DR JOHN TAYLOR OBE FREng FRS
DIRECTOR GENERAL OF RESEARCH COUNCILS, OFFICE OF SCIENCE AND TECHNOLOGY

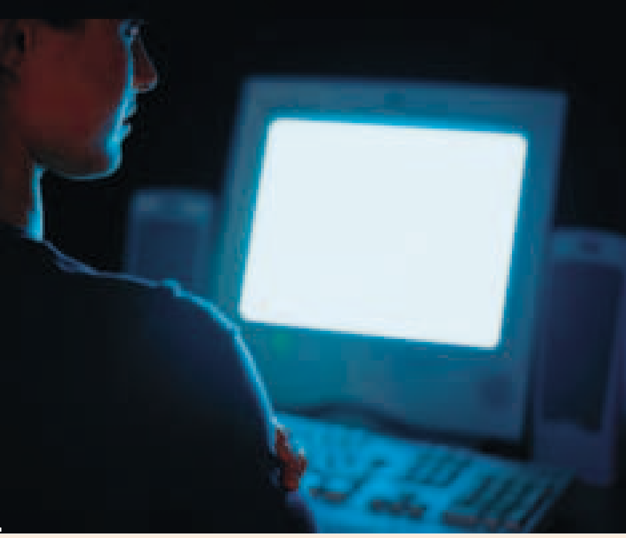

Funding for engineering and science research

Priorities for the future

Dr John Taylor, supported by the Science and Engineering Base Group of the Office of Science and Technology, is responsible directly to the Secretary of State for Trade and Industry for increasing the health and vitality of the UK science and engineering base. This is done principally through the UK science budget in basic and strategic research and related postgraduate training. Dr Taylor manages the government's relationships with the seven UK Research Councils (all independent non-departmental public bodies) through which most of the budget is channelled.

In addition he is responsible for ensuring that the knowledge generated by the science and engineering base is transferred out and exploited so that it has an impact on the economy and on society. John Taylor's role also involves taking the lead, through funding and encouragement of others, in the promotion of the public understanding of science and engagement of people in science-related issues and debate, including promoting greater participation by women at all levels of science. During his tenure as Director General, significant new investment has been made in the science and engineering base. In this interview we discussed some of these new science initiatives as well as more general issues.

What key issues must the UK address in order to maintain a strong science and engineering base?

From my perspective at the OST I need to consider what are the most important *outcomes* that the government expects for the public funding that it puts into science, engineering and technology research. I use a version of the 'right-hand rule' – I hold up three fingers at right-angles to represent the three most important outcomes, which are new knowledge and know-how, highly trained researchers, and impact on society and the economy ('knowledge transfer'). We look for a combination of these things and this brings some interesting challenges.

We have to strive to ensure that whatever we do is excellent: research needs to be first-class because if it is not, the rest tends to fall away. If you want excellent research then you must have excellent people, so a key challenge is how to recruit and retain the very best researchers in publicly funded research. This leads to issues of career structures, salaries and the research infrastructure. We have paid considerable attention to this; for example, we have recently made substantial increases to PhD stipends and the flexibility with which universities can fund them.

We need to provide much better facilities and infrastructure – that is one of the main things that attracts people to research. Essentially, for the past 10 years universities and research institutes have neglected the infrastructure and used all the money that comes in to pay for hiring more people and taking on more research. Schemes like JIF (the Joint Infrastructure Fund, which ran from 1998 to 2001) and now SRIF (the Science Research Investment Fund) are

> a key challenge is how to recruit and retain the very best researchers in publicly funded research

putting hundreds of millions of pounds into rebuilding infrastructure. We have to remember that research itself is getting more expensive 'per unit' (however you want to define a unit): it requires more equipment, more facilities, more infrastructure, more support.

Is there still perceived to be a 'brain drain'?

I don't think there is a significant 'net' brain drain, but there is a lot of mobility. This is to be encouraged since science is increasingly an international undertaking. Mobility around Europe and exchange with the US are both important, as of course is mobility between academia and industry, which can be a real issue when talking about salaries.

We in the UK find the idea of movement between academia and industry much harder than, for example, in the US. People do find it quite difficult to move because of the different career structures, reward systems, and so on: these tend to reinforce the idea that you should stay where you are, rather than taking the risk of moving across. If you move out of academia it is probably quite difficult to get back in – and the salary differentials mean you might not want to. Conversely, it is not easy for someone from industry to acquire enough academic 'brownie points' to move into academia.

These are issues on which we don't do well enough in the UK. The problem is compounded by the fact that we don't have enough industrial R&D labs in comparison with our main competitors – our level of industrial R&D is not as high as I'd like to see it, which means there aren't enough places to provide this interchange with academia.

How is the government helping to support emerging technologies?

My job is to consider the whole portfolio of science and technology that the public purse might fund. It is a top-down process: we're trying to look at the whole landscape of things that are likely to be important. Out of the total budget we might say: 'Let's allocate approximately this much for mathematics, this much for chemistry', and so on, but this is not enough because these traditional disciplines do not represent the real situation adequately. A great deal of the really interesting medium- and long-range research takes place between the disciplines, or where the disciplines come together. What we mean by something like 'chemistry' today is a hugely diverse area – from biosciences at one end to new materials at the other.

There is also a bottom-up process at work: we have to find the very best people to spend the public money in the best possible ways.

Somewhere in between the two we have to do a certain amount of steering – to say 'We'd like to have some new programmes'. For example, we think we are not doing enough in nanotechnology so we would like to strengthen that community. In the last Spending Review we looked at some major new programmes that cross traditional discipline (and Research Council) boundaries. The three new programmes we have set up are the post-genomics programme, the e-science programme and the basic technologies programme. Government doesn't try to pick winners but to ensure that there are enough resources in areas which people think will be fertile.

Why the emphasis on basic technology research?

In the UK basic science research is highly esteemed and we all see how excellence is publicly rewarded through

The 2000 Spending Review

The 2000 spending review outlined the science budget for the years 2001–02 to 2003–04. It includes an extra £1 billion to renew research infrastructure in partnership with the DfEE and the Wellcome Trust, and £252 million for three key areas: understanding the human genome, e-science and basic technology. The total funding in the science budget through the OST increases from £1.7 billion to approximately £2.2 billion by 2003.

Science budget additions (3 years)	£ million
E-science	£98
Genomics	£110
Basic Technology Research Programme	£44
Research infrastructure (SRIF)	£225
PhD stipends	£34
Knowledge transfer	£110
Research Council uplift	£104

Money is allocated to the Research Councils by the OST every two years for a three-year period. This means there is an overlap, but extra changes in this overlap are generally minimal – it is the time in which the Research Councils make arrangements for and design their programmes. Other major funding for universities is channelled through Higher Education Funding Councils (such as SRIF) or in collaboration with the DTI and other government departments.

Government doesn't try to pick winners but to ensure that there are enough resources in areas which people think will be fertile

Basic Technology Research

The Basic Technology Research initiative provides £44 million over three years. The programme is supervised by a steering committee covering interests across the Research Councils. It is supported by a Strategic Advisory Committee headed by Professor Alan McGregor of Kings College, London, with members from multi-disciplinary academic and industrial backgrounds. Although administered by the EPSRC it is strictly cross-council and it aims to get people thinking beyond disciplinary constraints.

The basic technology underlying a scientific discovery is often exploited across many disciplines. For example, atomic force microscopy is now widely used in materials science, biology, chemistry and engineering. The transistor is another obvious example. The direction in which research finds an application is often unpredictable, but the programme aims to provide seed-corn for new research to enable new capabilities. In turn, such capabilities can be exploited not only by industry but also to support new scientific advances.

There is huge scope for the application of fundamental technologies and sciences in new areas such as imaging, sensors, biomimetics and nanotechnology, but the programme is not prescriptive – good ideas are sought across the board.

Nobel prizes and so on. We want to emphasise that fundamental research in basic technology is also very important and deserves parity of esteem, as is the case (for example) in the US. I believe that fundamental research in science and technology forms a continuum. This research is driven by motivation – fundamental science research is done to acquire knowledge, fundamental technology research is undertaken to create a new capability. The researchers may have no idea of what they will do with it.

Think of a plasma physicist who might be researching the process of star formation in distant galaxies (in which case he is doing science) or might be trying to make fusion work in a reactor (in which case he is doing technology). Actually most of the time we probably couldn't tell which of the two he or she is doing.

So we have put serious new money into fundamental technology research and we have set up a single cross-council body to manage that money.

What is e-science?

What I mean by e-science is groups of scientists, technologists and engineers distributed around the world and collaborating using internet-type technologies to allow them to share resources of various kinds. This has already been happening over the last 10 years or so. Think of the sequencing of the human genome, which could not have been carried out without the internet because it was a giant global collaboration of at least 25 major institutions which shared the results of their experimental work via big databases on the internet.

Look at 'big' physics. We are building a new machine – the Large Hadron Collider – in Geneva which will begin operating in 2005. It is a giant accelerator ring that will produce a huge quantity of data. In order to analyse the hadron collisions the researchers will have about a petabyte a year of filtered data [1 petabyte = 10^{15} bytes or about 3 miles of CD-ROMS]. The data will be shared by physicists from all around the world. They will access it through an e-science infrastructure called the Grid which will allow them to move the data around the world and allow worldwide teams of physicists to collaborate.

Astronomy provides another example. We are going to build a 'virtual' observatory which will eventually put most of the data from most of the big telescopes on a Grid so that anyone, anywhere can look at the data from any of these instruments.

How does e-science improve on the technologies we use today, for example the web?

The web is a very primitive resource that can share data only in simple ways. It developed using html, a markup

E-science

The e-science programme will receive £98 million over three years. Of this, £15 million goes to develop core computing technologies, with further funding of up to £40 million provided by the DTI and industry under a Link scheme. A newly appointed Director of e-Science, Professor Tony Hey (on secondment from Southampton University), will lead this programme. The remaining funding will be used to develop applications technologies and is spread across the seven Research Councils to maximise the chances of novel applications and successes. Most applications are highly data-intensive, such as particle physics, health and bioinformatics, oceanology, climatology, meteorology, geology, fluid dynamics, national social science resources, in-silico design and testing.

The range of initiatives being developed to support the core e-science programme is based around six elements:

- A national e-science centre linked to a network of regional Grid centres;
- Generic Grid middleware and demonstrator projects;
- Grid Interdisciplinary Research Collaborations;
- Support for e-science test beds;
- Participation in international Grid projects and activities;
- The establishment of a Grid Network Team.

language for pages. The web exists today because of the research community – it was invented in CERN in order to help collaborative research on big physics. I believe the same thing will happen again in a major next iteration to set up the e-science infrastructure. This infrastructure will develop over quite a long time. The need for some type of solution to the Large Hadron Collider problem, for example, is defined by the fact that in 2005 data is going to come pouring out of the machine.

The basic way in which we use computers today is 'client–server'. The client is on your desk and you have to access the server – your organisation's server, a web server, or whatever it happens to be. You have to know where you want to go and how to get there. The new model now developing is 'client–utility'. The idea is that of an information utility, by analogy with the electric or water utilities. When you switch on a washing machine you don't have to worry which power station the electricity is coming from, or which reservoir the water is coming from – you have utilities 'behind the wall' that deal with those things.

The question now under development is how to provide such information and computing utilities that will allow you to access the resources you want without having to know where they are. A key aspect of the Grid model is the notion of doing 'in silico' experimentation, that is, carrying out many kinds of scientific experiment inside a computer instead of in the real world. For example, we can study protein folding or nuclear explosions or aerodynamics (wind tunnel experiments) by using very large computing resources.

The other aspect of the e-science infrastructure which is actually quite new to people is having very large collections of accessible data. This goes way beyond database management. For example, we might have one or two hundred large, complex databases of protein structures. How do we find our way around them? How do we know whether the data someone has put in there is good or bad?

We need to understand what kinds of electronic infrastructures help teams of widely distributed people to collaborate on very difficult problem solving. We need to work out how to organise the material not as passive web pages but as rich, complex data structures that they can add to and share. And we need to consider how to provide very powerful computational resources on which they can run experiments.

How will e-science change the way research is done?

The Grid will open up many fields to participation by relatively small research groups working almost anywhere in the world. If you want to do leading-edge astronomy you won't have to have a telescope on a mountain in Hawaii – you can use somebody else's. The whole concept of the Grid is that researchers will make their results available. People within defined communities will have to and will need to share their results. They will have to work out who owns the information and who is allowed to have access to it.

This raises many issues, not least that of intellectual property. Take post-genomics, for example. If information on DNA sequences is shared, the question becomes: 'if I use some of this information in the private sector to develop a new drug, whom do I reward?' Or: 'whom do I sue if the information is wrong?'

What is the government hoping to achieve by allocating funding to e-science?

We are trying to accelerate the emergence of the next generation of open source and open standards for utility computing. Just as the major achievement of the internet was the agreement to use TCP/IP as the protocol and the main achievement of the web was the agreement to use html, the question now is: what will everyone agree on as the platform for utility-style information systems?

> We need to work out how to organise material not as passive web pages but as rich, complex data structures

It is clear from the experiences with the internet and web that, if you establish open platforms that everyone supports, then people can build commercially valuable and sensible services and business models on top of those by adding value to the standards. There is considerable merit in getting everyone to agree quickly on open standards. Until that happens, there will be many vying factions and investors and application developers will not take much interest. A great deal of activity is going on in the US. We want to have the UK community participating in the international effort to get Grid-style utility computing paradigms and standards that are reasonably robust and can be agreed on as soon as possible. In fact the UK is ahead in Europe on e-science and we must now build on that success through recognising the implications of e-science for the whole science and engineering base. ■