Report on the British Colloquium for Theoretical Computer Science (BCTCS16) April 10-12 2000, Liverpool, United Kingdom

Martyn Amos* Paul E. Dunne†

The 16th Annual Meeting of the British Colloquium for Theoretical Computer Science was hosted by the University of Liverpool between the 10th and 12th April 2000. The colloquium was held in the excellent Swallow Hotel in the centre of the city. This meeting was the first occasion on which the colloquium returned to an institution (BCTCS7 was also held in Liverpool).

The colloquium was formally opened by a welcoming address from the Vice Chancellor, Professor Phillip Love. The meeting featured an interesting and wide-ranging programme of invited talks by by some of the most distinguished members of the community. Invited talks were presented by Leslie Valiant (Harvard), Grzegorz Rozenberg (Leiden), Klaus Weihrauch (FernUniversität-Gesamthochschule in Hagen), Xin Yao (Birmingham), Michael Wooldridge (Liverpool) and Paul E. Dunne (Liverpool). The invited programme was well complemented by a total of 30 contributed talks.

BCTCS17 will be held at the University of Glasgow from April 9th to April 12th 2001. Anyone wishing to contribute talks concerning Theoretical Computer Science is warmly encouraged to do so. Further information regarding BCTCS17 may be found on the Web at

http://www.dcs.gla.ac.uk/bctcs17/

The BCTCS Web pages, giving details of previous colloquia are located at

http://www.csc.liv.ac.uk/~ped/bctcs/summary.html

Abstracts of invited talks

Leslie G. Valiant
Division of Engineering and Applied Sciences, Harvard University

Robust Logic

It has been recognized for centuries that cognitive phenomena exhibit both inductive as well as deductive aspects. The processes of induction and deduction have been studied systematically though separately in the frameworks of computational learning and computational logic. Since cognitive computations appear to perform these processes in combination, a single framework is required within which the two can be discussed simultaneously. Robust logics are designed to serve just that purpose. They are based on the view that a knowledge-base can be made robust only if each assertion in it is verifiable empirically against and learnable from real world observations. The challenge then is to reconcile this with the advantages offered by conventional logics, in particular a sound basis for deduction. Robust logics are designed to bridge this gap while retaining computational feasibility. In this framework both the computational work as well as the accuracy of both learning and deduction are polynomially controlled.

Grzegorz Rozenberg
Leiden Institute of Advanced Computer Science, Leiden University

DNA Computing *in vivo* - gene assembly in ciliates

DNA Computing is one of the new exciting developments in computing science. One branch of this area, DNA Computing in vivo, studies computational processes in living cells. In our lecture we will discuss the computational aspects of DNA processing in ciliates. Ciliates, a very ancient group of organisms, have evolved extraordinary ways of organizing, manipulating, and replicating the DNA in their micronuclear genomes. Especially interesting from the computational point of view is the process of gene assembly, which will be the subject of this lecture. The lecture is of a tutorial style.

*School of Biological Sciences and Department of Computer Science, University of Liverpool, Liverpool L69 3BX, England
†Department of Computer Science, University of Liverpool, Liverpool L69 3BX, England

Klaus Weihrauch
Theoretische Informatik I, FernUniversität-Gesamthochschule in Hagen

Computable Analysis

Computable analysis is a branch of computability theory studying those functions on the real numbers and related sets which can be computed by machines like digital computers. The increasing demand for reliable software in scientific computation and engineering requires a sound and broad foundation not only of the analytical/numerical but also of the computational aspects of real number computation. Although many researchers have been active in computable analysis, it has never belonged to the main stream of research in computability, our knowledge of this field is remarkably insufficient and only very few mathematicians or computer scientists know at all a definition of computable real functions. At present computable analysis appears as a juxtaposition of several partly independent approaches which are more or less far developed. For the interested newcomer this situation is bewildering, since there are not even generally accepted basic definitions, and so learning the state of the art from the fragments is a laborious undertaking. The talk introduces to concepts of TTE (Type-2 Theory of Effectivity), one of the today's approaches to computable analysis.

Xin Yao
School of Computer Science, University of Birmingham

Some Theoretical Issues in Evolutionary Computation

Evolutionary computation is the study of computational systems which use ideas and get inspirations from natural evolution. It has attracted a lot of interests in recent years. This talk first gives a brief introduction to the field, including evolutionary learning, optimisation and design. Some of the recent work will be surveyed. In spite of many successful applications of evolutionary computation techniques. Theories of explaining why and how these techniques work are still in their infancy. This talk will discuss some of the theoretical work in evolutionary computation. In particular, some recent results in the average time complexity of evolutionary algorithms will be presented. These results are very useful because they provide some insights into the computational power of evolutionary algorithms in solving certain problems.

Michael Wooldridge
Department of Computer Science, University of Liverpool

The Verification Problem for Agent Communication Languages

In recent years, a number of attempts have been made to develop standardized agent communication languages. A key issue in such languages is that of conformance testing. That is, given a program which claims to semantically conform to some agent communication standard, how can we determine whether or not it does indeed conform to it? To show that a program which claims to conform to a standard does indeed conform to it is to verify it. In this presentation, I introduce the verification problem for agent communication languages, and highlight some problems with current approaches to the semantics of agent communication, which mean that their verification is not possible. I then outline an alternative, expressive agent communication language, and give a semantics for this language in such a way that verification becomes a realistic possibility. The techniques I develop draw upon those used to give a semantics to reactive systems in theoretical computer science.

Paul E. Dunne
Department of Computer Science, University of Liverpool

Computational Problems (Some Directions but No Solutions)

The study of algorithms dates back over 2000 years and is central to Computer Science. The last 50 years have seen the growth of a rich formal Theory of Algorithmics, whose concerns have have resulted in many issues that pose significant questions regarding the fundamental limits of computational processes. Many of these questions, cloaked in different language and terminology, provide a framework highly relevant to long-standing concerns of Formal Logic and Combinatorial Mathematics.
In this talk some critical open issues in the Theory of Algorithms are reviewed. For the questions examined, a brief outline of the historical context and development will be given, together with a personal view of why their resolution matters and some thoughts on the prospects for such resolution being achieved.

Abstracts of contributed talks

Martyn Amos
School of Biological Sciences and Department of Computer Science, University of Liverpool

DNA-based Logic

Complex natural processes may often be expressed in terms of networks of computational components, such as Boolean logic gates or artificial neurons. The interaction of biological molecules and the flow of information controlling the development and behaviour of organisms is particularly amenable to this approach, and these models are well-established in the biological community. However, only relatively recently have papers appeared proposing the use of such systems to perform useful, human-defined tasks. Rather than merely using the network analogy as a convenient technique for clarifying our understanding of complex systems, it may now be possible to harness the power of such systems for the purposes of computation. In this paper we review several such proposals, focusing on the molecular implementation of fundamental computational elements.

Tom Berry
School of Computing and Mathematical Sciences, Liverpool John Moore's University

A linear time string matching algorithm on average with efficient text storage

In this talk we describe an algorithm to search for a pattern in a efficiently stored text. The method used to store the text reduces it to $\log_2 s/8$ of its original size, where s is the size of the alphabet set S. We prove that the algorithm takes linear time on average. We compare the new algorithm with some existing string matching algorithms by experimentation.

Alexander Bolotov
Department of Computing and Mathematics, Manchester Metropolitan University

General Normal Form for Propositional Temporal Logic

Temporal logic is considered as an essential tool in the areas of Artificial Intelligence and Computer Science when structures to be described concern temporal aspects. The underlying models of time are usually either linear sequences (so called discrete linear-time temporal logic) or a choice of possibilities branching into the future (so called branching-time temporal logic). Varieties of propositional linear and branching-time temporal logics are characterised by specific syntactic restrictions resulting in the different levels of expressiveness - from linear-time temporal logic (PLTL) to full temporal fixpoint calculi. However, much of the research into these temporal logics has centered around the model-checking technique, with an obvious lack of research into efficient proof methods.
We present here a unified approach to analysing temporal logics based upon a general normal form. The normal form was initially developed by Fisher as a part of a clausal resolution method for PLTL. Any statement within some temporal specification language describing, for example, a program behaviour, is represented as a set of constraints that express the initial and global conditions (now $\rightarrow$ next-time or now $\rightarrow$ sometime). The simple structure of the normal form and its potential to capture recursion enabled the efficient application of the resolution procedure. Both the normal form and the clausal resolution technique were extended further to a branching-time setting (systems of Computation Tree Logic (CTL) and CTL*. Moreover, unlike the language of PLTL, CTL or CTL*, the normal form can represent such regularities as "a property 'P' occurs at every even moment in time". This follows from the ability to encode B'uchi automata in the normal form. Since B'uchi automata are known to be as expressive as (propositional) temporal fixpoint logic, then the normal form and, hence, the clausal resolution method, can be potentially applied to expressive fixpoint logics. This is the subject of current work.

Julian Bradfield
Department of Computer Science, University of Edinburgh

Yet Another Causal Logic

We describe some experiments with expressing causal or true concurrent properties in a logic without a real partial order semantics. Our logic is defined on an extended state space of Petri nets, or any other model with locality. Consequently, for finite safe nets the state space is still finite, unlike partial order or event structure based logics. Surprisingly, this simple idea seems to give strong distinguishing power, and many 'natural' causality properties are expressible.
Joint work with Angelika Mader, Javier Esparza and Michaela Huhn.

Tom Chothia
Department of Informatics, University of Edinburgh

A Distributed Calculus with Localised Areas of Communication

I will introduces the Local Area Calculus. This calculus is an extension of the pi-calculus with syntax for locations and a novel form of binding. It aims to capture the notion of names that are universally known but always refer to local information. This is achieved by binding channels so that they have a number of mutually exclusive local areas. A channel may be sent between its areas and it may be used for communication inside these areas but it may not be used to send data between them.
I discuss a semantics that will block any communication by a channel between its areas or by a channel that is outside its area. It is then shown how a non-recursive type system can guarantee that channels will not be used to make such communications.

Leszek Gasieniec
Department of Computer Science, University of Liverpool

Deterministic broadcasting in unknown radio networks

We consider the problem of distributed deterministic broadcasting in radio networks of unknown topology and size. The network is synchronous. If a node u can be reached from two nodes which send messages in the same round, none of the messages is received by u. Such messages block each other and node u either hears the noise of interference of messages, enabling it to detect a "collision", or does not hear anything at all, depending on the model. For the purpose of the presentation we adopt a weaker model without collision detection. We assume that nodes are completely ignorant of the network: they know neither its topology, nor size, nor even their immediate neighborhood. The initial knowledge of every node is limited to its own label. We study the time of deterministic broadcasting under this total ignorance scenario.
Previous research has concentrated on distributed randomized broadcasting algorithms working for unknown networks, and on deterministic off-line broadcasting algorithms assuming full knowledge of the radio network. Ours are the first broadcasting algorithms simultaneously distributed and deterministic, that work for arbitrary totally unknown radio networks.
Presented results include: optimal linear-time broadcasting algorithm for symmetric graphs, lower bound Omega(D log n), for arbitrary n-node graphs of diameter D, $O(n^{11/6})$, $O(n^{9/5})$, ..., $O(n^{3/2})$-time broadcasting algorithms working in arbitrary n-node graphs
Joint work with: B.S.Chlebus, A.M.Gibbons, A.Ostlin, A.Pelc, W.Rytter, and M.Robson.

Leslie Goldberg
Department of Computer Science, University of Warwick

A bound on the capacity of backoff and acknowledgement-based protocols

We study contention-resolution protocols for multiple-access channels (such as the Ethernet channel). We show that EVERY backoff protocol is unstable if the arrival rate, lambda, is at least 0.42. Thus, we show that backoff protocols have (provably) smaller capacity than "full-sensing" protocols. Finally, we show that the corresponding results, with the larger arrival bound of 0.531, also hold for every acknowledgement-based protocol. (The talk will not presume that the listener has any prior knowledge of contention-resolution protocols.) Joint work with Mark Jerrum, Sampath Kannan and Mike Paterson

Paul Goldberg
Department of Computer Science, University of Warwick

Learning Fixed-dimension Linear Thresholds From Fragmented Data

We investigate inductive learning in a situation in which examples (consisting of an input vector and binary label) have some of the components of the input vector concealed from the learner. This is a special case of Restricted Focus of Attention (RFA) learning. Our interest here is in 1-RFA learning, where only a single component of an input vector is given, for each example. 1-RFA learning is the most restrictive form of RFA learning (so that positive results apply in general), and it models a typical "data fusion" scenario, where we have sets of observations from a number of separate sensors, but these sensors are uncorrelated sources.
Within this setting we study the well-known class of linear threshold functions, the characteristic functions of Euclidean half-spaces. We give results concerning effects of the probability distribution over input vectors on the number of examples needed to learn, and a general- purpose algorithm which is efficient for any constant input dimensionality. The sample-size requirement is polynomial in input dimension, for various input distributions that satisfy a "smoothness" condition.

Graham Goulbourne
Department of Computer Science, University of Liverpool

Tree Algorithms For Data Mining

The problem of extracting all association rules from within a binary database is well-known. Existing methods may involve multiple passes of the database, and cope badly with densely-packed database records because of the combinatorial explosion in the number of sets of attributes which must be considered. We present a new class of methods founded on the concept of a P-tree (a partial support tree), generated from a single pass of the database in time linear to the size of the database; and in particular the mechanism for deriving associations rules from such a P-tree using the further concept of a T-tree (Total support tree). The T-tree, generated from a given P-tree then contains all the summation required for the final derivation of association rules. The advantage offered by the method is that it allows large (500+ columns), densely packed, databases to be analysed with less computation than other approaches.

David J. Haglin and Rupert W. Ford
Department of Computer and Information Sciences, Minnesota State University in Mankato (Currently visiting the Department of Computer Science, University of Manchester)

On a Message Minimizing Work Redistribution Problem

A Message Minimizing Problem is described which arises from a work redistribution and load balancing problem in parallel computing. Unlike many load balancing approaches that utilize local information and pass work to neighboring processors, we consider a global perspective and seek a redistribution plan that minimizes the overall processing time. This problem can Arise in many applications, such as in the U.K. Meteorological Office's weather forecasting model.
Given a situation where some processors have more work to do than others, redistributing the workload will decrease the total wall clock time required. Those processors with more work than the average workload are called overloaded, and the others are called underloaded. A solution to this problem is a plan for sending work from the overloaded processors to the underloaded processors in such a way as to achieve a perfect load balance. The cost associated with a solution is defined to be the number of packets needed to balance out the workload. The impact of potential collisions on the interconnection network as well as the size of the packets are ignored. This problem is shown to be NP-Complete, and various heuristics and approximation issues are investigated. Several theoretical results are shown that impact the design of a heuristic. Experimental results are also given.

James Heather
Department of Computer Science, Royal Holloway

The tragic tale of Alice and Bob

We live in an increasingly connected world. Our reliance on security protocols, that is, mechanisms to ensure the secrecy of sensitive communications and the authenticity of the participants involved, is becoming greater and greater. However, many of these protocols have never been proven to be watertight.
In this paper, we explain what a security protocol is, what it aims to achieve, and what it means for a protocol to be 'correct'. We illustrate with a well-known protocol, which was 'proved' correct when it was first proposed. However, it turns out to be fatally flawed; we give an attack on the protocol, and show how it may be fixed to prevent the attack.
We show how a network running a security protocol can be described in the process algebra CSP, and demonstrate how to model a powerful hostile intruder who may attempt to disrupt the operation of the protocol. We then show how this model may be used to find attacks on the protocol, or to prove correctness for a small system running the protocol.
Finally, we discuss the problem of analysing the protocol running on a network with an arbitrarily large number of agents. We introduce RankAnalyser, a tool designed to automate verification of security protocols on such a network.

Costas S. Iliopoulos, Laurent Mouchard and Yoan J. Pinzon
Department of Computer Science, King's College London

The Max-Shift Algorithm for Approximate String Matching

The approximate string matching problem is to find all locations which a pattern of length m< matches a substring of a text of length n with at most k differences. The program agrep is a simple and practical bit-vector algorithm for this problem. In this paper we consider the following incremental version of the problem:

given an appropriate encoding of a comparison between A and bB, can one compute the answer for A and B, and the answer for A and Bc with equal efficiency, where b and c are additional symbols? Here we present a simple bit-vector algorithm for answering these questions that requires only O(n/w) time, where n is the maximum length of A and B. We also present an O(nm/w) algorithm for the fixed-length approximate string matching problem: given a text t, a pattern p and an integer h, compute the optimal alignment of all substrings of p of length h and a substring of t. We also show how to apply this algorithm to the cyclic string comparison problem, given text $t=t_1t_2...t_n$ and the pattern $p=p_1p_2...p_m$. Let $cycle(t_1t_2...t_n)=t_2...t_nt_1$, and let $cycle^r(t)$ be the result of applying cycle exactly r times. The cyclic string comparison problem is to determine the integers r and s such that d=delta($cycle^r(t)$, $cycle^s(p)$) is minimal.

Robert W. Irving and Lorna Love
Department of Computing Science, University of Glasgow

The Suffix Binary Search Tree

Suffix trees and suffix arrays are classical data structures that are used to represent the set of suffixes of a given string, and thereby facilitate the efficient solution of various string processing problems - in particular on-line string searching. In this talk I will illustrate the potential of suitably adapted binary search trees as competitors in this context. The suffix binary search tree (SBST) and its balanced counterpart, the suffix AVL-tree are conceptually simple, relatively easy to implement, and offer time and space efficiency to rival suffix trees and suffix arrays, with some distinct advantages - for instance in cases where only a subset of the suffixes need be represented. Construction of a suffix BST for an n-long string can be achieved in O(nh) time, where h is the height of the tree. In the case of a suffix AVL-tree this will be O(nlogn) in the worst case. Searching for an m-long substring requires O(m + l) time, where l is the length of the search path. In the suffix AVL-tree this is O(m+logn) in the worst case. The space requirements are linear in n, intermediate between those for a suffix tree and a suffix array. Preliminary empirical evidence, illustrating the competitiveness of suffix BSTs in practice, will also be presented in the talk.

Stephen Lakin
Department of Mathematics and Computer Science, Leicester University

Groups with context-sensitive decision problems

The issue of decision problems in groups, particularly the word problem, has been widely studied. However much of this work has focussed mainly on the solvability of the problems and less has been done in studying the constraints on the resource bounds for these computations in terms of time and space, particularly space considerations.
Here we are interested particularly in context-sensitive algorithms, that is those that operate in non-deterministic linear space. We consider some closure properties of the class of groups that have context-sensitive word and conjugacy problem and demonstrate via the reduced and irreducible word problem the power of context-sensitive algorithms when dealing with subwords of a given word.
The aim of this talk is to give an general introduction to and overview of the research area rather than presenting a series of technical results.

Gerald Luettgen and Michael Mendler
Department of Computer Science, Sheffield University

What is in a Step, Exactly?

Pnueli and Shalev presented two equivalent formalizations of Statecharts' macro-step semantics in a seminal paper in 1991. This semantics, however, lacks compositionality. Huizing and Gerth showed that combining compositionality, causality, and the synchrony hypothesis cannot be done within a simple, single-leveled semantics. Some researchers subsequently devoted their attention to investigating new variants of Statecharts, obeying just two of the three properties. Other researchers combined all three properties by storing complex semantic information via preorders or transition systems. However, no analysis of exactly how much information is needed to achieve compositionality has been made so far. This paper provides such an analysis by taking a logical or axiomatic approach. We first study the compositionality defect under the traditional logical interpretation of Statecharts and trace it to the invalidity of the Law of the Excluded Middle. We then show that by adopting an intuitionistic interpretation a fully-abstract macro-step semantics falls out. The resulting model theory is based on a particular class of intuitionistic Kripke models, called behaviours. Our results yield, for the first time in the literature, a simple fully-abstract semantics for Statecharts, which also interprets Pnueli and Shalev's concept of failure naturally. As a by-product, our logical approach suggests a natural way of admitting disjunctions

in transition triggers, thereby solving a semantical inadequacy of Pnueli and Shalev's setting. The results not only give insights into the semantic subtleties of Statecharts, but also provide a basis for developing algebraic theories for macro steps and for comparing different Statecharts variants.

Saeed Mahfooz, Madjid Merabti and Reuben Peirera
School of Mathematical and Computing Sciences, Liverpool John Moores University

Differential Services Architecture

This paper comprises the motivation of our research-followed by summary of our proposed architecture, our findings. Distributed multimedia applications require guarantees on network performance, such as bounds on throughput, delay, jitter, and reliability. Services that guarantee such bounds include the Resource Reservation Protocol, (RSVP) and guaranteed service class that are currently used to support multimedia applications. These services consume too much of the Internet resources while providing service to the multimedia applications. There is need for such services that could offer all the guarantees and at the same time not putting too much pressure on the scarce Internet resources. Internet engineering task force (IETF) has assigned a special group, Differential Service group (DIFFSER, March 30th, 1998) to come up with new techniques for the future Internet services.
Sharing (RBS) differential architecture. In our architecture we proposed a three-tier mechanism having services classes (Premium, Standard and Best-effort) that fulfils users different applications needs. This architecture extends, both the traditional, best-effort and Differential Services models. Our framework considers two important aspects:

- Bandwidth allocation Scheme
- Quality of Service allocation

The RBS architecture describes the layout of the routers in the Internet Service Providers (ISP) domain containing Border Routers (BR), and Core Routers (CR) . Border router main function includes Relative Bandwidth sharing according to their Service Level Agreement (SLA). Other functions include Packets Classification, Shaping and Marking. The classifier in BR identifies packets of different service classes from the DS byte of the incoming packet stream. In our architecture we have focused more on the role of the border router because if traffic is shaped and policed well at the border it reduces the chances of congestion inside the ISP domain. Hence core routers can have more time for computation and optimal route selection for delay sensitive packets inside ISP domain in order to provide Quality of Service (QoS).
Core router performs optimal resources allocation for QoS to the service classes inside ISP domain. The packets that enter ISP domain after shaping and policing are offered QoS to the delay sensitive traffic through DS based-routing technique. Need for QoS to the delay sensitive traffic arise mainly from the uneven distribution of the traffic in the ISP routers by the dynamic routing protocol OSPF. We are using bandwidth and hop-count as the metrics for path selection to provide the required QoS to the service. The use of bandwidth as metrics also provides the services (premium and standard) the amount of bandwidth they require, but the problem is that it requires more computational time for the routers to select the path that has the required bandwidth. Therefore to further reduce the computational time, pruning could be a solution where all the routes where the bandwidth is less than the required bandwidth are pruned first, and then the required route is selected in the remaining routes. Some of the measures that could reduce the computational overhead are:

- Use of large value timer
- Choosing bandwidth and hop-count as metrics for path selection
- Pruning the unsuitable links before computing the routing table

Therefore the QoS mechanism is to select optimal routes for the service classes that are likely to meet the required QoS. This QoS mechanism complement differential services in general and our proposed mechanism in particular to deliver QoS to the delay sensitive traffic. In order to evaluate the performance of our scheme we conducted simulations . Where we presented a mechanism to provide Quality of Service (QoS) guarantees to different multimedia applications that share link bandwidth in IP-based differential services domain. In this mechanism weights are associated with each service class, providing link share for service classes and to individual users according to their priorities. The test data used portray different multimedia appli cations i.e. MPEG-2, IP telephony. The simulation results obtained show the effectiveness of our scheme for multimedia applications by allocating link share to each multimedia application and minimising end-to-end transmission delay by bringing them in line with the recommended standard acceptable transmission delay for multimedia applications.

Grant Malcolm and Ray Paton
Department of Computer Science, University of Liverpool

Computing with Proteins: Algebraic Models for Protein Signalling Networks

Protein interactions within a cell display various computational properties, such as memory storage, switching and communication through intracellular signalling systems. In this talk we apply some ideas from computer science to describe the information processing nature of signalling systems. We view the signalling ecology as a parallel distributed processing network of agents, and we give an overview of the algebraic and topological approaches that we feel will be useful in modelling such systems.

David Manlove
Department of Computing Science, University of Glasgow

Medical matching with indifference

The annual process of assigning graduating medical students (also known as residents) to their first hospital appointments is accomplished by a centralised matching scheme in several countries. Each resident provides a strictly ordered preference list of hospitals, and vice versa. The algorithm constructs an allocation of residents to hospitals known as a stable matching, based on these preference lists. However in practice, it is possible that a hospital might be indifferent between two or more residents (or vice versa) so that agents may wish to express ties in their preference lists. In this talk, I will present a new algorithm for matching residents to hospitals when ties are permitted.
Joint work with Rob Irving and Sandy Scott

Florent Renaud Mathias Madelaine
Department of Mathematics and Computer Science, Leicester University

Application of the theorem of Baker-Pixley to the constraint satisfaction problem (CSP)

CSP is hard in general. An interesting question consists of finding all tractable restrictions of the general problem. Schaeffer, in an article concerning Generalized-Sat (a restriction of CSP to Boolean domain) has shown that such restrictions can be characterised by some kind of closure property that have been studied in a particular domain of Universal Algebra, called clone theory. Since that, many results from Universal Algebra have been successfully applied to the study of CSP.
In this talk, I will introduce the Constraint Satisfaction Problem and briefly give an idea of clone theory. I shall present the theorem of Baker-Pixley (a generalisation of the Chinese Reminder theorem) and explain how it can be used in the study of CSP in connection with consistency algorithms.

Duncan W. Parkes and Richard M. Thomas
Department of Mathematics and Computer Science, University of Leicester

Groups, formal languages and string rewriting

There are many interesting connections between group theory and formal language theory. One of these is the comparison of the algebraic structure of a group with the complexity of its word problem, or with the complexity of some specified subset of the word problem such as the reduced word problem or the irreducible word problem. This last subset has a close relationship with finite special string-rewriting systems (a string rewriting system is said to be special if the right hand side of every rule is the empty word) which are confluent on the Thue-congruence class containing the empty word.
The purpose of this talk will be to explain what all these notions are and to mention some recent results in the area.

Ray Paton
Department of Computer Science, University of Liverpool

Is there a biology of quantum information?

This paper briefly considers the notion of a biology of quantum information from a number of complementary points of view. We begin with a very brief look at some of the biomolecular systems that are thought to exploit quantum mechanical effects and then turn to the issue of measurement in these systems and the concomitant generation of information. This leads us to look at the internalist stance and the exchange interaction of quantum particles. We suggest that exchange interaction can also be viewed using ecological ideas related to apparatus-object. This can also help develop the important notion of complementarity in biosystems in relation to the nature and generation of information at the microphysical scale.
Joint work with Koichiro Matsuno

Igor Potapov
Department of Computer Science, University of Liverpool

Logical analysis of communicating finite-state machines model under additional local conditions

The model of communicating finite-state machines (CFSM) underlies the formal description techniques Estelle and SDL, used for the specification of communication protocols. However, the CFSM model has the same expressive power as Turing machines, hence a lot of correctness problems connected with behaviour of network (is the communication bounded, free from deadlocks and unspecified receptions) are algorithmically undecidable in general. In the CFSM model, a protocol is specified as an oriented graph G, where nodes are process represented by finite-state machines which communicate asynchronously by messages and edges are unbounded FIFO channels which allows the communication. Much of the research has been devoted to identifying the classes of cyclic protocols with decidable logical errors. We will consider the protocols with topology is specified as an oriented graph G. In this paper, we show that logical errors are decidable for special classes with local restrictions. The possibility of applying previous results to analysis network communication with complexity topology was looked at. The main idea was to show the influence of local restrictions on the behavior of the network.

Naila Rahman and Rajeev Raman
Department of Computer Science, King's College London

Adapting radix sort to the memory hierarchy

We demonstrate the importance of minimising misses in the translation-lookaside buffer (TLB) for obtaining good performance on modern computer architectures. We focus on least-significant-bit first (LSB) radix sort, standard implementations of which make many TLB misses. We give three techniques for reducing TLB misses for LSB radix sort: reducing working set size, explicit block transfer and grouping by pre-sorting. We note that: 1) Choosing the radix size to minimise TLB misses greatly improves performance relative to radix sizes which are chosen on the basis of cache performance alone, even though doing so nearly doubles the number of operations performed. 2) All TLB-optimised algorithms outperform LSB radix sort tuned for cache alone and cache-tuned implementations of comparison-based sorting algorithms. Pre-sorting outperforms all other implementations, improving on the other TLB-optimised algorithms by over 20over 45the worst case. This is provably not true for standard LSB radix sort. We also apply these techniques to the problem of permuting an array of integers, and obtain gains of over 30

Naila Rahman and Rajeev Raman
Department of Computer Science, King's College London

Analyzing the cache behaviour of non-uniform distribution sorting algorithms

We analyze the average-case cache performance of distribution sorting algorithms in the case when keys are independently but not uniformly distributed. We use this analysis to tune the performance of the integer sorting algorithm MSB radix sort when it is used to sort independent uniform floating-point numbers (floats). Our tuned MSB radix sort algorithm comfortably outperforms a cache-tuned implementation of bucketsort when sorting uniform floats from [0, 1).

Stephan Reiff
Department of Computing Science, University of Glasgow

Identifying Resolution Choices for an Online Feature Manager

There are many approaches to resolving feature interactions when they have been detected. These can be characterised as online (at runtime) or offline (at design time). The latter usually make use of some formal method.
We aim to combine both these techniques, and follow the approach of Marples and Magill. They use a transactional method (that is, they explore all possibilities, then rollback and commit to one such possibility). The major drawback with their work is that they employ a weak resolution mechanism, i.e. the method to select the "best" behaviour to be commited.
In order to improve resolution, we have developed a formal model of feature behaviour (as sets of traces). Based on this model, we can then define the space of all possible resolutions, and select the "best" one. The former task in itself is highly interesting and turns out to be very complex (and combinatorially explosive). A particular challenge is eliminating traces which represent inconsistent (or simply wrong) behaviour.

Anthony Roy and John Stell
Department of Computer Science, Keele University

Spatial Relations Between Indeterminate Regions.

Systems of relations between regions are an important aspect of formal theories of spatial data. Examples of such relations are part-of, partially overlapping, and disjoint. One particular family of systems is that based on the Region-Connection Calculus (RCC). These systems of relations were originally formulated for ideal regions, not subject to imperfections such as vagueness or indeterminacy.
This paper presents two new methods for extending the relations based on the RCC from crisp regions to indeterminate regions. As a formal context for these two methods we develop an algebraic approach to spatial indeterminacy using Luk algebras. This algebraic approach provides a generalisation of the "egg-yolk" model of indeterminate regions. The two extension methods which we develop are proved to be equivalent. In particular, it is shown that it is possible to define part-of in terms of connection in the indeterminate case. This generalises a well-known result about crisp RCC regions. Our methods of extension take a relation on crisp regions taking values in the set of two Boolean truth values, and produce a relation on indeterminate regions taking one of three truth values. We discuss how our work might be developed to give more detailed relations taking values in a six element lattice.

W. Rytter
Department of Computer Science, University of Liverpool

On the complexity of solving word equations

Word equations are used to describe properties and relations of words, e.g. pattern-matching with variables, imprimitiveness, periodicity. It is known that the solvability problem for word equations is NP-hard, even if we consider (short) solutions with the length bounded by a linear function and the right side of equations contains no variables. Recently it has been shown by W. Plandowski that the solvability of word equations is in P-SPACE, this algorithm is the milestone achievement is this area.
The main open problem is to close the gap between NP and P-SPACE. We close this gap for a restricted version of the problem. A very simple extension of the Plandowski algorithm is shown to the case when the variables obey regular constraints: the string-value of each variable should be a member of a corresponding regular language, whose description is part of the input.
We show that such variation of the word equations problem becomes P-SPACE complete. This resolves completely the complexity characterization of the solvability of word equations with regular constraints. We also discuss the importance of the text-compression approach to the solution of word equations. Using this approach it has been shown that the unrestricted solvability problem is in NP if the length of minimal solutions is bounded by a singly exponential function (there is a general belief that such upper bound is correct, at least no counterexample has been found).

Peter Saffrey
Department of Computer Science, University of Glasgow

Model Checking and Abstraction

Abstraction is a method used to reduce the size of a state space search, typically by compressing large variables into smaller, simpler ones. This method is not always suitable to all problem domains. I will discuss one such problem domain, that of a telephone system, and discuss why these techniques are unsuitable. I will then suggest what I intend to do to adapt them to my purposes.

Paul Strickland and Dhiya Al-Jumeily
School of Computing and Mathematical Sciences, Liverpool John Moores University

Using term rewrite rules to generate intelligent advice for students of algebra

We present algorithms for guiding an intelligent tutoring system for simple algebra. Taking two consecutive terms in a simple algebraic argument, we attempt to construct an 'explanation' of any discrepancies encountered. To achieve this we build up two partial trees of terms, derived from those input by a selection of rules and malrules. The algorithm is guided by the use of a simplification ordering. When two equivalent terms are found we use them, together with a record of the malrules used, to construct intelligible feedback for the student.

M.J. Wooldridge and Paul E. Dunne
Department of Computer Science, University of Liverpool

Tractable and Intractable Classes of Agent Design Problems

The Agent Design problem involves determining if it is possible to construct an agent satisfying a given specification within a particular environment. The most basic examples of such specifications are Achievement Tasks (AD, where an agent is required to attain a particular goal state) and Maintenance Tasks (MD, where an agent is required to avoid entering a particular state). Within the most general definition of 'environment', the first author has shown both to be PSPACE-complete. This result has, recently, been extended by the present authors to categorise the complexity of Achievement and Maintenance tasks in a number of more limited settings, under which the problems are NP-complete or admit of polynomial-time algorithms. In this talk the effect on problem complexity of specifying tasks by arbitrary propositional formulae is considered: AD and MD being special cases. We show that, in those settings where AD is PSPACE or NP-complete the general problem is 'no more difficult', i.e. remains in PSPACE or NP. In contrast, a number of the settings for which there are polynomial-time methods, become NP-complete for general propositional formulae.