PROCEEDINGS OF THE INTERNATIONAL CONFERENCE
ON FIFTH GENERATION COMPUTER SYSTEMS 1988,
edited by ICOT.

UK IKBS PROGRAMMES

Timothy Walker

Information Engineering Directorate
United Kingdom

ABSTRACT

This paper reviews the progress of previous UK programmes in IKBS and describes the new programmes which have recently been started.

I am sure that all of us here have found this conference a fascinating experience. The quality and range of the papers, and of the discussion, has demonstrated the results of work in this area from around the world. As well as reports on the progress of the Fifth Generation Programme, the immediately preceding papers have brought you up to date with plans for the European programmes ESPRIT and for the MCC in the US. My talk reports on UK programmes, both Alvey and its successor.

Indeed it may be worth stressing that there is a successor to Alvey and a substantial one at that. We have new organisational arrangements which have brought together for the first time virtually all Government funding for research in IT R&D, covering support for long term basic research in universities, for collaborative research between industry and academia, for Government support for industrial research and awareness and for higher education and training in IT. This allows us to take an overview of all Government activity of this kind in IT and to produce a more coherent and balanced programme than might otherwise be possible. I am glad to say that at least one member of the committee overseeing these programmes, Professor Robin Milner, is playing a prominent role in this conference.

I would however like to start by commenting on the names of the various programmes and the way that they have changed over the years. It was, I think, Alan Turing in his paper "Computing Machinery and Intelligence" who first proposed a definition and test for "machine intelligence" but the term "Artificial Intelligence" was coined by John McCarthy in 1956. A considerable effort was devoted to developing computers which could play games, particularly chess, but the term artificial intelligence obviously has a wider application. In turn it has captured the imagination and incited hostility. While it is, I know, supposed to be unlucky to quote Shakespeare's Macbeth, I am reminded of Macbeth's question of the witches ``say from whence you owe this strange intelligence.''

It has provided a continuing subject for science fiction from the original intelligent computer novel by Olaf Stapleton in 1930 to Asimov's robots in the 1950s, Arthur Clarke's 2001 and most recently C3PO in Star Wars. At the same time the sometimes overstated claims of its protagonists have encouraged others to argue that artificial intelligence is a contradiction in terms. The philosopher John Searle is perhaps the most recent, although I have to confess that I found his argument tended to knock down straw men rather than those of more substance.

It is perhaps not surprising therefore that more recent programmes have avoided the term "artificial intelligence". In Japan the term "Fifth Generation Programme" was chosen and it is interesting that while the titles for papers at this conference often include the words "intelligent", "reasoning", "knowledge" (and even "ignorance"!) none of them refer to "artificial intelligence". The terms used now by practitioners tend to be less grandiose; Professor Aiso describes Fifth Generation Technology as "the nucleus of knowledge information processing"; ESPRIT I used the term "Advanced Information Processing" but its successor ESPRIT II has dropped the "advanced" to use "Information Processing Systems`. In the UK we have used "Intelligent Knowledge Based Systems" or IKBS for short.

Some of you may feel that dwelling on the names of these programmes is an irrelevance but I am not so sure. I am certain that there has been a deliberate attempt to distance the community from the more exaggerated claims of the proponents of artificial intelligence. I believe also that the names recognise that the new techniques are much less likely than had been supposed to replace humans, rather they will be used to assist human decision taking. It is from this change in emphasis that the new names have been formulated, most of which concentrate on the processing of knowledge or information.

Certainly this was the thinking that underlay the Alvey report and the programme in IKBS which stemmed from it. An IKBS was defined as a system which uses inference to apply knowledge to perform a task and the Alvey report and the subsequent strategy set out the research that was needed to develop such systems. At that stage, like the Fifth Generation Project, it was seen as a 10 year programme but unlike the Fifth Generation Project, the Alvey programme was not tied to specific goals for hardware or software. It concentrated on tools and methods, sometimes described as enabling technologies, that is those technologies that could then be applied by companies for the development of specific products.

Of course the UK had to start from where we were. At that time almost all the expertise was in universities so much of the development of the IKBS part of the Alvey programme had been done in universities. However the arguments about the nature of artifical intelligence and the usefulness of research in this area which followed the Lighthill report had led to a somewhat demoralised academic community.

Funds had been reduced and some of the most distinguished UK academics had left to work in the US or elsewhere. The overall aims of the IKBS part of the programme were first to revitalise and strengthen the academic groups in the UK and second to achieve some technology transfer into industry. This has been characterised in a recent review of the programme (by Erik Arnold of the Science Policy Research Unit of the University of Sussex) as a science policy rather than an industrial policy. While this has some basis, it neglects the industrial relevance of technology transfer.

The overall strategy of the Alvey IKBS programme therefore had three elements. These comprised a research and development element, an awareness and promotional element and a general support element. These were set out in 1983 as follows;

The Research and Development Programme

- "Show Me" projects providing immediate industrial demonstrations of existing IKBS technology.

- Short Term Development projects aimed at producing marketable products within 3-4 years based on low risk developments from current technology.

- Demonstrator projects involving industry-led collaborative R&D aimed at building complete prototypes of possible future systems requiring substantial R&D progress.

- Research Themes involving medium to long term directed research by collaborating teams on a limited number of carefully chosen topics likely to be crucial to success in building advanced IKBS.

- General Research accelerated progress in the directed part of the R&D programme will only be possible against a background of a well balanced portfolio of high quality, speculative research into all the important aspects of IKBS and related topics such as software technology and man-machine interaction.

The Awareness and Marketing Programme

In parallel with the R&D activities, this programme is aimed at educating and informing a wide range of possible users of IKBS about their potential, to create an informed marketplace for the products their potential, to create an informed marketplace for the products of the R&D programme as they develop.

In particular, the awareness programme will begin by addressing UK industries where IKBS could have significant impact on productivity, quality and performance in the near term, and will solicit the involvement of selected firms in the definition and implementation of particular demonstrator projects in the R&D programme.

The Infrastructure and Support Programme

- provision and support of common computing hardware for IKBS researchers in industry and academia

- development of common frameworks and standards for interchange of software,

including:

- communications network facilities
- software libraries
- software quality assurance
- updating and distributing catalogues of relevant software available from other countries
- an expanded education and training activity to increase the flow of high quality people available for IKBS research, development and production elsewhere which may be relevant to the UK IKBS community

The achievements of Alvey IKBS have been considerable. Significant progress has been made towards reinforcing existing centres of research excellence and a number of key individuals have been encouraged to return from the UK or have decided not to leave in the first place. Industrial interest and capability in AI has been raised; the number of people in industry with relevant experience has risen dramatically. Links between the universities and industry have been improved and academics have a better sense of industrial needs and the intellectual challenge of scaling up laboratory techniques into useable industrial systems. At the peak, about 350 researchers were employed on Alvey projects in addition to existing members of University faculties.

It may be worth saying a little more about the Research Clubs since these have stimulated very considerable industrial interest in expert systems. There were nine in all covering a wide range of activities:

ALFEX	Alvey Financial Expert Systems
ARIES	Insurance Community
DAPES	Data Processing
EMEX	The Econometric Model Building
PLANIT	The Planning IKBS Club
QSES	Quantity Surveying
RESCU	Real Time Expert Systems
TRACE	Transport and Travel
WEISC	Water Industry

In each case membership cost only a few thousand pounds and allowed members to specify an expert system for use in their own industry and to observe its evolution and development in parallel with others with a similar interest. They also have the right to use the resulting software without further payment. Most of these users had no previous experience of expert systems and therefore it provided very useful experience for them in de-mystifying the technology as well as learning how to deal with specification and use. Altogether the nine have more than 180 member firms and have spread awareness of expert systems through a significant proportion of UK industry and commerce. Many of the clubs are now continuing without Government support.

Another successful initiative was the Journeyman scheme whereby individual industrial workers were seconded into university teams for a few months in order to deveop their own understanding of IKBS techniques to a stage where they could develop industrial projects. Many people have passed through this scheme and this has been a key factor in extending the range and numbers of industrial researchers knowledgeable about IKBS.

At the same time, and perhaps surprisingly for so long term a research programme, there have been significant commercial outputs from the research, not just in expert systems but also in the IKBS aspects of speech and image processing. I have also been struck by the increased interest being shown in industry in natural language, although it is still clearly some time before computers will be able to deal quickly with anything other than a fairly restricted set of language.

In considering the success of Alvey, my own experience gives me a useful perspective. As some of you will know, although I helped to start both the Alvey and the Esprit programmes and worked on them during their first two years, I then moved to other work before returning as Director of the new programme. It is perhaps easier therefore for me to notice changes in the field of IKBS than for those who have been associated with it during the whole period. And indeed I have noticed a number of changes.

The first is that expectations of what might be achieved have been lowered. This has been emphasised by the development of expert systems many of which are now in use commercially as a result partly of the stimulus provided by the Alvey Clubs. But despite this growth, they are put to less ambitious uses than perhaps had been imagined a few years ago. In most cases they are used chiefly to store the knowledge or know how of experts in a form easily recalled by those less expert. This seems to me an eminently sensible use and one that should not be interpreted in any way as second best. I am conscious that I may well be using "knowledge" in a rather loose way but to define it here would take the rest of my paper at least. I am content to rely on

Dr Jowett, the Master of Balliol College, Oxford, during part of the 19 century. An apocoyphal verse runs:

First came I: my name is Doctor Jowett
There is no knowledge but I know it
I am the Master of this college
What I don't know, isn't knowledge

Second, in the same way that expectations of IKBS have been reduced, it is becoming increasingly recognised that it is one tool, to be used with others. When I first became involved with Alvey, I was conscious of fierce arguments between the proponents of IKBS and the software engineers about the best way to produce software. Each side regarded the other as misguided and either irrelevant, incompetent or clinging to the past. I was struck by the relative absence of these arguments when I returned to the area. There was a much greater recognition that both approaches had real advantages - and disadvantages - and that the way forward was to devise ways of combining the two approaches. Thus most IKSEs now aim to provide an environment that can support both kinds of activity. I see a similar situation in the proponents of LISP and PROLOG, who now seem also more interested in working together than pointing out the faults in the other approach.

Third, the area of human computer interface (HCI) is receiving much more attention and in a more systematic way. Five years ago it was recognised as important but no one really knew what to do about it. Now we accept that there is no point in devising means to store and process knowledge if we do not have effective techniques for extracting that knowledge from the experts. HCI is becoming an integral part of many different technical areas.

This also leads me on to one of the other major developments in the Alvey Programme, that of systems architectures. Computer architecture has, of course, always had an important part to play in computer science, but has become of increasing importance with the advent of parallel processing and distributed systems. Architectures also need to be more flexible - dare I say more intelligent - and to be able to relate to the system or organisation with which they are dealing. It is interesting that the Alvey Report hardly mentions architectures but that within 18 months the Alvey Directorate found it necessary to create it as a topic in its own right. While the research was naturally concerned at first principally with theoretical studies, it has now moved on to working with systems or subsystems. One of the most innovative projects, and indeed the only Alvey project where the collaborators all worked at the same location, was ANSA, Advanced Network System Architecture. Although this started as a communications project it has become possibly the single largest driving force behind the ISO model for "open distributed processing". Interestingly it too has become involved in issues such as requirements elicitation and organisational structure. Like many of the Alvey projects ANSA is moving into Europe and will, having taken on other collaborators, become an Esprit project.

All these factors have influenced the new UK programme in terms of both structure and content. The new programme is structured less around individual technologies and more by function. Thus instead of IKBS, software engineering, MMI and so on, we have one area devoted to systems architecture and another to systems engineering, the latter of which contains both IKBS and software engineering as well as human factors. There is of course some overlap between systems engineering and systems architecture but this would be inevitable wherever the boundaries were drawn.

Some of the work will of course be carried out in the European programme, ESPRIT. My paper earlier this week explained the criteria we use for deciding whether a project is more appropriate for a European or national programme and I will not repeat them here but will confine my comments to what is proposed for the national programme.

The major themes of the programme are:

the industrialisation of the techniques developed over the last few years, in Alvey, Esprit and elsewhere;

the promotion of interdisciplinary work; and

the integration of IKBS into the generality of approaches to the building of complex systems.

Thus while containing substantial element of academic research, the programme has clear industrial objectives. In particular it is designed to capitalise on the achievements of previous programmes. Conversely, those areas such as expert systems which are being exploited commercially, are not included in the programme. There will however be scope for more advanced work on particular classes of expert systems, for example these dealing with real time or cooperating or distributed

expert systems.

What then does the new IED programme concentrate on? The first point to note of course is that we do not have a IKBS programme as such. One of my first acts on taking over Director was to ask the three people in charge of the IKBS, software engineering and HCI parts of the Alvey programme to produce a joint strategy for the new programme rather than three separate ones. Despite some initial misgivings they did this and it has been welcomed by the community. I will however concentrate in this talk on those aspects of new programme that would normally be recognised as IKBS.

Like Alvey, it concentrates on basic science, methods and tools and in demonstrating their use in practical industrial environments. There are no specific hardware goals. The basic science will concentrate on cognitive science, particularly on inference and reasoning mechanisms and deep knowledge representation, and on recent developments in logic, including those necessary to deal with more advanced systems needed for concurrent system design or for real time and predictably reliable systems. Models also need to be developed for the entire life cycle of software and for simulation of human users.

The theme of bringing software engineering and IKBS together continues in the new programme's work on IKSEs. The Alvey programme supported the development of both second generation IKSEs and the Poplog environment, all of which are now commercial products.

However it is clear that more advanced IKSEs require very substantial investment and must provide acceptable migration paths for new users or those used to different approaches. Our work will emphasise the development of a better balance between the knowledge based and conventional elements, we need to incorporate the results of research conventional elements, we need to incorporate the results of research into life cycle models, project management and development methodologies into the next generation of knowledge based tools. A further area where this kind of interdisciplinary work is needed is in the combination of formal methods and declarative languages.

The more techniques are used which allow a system implementation directly from specification, the more important becomes the development of that specification and the requirements capture which precedes it. This is closely related to the knowledge elicitation and knowledge representation necessary for IKBS programme. Indeed it can be argued that requirements capture is the elicitation of user knowledge about the desired solution and that system design is the bringing together of that knowledge with a designer's knowledge about the design options available in a suitable knowledge representation framework.

Despite these similarities, there has been relatively little overlap between the two approaches although the use of formally based specification languages and automated reduction systems to reason about characteristics of the specification or to infer realisation of the specified system is common ground. These techniques will need to exploit the potential of massively parallel computing engines and therefore provide a natural link to the systems architecture part of our programme. We hope in the new programme to bring the various communities together to make use of the insights which each of them have gained in the process.

This leads one to another very important problem, that of validation. Traditionally this has been addressed in the software engineering programmes and rather less so in those dealing with IKBS. Yet there are some very fundamental problems which have significant commercial and legal implications. This is already pressing for expert systems. Indeed, as far as I am aware, there is still no way of establishing unambiguously that the contents of the system do indeed represent accurately the knowledge of the expert(s) concerned. (It is interesting to notice that there was a paper on this subject yesterday). The problems that stem from this relate to responsibility for errors produced by the expert system; are these due to inadequacies in the expert's knowledge; in the way the expert system has dealt with them; or in the use to which the system has been put? It is possible to argue that this is not really different from the arguments about whether faults arise from an inadequate specification rather than an inadequate implementation in a conventional piece of software, but the involvement of the expert, often from neither the buyer's or vendor's organisation, does cause extra difficulties. No doubt it will give lucrative employment to lawyers for a long time to come yet!

The final topic I would like to deal with is that of human factors. Although there have been few, if any, papers directly on this subject during this conference, the

area is almost always in the background, particularly when dealing with knowledge elicitation or, of course, natural language. There are constant arguments between those who regard human computer interaction as an independent subject and those who feel that it can only be investigated as part of other related topics. The new UK programme tends towards the second view, but not exclusively. Certainly we have chosen to produce a workplan which seeks to integrate HCI with IKBS and software engineering, but which does suggest some topics on HCI alone.

This is one change from Alvey; the other is the increased emphasis on HCI at a organisational level rather than with individual people. This reflects a recognition at a specification level of Donne's phrase "no man is an island" and is of particular importance in knowledge based approaches where much of what is called "knowledge" includes assumptions about how other people in an organisation will react. I am sure that we need to do more work on the effects on organisational structure of the continued increase in the automation of information processing.

To sum up therefore, the new UK programme for IKBS has deliberately attempted to integrate these techniques both with software engineering and with human factors and to accelerate the industrialisation of these methods. It recognises explicitly that these techniques will be used to help and not replace human beings, It will also continue to build up the pool of skilled manpower in this area.

We look forward to reporting the results at future Fifth Generation Programme Conferences.