioral properties when forces act on a material bond is vital when determining the life limiting factors of a system. At AWE numerous adhesives have been developed and incorporated into trials, manufacturing and production of assemblies. Before a batch of an adhesive is used it is subject to testing to ensure that all materials meet the required standard.

Adhesives can be defined as materials which when applied to the surfaces of materials can join them together and resist separation. An adhesive must do two things:[1]

- It must wet the surfaces, spreading completely over the surfaces so there is contact between molecules of the adhesive and molecules and atoms of the surface.
- It must harden to form a cohesively strong solid. The exception to this is pressure sensitive adhesives (PSAs).

There are six theories of adhesion:[1] physical adsorption involving formation of van der Waals forces across the adhesive-substrate interface, chemical bonding where covalent, ionic or hydrogen bonds are formed, diffusion where mobile polymers in contact can interdiffuse, electrostatic where an electrical double layer forms between two metals in contact, mechanical interlocking where the adhesive enters pores in the substrate before hardening and weak boundary layer theories based on removal of contaminants to produce clean surfaces which adhesives can attach to.

Depending on the materials involved several of these mechanisms can occur when an adhesive bond is made.

AWE's Adhesives group is responsible for understanding, testing and validating a wide range of adhesives. The Adhesives group also investigates process improvements to create better adhesives and test them. Current research into adhesive processes is aimed at methods of improving routine adhesive testing practices to make them more environmentally friendly, safer and to produce more consistent results. This is to ensure that a test is truly a test of the quality of the adhesive and not of the operators, testing conditions or quality of surface preparation.

When an adhesive is tested, a series of processes is carried out: surface pre-treatment, adhesive mixing, bonding, curing and mechanical testing. This article presents an overview of the process used for creating adhesive bonds, the testing of adhesives and improvements to these processes.

Surface Pre-Treatment

Surface pre-treatment of a material surface prior to bonding is fundamental to achieving satisfactory and consistent bond strength and durability. It enables formation of strong interactions between adhesives and surfaces in close contact with each other.

Surface pre-treatment improves adhesion by:

- Increasing surface free energy which makes the surface more easily wetted and therefore increases the substrate-adhesive contact area.
- Increasing surface roughness which enables mechanical interlocking between the adhesive and substrate.
- Removal of weak boundary layers. Weakly bound or structurally weak material on the surface, which the adhesive would bond to instead of the substrate, is removed.
- Enabling formation of chemical bonds. The surface is chemically activated or a functional molecule is used to form chemical bonds with both the substrate and the adhesive.

At AWE aluminium (Al) or stainless steel test pieces are typically used in adhesive tests. These substrates include rectangular plates for single lap shear and peel tests and cylinders for butt tensile tests.

Metallic substrates are often contaminated with grease and dirt from their production and storage. It is necessary to remove this contamination as it can form a weak boundary layer between the adhesive and substrate. Solvent cleaning is used before any subsequent chemical or abrasive surface pre-treatments. Solvent degreasing is carried out

FIGURE 1

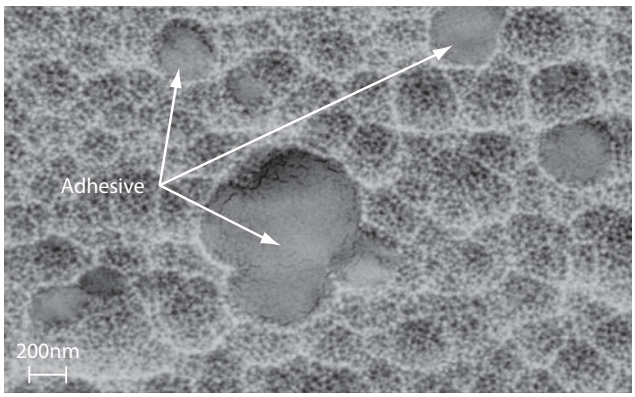

(a) SEM image of an Al surface after FPL etch process showing pores which the adhesive can fill.

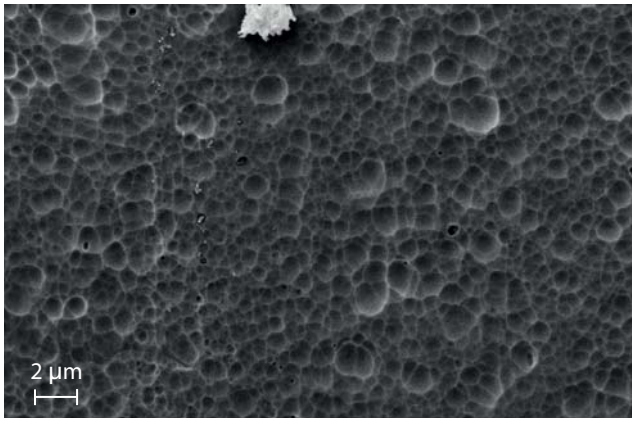

(b) SEM image of an Al surface after acidified ferric sulphate etch process.

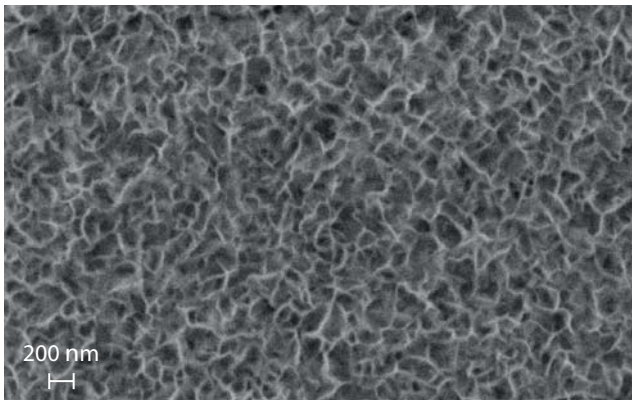

(c) SEM image of an Al surface after the hot water process at 80 °C for 15 minutes. The microstructure formed is more random without the round pores seen with the FPL and ferric sulphate etch methods.

on a small scale using methyl ethyl ketone (MEK) in a Soxhlet extractor. This cleaning method allows the specimens to be washed by a continuous cycle of hot, clean solvent.

After solvent cleaning the cleaned surface can be modified by mechanical (abrasive) or chemical means. These surface modifications offer varying levels of improvement to bond strength and durability compared to using untreated surfaces.

Abrasive pre-treatment is the mechanical roughening of a surface before bonding. This surface roughening improves mechanical interlocking and can remove weak boundary layers and loose material on the surface. Although it is a good method of increasing bond strength, surface roughening alone is unlikely to improve bond durability, as the metallic surface is not chemically modified in any way.

Chemical techniques, such as chromic acid etching are known to result in consistently high joint strengths and high bond durability. There are a number of variations of the chromic acid etch in use. The most common is the optimised Forest Products Laboratory Etch (FPL Etch).[2] This method involves immersion of the substrates in a 65 to 70 °C solution of sodium dichromate and sulphuric acid.

When etching Al, the original oxide layer surface is stripped away and a fresh oxide layer is formed, via reaction with the FPL Etch solution. The newly formed oxide layer is also attacked by the sulphuric acid resulting in a stabilisation effect, which controls the amount of oxide present on

the surface. The resulting surface has a fibrous morphology with an increased surface area. The result is good wettability and mechanical interlocking as well as a clean strong surface. Chromate also improves the corrosion resistance of aluminium, which increases bond durability. Figure 1a shows an SEM image of an Al surface after undergoing the FPL Etch.

Until 1999 AWE used a similar method for the pre-treatment of aluminium lap shear specimens used in the testing of epoxide and polyurethane adhesives. This was referred to as the 'AWE Etch' and was derived from Defence Standard 03-2.[3]

An increased awareness of the harmful effects of chromium prompted research into safer, more environmentally friendly alternatives to the chromic acid etch, while still achieving the same bond strengths. As a result of this research the chromic acid has been replaced at AWE by the acidified ferric sulphate etch. The decision to use this etch followed studies[4] that compared a wide range of surface pre-treatments and focussed in particular on pre-treatments offering improvements to safety. The work showed that ferric sulphate etching gave equivalent bond strengths to the AWE chromic acid etch. Figure 1b shows an SEM image of an Al surface after undergoing an acidified ferric sulphate process.

Whilst the ferric sulphate etch is safer than chromic acid, the process still contains potentially harmful chemicals and hot acid. Recent studies at AWE have investigated replacing the ferric sulphate etch with a pre-treatment using only hot water.[5] Such a treatment would make the process safer and more environmentally friendly than either of the acid etch systems. Figure 1c shows an SEM image of an Al surface after undergoing the hot water process at 80°C for 15 minutes.

Studies using both butt tensile and single lap shear tests, with both epoxide and polyurethane adhesives, have shown that in both types of test joint strengths can be achieved with hot water treated surfaces that are as high as those achieved with ferric sulphate etched surfaces. Although equally strong, the bonds are likely to be less durable than those with chemically etched surfaces as the metal is not chemically modified in a way that can improve corrosion resistance. This does not present a problem as the surface treatment is used in short timescale specification tests and not for long term bonding of components.

Mixing and Curing

Adhesives with two or more components need to be mixed prior to adhesive bonding. The separate parts of an adhesive are usually a functionalised polymer and a cross-linking agent. Accelerators and fillers may also be added. Often adhesives contain other additives such as emulsifiers, anti-oxidants and pigments, these additives are usually incorporated in one of the parts of the adhesive rather than as a separate ingredient added when the adhesive is used. Small amounts of adhesives, mixture weights of 20-30 g, are used for testing.

Traditionally the adhesive is mixed in a beaker by hand with a spatula but this method results in batch inconsistency. A recent study has investigated alternative mixing techniques with the aim of improving the reliability of results and reducing risks associated with mixing by hand. Box 1 explains some of the different approaches investigated.

FIGURE 2

(a) (b) (c)

Epoxide adhesive mixed by (a) hand for 30 sec, (b) dynamically at 6000 rpm for 3 minutes and (c) centrifugal at 2000 rpm for 3 minutes.

FIGURE 3

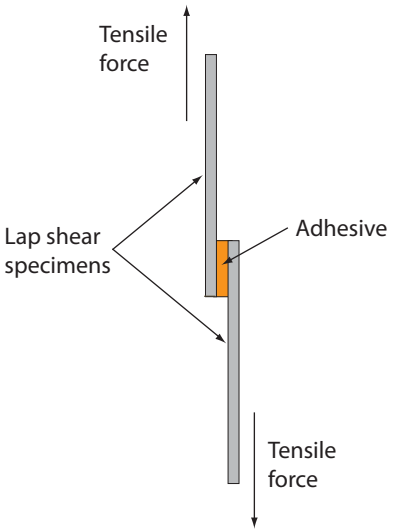

Single lap shear test setup.

Figure 2 shows the difference between hand, dynamically and centrifugally mixed samples of an epoxide adhesive. A blue dye was added to the adhesive to show how well mixed the adhesive was. In the hand mixed and dynamically mixed samples there are lighter and darker blue areas showing that the adhesive is not completely mixed, while in the centrifugally mixed sample the colour is homogenous. Also, air bubbles are clearly visible in the hand mixed and dynamically mixed samples and not in the centrifugally mixed sample. These results show that changing the mixing method can improve the quality of the adhesive.

Specified curing times for the adhesives used by AWE vary from 5 minutes, for the rapid-test for contact adhesives, to one week for a polyurethane adhesive. The majority are cured at ambient temperature and humidity. For polyurethane adhesives curing is ideally carried out under low humidity conditions. Otherwise reaction between the isocyanate ingredient of the polyurethane and atmospheric water could result in the formation of carbon dioxide gas causing the adhesive to foam and form significantly weaker adhesive bonds.

Testing

Adhesive testing at AWE is conducted to ASTM, ISO and British Standards developed to specific AWE material requirements. The adhesive is tested to prove that it will fulfil its intended purpose and to provide evidence of batch-to-batch consistency. The adhesive joint strength is usually the defining parameter of an adhesive. Sometimes other parameters such as elongation at break and flexibility are also defined. Adhesive tests used include the single lap shear, butt tensile and 180° peel tests, the set-ups for these are shown in Figures 3, 4 and 5. The polyurethane and epoxide adhesives routinely used at AWE are tested using single lap shear test and butt tensile tests, contact adhesives are tested using butt tensile joints, Pressure Sensitive Adhesives (PSAs) and tapes are tested using a 180° peel test.

The single lap shear tests use 100 x 25 x 1.5 mm metallic substrates. Adhesive is spread across the full width of the specimen to a distance of at least 13 mm in from the treated end. The specimens are then mounted in jigs in pairs with a 12.5 mm overlap to form single lap shear joints. A weight is placed on top of each set of six joints in the jig such that the adhesive cures under an applied load of 1 kg. The jig holds the specimens aligned so that the joint is straight and the overlap area is constant while the adhesive cures.

For butt tensile tests, 16 mm diameter cylinders are used. The cylinders have a screw thread inside the non-bonded face and therefore can be screwed into the fixtures of the tensile testing machine. Adhesive is spread on the flat ends of the cylinders which are then paired up and placed in a spring mounted jig. This arrangement applies a force of approximately 100 N to the samples whilst curing.

The bonded adhesive test samples are pulled using a mechanical test machine. Slow separation speeds of 1 or 1.3 mms^{-1} are usually used. Force-distance curves are produced which show the sample's ultimate strength and information relating to its flexibility.

The output from the single lap shear and butt tensile tests

FIGURE 4

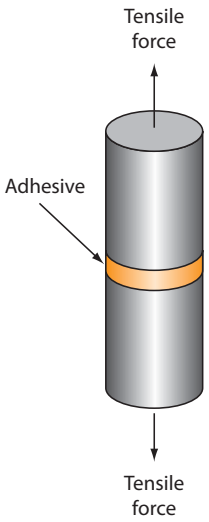

Butt tensile test setup using cylinders.

FIGURE 5

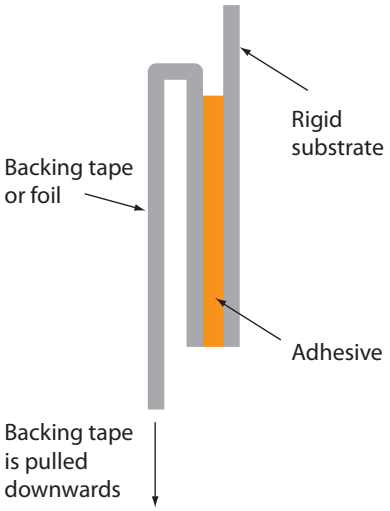

180° peel test setup.

provides data that can be described by a force-distance curve. Figure 6 shows some typical tensile test results. The curves tend to follow the same pattern; an initial linear slope, due to the force felt by the joint steadily increasing as the sample is pulled by the crosshead, then, if the adhesive is brittle, a sudden decrease to zero as the adhesive breaks. More 'rubbery' adhesives yield before complete failure, in these cases the non-linear part of the trace is more rounded, as the adhesive stretches with little increase in force when it yields.

For 180° peel tests tape or adhesive is applied to a metallic plate. For liquid pressure sensitive adhesives a template is used to control the width and thickness of the adhesive on the metal substrate. Strips of a flexible material, such as polyester film, are then laid on the adhesive or tape with a 'tail' off the end which is at least as long as the bonded area.

During peel tests two surfaces are peeled apart, the adhesive between them stretches and forms columns of adhesive. The columns are elongated in the direction of the peel force. In addition the curvature of the material pieces being peeled apart causes a compressive zone and induces shear stresses. Peel strengths are typically calculated as force averaged over a peel distance between two limits as shown in Figure 7.

A number of peel test configurations are available, the simplest being T-peel tests for flexible to flexible substrate bonding or 90° and 180° peels for flexible to rigid assemblies. These tests are standard and are described by the appropriate ASTM and British standards.[6-9]

At AWE tapes and pressure sensitive adhesives are tested using a so-called 'flexible to rigid' assembly, where a flexible material, usually a polyester film, is bonded to a rectangular steel plate, this is shown in Figure 5. The plate and tape are clamped into the grips of the mechanical test frame and pulled apart at a constant separation speed.

The result of such a peel test is produces a force-extension curve, indicating the force measured by the load cell against the distance the sample is peeled i.e. the distance the grips are pulled apart.

The force-extension plots usually have a sharp peak at the beginning, where the peel starts, followed by a levelling off as the force becomes relatively constant

FIGURE 6

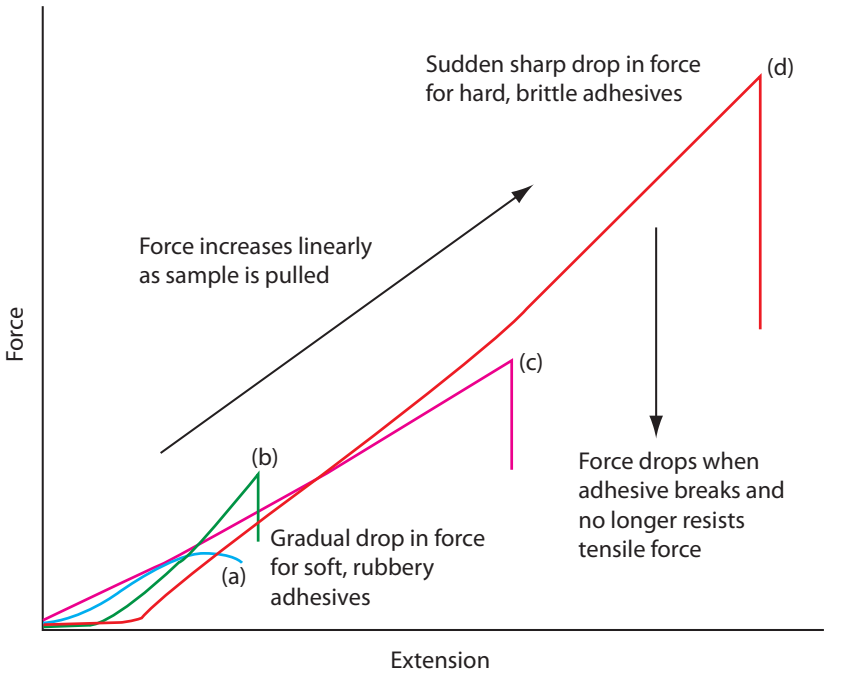

Force-Extension curve for tensile test of adhesively bonded single lap shear or butt joint for (a) a weak rubbery silicone adhesive, (b) a weak brittle contact adhesive, (c) a medium strength epoxide or polyurethane adhesive and (d) a strong epoxide or polyurethane adhesive.

through the rest of the peel. The average force is determined over 100 mm of peel, the first 25 mm of peel is disregarded to ensure that the early anomalous peak is excluded. In standard test reports the average, maximum and minimum peeling force for each specimen are recorded. This methodology allows easy comparison between the average peel strengths of specimens and also highlights any unusually strong or weak points in the peel test.

Samples are tested in sufficiently large quantities to give a high level of certainty to the results. For one part PSAs or adhesive tapes, a single set of six peel tests is usually considered sufficient. For the multi-component epoxides and polyurethanes at least two operators each prepare a set of six butt joints or single lap shear joints.

Summary

The adhesive research and development programmes presented here aim to improve specification testing methods making the results more consistent and improving the health, safety and environmental effects of the processes.

Future work is aimed at investigating new adhesives, advanced testing methods and ways of controlling adhesive strength to aid disassembly and recycling.

Acknowledgements

The authors wish to thank past and present members of the Adhesives, Pottings and Encapsulants Team who have contributed to the work presented here and the Materials Department at the University of Loughborough.

FIGURE 7

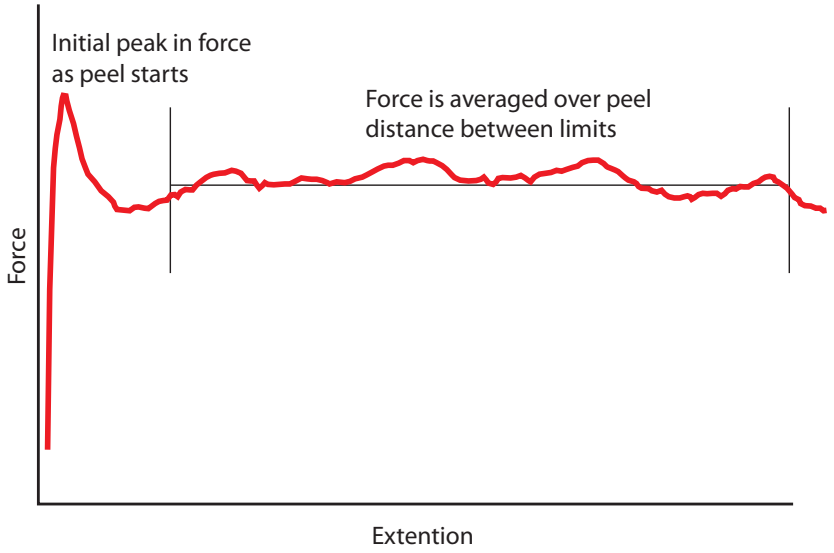

Force-Extension curve for peel test.

References

1 Comyn J, Adhesion Science, RSC Paperbacks (1997)
2 Wegman R F, Surface Preparation Techniques for Adhesive Bonding, Chapter 2, Noyes Publications, USA, (1989)
3 Ministry of Defence Defence Standard 03-2, Cleaning and Preparation of Metal Surfaces, Process O – Oxide Removal Treatment for Aluminium and Aluminium Alloys – Bath Treatment, Issue 1 (1970)
4 Williamson J R, Surface Pretreatment of an Aluminium Alloy using Alternatives to Chromic Acid Etching, MSc Thesis, University of Surrey, (2000)
5 Ford K M, Etheridge J A, Boyd M, Environmentally Friendly Surface Pre-Treatments, presented at Euradh 2006, 2006, St. Catherine's College, Oxford, UK
6 BS EN ISO 11339:2005, Adhesives – T-peel test for flexible to flexible bonded assemblies, British Standards Institute, (2005)
7 BS 5350: Part C10: 1991, ISO 8510-1:1990, Methods of test for Adhesives Part C10. 90° peel test for a flexible-to-rigid assembly, British Standards Institute, (1991)
8 ASTM D 6862-04, Standard Test Method for 90 Degree Peel Resistance of Adhesives, ASTM International, (2004)
9 BS 5350: Part C11: 1991, ISO 8510-1:1990, Methods of test for Adhesives Part C11. 180° peel test for a flexible-to-rigid assembly, British Standards Institute, (1991)

BOX 1

Mixing study

For many years AWE has used hand mixing to mix two part adhesives such as epoxide adhesives; a very user dependent process. Effective mixing can be difficult if the adhesive components are of similar colour, due to the difficulty of judging when the adhesive mixture has reached homogeneity.

This variability has brought about a requirement to improve methods of testing so that the quality of mixing is not a contributing factor to the adhesive joint strength. Consequently, work is ongoing at AWE to identify 'foolproof' methods of mixing certain adhesives, i.e. to determine standard testing methods that will produce reliable results every time, with any user. Two epoxide adhesives, one filled and one unfilled, have been used in a series of tests to establish the best mechanical mixing technique for each type of adhesive.

An epoxide adhesive is formed when a polymeric amine chemically reacts with an epoxide resin, as shown in Figure 8a. To achieve a fully cured adhesive both parts of the adhesive must be well mixed so the maximum possible number of chemical bonds are made between the amine and epoxide resin polymers. The fully cured adhesive contains epoxide linkages which link the polymers together as in Figure 8b.

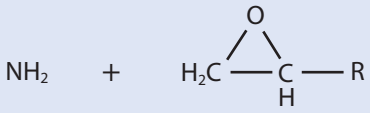

Figure 8a: Amine and Epoxide Resin

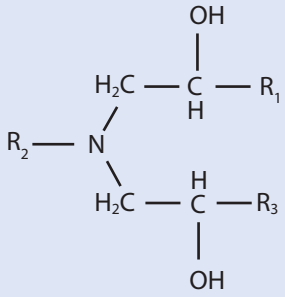

Figure 8b: Epoxide Adhesive

Mechanical mixing

Both dynamic mixing, (using a rotary mixer) and centrifugal mixing (using a Speed Mixer™, which spins the samples in a sealed pot), have been investigated as potential improvements for the mixing technique to use in specification testing. A series of tests were carried out mixing epoxide adhesives under various conditions to determine the effect of mixing on the adhesives' appearance, joint strength, surface tack, hardness, chemical properties and thermal properties.

BOX 1

Comparison of mixing techniques

Hand mixing resulted in strong adhesive joints for both the filled and unfilled adhesive but showed a significant variation in results. Hand mixed samples had high hardness and low tack values, indicating that the samples had cured fully. Visual examination suggested that samples mixed for only 30 seconds showed signs of inhomogeneity whilst those mixed for longer were homogenous.

Dynamic mixing gave samples with bond strengths similar to those for hand mixed samples, although there was a significant variation observed in results using lower mixing speeds. The variation was less for samples mixed with the higher mixing speed. These samples had high hardness and low tack values, again indicating that the samples had cured fully. Visual examination suggested that samples mixed at slow speeds for only 30 seconds showed signs of inhomogeneity, whilst those mixed for longer did not. Dynamically mixed samples also foamed a lot when mixed. Problems were encountered incorporating the filler as the powder tended to become airborne as soon as the mixing started, causing a hazard for the user.

Centrifugal mixing, unlike hand mixing and dynamic mixing, did not cause the adhesives to foam. Samples that were centrifugally mixed at 1000 rpm gave poor results and were tacky, weak and showed signs of inhomogeneity. In general, samples centrifugally mixed at higher speeds performed better as adhesives. However, samples of the filled version of adhesive warmed up during mixing; this internal heating may effect the resultant pot life and the curing of the adhesive.

Visual examination of the centrifugally mixed samples indicated that all showed good homogeneity. Chemical and thermal analysis showed no major difference between adhesive samples mixed using the three techniques.

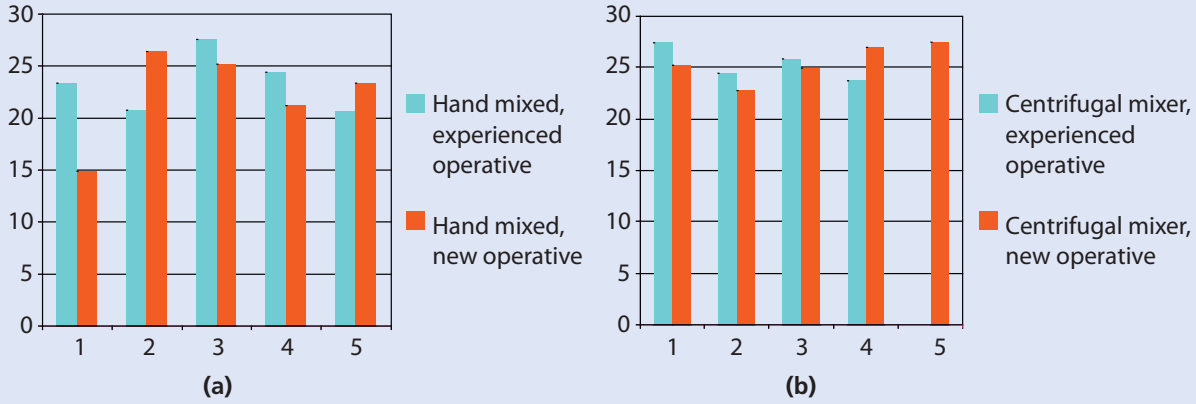

Figure 9: Tensile strength (MPa) of different adhesive batches undergoing butt joint testing when mixed by experience and new operatives using a) hand mixing b) centrifugal mixing.

These results have shown that, of the three methods investigated, centrifugal mixing at 2000 rpm for five minutes was the best method. It was found to be safe (not a dust hazard), to give good sample homogeneity, to provide high strength bonds with low variation between users and did not foam the sample. Figure 9 shows that while changing the mixing techniques did not have a dramatic effect on the average joint strengths measured, it removed the user variation in the process.

Consequently, the recommendation of the mixing study was that it would be advantageous to incorporate these conditions into the specification tests to reduce variability, make mixing easier and prevent loss of powdery fillers in the mixing process.

AUTHOR PROFILE

Kathryn Ford can be contacted on e-mail: kathryn.ford@awe.co.uk

Kathryn Ford

Kathryn graduated from the University of York in 2000 with an MChem and joined the graduate scheme at AWE. She is now part of the Adhesives Pottings and Encapsulants Team where her work focuses on potting materials, surface pre-treatments and chemical analysis.
In 2007 Kathryn completed a part time MSc in Materials for Industry at Loughborough University with a dissertation on methods of reducing the strength of adhesive bonds to aid their disassembly.
Kathryn is a member of the Institute of Materials and the Royal Society of Chemistry and is a Chartered Chemist.

AUTHOR PROFILE

Sarra Boyd can be contacted on e-mail: sarra.boyd@awe.co.uk

Sarra Boyd

Sarra graduated from Loughborough University with a BSc (1st class honours) degree in Chemistry and Forensic Analysis. In April 2006 Sarra joined the AWE Organic Materials Team within the Materials Development Group. Since starting work at AWE, Sarra has worked on adhesives and filled systems. She is currently focusing on filled systems and paints. Sarra is a chartered chemist and a member of the Royal Society of Chemistry and Institute of Materials. Sarra is currently studying part time for an MSc in Materials for Industry at Loughborough University.

AUTHOR PROFILE

Julie Etheridge can be contacted on e-mail: julie.etheridge@awe.co.uk

Julie Etheridge

Julie joined AWE as a trainee technician before joining the Electroceramics section where she assisted in research into a wide range of ceramic materials including Piezoelectrics and Cermets. Whilst in this section Julie obtained her degree in Applied Chemistry at Kingston University, London.
In 2001 she joined the Organic Materials Team specialising in adhesives. Julie completed a Post Graduate Diploma in Material Science at the University of Surrey and is now Team Leader of the Adhesives, Pottings and Encapsulants Team.
Julie has presented her work at international conferences including 'Polymers in Defence and Aerospace Applications' in 2007 and the 'International Conference of Structural Adhesives in Engineering' in 2004.

High Performance Computing for National Security

Since the Comprehensive Test Ban Treaty (CTBT) came into effect in 1996, a CTBT certification method has been in place to underwrite the safety and performance of the United Kingdom's nuclear deterrent. This method is based on physics and engineering calculations, supported by data from plasma physics, engineering and hydrodynamics experiments, the nuclear test database and from ab initio computational modelling of material properties. High Performance Computing (HPC) is a fundamental requirement for the CTBT certification method as the calculations are too large and complex to be performed by any other means. The areas of certification enabled by HPC are illustrated in Box 1.

To provide the required levels of assurance, an increasing level of sophistication in modelling is required with a defined certification programme placing a growing demand for calculations in future years. The vision is to gain the ability to certify using predictive, fully-resolved 3D full-physics and engineering calculations.

The requirement for calculations is met by the provision of an integrated HPC solution, involving supercomputer platforms operated by experienced staff and with parallel algorithms and visualisation tools developed by dedicated computer scientists.

BOX 1

Use of HPC in CTBT certification

High Performance Computing is essential for all aspects of the CTBT certification programme. It is required for physics and engineering code calculations, as well as for supporting material theory calculations needed to provide input data for the codes. Experiments on the Orion laser facility and from the Hydrodynamics area provide required data and integrated trials, but even these experiments would be impossible without HPC for trial design and diagnosis.

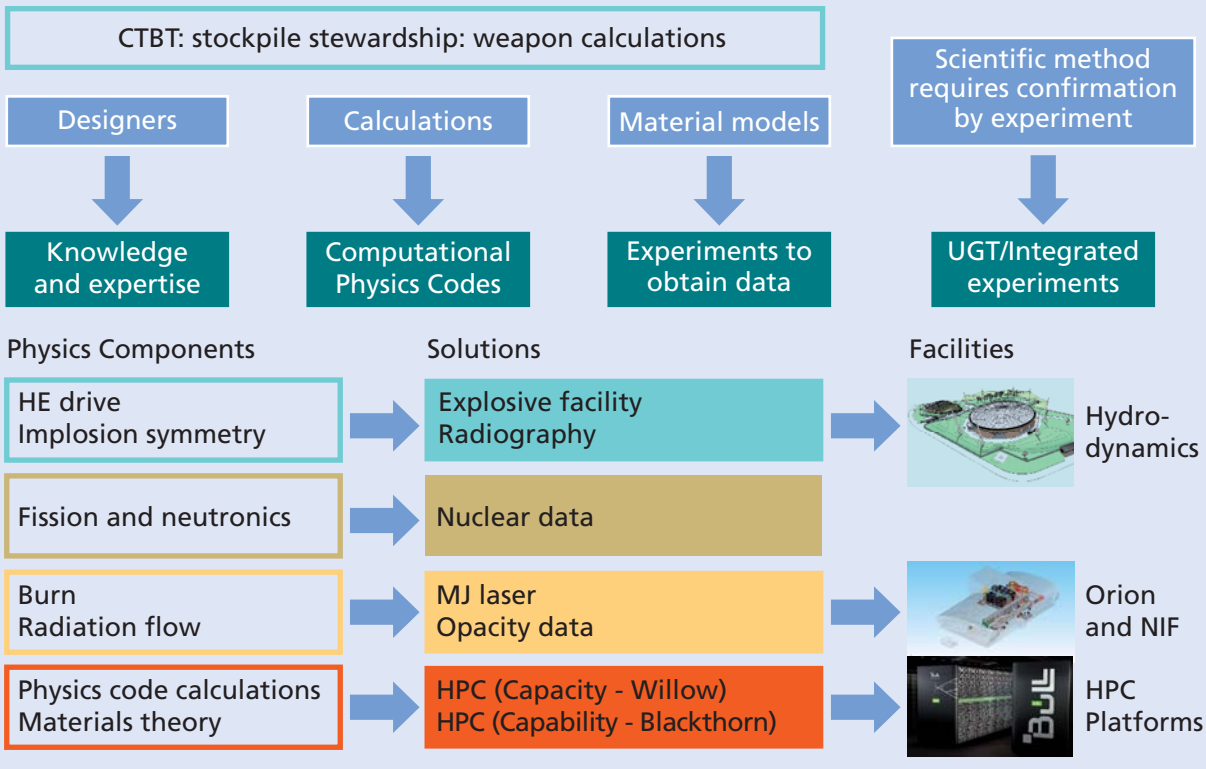

High Performance Computing Group

The HPC group acquires and operates supercomputer platforms to meet AWE's scientific and engineering requirements. The workload a supercomputer can process is usually defined in trillions of floating-point operations per second ($Tflops^{-1}$). The workload is a combination of capacity and capability calculations.

Capacity refers to the bulk workload, small calculations where throughput is the key concern. Capability refers to the largest calculations, where a large fraction of a supercomputer is dedicated to a single calculation. In reality the calculation requirements are much more complex e.g. different amounts of memory per operation are required for different applications.

In essence a supercomputer is a collection of processors connected together with a very High Speed Network (HSN) for inter-processor communications, a high speed disk network for Input/Output (I/O) and with the operating system minimised in order to reduce interference. In practice the machines have very complicated specifications requiring detailed assessment involving theoretically modelled and actual benchmarked performance.

The best overall solution requires computer scientists to work with the supply chain to optimise the hardware solutions deployed. The application of these processes to the current supercomputer uplift at AWE is the focus of this article.

The HPC group is also responsible for the provision of other dedicated HPC hardware systems to support the supercomputers. This includes a persistent file-store and tape archive, a visualisation cluster which provides the capability to interactively visualise the calculated data sets from the supercomputers and a computer optimised for engineering applications with high memory. All of these systems must be updated in parallel with improvements to the core calculation capability.

It was recognised that supercomputing requirements at AWE will continue to evolve, as will available technologies. The strategy is to implement incremental changes to the HPC capability at AWE through short term projects. A team was put in place to deliver the first of these projects, which aimed to meet AWE's engineering and scientific needs up to 2012.

The project team determined that 210 $Tflops^{-1}$ of computing was required, of which one platform must support at least 100 $Tflops^{-1}$ of capability class calculations. A solution of one 145 $Tflops^{-1}$ capability computer, named Blackthorn, and twin 36 $Tflops^{-1}$ capacity computers named Willow was selected. The three computers together provided an almost perfect match to the 210 $Tflops^{-1}$ requirement.

Before the design specification and invitation to tender could be issued, appropriate benchmarking codes needed to be selected. These benchmark codes are algorithmically representative of the code calculations required by AWE and are used to demonstrate the relative performance of the computers offered by each vendor.

"Although the benchmarks were to be run on up to 1,024 processor cores, AWE was interested in how the jobs would scale on up to one third of the total machine size (~4,500 processor cores)."

BOX 2

Performance Modelling

1. Application

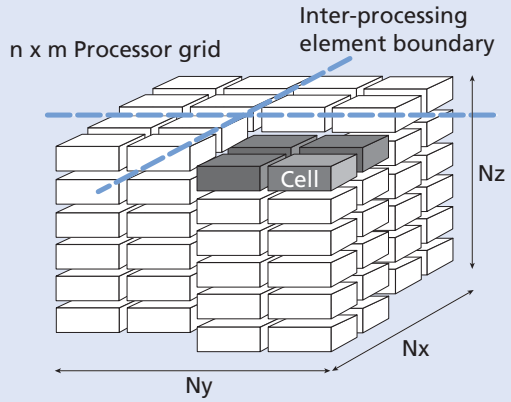

2. Model parameters and equations

W_g = Compute time per cell
NL_c = Network latency (core to core)
NL_p = Network latency (proc to proc)
NL_n = Network latency (node to node)
NB_c = Network bandwidth (core to core)
.
.
.

$$W = W_g \times \text{Tile depth} \times \frac{Nx}{n} \times \frac{Ny}{m} \times \ldots$$

$$StartP_{(n,m)} = max\,(startP_{(n-1,j)} + W_{(n-1,j)} \ldots.$$

3. Benchmarks

W_g = 1.67 x 10^{-7} seconds
NL_c = 0.655 micro seconds
NB_c = 2.7 Gbytes per second

W = (1.67 x 10^{-7}) x 1 x 2 x 2 x .....

4. Predictions and validation

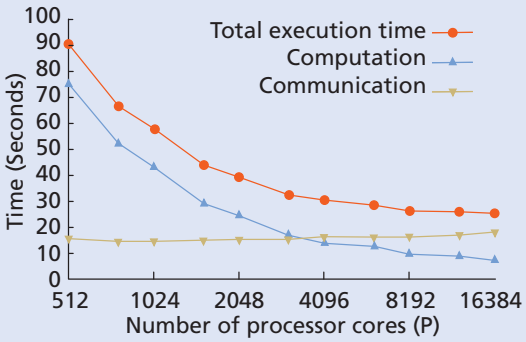

(1) Application is characterised in terms of its algorithmic properties and their relationship with the underlying computing architecture.

(2) Low-level model parameters are identified (e.g. computation, communications, I/O) and combined through model equations (that capture the algorithmic behaviour).

(3) The low-level parameters are benchmarked on the candidate target architectures, these results are used to prime the models.

(4) Models can be executed to mimic run-time behaviour, parameters and model equations can be adjusted to investigate alternatives; the cycle is closed by the models being used to provide recommendations for how to re-engineer the applications.

FIGURE 1

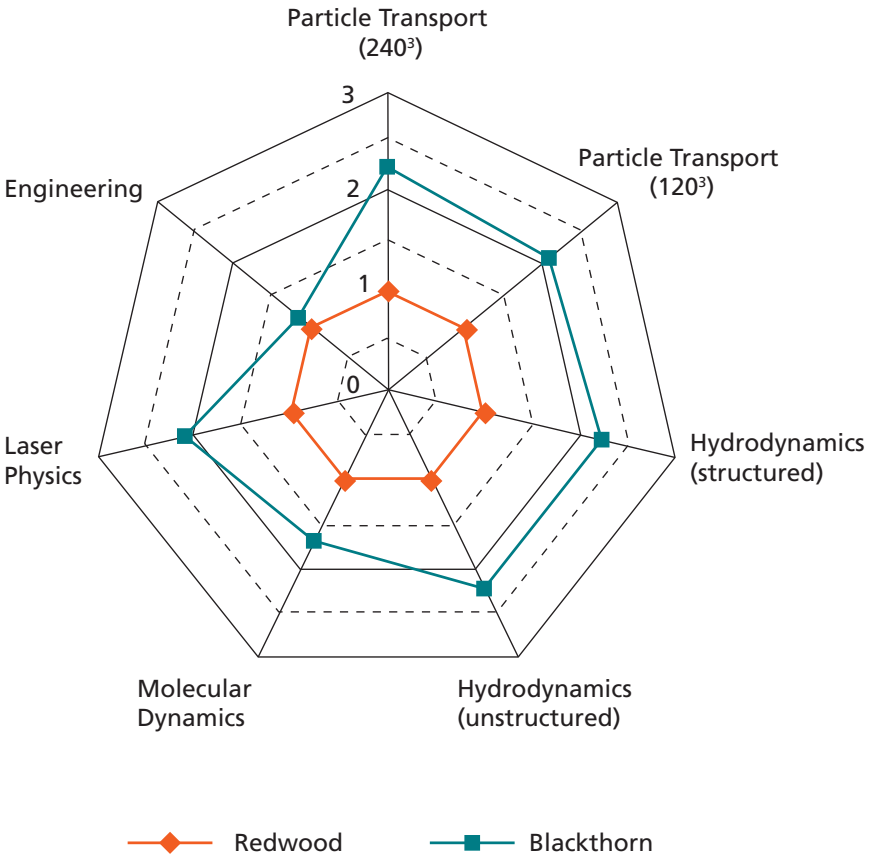

Performance comparison of Redwood and Blackthorn showing relative speed-up at 1024 processor cores. Blackthorn also has more cores than Redwood, allowing more jobs to be run simultaneously.

Benchmarking HPC Solutions

The High Performance Systems Group at the University of Warwick has been working with AWE since 2007 on the benchmarking and performance modelling of supercomputing systems and their applications. Comparing machines through computer benchmarking typically involves exercising some aspect of the proposed system (micro-benchmarking) or running applications of interest with small sub-samples of test data (application-benchmarking). AWE used both types of benchmarking when assessing options for Willow and Blackthorn.

Selecting the right combination of micro- and application-benchmarks and capacity and capability problems is crucial when evaluating machines. In the case of Willow it was anticipated that the maximum job size would be relatively small (~256 processor cores) which could simply be run on many existing test machines. It

"The models are highly complex, based either on mathematical methods or on application and architecture simulation."

was also important to understand how potential schedulers would pack jobs to ensure effective throughput when numerous jobs were executing concurrently. The benchmarking for Blackthorn represented a different proposition. Although the benchmarks were to be run on up to 1,024 processor cores, AWE was interested in how the jobs would scale on up to one third of the total machine size (~4,500 processor cores).

FIGURE 2

	Blackthorn runtime performance relative to Redwood	
Mesh Size	1024 Cores	1/3 of Total Machine
500^3	1.91x	2.90x
1000^3	1.93x	3.12x

Performance modelling projections for future AWE particle transport workloads.

This highlighted a traditional problem with computer benchmarking for capability-class machines; benchmarks are invariably run on a smaller number of processors than the full application will occupy. This problem is compounded when vendors benchmark on one machine and then tender on another e.g. with higher processor clock-speed, more processor cores or improved network latency. It is important to understand these issues and be able to verify vendor predictions.

FIGURE 3

The Bull capability computer, Blackthorn, installed at AWE.

FIGURE 4

Platform	Redwood (Cray XT3)	Willow (each) (Bull Inca Blade)	Blackthorn (Bull Inca Blade)
Peak Performance ($Tflops^{-1}$)	42	36	145
Cores	7888	3724	12960
Processor (GHz)	AMD Opteron Dual Core (2.60)	Intel Nehalem Quad Core (2.4)	Intel Westmere Hex Core (2.80)
Memory (TeraBytes)	34	11.2	52
Disk (TeraBytes)	100	160	460
Interconnect	Cray Torus	Quad Data Rate Infiniband fully non blocking	
Cabinets	45	14	42
Power (kW)	850	<150	<500

A comparison between AWE's new supercomputers and the Cray which was purchased in 2006.

Performance modelling (Box 2) takes computer benchmarking a step further. A performance model allows a benchmark to be accurately adjusted to answer 'what if' questions:

- What if the tendered processor is faster?
- What if the network latency offered by a vendor is improved?
- What if the data are increased by a factor of ten on the proposed computer?
- What if a thousand more processor cores are purchased?

The models are highly complex, based either on mathematical methods or on application and architecture simulation. Much of the work on performance modelling has come out of the Performance and Analysis Laboratory (PAL) at Los Alamos National Laboratories (LANL). Their work gained considerable notoriety in the 1990s through two classic papers[1,2] which showcased the use of performance modelling. One of these, the ASCI Q[1] case, highlighted LANL confidence in performance modelling and its ability to predict what a machine should do. When the computer initially performed more slowly than expected, their faith in the accuracy of their modelling made them re-examine the hardware and, as a result, saved millions of computing hours.

AWE's New Workhorses

The benchmarking exercises for Willow and Blackthorn were able to capture and compare the variety of machine offerings. The micro- and application-benchmarks exposed machine characteristics which were then

"The capacity to carry out large numbers of calculations will help to develop the understanding of the science of weapon safety and performance, and support the CTBT certification approach."

used to prime the performance models. The performance models were able to substantiate any claims made by the vendors concerning performance and scalability. They could also be used to investigate future scenarios that the machines might serve.

Figure 1 shows the predicted performance of Blackthorn against AWE's previous HPC machine, the 42 Tflops^{-1} Redwood supercomputer,[3] for a range of applications. There is a greater than two times speedup for particle transport and structured- and unstructured-hydrodynamics codes running on 1024 processor cores. Figure 1 also highlights where future performance hotspots are likely to manifest themselves.

Figure 2 is a table showing the projected performance of an AWE particle transport code for large (500^3) and extra large (1000^3) data-sets. These projections were derived from AWE's HPC simulator,[4] which mimics the behaviour of complex applications running on a range of supercomputing architectures.

The performance modelling showed that Blackthorn would deliver considerable speedups over the previous Redwood machine. With these new benchmarking and performance modelling capabilities, AWE is able to accurately assess the computing landscape in light of its own scientific requirements. It was determined that the best technical solution, at the best value for money, was offered by the French manufacturer Bull, as shown in Figure 3. The first pilot computer arrived at AWE in November 2009 with the Willow platforms arriving in January and March 2010. Following transfer of the users to Willow and the decommissioning of Redwood, the Blackthorn computer was installed in June 2010 and is now in full production.

> "The HPC group has established a strategy and is working with vendors and the University of Warwick to understand the technology roadmap over the next five to ten years."

Within the UK, only the national academic system Hector, in Edinburgh, has a greater capability than AWE. On total aggregate capacity AWE is second only to ECMWF's dual IBM systems. AWE's new supercomputers are not only more computationally powerful, but also much greener than the old Cray, delivering almost six times more compute performance per kilowatt and within a lower overall power budget. Figure 4 shows a comparison between the Cray supercomputer Redwood and the Bull Willow and Blackthorn machines.

Use of the Supercomputers

Installation of Willow and Blackthorn has already driven benefit to the users in moving forward the technical programmes at AWE. Blackthorn will enable several strands of research to access regimes previously unachievable on Redwood.

For material modelling this includes: quantum mechanical molecular dynamics calculations of Actinide hydriding,[5] modelling of shocks in crystals[6] and theoretical modelling of hot plasmas.[7] For plasma physics, the understanding of the short pulse (picosecond) Orion laser[8-10] would be impossible without the improved calculations of laser-plasma interactions that Blackthorn allows.

Blackthorn will also be used to design and diagnose plasma physics experiments on the US National Ignition Facility (NIF). 3D direct numerical simulation of mixing interfaces[11] will be possible with larger domains and increased resolution, improved 2D and 3D physics codes will also improve understanding of shock

waves in solids, liquids and gases under conditions of high temperature and pressure.[12,13]

The capacity to carry out large numbers of calculations will help to develop the understanding of the science of weapon safety and performance, and support the CTBT certification approach.

The Future

The project team has successfully implemented the computing uplift required to meet AWE's needs for the short term. The HPC group has established a strategy and is working with vendors and the University of Warwick to understand the technology roadmap over the next five to ten years.

The HPC group will continue to monitor the supercomputing need at AWE and investigate important issues such as power and resilience. AWE will continue to provide the manpower, facilities and integrated supercomputer solutions to meet the needs of the United Kingdom's nuclear deterrent.

Acknowledgements

The authors gratefully acknowledge the HPC staff who successfully delivered the HPC short term project described herein. In particular the operations staff who worked beyond the call of duty on system installation and configuration, Ken Atkinson who led the procurement from the HPC standpoint and the parallel technology team who ensured a more coherent approach to porting than ever seen before and unravelled MPI issues that threatened the success of the project.

Also key were Iain Miller and, from the University of Warwick High Performance Systems Group, Simon Hammond for the benchmarking and performance modelling effort. The work by the University of Warwick was conducted under grants CDK0660 (The Production of Predictive Performance Models for Future Computing Requirements) and CDK0724 (AWE Technical Outreach Programme) as well as by the TSB Knowledge Transfer Partnership grant number KTP006740.

The project would not have been successful without the support of the HPC User Requirements Group who defined the requirements and also the code custodians who provided benchmarks and ported codes very promptly. The project team, in particular Emma Kerswill who led the front-end development and Gordon Best who project managed the installation, were vital to the success of the endeavour, as was the support of the Investment Sanction and Scrutiny Department, most notably Bob Perridge.

Finally the computers could not have been delivered without some excellent work by the Supply Chain Management Department, principally represented by Ray McArdle and Fiona McGowan.

References

1 F. Petrini, D.J. Kerbyson, S. Pakin, The Case of the Missing Supercomputer Performance: Achieving Optimal Performance on the 8,192 Processors of ASCI Q, SC'03, November 15-21, 2003, Phoenix, Arizona, USA

2 D.J. Kerbyson, A. Hoisie, H.J. Wasserman, A Comparison between the Earth Simulator and the AlphaServer Systems using Predictive Application Performance Models, IEEE International Parallel and Distributed Processing Symposioum (IPDPS) 2003, April 22-26, 2003, Nice, France

3 D. Ball, B. Maskell, Supercomputing at AWE, Discovery 14, February 2007

4 S.D. Hammond, G.R. Mudalige, J.A. Smith, S.A. Jarvis, J.A. Herdman, A. Vadgama, WARPP – A Toolkit for Simulating High-Performance Parallel Scientific Codes, 2nd International ACM Conference on Simulation Tools (SimuTools '09), March 20-23, 2009, Rome, Italy

5 D. Price, Molecular Modelling at AWE, Discovery 9, August 2004,

6 S. Case, Numerical Modelling of Shock Response at the Mesoscale, Discovery 14, February 2007

7 B. Crowley, S. Davidson Radiative Opacity ,Discovery 2, Jan 2001

8 E. Gumbrell, A. Moore, J. Lazarus, High Energy Density Plasma, Discovery 15, September 2007

9 R. Evans, Computer Simulation of High Intensity Laser Plasma Interactions, Discovery 6, January 2003

10 J. Foster, Supersonic Plasma Jets, Discovery 2, January 2001

11 D. Youngs, D. Holder, Turbulent Mixing, Discovery 8, February 2004

12 A. Barlow, Arbitrary Lagragian Eulerian (ALE) Methods Solving Shock Hydrodynamic Problems, Discovery 4, January 2002

13 A. Barlow, AWE Sponsored Workshop on Numerical Methods for Multi-material Flows, Discovery 12, February 2006

AUTHOR PROFILE

Stephen Jarvis can be contacted on e-mail: stephen.jarvis@warwick.ac.uk

Stephen Jarvis

Stephen is Professor of High Performance Computing at the University of Warwick and co-organiser for one of the UK's High End Scientific Computing Training Centres. Stephen has authored more than 130 refereed publications.
Stephen is internationally recognised for his research on software/hardware performance analysis. In 2009 Stephen was awarded a prestigious Royal Society Industry Fellowship in support of his industry-focussed work on high-performance computing.

AUTHOR PROFILE

Andrew Randewich can be contacted on e-mail: andrew.randewich@awe.co.uk

Andrew Randewich

Andrew received his PhD in Plasma Physics and is currently Group Leader High Performance Computing at AWE. Prior to taking up this position in summer 2009, Andrew was sponsor for the Willow and Blackthorn project.
Andrew has also been the Physics Programme Manager, a team leader in the Computational Physics Group and technical assistant to the Director of Research and Applied Science. Andrew is a Fellow of the Institute of Physics.

AWE's Outreach, Major Events and Collaborative Activities

2010 was another successful year for AWE's events and outreach programme. There were a number of AWE hosted and sponsored events representing the diversity and richness of our scientific activity. A selected few are highlighted in this article.

Nitrocellulose capability

The 4th Nitrocellulose Workshop was held on 11 – 12 May 2010 at the Defence Academy, Shrivenham (Cranfield University). The event was jointly sponsored by AWE and the MOD's Defence Ordnance Safety Group (DOSG). DOSG have funded active research projects in the field of nitrocellulose, both at AWE and Cranfield University.

Professor Ian Wallace, Head of Cranfield Defence and Security, opened proceedings and said: "The relationship we have with AWE, through our Strategic Alliance, and DOSG is an important one, in support of a number of scientific and defence sector disciplines. We are very grateful for their kind support to this workshop."

AWE materials scientists Alan Macdonald, Ruth Knight and Dr. Paul Deacon gave interesting and thought-provoking presentations. They spoke about nitrocellulose analysis, ageing and the benefits of international collaboration. There were also presentations given by internationally renowned nitrocellulose experts and scientists, and the meeting was attended by almost 50 delegates.

The event opened up a number of research opportunities for the future, and AWE looks forward to the next conference which was tentatively proposed for May 2012 in Switzerland.

Thermal Analysis and Calorimetry conference

The Thermal Analysis and Calorimetry conference, a twice-yearly forum supported by the Royal Society of Chemistry's Thermal Methods Group, took place on 30 – 31 March 2010. A strong delegation from AWE's

NITROCELLULOSE WORKSHOP DELEGATES

materials science community attended along with a number of internationally renowned experts in thermal analysis techniques.

Dr. Mogon Patel, materials scientist and lead organiser, opened the event and said:

"This is clearly a very important event for AWE as many thermal analysis techniques are fundamental – particularly in materials life assessment, and new materials validation/verification activities."

The Cyril Keatch Young Scientist competition and poster session accompanied the event and Dr. Paul Nevitt, AWE materials scientist, won the Young Scientist Award.

Explosive event

The International Detonation Symposium; a major four-yearly meeting of the explosives community to discusses detonation chemistry and physics along with related phenomena, was held on 11 – 16 April 2010 in Coeur d'Alene, Idaho, US. Over 300 delegates, representing 87 institutions from 14 countries attended the event and over 180 presentations were given.

The symposium started as a meeting organised by the Office of Naval Research in 1951. Since then, it has rapidly developed into a major meeting and is sponsored by both the US Department of Defense and the Department of Energy. The symposium is now organised by various US labs, which in turn, supply an organising chair. The US Navy provides a continuing chair to oversee proceedings.

INTERNATIONAL DETONATION SYMPOSIUM DELEGATES

AWE AND DSTL DELEGATES

The meeting successfully provided AWE with the opportunity to promote its detonation research work to a truly wider community. It is hoped that AWE will benefit significantly through its links with a host of institutions, enhancing collaborations and introducing the potential for business development.

AWE-Dstl inter-organisational conference

Over 70 AWE and Dstl staff attended a conference on 29 – 30 July 2010 to improve working links between the two defence businesses. Jointly hosted, the conference was designed to foster, develop and strengthen the relationship between the two companies. The forum also provided a springboard for training and development opportunities for new starters.

William Harding, AWE, opened proceedings alongside Andrew Nelson from Dstl. The two-day agenda included interesting technical presentations by promising graduates from both companies. Subjects included nuclear weapons assessment, high performance computing, national nuclear security and ultra high temperature ceramics. Thought-provoking teamworking exercises focussed on common goals for the collaboration.

Also discussed was the Exchange of Early Career Scientist programme, a technical exchange activity in which Dstl graduates are invited to AWE and vice versa. Poster sessions highlighted previous successful examples of where the collaboration has yielded improved performance, benefit and value.

Guest speakers at the conference included Peter Sankey, Director Strategic Technologies MOD; Frances Saunders, CEO Dstl and Rob Fletcher, Director Commercial AWE. Gordon Arthur, Head of Technical Outreach AWE, said:

"The energy at this event was electric; it was great to see these communities from AWE and Dstl interact and start to build links. We emphasised that this group was seed-corn in both organisations to build future partnerships and interaction – which is clearly necessary in order to achieve the MOD's expectations of its technology communities."

Nuclear conference delivers engaging debate

AWE hosted the Project on Nuclear Issues (PONI) Conference on 21 – 22 September 2010. The aim of PONI is to bring together international nuclear experts to discuss the 'big nuclear issues' facing the world. Some 100 delegates from the UK and US attended the event.

PONI is a forum managed by the US-based Center for Strategic and International Studies, on behalf of and supported by the US Department of Energy and Department of Defense, and both the MOD and AWE participate.

AWE CEO, Robin McGill, opened the event and said:

"AWE is going through a period of strategic transformation to support our challenging nuclear warhead capability sustainment programme, and the positive impact of the recently recruited and early career personnel is clearly visible, PONI plays a major part in this endeavour."

Guest of honour, the Rt Hon. Lord David Owen, said:

"I am delighted to be invited here this evening and to be asked to contribute to the debate on some of the most complex global nuclear weapons issues, it is through the impact of this forum [PONI] that we are able to openly discuss ensuing issues and threats that affect us all."

AWE's PONI equivalent forum, the Nuclear Weapons Policy Discussion Programme (NWPDP) group, is formed of some 30 scientists and engineers from across AWE. The NWPDP provided a number of presentations to the PONI conference. Richard White, chair of the NWPDP, said:

"The breadth of the topics discussed allowed a fully inclusive debate and underlined the numerous complex issues that face us."

PROJECT ON NUCLEAR ISSUES CONFERENCE

MOD CHIEF SCIENTIFIC ADVISER COMMENDATION AWARD WINNERS

AWE wins MOD Award for scientific excellence

Work led by the Arms Control Verification Research team at AWE won a Ministry of Defence Chief Scientific Adviser Commendation Award, presented by Professor Mark Welland Chief Scientific Adviser to the MOD. The award is for the important role AWE played in a unique international collaboration with Norway on the science of verifying warhead dismantlement.

The project team undertook a highly successful exercise in June 2009, which looked at how a non-nuclear weapon state could verify the dismantlement of a device belonging to a nuclear power, without divulging sensitive information and compromising national security.

AWE leads on the technical work of this collaboration on behalf of the UK, working with the MOD, Foreign and Commonwealth Office, the Department for Energy and Climate Change and VERTIC (Verification Research, Training and Information Centre).

AWE project leader, Dr. David Chambers, recently presented the results of this exercise at the United Nations to the Non-Proliferation Treaty Review Conference on behalf of the UK.

Graeme Nicholson, AWE Director for Science & Technology, said:

"The initiative with Norway is the first time a nuclear weapon state has worked with a non-nuclear weapon state and this has placed both countries at the forefront of nuclear disarmament issues."

The work has generated much international interest, with a recent article in the journal Nature referring to it as 'a model of how to conduct international collaboration in this field'.

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 & 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

Printed on Cocoon 100 silk

Please place FSC logo here

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 21 • December 2010