

International
Review of
UK research in **COMPUTER SCIENCE**

01000100101100001110110101001001010010101001001010000110010101
00101111000101011110010110001010100100101100010101001010010100

00101010000010101011110001000
10010001010100100011110000110

Edited by Fred B. Schneider and Mike Rodd

EPSRC



BCS
THE BRITISH COMPUTER SOCIETY

INTERNATIONAL REVIEW OF UK RESEARCH IN

COMPUTER SCIENCE

Fred B. Schneider and Mike Rodd, editors

EPSRC BCS IEE

Contents

Foreword

Executive summary

1.	Introduction	1
1.1.	The nature of computer science	1
1.2.	About this study: Charge and inputs to the Panel	2
2.	Importance of computer science research for the UK	3
3.	Support for UK research in computer science	5
3.1.	Magnitude of international funding for computer science research	7
3.2.	Methods of UK funding for computer science research	8
4.	UK computer science research environment	12
4.1.	Attracting and retaining talent	12
4.2.	Engaging with the community	13
4.3.	Evaluating the impact of computer science research	15
5.	Quality of UK computer science research	17
	Comments on some specific broad areas of study	18
6.	Conclusions	26
7.	Acknowledgements	27

Appendices:

A.	Data supporting the Review	28
A.1.	Main data document	28
A.2.	Presentations	28
A.3.	Bibliometrics	29
A.4.	Site visits	29
B.	Terms of reference	30

FOREWORD

This review of computer science research in the United Kingdom is the third in a series which sought to present international researcher perceptions for core fields of the UK science and engineering research base. The first, undertaken in 1999, concerned engineering; it was followed, in 2000, by a study of physics and astronomy.

The Office of Science and Technology (OST) was interested in receiving an international assessment of the standing of British computer science research. The Engineering and Physical Sciences Research Council (EPSRC) sought to understand the strategic position of computer science, as perceived by international experts. As representatives of the relevant research communities, the Royal Society, the IEE, and the British Computer Society were keen that the survey should be undertaken with rigour, accuracy, and sensitivity. All agreed that to be useful, the review would need to command the respect of the UK computer science research community.

The Terms of Reference for this review were to:

- report on the standing and potential of computer science research being undertaken at UK universities,
- discuss its likely impact on the UK science base, and on the nation's wealth and well-being, and
- provide comparisons with computer science research internationally.

The review would be based on existing data, international assessments, and site visits. The IT/CS Programme within the EPSRC was an obvious focus, though research funded by other sources should also be studied.

The Review was overseen by a Steering Group comprising:

Professor John Midwinter, IEE (Chair);
Professor Richard Brook, EPSRC;
Professor Wendy Hall, BCS;
Professor Roger Needham, The Royal Society.

The IEE provided the Secretariat for the Steering Group and for the International Panel: Dr Mike Rodd (Director, Knowledge Services), Mr Martin Barrett (PN Manager) and Mr John Coupland (Head of Library Services). This Secretariat was supported by colleagues from the other three sponsors: Mr Vince Osgood and Dr James Fleming of the EPSRC, Dr Mark Scott of the Royal Society, and Mr Malcolm Sillars and Mrs Peta Walmisley of the BCS.

The Steering Group, working closely with the Panel Chairman, determined the membership of the International Panel. Every effort was made to cover the relevant sub-fields of computer science, to represent perspectives from industry, government, and academia, and to achieve both international and gender balance. The Panel members were:

Professor Fred B. Schneider, Cornell University, USA (Chairman);
Professor Susan Davidson, University of Pennsylvania, USA;
Professor Susan Graham, University of California, Berkeley, USA;
Professor David Harel, The Weizmann Institute of Science, Israel;
Professor Juris Hartmanis, Cornell University, USA;

Dr Butler Lampson, Microsoft, USA;
 Mr Michael Lesk, National Science Foundation, USA;
 Professor Dr Kurt Mehlhorn, Max-Planck-Institut für Informatik, Germany;
 Professor Moira Norrie, ETH Zurich, Switzerland;
 Dr Raymond Perrault, SRI International, USA;
 Professor Larry Peterson, Princeton University, USA;
 Mr Justin Rattner, Intel Corporation, USA;
 Professor John Stankovic, University of Virginia, USA;
 Professor Axel van Lamsweerde, Université catholique de Louvain, Belgium.

The Steering Group is grateful to the Panel and especially to the Chairman, Professor Fred B. Schneider, for the time committed to this review and for the thorough approach taken to the work. The findings and recommendations will be important to all the sponsors in helping to shape strategies for future support of computer science research in the UK. The report is commended to the British computer science community for its consideration and comment.

Professor Richard Brook	Chief Executive, EPSRC
Professor John Midwinter	President, IEE
Mr Alastair Macdonald	President, British Computer Society
Professor Lord May AC Kt	President, The Royal Society

June 2001



The picture has the panel from left to right as follows:

David Harel, Michael Lesk, Butler Lampson, Moira Norrie, Axel van Lamsweerde, Fred B. Schneider, Ray Perrault, Susan Graham, Kurt Mehlhorn, Susan Davidson, Justin Rattner, Juris Hartmanis, John Stankovic, Larry Peterson.

EXECUTIVE SUMMARY

Computer science in the UK has traditionally been of the highest quality. However, while the UK remains a world leader in some research areas and is a strong participant in many others, this position is by no means assured. Declines in certain fields are already evident; more will follow, given the current levels of support and the nature of today's university research environment. The consequences could be far-reaching. Computer science is not only an academic discipline offering deep intellectual challenges; it is also a discipline where research results can translate into competitive advantage and economic well-being on a local, national, and international scale.

The market for skilled computer scientists is keenly competitive at all levels, from postgraduate students to senior professors. Industry in the UK and elsewhere, and universities in the US and Europe, are eager to hire talented people. Academic salaries and working conditions in the UK, however, are not competitive. So the talent pool in UK universities is shrinking, with senior researchers leaving and students forgoing advanced degrees for lucrative industrial positions. Some means must be found to reverse this exodus. Increased salaries should be seriously considered—especially for junior positions, where there is limited opportunity to augment one's income from outside sources.

The level of EPSRC research funding and the manner in which it is distributed make it unnecessarily difficult to be an effective researcher in a UK university. Funding is low by international standards, and responsive-mode grants—the primary means of financing faculty research—are too small and lack the flexibility that would permit investigation of questions that had not originally been anticipated. Research infrastructure (both staff and equipment) is inadequately supported; funding to build research platforms, so that artefacts can be distributed and studied by a community, is also difficult to obtain. Moreover, the process by which funding decisions are made discriminates against proposals in new areas, proposals involving larger experimental projects, and proposals that describe high-risk or inter-disciplinary projects. Expectations concerning industrial tie-in for funded research are counter-productive; truly significant industrial impact usually requires establishing new companies.

Dividing the IT/CS Programme into one programme for computer science and one for physical layer technology would allow increased visibility, better control, and independent budgeting for each of the disciplines (which are quite different in character and mode of research). This would bring clear benefits. But even if structural changes are not made, decisions about funding levels, research programmes, and the funding of individual grants would benefit from the EPSRC having significantly more input from computer scientists. Participation by computer scientists should be mandatory on advisory panels that set EPSRC priorities; computer scientists should also be involved on a day-to-day basis in the management of both the EPSRC and the IT/CS Programme, so that the role, the potential, and the culture of the field are accurately represented and best leveraged within the EPSRC. A form of secondment seems a particularly attractive means of staffing the IT/CS Programme with senior computer scientists.

The UK's research strength and international leadership—now at risk—were noted by the Panel in a broad set of areas:

- logic, semantics, and formal methods within the theory community, as well as activities in quantum information processing;
- programming language design, enhanced by strength in logic, semantics, and software engineering;
- technologies that contribute to system trustworthiness, including software engineering, security, and dependability;
- architectures and algorithms for the design of real-time and embedded systems;
- artificial intelligence areas: speech engineering and computational linguistics, machine learning, artificial neural networks, computer vision, and automated reasoning;
- human-computer interaction;
- bio-informatics.

Two noteworthy imbalances should be redressed by encouraging research in algorithms and in experimental computer systems. The UK has never exhibited broad strengths in algorithms and complexity research, yet expertise here will become increasingly important for applications of computing; research in experimental computer systems has been a strength but is now being eroded. In addition, given the opportunities that the recently announced 3-year e-Science initiative will create, the decline in UK research activity in high-performance scientific computing and the absence of funding within e-Science for longer-term computer science research seem ill-considered.

“The high state of wealth and civilization which the English people have attained within the last half century, has been greatly promoted by the application of the power of the steam-engine.”

John Farey, *A Treatise on the Steam Engine*, 1827.¹

1. INTRODUCTION

A nation’s economic prosperity need not be hostage to natural resources or the availability of cheap labour. Technology can bring economic well-being. Steam power once played that role, as John Farey notes; today, most believe that computing and information technology reign and will continue to do so well into the next century. Nations that dominate this sector can ensure their businesses a competitive advantage and provide their citizens with a higher standard of living. Research in computer science—the mechanisation and representation of processes, knowledge, and information—is crucial for achieving and maintaining that dominance.

The UK has a proud history of accomplishments in computing and information technology. But the field moves rapidly, and leads erode quickly. So investments in research and education are critical for supplying the ideas and the staff needed to compete technologically and economically. Identifying areas in which investments should be made, providing sufficient levels of funding, and maximising the leverage from those investments are crucial. It is sensible to seek outside expert advice on these matters; hence this report—a synthesis of one week’s deliberations and university site visits by an international panel of fourteen computer scientists in June 2001.

1.1. The nature of computer science

The Terms of Reference (see Appendix B) provided for this study characterise computer science by means of a list of research areas. Though such a listing might today be accurate, it misleads. Ours is decidedly a new kind of discipline—one in which, for example, theory is not concerned with explaining extant physical phenomena, and experimentation is not necessarily concerned with testing whether theory predicts reality. Viewing computer science through the lens of traditional academic disciplines can lead to misconceptions of what is important and of how research in computer science is done, even though aspects of computer science can be traced to the natural sciences, mathematics, or engineering.

Theoretical research in computer science concerns the power, limits, and costs of computation. A theorem might, for example, characterise a class of problems that cannot be computed by a digital computer, thereby making a statement not only about present-day hardware but also about any digital computing device that might ever be built or conceived. Theoreticians also derive bounds on the time or memory required by any program that solves problems from a given class (such as searching, sorting, or scheduling). Sometimes the cost or intractability of a problem is reason for dismay;

¹ This quotation is prominently displayed amongst the steam engines in the East Hall of the Science Museum in London. Artefacts exhibited elsewhere in the museum presage the revolution in computing to come. One can find portions of Charles Babbage’s Difference Engine No. 1 and his Analytical Engine (under construction at the time of his death in 1871), as well as a full-sized operational Difference Engine No. 2, constructed recently from drawings prepared by Babbage.

sometimes not—the premise of modern cryptography is that reversing certain methods of encryption would be intractable. Much of the work in theoretical computer science is mathematical in character. And formal logic is central. By definition, each step in a formal proof must be mechanisable, creating an intimate connection between proof and computation. The study of programming languages, like the study of logic, is concerned with the expressive power of formal notations, with correspondences between syntax (programs) and semantics (what they mean), and with the means by which texts in a formal language can be analysed (automatically or manually) in order to extract truths.

Experimental work plays a very different role in computer science research from the part it plays in the natural sciences. Rather than attempting to understand an existing reality, experiments in computer science are often intended to explore new approaches or abstractions. Here, a research prototype might be built, instrumented, deployed, and measured in order to evaluate the strengths and weaknesses of something that the prototype embodies. Experiments are pivotal in the natural sciences when they demonstrate aspects of reality that depart from what current theory predicts and thus defy our understanding; in computer science, it is prototypes that launch paradigm shifts. Building a prototype might expose implicit assumptions, prevent key sub-problems from being ignored by its builders, or allow its users to discover synergies and unanticipated uses—issues that would not be addressed if analytical techniques were applied to a paper design. Timeshared computing, the personal computer (with all its productivity-enhancing software), and the Internet itself, all started as experimental prototypes.

Computer science is also distinguished by the manner and extent to which tomorrow's results and artefacts are obtained by using today's. A civil engineer rarely builds a highway by using repeated traversals to strengthen a footpath. But in computer science, we find researchers developing an operating system by running that operating system, and researchers in programming languages writing compilers for a new language in that same language. Indeed, one of the fundamental models of computation studied by theoreticians—the Turing Machine, named after British computer scientist Alan M. Turing—is a device that can simulate itself or any other digital computing device. Abstractions and artefacts exhibiting such universality are highly prized in computer science, and the study of the intellectual glue by means of which tools—mental and developmental—can be composed is among the field's highest callings.

1.2. About this study: Charge and inputs to the Panel

The Charge to the Panel (see Appendix B) was, in essence, to evaluate the standing and potential of computer science research in the UK. The size of the Panel and the limited time available for discussion made it impossible to conduct a thorough analysis of each sub-field of computer science or each UK university department in which researchers are active. So the Panel was forced to draw upon its members' personal experience and knowledge of the UK and international research communities, sharpened by in-depth briefings about the UK system and by visits to selected universities with 1996 RAE Grade 5*, 5 and 4 computer science departments.²

² Bristol, Cambridge, Cardiff, Edinburgh, Heriot-Watt, Imperial College, Manchester, and UMIST were visited. The panel was split into four small groups. Each group visited two universities.

The Computer Science portion of the IT/CS Programme in the UK Engineering and Physical Sciences Research Council (EPSRC) was studied with some care, because this is a primary source of funds for university research in computer science. The IT/CS Programme exerts considerable influence on what research is pursued, not only by choosing which areas to support but by its management structure, by the way it interacts with UK computer science researchers, by the process it uses to review proposals, and by the duration and size of the proposals that it tends to fund.

In conducting this study, the Panel had access to various inputs, including:

- a body of data provided by the secretariat, concerning funding, staffing, and outputs associated with computer science research in UK universities;
- a bibliometric analysis, undertaken by the IEE using its INSPEC database and covering a broad spectrum of computer science journals and conference proceedings;
- position papers (“UK Computer Science Research: Vision and Opportunities” and “Input to the Quinquennial Review of Research Councils”) prepared by the UK Computing Research Committee, a body of senior computer scientists from academia and industry;
- oral presentations from members of the steering committee, the Chief Executive of the EPSRC, and the IT/CS Programme management, regarding the funding scene in the UK, the UK academic system, and the EPSRC’s funding mechanisms and programmes;
- on-line searches to access specific information, including data on funding and research programmes in other countries.

The Panel did not validate this information or what was offered at the various briefings and visits. Thus, there are places in this report where the Panel is generalising from anecdotes. In reporting what is believed by the researcher community, the report samples a community’s perceptions; misconceptions widely held by the community obviously need to be corrected. Finally, it is worth emphasising that the report is a record of the views of fourteen individuals, who each participated in brief visits to two UK universities. The report is not an exhaustive study of the computer science research being conducted at UK universities—there seemed to be little point in duplicating the Research Assessment Exercise (RAE) currently in progress, and the raw data gathered for that RAE was not made available to the Panel.

2. IMPORTANCE OF COMPUTER SCIENCE RESEARCH FOR THE UK

Success today in a developed nation’s industry, education, and commerce seems to depend increasingly on bringing the latest computing technology to bear. Failure to keep up brings the risk of failure in the global marketplace. And with barriers to international commerce falling (especially in western Europe), failure in the global marketplace brings failure in the local marketplace. In addition, we find critical national infrastructures—communication, finance, energy distribution, and transportation, not to mention civil and national defence—also coming to depend more and more on networked computer systems. Thus, at least in developed nations, quality of life is affected by access to computing technology and expertise in deploying it. Moreover, the providers of computing technology can profit handsomely in this environment, and that further buttresses the economy. In the United States, a world leader in this arena, computing technology and related

services made up 8% of its GDP in 1998 and provided 35% of its economic growth; workers' wages in this sector were 85% higher than the private-sector average.³

With the price of computing power decreasing exponentially⁴, the feasible domains for the application of computing technology are spreading at an accelerating pace. Today, corporate computing and storage capabilities are extraordinary, powerful desktop and laptop PCs have become ubiquitous, and handheld personal digital assistants are found in increasing numbers of pockets and handbags. Embedded processors and networked devices are already becoming prevalent. In some cases, they enable new functionality; in other cases, they re-implement old functionality in more cost-effective ways.⁵ New industries will be created. They, along with the old industries, will require staff who are conversant in the new technologies—staff who risk becoming irrelevant unless they stay current with a steady stream of technological developments being made in “Internet time”.

The UK's ability to educate its young people in computing technology is clearly crucial, both in order to satisfy staffing needs and as a means to inject current knowledge into the practice. A strong programme of computer science research is also critical, for the following reasons.

- It virtually ensures a steady stream of ideas that can be monetised by incorporating them into products and services.
- It shortens the time interval between when a discovery is made and when that discovery can be exploited, because less geographical and cultural distance must be covered from the producer to the consumer. Intel's Andy Grove said “We are moving from a world in which the big eat the small, to a world in which the fast eat the slow.”
- It is an efficient and effective means of educating teachers, because some fraction of the postgraduate students in the field can be expected to remain in academia, although not necessarily as active researchers.

The structure of the computing industry has changed over the last decade, making it feasible for virtually anyone to compete on a global scale. Long gone are the days when software and hardware were bundled so that computer manufacturers had an inside track for selling their software. Today, the industry is segmented horizontally, with public, rather than proprietary, interfaces defining the segment boundaries. Having a resident hardware industry is no longer essential. Individual companies obtain the greatest leverage by dominating an existing horizontal segment or by creating a new one. Success now becomes possible, for example, by building a popular Web browser or by designing and licensing the format for encoding and transmitting music files. Notice, though, that the earliest entrants into a market gain a strong competitive

³ US Department of Commerce, Economics and Statistic Administration. Digital Economy 2000. US Department of Commerce, Washington, D.C., June 2000.

⁴ In particular, the price of processing power and memory halves every 18 months, the price of disk storage halves every 9 months, and the price of optical bandwidth halves every 10 months.

⁵ Among other things, software is easier to change than hardware. Innovation that can be implemented entirely in software is thus quite cheap.

advantage. Agility in recognising opportunities and speed in deploying solutions are what counts, so a strong computer science research programme is a clear asset.

Examples of computer-related industries that come from smaller countries and compete successfully are well known: mobile phones from Finland and display screens from Taiwan, to mention two. A nation like the UK, which already has an extremely strong computer science research programme, has an advantage over countries who have not invested and, therefore, have not built up a culture and infrastructure within their universities for research excellence in computer science. While the computing industry may run on “Internet time”, universities are relatively slow to change. So, considering the economic payoff that a strong UK presence in computing technology could bring, investments in computer science research are easy to justify. Even to assume a purely defensive posture and protect current UK market strengths would seem to justify the investments: there is likely to be a revolution in the recorded music industry brought about by on-line access, a revolution in medical care is being brought about by new sensing techniques and robotic assistance, the Internet is globalising education (by facilitating “distance learning”), and on-line content delivery is coming of age.

3. SUPPORT FOR UK RESEARCH IN COMPUTER SCIENCE

The IT/CS Programme in the EPSRC is responsible for supporting research in both computer science and the so-called “physical layer” technology: electronic and photonic devices, VLSI design, digital signal processing, and RF and microwave technology. Physical layer technology research receives 55% of the IT/CS Programme budget; computer science research gets 45%.

The IT/CS Programme budget⁶ for 2000/01 is £70.3 million, allocated as follows:

- £53.3 million for new projects;
- £8.5 million for doctoral training (3 years of training for the equivalent of 235 new PhD starts);
- £8.5 million for masters training (1 year of training for the equivalent of 900 new starts).

Thus, £31.75 million of the IT/CS Programme’s planned expenditures (45% of £70.3 million) is available to satisfy new computer science research baseline programme commitments (research grants, doctoral training, and masters training), and of that, £24.075 million—approximately 5% of the EPSRC budget—can be allocated to new computer science research projects. In addition, the EPSRC recently made available an additional £45 million to support five six-year interdisciplinary research collaborations (IRC).

The funding level of the EPSRC is currently planned to remain constant over the next three years, as is funding for the IT/CS Programme. The EPSRC 2000/2001 Annual Report⁷ shows that the number of proposals considered by the IT/CS Programme decreased over the three-year period 1998–2000, as did the number and value of the

⁶Vince Osgood. “The EPSRC IT and Computer Science Programme”, Computing and Control Engineering Journal, August 2000, pp. 168–172.

⁷ Engineering and Physical Sciences Research Council Annual Report 2000/2001. Laid before Parliament on July 19, 2001.

proposals that were funded. Figures for the computer science research portion of the budget for 1998–2000 were provided by IT/CS Programme management, and they show: (1) that the number and value of submissions and proposals increased during this three-year period, and (2) that the value of the grants awarded increased in the first year and remained approximately constant for the next two.

Industry in the UK also funds research in computer science (£7.5 million per year on average, during the 1996–2000 period), as do the ESPRIT/IST and other EU programmes (£11.5 million per year on average during the 1996–2000 period).⁸ Only certain, more-applied research areas are able to attract this non-EPSRC funding. In addition, the Panel was told that obtaining funds from EU sources is becoming an increasingly complex and expensive process, and the funds are then often tightly constrained, making this source of support progressively less attractive.

At what level should computer science research be supported in the UK? This question might be attacked in a variety of ways, each likely to yield a different answer. The Panel considered two lines of inquiry for which the data was readily at hand:

- compare the level of UK expenditures for computer science research with expenditures made elsewhere, while acknowledging the difficulty in drawing conclusions because of differences in budgeting rules, educational structures, and so on;
- identify the impact that changes to the level of research funding might have on the nature, quality and quantity of computer science research being done in the UK, given the current staffing levels.

A more careful analysis, along the lines of the recent US National Research Council CSTB⁹ study, might prove valuable to some UK policy makers. Nevertheless, the Panel's conclusions, developed in detail below, are hard to contest.

- By international standards of comparison, the UK significantly under-funds its computer science researchers.
- This under-funding affects the nature of the research being performed, discriminating against researchers who attempt to run larger groups and against experimental research efforts.
- There is insufficient funding to support the much-needed general-purpose research infrastructure or to support the construction of special-purpose research platforms.

In short, not only can increased funding for computer science research be justified in terms of parity with the investments being made by international competitors, but such increases would also leverage the UK's best researchers and, therefore, enable them to conduct research that could have industrial (and hence economic) impact.

⁸ This data, provided by Ian Wand, is derived from information collected for the 2001 RAE.

⁹ Computer Science and Telecommunications Board (CSTB), National Research Council. *Making IT Better: Expanding Information Technology Research to Meet Society's Need*. National Academy Press, Washington, D.C., 2000.

3.1. Magnitude of international funding for computer science research

Computer science research in the United States is funded by a number of governmental agencies. According to the CSTB study cited above, the total US government computer science research expenditures in 1998 were \$1,399 million, with approximately a third (\$419 million) devoted to what was termed “basic research”. So almost ten times as much funding (\$419 million versus £31.75 million) is spent on basic computer science research by the US, a country five times more populous than the UK. The Panel believes this two-fold difference is actually an underestimate; certainly, the difference is very significant. Also, when compared with total US government expenditures for computer science research, the population-corrected difference is closer to fifteen-fold, although admittedly this overstates the discrepancy because the US figures include support for high-performance computing centres, and the UK figures do not.

Three agencies (the National Science Foundation (NSF), the Defense Advanced Research Projects Agency (DARPA), and the Department of Energy (DOE)) in the US together support \$365 million of the basic research in computer science, and the NSF is responsible for funding the lion’s share (\$258 million). The NSF Directorate for Computing and Information Science and Engineering (CISE) accounted for 10.5% of the NSF budget in 1998; this year (2001) it accounts for 13.6%, an increment of roughly one third. This increment can be directly traced to a report prepared by an Executive Branch advisory board, the President’s Information Technology Advisory Committee (PITAC). That a new US President has renewed the PITAC charter indicates a continuing belief that computing technology is important for the economy and well-being of the United States.¹⁰

According to the CSTB report cited above, industry’s support for university research in the US is growing but represents less than 10% of all university computer science research funding. A select set of the top-ranked US universities, however, actually receive 20–30% of their research funding from industry.

International comparisons of research funding are difficult to make, since different countries use different funding methods. France and Japan, for example, rely heavily on national research institutes and laboratories, rather than expecting most research to be done in university departments. However, virtually every major country that the Panel investigated has identified computer science and information technology as a major research area. In Japan, for example, a recent reorganisation created a major new research organisation, the National Institute of Informatics, to pursue such research.

Computer science receives a relatively small portion of the research funding in many countries, but anecdotal evidence suggests that the fractions are increasing. Perhaps most ambitious is the EU Framework VI proposal for spending 16.27 billion euros over 2002–2006, of which 27% is destined for information technology (although not all of it for fundamental research). The following figures are also significant.

¹⁰ The role of PITAC also represents an interesting point of contrast with the UK, where there is no standing government advisory panel comprising computer scientists—not even within the EPSRC—to provide advice on funding computer science research.

- The DFG in Germany last year spent 841 million DM on projects, of which 92.5 million went to projects in the area of electronics and information. Since the DFG funds a broad spectrum of research, including the humanities and social sciences as well as biomedical fields, it might be more representative to look at computer science funding as a fraction of the natural sciences and engineering budget: it is 19% of that total.
- In France, a recent new “Information Society” programme has been funded, with 45.7 million euros. The Ministry of Research there concentrates its funding at a few major research institutes; a new programme¹¹ is adding 885 million FF in the fiscal year 2001 for special research topics. From this, IT/CS (that is, computer science and physical layer technology) will receive an additional 52 million FF.
- In Scandinavia, 15% of Denmark’s budget for the physical sciences and engineering is spent on research in computer science. Norway invests 250 million NOK into information and communication technology, 14% of a total budget of 1743 million NOK in the contributing research councils.

With the notable exception of Scandinavia (which is well advanced in the use of computing and information technology and clearly benefits from those investments), smaller countries are not making such substantial commitments to computer science research.

- Italy is spending only about 9% of its physical sciences budget on “informatics”. All told, a relatively small amount of the total scientific research budget in Italy goes to computer science—only a few percent.
- Computer science research in Canada and Australia receives about 6% of the research budget.
- In Hong Kong, the Research Grants Committee puts about 7% of its budget into computing.

3.2. Methods of UK funding for computer science research

Responsive-mode grants. A majority (65%) of IT/CS Programme-supported computer science research is funded through so-called “responsive-mode” grants. Here, a researcher writes a proposal describing some well-defined research problem or project. That proposal is then peer-reviewed and ranked, and a funding decision is made, based on the ranking and the availability of funds. The Panel was told that the success rate for responsive-mode grant proposals is 30–35%. The Panel sampled proposals that had just missed getting funded, to investigate whether high-quality proposals were being rejected—it did not seem that they were, although the Panel heard anecdotal accounts of good proposals being declined for lack of funds. A small number of funded proposals were also reviewed by the Panel, and there was no evidence that poor-quality research proposals are being selected for funding.

¹¹ See <http://www.recherche.gouv.fr/recherche/finance/fns.htm>.

The typical responsive-mode grant tends to be narrowly scoped and supports one Research Assistant (RA) and possibly one PhD student, together with a small allotment of money for travel and equipment. This has the following consequences.

- The most active UK researchers are forced to seek multiple grants in order to maintain even a moderate-sized research group. These researchers must pay a high price in proposal preparation, periodic status report preparation, and managing multiple budgets. Moreover, most of this price cannot be seen as enhancing the bottom line—that is, the quality and amount of research.
- Unanticipated research questions, which inevitably arise in the course of an investigation, may go unexplored while awaiting a subsequent grant, because a researcher's current (responsive-mode grant) funding has been so narrowly scoped. In theory, a researcher could explore anything, once a responsive-mode grant has been awarded. But when a grant ends, the researcher must write a summary of accomplishments, which is forwarded, along with the original proposal, to reviewers for evaluation. Few reviewers ignore the original proposal in writing their evaluations, so researchers are not comfortable in exploring issues that are outside the scope of the original proposal.
- It is difficult to get support for experimental research, since such efforts can seldom be described at the granularity of the typical responsive-mode grant proposal. The computer science experimentalist might start with a vision for a system or a point of leverage begging to be explored; only in the course of building prototypes would specific technical problems emerge (and be addressed). So, while a vision could be articulated for a proposal before the project starts, the specific problems—so important for getting a responsive-mode proposal funded today—could not.

The Panel believes that these restrictions on responsive-mode grants are having an adverse affect on the nature and quality of computer science research throughout the UK. Responsive-mode research support should be expanded to allow projects with larger budgets, so they can support an active senior researcher's entire team, and to allow the undertaking of experimental projects, where a compelling vision (with sound technical justifications) should suffice to justify the grant.

Research infrastructure. More than ever before, computer science departments engaged in research require an infrastructure comprising hardware (desktop computers, networks, and servers), software, and the technical staff to keep these systems running. The lifetime of a typical desktop PC today is three years, and system software must be updated far more frequently. The rate of change, which can be directly attributed to Moore's Law¹², is not likely to moderate any time soon. Without up-to-date facilities and technical staff to oversee them, researchers are at a considerable disadvantage.

- Time spent by researchers on system configuration issues is time taken away from doing research.

¹² Gordon Moore observed in 1965 that the number of transistors that fit on a computer chip seemed to double every 18–24 months. Still true today, this is now known as "Moore's Law".

- Systems built elsewhere may not run, or may not run well, on out-dated equipment, depriving researchers of an opportunity to analyse, build on, or leverage the work of others.
- Systems research done on old equipment risks irrelevance, because performance trade-offs and resource limitations change in successive generations.

Site visits conducted by the Panel suggested that strong UK computer science departments still lack funding for infrastructure. Researchers (often Research Assistants) are being diverted from doing research into doing facilities management; and an environment where each piece of equipment is associated with a particular funded responsive-mode proposal is inhospitable to the serendipitous research excursions that can lead to entirely new research directions and revolutionary innovation.

What is needed is a means of funding computer science infrastructure over reasonable time frames (say, five years) for groups of projects that could benefit from the economies that shared equipment and support staff allow. “Quality Research” funding from the Higher Education Funding Council (HEFC.QR) is intended for this purpose.¹³ But these funds seem to be required elsewhere—notably to reduce the student/staff ratio by supporting staff salaries. Moreover, HEFC.QR funds are seldom returned directly to the department’s researchers who bring in those funds.¹⁴ The recently announced EPSRC Platform Grants programme (which commenced in April 2001) does not seem to address the need for infrastructure either, since Platform awards are intended for single groups, primarily to provide key researchers with funding stability over longer periods than responsive-mode grants now permit.

Research platform support. In some areas of computer science research, results are communicated not by publications that are read but by software that is used and modified. Such “research platforms” might be built by a single research group or by a collaboration. The research platform is then adopted by a broader community and may well end up defining the context within which research in a sub-area is conducted—extensions are contributed, unanticipated uses of the research platform are investigated, limits are explored, and so on. The result can be catalytic and integrative for a research community, leveraging an investment so that all can profit and so that technology transfer is facilitated. The early visibility in hardware and software systems of the Cambridge and Manchester departments can be attributed, in part, to the research platforms that they produced.

But building a research platform is different from doing ordinary experimental computer science research, where prototypes are not expected to be long-lived. A community is unlikely to embrace a research platform without the promise of on-going maintenance, if documentation is inadequate or if the construction is not of production quality. EPSRC Platform Grants could have become an appropriate instrument for supporting the construction of research prototypes, although the current programme

¹³ John Midwinter. “The funding of British University Departments” (briefing paper prepared for the Panel).

¹⁴ In many cases, HEFC.QR funds are not even returned to the department that brings the money in, so a strong computer science research group or department may end up subsidising work elsewhere.

formulation¹⁵ somewhat misses the mark. For example, listed among the criteria for obtaining a Platform Grant is a track record “including strong support from industry”, and that seems irrelevant. Also, funding for longer-term technical support staff is not listed as an intended use for Platform Grants, yet that seems crucial for supporting a research platform.

Issues regarding peer review. EPSRC grants are evaluated by panels populated by the proposer’s scientific peers. Experts in an area review and rate proposals in that area. While one might expect this process to lead to sensible funding decisions, other dynamics intrude. In particular, Panel members’ interviews with UK computer scientists revealed two major drawbacks.

- Peer reviewers tend to favour support for incremental research in existing areas. This behaviour is not surprising, having been observed elsewhere in connection with paper selection by conference programme committees and journal submission reviewers. However, when the ultimate decision is made by a subject specialist who uses the peer reviews (along with some knowledge of each reviewer’s scientific predilections) as input, the tyranny of the conservative majority can often be overcome. The current culture of the EPSRC, however, does not seem to allow subject specialists to make the final decisions on proposal funding.
- Peer review panels recommend reductions in the requested levels of support, even for world-class researchers proposing exciting new work, because of the reviewers’ very real perceptions that their own support must also come from that same fixed-size pool. Clearly, if the funds available for computer science research were increased (as the Panel advocates), then this phenomenon should become less pronounced. A belief by some university faculty that all university computer science teachers should be doing research (and therefore that any computer science teacher could be qualified to serve as a peer reviewer) further exacerbates the problem, for it lends credibility to the belief that funding should be shared equitably by the entire university computer science community.

The Panel is particularly concerned with a widespread perception among UK researchers that proposals in new areas or in interdisciplinary subjects, and proposals that involve radical new ideas, are unlikely to be funded. Such research can lead to a far greater return on investments than the funding of incremental advances. Some means must be devised to ensure that UK researchers are encouraged and supported in undertaking high-risk, inter-disciplinary, or so-called “blue-sky” projects. This may require significant changes to the process used for deciding whether to fund certain proposals.

Prescription for industrial impact. Much computer science research will not, and cannot, profit from industrial collaboration (if only because practical deployment of the work is far distant), and scientists forced into such marriages are being diverted from more productive uses of their time. Moreover, a researcher who, in the name of industrial collaboration, attacks and solves “toy” problems will have little chance of affecting practice—for the time and effort generally required to understand the real problems

¹⁵ “Platform Grants Guidance: IT and Computer Science (ITCS) Programme”, http://www.epsrc.ac.uk/documents/programmes/structure/itcs/pgrantitcs_nw0.htm.

facing practitioners dwarfs the funding-scales of a typical grant. Thus, requiring industrial collaboration of computer science researchers is rarely effective, and sometimes even counter-productive. Today's exhortations towards industrial collaboration are a misguided criterion for determining computer science funding.

Furthermore, research that ultimately has had a major impact on industry is usually work that, at the start, was considered irrelevant by industry itself: RISC processors, the Internet, and the Web are only three of many examples. Usually, new companies must be started to exploit such work. As Gordon Moore said, "It's much easier to get a venture capitalist to give you money than to persuade the management of a successful company to try something new." So, up-front industrial sign-on is often a negative sign.

For these reasons, an organisation attuned to the immediate needs of industry, such as the Department of Trade and Industry in the UK or the Department of Commerce in the US, is, by its nature, likely to make the wrong decisions in choosing what computer science research to fund. This is not to say that the DTI still has no valuable role to play; its input is still vital in making it easier for new UK companies to start and especially in making it possible for them to grow and thrive. The UK has had many successful startups in computing technology, but none has grown to the size of Microsoft, Oracle, Cisco, Dell, or Sun.

4. UK COMPUTER SCIENCE RESEARCH ENVIRONMENT

4.1. Attracting and retaining talent

The demand for computer scientists at every level far exceeds the supply in virtually all the developed nations. Industrial salaries in the UK reflect this discrepancy and are relatively high, whereas UK academic salaries for computer scientists do not. The UK academic culture, with its salaries that are uniform across fields of study, location and stature, makes the situation especially grim. Moreover, working conditions in UK academia are, if anything, less attractive to the research computer scientist than industrial positions—what industrial positions lack in freedom, they make up for in lower overall workloads and a better research environment (with an ample supply of interesting, important research problems, hardware, and support staff).

Whilst salaries in the UK at the professorial level are lower than those found in industry, senior academics do negotiate their salaries and have the means to supplement their incomes through consultancies, by participation in their own companies, or by writing. Younger academics have fewer such opportunities and, as they will provide the UK's next generation of computer science leaders, it is their salaries that most concern the Panel. Salaries for lecturers and senior lecturers are extremely low—often less than half of what could be expected when working in UK industry or in another country. This situation is compounded by the high teaching and administrative loads normally required of people in lecturer and senior lecturer positions, as well as by the higher living costs associated with the UK's larger cities.

Post-doctoral Research Associates (or Assistants) are crucial in any research programme, and the Panel was impressed by the quality and commitment of those whom they met. There is, however, a discernible problem with recruiting suitable persons, especially

“home-grown” PhD graduates. The temptations of significantly higher salaries in industry and the debts now being incurred by undergraduates result in an ever-increasing reliance on foreign students to fill this role.

The recently enhanced level of EPSRC Studentships has not appreciably helped in persuading top-flight graduates to study for PhDs. Additionally, the expectation that a student will complete the PhD in three years (compounded by punishing a department where students exceed a four-year barrier) is counter-productive. PhD students, mindful of the time limit, find themselves:

- embarking on less ambitious dissertation projects;
- unable to participate in broadening activities (for example, additional courses, seminars, pursuing unrelated technical problems they encounter) that would prepare them to be better researchers; and
- removed from their research advisor’s mainstream research projects, where they might better learn about research per se as apprentices.

So, because of the three-year limit, which is often exceeded anyway, these newly educated scientists tend to lack breadth, experience, and important research skills. In the US, a student spends between four and seven years working on a doctorate after receiving an undergraduate degree. Systems students and experimentalists take longer; theoreticians take less. The situation is similar in other European countries.¹⁶

All told, today’s UK academic researchers have many incentives to leave. They can move to industry, or they can take an academic position abroad, where salaries are high and working conditions are better. Post-doctoral Research Associates (or Assistants) similarly have few incentives to work toward senior university positions in the UK. In discussions with Research Assistants, the Panel found that few saw themselves moving into full-time UK academic posts; the majority will seek research posts outside the university system or on foreign shores. The Panel also noted an increase in the number of Far Eastern and Indian computer scientists being recruited into academic posts. Foreigners do not have the same ties to the UK as residents do and, having moved once to better themselves and exploit wider opportunities, might be more inclined to do so again.

4.2. Engaging with the community

The Panel applauds the formation of the UK Computing Research Committee (UKCRC), in part, because it is an organisation that can articulate the value of computer science research within the UK. Policy makers can now get information that will allow them to choose sound and sensible national and research-spending priorities. The UKCRC should, however, be seen by the computer science community as only a first step. Ultimately, it will be necessary for UK computer scientists’ voices to be heard at the highest levels of government, and UK computer scientists must be prepared to vacate their classrooms and laboratories (temporarily) to answer this call.

The setting of research priorities in computer science and the setting of priorities within the EPSRC beg participation by UK computer scientists. The Panel was quite disturbed to note that the EPSRC IT/CS Programme is, by design, not managed by a

¹⁶ In France, the average PhD completion time is reputed to have decreased to 3–4 years, from 4–6 years, but concern is now being voiced about the quality of these graduates.

computer scientist and that there is no mandatory representation from computer science on the advisory panels (viz. TOP and UP) that set EPSRC priorities.¹⁷ The now-abolished “Subject Standing Committees” once provided a window for the EPSRC management into our discipline, but that window has now been shut. Nobody should be surprised that UK computer science research is under-appreciated and under-funded, since the people with a first-hand understanding of the subject are not at hand on a day-to-day basis to discuss the role and potential of the field within the EPSRC. The computer science community is consulted by EPSRC staff at meetings and special panels, but computer scientists are being engaged at only a small number of well-defined points in the debate and as outsiders.

Here are two examples of EPSRC policy decisions that the Panel questions; had EPSRC management been working closely with computer scientists during the discussions that led to each, different decisions are likely to have been reached.

- Welcome and important though the e-Science initiative may be, the Panel believes that its long-term success requires a scope that includes broad coverage of relevant but more speculative computer science research. The current budget for the 3-year e-Science initiative is unable to fund such efforts.¹⁸
- Work on applications of computer science is important, but there needs to be a balance with the study of core topics. EPSRC research funding opportunities and grant evaluation criteria seem to be designed to favour applications at the expense of the core, and the UK research community is following the lead. Ignoring the core jeopardises future innovation and deprives tomorrow’s researchers of the tools with which they will innovate.

In settings where a single agency is responsible for all the funding in a particular field of study, there can be a reluctance to assign control for an extended period to a single scientist from that field. Power, it is argued, would be too concentrated. The usual solution is to rotate scientists into this position of control for fixed (say, three-year) terms. By choosing senior scientists for the position, the funding agency benefits from the occupant’s broad experience and respect within the community. One might choose scientists in their prime or scientists towards the end of their careers. The latter would be particularly easy to implement today in the UK, because the mandatory retirement policy for academics produces a steady stream of candidates, many of whom are looking for new challenges and would enthusiastically embrace a programme management position at the EPSRC. The Panel therefore urges the EPSRC to reconsider its current policies and to introduce secondment as a means of staffing the IT/CS Programme with senior computer scientists.

Finally, the Panel feels that the IT/CS Programme should be split into two: one programme to fund work in computer science, and a second to fund physical layer technology. That our Panel was instructed to evaluate only the computer science portion of the IT/CS Programme is very telling. Certainly, the current organisation as a single IT/CS Programme obfuscates the level of commitment that the EPSRC makes to

¹⁷ However, at present, there are computer scientists on the EPSRC Council, TOP and UP.

¹⁸ This information was obtained from discussions with Tony Hey, Director of the e-Science initiative.

each portion and makes independent control and budgeting of the two portions difficult.¹⁹ The modus operandi and levels of funding for physical layer technology are different from those for computer science; research in physical layer technology is closer to the model of the natural sciences, whereas computer science follows a different model (see §1.1). So it is easy to imagine funding instruments and programmes that are suitable for one but not the other portion of the current IT/CS Programme. In partitioning the programme (even without augmenting either budget), the EPSRC would also send a powerful message of support to the computer science community—that computer science is being recognised as a first-class discipline—and the EPSRC would be acknowledging the reality that computer science has its own research agenda, methods, and paradigms. Partitioning the programme would also allow the EPSRC management to start thinking about computer science research at a level of the organisation befitting the strategic significance and promise that this discipline holds for all of the science and engineering fields.

4.3. Evaluating the impact of computer science research

Funded scientific research is meant to have an impact. In computer science, impact occurs through publications, software artefacts, and commercialisation. Each is a distinct avenue by which one researcher's ideas can come to affect others; the different modes are better suited for communicating different kinds of ideas.

Publications. Publications are widely used for communication in all academic disciplines. In the natural sciences, peer-reviewed journal articles play the major role, and citation counts are a good measure of a publication's significance. In computer science, conference proceedings (many of which are highly selective in accepting papers) have become an important means of conveying information, rivalling journals in some experimental sub-fields. The speed with which an idea gets into print through a conference proceedings volume is very appealing to a community where technological innovation occurs on "Internet time".

Computer science publications that have high citation counts are, by definition, influential. But a paper with a low citation count can be influential, too. Such a paper might not be widely cited because it solved a particular and important problem (allowing some research community to move in a new direction) or because the ideas in the paper have entered the lexicon and now represent commonly held views. Both phenomena can be observed in the computer science literature.

Software artefacts. Some ideas—notably in connection with experimental systems—are best conveyed by building them into a software system and distributing it. Only by seeing an idea in operation or by using it oneself (perhaps in unanticipated ways) can one come to understand and appreciate that idea. A journal or conference publication might describe a system and offer examples of its uses, but that is no substitute for using the system oneself. "Building on somebody else's ideas" is more than just a metaphor in computer science research.

¹⁹ The Panel, for example, was never able to get a good, high-level breakdown of EPSRC expenditures in the various sub-fields of computer science research.

The impact of a software artefact is not only measured by the numbers of its users. One must also note the systems whose design or construction was influenced by the artefact, directly or indirectly. For example, the key ideas found in user interfaces of today's computing systems (notably, the window and the mouse) can be traced to software artefacts that were publicly demonstrated but were never distributed and had very few users.

Commercialisations. The inertia of communities and large corporations makes it difficult for them to embrace ideas whose adoption requires radical change. Often, a new technology (3.5" disks are a classic example) is initially attractive to new customers but not to existing ones. A start-up company might then provide the best vehicle for spreading the new idea. Since the road from idea to product involves more than technical innovation, failure in the marketplace does not necessarily mean that the original idea was flawed. It is, however, difficult to argue with success—if people are paying to use an embodiment of the idea, the idea must be regarded as having had an impact.

The US economy has substantially benefited from companies started to commercialise research results in computer science. Workstations (Sun), internetworking (Cisco), LANs (3Com), user interfaces (Apple), graphics (Evans and Sutherland, Nvidia), hypertext (Netscape) and relational databases (Oracle) are only a few examples. The UK has not been as successful in capitalising on its computer science research, and it is tempting to wonder why. For one thing, over the last decade or so, computer science research in the UK has been less concerned with experimental systems (and some reasons for this have already been suggested). But many factors affect how easy it is to start a successful company in a particular country, and commercialisation ought to be promoted as a viable path by which UK computer science research can make an impact.

Metrics. Various metrics are employed in evaluating the success of research investments. Computer science is different enough from the other disciplines under the auspices of the EPSRC that the choice of metrics for evaluating computer science research requires careful consideration. For example, numbers of computer science journal publications reveal little about an investigator's research impact, and citation counts give an incomplete story about a paper's impact. Ideally, a set of metrics should accomplish two things:

1. provide some measure that is correlated with research impact; and
2. cause researchers who optimise for that metric to engage in desirable behaviours.

Because of this second effect, the Panel cautions the EPSRC and others to choose their metrics carefully. Research impact is the desirable behaviour to be encouraged. Researchers should publish results, not to increase their publication counts but because that is the avenue that will have an impact (because the results are significant and the publication allows the community easy access to them). Similarly, the creation and dissemination of software artefacts should be available and appreciated as a mode of dissemination (which has implications about types of funding).

It is also important that research programmes be evaluated by the EPSRC as a whole and not on an individual, project-by-project basis. A research programme should be

managed to encourage risk-taking, even though that lowers the chance of success for each individual project. When it is the individual projects that are being evaluated, investigators are driven to adopt conservative goals that they are confident can be achieved. But this reduces the chances of breakthroughs. For example, if each of 30 projects in a portfolio has a 5% chance of success, then one can expect one or two triumphs; 30 projects with the conservative goals that would enable a 95% chance of success for each are unlikely to produce any breakthroughs.

5. QUALITY OF UK COMPUTER SCIENCE RESEARCH

In a number of areas, the UK is a world-wide leader, demonstrating an outstanding record of innovation and first-rate science. This is best exemplified by the number of UK researchers who have received the ACM Turing Award, the field's equivalent of a Nobel Prize. Since the award's inception in 1966, there have been five UK recipients (Maurice V. Wilkes, J. H. Wilkinson, Dana S. Scott, C. Anthony R. Hoare, and Robin Milner), and that equals the total number of Turing Awards made to researchers in all the rest of Europe and Asia.

Two historically important areas of UK computer science research strength and long-standing impact are as follows.

- **Semantics, logic, and formal methods.** Programs prescribe sequences of instructions for computers to execute. We hope that these sequences compute what the programmer intended, whether that is predicting the trajectory of some sub-atomic particle, rendering an image, or translating actions in a cockpit into signals that move the rudder and elevators of an aeroplane. The field of semantics is concerned with establishing the correspondence between what a program computes and the program text itself. This gives a basis to understand, using logic and methods based on logic, whether a program implements its specification and whether one program (such as the input to a compiler) is equivalent to another (the machine-language output of that compiler). UK researchers wrote the book on semantics, having introduced denotational semantics, axiomatic semantics, and process algebras such as CSP and CCS. So-called “Hoare logic” and “higher-order logic” pioneered the use of logic for reasoning about programs and hardware.
- **Programming languages.** The language in which a program is written affects how computations are described. A well-designed programming language gives enormous leverage to the programmer by providing powerful tools for designing suitable abstractions (so that irrelevant details can be ignored) and by supporting automated analysis to detect inconsistent uses of variables and operations. Building from a strength in semantics and logic, UK researchers have been responsible for developing new languages, like BCPL, ML, Occam, Lotus, and Haskell, which are widely used and whose influence is felt in virtually every modern programming language. For example, BCPL (developed in Cambridge around 1970) is the basis for C, the most widely used language in the world today; and Haskell (designed collaboratively at Glasgow and Yale around 1990) is currently probably the leading functional language. UK researchers have also developed control structures—such as monitors and synchronous message-passing—and type systems that have been widely instantiated in programming languages in general use.

The discipline of computer science today is quite broad. What follows merges individual Panel members' perceptions about various specific areas of UK computer science research. The UK's strength and leadership are apparent throughout, and the Panel believes that a comprehensive survey (like the current RAE) would affirm the views given below. However, the Panel did find evidence of some imbalance, and it therefore recommends that research efforts be expanded in algorithms and in experimental computer systems. With a few notable exceptions, algorithms and complexity theory have never been emphasised by UK researchers, and the Panel feels this will increasingly become a liability as computer science techniques become pervasive in the natural sciences, the social sciences, and the humanities. Also, whereas the UK has been a world leader in experimental computer systems research, the Panel perceives that strength now to be eroding.

Comments on some specific broad areas of study

Theory. Algorithms are at the heart of computing. Any system, be it software or hardware, will have algorithms (and data structures) as its basic building blocks. Web search-engines, routing in the Internet, caching, genome analysis, cryptography, CAD for circuit design, mechanical engineering and architecture, image analysis, and optimisation tools, all depend critically on suitable choices of algorithms and data structures. The field of algorithms is today an extremely important enabling technology, and, if anything, its importance is increasing.

The UK has made small but extremely significant contributions to the field of algorithms and complexity theory. Currently, Warwick is notable for covering the full breadth of the subject. Here, outstanding contributions have been made to circuit complexity, computational geometry, distributed algorithms, graph algorithms, and algorithmic bio-informatics. Other UK departments have made key contributions to specific parts of the field. Examples include seminal work on rapidly mixing Markov chains in Edinburgh, work in computer algebra in Bath, and work on software libraries for parallel computing in Oxford.

Outside the UK, universities with credible computer science research efforts tend to have substantial and broad research activities in algorithms and complexity theory.²⁰ In the US, research in algorithms and complexity theory dominates other forms of theory; in the rest of Europe, work in theory is evenly split between that and work in semantics and logic. Consequently, UK computer science graduates are less knowledgeable than their international peers about current developments in the design, analysis, and implementation of algorithms. The Panel believes that the overall UK effort in algorithms is too small in absolute terms (when viewed against international peers), and also when compared to UK research effort devoted to semantics, logics, and types. Failure to strengthen work in the algorithms area may mean that the UK will lose competitiveness for a wide range of applications, including information retrieval, medical applications of computing, scientific computing, and so on.

²⁰ Examples include the top US Computer Science departments (Stanford, MIT, Carnegie-Mellon, University of California at Berkeley, Cornell, Princeton, University of Pennsylvania), as well as top departments elsewhere (for example, Saarbrücken and just about every German university, the Weizmann Institute and just about every Israeli university).

With regard to research in semantics, logic, verification, and the theory of programming languages, the UK is without peer. These areas will continue to be of central importance to computer science, and there is no doubt that UK research can and should continue to play a leading role in them. Thus, this research should continue to be supported.

Quantum information processing is a new paradigm, which proposes to exploit quantum mechanical effects for computing and information transfer. Already well established in the UK (much better than, for example, in Germany)—with active and highly visible groups in Bristol and Oxford, among other places—it has fantastic potential. For instance, efficient quantum algorithms exist for factoring and computing discrete logarithms, which enables efficient attacks against RSA and most of the other public key cryptography currently used on the Internet, but there also exist quantum cryptographic methods for unconditionally secure information transfer (that is, secure by virtue of the laws of physics), in contrast to the cryptographic algorithms being used today, which remain secure only until significantly faster means of computing are developed.

Realising the promise of quantum computing requires two things: (1) showing that quantum computers are actually useful by exhibiting tasks that they can do significantly better or faster than classical computers, and (2) actually building a quantum computer. Computer science research in quantum computing concentrates on the former, requiring close and intensive interaction with physics and mathematics. The goals here include the development and analysis of new algorithms for quantum computers, new forms of communication and cryptography that make use of quantum mechanics, and a quantum information theory, along with methods for error correction in quantum computers.

Software and systems. Research in experimental systems is concerned with the design, implementation and evaluation of computer systems. In the past, research has centred around operating systems, computer networks, and distributed systems, but nearly every area of computer science (including databases, multimedia, programming systems, and networks of embedded devices) has the potential to involve this sort of experimental work. The research is often opportunistic in nature, leveraging shifting tradeoffs involving the capabilities of the underlying technologies. Once an emerging technology trend or user need is recognised, the researcher will propose new system designs, implement prototypes, and experiment with those prototypes in an effort to evaluate the performance and functionality of the designs. Ideally, the process identifies implementation techniques that apply across a wide range of systems and technologies. The process also has an important side effect: the research prototype itself. Such a prototype can be transferred to industry for commercialisation.

The UK has historically been a world leader in systems research, with early work at Cambridge and at UCL leading the way in operating systems, distributed systems, and computer networks research. In fact, UCL, was the first European node on the Arpanet (which ultimately evolved into today's Internet). However, this strong position in systems has eroded, with particularly acute losses over the last three to four years—industrial demand (including start-up ventures) and the retirement of senior researchers have drained the research leadership in this area from UK universities. Today,

the Panel can point to research groups in only a handful of universities (Cambridge, Lancaster, and Newcastle, for example) that undertake leading-edge systems research. The Panel cautions that the decline in UK systems research will continue unless steps are taken to improve the environment for conducting systems research, a software-intensive activity requiring funding for both long-running projects and significant support staff—issues discussed above in §3.2.

One enduring UK strength in the general software and systems area concerns research into technology associated with achieving software system trustworthiness. The work in semantics, logic, and formal methods can be seen as addressing the problem of programming errors. A trustworthy system must also tolerate malicious users and hardware failures. Moreover, for deployment in settings that involve sensing and controlling the physical world, the operation of these systems must satisfy real-time constraints. In the UK, one can today find world-class research groups working separately on each of these elements, as well as investigating how the elements interact.

Computer security is a broad subject, ranging from physical security, control of electromagnetic emissions, and cryptography, through access control, authentication mechanisms, and software assurance, to legal and social mechanisms for discouraging or punishing misbehaviour. Historically, research in this area has been dominated by the military's need to control classified information. This is now changing, because as the world becomes ever more networked, malicious people can do more harm, and security measures that can limit this harm become more important. The UK has one of the leading academic research groups in security, at Cambridge. It also has people, especially at Oxford, who use formal methods to validate the security of mobile code and of communication protocols.

The Dependability Group at the University of Newcastle upon Tyne pioneered the view that the various elements of trustworthiness need to be considered together; this group has remained a leader in the field, bringing together researchers interested in fault-tolerance, security, and operating systems. The influence of Newcastle is felt, not only through its technical contributions but also through its graduates, who have themselves assumed leadership positions in the field. Specific system constructs developed and studied at Newcastle include recovery blocks and various replication management protocols.

Finally, the University of York has one of the world's leading research groups in real-time and embedded systems. These researchers are credited with a significant body of research in static real-time scheduling theory. They have also demonstrated how to employ these theoretical results in practice, by accounting for networking and operating-system overheads. This combination of theory and practice has resulted in important and practical applications of their work. The group is strong in applying formal methods to the design and analysis of safety-critical systems.

Software engineering. The UK has a strong record of research contributions to software engineering. Research on the more formal and theoretical aspects of software development has traditionally been centred in Europe, and the UK has played a major role in this, building on its strength in semantics and logics. As examples, UK researchers have been strongly involved in the definition of specification techniques and

languages, such as Z (developed at Oxford) and VDM (from Manchester), which are now used world-wide. UK researchers proposed the first programming logics, and UK researchers at City University are the leaders in scientifically-based work on metrics and measurements for software products and processes, and for software reliability analysis.

On the requirements-engineering side of software engineering, UK researchers (for example, at Imperial and UCL) are responsible for leading work in inconsistency management, with their techniques for capturing multiple viewpoints, consistency checking, and reasoning about requirements in spite of inconsistency. Over the last decade, there has also been much UK work on specification reuse (at City and UMIST), scenario-based requirements validation (also at City and UMIST), and reasoning about responsibilities for system components (at Imperial and KCL).

The UK has one of the strongest groups world-wide in software architecture and distributed software engineering (at Imperial), based on work in techniques for describing, analysing, animating, and dynamically reconfiguring software architectures. UK researchers are also well known for their contributions to software process engineering (at Manchester) and software re-engineering (at Durham). And UK researchers (at Lancaster) are leaders in middleware.

The leadership position held by UK researchers in software engineering should not be allowed to erode. All research in software, as well as commercial software development in the UK, benefits thereby.

Databases, knowledge management and information retrieval. With regard to databases and knowledge management, the UK has for several decades shown considerable strength in research involving models and languages (for example, PFDM at Aberdeen, and DOOD at Heriot-Watt), and has had an international impact by virtue of pioneering work on persistence in programming languages (for example, PS-Algol, Napier and PJAMA at Glasgow and St. Andrews). More recent pioneering work in description logics and ontologies pertaining to bio-informatics is being pursued at Manchester, and work in hypermedia systems is being done at Southampton. However, these contributions have primarily been at the periphery of the subject, rather than at the core. Notably missing are significant efforts in experimental systems building and work in new database paradigms (for example, ubiquitous information).

Nevertheless, the UK has traditionally been a leader in information retrieval, with some of the founding researchers. Today, the bulk of the activity in that field has shifted to the United States, but UK text retrieval research is still internationally first-rate. The idea of statistical tagging and parsing was first explored at Lancaster, and this is now being pursued at Sussex. The research frontier in this field has moved to image and sound retrieval, and away from text handling, and the group at Cambridge has a leading project in speech (sound) searching. However, by and large, there are only scattered pockets of researchers in the UK, tracking the new centre of gravity of research in this area: the US. This needs to be thoughtfully addressed if the UK is to benefit from the many new opportunities emerging in this field, with its immense impact on all aspects of bulk information handling.

Artificial intelligence (AI). The UK has been a major contributor to the development of

artificial intelligence (which we take to include natural language and speech processing, computer vision, artificial neural networks and related aspects of cognitive science, and neuroscience) from the start in the 1960s. The first AI departments in Europe were housed in the UK, and London hosted the second International Joint Conference on Artificial Intelligence (IJCAI) in 1971. Although the UK no longer dominates the AI scene in Europe as it did in the 1970s and early 1980s, this is largely due to growth elsewhere—especially in Germany, France, and Italy.

The UK is a world leader in speech engineering and computational linguistics. The HTK speech recognition system (developed at Cambridge) is consistently ranked among the world's best on standard benchmarks; it was also the basis for the creation of Entropic Ltd., acquired by Microsoft. The Festival speech synthesis system (from Edinburgh) is widely used internationally and has now been commercialised by Rhetorical Systems. Bayesian methods for natural language processing and data mining are the basis for the offerings of Autonomy, another UK university spin-out, and Sheffield's toolkit for information extraction, GATE, has been widely distributed. Strong groups developing both symbolic and statistical approaches to natural language processing and its applications to interactive systems, information retrieval, summarisation, and machine translation can be found in Brighton, Cambridge, Edinburgh, Essex, Sheffield, Sussex, and UMIST.

Major contributions to machine learning have come from the UK, notably the early work by Michie at Edinburgh, on learning control and on structured induction. Inductive logic programming originated in early work by Edinburgh researchers, and UK researchers continue to dominate. UK researchers have also made major contributions to the field of artificial neural networks, with commercial developments including Integral Solutions' Clementine software.

In the field of computer vision, leading work is being done in the areas of medical imaging (Oxford and Imperial are among the leaders), and work in active vision was recently being pursued at Oxford, though recent defections to industry may have an impact on those efforts.

Finally, UK groups are world leaders in certain aspects of automated reasoning, including theorem-proving, spatial reasoning, search methods for the solution of combinatorial problems, and planning and scheduling.

Computer architectures. Computer architecture research in the UK has dramatically declined in the last decade, from its historically competitive levels. This problem is not unique to the UK—international research in the area has generally narrowed in both range and focus. Nevertheless, the Panel does feel that opportunities for research into innovative architectures abound in the areas of communication and media processing. Given that the UK has produced leading computer architectures, from the earliest days of electronic computing (Colossus and ACE), through the mainframe era (Manchester), and into the microprocessor era (Transputers and ARM), it could conceivably continue at the forefront of computer architecture design for decades to come. It is not healthy for either UK computer science research or the related computer industries simply to rest on the success of these important, but now historic, machine architectures.

Internationally, communication and network processor research is focused on the increasing amount of computing required to route packet traffic across the Internet at data rates reaching into the tens of gigabits per second. Computer architectures for media are required to address the coding, decoding, and transcoding of the streams of video, audio, and graphic data that will allow the Internet to become an increasingly compelling medium for dynamic audio and visual content. In both cases, the choice of instruction set and the internal organisation of machine resources (such as arithmetic elements, register files, datapaths and cache memories) are dramatically different from those of general-purpose machines. The UK's strong industrial activity in communications and the growing presence of electronic media companies suggest that computer architecture research in these areas would have a rapid uptake and a timely impact in these national industries.

There are other, more fundamental, aspects of computer architecture that are driving international research. First is the increasing use of parallelism, both implicit and explicit, at the hardware thread level within both general- and special-purpose architectures. Second is the rapidly growing focus on power-aware architectures that are sensitive to the environmental limitations imposed by battery-operated computers, which underlie both mobile computing and the anticipated pervasive computing environment. Power consumption and heat output will be critical factors in the long-term success of these areas of information technology.

The quality of UK computer architecture research can and should be substantially improved, as evidenced by the low level of UK participation at international conferences, the apparent high rejection rate for the UK papers being submitted, and various other points that emerged during direct discussions with academic and research staff in this area. Two factors contribute to the weakness. First, the access to and the use of internationally accepted research tools is quite limited. Second, the use of locally created performance benchmarks, rather than well-understood international benchmarks, makes it difficult for non-UK peer reviewers to compare UK researchers' results directly with those of the international community. Internationally, architecture research has become intensely quantitative over the last decade; UK researchers must now move quickly to embrace the new quantitative measurement frameworks.

A few UK companies are succeeding with their own computer architectures (ARM and STM), and there are many talented researchers within the university community who are currently working in related, primarily applied, fields. In addition, several of the country's leading industrial computer architects are now on the staffs of top universities (for example, Bristol and Manchester). This suggests that it would be relatively straightforward to restore research in computer architecture to internationally competitive levels within a short period of time. Because computer architectures are now viewed as distinct intellectual property, quite separate from their implementations, it is plausible for the UK to become an international leader in the field. Indeed, the success of the ARM processor is proof positive that an internationally successful computer architecture does not require a strong national computer industry.

Human-computer interaction (HCI). HCI is important for a wide variety of information technologies, in which humans use computing devices or communicate with one another using information technology as a facilitator. It is one of the most rapidly

growing areas in computer science. Major areas of HCI research include natural language interfaces, collaboration and conferencing systems, information search and utilisation, human capabilities and performance, software system interfaces and toolkits, and graphics and visualisation.

Historically, the UK has been strong in many areas of HCI, both because of UK expertise in physical ergonomics and psychology, and because of a few strong research groups. The leading HCI conference is arguably the ACM CHI conference, held annually in North America; papers from that region dominate. Although most published papers originate in the US, the UK contends for second place with Canada; 7% of the published papers are from UK authors. First-class HCI research has been done at the Xerox Europarc lab and at UK universities (including Imperial, UCL, Loughborough, Glasgow, and UMIST).

The UK's strengths in HCI are diverse. Groups at Edinburgh and Cambridge are widely respected for their work in natural language interfaces, especially speech recognition and synthesis. A group at Lancaster could reasonably claim to have originated statistical tagging and parsing; the area has since been taken over by US researchers, but there is also a strong, narrowly focused, group at Cambridge. Cambridge has also been influential in collaboration and conferencing systems, as have Imperial College and UCL. The UK research on this topic is noteworthy for its effort to be scientific, rather than simply writing and hyping software, regardless of its utility. Research in human capabilities at York and elsewhere is also particularly noteworthy, rivalling that in the US. This area is closely tied to psychology, and the collaboration with psychologists is an important UK strength.

The UK was present at the creation of research in software system interfaces and toolkits; a UK researcher is credited with the first user interface-management system. Today, most of this research is tied to software systems research, in which the US dominates. The US also leads in graphics and visualisation research, discussed in the next section.

In summary, the UK had an early start in HCI but has lost ground in some areas and has neglected to build expertise in others. Even the current strengths will not persist without a broad, well-supported research base. The importance of HCI—both as a research area and as a component in systems with which humans interact—leads the Panel to urge that the area be well supported. Advances in collaboration, natural language processing, and interfaces could lead to major industrial activities and significant societal benefit.

Graphics and visualisation. The fields of graphics and visualisation have a broad research agenda, motivated both by realism and by artistic expression, as support for the understanding of information and as entertainment.

Computer graphics in the UK has largely been associated with practical applications, involving various user groups. Thus, for example, as a result of research in graphics and CAD in the 1970s, there was a significant growth of CAD and animation companies around Cambridge. The UK graphics community has also been active in developing standards such as GKS, CGM, PHIGS, and CGI, and in refining the original Virtual Reality Modelling Language (VRML) specification for promulgation as an ISO/IEC standard.

Work at Sussex and Bath in the 1970s established a group for rendering 3D scenes in real time, leading to the full-colour raster operation chip for the Atari Transputer Workstation. Sussex were also significant partners in developing the SPIRIT Workstation.

Virtual reality (VR) has matured to the point where a deeper understanding of the principles of its application is required. The work at UCL is notable here, involving interaction between real users and virtual ones, the sense of immersion, and evaluating what makes VR effective. Recent developments have made complex real-time virtual realities possible. A team at Manchester is the first both to understand the need to integrate a VR kernel into networks of workstations, and to provide a solid implementation (Maverik, available over the Web). Work in Nottingham is focused on collaborative virtual environments.

Leadership in the presentation of data and the results of simulation, as well as the use of animations and graphic models for data, has come from the US, with its strong links to the entertainment industry and the large-scale computing industry.

The UK is the acknowledged world's leading producer of computer games—a clear indication of the opportunities offered in new market sectors if the appropriate skills are available. Given the success that has already been achieved, and also the continuing need for physical representations of all forms of information, there is good reason for the UK to expand its funding of fundamental research in the graphics field.

High-performance and scientific computing. The UK has a history of contributions to high-performance computing, particularly as it is used for science and engineering research. Numerical algorithms and libraries have been an internationally recognised strength, although that work has not always taken place within computer science departments. UK researchers have participated in the international high-performance Fortran effort, in making fundamental contributions to distributed computing, and in the now-standard Message Passing Interface for communication among machines.

More recently, research in the computer science aspects of high-performance scientific computing appears to have diminished. There are scattered research efforts in a few academic departments, notably work on load balancing, benchmarking, routing, and related issues concerning performance. There is also recent work on the bulk-synchronous approach to parallelism. However, there appear to be neither the collaborative efforts nor the research platforms required to tackle the topics of the day—questions such as the appropriate system software to manage clusters of off-the-shelf multi-processors, the appropriate programming models and tools to build high-performance applications on those systems, the middleware needed to exploit the vision of grid-computing, the methods for organising, managing, and analysing huge collections of diversely represented heterogeneous observational and experimental data, and so forth.

Having so little research to build on is particularly unfortunate because of the opportunity presented by the recently announced e-Science Programme, which will provide a fertile source of real-world computational challenges, not only for computational biologists, computational chemists, and computational physicists, but for the advancement of the infrastructure and tools that constitute their electronic

laboratory. The UK could attempt to meet e-Science goals by using the fruits of computer science research done abroad, but the more sensible course would be to use local talent and enhance the local research expertise in this important area.

Bio-informatics. An extremely promising area world-wide, bio-informatics encompasses the management, visualisation, and analysis of biological data. It poses challenges in many aspects of computer science (including, for example, algorithms, databases, software engineering, machine learning, and knowledge representation), requiring computer scientists to work in close collaboration with researchers in biology and statistics. It is widely believed that expertise in bio-informatics will be crucial to the way in which biology will be performed in the future.

The UK has taken a leading role in bio-informatics, notably through the world-class research being performed at the Sanger Centre, where researchers have made fundamental advances in areas such as databases for semi-structured and incomplete data (AceDB) and algorithms for similarity searching and feature recognition. Manchester University has led an international effort to construct ontologies for biological data, as a crucial step in integrating multiple heterogeneous data sources world-wide (TAMBIS). Another very strong project involves the combination of engineering, computer science and medicine at Imperial College, London, to build real-time MRI systems that diagnose cardiac problems. The attempt by this group to use eye tracking to guide computer learning for image analysis is unique and quite promising. Other bio-informatics efforts are just getting started at places throughout the UK, including Edinburgh, Imperial College, UCL, and Oxford.

This is clearly an area in which the UK is well positioned, thanks partly to its leading researchers in biological research (such as the MRC's Laboratory for Molecular Biology, with its multiple Nobel prizes). Moreover, this is an area where a strong computer science base will leverage other disciplines, presenting opportunities for broader sources of support.

6. CONCLUSIONS

Although nominally about academic computer science research, this report is really about opportunities—understanding them, and then leveraging existing strengths in order to best exploit the opportunities. The message is simple, and it involves three themes.

First, computer science is a new kind of discipline. It differs in both character and culture from the pure and applied sciences, engineering, and mathematics. Traditional methods of supporting and evaluating research are not always appropriate for this new field.

Second, not only does computer science offer deep intellectual challenges, but research and expertise in this field will almost certainly translate into competitive advantage and economic well-being for nations, as it has in the past. Strength in computer science will be imperative for the UK in the 21st century.

And third, the UK must change the environment in which academic computer science research is being conducted. It must increase salaries and the level of support for research in computer science, change the vehicles used to fund that research, allow computer scientists to play a more active role in defining and managing the nation's computer science research programmes, and encourage growth and strength in two specific areas of study (algorithms and experimental systems). The UK has been a world leader in computer science research, but this position is currently being allowed to erode and will continue to do so unless changes are made.

7. ACKNOWLEDGEMENTS

The panel is grateful to the Steering Group for its excellent preparatory work and for its assistance throughout the review. The Panel is also very appreciative of all the hard work done willingly and well by members of the secretariat—one could not imagine better scientific and logistical support. Sue Rodd is acknowledged for her assistance in editing this report. The Panel gratefully acknowledges the time and effort invested by Vince Osgood and Dr. Jim Fleming on behalf of EPSRC, Ian Wand on behalf of the 2001 RAE, and Cliff Jones on behalf of the UK Computing Research Committee, in briefing and in answering follow-up questions. Finally, the Panel would also like to express its sincere appreciation to the universities visited, where both staff and students invested considerable efforts in ensuring that the visits were informative and enjoyable.

APPENDICES:

A. DATA SUPPORTING THE REVIEW

A.1. Main data document

All members of the International Panel were provided with documentation compiled especially for the Review. This included the following information:

- Title and Terms of Reference of the Review;
- an overview paper on the funding of British university departments, including salary scales for 2000/2001 and 1996 Research Assessment Exercise (RAE) ratings of computer science institutions;
- an article on the EPSRC IT and Computer Science Programme;
- UK Computer Science Research: Vision and Opportunities (a submission specifically produced for the Review by the UK Computing Science Research Committee (UKCRC));
- the UKCRC input into the Quinquennial Review of the Research Councils;
- overview of the EPSRC Balance of Programme Exercise;
- EPSRC Business Planning Processes—Information Technology and Computer Science;
- the combined comments from the Technical Opportunities Panel (TOP) and the User Panel (UP) that were used in formulating the 2000 plan for the Information Technology and Computer Science Programme;
- business plan for the EPSRC Information Technology and Computer Science Programme for 2000;
- EPSRC grant data;
- EPSRC contextual data;
- list of computer scientists working in the UK who are Fellows of the Royal Society (FRS), Fellows of the Royal Academy of Engineering (FREng) or Fellows of the British Academy (FBA);
- Higher Education Statistics Agency (HESA) data showing qualifiers in computer science by level of qualification and institution for each year from 1995 through to 2000;
- a report on International Perceptions of UK Research in Physics and Astronomy;
- bibliometrics.

A.2. Presentations

The Panel was given presentations, followed by discussion/question sessions, on the following topics:

- an overview of the funding of British university departments (an elaboration on the aforementioned overview paper listed in A.1.);
- an overview of research funding in the UK, including an insight into industrial funding;

- a briefing on the computer science part of the EPSRC IT and Computer Science Programme, including:
 - balance of programmes,
 - evaluation,
 - business planning,
 - how decisions are made on grant proposals including specific examples of proposals that just failed to receive funding,
 - e-Science,
 - computer science in the humanities and social science;
- a briefing on specific parts of the IT and CS Programme—namely, image and vision computing and distributed information management;
- a briefing on the Research Assessment Exercise (RAE) 2001;
- a briefing on the UK Computing Research Committee's submission to the Panel (dealing with specific issues related to the aforementioned report in A.1.).

In addition to copies of the slides from the presentations, the following information was also produced after the presentations, at the request of the Panel:

- job description of the EPSRC Programme Manager for Information Technology and Computer Science;
- the 1999 and 2000 annual reports of the EPSRC Information and Technology and Computer Science Evaluation Panels;
- a copy of the EPSRC Individual Grant Review (IGR) Form;
- a copy of the EPSRC IGR Assessor's Form;
- summaries of the 1999 and 2000 Quality, People, Impact and Exploitability (QIPE) scores for the Mathematics and Physics programmes.

A.3. Bibliometrics

The bibliometrics comprised a spreadsheet showing papers indexed by the INSPEC (Information Services for Physics and Engineering Communities) database in specific journals for the period 1995 to 2001, broken down by country of domicile of author. INSPEC is recognised as the leading English-language supplier of services, providing access to the published information in the fields of physics, electronics, and computing. The INSPEC database currently contains records for nearly seven million scientific and technical papers. These are being added to at a rate of over 350,000 per year.

A.4. Site visits

In order to provide contextual information to support the Review, the Panel members visited eight universities. The Panel was split into four subgroups—two consisting of four members, and two consisting of three. Each subgroup visited two universities. The universities invited to participate were selected to illustrate substantial research departments (RAE grades 4, 5 and 5*) and to illustrate the regional diversity of research departments.

To support the site visits, Panel members were provided with complete programmes by the universities, with Web-links to relevant information for each university department. The universities were encouraged:

- to provide an insight into all computer science research work within the university as a whole, and to put it within the context of the UK/international scene more generally;
- to address particular issues (for example, recruitment, retention, links with other university departments, and links with industry);
- to allow Panel members to enter into discussions and ask questions, and to meet with a range of people, including lecturers, research assistants, and students.

B. TERMS OF REFERENCE

The brief given to the Panel was:

To report on the standing and potential of computer science research* in the UK and provide a comparison with computer science research internationally.

- a) The study will cover all computer science research undertaken in UK universities.
- b) The study will assess the current standing of computer science research and its likely impact on the UK science base, and on the nation's wealth and well-being.
- c) The study will draw on existing data, international assessments and site visits.
- d) The study will result in a report to be published and made available to the sponsors and the UK computer science community.

* The Research Assessment Exercise 2001 (UoA 25) uses the following definition of computer science.

For the purposes of this study, the coverage of computer science includes the theoretical and practical study of the following: adaptive systems, algorithms, artificial intelligence, computer architecture and engineering, computer graphics, computer vision, databases, dependable systems, distributed systems, formal methods, high performance computing, human-computer interaction, information retrieval, information systems, machine learning, multimedia, networks and communications, operating systems, pattern recognition, programming languages, software engineering, speech and language technology.