UK COMPUTING RESEARCH COMMITTEE

UKCRC Submission to the EPSRC International Review

September 2006

This document presents the views of the UKCRC on a number of different areas that we believe are germane to the International Review of ICT.

The UK Computing Research Committee (UKCRC)[1] is a body of internationally recognized researchers in Computer Science. It is an expert panel of the professional institutes, the British Computer Society (BCS) and the Institution of Engineering Technology (IET), and also of the Committee of Professors and Heads of Computing (CPHC), which represents the computing departments of UK Universities. It comments on a wide range of issues that are relevant to the UK research community and has been instrumental in establishing a series of UK Grand Challenges in Computing which set out a number of long-term goals for challenging research in the discipline. Currently, members of the UKCRC are drawn from over 30 research-active universities and industrial research institutions.

The UKCRC has, so far, made 29 different submissions and reports on a range of topics. These submissions are available from: http://www.ukcrc.org.uk/resource/reports/index.cfm

This submission relates to that part of the EPSRC ICT Review concerned with computer science (broadly interpreted to include theoretical and practical aspects of the discipline) rather than devices or communications hardware.

This submission is in two parts:

- **UK Universities Computing Research: Current Issues and Research Landscape**
  In this part, we discuss a number of current issues that are relevant to the review and present a 'research landscape' – an overview of some of the principal research activity in computer science in the UK.
- **UK Universities Computing Research: From the 2001 International Review to 2006**
  In this part, we look back at the conclusions of the 2001 international review and discuss a range of factors that have influenced computer science research since 2001.

There is inevitably some overlap between these documents but, in general, they offer different perspectives on UK computer science research (the first part concentrates on the present and future, the second part concentrates on the recent past).

The submission covers a broader range of issues than that set out in the review framework. However, to relate this submission to the published review framework, the following table points to the sections of each document that wholly or partially address the framework questions. In the table, we refer to the current issues document as Part A, and the 2001 to2006 document as Part B.

Framework question	Document section
1. To what extent is the ICT community addressing key technological/societal challenges and engaging in new research opportunities?	Part A: 5 Part B: 6
2. To what extent is the ICT research base contributing to other disciplines and multidisciplinary research?	Part A: 2.3, 5
3. What is the level of knowledge exchange between the research base and industry that is of benefit to both sides?	Part A: 4 Part B: 7
4. To what extent is the UK ICT research activity focussed to benefit the UK economy and global competitiveness?	Part A: 4 Part B: 6.2
5. To what extent is the UK able to attract young scientists and engineers into research, nurture and support them at every stage of their career to benefit the UK?	Part A: 2.4, 2.5, 2.9, 3 Part B: 4.3
6. To what extent is the UK able to attract and retain overseas scientists and engineers to the UK?	Part A: 2.4, 3 Part B: 5
7. What is the impact on a global scale of the UK ICT research community both in terms of research quality and the profile of researchers?	Part A: 5
8. What evidence is there to support the existence of a creative and adventurous research base and portfolio?	Part A: 1,2,5 Part B: 3, 6.2

[1] http://www.ukcrc.org.uk

UK Universities Computing Research: Current Issues and Research Landscape

Since the previous international review of computer science in 2001, the UK computer science research community has flourished and continues to make research contributions at the highest international level. Notably, in the 2001 research assessment exercise, the number of computer science units performing at the highest levels (5 and 5*) increased significantly, reflecting the fact that more departments were achieving an international standard of excellence.

In 2001, the community was emerging from an extended period of limited funding and government initiatives to support science in general have meant that EPSRC funding for computer science has increased. Initiatives such as the creation of inter-disciplinary research collaborations have been immensely successful as we discuss later in this document and computer science researchers have benefited from very large sums devoted to e-science. A generally more benign university funding environment has reduced some of the pressures on academics in general although computer science departments have (mostly) had to cope with falling numbers of undergraduate students.

In the other part of the submission, we examine the factors that have influenced computer science since 2001 so, in this part, we focus on a range of issues concerning interaction with the EPSRC that have been raised by members of the computer science research community. We also provide some more general background information for the panel about the funding of research in the UK.

A key component of this part of the submission is the 'Research Landscape' (section 5) where we discuss the contributions of the UK community in a number of areas of research strength. Over the last five years research in Computer Science has grown in breadth and depth. It has become more inter-disciplinary (e.g. joint research with life sciences, medicine, electronic engineering, psychology, sociology and arts) and more applied (e.g. aerospace engineering, games and creative media, healthcare) whilst making significant progress on fundamental, core research in areas like networks, distributed systems, human computer interaction, algorithms, foundations, and formal reasoning. Two new themes running throughout all the research topics are ubiquitous, or pervasive, computing, and the internet and internet applications. Both have gained significance dramatically over the last five years. Not only are they the direct result of earlier Computer Science research (e.g. internet search engines are based on early information retrieval research), but they but they also present new, exciting research challenges for the whole research community. This summary, of course, cannot cover all high-quality research and, in addition to the work and institutions discussed here, there are many individuals who have an international reputation for their work in different areas. Space has not allowed us to cover these here.

Overall, the UKCRC's view is that the EPSRC has been generally supportive of computer science research and the EPSRC has played an important role in the improvement of the research environment. Interactions with members of EPSRC staff have generally been constructive and many UKCRC members have commented on the support provided by individual EPSRC staff members. We believe that the EPSRC and the community work together effectively to maintain the high standard of UK computer science research.

1 Background

In this section, we present background information about a number of areas which affect the health of the UK computer science research community.

1.1 Dual funding

Research in the UK operates according to a dual funding system with the principal national funders being the Research Councils and the Higher Education Funding Council (HEFCE). Scotland, Wales and Ireland have equivalents to HEFCE and there are distinct differences in policy. However, these distinctions are not significant here. For some disciplines (but not computer science), charity funding is significant.

The Research Councils fund research according to competitive proposals and do not cover the full costs of research. They currently cover 80% of the full economic costs with the remainder covered by the HEFCE research funding to the university. The majority of HEFCE research funding (so-called QR funding) is distributed formulaically and, for each department, the QR funding is a function of its RAE rating (see section 1.3).

In the past few years, long-term underspending on research infrastructure has been recognised as a serious problem and HEFCE have provided universities with earmarked capital funding (so-called SRIF funds). Computer science has benefited as have all science disciplines from this with benefits including new buildings, labs, etc.

1.2 The Research Assessment Exercise

The Research Assessment Exercise (RAE) was first introduced in the 1980s as a means of judging the quality of research going on in UK Universities. Initially, this was simply an information exercise but, since 1992, the results of the RAE have been used in a formulaic way to compute the amount of discretionary research funding (i.e. funding that is not associated with specific projects) allocated to different subject areas in a University. However, each University may redistribute their overall research funding according to local priorities so the research support assigned to computer science may be more or less than that calculated by the funding council formula.

In the last RAE, units of assessment (broadly a discipline such as computer science) were assessed on a 7 point scale and units on the top 3 points on this scale (called, confusingly 4, 5, and 5*) were allocated significant research funding. The amount of funding concerned is very significant and the allocation algorithm is non-linear, rewarding the units with the highest level of research rating.

The next RAE will take place in 2008 and assessment panels have been set up. A somewhat different approach to assessment will be used so that units will be graded on a continuous rather than a discrete scale. Until recently, the general principle that research funding would be assessed formulaically depending on RAE rating was maintained. However, in a recent government announcement, it was stated that research funding after 2008 would be simplified and based on metrics which have still to be agreed. The 2008 RAE is still going ahead but it is unclear to what extent their conclusions will be used to inform the long-term research funding process.

It is probably the case that, in the early years of the RAE, it had a positive effect on research as it helped focus the attention of the community on research quality. However, as in all such exercises, the community rapidly learned to 'play the assessment game' and changed its behaviour to reflect the needs and priorities of the RAE. Primarily, this requires demonstrable outputs in the assessment period so there has been a move away from long-term, adventurous research towards shorter-term, publishable research. Furthermore, there is less to gain in RAE terms from systems projects that require significant effort in 'platform engineering' and so this type of research is being carried out in fewer and fewer institutions.

A significant proportion of eminent UK researchers spend a great deal of time as assessment panel members and the attention of many other researchers is focused in their institutions on RAE preparation, so diverting them from active research. The change of behaviour towards shorter-term research conflicts with the stated priority of the EPSRC to support more adventurous 'blue skies' research.

However, the community is unconvinced that sufficient thought has been given to a metrics based alternative. The current proposals in this respect propose that the most significant metric would be the amount of project funding gained from bodies such as the EPSRC. This would undermine the principle of dual funding and all of the current projections of the effect of this on different disciplines show that it would lead to a diversion of funding away from computer science to medicine and other sciences with expensive equipment requirements.

Furthermore, such an approach is likely to lead to increased pressure on researchers from their institutions to submit proposals for external research funding. Any significant increase would, almost certainly, place intolerable pressure on the EPSRC and the increasingly strained peer review system for assessing research proposals (see section 2.1).

1.3 Pressures on research

The changes in UK universities as a result of government policy initiatives have resulted in serious detrimental pressure on research and research scientists. There is increasing external regulation which has spawned an associated bureaucracy within universities. This bureaucracy makes increasing demands on the time of (especially) senior staff to complete paperwork and follow procedures, which have little discernable academic value.

Many (although by no means all) universities are still in financially precarious positions which means that pressure is placed on researchers to give priority to activities which potentially raise income, without necessarily assessing the academic benefits that might ensue from these activities. The significance of the QR income associated with the RAE rating also means that virtually all universities are demanding a significant amount of time from senior staff for RAE preparation activities.

In the previous review, we made the point that the CS research community was under pressure from increasing number of students taking computer science degrees, with the consequent demands on time through increased teaching. There was a lag in appointing new staff to cope with increased student numbers. The worldwide decline in CS numbers has partially alleviated this pressure but, paradoxically, staff are now under pressure to devise new ways to maintain the income previously associated with undergraduate students. A number of universities have been so seriously affected by the decline in student numbers that they have significantly reduced their activities in computer science. A particular, serious, problem is that, in some universities, the funding available for teaching support has declined and so staff have correspondingly less time for research.

The consequences for PhD student numbers of the overall drop in CS degrees awarded is unclear but we are obviously concerned that the number of PhD applicants will fall. This will, in time, lead to a reduction in the numbers of potential candidates for academic posts.

2 EPSRC support for the community

The general experience of members of the UKCRC who have all had extensive interactions with EPSRC staff is that, virtually without exception, these staff have been helpful and constructive and supportive of computer science research.

However, over the past few years there has been a rapid turnover of EPSRC staff due to both natural job movement and the EPSRC's policy of rotating staff around the different research areas. This means that there is a continual loss of organisational memory within the EPSRC, with subsequent instability that affects the research community. While we recognise the benefits of staff rotation, we feel that the current communication mechanisms do not always work effectively. A recent example of this was proposed changes to the review criteria for long-term research projects such as the IRCs. As the EPSRC's associate programme manager (APM) had changed (twice) over the lifetime of these projects, the proposed criteria were different from those discussed at the beginning of the project. EPSRC adopts a flexible, hands off style of research project management. For example, investigators can usually vire between cost headings and most importantly, they can alter research directions during a project (with justification in the final report). This is very much appreciated by the community who are in general trusted and supported to carry out their research, with few bureaucratic over-heads (as compared to other funding bodies).

2.1 The peer review system

Research proposals to the EPSRC are sent for peer review, with reviewers selected by EPSRC staff. A 'prioritisation panel', drawn from members of the research community, then meets to consider the review reports and to rank the projects on the basis of these reports for funding. The panel members are explicitly instructed not to review proposals themselves but to base their assessment on the reviews received. EPSRC staff then decide, on the basis of the available budget, how many proposals from the top of the ranked list can be funded.

It is the opinion of the UKCRC that this system is under considerable pressure. EPSRC staff find it increasingly difficult to find reviewers with the time to devote to providing a thorough and constructive review. This has the consequence that projects may have quite different numbers of reviews (so the panel must somehow prioritise a project with 4 reviews alongside one with 2 reviews) and too many reviews are of dubious quality – the reviewer may not have the appropriate background or may only have had time to write a cursory review. The situation is exacerbated by requests to review final reports from completed projects as well as proposals for support.

To help detect factual errors, proposers are sent reviews before a panel meeting and may prepare a single page response to these. The panel may then choose whether or not to disregard a review. However, this does not necessarily help with situations where reviews are perfunctory or where a reviewer has misunderstood or misread a proposal.

Furthermore, the nature of the peer review process is that it tends to favour incremental rather than adventurous research. It is much easier to judge the quality of a proposal which proposes some advances in a known area than it is to judge a proposal which proposes a radical new research direction. Reviewers seem to be inherently conservative and this means that incremental research tends to be more highly rated.

We note that the EPSRC is aware of this latter problem and has put a number of initiatives in place (platform grants, large projects assessed by a separate panel, etc.) that partially address this problem. However, the problem persists for smaller-scale projects.

The UKCRC recognises that there are no easy answers to the problems of peer review. As we have said, other pressures on academics means that many reviewers are unable, through lack of time, to review proposals properly. More active panel member involvement would, we believe, be helpful but, of course, with a move away from specialised panels, it is difficult to ensure that the panel has the breadth of expertise to cover all proposals being considered.

However, something must be done as there is concern across the community about the peer review process. The UKCRC suggests that the best way forward might be to establish a working group of EPSRC staff and active researchers to address this problem and propose a mechanism that retains the integrity of peer review yet is perceived as effective by the research community.

2.2 Panel structure

A recent EPSRC decision was to move to a single ‘Computer Science’ (ICT) panel which considers for prioritisation all responsive-mode ICT proposals alongside. This contrasts with the previous situation where there were a number of sub-panels concerned with different aspects of the discipline; specifically, there were sub-panels for CS, people and interactivity, communications, photonics and electronics& functional materials.

Members of the UKCRC community have raised two concerns over this move:
1. The breadth of expertise on the panel will, inevitably, be more limited than when more specialist panels were used. Given the problems with the peer review system highlighted above, it may then be difficult for a panel to assess the quality of reviews in an area that falls outside the expertise of all panel members.
2. There may be an unintended shift of resources from one area of the discipline to another. This may result because of the different numbers of proposals in different areas and different reviewing practice.

We note that EPSRC has recognised these concerns in the area of large proposals requesting support of more than £1 million and has established a separate panel to consider such proposals.

2.3 Inter-disciplinary research

Inter-disciplinary research within the UK has flourished over the last five years, in response to (cross) research council led initiatives and changing cultures within universities. Significant new collaborations and consolidations have been in the areas of cognitive systems, e-science, health informatics, bioinformatics, systems biology, complex and hybrid systems, computational algebra, novel computation, technologies in support of creative arts and social informatics.

Cross cutting research initiatives by the EPSRC, in collaboration with other councils, have clearly stimulated e-science, bioinformatics, systems biology. Computer science research was stimulated by calls for work in fundamental computer science for e-science, the semantic grid and autonomic Computing and in HCI and security issues. There is a perception that the e-science programme has suffered from its short timescale and the need to commit large amounts of money quickly. This has not allowed researchers from the natural sciences time to explore opportunities to exploit emerging computer science research results and create inter-disciplinary proposals that contribute to both natural and computer science.

An extremely positive development for the computer science community has been the funding of five "IT-Centric" IRCs (inter-disciplinary research collaborations). Five IRCs were established following a call by the EPSRC which attracted over 100 initial 'expressions of interest':

- AKT: Advanced knowledge management (Director Professor Nigel Shadbolt, Southampton)
- DIRC: Dependability of computer based systems (Director Professor Cliff Jones, Newcastle)
- Ultra fast photonics for data communications above terabit speeds (Director Professor Wilson Sibbett, St. Andrews)
- Equator: technological innovation in physical and digital life (Director Professor Tom Rodden, Nottingham)
- MIAS: From medical images and signals to clinical information (joint EPSRC/MRC) (Director Professor Sir Mike Brady, Oxford, then Professor David Hawkes, UCL)

The IRCs are just approaching the end of the six year funding and all have been outstandingly successful. A challenge for all the groups is how to maintain the inter-disciplinary teams and sustain research momentum.

The significantly increased volume of inter-disciplinary research in the UK has, mostly, been funded through initiatives such as the bioinformatics initiative or the IRCs where projects had to include inter-disciplinary components. The perception of the community is that smaller, responsive-mode inter-disciplinary research proposals are less well-received. We believe that this is probably a problem of the peer review system where there are insufficient reviewers with inter-disciplinary expertise. It is often the case that the original contribution is at the interface of the disciplines rather than at the core of one of the disciplines involved. Reviewers without knowledge across disciplines who consider only the core research proposed may, understandably, rate inter-disciplinary proposals less highly than more focused, discipline-specific proposals.

The range of inter-disciplinary work in which computer scientists are involved highlights the fact that 'ICT' as a single area covering computer science plus associated hardware and materials research is increasingly outdated. We believe that both of the principal areas covered by 'ICT' would benefit from being separated.

2.4 PhD studentships

PhD students in the UK are classed as students rather than employed by the institution as research or teaching assistants. They must pay fees to the university (approximately £3000 p.a. for students domiciled in the EU; a significantly higher rate (£9000-£13, 000). for students domiciled elsewhere). EPSRC provides support for research students in the form of a payment to cover fees and a living allowance of about £1000 per month. As well as EPSRC-funded students, most universities also have research students who are self-funded or funded by industry or in many cases, funded by overseas governments and agencies.

EPSRC funding comes in the form of Doctoral Training Accounts with an annual grant for research student support that is computed as a function of the value of EPSRC project funding.

The introduction of DTAs (doctoral training accounts) has been generally welcomed by the community, providing much needed flexibility within institutions. The DTA grants are awarded (annually) to universities and calculated by means of an algorithm based on the EPSRC research grant income received by the departments within the institution. The nominal EPSRC stipend has increased incrementally in value over the last few years, and also in duration (nominally 3.5 years), nevertheless in some institutions it remains difficult to attract, retain and fund good research students from within the UK. Many excellent students from the EU and further afield wish to study in the UK, but funding is problematic as the rules for EPSRC's Doctoral Training Accounts only allow for the payment of fees, not stipends for EU students. This is in contrast to project studentships where non-UK residents may be appointed. We appreciate that the reasons for this discrepancy are outside the control of the EPSRC but finding some way of addressing this would, we believe, improve the pool of research student candidates.

Over the last few years, following government recommendations (e.g. QAA code of practice 2003/4), institutions have tightened up processes and procedures with regard to monitoring progress, development of transferable skills and the duration of PhD studies. In order to "count" in national statistics, theses must be submitted within a maximum four years after first matriculation (full-time student); this inevitably has an effect on what can be achieved during a PhD. EPSRC, however, have (sensibly) recognised that a 4-year limit is counter-productive and have removed sanctions that were previously applied for non-submission.

2.5 The first grant scheme

EPSRC operates a 'first grant' scheme that is open to recently appointed members of academic staff. Proposals to this scheme are considered separately from proposals made by more experienced researchers and are ranked relative to other first grant proposals. The UKCRC wholeheartedly supports the notion of providing support for early career researchers but has some concerns over the current eligibility criteria:

1. Proposers may not resubmit proposals to the first grant scheme if unfunded even if the peer reviews for a proposal are very positive (although, of course, these may be resubmitted as normal responsive mode proposals). This seems to us to be anomalous given that resubmission is allowed (and sometimes encouraged) for other responsive mode proposals. It would be more consistent if proposers for first grants were permitted to resubmit proposals as first grants that were deemed to be 'fundable' by the panel but which fell below the funding threshold.
2. The time limit for submissions of first grant proposals of 24 months is short given that early career academics are subject to increasing pressure to complete formal teaching-related assessments (e.g. the Certificate of Learning and Teaching in Higher Education) and to develop new courses. They may also be concerned with preparing results from their PhD for publication. This leaves them little time for proposal preparation and we suggest that a longer time limit for first grant proposals should be allowed.

2.6 Research infrastructure

We have already made reference to the fact that 'playing the RAE game' discourages investment in systems research projects where there is a need to build and maintain an experimental platform to carry out the research. For example, a research group concerned with new database technology may wish to develop a large-scale database with 1 million+ items to serve as an experimental testbed. However, while this requires significant effort, few research outputs stem directly from this work although, we argue, it may lead to more effective evaluation of new research results and, hence, better quality research.

Similarly, developers of large-scale tools such as theorem provers (e.g. Isabelle) require long-term support to maintain these tools as a resource for the community. Again, significant effort is required with few direct research results.

We believe that to maintain and develop research in experimental computer science, consideration needs to be given as to how such platforms and tools can be funded in a predictable way. Otherwise, we are faced with a situation where research that shows promise cannot be tested in a realistic way and hence cannot be developed beyond the demonstrator stage.

This is not simply an issue of peer review not accepting such proposals. The current review mechanism is explicitly designed to support new research rather than support the maintenance of research tools. Under the current criteria, any such infrastructure proposals would not satisfy essential requirements and so will be rejected. Different review criteria for such proposals are required. The creation and maintenance of such platforms plays a similar role in computer science to that played by major capital equipment in other areas of science and engineering, for which there are various forms of special provision. We recommend that consideration be given to the creation of mechanisms for the creation and long-term support of large-scale experimental platforms for computer science research.

Finally, we refer to the comment in the 2001 International Review where the point was made that researchers should not have to spend significant periods of time in activities such as the management

[8] www.smithinst.ac.uk/PressRelease.html

[9] An exception has only been made for medicine where starting salaries for academics might be a third higher; the non-clinical (full) Professorial minimum is £37K whereas for clinical it is £49K.

[10] Other interesting attempts to determine a research agenda are given in [Ref-12, Ref-13].

of evolving configurations of research software. It should be noted that this falls outside the normal computing infrastructure support that is available in all universities. While applicants may make a case for such support, such cases have not generally been accepted by the EPSRC, perhaps because the effort required for software support is often not apparent to a non-specialist. The lack of support leads to poor software engineering practice and may contribute to researchers being unable to gain maximum benefit from previous work in their own and other groups.

2.7 Evolution of EPSRC research funding

Since the previous International Review, there has been a significant change in the EPSRC research funding portfolio for Computer Science and ICT. Rather than simply focus on providing research support for individual, relatively small projects, new funding mechanisms such as Platform Grants and Portfolio Grants have been introduced. These are based on research reputation and previous success in gaining EPSRC funding. As we discuss in the companion document, data on the percentage of computer science research funded by these mechanisms is not easy to find.

These approaches were introduced to help promote inter-disciplinary and adventurous research as it was recognised that the peer review system was not always effective in evaluating such proposals. In general, the UKCRC supports the intention behind such funding mechanisms.

Recently, additional mechanisms for funding larger scale projects have been introduced and significant funding has been allocated to these areas. For targeted areas, such as autonomous systems and large-scale complex IT systems, a process of consortium building has been supported with the EPSRC advertising for and appointing programme directors. Their responsibility is to create a consortium and develop a proposal to carry out the work. The first proposal under this approach (large-scale complex IT systems) is intended for submission by the end of September 2006. In addition, it has been recently announced that large projects (more than £1 million) will be considered separately.

The UKCRC recognises that there is a need for large-project support and is supportive of the notion that special mechanisms are required to properly assess large projects.

2.8 Research planning

The majority of the EPSRC research project funding is devoted to responsive mode proposals which may be in any area of computer science. These are typically relatively small projects although, as discussed in the previous section, large responsive mode projects are also considered. The UKCRC's view is that the majority of funding for research projects should continue to be allocated to responsive mode projects.

Other funding is allocated according to strategic areas and the EPSRC is advised in these areas by the Technical Opportunities Panel (TOP) , the User Panel (UP) and the Strategic Advisory Teams (SATs). UKCRC members participate in these, particularly TOP and the ICT SAT. The introduction of SATs is particularly welcome. Suggestions from these panels are assessed and may be taken up as targeted research programmes of different types. The process of moving from a panel suggestion to a funded programme can be a lengthy one, often lasting several years.

The view of the UKCRC is that it is critical that long-term research challenges should be the primary influence on research planning. As discussed in the companion document (UK Universities Computing Research: From the 2001 International Review to 2006), the UKCRC has proposed a number of Grand Challenges which serve as a basis for stimulating and focusing future a shared vision of directions for computer science research, with a mechanism in place for proposing new challenges. These have been well-received by the research community and the EPSRC and have the potential for focusing research in different areas around a common, long-term goal.

2.9 Public engagement

Until recently, public engagement has not been a high priority for Computer science, since the public perception of the subject was quite good. But following the dotcom bubble, the introduction of ICT curricula (as opposed to Computer science) in schools, and general poor media publicity concerning IT job off shoring, public perception has changed for the worse. Since the last review, the research community has responded, with EPSRC's help, and is now actively engaged in at least 10 EPSRC public engagement awards and there have been two Senior Media Fellows in Computing (Prof. Noel Sharkey and Prof. David Howard). More Media Fellows would encourage student recruitment and ultimately the attraction of young researchers.

3 Recruitment of research staff

The research community is highly dependent on the recruitment of high-quality research staff. While the preference is generally to employ post-doctoral research staff, the lack of qualified candidates in some areas (particularly practical rather than theoretical areas) means that pre-doctoral candidates are sometime employed.

In some respects, the situation has improved over the past 5 years. While there is no shortage of industry positions, graduating PhD students no longer see industry as offering high salaries and a secure job. University salaries have slowly improved and recent PhD's often see a post-doctoral appointment as more attractive than a junior academic post. They recognise the benefits of continuing to work in an active research community. The EPSRC's provision of research fellowships for exceptional post-doctoral researchers has been very helpful in retaining staff and in promoting research.

However, there remain a number of problems in recruiting staff which we would like to bring to the attention of the panel, although these are not specific to computer science:

1. European employment legislation which requires employers to offer staff who have been on temporary contracts for 4 years to be offered the same employment rights as staff on indefinite contracts. Some universities have reacted to this by discouraging the re-employment of qualified and experienced researchers. Others have attempted to ensure continued employment of staff who have already passed the 4 year threshold by redeployment, even when the research is in an unrelated area.
2. Although salaries have improved relative to industry at the lower end of the scale, salaries in the middle and upper scale points remain relatively low. This makes it difficult to retain experienced staff who wish to pursue a university research career.
3. The best applicant for a post may be more experienced than anticipated, and consequently require a higher salary than allocated in the original research proposal. There is no flexibility in the system which can accommodate a more expensive RA, i.e. projects are cash limited. Therefore it is often the case that the only way the best RA can be appointed is to make the appointment for a period of employment which is shorter than that which was originally proposed (and advertised). This can seriously affect the ability to carry out the full research programme.
4. There is a refusal in many universities to issue employment contracts to researchers for the full length of the project because grants are awarded on a cash limited basis, and do not necessarily incorporate all nationally agreed pay rises. For example, the EPSRC may award a project for 36 months, but the university (who is the employer) will only grant an employment contract for 33 months because funding for 36 months is not guaranteed; the contract may be extended after this date, but only if there is sufficient funding. This places both the PI and the RA in a difficult position, the PI because they have a research programme lasting 36 months, and the RA because of financial uncertainty.

We recognise that (3) and (4) above are difficult problems to fix, since both the funders and employers have budgetary constraints. Currently, proposers cannot allow for contingencies in their request for EPSRC funding – perhaps the introduction of such a heading, limited to a small percentage of the total staff costs, might help the situation.

4 Interactions with industry

Interactions between academe and industry have been very healthy, despite difficulties within the IT sector. Many of the interactions are not recorded by the usual EPSRC mechanisms, which record only formal collaborations on EPSRC funded projects. In particular, research exploitation through the formation of spin-out companies is particularly common in CS, but not captured by the EPSRC mechanisms. Long-term collaborations often arise as a result of a research project, but these may fall outside the accounting period used by the EPSRC and are therefore unrecognised. In April 2005, UKCRC submitted a report to the EPSRC entitled "Interactions with Industry". This report outlined the results of surveying interactions between 11 5/5* Departments and industry. The survey showed that across the 11 Departments surveyed, there were interactions with 280 UK based companies and 70 companies based abroad, totalling, during the period 2000-2004, £17M in cash and £57M in kind contributions.

External targets placed on EPSRC to have formal industrial involvement in a certain percentage of projects places differential strain on the discipline as some areas are more appropriate for industrial involvement than others. UKCRC believes that industrial sponsorship should not be seen as some quality indicator; furthermore, industry has a patchy record of exploiting research and predicting what research will be useful. The definition of 'industry' tends to be narrow and it excludes many other sectors, for example, local government. Computing researchers interact increasingly with large parts of the public sector such a local authorities and the health service.

5 The Research Landscape

This section gives an outline of UK research strengths and an indication of some future directions. To help the reviewers, we have mentioned a selection of researcher and universities, but only as exemplars. It is very important to note that there is no attempt to list everybody active in each area, nor does mention/absence indicate a quality judgement. The results and goals, whether scientific or technology transfer, are only indicative, not definitive and there are some overlaps in the descriptions. The topics covered (in no particular order), are:

1. Hardware and Architecture
2. Communications and Networks
3. Middleware and Distributed Systems
4. Software Engineering
5. Formal Methods
6. Dependability and Security
7. High Performance Computing
8. Information Engineering
9. Artificial Intelligence
10. Human Computer Interaction
11. Ubiquitous Computing
12. Theory and Foundations
13. Algorithms and Complexity
14. Vision and Medical Imaging
15. Bioinformatics & Systems Biology
16. Decision Support Systems and Evolutionary Computing

For each topic, an indication is given of the relevant EPSRC ICT area(s).

5.1 Hardware and architecture

Over the last 5 years there have been major shifts in the global industrial outlook for computer hardware and architecture, most notably announcements by the leading US high-end microprocessor manufacturers that their future roadmaps would be based on multi-core chips rather than ever-more complex uniprocessors. There is now an imperative to address the problem of running general-purpose programs on shared memory chip multiprocessor machines in addition to the established need for improved support for particular data types such as streaming media.

Current UK academic research in the computer architecture area includes work on chip multiprocessor architecture (Watson, Manchester) and, addressing the increasing interest in hardware for ubiquitous computing, event-driven architectures for very-low-power applications (May, Bristol). Much UK architecture research continues to focus on the need for extended data type support. This has been approached through SIMD instructions set extensions (Topham, Edinburgh), the efficient synthesis of dedicated hardware (Brown & Zwolinski, Southampton) or reconfigurable architectures (Luk & Cheung, Imperial; Arslan, Edinburgh), and has been accompanied by the formation of several start-up and university spin-out companies in this area.

Systems-on-Chip have continued to increase in complexity and ARM's ascendancy in this market has continued, with shipments of ARM processors in 2006 expected to exceed 2 billion. UK academic research strengths in asynchronous design (Furber, Manchester; Yakovlev, Newcastle; Moore, Cambridge), and in particular in asynchronous microprocessors, paved the way for the ARM966, a commercial asynchronous ARM processor developed by Handshake Solutions, a Philips company. Other UK research addresses on the growing complexity of on-chip communications through the development of globally-asynchronous locally-synchronous (GALS) techniques such as self-timed Networks-on-Chip, also leading to commercial exploitation.

University research in the 1990s into tools to support architecture migration through dynamic binary translation led to the establishment in 2000 of Transitives (Rawsthorne, Manchester), a market-leading spin-out company who played an important role in the high-profile transition by Apple from PowerPC to Intel processors in 2005. Progress in dynamic translation may point the way to dynamic parallelisation techniques, potentially addressing the need for general-purpose multiprocessing noted above.

Longer-term research in the architecture area has been promoted through EPSRC initiatives in Novel Computation and joint Research Councils' follow-up to the DTI Cognitive Systems Foresight programme. This has led to a marked increase in funding and activity in hardware and architectures for radical approaches to computation such as neural computing (Austin, York; Furber, Manchester; Murray, Edinburgh), bio-inspired computing (Tyrrell, York), and evolutionary multi-objective computing (Erdogan, Edinburgh). An EPSRC Inter-disciplinary Research Collaboration has been established in Quantum Information Processing (Briggs, Oxford). UK strengths in quantum computing are on the theory side; currently there is strong interest in 'flying qubits' as a route to scalability through networking small-scale subsystems.

ICT areas: Electronic Devices and Subsystems, Systems on a chip

5.2 Communications and networks

UK researchers have made significant contributions to the interconnected fields of communications and networking. Continuing a long tradition of international excellence, the UK researchers in these areas have tackled many of the significant issues that have emerged since 2001. Many of these efforts are collaborative in nature, and strong linkage with industry and EEE continues to be a dominant theme. There is a particularly strong linkage of these activities to standardization efforts within the IETF.

Three universities have created new programmes in networking research since the last international review: Glasgow (Sventek) and Kent (Marshall) have attracted strong research leaders from industry (HP/Agilent and BT, respectively), and St Andrews has attracted a leading network researcher from UCL (Bhatti).

In 2002, Bhatti (St Andrews) and others successfully lobbied for the creation of UKLight, a national high-speed optical facility to support research into next-generation network architectures and protocols. UKLight has now been deployed widely across the UK, and has led to the provision of separate research connectivity for the UK as part of SuperJanet 5, substantially predating the current US effort in GENI.

There has been a significant amount of activity in network routing, transport protocols and networked application performance, especially in light of the emergence of new forms of networks, specifically wavelength-agile optical, GPRS/UMTS packet, delay-tolerant, and mobile ad hoc networks. These studies are based upon modelling and simulation (Ould-Khaoua/Glasgow), emulation (Bhatti/St Andrews, Crowcroft/Cambridge, Handley/UCL). The XORP modular router open source activity is led by Handley (UCL). The open-source NRS platform (Bhatti/St Andrews) is being adopted as part of a Europe-wide network QoS control experiment.

Lancaster (Hutchison), Southampton (Chown) and UCL (Kirstein) have led experimentation in the deployment, use, and performance characterization of IPv6 networks. Lancaster and Southampton have been particularly active in the mobile IPv6 activity in the IETF.

Security and trust in networks, especially mobile networks, has been studied at UCL (Hailes).

There is substantial activity in sensor networks. Kent (Marshall) and Southampton (de Roure) have constructed and deployed a number of environmental sensor networks through the DTI's EnviSense programme; these deployed networks are being monitored to guide future design and development. Additionally, through the WINES programme of the EPSRC, there are a number of current projects actively building, deploying, and evaluating sensor networks to understand how their requirements will affect future network architectures.

Several groups have focussed on low-level measurement of network activity. Cambridge (Pratt & Moore) developed the nProbe, a full line-rate packet capture device, using commodity processors, busses, and network interface cards. Loughborough (Philips), QMUL (Moore) and St Andrews (Bhatti) are constructing monitoring systems for the 10 Gbit/sec lightpaths provided by UKLight. QMUL (Moore) has developed a number of statistical methods for analysis of packet traces collected using devices such as the nProbe. Glasgow (Sventek) has focussed on real-time classification of network flows to the corresponding generating applications, using both innovative payload inspection schemes and wavelet-based statistical schemes.

Finally, network management research has focussed primarily upon policy-based approaches, and their integration with traditional operational support systems. Imperial (Sloman, Lupu, Dulay) have explored a policy basis for management at all levels of abstraction, while Stirling (Turner) has focussed specifically on tuning policy-based management to the management of networks.

ICT areas: Networks and Distributed Systems, Systems Integration, Mobility

5.3 Middleware and distributed systems

UK researchers have made significant contributions to the fields of middleware and distributed systems. Continuing a long tradition of international excellence, the UK researchers in these areas have tackled many of the significant issues that have emerged since 2001. Many of these efforts are collaborative in nature, and strong linkage with industry continues to be a dominant theme.

Throughout the period since the last international review, Cambridge (Bacon and Moody) has been one of the leading institutions world-wide in the design, implementation, and characterization of event based systems; their recent work looks at providing this type of infrastructure in mobile networking environments. Glasgow (Sventek) has adapted publish/subscribe mechanisms to the constrained environments of environmental sensor networks. UCL (Hailes and Mascolo) has focussed on other aspects of mobile middleware, particularly optimal routing and security/trust. The work on reflective middleware at Lancaster (Blair and Coulson) was the first of its kind. Blair and Sventek (Lancaster and Glasgow) are current and past chairs of the steering committee for the ACM/IFIP Middlware Conference, the primary venue for reporting the latest middleware research results.

There has been significant activity in the area of Pervasive Computing in the UK during this period. The Equator Inter-disciplinary Research Collaboration, funded by the EPSRC, addressed fundamental research issues arising from the interweaving of physical and digital interactions; one of its research challenges focussed on the infrastructure required to support the dynamic assembly of new devices into coherent user experiences. Nottingham (Rodden), Chalmers (Glasgow), and Southampton (De Roure and Hall) were the principal institutions focussed upon this particular challenge, and their results have been widely disseminated through Ubicomp, Percom, and other pervasive computing research channels. St Andrews (Morrison and Dearle) studied the construction and deployment of Global Smart Spaces, to devise new software infrastructures to facilitate requisite low-level interactions by both explicit and implicit high-level context of users on a large scale. Many of the projects funded by the EPSRC's WINES programme are focussed on various aspects of pervasive computing, as well. There has also been substantial funding from DTI as part of their Next Wave Technologies Programme with focus on Healthcare and Body Sensor Networks at Imperial (Yang, Sloman), Environmental Monitoring at Southampton (De Roure) and Smart Homes at Loughborough (Kalawsky).

For pervasive computing systems to become ubiquitous, it is essential that these systems become self-managed or autonomic. Work at Imperial (Sloman, Lupu, Dulay) and Glasgow (Sventek) is focussed on autonomic management of such systems, and particularly the dynamic federation or composition of two or more such systems. Other efforts at St Andrews (Morrison and Dearle) have focussed upon how closed loop control mechanisms are incorporated into software by the software engineering tool chain.

There has been significant work in the area of middleware in support of e-Science. The UK National e-Science programme has sponsored many research and development activities to improve the e-Science infrastructure for use by physicists, astronomers, biologists, and other experimental sciences. Many of these efforts are strongly linked to standardization efforts within the Global Grid

Forum. Of particular note has been activity in data storage and access, manifested in the OGSA-DAI specification produced by the National e-Science Centre (Atkinson/Edinburgh&Glasgow) and adopted by the GGF. There is ongoing work on digital curation of scientific data (Buneman/Edinburgh). Both of these efforts manifest themselves in augmented Grid middleware. The e-Science programme has also set up the Open Middleware Infrastructure Institute, a joint enterprise at Southampton, Manchester and Edinburgh, to develop, harden and support software and middleware for the UK's Scientists.

ICT areas: Networks and Distributed Systems, Systems Integration, Mobility

5.4 Software Engineering

Since the 1960s, when Randell and Buxton arguably coined the term 'software engineering' at the early NATO conferences, there has been significant UK involvement in software engineering research. UK researchers are an active part of the international community and have chaired major international conferences such as ICSE. One of the two leading texts in the area has a UK author (Sommerville). Areas of strength include requirements engineering, software architectures, distributed systems engineering, software evolution, real-time systems, aspect-oriented software development, empirical software engineering and testing.

UK researchers in requirements engineering (RE) have made important contributions in a number of areas. Nuseibeh (Open) along with Finkelstein (UCL) were involved in early work in inconsistency management, with Nuseibeh now a leading figure in security requirements engineering. Kramer (Imperial) is working on the use of models in requirements elaboration. Maiden (City) has extended requirements engineering with creativity theories and models from AI and psychology and is extended RE towards service-centric systems. Inter-disciplinary approaches have been pioneered in the UK - Bustard (Ulster) has linked RE with soft systems analysis, Sommerville (St Andrews, previously Lancaster) pioneered the use of ethnography in RE and Loucopoulos (Manchester) has investigated the links between organizational modelling and RE.

The team at Imperial (Kramer, Magee) have been involved in software architecture research for many years, starting with the CONIC system. Current work is concerned with behaviour analysis, support environments and architectural approaches to self-organising systems. Finkelstein and Emmerich (UCL) are examining relations between RE and software architecture. Work in architecture is closely related to distributed software engineering with Kramer, Magee, Finkelstein and Emmerich all active in this area. Rosenblum who has joined the UCL team from the USA, works in the design and validation of distributed component-based software and Henderson (Southampton) in middleware engineering. City (Maiden) and Lancaster (Sawyer, formerly Sommerville) are partners in SeCSE, a European project concerned with service-oriented systems engineering.

Warboys at Manchester has been active in software process modelling research for many years and is currently working with Morrison (St Andrews) in compliant system architectures which is concerned with supporting integrated process and architectural evolution. Munro (Durham) established the research institute for software evolution which has made important contributions in the area of software visualisation.

The real-time systems group at York led by Burns and Wellings is world-leading in the world and is currently involved in research in scheduling theory, distributed and embedded systems, real-time programming languages and operating systems, SoC and FPGA platforms, executing time analysis and power-aware computation. In the emerging area of aspect-oriented software engineering, Rashid (Lancaster) leads a European network project and has made contributions in aspect-oriented databases and the relationships between aspects and requirements. Groups at the Open University and UCL are also active in this area.

A group at Durham and Keele (Budgen, Brereton and Kitchenham) focuses on empirical software engineering and are currently working to establish an infrastructure to support evidence-based software engineering. Holcombe (Sheffield), using projects from his innovative student-run software company, is leading a project (the Observatory) concerned with the empirical assessment of agile approaches to software engineering. Shepperd (Brunel) and Fenton (QML) are also active in this area.

From a background in formal methods for software testing and evolutionary computation, Harman, (Kings), Hierons (Brunel) and Clarke (York) have extended their work into a new field of search-based software engineering where they are leading a large EPSRC-funded project. Holcombe (Sheffield) is also active in testing and was instrumental in developing the X-machines based approach.

ICT areas: Software Engineering, Systems Methodology and Architecture, Systems Integration

5.5 Formal Methods

The UK plays an internationally leading role in formal methods research because of its strength in theory, tools, and industrial application. It builds on successful research that has had a major impact on academia and industry, and is well connected to work on software engineering and dependability. Fitzgerald (Newcastle) chairs FME, the world's leading formal methods dissemination organisation, and UK researchers regularly chair programme committees for all the key conferences. The UK is taking a major role in the international grand challenge on verified software, helping to set the research agenda and driving the development of the repository and pilot projects (Bicarregui (RAL) and Woodcock (York)).

Important UK work on specification languages and development methods (B, CCS, CSP, VDM, Z) continues, emphasising tool support and combining different paradigms. Henson (Essex) works on improved semantics and refinement techniques. Jones (Newcastle) and Butler (Southampton) lead the European RODIN project, developing methods and tools (particularly for Event-B) for developing fault-tolerant systems with unpredictably changing environments. Schneider and Treharne (Surrey) have a well-established programme of work in CSP||B. Woodcock and Cavalcanti (York) combine Z, CSP, and the refinement calculus in Circus. Hillston (Edinburgh) won the prestigious Needham Award for the stochastic process algebra PEPA that studies behavioural and performance properties. Butler (Southampton) combines B and UML, Poernomo (King's) and Paige (York) integrate formal methods into model-driven and agile development. Hustadt (Liverpool) and Schmidt (Manchester) work on practical reasoning for web ontologies and multi-agent systems. Jones (Newcastle) tackles the problems of atomicity. Kwiatkowska (Birmingham) is developing systematic methods for ubiquitous computing. Martin (Queen Mary) has made important advances in a computational logic framework for reasoning about dynamical systems. Cavalcanti (York) is contributing to this emerging area with formal methods, refinement techniques, and tools for control systems. Calder (Glasgow) and d'Inverno (Westminster) work on biological modelling.

UK work on tools (FDR, HOL, Isabelle) continues with temporal logic provers (Liverpool) and Voronkov's Vampire theorem prover (Manchester), a regular winner of the annual CASC competition for automated, first-order provers. Kwiatkowska (Birmingham) is prominent in probabilistic model checking with the leading tool PRISM. Roscoe and Lowe (Oxford) are making important progress in protocol security and information flow using CSP and FDR. Paulson (Cambridge) is linking interactive and automatic proof tools. Melham (Oxford) works on model checking partially ordered spaces. Derrick (Sheffield) and Thompson (Kent) work on formally based tool support for developing Erlang programs.

The UK has considerable strength in static program analysis (Cambridge, Imperial, Queen Mary, King's, City, Kent, Oxford, Birmingham, Leeds), an area whose focus has changed from compiler optimisation to automatic program verification. Hankin and Wicklicki (Imperial) have pioneered the rigorous use of quantitative methods in static program analysis in the development of probabilistic abstract interpretation, and Malacaria (Queen Mary), Hunt (City), and Clark (King's) have applied information theory to program analysis. Separation logic (O'Hearn, Queen Mary; Imperial; Cambridge) has provided a significant breakthrough in the 35 year-old problem of reasoning about linked data structures in memory, and the UK lead is being followed by many theoreticians and tools researchers worldwide. Microsoft is developing separation logic tools, which have already discovered heap and termination bugs in Windows device drivers. The application of game semantics has led to significant advances in model checking concurrent imperative programs and open systems (Ghica, (Birmingham), McCusker (Bath), Ong (Oxford)).

Bowen (LSBU), Harman (King's), Hierons (Brunel), Bogdanov (Sheffield), and Clark (York) have explored the fruitful interplay between formal methods and testing.

ICT areas: Software Engineering, Systems Methodology and Architecture, Systems Integration

5.6 Dependability and Security

UK researchers from Newcastle (Randell, Rushby, etc.) were instrumental in establishing the research area of software dependability and Newcastle remains a centre of excellence in this area. Work in Newcastle (Randell, Anderson, Ryan has evolved to include security-related issues (e.g. in the MAFTIA project) and issues of service dependability. Newcastle, along with City University established the centre for Software Reliability (CSR) which has a major industrial influence. Bloomfield who has extensive industrial experience in safety-related systems has recently taken over as Head of CSR at City where work on modelling dependability (Strigini, Littlewood) and safety-cases (Littlewood, Bloomfield) are particular strengths. Both Newcastle and City (along with Lancaster, York and Edinburgh) were partners in DIRC, which has extended technical notions of dependability to cover socio-technical issues. The RODIN project (Romanovsky and Jones, Newcastle) is combining the research on fault-tolerance and formal methods; the EU NoE ReSIST extends that research to cover ubiquitous computing systems.

In high-integrity aerospace systems, McDermid at York has established what is probably the leading research centre in the world in this area. Research includes safety analysis and formal modelling, systems of systems, formal development, systems and software architectures and safety cases. York work in formal methods (Woodcock) and real-time systems (Burns, Wellings) is complementary to this. Johnson in Glasgow has established an international reputation for his work in accident analysis.

The UK has two large security research teams, at Cambridge and Royal Holloway. At Cambridge, Needham initiated the study of protocols in the 1970s and 1980s while Royal Holloway’s Piper wrote an influential textbook on cryptography in 1982. Since then Royal Holloway has done much work on the discrete mathematics that underpins modern cryptology, while Cambridge has worked in many systems areas.

Cambridge research contributions include iris scanning (Daugman), the use of theorem provers in protocol verification (Paulson), early contributions to peer-to-peer and censorship-resistant systems (Anderson), the analysis of hardware tamper-resistance (Kuhn, Anderson, Skorobogatov), pervasive computing and location privacy (Stajano) and most recently the establishment of the economics of information security as a discipline (Anderson).

Royal Holloway contributions include key management schemes for public sector email (Mitchell), secret sharing and authentication schemes (Martin) and cryptanalysis (Murphy, Paterson). Both teams have a long record of working with industry on information security problems.

As CPUs and communications are built into ever more devices, software will make up ever more of the value added by UK industry, and ever more industries will come to resemble the software industry. There will be the good, the bad and the ugly. The good includes flexibility, adaptability, productivity and personalisation. The bad will include usability and maintainability problems. The ugly will range from monopolies to vulnerabilities that can be exploited by criminals and others. Getting the maximum benefit from technological progress will necessitate more attention to information security issues. As more and more value, and more and more vulnerability, becomes digital, infosec problems will in time come to dominate more conventional security problems of the kind currently dealt with by the police and the courts. Dealing with them will require a mix of technological and policy tools, to whose development the UK makes a disproportionate contribution.

ICT areas: Networks and Distributed Systems, Systems Integration, Fundamentals of Computing and AI

5.7 High Performance Computing

High Performance Computing (HPC) aims to contribute to the most ambitious scientific goals by widening the scope and extending the boundaries of researchers’ ambitions and by stretching the development of hardware and software technologies in directions that often yield unexpected benefits outside the HPC research arena.

UK computer science has been heavily involved in all aspects of HPC ranging from the development of formal definitions of languages, standards and systems through to the construction of tools and applications. Of particular significance has been the work and development in the area of parallel computation. The pioneering work of Hoare and Milner with notations such as communicating sequential processes (CSP) and a calculus for communicating systems (CCS) laid the foundations for the rigorous treatment of synchronisation and communication between processes, fundamental in the development of the distributed memory computation model of parallel computation. Other notable highlights in this area have been the work in languages and applications. Not only has the design and implementation of particular languages led to clarification of parallel computing concepts but UK computer scientists have been at the forefront of international efforts to establish parallel software standards and benchmarks. A notable example is MPI, the message-passing interface which is widely accepted by vendors and researchers as a basis for promoting portability across high performance platforms. As for MPI there have been important contributions at the international level to research on other standards such as OpenMP.

At the same time significant experimentation was carried out on the development of new compiler techniques to ensure efficient implementation of the linguistic concepts on high performance machines. Research on parallel algorithms, the basic building blocks of applications software, ranges from work based on PRAMs and Valiant's BSP model to more pragmatic research on parallel numerical algorithms. Applications range from engineering simulations such as computational fluid dynamics to commercial applications such as data warehousing, data mining and financial modelling. Experience gained from these applications has been stimulated by research into the provision of appropriate software tools and the discovery of new parallel algorithms. There has been a long history of successful collaboration between academic researchers and the emerging HPC industry in the UK. While UK HPC companies have found it difficult to establish themselves as international leaders in broad market sectors, there are a number of small scale successes, such as NAG, Quadrics and Clearspeed, who continue to thrive in their niche areas.

One of the significant initiatives to emerge from the UK in recent years has been the e-Science Initiative which was structured as an inclusive initiative embracing multidisciplinary activity. An important feature of the initiative was the national infrastructure based on a number of Regional e-Science Centres. In general the Centres were awarded to Universities which had well established research records of activity in high performance computing. Thus the research activity in the HPC area contributed to the underlying structure of the e-Science Initiative and can claim some credit for its success nationally and internationally. In addition UK HPC research has made many contributions in the development of middleware and in the Global Grid Forum (now the Open Grid Forum) which helps encourage standardisation activities, and has underpinned the UK's leading role in applications-led e-Science research One of the important outputs of the UK e-Science Initiative has been the Open Middleware Infrastructure Institute which provides a web service infrastructure for building grid applications. The e-Science initiative has also been the catalyst in interactive use of HPC where research in computational steering and visualization has increased scientists' ability to work 'within a simulation' - leading to better insight on the part of the scientist, and to more effective use of scarce HPC resource.

In recent years, interest has burgeoned in the possibility of using grid computing for HPC. This entails the complications of handling heterogeneity and dynamic workloads in parallel computations, and represents a step change in complexity. UK computer scientists are finding this extended capability a fertile field in the context of performance modelling, performance prediction and performance control of Grid-based computations.

Future computer science research challenges are intimately related to the e-Science programme and the development of a reliable computation and information Grid infrastructure. Research into tools, algorithms and environments will result in significant contributions to the development of Grid technology. In addition, in the future there is an opportunity for computer science research into new areas such as in reconfigurable systems in HPC based on FPGAs, the impact of the shift to multicore processors, and the impact of lambda networks and UK-Light type systems on distributed HPC applications. The biggest challenge in HPC is in the mechanisms and methodologies for software construction and this is where UK computer scientists can continue to make major contributions.

Note: In 2005 an International Review of HPC was commissioned by the Research Councils

http://www.epsrc.ac.uk/ResearchFunding/FacilitiesAndServices/HighPerformanceComputing/InternationalReview/default.htm.

The Panel found that research which relies on high performance computing is of the highest quality and is competitive at the international level. However, the UK cannot afford to be complacent and tread water in a rapidly evolving field such as computational science and engineering.

ICT areas: Parallel Computing, Modelling and Simulation

5.8 Information Engineering

The UK community has an internationally leading position in Information Retrieval (IR). Four of the eight winners of the ACM SIGIR's prestigious Salton Award are British of which three remain active (van Rijsbergen, Glasgow, Sparck Jones, Cambridge, Robertson now Microsoft Cambridge Research, formerly City). Glasgow is one of the world's leading centres for IR research, especially through Terrier, its experimental IR engine, and van Rijsbergen's continuing influence on developments in IR theory. Other major groups in the UK include: Queen Mary, London (Lalmas, theory of IR and XML retrieval); Sheffield (Sanderson, still image and cross language retrieval); Strathclyde (Crestani, digital libraries; Ruthven, formal methods and interaction); Robert Gordon (Harper: language modelling approach to IR, within document retrieval); Sunderland (Tait, still image retrieval and word sense disambiguation); the Open University (Rueger, formerly Imperial, image retrieval).

In Knowledge Management the UK also has had a world leading place in the development of Ontologies through the work of Horrocks (Manchester) on OWL and other work in collaboration with Paton and Goble (also Manchester). Horrocks work led to him being award the prestiguous Needham prize. Work on semantic web applications and life cycles at Southampton by Shadbolt, Hall and others was awarded the International Semantic Web Challenge and was instrumental in Shadbolt being awarded a fellowship of the Royal Academy of Engineering. Other major groups include the Open University (Motta, large scale ontology engineering) and Ciravegna (Sheffield, knowledge acquisition, sharing and reuse). (Note there is some overlap between IR, KM and AI, see next section.)

The database research community has been strengthened by international appointments at Edinburgh (Buneman, Libkin) and Oxford (Gottlob). The Edinburgh Group (Buneman, Libkin, Fan, Viglas) has a world leading position on provenance and semistructured data management, and Gottlob has produced seminal results on the foundations of graph languages and information extraction. The longer established database community in the UK has been closely associated with the e-Science programme, investigating service-based data access and integration at Edinburgh, Manchester and Newcastle (Atkinson, Paton, Watson), and leading the development of several data-oriented standards in the Global Grid Forum. The recent emphasis in scientific data management is reflected in work on archiving of scientific data in Edinburgh (Buneman), biodiversity informatics in Cardiff (Gray), taxonomic data management at Napier (Kennedy), data integration and ontologies for genomics at Manchester (Embury, Goble, Paton, Stevens), sequence indexing at Glasgow (Atkinson), and scientific data mining at Imperial (Guo).

ICT areas: Information and Knowledge Management, Systems Methodology and Architecture

5.9 Artificial Intelligence

The UK has been a pioneer and world leader in Artificial Intelligence and Cognitive Science research since foundation of the field. For current purposes, we will interpret these areas broadly to include work on computer vision, neural nets, computational neuroscience and cognitive systems. During the period since the last EPSRC International Review of Computer Science in 2001, UK AI has not only maintained but strengthened its world-leading role. There are AI research groups at many UK universities, some of which are major groups of long standing, going back to the mid-60s. UK AI researchers are playing central roles in two of the seven UKCRC Grand Challenges: Memories for life and The architecture of brain and mind. AI also played a central role in two of the five EPSRC Inter-disciplinary Research Collaborations: Advanced knowledge management (AKT) and From medical images and signals to clinical information''(MIAS).

The high standing of UK AI research is indicated by the IJCAI awards, which are the major AI esteem indicator. These have been dominated by North American winners, but the UK has provided the only two non-North American winners of the Research Excellence Award: Michie (2001) and Bundy (2007), and one of the three non-North American winners of the Computers and Thought Award: Jennings (1999). UK AI researchers are frequently invited to write the definitive articles on their research area for handbooks and similar survey publications.

During the review period we can point to a number of major impacts of AI research.

Knowledge Representation and Inference: There are major centres at Southampton (lead site for the AKT IRC) Birmingham, Cambridge,Glasgow, Edinburgh, Imperial, Leeds, Liverpool, Manchester, Open, Strathclyde, Sheffield and York. The APES consortium has a world-reputation for its work on constraint problem solving, etc. Four major, European, next-generation internet projects are being coordinated from the UK: SEKT, NEON, X-MEDIA, and OpenKnowledge, building on strengths in ontologies, multi-agents and knowledge-based systems. Ground-breaking work on spatial representation and reasoning is being done at Manchester, Leeds, London (Imperial, Birkbeck) and Liverpool. The plan formation and robotics communities are transforming approaches to emergency response (Edinburgh) and space exploration (Strathclyde, Aberystwyth). Engagement with user groups in formal methods, web applications, systems biology etc. has grown significantly over the review period. Groups are involved in many large-scale international collaborations with researchers, government agencies and businesses in Europe, USA, Japan and elsewhere.

Speech and Natural Language Processing: There are major centres at Cambridge, Edinburgh, KCL, Leeds, OU, Oxford, Sheffield, Sussex, and a high proportion of papers at the premier international Annual Conference of the Association for Computational Linguistics (ACL), and IEEE on Acoustics, Speech and Signal Processing (ICASSP), are from these groups. The Cambridge HTK toolkit is the most widely used freely available research speech recogniser. The Cambridge/Sussex RASP and Edinburgh/Oxford C&C parsers, and the Sheffield GATE language technology tools, are widely used and free to academic researchers. The Edinburgh FESTIVAL speech synthesiser is the most widely-used freely available research synthesiser, and was the foundation for a spin-out company that developed a commercial version, used by Toyota and others. The UK has a very strong tradition in information retrieval, and has active research in methods for text management and retrieval at Glasgow, Sheffield, Cambridge and City University. UK researchers in speech, language, and information processing have been significant participants in major international evaluation programmes and bakeoffs e.g. the DARPA and NIST competitions in spoken language systems, text retrieval (TREC), machine translation (GALE), etc.

Robotics: There are major centres at Aberystwyth, Edinburgh, Essex, Imperial, Oxford, Surrey, Sussex and West England. The ground-breaking Robot Scientist project at Aberystwyth and Imperial (King and Muggleton) built a totally automatic biologist where hypotheses were formed by inductive logic programming and tested by robots, with the results of the experiments providing the input for the next round. Imperial has developed technology that can provide cheap real-time localisation for domestic robots, humanoid robots, wearable sensors, game interfaces or other devices. Edinburgh has applied machine learning to real time, online acquisition of sensorimotor maps in humanoid robots for dynamic control and planning, which is critical for human robot interaction.

Machine Learning: There are major centres at Bristol, Cambridge, Edinburgh, Gatsby/UCL, Imperial and Southampton. Through the national e-Science Programme, and otherwise, UK researchers have applied data-mining techniques to both scientific and industrial very large databases. Southampton's work on biometrics, including gait analysis, is very relevant to surveillance and anti-terrorism. Bristol, Edinburgh and UCL are key players in an international consortium developing kernel machines. Systems biology has become an important application area for both learning and reasoning techniques, with major new centres at Edinburgh, Glasgow, Imperial, Manchester, Newcastle, Nottingham and Oxford. Groups at Essex, Plymouth, UCL and Westminster have found widespread applications of genetic algorithms, including creating artworks in music and multimedia.

Computational Neuroscience and Cognitive Systems: There are major centres at Edinburgh, Gatsby/UCL, Plymouth and Southampton. These groups were strongly represented at the recent FP7 consultations for the forthcoming EU Framework 7 Artificial Cognitive Systems call. They have a high proportion of papers at leading international conferences in the field, such as Neural Information Processing Systems (NIPS). AI researchers have played a key role in the Foresight Cognitive Systems Programme.

Multi-agent systems: There are major centres of activity in this recent offshoot of artificial intelligence at Aberdeen, Imperial, Liverpool, and Southampton. The prominence of UK research in this area is evidenced by the fact that the only two non-US winners of the annual ACM Autonomous Agents Research Award are from the UK: Jennings (Southampton) in 2003, and Wooldridge (Liverpool) in 2006. Major areas of UK strength are computational auction design, automated negotiation, logical and theoretical foundations of multi-agent systems, agents and the semantic web, logic programming and multi-agent systems, argumentation, and agent-oriented software engineering.

ICT areas: Fundamentals of Computing and AI, Neural Computing, New and Emerging Computer Paradigms, Artificial Intelligence Technologies, Cognitive Science and its Applications

5.10 Human Computer Interaction

The UK has held a strong position in the international HCI community since the emergence of HCI as an area or ‘discipline’ in the early 1980s, helped in no small part by the MMI strand of the Alvey programme. In particular, the UK, specifically Brian Shackel, led the formation of IFIP TC13 and UK researchers were involved in the early ACM CHI conferences. UK involvement in these two major strands of international HCI (accepting the US ACM as de facto ‘international’) continues for example Cockton (Sunderland) is currently vice-chair of TC13 and 3 out of the last 4 CHI conferences have had British papers’ chairs: Monk (York), Rodden (Notts) and Payne (Man). The international recognition of UK HCI is also recognised in that of 37 members of the SIGCHI academy (awarded for recognition of long-term research contribution) 6 of the 8 non-North Americans are British or work in the UK. It is also interesting to note that the two principal international HCI textbooks have largely British authors.

As well as having a strong international standing, HCI has a broad internal UK base. The British HCI Group comprises one of the largest (and profitable!) BCS SIGs and hosts an annual HCI conference with substantial international attendance. Industrial involvement is also strong both professionally with a large UPA (usability professionals association) chapter and academically where the loss of Xerox’ Cambridge presence has been met by growth in other UK labs, for example HP labs’ academic collaborations in mobile applications. Notable also are a number of recent movements from UK and US industrial labs to senior academic positions such as Frohlich (from HP to Surrey) and Whittaker (from IBM/Lotus to Sheffield).

The long-term strength in social and ethnographic analysis (e.g. Rouncefield (Lancaster) and Proctor (Manchester)) has been important in allowing DIRC to build rich socio-technical understanding of dependability. The EQUATOR IRC has also established a key international presence, especially in US conferences. EQUATOR and other UK research groups such as Perry (Brunel) and Brewster (Glasgow) embody a strong interaction focus to areas of ubiquitous computing, situated displays and mobile computing. Indeed mobileHCI, now a major international conference in its own right began with a series of workshops in the UK.

In general UK HCI has particular strengths in areas at the edges of mainstream HCI, due in part to the fact that the UK has one of the least conservative academic environments in the world. This has allowed new areas to blossom and older areas to mature.

Over the last few years interactions between HCI and ‘non-functional’ areas such as games, arts and design, creativity, fun and user experience have become important internationally and the UK has been a strong player and often led these emerging areas, for example Monk (York), Wright (SHU), EQUATOR and the nexus of Edinburgh universities. At the opposite extreme, the international community in formal methods in HCI, particularly influential in Europe, grew largely from early researchers then in York (e.g. Harrison (Newcastle), Thimbleby (Swansea), Dix (Lancaster)) and continues in work such as model checking (Harrison (Newcastle)) and safety critical systems (Johnson (Glasgow)). Continuing the ‘edginess’ theme, HCI for non-WASP users is also strong including accessibility and the elderly (e.g. Newell (Dundee)), children’s interaction (e.g. Read (UCLan)) and development issues (e.g Dearden (SHU)).

As mentioned one of the reasons for the strength of UK HCI has been the relatively open attitudes within computing departments. (In contrast, HCI researchers in much more rigidly disciplinary psychology departments have either struggled or moved to computing.) This openness has allowed the development a cross-disciplinary community that is unusual internationally. For example, it is notable that the bulk of US work at CHI fits much more obviously into disciplinary niches. Work that cuts across disciplines always faces hurdles in recognition and funding, even in well established areas like HCI, but looking outward to other countries and other disciplines it is clear that we have a much more open academic environment, which in turn offers an opportunity to continue and to build upon our international prestige in the area.

ICT areas: Human Computer Interactions, Cognitive Science and its Applications, Multimedia, Vision, Hearing and Other Senses

5.11 Ubiquitous Computing

The period since the last international review of computer science in 2001 has seen Ubiquitous Computing emerge a distinctive research focus for many researchers groups within the UK engaging researchers from across the broad spectrum of the ICT remit. The growing interest in Ubiquitous Computing reflects a broad international trend fuelled by the rapid growth in wireless communications and the increasing availability of digital devices that offer the chance to make Ubiquitous Computing a real possibility in everyday lives.

The emergence of Ubiquitous Computing is reflected in large-scale initiatives such as the EU ambient intelligence programme, NSF initiatives, and industrial research programmes from Microsoft Research, Intel, and HP. Ubiquitous Computing provides a central focus for a number of research activities supported by EPSRC including the Equator IRC, UK EPSRC Wines programme and the UK-Ubinet network of excellence.

Indeed, Marc Weiser's original vision ("The World Is Not a Desktop," ACM Interactions,1(1), Jan. 1994) and the computing research agenda that underpins it has started to emerge as one of the most significant research challenges for computing science at this moment in time. Within the UK this is also reflected in the rapid growth in the UK-Ubinet activities and the major interest in initiatives such as the UK Ubiquitous Computing Grand Challenge. Ubiquitous computing activities can be considered in terms of three currently distinct perspectives:

The *theoretical perspective* is reflected by those in the theoretical computing community who seeks to understand the consequences of the shift to a world where an interconnected set of heterogamous computing devices are embedded into the world. This includes the work of researchers at Cambridge, Birmingham and Southampton who seek to reason about the consequences of this shift to a ubiquitous computer composed of this collection of devices and to understand the impact on our current theoretical approaches and models.

The *systems perspective* is reflected in the activities of a number of communities. The *communication systems and middleware community* including researchers at Imperial, Cambridge, and Lancaster seek to address the architectural and network challenges posed by the large scale, heterogeneous and dynamic nature of ubiquitous computing. The *sensors, novel devices and wearable systems community* including those at Cambridge, Lancaster and Bristol who seeks to construct digital devices to sense and interact with the world we inhabit. These devices may be embedded in the physical spaces we inhabit, carried with us as we move through the world or worn on the person.

The *experience perspective* is reflected by a broad community that wishes to understand how to realise ubiquitous computing environments that meet the needs and desires of users. Within the UK much of this work has been driven by members of the Equator IRC involving 8 different institutions and coordinated by Nottingham University which brings together researchers from a broad set of traditions. These researchers consider how people might live and interact with ubiquitous computing devices, what the interactive principle needed to underpin these devices should be and how a ubiquitous computing society might be shaped from a socio-technical perspective.

Researchers from the UK play a leading international role in Ubiquitous Computing. This is reflected in terms of the proportion of research outputs in leading research venues where the UK is second only to the US. It is also reflected in the organisation of international conferences and workshop where UK researchers routinely chair the leading international conferences.

A key feature of the UK Ubiquitous Computing community has been its broad engagement with real world users across a number of sectors. This includes domestic environments, healthcare, environmental monitoring, tourism and the entertainment and leisure sectors to name only a few. In addition to brining disciplines such as sociology, psychology, art and design closer together and blending both hardware and software disciplines within ICT the work in Ubiquitous Computing has also linked with a range of disciplines through initiatives such as the eScience programme. The commitment to user-focused research within ubiquitous computing is reflected in the development of methods and techniques that support the construction, deployment and analysis of Ubiquitous Computing in real world settings rather than research labs.

The importance of Ubiquitous Computing is also reflected within the UK grand challenges initiative where it forms one of six grand challenges identified by the UK community. The authors of the Ubiquitous Computing Grand Challenge are drawn from a number of previously quite separate computing research communities reflection that the core of Ubiquitous Computing lies in the convergence of the three different perspectives outline above. Real advances in Ubiquitous Computing depending on the successful blending of perspectives drawn from the science of computing, the engineering of complex systems and the understanding of their use in social settings. This blend of disciplines has now become a key feature of most UK ubiquitous computing projects with the inter-disciplinary approach initially adopted by the Equator IRC reflected in the majority of consortiums tackling EPSRC WINES projects.

ICT areas: Human Computer Interactions, Cognitive Science and its Applications, Software Engineering, Systems Methodology and Architecture, Systems Integration, Mobility

5.12 Theory and Foundations

The study of logical and semantical methods plays a key role in the foundations of Computer Science, and provides a basis for programming language design and specification, verification, program analysis, and structural methods in Computer Science in general.

U.K. researchers have played a major leading role in this area, arguably second to none world-wide. The lineage of major figures starts with Alan Turing, and continues with Strachey, Landin, Scott (who did his seminal work on Domain theory while at Oxford), Hoare, Milner and Plotkin. Fundamental contributions to the field include: denotational semantics and domain theory (Strachey, Scott, Plotkin), lambda calculus as a fundamental tool for the study of computation (Landin, Strachey), structural operational semantics (Plotkin), program logics (Hoare logic: Hoare), algebraic process calculi (CCS: Milner, CSP: Hoare), mobility and the foundations of global computation (pi calculus: Milner), type-checking in programming languages (Milner), interactive proof systems (LCF: Milner, HOL: Gordon, Isabelle: Paulson). This internationally leading role in the field continues into the present day. For example, the International Symposium on Logic in Computer Science (LiCS), the major conference in its field, has had UK-based Program Chairs in 2002 and 2007, and the only non-US based based General Chair to date was from the UK. Major centres in the UK include: Edinburgh (Plotkin, Stirling, Wadler, Buneman, Sanella), where the Laboratory for the Foundations of Computer Science is one of the largest groups in this field world-wide; Cambridge (Milner, Hyland, Winskel, Pitts, Fiore); Oxford (Abramsky, Gottlob, Roscoe, Ong, Baltag, Lowe, Coecke); Birmingham (Jung, Kwiatkowska, Reddy, Escardo, Ryan, Ghica, Levy); QMUL (O’Hearn, Martin, Honda); Imperial (Hankin, Gardner, Huth, Yoshida, Wicklicky); Swansea (Tucker, Berger, Mosses, Moller, Setzer). There is also a major group in this area at the Microsoft Research Laboratory in Cambridge (Hoare, Cardelli, Gordon). Some research highlights are listed below.

Mobile, Ambient and Ubiquitous Computing: From the foundational point of view, mention should be made of the fundamental contributions made by UK researchers to this area, which has had an enormous impact on the field internationally. Major landmarks include the work by Milner, Parrow and Walker on the pi-calculus, and Milner’s subsequent work both on this topic, and on action calculi and now on bigraphs; and Cardelli and Gordon’s work on Ambient Calculus.

Separation Logic: (O’Hearn, Reynolds, Bornat et al) has grown out of the work on Bunched Logic by Pym and O’Hearn, and has developed as one of the most interesting current approaches to verification of heap-based imperative programs, and now also of shared-variable concurrent programs. It has found traction on verification problems which have proved resistant to other approaches over the past four decades.Coming from a different direction, the Spatial logics arising from the Ambient Calculus of Cardelli and Gordon exhibit some similar ideas, and have led to interesting developments such as the Context Logic of Gardner et al.

Game Semantics: This was a major development in semantics in the 1990's which gave new concepts and tools for analyzing the space of programming languages, and constructing fully abstract models for a wide range of programming languages, an issue which had been intractable previously. Following the construction of fully abstract models for PCF by Abramsky, Jagadeesand and Malacaria and Hyland and Ong, Abramsky and his students showed that languages with state, control and other computational effects could be analyzed in a systematic fashion by varying the conditions on strategies. There was also a cognate use of game semantics to study logical systems Altogether this has grown into a substantial area of research, which is well represented in the leading conferences and journals. A recent workshop in Marseille had over 90 participants. U.K. research has played a central, and indeed dominant role; there are also strong French and Italian contributions.

Algorithmic Game Semantics A major development over the past five years has been the development of Algorithmic Game Semantics, in which the concrete nature of Game Semantics is exploited to yield algorithmic representations of strategies as automata. This can be used on the one hand to yield implementations of compositional software model-checkers, in which programs, possibly with free variables, are compiled into algorithmic representations of their (fully abstract) game semantics; and also, to yield results characterizing the complexity of deciding program equivalence, or model-checking problems, for various classes of problems. This work was pioneered at Oxford (Abramsky, Ong, Ghica, Murawski). The model-checking side continues both there, and at Birmingham (Ghica) and Warwick (Lazic). Ong and Murawski, working in part with leading Algorithmic Verification researchers such as Walukiewicz, have developed a stream of beautiful results using game-semantical methods to classify the complexity of deciding equivalence for various fragments of higher-order programming languages. There has also been a fruitful interplay between Ong and Stirling, who has used game semantical methods in producing a remarkable solution to a famous open decision problem from the 1970's, Huet's Higher-Order Matching problem.

The Nominal Paradigm: Pitts, Gabbay, Stark, Urban et al. have extracted the use of names from such contexts as the pi-calculus to give a fresh look at fundamental issues of variable binding, and how to embed reasoning about bound-variable notions in theorem provers. They have developed nominal logics, type theories and programming languages. Moreover, there are beautiful connections with permutation models in set theory, and associated ideas in categorical logic.

Quantum Information and Computation. The U.K. has had a major presence in this field on the Physics and Mathematics side (Deutsch, Josza, Ekert et al). A pioneering role has been played by the group at Oxford (Abramsky and Coecke), who have shown how categorical and diagrammatic methods can be used very effectively in modelling and reasoning about quantum protocols. A major European project led by Oxford, and involving both Computer Scientists and Physicists due to start in January. A workshop in this area at Oxford attracted over 100 participants. This is a very exciting interdisciplinary area, which shows the potential for Computer Science ideas and methods to be applied more widely, and for fruitful interactions between computer scientists, physicists and mathematicians.

ICT areas: Fundamentals of Computing and AI

5.13 Algorithms and Complexity

UK algorithms and complexity research has traditionally been of very high quality but the number of researchers has been relatively low with these researchers often in small research groups. The most recent International Review explicitly noted that "the UK effort in algorithms is too small in absolute terms" and that "failure to strengthen work in the algorithms area may mean that the UK will lose competitiveness".

However, since 2001, there have been changes to the UK algorithms and complexity landscape. Perhaps the most significant change is the establishment of one of the largest UK algorithms and complexity research groups at Durham, consisting of 9 permanent academic staff, when prior to 2002 there was no algorithms and complexity research undertaken at Durham. The Durham group has interests across the algorithms and complexity spectrum and including: exact, approximate, parallel and randomised algorithms; graph theory, combinatorics and discrete mathematics; computational, proof, parameterized and descriptive complexity; constraint satisfaction, satisfiability and other theoretical aspects of AI; and applications of algorithms in phylogenetics and communication. The algorithms group at Liverpool has also expanded since 2001 and is comparable in size to Durham. Its interests include: the design and analysis of sequential, parallel, randomized and distributed algorithms; algorithmic game theory and mechanism design; computational learning; scheduling, distributed computing and network algorithms; graph theory and computational complexity; and the applications of algorithms in computational biology and computational mathematics.

However, since 2001, there have been changes to the UK algorithms and complexity landscape. Perhaps the most significant change is the substantial growth in algorithms groups at Durham and Liverpool. At Durham, the group now has 9 permanent academic staff where prior to 2002 there was no algorithms and complexity research undertaken at Durham. The Durham group has interests across the algorithms and complexity spectrum including: exact, approximate, parallel and randomised algorithms; graph theory, combinatorics and discrete mathematics; computational, proof, parameterized and descriptive complexity; constraint satisfaction satisfiability and other theoretical aspects of AI; and applications of algorithms in phylogenetics and communication. The algorithms group at Liverpool, which was already large, has also expanded significantly since 2001. It currently consists of 10 permanent academic staff, with another joining in January. Its interests include: the design and analysis of sequential, parallel, randomized and distributed algorithms; algorithmic game theory and mechanism design; computational learning; scheduling, distributed computing and network algorithms; graph theory and computational complexity; and the applications of algorithms in computational biology and computational mathematics. Warwick has been given the opportunity to secure its traditionally strong reputation in algorithms by the award of a £3.8 million Science and Innovation grant to set up a Centre for Discrete Mathematics and its Applications. The Warwick project will involve the Departments of Computer Science and Mathematics and the Business School. It will focus on both the interface between mathematics and computer science and the fundamentals of operational research. Moreover, it will mean the establishment of three new lectureships and a new chair. Warwick's research interests include: the design and analysis of sequential, parallel, randomized and distributed algorithms; algorithmic game theory; and graph theory and combinatorics. Other research groups, such as Edinburgh (probabilistic computation; complexity of combinatorial enumeration; information- and complexity-theoretic aspects of machine learning; combinatorial optimisation; game theory; computational biology), Glasgow (algorithmic graph theory; matching; complexity and approximability of optimisation problems; stringology), Kings College London (algorithmic graph theory; probabilistic combinatorics and randomized algorithms; stringology; combinatorial optimisation and network flows; applications of algorithms in music analysis, molecular sequences, web-graphs, peer-to-peer computing and data compression), Leeds (sequential, parallel and randomized algorithms; graph theory, combinatorics and discrete mathematics; scheduling and combinatorial optimisation; computational complexity; applications of algorithms in computational biology), Leicester (design and analysis of algorithms; algorithm engineering; formal language theory, computational complexity and decidability; algorithmic problems in algebraic structures; applications of algorithms to optical and ad-hoc networks, railroad optimisation, bioinformatics, text indexing, representing semi-structured data and network analysis), Oxford (sequential, parallel and quantum algorithms; graphs and hypergraphs; computational and descriptive complexity; constraint satisfaction; applications of algorithms to computational biology) and RHUL (design and analysis of algorithms; computational complexity; graphs and combinatorics; combinatorial optimisation; linear and integer programming; constraint satisfaction) have remained more or less as they were in terms of numbers, but are currently smaller than Durham, Liverpool and Warwick. There are also research groups in Mathematics departments, primarily at LSE, Oxford and QMUL, whose research (in discrete mathematics, combinatorics and graph theory) is strongly related to that undertaken by the algorithms and complexity community residing in UK Computer Science departments.

The profile of algorithms and complexity in the UK has been recently enhanced by the hosting of international conferences in the UK, e.g., ACiD at Durham in 2005 (with WG to come in 2008), SPIRE at Glasgow in 2006, SAGA at Kings College London in 2005, SIROCCO at Liverpool in 2006, and so on. Also, the Journal of Discrete Algorithms, now published by Elsevier, originated in the algorithms and complexity research group at Kings College London and is strongly represented on its editorial board with UK-based editors. In summary, the research quality of the UK algorithms and complexity community remains high and there has been a general expansion in the number of algorithms and complexity researchers in permanent positions in UK universities. However, this expansion does not mirror the international growth in the subject; for example, the number of submissions to many algorithms and complexity conferences has doubled in recent years and numerous new conferences in algorithms and complexity and its applications have recently been established.

ICT areas: Fundamentals of Computing and AI

5.14 Vision and Medical Imaging

The UK has a strong track record in computer vision (CV) and medical image analysis (MIA), which has been further enhanced over the review period. In terms of international impact, Oxford (Brady+, Zisserman*), Manchester (Taylor*, Cootes*), Surrey (Kittler*, Matas), UCL (Hawkes+, Hill+, Cox*), and Microsoft Research (Blake*, Bishop*, Fitzgibbon) dominate, with 7 (*), out of 42 from the field, in the top 0.1% most cited authors in computer science (CiteSeer: http://citeseer.ist.psu.edu/, weighted or unweighted, validated by ISI citations), and 3 more (+) who are comparable in terms of journal citations (> 1000 citations, ISI Web of Science: http://portal.isiknowledge.com). There is also internationally leading CV/MIA research (CiteSeer top 1.0% and/or >200 ISI citations for top ten publications) at Imperial (Rueckert, Hajnal, Yang), York (Hancock), Birkbeck (Maybank), KCL (Petrou), Edinburgh (Fisher, Williams), Cambridge (Cipolla), Oxford Brookes (Torr), Southampton (Nixon), Leeds (Hogg), and QMW (Gong). UK authors have won best paper prizes at 4 of the 7 leading international vision conferences (ICCV/ECCV alternating) since 2000. Since the last review, important new drivers have emerged, creating new opportunities and challenges. Advances in processor and memory technology have finally brought advanced computer vision to the desktop, creating the opportunity for widespread deployment. The rapid expansion in digital media has created a demand for intelligent methods for searching and indexing images and video. At the high-end, there is also an increasing demand for vision-enabled augmented reality in film and TV production. A growing preoccupation with public safety and crime prevention, together with extensive deployment of CCTV, has created a strong demand for automated surveillance and monitoring. The increasingly pervasive role of imaging in clinical medicine and biomedical research has created a demand for methods of quantitative analysis to elucidate disease mechanisms and provide measures of disease progression, particularly in drug development where image-based biomarkers are now of major importance to the pharmaceutical industry.

Although CV and MIA have always benefited from application pull, world-class research in the field is characterised by the development of generic theory and methods. Important trends over the review period are a significant convergence of both CV and MIA with machine learning, convergence between CV and graphics and, internationally, a regrettable divergence between CV and MIA (though this has been addressed in the UK – see below). There has also been a growing recognition of the importance of large-scale experimental evaluation, leading to more scientifically grounded research. The UK has provided leadership in all these areas.

In CV, statistical learning has become established as the dominant paradigm, with important contributions from most of the leading UK groups. Tracking and recognising faces and facial behaviour ha become major areas of research, with UK groups taking a leading role (Surrey: Kittler, Matas. Hilton; Manchester: Cootes, Taylor; Cambridge: Cipolla; QMW: Gong; Leeds: Hogg). Similarly whole-person tracking and behaviour recognition have received significant attention, again with significant contributions from UK researchers (Microsoft: Blake; Leeds: Hogg; Oxford Brookes: Torr; Surrey: Hilton; Southampton: Nixon. The expansion of digital media has led to renewed interest in generic object recognition. Again, UK groups have taken a leading role (Oxford: Zisserman; Oxford Brooks: Torr; Surrey: Matas, Kittler). Scene reconstruction from uncalibrated video has continued to receive considerable attention, particularly for applications in augmented reality, with important contributions from UK groups (Oxford: Zisserman; Microsoft: Blake, Fitzgibbon; Oxford Brookes: Torr; Cambridge: Cipolla; Surrey: Hilton). A major programme has also been funded under the Basic Technology programme, with the aim of developing comprehensive computational models of natural vision, building on insights from CV (Surrey/KCL: Petrou).

In MIA, EPSRC and MRC have funded an Interdisciplinary Research Collaboration (IRC), bringing together the strengths of Oxford, UCL, Manchester and Imperial in medical image and signal analysis, with long-term funding (6 years) and shared scientific goals. This has helped to keep the UK at the forefront of a rapidly evoloving field, and has led to substantial knowledge/technology transfer activity (see below). It has also served to keep a close connection between MIA and CV in the UK, in contrast to the international trend. Major scientific contributions from the IRC have been in unifying image segmentation, registration, and statistical modelling (Manchester: Taylor, Cootes; Oxford: Brady; UCL: Hawkes, Imperial: Rueckert), intelligent acquisition (UCL: Hill; Imperial: Hajnal; Oxford: Noble), and multiscale model-based interpretation of brain physiology (UCL: Delpy; Oxford: Tarassenko). There is also important work, associated with the IRC, on inferring the wiring of the brain from diffusion tensor imaging (UCL: Alexander; Manchester: Parker; Oxford: Brady/Behrens), and image-based biomarkers (Manchester, UCL, Imperial, Oxford).

ICT areas: Image and Vision Computing

5.15 Bioinformatics and Systems Biology

Since the last review, there are now many Bioinformatics groups in the UK, and several Chairs in Bioinformatics including the four originally funded by the EPSRC (Imperial, Manchester, Oxford and UCL).

Several UK groups are actively involved in the development of database standards and associated databases, for example for the ArrayExpress database which conforms to the MIAME standard for microarray data (Brazma at the EBI) and the FuGE model for functional genomics (Jones at Manchester). Many groups work on the development of database technologies applied to biological data, including query languages and data integration, for example Poulovassilis at Birkbeck and Buneman at Edinburgh. The development of ontologies for biological data has been led by the Manchester group (Goble, Paton and co-workers). Distributed information management, with a focus on the GRID is a strength (Paton at Manchester). Activities in biomedical text-mining are also strong (Rebholz at EBI, Webber at Edinburgh).

Groups which are active in the area of algorithms and methods for sequence analysis include KCL (Iliopoulos et al), and Royal Holloway College (Gammerman and co-researchers). Research on protein structure function and classification include using uncertainty (Bulpitt and Pickering, Leeds), the Threader robust structure prediction and classification (Jones, UCL), and the TOPS system using graph data structures and motif discovery (Gilbert and Viksna, Glasgow).

Research in the development and application of the Web and Grid to support the analysis of biological and biomedical data is a particular strongpoint in the UK, partly driven by funding initiatives by the Research Councils. Major achievements are results from the Mygrid project, which has developed a comprehensive loosely-coupled suite of middleware components specifically to support data intensive in silico experiments in biology, led by Manchester (Goble and co-workers) and including work at Newcastle (Wipat) and Nottingham (Greenhalgh). S Work at the University of Ulster (Dubitzsky and group) includes the development for data-mining tools and services for Grid environments.

Machine learning approaches have been heavily used in Bioinformatics, e.g. support vector machines for the comparison and classification of protein and DNA sequences (Gammerman and group, Royal Holloway College) and Inductive Logic Programming for the prediction of 3-dimensional structure of molecules and their biological activity (Muggleton, Imperial) (King, Aberystwyth). Artificial neural networks have been used by Ball at Nottingham Trent to analyse mass spectrometry profiles for the identification of biomarkers representing strains of microbial pathogens. King and his team at Abersytwyth have pioneered the Robot Scientist, which uses Active Learning in an iterative approach to experimentation where knowledge acquired from a previous iteration is used to guide the next experimentation step. Girolami at Glasgow has developed probabilistic modelling and inference techniques for interacting groups of genes implicated in the onset of certain breast cancers.

Process calculi for computational biology have been developed by Cardelli at Microsoft Research Cambridge; nodels of signalling pathways have been developed using the PRISM model checker (Kwiatowska, Birmingham; Calder, Glasgow), the performance evaluation process algebra PEPA (Hillston, Gilmore, Edinburgh; Calder, Glasgow) and Petri nets (Gilbert, Glasgow). A significant emerging area of interest is multi-scale modelling of biochemical systems from the individual to population (Calder and Hillston; Lio', Cambridge), including spatial dimensions at the intra and intercellular levels, for example the use of agents to describe the epitheliome (Smallwood, Holcombe and Pogson at Sheffield). Work at the interface with mathematical biology includes modelling tissues such as the heart and breast as well as the respiratory system (Gavaghan, Oxford) and modelling plant growth (Bangham, East Anglia).

Computer scientists will be involved in the recently established (BBSRC and EPSRC funded) Integrative System Biology centres at Manchester, Imperial, Newcastle, Edinburgh, Nottingham and Oxford.

ICT areas: Fundamentals of Computing and AI, Information and Knowledge Management

5.16 Decision Support Systems and Evolutionary Computing

The UK is world leading in the investigation of computational methods and techniques to underpin automated decision support. The field lies firmly at the interface between Computer Science and Operational Research and its application areas cut across many disciplines including engineering, computational chemistry, manufacturing, business, bioinformatics and medicine. It is primarily concerned with the investigation of innovative modelling approaches, search methodologies and algorithm development for complex problem solving environments. Typical problems addressed by UK researchers in this field include transport scheduling, production scheduling, maintenance scheduling, call-centre scheduling, personnel rostering, timetabling, cutting/packing, protein folding, routing and computational finance. The UK has been particularly strong in the development of heuristics, metaheuristics (including evolutionary methods), multi-objective approaches, constraint based methods and mathematical programming techniques. The UK's work in this area (on scheduling, heuristics and mathematical programming) was highlighted by the international review of Operational Research as a particular strength of the UK. The field is well supported, mostly by three programmes: ICT, Engineering and Mathematics which is further evidence of the inter-disciplinary nature of the area. For example, the Automated Scheduling, Optimisation and Planning group at Nottingham currently has over £9M of active awards including approximately £7M of EPSRC funding and other significant funding from BBSRC, industry and the EU. This includes the largest ever responsive mode grant awarded to a single computer science department (£2.6M) and funding for teams at Essex, Leeds and Coventry. UK researchers in this field are actively involved with industry. For example, Nottingham has collaborated with National Air Traffic Services Ltd, KLM Airlines, Tesco; its nurse rostering algorithms have been used in over 40 hospitals in Belgium and it has two spin out companies (event Map Ltd and Aptia Solutions Ltd). The School of Computing Sciences at UEA has worked closely with Lanner Group Ltd to develop a seat planning and scheduling application for a commercial call centre, currently in use by one of the major banks. Leeds has worked for many years on vehicle and crew scheduling problems in the transport industry, e.g. with over 40 UK rail and bus companies including First, GNER, Arriva, Virgin, and National Express. The University of Essex has worked closely with BTexact for a number of years and the team at Coventry has worked on production scheduling problems with industry and on radiotherapy scheduling with a large local hospital (both initiatives being funded by EPSRC).

There is a significant amount of international editorial and conference organising activity in this area in the UK. For example, the Journal of Scheduling was launched (almost 10 years ago) from the team at Nottingham and UK researchers have played a major role in the organisation of a number of international conferences and workshops including the international conferences on Computer-Aided Scheduling of Public Transport (CASPT), Practice and Theory of Automated Timetabling (PATAT) and the Multi-disciplinary International conference on Scheduling: Theory and Planning (MISTA).

The field of evolutionary computation has some overlap with automated decision support systems and artificial intelligence; the UK has a very strong community of researchers in this area. The leading international journal in the area (IEEE Transactions on Evolutionary Computation) has 9 Associate Editors (from at total of 47) from the UK in addition to the Editor-in-Chief. Examples of activity include a major collaborative grant (£2.5M) between King's College London, Birmingham and York on search based software engineering and University of Birmingham established the Centre of Excellence for Research in Computational Intelligence and Applications (CECIA), funded to the value of over £2M by the West Midlands Development Agency.

ICT areas: Fundamentals of Computing and AI, Information and Knowledge Management, New and Emerging Computer Paradigms, Artificial Intelligence Technologies

6 Concluding remarks

The UKCRC appreciates the efforts made by EPSRC staff to engage with the community and, as already discussed, UKCRC members generally find interactions with the EPSRC at all levels to be positive and constructive. We recognise that the EPSRC is subject to external political and financial direction which it must respond to and that some of the issues raised here stem from such directives rather than internal decision making processes.

The UKCRC and its individual members are delighted at the opportunity to assist the work of members of 2006 International Review. This submission is intended as a preliminary survey and summary of the points which the review may wish to cover. We look forward with interest to the conclusions of the current review, and hope to continue our functions in assisting in their effective implementation.