



Mathematics at the Interface of Computer Science and Industry



The Smith Institute for Industrial Mathematics and System Engineering

EXECUTIVE SUMMARY

The design of new products, the delivery of new services, and improvements in operational efficiency rely in large part on the successful industrial exploitation of computing and information technology. These capabilities rely in turn on the underpinning disciplines of computer science and mathematics, coupled to effective systems engineering.

This report explores significant areas of opportunity, where computer science and mathematics can be brought into closer collaboration, enabling the creation of tools and technologies that will boost business innovation and efficiency. Valuable scientific benefits will also accrue, through using a set of industrial challenges as a catalyst to correct the current disconnection that exists in the UK science base between mathematics and computer science.

An extensive survey of leading industrialists has highlighted three broad themes to which there should be attached a particular priority: **network behaviour, algorithmics** and **information management**. These are the most prominent areas where industrial requirements for the next few years cover a range of business sectors and demand a more coordinated response than is presently available. They are also areas in which there is evidence that the UK is falling behind its international competitors in North America and in Europe. The report does not dwell so heavily on areas such as scientific computing and cryptography, which are equally important but where there are already strong linkages between computer science and mathematics and between the science base and industry.

The tripartite interface of industry, computer science and mathematics is diverse and inherently multidisciplinary. It will become much more effective if closer working relationships can be fostered. Our findings have led to the following recommendations for action in the near term, which are designed to be important first steps in establishing new collaborations.

- There should be closer interaction between the relevant organizations that help to steer strategic direction in the science base. In particular, the Smith Institute and the UKCRC are eager to develop closer ties, which will encourage computer scientists and mathematicians to participate jointly in industrial workshops and briefings in the priority themes.
- In conjunction with its industrial and academic partners, the Smith Institute will identify a set of industrial topics to be the initial focus of these activities. These could include bioinformatics in the pharmaceutical sector, brand management in the retail sector and traffic management in the transport sector, along with cross-sectoral topics such as information retrieval and remote sensing.

The aims in the longer term must be to extend and strengthen significantly the links that exist in the UK science base between mathematics and computing, and in doing so to establish a critical mass of activity by

international standards, coupled to emerging business opportunities. Companies will then be in a correspondingly strong position to exploit the computing and information technologies for competitive advantage.

CONTENTS

EXECUTIVE SUMMARY	i
1 INTRODUCTION	1
2 CURRENT LANDSCAPE	5
3 CHALLENGES AND OPPORTUNITIES	11
3.1 Network behaviour	11
3.2 Algorithmics	16
3.3 Information management	21
4 ESTABLISHING EFFECTIVE COLLABORATIONS	27
<hr/>	
THE MICSI SOURCES	31
USEFUL WEBSITES	32
ABBREVIATIONS	33
ACKNOWLEDGEMENTS	34
CONTACT DETAILS	36
<hr/>	
SNAPSHOTS	
Computational physiology	3
Performance models for software design	4
Structural safety	8
Bandwidth sharing	13
Load management in distributed systems	14
Network diameters	19
Decision support for cutting and packing	22
Recognizing subgraphs	24
Language modelling for information retrieval	26
On-line bin packing	28

1 INTRODUCTION

Background to the MICS I Study

The Smith Institute for Industrial Mathematics and System Engineering works with companies from sectors as diverse as food, materials, transport, aerospace, telecommunications, healthcare and energy, enabling them to acquire new knowledge and expertise from the exploitation of mathematical modelling and analysis.

In 2002, the Institute published the roadmap *Mathematics: Giving Industry the Edge*¹, which explored the areas in which the mathematical and computational sciences will, if properly developed, drive future innovation in the UK economy. The roadmap highlighted six key themes: simulation, multidisciplinary design, uncertainty and risk, data, complex systems, and market behaviour. The roadmap themes cut across traditional academic boundaries and have provided a valuable framework for assembling new collaborations.

As we shall see, there is evidence that by international standards the UK is weak in the coupling between mathematics and computer science. A stronger interface would allow the UK to derive greater competitive advantage from its science base. The MICS I study (Mathematics at the Interface of Computer Science and Industry) examines this challenge more closely, by addressing the following questions:

- In which areas can mathematics come together with computer science in multidisciplinary partnership to achieve greatest industrial benefit?
- What are the barriers to innovation in areas that require expertise in both mathematics and computer science?
- What mechanisms can enable improved coupling at the three-way interface between computer science, mathematics and industry?

Innovation in this context has the meaning of the DTI's 2003 Innovation Report², namely 'the successful exploitation of new ideas'. It encompasses the development of new products, processes and services, and also improvements to existing ones. The MICS I study has looked at a wide range of business sectors, which together make up a community of end-users for new technologies based on mathematics and computer science.

In the areas of interest to the MICS I study, innovation is intimately related to successful systems engineering. Systems engineering is generally an iterative process of requirements capture and analysis, concept generation and analysis, systems design, systems development and assembly, integration, testing, verification and validation. Mathematics has traditionally contributed most heavily at the stage of concept

¹ *Mathematics: Giving Industry the Edge*, <http://www.smithinst.co.uk/News/Roadmap>

² *Competing in the Global Economy: the Innovation Challenge*, <http://www.dti.gov.uk/innovationreport/>

analysis, using mathematical models to evaluate and refine possible ways forward, and to understand the interactions between key parameters. There is also a rapidly growing, more strategic role in requirements analysis and concept generation. A closer interaction with computer science will lead to new capabilities at the other end of the systems engineering lifecycle, to systems design, testing, verification and validation.

We are often dealing with a knowledge supply chain, in which mathematics and computer science enable the IT and telecommunications sectors to supply end-users with new tools and services. It is important to consider such supply chains as a whole, and to recognize that points of weakness can lead to dramatic failures. For example, in August 1991 the Sleipner A offshore platform sank during a controlled ballasting operation off the coast of Norway, causing an event measuring 3.0 on the Richter scale, leaving a large pile of debris on the sea-floor, and costing \$700 million. The cause was traced to a flawed finite-element implementation of a model for linear elasticity. When corrected, the code predicted precisely what had happened¹.

Structure of the Study

The MICS I team has conducted interviews with over 50 leading industrialists and academics, to clarify the current UK landscape and the issues that are currently driving different business sectors. Three themes have been identified where stronger links between computer science and mathematics have greatest potential for boosting innovation. They are **network behaviour**, **algorithmics**, and **information management**.

This report examines these themes and suggests how progress can be made in meeting the challenges that they present. It is illustrated throughout using 'snapshots' of recent research, highlighting for the most part work that is taking place in the UK. These snapshots provide a taste of the growing body of relevant expertise in the science base. They represent the ideas that will be exploited as part of the next generation of industrial innovation. Each snapshot is introduced by indicating the know-how on which it is based, the capabilities that it is developing and the applications that it addresses.

¹ See <http://www.ima.umn.edu/~arnold/disasters/sleipner.html> for further details

Computational physiology

- **Application:** Clinical diagnosis; in-silico drug testing
- Capability:** Hierarchical and multidisciplinary modelling
- Know-how:** Systems biology; multiscale modelling; visualization; simulation

By bringing together recent developments in computational simulation, multiscale modelling, ontologies and visualization, there is now the prospect of linking the physiology of cells, tissues and organs with genomic and proteomic databases. The result will be so-called 'physiome models', which provide integrated quantitative descriptions of gene function, tissue properties, complex geometries and fluid flow. Recent work on the heart and lungs will be extended to other areas of physiology. Physiome models can be used as a computational setting for drug testing, and by connecting them to the gene sequences, CT scans and MR images for individual patients, they also hold out the possibility of model-based patient-specific diagnosis. The multidisciplinary effort required extends far beyond mathematics and computer science, to include physiology and bioengineering in a programme that has all the hallmarks of a grand challenge.

Reference:

● Edmund Crampin, Matthew Halstead, Peter Hunter, Poul Nielsen, Denis Noble, Nicolas Smith and Merryn Tawhai, Computational physiology and the physiome project, *Experimental Physiology*, **89**, 1-26 (2004)

Performance models for software design

Application: Software design and validation

Capability: Performance modelling

Know-how: Universal modelling language; queueing networks

As software systems become more complicated, reliable tools in support of software performance engineering become more valuable. New opportunities are presented by the emergence over the last few years of the Universal Modelling Language (UML) as a standard for the specification and design of object-oriented systems. UML is based on state diagrams that can be extended with timing information, to allow their performance analysis using the theory of stochastic processes. These techniques are making it possible to validate and compare the performance of alternative software architectures at the design stage, using measures such as response times, queue lengths, throughput and loss probabilities.

Reference:

Christoph Lindemann, Axel Thümmler, Alexander Klemm, Marco Lohmann and Oliver Waldhorst, *Performance analysis of time-enhanced UML diagrams based on stochastic processes*, 3rd ACM Workshop on Software and Performance, 25-34 (2002)

2 CURRENT LANDSCAPE

In order to understand the current landscape, one must look at three relationships: between mathematics and computer science, between industry and mathematics, and between industry and computer science.

The wider innovation landscape in the UK is addressed in the DTI's Innovation Report of 2003 and the Government's 10-year Science and Innovation Investment Framework, published in 2004¹. The international picture is one in which the innovation performance of the UK economy is close to the European Union average, but significantly lagging its major competitors such as France, Germany and the USA. The Government's stated target is to increase the combined public and private sector investment in R&D from 1.9% of national income in 2002 to 2.5% by 2014.

The interface between mathematics and computer science

The boundary between mathematics and computer science within the overall scientific landscape is drawn in different places in different parts of the world. These differences reflect the different ways in which national research strengths have evolved. The *de facto* boundary is determined by the respective activities of mathematics and computer science departments in universities.

The recent international reviews of UK research in mathematics (the IRM²) and computer science (the IRCS³), commissioned by the EPSRC, have identified the interface between mathematics and computer science as an area of general weakness in the UK. They have highlighted research in algorithms as a particular area on the interface that should be vigorously developed if the UK is not to fall behind its international competitors.

In the UK, mathematics has evolved to have much stronger links with physics, engineering and (more recently) the life sciences than it does with computer science. Conversely, the UK's long-standing strengths in computer science cover areas such as logic, semantics, formal methods and programming languages, and while much of this activity is mathematical in nature, it does not involve widespread interaction between academic departments. As a consequence, the interface between mathematics and computer science is weak by international standards, although there are notable exceptions in operations research, cryptography and information security, and emerging strength in areas such as bioinformatics, neuroinformatics and systems biology.

Operations research deserves particular mention, being an area where ideas from mathematics and computer science have been combined with strong industrial links over several decades to form a separate discipline, which is supported by its own learned society⁴ and was recently the

¹ *Science and Innovation Investment Framework 2004-2014*

http://www.hm-treasury.gov.uk/spending_review/spend_sr04/associated_documents/spending_sr04_science.cfm

² *International Review of UK Research in Mathematics*, <http://www.cms.ac.uk/irm/>

³ *International Review of UK Research in Computer Science*, <http://www.iee.org/Policy/CSreport/>

⁴ The OR Society, <http://www.orsoc.org.uk/>

subject of a separate international review by EPSRC¹. In many ways, it is a microcosm of the type of three-way interface that the MICS study aims to encourage more widely. On the other hand, perhaps because of its emergence as a separate discipline, operations research does not have the links to the science base at large that one might expect, including to the rest of mathematics and computer science. So even here there is a scientific interface to be developed.

The important question for this report is how the interface between mathematics and computer science can be strengthened in areas that will catalyse increased levels of innovation. These interactions can be partially characterized according to whether the capabilities being developed have their roots in mathematics or in computer science. For example, machine learning is an area where mathematics is adding value to computer science, and large-scale simulation is an area where computer science is adding value to mathematics. The UK is a leader in the development of computing tools to support the use of mathematics through computer algebra², statistical packages and numerical algorithms.

The importance is widely recognized of developing the mathematical and algorithmic aspects of computation alongside the information technologies themselves. As noted in a presentation by the President of SIAM to the US Computing Research Association³, new techniques are

needed to solve models that incorporate interacting subsystems, multiple scales and the effects of uncertainty. In large-scale simulation, therefore, mathematics contributes not only to the underlying models, but also to the algorithms and techniques by which those models are solved. As the need grows for simulations of complex distributed systems, the range of required techniques also grows. Over the next few years, agent-based simulation will become more important, and algorithms will draw increasingly on ideas from microeconomics.

The recent joint programme of the EPSRC and London Mathematical Society (LMS) in 'Mathematics for Information Technology' (MathFIT) provided an important stimulus to the interface of mathematics and computer science. Although not directed specifically at industrial application, some of the research within that programme has clear industrial potential, such as radio spectrum assignment, network reliability, the study of internet traffic statistics, and constraint programming. Building on the IRCS, IRM, and the Government's Investment Framework, there are now new initiatives to strengthen the interface. EPSRC has implemented 'discipline hopping' awards, and the interface of mathematics and computer science is among the areas earmarked for strategic development through the pilot implementation of Science and Innovation Awards, which are a joint initiative of EPSRC, HEFCE and SHEFC.

¹ *Review of Research Status of Operational Research in the UK*, <http://www.epsrc.ac.uk/CMSWeb/Downloads/Other/OpResReview.pdf>

² See, for example, GAP: *Groups, Algorithms, Programming*, <http://www.gap-system.org>

³ John Guckenheimer, *Numerical Computation in the Information Age*, http://www.cra.org/CRN/html/9803/affiliate/jg.3_1_t.shtml

The interface between industry and mathematics

It is recognized in the IRM that 'the UK has world-class activity in the organization of industrial mathematics.' The interface between industry and mathematics was strengthened significantly by the creation of the Faraday Partnership for Industrial Mathematics in 2000, under the management of the Smith Institute. The Faraday Partnership has enabled a much broader range of companies to access a much broader range of mathematical capabilities, in a facilitated environment. The consequence has been that the benefits of a mathematical way of thinking are leading to competitive advantage for companies more quickly and more consistently.

A mathematical approach often provides the lateral insights that are the catalysts for innovation. The formulation of industrial problems in mathematical terms involves a degree of abstraction that often makes clear how know-how can be transferred to great effect across business sectors. For example, in *Mathematics: Giving Industry the Edge*, it was described how knowledge of two-phase flow in the oil industry was used to design a new product for use in milking parlours.

The value that computer science adds to mathematics is derived in largest part from the numerical simulation of physical, biological or information systems. Simulation begins with a

mathematical formulation, or 'model' of a real-world phenomenon. The model is then solved using analytical and numerical techniques to produce quantitative predictions and understanding. The widespread availability at low cost of powerful computers has enabled most sectors of industry to use computational simulation as a matter of routine, in order to support and even replace experimentation. The results of simulations increasingly guide public policy as well as industrial design. Current subjects of large-scale computational simulation include global climate change, whole automobiles or aircraft, power networks, information networks, military conflicts, financial markets, and human organs. Many of these applications are highly multidisciplinary, requiring the additional integration of expertise beyond mathematics and computer science.

The information economy has provided particular opportunities for mathematics within the ICT industries, particularly in the design and operation of large-scale IT systems. Application areas such as information retrieval, performance modelling, resource allocation, coding and cryptography draw increasingly on ideas from discrete mathematics, statistics, combinatorial optimization and applied probability (and especially also number theory in the case of cryptography). Major successes include the deployment in the mid-1990s of trunk reservation and sticky random routing in BT's telephone network. Mathematical

modelling showed that these simple ideas increase network capacity almost to its theoretical limit. Underpinning mathematics is also well embedded in areas such as computer graphics.

The successful use of mathematics in industry relies much more on companies having personal contact and access to know-how than on the control of intellectual property. The UK's annual Study Group for mathematics in industry¹ exploits this fact by providing a unique forum in which industrialists benefit at first hand from the unparalleled versatility of industrial mathematics. In doing so, companies acquire entirely new perspectives

on some of their most difficult problems. Having originated in the UK over 35 years ago, there is now a world-wide network of Study Groups and they were singled out as a 'major success' in the IRM.

The whole notion of industrial mathematics is developing, and it is certainly not appropriate to see it, loosely speaking, as 'the most applied end of applied mathematics'. The challenges facing industry over the next 10 years, as identified in *Mathematics: Giving Industry the Edge*, call for a much wider view, drawing on any branch of mathematics that is appropriate to the challenge at hand. The concentration of industrial mathematics on

Structural safety

Application: Design of tall structures

Capability: System design

Know-how: Extreme event statistics; large-scale engineering simulations

The wind response of structures such as tall buildings, cranes and wind turbines is an inherently random process, owing mainly to the randomness in wind excitation, as well as the dynamic behaviour of the structure itself. Current approaches in wind engineering involve measurement of wind characteristics (speed, turbulence, etc.) for limited durations, and using these to generalize the wind action on various types of structures. While this approach works reasonably well, there is a need to improve the estimation of uncertainties involved in this process. This requires the use of advanced mathematical and statistical tools. Combining methods for extreme value statistics with wind engineering simulations will add great value to the estimation of uncertainty and have the potential to increase the reliability of structure design and operation.

Reference:

S.M.C. Diniz, M. Iancovici, M.A. Riley and E. Simiu, *Probabilistic performance criteria for tall buildings subjected to wind*, 9th ASCE Joint Specialty Conference of Probabilistic Mechanics and Structural Reliability (2004)

¹ See <http://www.smithinst.co.uk/StudyGroups>

simulation is giving way to a landscape in which simulation remains a prominent theme, but alongside multidisciplinary design, risk and uncertainty, data, complex systems, and market behaviour. The Smith Institute has correspondingly expanded its original core strengths in the modelling and analysis of physical processes, to incorporate increasing activity in areas such as operations research, discrete mathematics, data analysis and statistics.

Most companies see mathematics as offering an underpinning analytical tool, which can be used in a tactical way to address problems that arise in engineering. Many of them access these skills through hiring mathematically competent engineers. However, there is a distinct trend for companies to use mathematics more strategically, in making early-stage decisions about the value of design options and the potential threats of new technologies to existing products and services. This shift of emphasis is highlighting the key role of mathematical modelling, alongside analysis but distinct from it.

The interface between industry and computer science

Of all businesses, companies in the IT and telecommunications sectors generally seek the closest and broadest links with academic computer science, in ways that complement their own R&D programmes. Engineering companies adopt a

different approach, in which, like mathematics, computer science is a tool that contributes to the solution of engineering problems. In these cases, the majority of companies are looking to work with well-supported and well-documented codes.

The most successful modes of operation on the industrial interface differ between mathematics and computer science. Interaction between computer science and industry often requires a focus on detailed commercial practices and hence a need for privacy. Therefore personal consultancy and the establishment of spin-off companies from research programmes are more important routes to knowledge transfer than in mathematics; and the very open Study Group mechanism that works so well in mathematics is not easily adapted for use in computer science.

The UK Computing Research Committee (UKCRC) is currently engaged in an exercise to formulate a set of Grand Challenges in computer science¹. The objective is to advance the underlying science on broad fronts, and some of these areas, notably global ubiquitous computing, will lead to new technologies that can be exploited in partnership with industry over the long term. Looking also at the longer term, the recent report of Bullock and Cliff² as part of the UK Foresight Programme has mapped out opportunities for the exploitation of research advances in complexity and emergent behaviour.

¹ See http://www.nesc.ac.uk/esi/events/Grand_Challenges/background.html

² Seth Bullock and Dave Cliff, *Complexity and Emergent Behaviour in IT Systems*, <http://www.foresight.gov.uk>

In both computer science and mathematics, the main emphasis in universities is on developing the scientific leading edge, and consequently companies often have difficulties in locating the expertise that can provide value to their businesses on relatively short time scales. The MICS I interviews revealed an industrial perception in some quarters that computer science research does not focus on industrial impact, but overall we saw widespread evidence of very fruitful interactions between industrialists and computer scientists. Difficulties in reaching agreed positions on intellectual property are a barrier to collaboration, and are perhaps sometimes misinterpreted as signalling a disinterest in collaboration. In general there should be a proper balance between basic research, such as that being stimulated by the UKCRC's Grand Challenge exercise, and research that is industrially driven.

The major public investment in e-Science over the last few years has involved large segments of the computer science community. It is developing the infrastructure needed to support distributed global collaborations, which will involve coordinated access to large-scale data storage and retrieval facilities, high-performance computing, and visualization capabilities. The pilot projects have had significant industrial involvement and the expertise that they have generated is now being spread to a wider industrial audience through the DTI's Technology Programme¹.

¹ See <http://www.dti.gov.uk/technologyprogramme/>

3 CHALLENGES AND OPPORTUNITIES

The MICSI interviews with industrialists clearly identified a set of challenges that fall naturally into three themes: **network behaviour**, **algorithms** and **information management**. The following sections explore these three themes in turn, highlighting the main content of each. Many of the industrial applications that were covered during the interviews require expertise from two or three themes in combination, confirming the importance of a multidisciplinary approach.

3.1 Network behaviour

Technological trends towards the personalization of services, mobility and connectivity are leading to the emergence of very large-scale information networks that have dynamic constraints and are highly distributed and decentralized.

The overarching challenge in the design of networks is to develop models that connect emergent behaviour with network structure and the local rules of interaction between components. Such models can then be incorporated into design tools, capable of taking system requirements in the form of desired emergent behaviour and from them determining appropriate local structure and rules. A deeper mathematical framework is required to make better use of the results of current empirical work on the simulation of large networks.

The focus in this report is on computing and communications networks, but it should be recognized that progress in understanding and exploiting

network behaviour is being made on a wider front. Many valuable insights, generated by experiment, simulation and development of the underlying theory, are coming from social networks and biological networks in particular. The management of resources in networks and their use to deliver secure, scalable and reliable services are major challenges in many application areas.

Graph theory¹ is the natural branch of mathematics by which to study networks. From its roots in pure mathematics, graph theory is now becoming an essential tool in many applications, where it is being used alongside queueing theory, stochastic processes, dynamical systems, geometry and microeconomics. The most important issues cut across application areas.

- The **quality of service** experienced by users of networks is expressed in terms of performance measures such as throughput, response times and blocking fractions. Queueing theory provides a set of mathematical tools to predict these measures from knowledge of the network structure and the workload being placed upon it. The experience of an individual user depends on the concurrent actions of other users. **Performance modelling** guides strategies for network management, such as access control at the periphery of the network, flow control within the network, and priority techniques to offer different grades of service.

¹ A graph in this context does not refer to the display of data on a set of axes, but instead to the properties and structure arising from a set of interconnections (called links or edges) between a set of objects (called nodes or vertices). Graph theory is a branch of the wider area in mathematics known as combinatorics.

- Networks must be able to support **distributed and dynamic workloads**. Performance is often driven by the interactions between components in the network, rather than by the behaviours of the components themselves. These interactions lead to emergent behaviour that cannot easily be understood in terms of individual components. The key challenge in **network management** is to design control mechanisms that deliver high levels of performance by avoiding unwanted behaviour and preserving desirable behaviour.
- Large, dynamic heterogeneous networks do not perform well under centralized control. The alternative is a peer-to-peer architecture, with appropriate **distributed control mechanisms**, acting on the basis of local information. Distributed control is generally implemented through intelligent software agents, which interact with each other. The constraints and rules of interaction of the agents are governed centrally, but the decision-making within the network is local. Local decisions can perform very well (for example the sticky random routing mentioned in Section 2), but in general the trade-off between the speed and quality of local decisions is not sufficiently well understood.
- In networks faced with demands for service from many independent users, performance is maximized by **allocating resources** such as bandwidth, buffer storage and processing capacity in an intelligent way, and by balancing the overall load across the available resources. There have been significant advances over the last five years in understanding these issues and their operational implications. Where allocation decisions are decentralized, economic theory is being exploited in the design of internal 'markets' for resource allocation. Among the key ideas are fairness of allocation, stochastic methods, load-stealing and load-sharing.
- Grid infrastructures, or generically '**the Grid**', are networks in which there is access on demand to highly reconfigurable resources. It is a major challenge and also a major opportunity for industry to exploit these technologies successfully. The Information Age Partnership (IAP) have put forward a business-oriented action plan for the Grid¹, to stimulate the growth of new, high-value markets in which computing is a utility and software is a service. The IAP also recognizes that grid technologies are not yet at industrial strength, and 'remain largely unproven in terms of security, reliability, scalability and performance.' These four issues, along with trust and privacy, are important in all large-scale networks and mathematical modelling contributes centrally to their understanding.
- The ongoing e-Science programme² foreshadows some of the **new industrial applications** that can be supported by grid infrastructures

¹ The Information Age Partnership Grid Taskforce, *Unlocking the Grid*, <http://www.iapuk.org>

² See, for example, <http://www.allhands.org.uk/2004/>

when they become widely available, and the MICSI interviews gave the following additional pointers. Data from geographic information systems could be combined with simulations of people and vehicles to provide a tool for town planning. At a finer scale, the safety planning of large facilities such as shopping centres could be aided by the modelling and simulation of crowd behaviour in emergency situations, taking account of emotional and physical responses at the level of the individual. There are also suggestions that policy evaluation in both public and private sectors can be supported by large agent-based simulations.

- The established view of IT networks is one in which the assignment of processes to processors is fixed. The future is more open, with both **mobile processes and data** within the network; and the rapid growth of wireless technologies means that the network infrastructure itself is also dynamic. These developments raise serious issues of process security, in addition to data security, which must be integrated into decentralized network management systems. If the Grid vision of software services is to come to fruition, then customers must be provided with access to protected computing environments.

Bandwidth sharing

Application: Customer-related performance; network dimensioning
Capability: Performance analysis
Know-how: Stochastic processes; theory of networks

Modern communication networks are able to respond to randomly fluctuating demands and failures by adapting rates, by rerouting traffic and by reallocating resources. They are able to do this so well that, in many respects, large-scale networks appear as coherent, self-regulating systems. The design and control of such networks present challenges of a mathematical, engineering and economic nature. Early work focussed on studies of packet-level statistics but recent work has begun to consider flow-level representations. These approaches model the randomly varying number of flows in a network where bandwidth is shared fairly between concurrent flows. They are useful in assessing end-to-end (customer-related) performance measures, as well as providing dimensioning tools for the network infrastructures maintained by Internet Service Providers (ISPs).

Reference:

Frank Kelly and Ruth Williams, *Fluid model for a network operating under a fair bandwidth-sharing policy*, *Annals of Applied Probability*, **14**, 1055-1083 (2004)

- **Ad hoc networks** are wireless networks formed by the temporary cooperation between mobile nodes as they move in and out of communication range with each other. Their topology is highly dynamic, driven by the motion of the nodes. The motion of individual nodes may be arbitrary, or partially correlated in cases such as convoys of unmanned autonomous vehicles (UAVs), where the nodes are participants in a common objective. Ad hoc networks present new research challenges. In particular, the physics of radio wave propagation imposes spatial constraints on the network structure, stimulating research into so-called random geometric graphs. Economic theory is used in

network management, to give nodes the correct incentives to cooperate as relay points for the transmissions of others.

- **Scalable networks** maintain their levels of performance as they are scaled up to larger numbers of nodes and higher levels of demand. Companies need confidence that today's solutions can be upgraded in a smooth and cost-effective way to meet tomorrow's demands. The challenge is to provide scalable design principles and network management strategies, which avoid instability or failure appearing as emergent properties of the network. In multi-player networked games, scalability is an

Load management in distributed systems

Application: Distributed systems with heterogeneous dynamic loads
Capability: Performance optimization
Know-how: Stochastic processes; Markov decision theory

The emergence of Grid environments for inter-enterprise computing creates a new paradigm, in which the location and nature of computing resources are hidden from the user. Processing resources must be made available seamlessly and simultaneously to a range of heterogeneous demands. By building stochastic models of Grid systems, it is possible to calculate optimal policies for the management of resources, and also to design heuristics that are close to optimal and straightforward to implement in practice. This work is based on Markov decision techniques, which are at the heart of many important applications in job scheduling, network routing, load balancing and resource allocation.

Reference:
 Jennie Palmer and Isi Mitrani, *Optimal server allocation in reconfigurable clusters with multiple job types*, University of Newcastle technical report CS:TR823 (2003)

intrinsic part of the customer experience. Companies are also looking to design their internal structures from a network perspective, through scalable networks of manufacturing cells. Predicting the way in which network behaviour scales with network size is a question that is ideally suited to a mathematical approach. Much of the theory of random graphs becomes more directly applicable as the graphs (networks) become large. Nevertheless, achieving desired global behaviour in a scalable environment remains a general challenge.

- Inexpensive sensor devices and **radio frequency identification** (RFID) tags are creating new types of information networks, in support of operations in manufacturing, supply chains and logistics. Sensor networks built around low-cost components present new challenges in scalability, especially in the ability to handle high data volumes. However, the benefits are already being demonstrated. The UK brewing industry loses about 4% of its inventory of beer kegs in the supply chain annually, representing a cost of about £10 million. Scottish and Newcastle have responded by introducing RFID tags on 1.9 million kegs. Wal-Mart, the world's largest retailer, has required its top 100 suppliers to incorporate RFID technology into all pallets that are shipped to its distribution centres from January 2005.
- **Network topology** (the layout of connections within a network) is fundamental to determining emergent properties and network performance. The theory of random graphs is being extended to study different broad types of network topology, and there is also an important connection to statistical physics. A great deal of current interest has been directed at so-called 'scale-free' networks, such as the Internet, which have a relatively large number of short-range connections and a relatively small number of long-range connections. In these networks, any node can be reached from any other by following a small number of links (the 'small-world' phenomenon).
- The behaviour of networks can evolve because of the addition, removal or mobility of nodes and links, or because of changes in the workload being applied. How should dynamic networks be controlled and how can their stability of behaviour be ensured? The more general study of dynamical systems offers useful techniques that are applied to fluid analogues in the case of heavily loaded networks. An everyday example of well-controlled **network dynamics**, underpinned by a careful mathematical analysis, is the stability afforded by the paradigm of 'additive increase, multiplicative decrease' for transmission rates in TCP connections.

- Transport networks provide an example in which the operation of a physical network can be made more reliable through the use of new information networks embedded in it. Rail networks, for example, require ongoing **maintenance** and the timely deployment of maintenance resources can be supported by sensor networks attached to the rail infrastructure.
- In networked environments, quality of service is greatly influenced by the network's **resilience** to random failure or deliberate attack and its subsequent **recovery**. Recent failures of the electricity grids that supply London, the eastern US and Italy have highlighted the dramatic consequences in networks that are not resilient. Scale-free networks such as the Internet have topologies that make them resilient to random failure but vulnerable to a coordinated attack on key hubs. We are all familiar with problems posed by malicious mobile code such as worms and viruses. Networks that are deployed in military environments, for example in support of UAVs, must be especially resilient to deliberate attack. It is a key challenge to ensure that safety-critical systems are sufficiently robust in their emergent behaviour. The ability of networks to withstand the removal of nodes or links is another natural question for graph theory, and strategies for defending against such attacks can be analysed in terms of game theory.

3.2 Algorithmics

Algorithms are the basic mathematical building blocks from which computing systems are constructed. Algorithmics is the discipline that covers the design, analysis and implementation of new algorithms and their associated data structures. It is essential for the innovation of computing infrastructure, tools and applications.

The discipline of algorithmics connects application domains to the design and implementation of suitable algorithms, usually through the intermediate step of a mathematical model¹. It involves a mix of underpinning theory and computational practice, and it is important that these two strands are developed in unison, as is commonly the case in operations research.

The increasing power of computational simulation and optimization over the last 20 years has been as much due to algorithmic developments as to advances in microprocessor speed and data storage capacity. Moreover, many applications are constrained by factors such as cost, battery life, or communications bandwidth, and as a result rely on algorithms tailored to their particular environments.

- The performance of any algorithm is determined to some extent by **resource constraints**: processing power, memory capacity, disk storage, bandwidth and latency. Hand-held and mobile devices have the additional limitation of

¹Strictly speaking, an algorithm is a prescription for finding an exact solution to the problem in question. For computationally intractable problems, there is also great interest in heuristics, which are more general procedures aimed at finding rough but satisfactory solutions. Here we include the study of heuristics under the umbrella of algorithmics.

battery life. In real-time applications, such as search engines or computer games, resource constraints become a key factor in systems engineering. New computing paradigms will introduce new resource constraints, particularly on bandwidth and latency in the case of grid technologies. Efficient compression techniques will be required to transfer data efficiently between processing units.

- **Algorithmic performance** is measured by the computational resources that are consumed, which are partly temporal (execution time) and partly spatial (storage requirements). These measures depend on the detailed implementation of the algorithm, and so comparisons usually look at increases in the required resources with the size of the problem. There are many tasks, for example in scheduling and facility location, where the required resources increase exponentially with the problem size, making them effectively intractable. In these cases, one can turn to approximation algorithms. These generate approximate results, whose closeness to the true result is then an additional measure of performance.
- It is perhaps surprising that performance can be improved by allowing algorithms to include random choices as part of their execution. These **randomized algorithms** often deal better with 'worst-case' problem

instances and can streamline execution through sampling the input data. A flexible paradigm occurs in load balancing, where traffic is assigned to links in a network, or computational tasks to parallel processors. Choosing the more lightly loaded of two randomly chosen links or processors gives an exponential reduction in the maximum load when compared with making a single choice.

- **Heuristics** are procedures, generally with randomized aspects, that are heavily motivated by accumulated computational experience on broad classes of problems. They are especially useful in attacking problems beyond those for which an algorithmic theory has been fully developed. Although heuristics do not have performance guarantees, they often lead to dramatic improvements over current industrial practice, and sometimes also far in excess of what currently available theory might suggest. There is much to be gained from developing general-purpose heuristics that are robust, flexible and easy to implement.
- The computer games industry has particular algorithmic challenges, arising from the need to deliver **interactive simulation environments** at very low cost. Current algorithms for physical simulation in games engines do not adapt well to real-time demands. Resource constraints call for a complete reassessment of the modelling process. For example, it is out

of the question to provide a solver for the Navier-Stokes equations in games engines, so what should fluid modelling consist of in these circumstances? Moreover, the user experience calls for high-end graphics on low-end machines and increasingly on multiple machines. The industry also requires adaptive algorithms that can compress and decompress images for delivery on multiple devices.

- Data gathered from sensing and monitoring devices should be processed and acted on as necessary in real time. More use is needed of **hybrid models** for such systems, in which the gathered data is assimilated into the parameters of simulation models. In this way, one achieves data reduction on to the domain of interest, and uses a mixture of data-driven and physics-based simulation to make predictions. Weather forecasting is just one area where such approaches are being actively pursued. Different hybrid models might reduce the same data streams to knowledge in different domains. Parameterization within hybrid models can also be used where computing resources are insufficient to model all the relevant processes in full physical detail. There is then the question of what is a sufficient or optimal hybrid model.
- There is a general trend for companies to deploy systems using **commercial off-the-shelf (COTS) hardware**, but this is not always compatible with the requirements of software

dependability. Safety-critical systems often use bespoke architectures, which have high maintenance costs. Use of commercial hardware in tandem with bespoke software would reduce costs significantly, but such systems would need to maintain the current low levels of undetected failures. Companies would benefit greatly from dependable algorithms that are compatible with COTS hardware.

- Companies wish to have efficient software implementations without **dependence on architecture features** such as the complicated memory hierarchies in parallel and distributed environments. Of particular importance, because of their ubiquitous use in computational modelling, are efficient algorithms for matrix computations. Their performance is greatly affected by the choices of matrix partitioning and block sizes in relation to the available memory caches. If compilers made such decisions internally then modellers could concentrate on the fidelity of their simulations, knowing that the available hardware was automatically being put to best use¹. Grid services have a deeper issue to face, namely that the architecture on which algorithms are executed may not be decided until run-time.
- In **real-time systems**, the load to be processed often consists of a sequence of 'jobs', which must be carried out as they arise, without knowledge of the jobs that will come in the

¹Erik Elmroth, Fred Gustavson, Isak Jonsson and Bo Kagstrom, *Recursive blocked algorithms and hybrid data structures for dense matrix library software*, SIAM Review, **46**, 3-45 (2004)

Network diameters

Applications: Peer-to-peer networks; energy distribution; consumer marketing

Capability: Network design

Know-how:

Random graphs; scale-free networks

The Internet, telecommunications networks, energy distribution networks and networks of consumers are all key drivers of economic activity. Understanding the impact that network structure has on the efficient distribution of information or other resources is essentially a mathematical problem. For example, the diameter of the network, defined as the largest separation of any pair of nodes, is known to be related to the degree distribution, which measures the relative number of neighbours (or degree) of each node. A small diameter means that information or resources can be propagated network-wide at low cost. It has recently been shown that many scale-free networks have a diameter that grows like $(\log n)/(\log \log n)$, where n is the number of nodes, in contrast to $\log n$ for 'classical' random graphs. Scale-free networks seem to occur frequently in practice, and their smaller diameters open up opportunities for faster peer-to-peer communications, shorter energy distribution routes and more effective marketing strategies.

Reference:

Béla Bollobás and Oliver Riordan, The diameter of a scale-free random graph, *Combinatorica*, **24**, 5-34 (2004)

future. For example, operating systems generate jobs in the form of page requests. There is a fast-access cache, which can hold a limited number of pages, while the remainder must be retrieved from slow memory. The aim is to use the cache to minimize the number of times that the slow memory must be accessed (these events are called page faults¹). **On-line algorithms** are at the heart

of all such challenges. Their quality is usually measured by the competitive ratio, which is the ratio of the relevant performance measure to that of an 'adversary' who has the advantage of being able to see the future sequence of jobs.

- Hardware verification using formal methods² is becoming the norm for the market leaders in microprocessor design (companies such as

¹ Wolfgang Bein and Lawrence Larmore, *Trackless and limited-bookmark algorithms for paging*, ACM SIGACT News, **35**, 38-48 (2004)

² Robert Jones, John O'Leary, Carl-Johan Seger, Mark Aagaard and Thomas Melham, *Practical formal verification in microprocessor design*, IEEE Design and Test of Computers, **18**, 16-25 (2001)

Intel, ARM, IBM and AMD). In contrast, **software verification** is still an emerging area for industry and will become an increasingly important requirement. The field of semantics is concerned with verifying the correspondence between what a program computes and the program text itself. It is well developed for computational logic and for integer-valued calculations, but much less so for floating-point algorithms, which is where the bulk of engineering applications lie. There is a trend towards closer dialogue and iteration between design and verification in the systems engineering lifecycle. Better systems design methods will allow the consistency checking of large sets of requirements and earlier checking of designs against requirements. Given the commercial pressures to keep the time to market of new products as low as possible, there would be great value in tools to predict the time needed to verify a system before it is built. Emerging directions include the systematic analysis of symmetries in order to streamline the verification process.

- Many systems are described by **hierarchical models**, in which different aspects are understood at different scales, either in space or time. The relevant scales in a problem can cover many orders of magnitude, and their simulation requires further development of **multiscale algorithms**. An important multidisciplinary

application is the integration of chemistry with multiphase CFD models.

- **Quantum computing** will probably not have an imminent widespread impact on industry, but its potential to do so in the future, coupled to its strong mathematical foundations, is sufficient to warrant inclusion in this report. The size of a quantum computer is measured in 'qubits'. Current experiments are at the level of two or more qubits, and are exploring about 20 different technologies. Approximately 30 qubits are required to match a conventional computer, while a quantum computer of 50-100 qubits will be able to simulate quantum systems that are in principle beyond the reach of conventional computers. Such levels might be reached 10 years from now. Potential applications are molecular modelling in drug design, and generic applications such as database search. The construction of a quantum computer with tens of thousands of qubits would fatally undermine today's 'secure' communications that are based on public key cryptography. Continued investigation into quantum algorithms and protocols will help to establish the market promise needed to drive the development of large-scale quantum processors.

3.3 Information management

All organizations generate data, in quantities that are unprecedented, thanks to breakthroughs in sensor technology and storage capabilities. Datasets running into terabytes are commonplace and petabyte instances are emerging. To avoid 'data overload', a crucial challenge is the processing of data into information that can be used for decision support in operations, design and strategy formulation.

Information management includes data retrieval, communication, storage and visualization. In recent years, advances in computing technology have brought about a step-change in the capability of organizations to collect, transfer, store and manipulate data. The MICS interviews pointed to a number of areas where a more sophisticated mathematical approach can increase the exploitation of these advances in improved operational efficiency and better management decisions. Multimedia applications in particular are highly dependent on signal processing.

It is useful to think of turning data into information through the incorporation of context and then using that information to make decisions. The information held in data is context-dependent and it is not possible to separate useful from useless data in isolation. The exchange between data and information is often also iterative, for example when the information gained from one set of experiments guides the design of

the next set, or the output from a search engine leads to a more refined query. It is a major challenge to design systems that are **context-aware** and **task-dependent**.

Given that data is the raw material of information management, there must also be reliable methods for its storage, coding, communication, protection and authentication. The UK benefits from healthy industrial connections with the science base in the areas of cryptography and information security, and this strength needs to be sustained so that future requirements are well served.

- Increasing data volumes necessitate better **automated pattern recognition** and interpretation, in support of human interpretation. Pattern recognition is a means of mapping data on to information, given the context of a particular application. Different applications require different approaches, especially when data is not numerical. Some pattern recognition problems are computationally intractable, for example the recognition of subgraphs in large networks, and so in these cases there is a need for approximation algorithms or heuristics¹. Standard tools that handle single relations in single databases need to be replaced with a capability to access several databases simultaneously, looking across them for multifactor links and correlations. The pharmaceutical industry is a key emerging market for such technologies and

¹ Jun Huan, Wei Wang and Jan Prins, *Efficient mining of frequent subgraphs in the presence of isomorphism*, 3rd IEEE International Conference on Data Mining (2003)

there are wider applications in medical imaging, remote sensing, credit rating and fraud detection. Visualization techniques also play a vital role, in enabling pattern recognition by humans.

guide **simulation strategies** and allow simpler causal models to be developed. Major applications here include climate change and network behaviour.

- It is a challenge to evaluate and compare results across a set of simulation scenarios, each of which may produce gigabytes of data. Techniques to highlight the major differences and similarities among output datasets will
- **Data quality** must be managed so that data reflects the current state of the system to which it relates. Tools that are used for planning or real-time decision-making cannot function properly if there is unmanaged uncertainty about datasets being correct and up-to-date.

Decision support for cutting and packing

- **Application:** Waste reduction in steel, textiles, polycarbonates, glass and fibre glass
- **Capability:** Stock cutting
- **Know-How:** Heuristics

The inefficient cutting of raw material or finished product is a significant source of waste in many manufacturing industries. The optimum packing of shapes on to stock sheet of material is an intractable problem to solve exactly and has been the subject of approximate and heuristic approaches for over 40 years. A new heuristic has recently been developed, called the 'best fit' heuristic, which outperforms all previous methods for the packing of rectangles. Similar advances have been made for the problem of packing irregular shapes, including arcs and holes, which also outperform the previous state of the art.

This new technology has been exploited with industrial collaborators for use in the steel industry, using the Teaching Company Scheme, and is being prepared for other markets through a new start-up company.

Note:
The Teaching Company Scheme has been replaced with Knowledge Transfer Partnerships; see <http://www.ktponline.co.uk>

Reference:
Edmund Burke, Graham Kendall and Glenn Whitwell, A new placement heuristic for the orthogonal stock cutting problem, Operations Research, 52, 655-671 (2004)

The problem is greater when data comes from several sources of diverse character and reliability. Very similar challenges arise in the supply of reliable data, especially if the same datasets are used in multiple applications. These include consistency checking, the replication of changes, and responsibility for data quality. Consistency checking involves checks between datasets and checks over time within the same dataset. In future, intelligent sensor networks will provide these capabilities in applications such as weather forecasting.

- When data can be set rapidly in context, it is a powerful aid to **early-stage decision-making**. Intelligent information retrieval can make use of the context in which a query is being made. Data streams from the monitoring of high-integrity systems can interact with models of the system in order to provide early detection of defects, within the range of normal scatter. Even the model design process itself can be accelerated through the targeted use of existing data and experimental programmes in order to assess and refine initial modelling assumptions.
- The need to process very large volumes of data and to communicate the information derived from them is driving new developments in **data compression**. For example, companies wish to compress the output from simulations while

retaining the information required for surface rendering or the visualization of complicated fluid flows. Wireless transmissions of data, using either satellites or point-to-point terrestrial links, have a particularly acute need for compression, owing to the scarcity and cost of radio spectrum. The **dimensional reduction** of large datasets is a recurrent theme in information management. Some applications, along with their associated algorithms are now standard, for example the use of digital formats in the music and video industries. More generally, the development of new methods is an area of great opportunity for statistics and other branches of mathematics. For example, a multidisciplinary team at Stanford¹ has produced a preliminary software tool that uses topological methods for the dimensional reduction and parameterization of high-dimensional datasets.

- **Real-time decisions** use data streams as input to on-line algorithms. Many applications are fully automated, for example process monitoring, resource allocation in networks, and the detection of denial-of-service attacks, and they will become more widespread with the increasing deployment of sensor networks. Among the challenges here are to ensure that the data being streamed is relevant and that the on-line algorithms have execution times that are compatible with the data rate.

¹ See <http://math.stanford.edu/comptop/>

Recognizing subgraphs

Application: Bioinformatics; drug design; counter-terrorism
Capability: Information management; algorithmics
Know-how: Subgraph isomorphism; network structures

Information is often derived from finding patterns of relationships in large datasets. By modelling items in a dataset as nodes in a graph, and the relationships between them as the corresponding edges, the question of finding patterns is naturally turned into a subgraph isomorphism problem. Such problems are known to be computationally challenging, and there is much current interest in designing and implementing improved algorithms, motivated by high-profile applications in the life sciences and in intelligence analysis for counter-terrorism (where the relevant graphs are often social networks). New techniques must be able to fuse data sets that are incomplete, uncertain or compiled from diverse sources.

Reference:

Thayne Coffman, Seth Greenblatt and Sherry Marcus, *Graph-based technologies for intelligence analysis*, Communications of the ACM, **47**, 45-47 (2004)

- The general searching of datasets for special features and causal connections is called **mining**, in contrast to the focussed issues of information retrieval and pattern recognition. A key objective is the identification of clusters of similar objects. Established approaches for numerical data are now being supplemented with new ideas for cluster analysis in qualitative data¹ and clustering algorithms for data streams², in which the algorithm can make only one pass through the data and has access to limited memory. The latter techniques are important in situations where cluster analysis feeds into real-time decision-making, or where

resource constraints prevent large datasets from being stored and analysed off-line. A different line of current development for general search problems is inductive logic programming (ILP³), which has its roots in machine learning and computational logic.

- **Information retrieval** matches search queries to documents and ranks the results. The most popular methods for searching the web, as exemplified by Google's PageRank algorithm, judge the relevance of a document by analysis of its hyperlinks. However, these methods are much less effective when applied to the 'deep

¹ Moses Charikar, Venkatesan Guruswami and Anthony Wirth, *Clustering with qualitative information*, 44th IEEE Symposium on Foundations of Computer Science, 524-533 (2003)

² Moses Charikar, Liadan O'Callaghan and Rina Panigrahy, *Better streaming algorithms for clustering problems*, 35th ACM Symposium on Theory of Computing, 30-39 (2003)

³ David Page and Ashwin Srinivasan, *ILP: A short look back and a longer look forward*, Journal of Machine Learning, **4**, 415-430 (2003)

web', meaning company intranets and other resources with restricted access¹. The diversity of content on intranets and the need for secure access require new approaches, which are likely to be hybrids, for example using language modelling in conjunction with an analysis of document links.

- Approximately ten times as much data is stored on the deep web as on the public web. Companies recognize that there can be great benefit in **sharing data** with others, especially in the planning of new systems, but in practice they are often constrained by concerns over

compromising their commercial positions. The development of better security and encryption techniques can aid secure sharing. These issues will become increasingly relevant as Grid technologies mature. Government and market pressure force some cooperation, for example in the rail industry, and where there is political will the European Union is putting in place common specifications, but this is a long process and the consequences for system procurement are not yet apparent.

¹ Rajat Mukherjee and Jianchang Mao, *Enterprise search: tough stuff*, ACM Queue, 2, 37-46 (2004)

Language modelling for information retrieval

Application: Data mining; search engines

Capability: Information retrieval

Know-how: Statistics; mathematical logic; theory of networks

Current search engines are based on research from the 1970s and 1980s, since when there has been an exponential growth in on-line information, both locally within businesses and globally on the web. Language modelling can be used to improve the performance of search engines, in order to meet the scale of present and future demands. The general idea is to model a document as a stochastic process that will generate a query with a certain probability. Documents are ranked with respect to a query, in order of their probabilities. One of the main challenges in application is to integrate language modelling with the network-based approach as exemplified by Google. The theoretical basis for capturing the scale-free nature of the web is statistical mechanics. To build a new generation of search engines, combining language modelling and network theory, will require extensive collaboration between computer scientists and mathematicians.

Reference:

Bruce Croft and John Lafferty (eds.), *Language modeling for information retrieval*, Kluwer (2003)

4 ESTABLISHING EFFECTIVE COLLABORATIONS

Having identified themes of agreed scientific importance, it remains to address the question of how to organize the required three-way collaborations between computer scientists, mathematicians and industrialists. Business and academia have very different metrics for evaluating performance. In business, profitability is the main criterion, on timescales that vary between companies and across sectors, from a few months to several years. In academia, reputation and influence are the main motivations. Effective collaborations balance the achievement of specific business objectives with strengthening the science base. They need a carefully planned package of coordinated mechanisms, especially when there is an intrinsically multidisciplinary element as in the MICS themes.

Mechanisms

The experience of the Smith Institute is that collaborations proceed best through face-to-face engagement, even if the ultimate aim is the creation of new computing tools. In the early stages, the emphasis should be firmly on understanding each partner's expectations and expertise. The following are valuable mechanisms to use at different stages as relationships develop.

- **Briefings** Many companies find it difficult, or at least burdensome, to locate new sources of relevant expertise in what is

a complicated academic landscape, and academics face similar difficulties in making new industrial contacts. Briefings provide an early-stage mechanism for extending existing circles of contact. They provide companies with a means of learning at first hand about current capabilities and emerging research, and provide academics with an overview of industrial priorities. Central coordination through some sort of knowledge transfer network enables multiple organizations to benefit simultaneously. Companies contributing to the MICS study commented on their desire to work with familiar partners, and briefings can build up the required level of confidence to establish strong new collaborations.

- **Study Groups**¹ Study Groups work well when an initial contact has been established. The agenda is set by industrial presentations of target problems, which are followed by several days of intensive brainstorming in a dynamic and multidisciplinary environment. The success of Study Groups in mathematics derives from the widespread applicability of modelling techniques and the ability to evaluate, compare, refine or discard ideas very quickly. The output from each Study Group problem is a recommended solution, suggestions for further development, or a clarification of potential barriers to progress. All three are valuable to

¹ See <http://www.smithinst.co.uk/StudyGroups>

On-line bin packing

Applications: Multimedia on demand; data curation; job scheduling
Capability: Resource allocation
Know-how: On-line algorithms

Allocating resources to minimize waste is a challenge that arises in settings as diverse as the scheduling of tasks among a workforce (where the resource is time), the streaming of multimedia data (where the resource is bandwidth) and the archiving of data sets (where the resource is magnetic or optical storage). Resources are often available in chunks of fixed capacity, for example the length of a work shift, the total bandwidth of a communications channel, or the capacity of a magnetic tape. In these situations, the problem is naturally one of 'bin packing', where a 'bin' is the natural unit of resource. Furthermore, in real-time applications there is the on-line challenge of assigning each task, data stream or data set to a bin as it arises. There is no information available about the size and number of items to be assigned subsequently, and once assigned, items may not be moved to another bin. Recent research has produced important new algorithmic capabilities for this type of problem.

Reference:

Edward Coffman Jr, Costas Courcoubetis, Michael Garey, David Johnson, Peter Shor, Richard Weber and Mihalis Yannakakis, *Perfect packing theorems and the average-case behavior of optimal and online bin packing*, SIAM Review, **44**, 95-108 (2002)

the company. Although their open format is not suited to many areas of computer science, there is greater scope for using Study Groups on the interface with mathematics.

- **Secondments** The MICSII study revealed a widespread support for the notion of secondments, but at the same time a recognition that career structures and day-to-day

demands make them difficult to put into practice. Secondments allow a deeper level of relationship to be established than either briefings or Study Groups and also provide time for solutions to be implemented. It is important that the secondment is designed around a specific strategic objective. Knowledge Transfer Partnerships (KTPs)¹ are a closely related alternative, without the

¹ See <http://www.ktponline.org.uk>

difficulties of implementation. They provide support for a recent graduate to be recruited as an Associate, working in a company on a project of strategic importance for between one and three years, with a half-day of supervision each week from a knowledge-base partner.

To see how connections within the research base and between industry and the research base can build into a critical mass, one can look to leading centres overseas. In the US, there is the DIMACS (Discrete Mathematics and Theoretical Computer Science) centre, which since 1989 has been a Science and Technology Center of the National Science Foundation, bringing together Rutgers and Princeton universities with a group of leading corporate research laboratories. In Europe, a prominent example is the Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), which was founded in 1984 as a non-university research institute concentrating on application-oriented algorithmic mathematics. There are no similar centres in the UK, making it difficult to compete internationally.

Leveraging existing activity

Over the last few years, the Smith Institute has used its position as an intermediate organization to build the Faraday Partnership for Industrial Mathematics into a highly connected network of approximately 1,000 academics, industrialists and policy makers, which enables innovation across a broad range of business sectors. The Smith Institute's team of Technology Translators facilitate and support all aspects of industrial-academic collaboration within the network, using mechanisms chosen to suit individual circumstances.

The three MICS themes represent opportunities where this experience can be used to stimulate new areas of business activity. Such a programme will require a collaborative effort that brings together computer scientists and mathematicians in conjunction with industrialists and policy makers. To establish the necessary connectivity in the science base will need the support of organizations such as the UKCRC and learned societies including the IEE, IMA and LMS. As first steps, we make the following recommendations:

- There should be closer interaction between the relevant organizations that help to steer strategic direction in the science base. In particular, the Smith Institute and the UKCRC are eager to develop closer ties,

which will encourage computer scientists and mathematicians to participate jointly in industrial workshops and briefings in the priority themes.

- In conjunction with its industrial and academic partners, the Smith Institute will identify a set of industrial topics to be the initial focus of these activities. These could include bioinformatics in the pharmaceutical sector, brand management in the retail sector and traffic management in the transport sector, along with cross-sectoral topics such as information retrieval and remote sensing.

These steps will lay the foundation for long-term collaborative activity, which derives competitive advantage for the participating companies alongside the strengthened coupling of computer science and mathematics within the science base.

THE MICSI SOURCES

The MICSI study drew on four types of information source:

- Existing landscape or strategy documents in relevant fields;
- Interviews with industrialists and academics;
- Expertise accumulated by the Faraday Partnership for Industrial Mathematics;
- The recent research literature.

The primary source for information on industrial priorities and opportunities were interviews with approximately 30 leading industrialists. The companies involved were chosen to represent major providers of information and communications technologies (for example BT and Hewlett-Packard) and major end-users (for example Airbus, Rolls-Royce, GlaxoSmithKline and BP). A further set of interviews was held with relevant academics, to provide a wider perspective of the science base than is possible in a purely desk-based survey and to identify specific examples of where a closer interaction between mathematics and computer science will have major impact.

USEFUL WEBSITES

The Smith Institute for Industrial Mathematics and System Engineering

<http://www.smithinst.co.uk>

UK Computing Research Committee

<http://www.ukcrc.org.uk>

The Institute of Mathematics and its Applications

<http://www.ima.org.uk>

The London Mathematical Society

<http://www.lms.ac.uk>

The Institution of Electrical Engineers

<http://www.iee.org>

The OR Society

<http://www.orsoc.org.uk>

The Royal Statistical Society

<http://www.rss.org.uk>

Department of Trade and Industry

<http://www.dti.gov.uk>

Engineering and Physical Sciences Research Council

<http://www.epsrc.ac.uk>

UK Foresight Programme

<http://www.foresight.gov.uk>

MathFIT - Mathematics for IT

<http://www.mathfit.ac.uk>

National e-Science Centre

<http://www.nesc.ac.uk/>

DIMACS: Center for Discrete Mathematics and Theoretical Computer Science

<http://dimacs.rutgers.edu>

ZIB: Konrad-Zuse-Zentrum für Informationstechnik Berlin

<http://www.zib.de>

ABBREVIATIONS

CFD	Computational Fluid Dynamics
COTS	Commercial-off-the-shelf
DIMACS	Center for Discrete Mathematics and Theoretical Computer Science
DTI	Department of Trade and Industry
EPSRC	Engineering and Physical Sciences Research Council
HEFCE	Higher Education Funding Council for England
IAP	Information Age Partnership
ICT	Information and Communication Technologies
IEE	Institution of Electrical Engineers
ILP	Inductive Logic Programming
IMA	Institute of Mathematics and its Applications
IRCS	International Review of UK Research in Computer Science
IRM	International Review of UK Research in Mathematics
LMS	London Mathematical Society
MathFIT	Mathematics for Information Technology
MICSI	Mathematics at the Interface of Computer Science and Industry
QIP	Quantum Information Processing
RFID	Radio Frequency Identification
SHEFC	Scottish Higher Education Funding Council
SIAM	Society for Industrial and Applied Mathematics
TCP	Transmission Control Protocol
UAV	Unmanned Autonomous Vehicle
UKCRC	UK Computing Research Committee
ZIB	Konrad-Zuse-Zentrum für Informationstechnik Berlin

ACKNOWLEDGEMENTS

The MICS study was supported by a grant from the Department of Trade and Industry. The Smith Institute acknowledges with warm thanks the contributions of all those who made time available to contribute information, opinion and ideas:

Prof Samson Abramsky (University of Oxford)
Mr Andrew Allsop (Arup)
Dr Alexandra Bach (AEA Technology)
Mr Steve Baker (BNFL)
Dr Steven Barrett (GlaxoSmithKline)
Mr Nigel Barry (Airbus)
Dr Keith Briggs (BT Exact)
Prof Edmund Burke (University of Nottingham)
Dr Dave Cliff (HP Labs)
Mr Andrew Cliffe (Serco Assurance)
Mr Peter Coleman (Airbus)
Dr Peter Cowley (Rolls-Royce)
Dr Mike Cullen (Met Office)
Dr Gary Davies (GlaxoSmithKline)
Dr Jeremy Everett (Pfizer)
Dr Aiden Flynn (GlaxoSmithKline)
Dr Richard Gibbens (University of Cambridge)
Prof Mike Giles (University of Oxford)
Prof Sven Hammarling (NAG)
Dr Gavin Harper (GlaxoSmithKline)
Dr Nick Hazel (BP Chemicals)
Prof Tony Hey (UK e-Science Core Programme)
Prof Sir Antony Hoare (Microsoft Research)
Dr Richard Jarvis (BNFL)
Prof John King (University of Nottingham)
Prof Andrew Lacey (Heriot-Watt University)
Dr Jonathan Lawry (University of Bristol)
Mr Gordon McConnell (Airbus)
Prof Isi Mitrani (University of Newcastle)

Mr Ken Mylne (Met Office)
Dr Ender Ozkan (Arup)
Prof Kevin Parrott (University of Greenwich)
Dr Kenny Paterson (Royal Holloway, University of London)
Mr Daren Payne (Airbus)
Mr John Ponting (Met Office)
Dr Nick Randon (University of Bristol)
Ms Padma Reddy (CERCIA)
Dr Thorsten Schnier (CERCIA)
Mr Sanjiv Sharma (Airbus)
Prof Frank Smith (UCL)
Dr Tim Spiller (HP Labs)
Dr Alan Stevens (Rolls-Royce)
Prof Iain Stewart (University of Durham)
Prof Peter Stow (Rolls-Royce)
Mr Martyn Thomas (Railway Safety and Standards Board)
Prof Keith van Rijsbergen (University of Glasgow)
Ms Faith Wainwright (Arup)
Mr Vibhu Walia (CERCIA)
Dr Lincoln Wallen (Criterion Software)
Dr Eddie Wilson (University of Bristol)
Dr Keith Winters (AEA Technology)
Prof Jim Woodcock (University of York)
Prof Xin Yao (CERCIA)

Particular thanks are due to the members of the Working Group convened by the UK Computing Research Committee, the Computer Science Committee of the London Mathematical Society, and the Scientific Committee and Faraday Advisory Board of the Smith Institute, who provided valuable guidance during various stages of drafting the report.

The MICSi interviews, background research and report compilation were carried out by staff of the Smith Institute: Dr David Allwright, Dr Tim Boxer, Dr Melvin Brown, Dr Robert Leese, Dr Paul Moseley and Dr Heather Tewkesbury.

CONTACT DETAILS

The Smith Institute for Industrial Mathematics and System Engineering

Surrey Technology Centre

Surrey Research Park

Guildford GU2 7YG


United Kingdom

Telephone: +44 (0) 1483 579108

Fax: +44 (0) 1483 568710

Email: office@smithinst.co.uk

www.smithinst.co.uk



The Smith Institute for Industrial Mathematics and System Engineering
Surrey Technology Centre, Surrey Research Park
Guildford GU2 7YG
United Kingdom
Telephone: +44-(0)-1483-579108
Fax: +44-(0)-1483-568710
Email: office@smithinst.co.uk
Web: www.smithinst.co.uk

Designed in the United Kingdom by Left Edge Marketing (07747) 842348