# HIGH POWER OPERATIONAL EXPERIENCE AT ISIS

D J S Findlay, on behalf of all staff in the ISIS Accelerator and Target Divisions,
ISIS, Rutherford Appleton Laboratory, STFC, Oxfordshire OX11 0QX, UK

*Abstract*

Since 2008 ISIS has been running a second target station (TS-2) optimised for cold neutron production while continuing to run the original target station (TS-1) which began operating in 1984. The ISIS 800 MeV proton synchrotron cycling at 50 Hz produces a total beam power of 0.2 MW which is split between TS-1 and TS-2, 40 pps to TS-1 and 10 pps to TS-2. ISIS operations are described, including the first years of the new two-target-station operational régime.

## INTRODUCTION

Although J-PARC [1], PSI [2] and SNS [3] are spallation neutron sources with higher power proton beams, ISIS [4] may still be the world's most productive spallation neutron facility in terms of science delivery, and since 2008 there have been two operational target stations at ISIS. Currently each year on average ~750 experiments are carried out involving ~1500 visitors who make a total of ~4500 visits (on average, very roughly, each visitor visits ISIS three times a year). These numbers include ~100 experiments and ~300 visits for the ISIS muon facility on TS-1. This paper summarises the experience at ISIS of running two target stations — experience that may be of interest to other facilities considering a second target station.

The ISIS First Target Station (TS-1) began operations in 1984, and since then neutron scattering work carried out on TS-1 has resulted in a total of ~9000 scientific publications.

The ISIS Second Target Station (TS-2) began operations in 2008. TS-2 was built to facilitate neutron scattering measurements on soft matter, biological samples, and advanced materials, and the target station is optimised for the production of high peak fluxes of cold neutrons in a way that was not possible on TS-1.

The key elements of the accelerator system at ISIS are as follows: $H^-$ ion source at –35 kV, 665 keV 4-rod 202.5 MHz RFQ, 70 MeV 4-tank 202.5 MHz $H^-$ drift tube linac, 52 m diameter 800 MeV proton synchrotron with six 1.3–3.1 MHz fundamental RF ferrite-loaded cavities and four 2.6–6.2 MHz second harmonic ferrite-loaded cavities. The key elements of target systems are as follows: a tantalum-coated tungsten plate primary target with two water moderators, a ~100°K liquid methane moderator and a 20°K liquid hydrogen moderator for TS-1; and a tantalum-coated tungsten cylinder primary target with a coupled hydrogen / solid methane moderator and a decoupled solid methane moderator for TS-2. There are twenty-six beam line instruments on TS-1 (both neutron and muon instruments), and currently seven neutron beam line instruments on TS-2; an additional six or seven instruments for TS-2 are foreseen under Phase 2 of the overall TS-2 project. ISIS is also host to MICE [5], the Muon Ionisation Cooling Experiment, an important step on the road to a practical neutrino factory. A schematic layout of ISIS is shown as Figure 1.

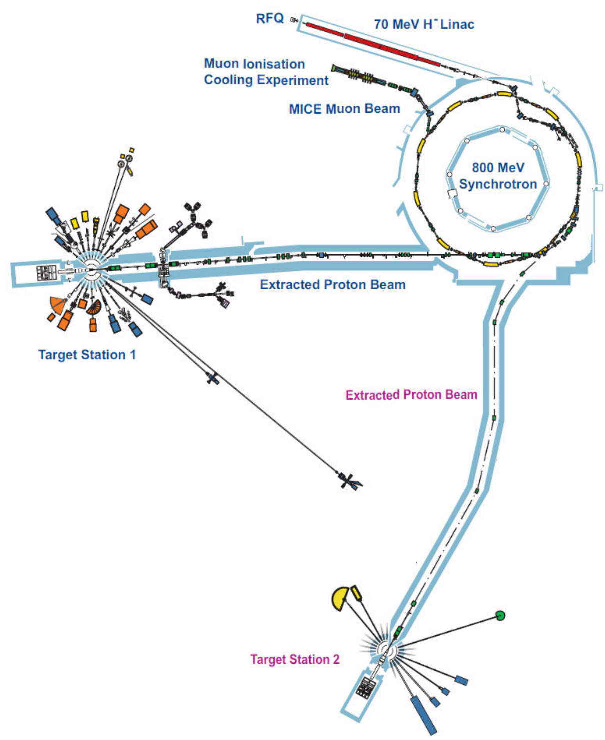


Figure 1: ISIS schematic layout.

## AVAILABILITIES

Figure 2 (upper half) shows availabilities of the ISIS accelerator and target system over the past twelve years. (For each user cycle, ISIS machine availabilities are defined as (total number of beam pulses actually delivered to target) ÷ (total number of beam pulses originally scheduled to be delivered to target); everything that prevents beam from being delivered to target, *e.g.* off-time for re-tuning, accelerator faults, target faults, plant faults, and RAL site electricity supply faults, counts towards machine non-availability.) The average of the set of availabilities is 86%, and the standard deviation is 8%; availability appears to have become gradually worse with time. However, until and including 2003 there used to be the opportunity to add "run-on" to cycles with poor availabilities — whereby several "bad" days could be replaced by additional "good" days added to the end of the cycle — but this opportunity no longer exists. Adding run-on could lead to noticeable improvements in availabilities, as several days in a cycle several tens of days long can represent a ~10% effect. In order to make a fair comparison of the availabilities over the twelve years covered in this paper the run-on effect has been removed (by adding the "bad" days to the duration of the cycle, and assuming that the beam was off during the bad days). The resultant data are also shown in Figure 2 (lower half). The availabilities can now be seen to be essentially constant between 1998 and 2006 inclusive, and then are slightly lower from 2007 onwards. So the upper half of Figure 2 probably gives an unhelpful impression; it seems likely that on average availabilities did not change until TS-2 plant and equipment was incorporated in ISIS in 2007, whereupon there was simply more plant and equipment to go wrong.

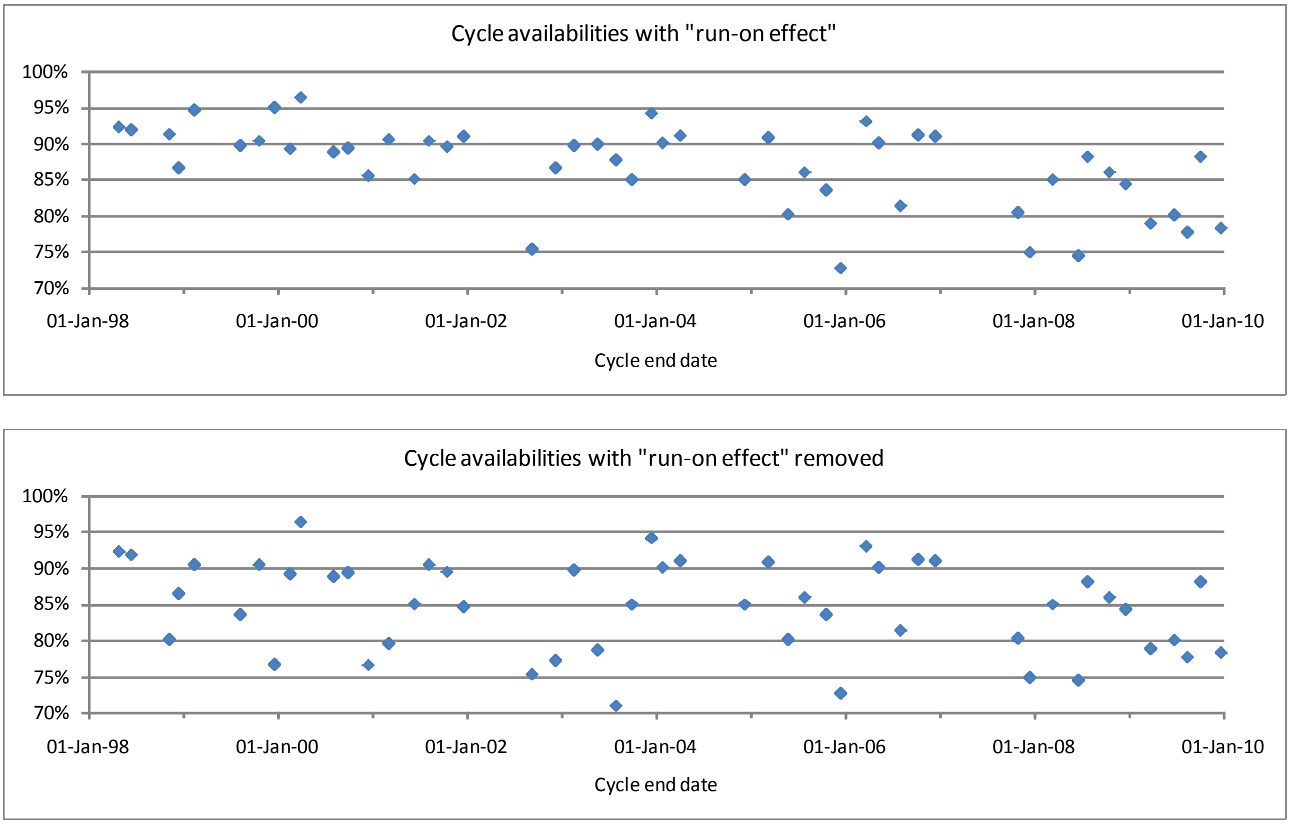


Figure 2: Availabilities of ISIS accelerator and target system since 1998, with and without run-on effect.

But how should the availabilities be best presented? Probably as in the upper half of Figure 2, *i.e.* without removal of the run-on effect, as that was how ISIS actually ran at the time, but also including the comment that the apparent worsening of performance is simply a consequence of how the machine was *scheduled*, *not* a consequence of how it *ran*. It is probably also true that it is the availabilities in the upper half of Figure 2 that should be compared with availabilities of other facilities, especially as many of them operate run-on regimes (even ILL, for example, in 2008 added five days of running to compensate for "minor pre-start-up testing woes" and "a cut in the mains electrical supply" [6]).

Figure 3 shows the frequency distribution of the cycle availabilities — but plotted in terms of "non-availability". Also shown in the figure is a fit by the log-normal distribution. The log-normal distribution is used to represent the multiplicative product of many independent random variables each of which is positive (in effect, the distribution is a sort of "multiplicative equivalent" of the central limit theorem for additive quantities), and the consistency of the fit and the data tends to support the idea that down-time is due not to any one particular cause but to a large number of causes.

For the time distribution of "off-times", see [7].

On ISIS machine down-time is divided into a great many plant and equipment categories — too many for immediate appreciation. But the periods of down-time highlighted in the operations reports can be attributed to twelve overall categories*, and an illustration of the change in distribution of the "headline" faults† with time is shown in Figure 4. Apart from the tall "Moderators" column and the less tall "Vacuum" column in 2008 and

* There is always a degree of arbitrariness about such representations. For example, should the failure of an RF window in a linac tank be categorised as an RF failure or as a vacuum failure?

† The "headline" faults are the faults emphasised in the operations report compiled after each cycle. There is a "chronic background" of faults which together with the headline faults make up the total number of faults.

2009 respectively (which, it is defensibly hoped, are simply anomalous), the number of tall columns does decrease with time, suggesting that the most significant issues are indeed being overcome.

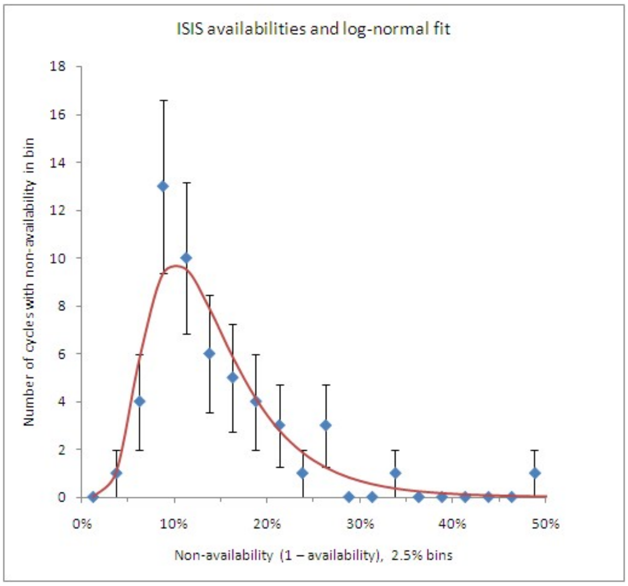


Figure 3: Frequency distribution of availabilities presented as "non-availabilities". The error bars have been taken as the usual square roots, and the $\chi^2$ of the fit per degree of freedom is 0.58.

## OPERATIONS

The ISIS running pattern is roughly as follows. Typically each year there are five sequences as follows: maintenance and/or shutdown period; ~7–10 days for run-up and machine physics; ~35-day user cycle (operating twenty-four hours a day, seven days a week); ~3-day machine physics period. Because of problems encountered during shutdown/maintenance periods or as equipment is brought back on again or because of problems encountered during user cycles, roughly one in every three machine physics periods has been lost.

Since TS-2 has become operational, the accelerators have been run up and the beam optimised while delivering beam to TS-1 alone. Machine physics has also been carried out while running to TS-1 only. Once the machine is running well to TS-1 it takes typically only an hour or two to set up the proton beam line to TS-2.

The ~140-metre-long proton beam transport line to TS-2 (EPB2) has proved to behave very reliably in practice. The beam spot on the target is approximately Gaussian with a diameter at one-hundredth maximum of 36 mm, and its position on the target is constrained by a 300-mm-long copper collimator with a tapered bore and with its downstream end 1.05 m in front of the target. The position of the beam spot on the target is monitored by intercepting small fractions of the halo around the beam and measuring the resultant temperature increases using thermocouples. More details of the EPB2 proton transport line are given in [8].

When TS-2 was fully incorporated into ISIS, the entire ISIS machine interlock system was replaced and upgraded. The extension of the scope of the interlock system to accommodate TS-2 resulted in a surprisingly large amount of extra complexity, largely as a result of the need to be able to run to both target stations simultaneously or to each of the two target stations on its own. Again, more details are given in [8].

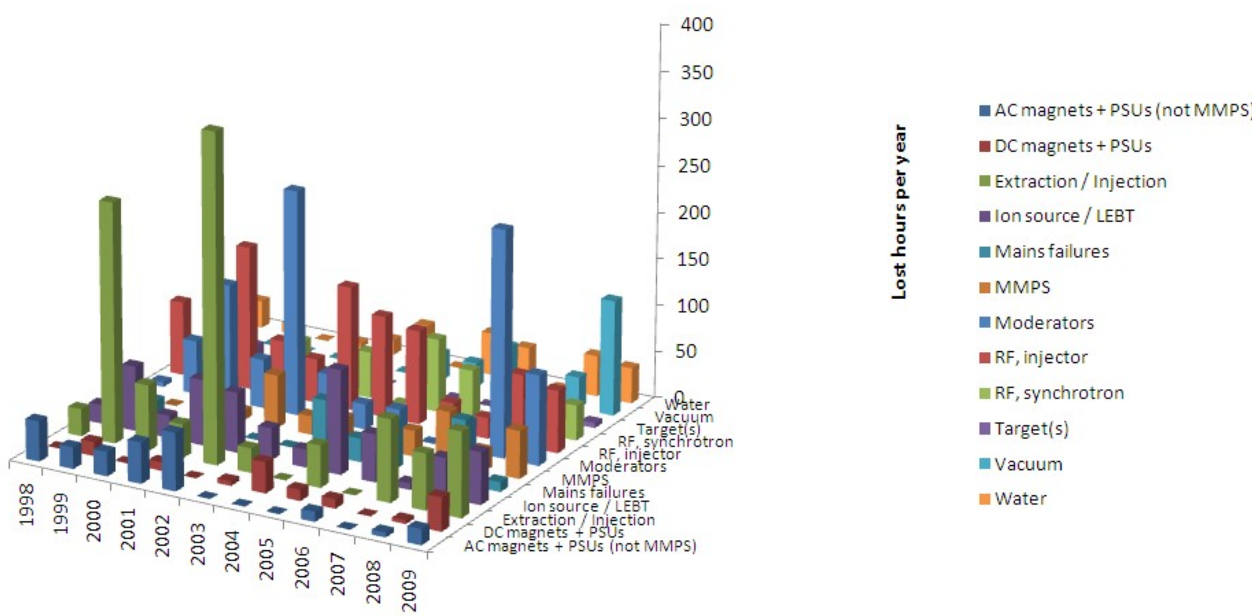


Figure 4: Distributions of lost hours in headline faults as a function of major fault category for the last twelve years.

## *Beam Losses*

At ISIS beam losses are measured and controlled through two systems, beam intensity monitors (resonant current transformers which measure the beam intensity to an accuracy of $\pm 3\times10^{10}$ protons per pulse), and long argon-filled coaxial ionisation chamber beam loss monitors (typically several metres long, and with a sensitivity (after amplification) of roughly $4\times10^{-5}$ $E^2$ femto-volt-seconds per proton of energy E lost (where E is in MeV)). The systems incorporate dedicated microprocessors to measure pulse-by-pulse values of beam loss and compare these values against preset tolerance levels which are different for different parts of the machine (*e.g.* the tolerance levels are highest near the collector straight in the synchrotron where beam losses are inevitably the highest). Both systems can issue trigger signals to the beam inhibit system: the beam inhibit system switches the beam off at the ion source for 1 second and then switches the beam back on again; but if another trigger signal to the inhibit system is issued from anywhere on the machine within the following 10 seconds then the beam inhibit system switches off the beam, inserts the beam stop in the low energy beam transport line before the RFQ, and calls for operator intervention to re-establish the beam. The tolerance levels were defined and set many years ago and have never been changed since. Of course, the beam loss triggering system can be overridden at low repetition rates to allow the beam to be set up.

After every user cycle (which typically lasts ~35 days), dose rates around the ISIS machine are measured. Typically, after a few days' cooling, the average dose rate around the synchrotron (excluding the collimation straight which is shielded locally) is 2 mSv/hour on contact and 0.2 mSv/hour at 0.5 metres (although there are large point-to-point variations — the standard deviations of the two sets of dose rates exceed the averages by ~20%). For cooling times of at least 1 day, the average activity around the synchrotron decays approximately as $\text{time}^{-0.25}$.

While in general the legal annual limit for radiation doses to people in the UK is 20 mSv, the formal investigation level at the Rutherford Appleton Laboratory (RAL) is 6 mSv, and a dose constraint of 3 mSv prevails at ISIS. For the ~300 ISIS staff who wear radiation badges annual collective radiation doses are typically ~50–100 mSv. Clearly it is very important to reduce beam losses as much as reasonably possible in order to minimise dose to maintenance workers, but it is also equally important for engineering designers of hardware to take into account radiation doses to people from the outset — dose rates *per se* are much less important than annual doses to people.

## *Two-Target-Station Experience*

As far as machine operations are concerned, the following ten points may be made as regards experience at ISIS of running two target stations. Some of the points are, of course, very obvious.

If beam is not to be "stolen" from an existing target station, the accelerator system has to be upgraded to produce more beam current before operations begin on the new target station. This was one of the reasons for upgrading the ISIS synchrotron RF systems to dual harmonic operation over the past few years [8].

Commissioning a new target station inevitably interferes with the continuing user programme on the existing target station. Although ISIS has independent timing pulse trains for TS-1 and TS-2, for safety reasons the beam to both target stations is immediately tripped if any primary interlock is breached in the machine areas, the neutron instruments, or the target systems on either TS-1 or TS-2, and so problems on TS-2 reduce the availability of TS-1. Of course, it is possible to change the running mode so that beam is delivered to TS-1 only instead of to TS-1 and TS-2, but the changeover is not instantaneous, and neutron users have to be given sufficient notice of the change.

ISIS can run to its target stations in three modes: to TS-1 alone, to TS-2 alone, and to both TS-1 and TS-2 simultaneously (and, incidentally, also to a low-Z full-energy beam dump in the synchrotron room). But, especially during commissioning of the new target station, it is important that the process for switching between modes be as quick as possible. On ISIS there is a comprehensive mode-switching system involving mechanically interlocked keys, electrically interlocked pulse train switching, magnet power supply isolation and magnet earthing, but initially the mode-switching process took a surprisingly long time.

In general, of course, the accelerators have to run harder to produce more beam for an additional target station; even though parts of the accelerator system may have been upgraded (*e.g.* the synchrotron RF systems on ISIS), other parts may not have been and may suffer accordingly. On ISIS a decrease in the lifetime of ion sources may be becoming perceptible, although previously average lifetimes have been ~30 days and it takes only ~3 hours to change an ion source.

The numbers of additional staff required to run an additional target station should not be underestimated. On ISIS staff numbers were increased to run TS-2 as well as TS-1 (including an increase from three to four in each of the five machine crew shifts), but the extra effort required to accommodate the new ancillary plant for TS-2 was more than expected. Partly this was because overheads of regulatory compliance had become more onerous by the time TS-2 was being commissioned.

When ISIS is running to TS-1 at 40 pps, the neutron users see an irregular series of neutron pulses from the TS-1 target and moderators. However, there have been few or no objections from the neutron users, and in fact some users have taken advantage of the 40 ms gap occurring every 100 ms to extend their data-taking to lower neutron energies than would otherwise be possible. Similarly, when ISIS is running to TS-1 at 40 pps but not running to TS-2, the beam loading in the accelerators is also irregular since the RF systems still run at 50 pps. But there seem to have been no obvious problems in the accelerator systems caused by the irregular beam loading.

The advent of a substantial new addition to a large facility can suddenly highlight the cumulative effect of gradual advances in technology over the period of time since the facility was originally built. Staff who construct and commission a facility are well placed to continue to operate the facility because of their intimate knowledge gained throughout the construction and commissioning phases. But it is desirable to keep technical knowledge up to date through appropriate training so that no surprises arise when the substantial addition to the facility is made.

Designing, constructing and commissioning a new target station tends to expose the relentless onward march of regulatory rigour. There is no reason to believe that in the 1980s the original ISIS target station ran unsafely in any way whatsoever, but over the quarter-century separating TS-1 and TS-2 the rigours of regulatory compliance have become ever more onerous. Of course, the increased regulatory compliance for the new target station encourages increased regulatory compliance for the existing target station.

The increased rigours of regulatory compliance have significant implications for staff training, and especially for the five machine crew shifts running the machine. At ISIS it has proved difficult, and continues to prove difficult, to deliver training to the shift crews, as on average only one-fifth of the total crew complement are present during normal hours on Monday–Friday. In addition, the crew shifts rotate every few days so that there is an inevitable mismatch between the irregular patterns of attendance of particular members of the crew and the regular patterns in which training sessions are usually most easily organised. In practice, training to meet increased regulatory requirements can represent a surprisingly heavy overhead.

Finally, of course, running a new target station increases the cost of electricity consumed. When running to TS-1 alone ISIS consumed ~10 MW of electricity, but running TS-2 consumes an additional 2–3 MW, mostly for the power supplies for the magnets in the proton beam transport line to TS-2[‡]. At present the ISIS electricity bill is ~12% of the total operating budget (including staff costs).

## SUMMARY

ISIS is the first spallation neutron source to run to two target stations. It proved challenging to construct and commission a second target station without degrading the service delivery to users of the first target station, and inevitably there was some worsening of the overall machine availability during the commissioning process. But, overall, the new TS-2 target station has proved to be a very successful addition to ISIS, and the neutron beam line instruments on TS-2 are delivering the excellent performances expected of them.

[‡] The advent of TS-2 on ISIS has coincided with the introduction of a UK-Government-inspired "carbon reduction commitment" scheme under which ISIS will suffer financially if it does not reduce its consumption of electricity. A procedure has been put in place to ramp down all the magnets in the proton beam lines to both target stations after the beam has been off for presettable time, since altogether the power supplies for the two proton beam lines consume ~3 MW. Unfortunately, the additional ramping down and up seems to be putting some of the older magnets at risk from earth-leakage problems, and the procedure may have to be abandoned for the proton beam line to TS-1.

## REFERENCES


[1] http://j-parc.jp/

[2] http://www.psi.ch/

[3] http://neutrons.ornl.gov/facilities/SNS/

[4] http://www.isis.stfc.ac.uk/

[5] http://www.mice.iit.edu/

[6] ILL Annual Report 2008, page 102.

[7] Time when the beam is off is made up of "trips" and "inhibits". Trips are events when the beam is switched off automatically but can only be switched on again by the machine crew. Inhibits are events when the beam is automatically switched off and on again after ~1 second, typically in response to an anomalous beam loss event. Trips are recorded in four categories: longer than one second, and longer than 1, 3 and 6 hours. Twelve-year means ± standard deviations are: inhibits, 74±45 per day; trips >1 s, 29±17 per day; trips >1 hour, 0.45±0.15 per day; trips >3 hours, 0.21±0.09 per day; trips >6 hours, 0.10±0.07 per day.

[8] J W G Thomason, Accelerator Development for Operating Two Target Stations at ISIS, ICANS-XIX, Grindelwald, 8–12 March 2010.