REPORT ON BCTCS 2011

The 27th British Colloquium for Theoretical Computer Science
18-21 April 2011, University of Birmingham

Achim Jung and Paul Levy

The British Colloquium for Theoretical Computer Science (BCTCS) is an annual forum for researchers in theoretical computer science to meet, present research findings, and discuss developments in the field. It also provides an environment for PhD students to gain experience in presenting their work in a wider context, and benefit from contact with established researchers.

BCTCS 2011 was hosted by the University of Birmingham, and held during 18–21 April 2011. The event attracted just over 50 participants, and featured an interesting and wide-ranging programme of 6 invited talks and 24 contributed talks; the majority of the contributed talks were given by PhD students.

The programme contained a pleasant spread of topics from Algorithms to Semantics, from Security to Natural Computation, and from Formal Methods to Quantum Computation. We are grateful to our invited speakers David S. Johnson, Cliff Jones, Prakash Panangaden, Peter Selinger, Nigel Smart, and Carsten Witt for recognising the interdisciplinarity of the event and for pitching their talks accordingly.

The meeting was smaller than in previous years since we were not able to secure support for PhD students from EPSRC, and many students had to withdraw their preliminary registration. On the positive side, we were able to schedule all talks without parallel sessions, and perhaps achieved increased interaction between communities in this way. It was certainly pleasing to note that most participants attended all presentations and there was a lively question-answer session after every talk.

Abstracts for all of the talks are provided below. The financial support of the London Mathematical Society (LMS) is gratefully acknowledged.

BCTCS 2012 will be hosted by the University of Manchester from 2–5 April 2012. Researchers and PhD students wishing to contribute talks concerning any aspect of theoretical computer science are warmly welcomed to do so. Further details are available from the BCTCS website at `www.bctcs.ac.uk`

Invited Talks at BCTCS 2011

David S. Johnson, AT&T Laboratories, New Jersey
Bin Packing: From Theory to Experiment and Back Again

In the bin packing problem, one is given a list of 1-dimensional items and asked to pack them into a minimum number of unit-capacity bins. This was one of the first NP-hard problems to be studied from the "approximation algorithm" point of view, and over the years it has served as a laboratory for the study of new questions about approximation algorithms and the development of new techniques for their analysis. In this talk I present a brief survey of this history, highlighting the many surprising average case behaviour results that have been obtained. Several of these surprises were first revealed by experimentation, which led to conjectures and then to proofs, and I will describe this interplay between experimentation and theory. I will also highlight some as-yet-unproven conjectures suggested by the experimental data.

Cliff Jones, University of Newcastle
AI4FM - How To Say "Why" In Proofs

The AI4FM project is exploring how to use AI thinking in discharging proof obligations (POs) that arise in the formal development of programs. Of course, developing better and better heuristics for theorem provers has been explored by AI researchers since its early days. The approach being taken in this project (with Alan Bundy providing the main AI input) is to design a system that can learn from how an expert tackles POs that are beyond current heuristics. We have industrial data that suggests as few as five ideas might kill off a hundred undischarged POs. The first step is to design ways of recording the insights which are often about "why" the expert chose a strategy; we will move on to learning such strategies in the next phase of the AI4FM project.

Prakash Panangaden, McGill University, Montreal
Epistemic Strategies and Games on Concurrent Processes

We develop a game semantics for process algebra with two interacting agents. The purpose of our semantics is to make manifest the role of knowledge and information flow in the interactions between agents and to control the information available to interacting agents. We define games and strategies on process algebras, so that two agents interacting according to their strategies determine the execution of the process, replacing the traditional scheduler. We show that different restrictions on strategies represent different amounts of information being available to a scheduler. We also show that a certain class of strategies corresponds to the syntactic schedulers of Chatzikokolakis and Palamidessi, which were developed

to overcome problems with traditional schedulers modelling interaction. The restrictions on these strategies have an explicit epistemic flavour. This is joint work with Konstantinos Chatzikokolakis and Sophia Knight, both at INRIA Saclay.

Peter Selinger, Dalhousie University, Halifax
Logical Methods in Quantum Information Theory

I will talk about some recent applications of logical methods to quantum information theory. In computing, a higher-order function is a function for which the input or output is another function. I will argue that many of the interesting phenomena of quantum information theory involve higher-order functions, although that is not how they are usually presented. I'll talk about the quantum lambda calculus as a possible framework to describe such phenomena.

Nigel Smart, Bristol University
Homomorphic Encryption

The big story in cryptography over the last two years has been the discovery of a fully homomorphic encryption scheme. Such schemes, if made practical, could provide a fundamental paradigm shift in how we build secure online services, such as envisaged by the shift to cloud computing. However, even now we can practically perform interesting advanced applications using standard homomorphic encryption. In this talk I will look at some standard homomorphic encryption schemes, and describe a voting application. I will then describe in simple terms how the new fully homomorphic encryption schemes work.

Carsten Witt, Technical University of Denmark, Copenhagen
Bio-Inspired Computation Meets Theoretical Computer Science

Biologically-inspired methods of computation such as evolutionary algorithms, ant colony optimisation, particle swarm optimisation etc. are well established in numerous engineering disciplines. As their computational complexity was generally unknown, these approaches lived in the shadows of theoretical computer science for a long time. This dramatically changed in recent years, when bio-inspired computation was finally perceived as a family of algorithms and analysed using methods from classical algorithms and complexity theory. This talk will present some of the most exciting results that have been obtained in this new research area. We focus on evolutionary algorithms and ant colony optimisation in combinatorial optimisation and prove how they efficiently find optimal or approximate solutions to problems known from the theory of algorithms.

Contributed Talks at BCTCS 2011

Florent Balestrieri, University of Nottingham
The undecidability of pure stream equations

Polymorphic stream functions can be defined by pure stream equations. While there always exists a solution for them, it is not always the case that it is unique: the equations may provide only an under-specification for a function. Such a definition in a functional programming language would be non-productive, and the program might freeze forever, waiting for further data that won't be computed. In this article, we show that unicity of solutions is undecidable by giving a reduction from the generalised Collatz problem, which has been proved Π_2^0-complete by Stuart A. Kurtz and Janos Simon.

Liang-Ting Chen, University of Birmingham
A final Vietoris topology coalgebra construction

The construction of final coalgebras in Set is mysterious. In general the construction involves a universal quotient in the super-large category of classes not in Set. With the cardinality condition, the question is still not simple. Take the finitary powerset functor as an example. We know it converges at $\omega + \omega$, but we do not know where it stops exactly. However, with topology the construction becomes simpler. By applying Stone duality, the construction of final coalgebras in topological spaces is equivalent to the construction of initial algebras in the category of open sets, i.e. frames (with mild conditions). This approach gives a simpler way to calculate and provides applications in coalgebraic logic and domain theory.

Konrad Dabrowski, University of Warwick
Parameterized complexity of finding induced matchings

Many graph-theoretic problems are difficult to solve in general. However, this is not the end of the story. Sometimes we only want to solve the problem on certain types of graphs. This gives us some extra knowledge about the structure of the graph, which we can use to our advantage. We introduce a parameter which encodes some of this structure. If we are lucky, we find that this parameter somehow restricts all the "non-polynomial behaviour" of the problem. I will give some examples where this approach works, using some Ramsey-theoretic results. In particular, I will talk about the induced matching problem. This is the problem of finding an induced subgraph H in a graph G such that all the vertices in H have degree exactly 1 and H is as large as possible. I will also give a brief introduction to parameterized complexity.

Thomas Davies, Swansea University
An analysis of CSP implementation techniques

The formal language CSP has many applications in the area of critical systems, thanks to its handling of concurrent processes. We seek to study and compare four CSP implementation techniques, which take a CSP model and create an executable program from it, and perform an analysis of the similarities and differences between them. We will use the Byzantine Generals problem, a problem embedded in the heart of safety-critical computer systems, as a master example in this study.

Laurence E. Day, University of Nottingham
Towards modular compilers for effects

Compilers are traditionally factorised into a number of separate phases such as parsing, type checking, code generation etc. However, there is another potential factorisation that has received comparatively little attention: the treatment of separate language features such as mutable state, input/output, exceptions, concurrency and so forth. In this talk I will focus on the problem of modular compilation, in which the aim is to develop compilers for separate language features and prove their correctness independently, which can then be combined as required. I will illustrate how Haskell addresses this concept by way of example, before demonstrating a solution for modular syntax in Haskell. I will then show how this syntax can be used for flexible language construction, and conclude by briefly discussing modular semantics and future work.

Murdoch J. Gabbay, Heriot-Watt University
Metamathematics based on nominal terms: first-order logics over nominal sets

Nominal sets have enabled the development of extensions of first-order logic with term-formers that can bind. This in turn has enabled the development – and proof of correctness – of finite first-order axiomatisations of systems like the λ-calculus, first-order logic, and arithmetic. First-order logics over nominal sets can specify and reason about systems that would otherwise require infinite axiom schemes or higher orders. In this talk I will sketch nominal algebra and permissive-nominal logic, why we created them, and how I envisage them being applied.

Vashti Galpin, University of Edinburgh
Stochastic HYPE: modelling stochastic hybrid systems

Stochastic HYPE is a process algebra for modelling natural and artificial systems that exhibit stochastic, continuous and discrete behaviour. It extends the hybrid process algebra HYPE with events that have exponentially-distributed durations. The semantics of HYPE systems are given by Transition-Driven Stochastic Hybrid Automata, a subset of Piecewise Deterministic Markov Processes. This presentation will introduce stochastic HYPE, and show how it is possible to ascertain

that a stochastic HYPE model does not have infinite sequences of discrete behaviour through a syntactic analysis of the model. Such models are described as well-behaved. This technique will be illustrated with the example of a railway crossing. This is joint work with Jane Hillston, University of Edinburgh and Luca Bortolussi, University of Trieste.

James Gate, Durham University
The expressibility of fragments of hybrid graph logic on finite digraphs

Hybrid modal logic is an extension of modal logic that introduces new structural and language features that allow for the use of nominals (named points). By applying hybrid modal logic to graphs and adding the ability to verify the existence of a path in the graph we obtain Hybrid Graph Logic, of which we study the finite model theory. This logic is of interest as it can represent problems such as connectivity and even some co-NP-complete problems whilst still remaining decidable. We have developed a pebble game for Hybrid Graph Logic and use it to demonstrate infinite expressibility hierarchies over finite digraphs in fragments of the logic that are parameterized by the quantifier-rank of formulae, the number of propositional symbols and the number of nominals symbols that are available. Here we present these results by playing the pebble game on two structures and then arguing that whilst one player wins the game played over r rounds, the other player wins the game played over $r + 1$ rounds. This means that the two structures are r-equivalent but not $(r + 1)$-equivalent and that they separate the fragments of Hybrid Graph Logic consisting of formulae of quantifier-rank r from $(r + 1)$. We also present structures that have this property for varying numbers of propositional symbols or nominals.

Julian Gutierrez, University of Edinburgh
Concurrent logic games on partial orders

Most games for analysing concurrent systems are played on interleaving models, such as labelled transition systems or infinite trees. However, several concurrent systems have partial order models (e.g., Petri nets or event structures) rather than interleaving ones. As a consequence, a potentially algorithmically undesirable translation from a partial order setting to an interleaving one is required before analysing them with traditional techniques. As a first attempt to address this problem, we study a game played directly on partial orders and describe some of its algorithmic applications. The game provides a unified approach to verification which applies to different decision problems (e.g., bisimulation or model-checking) and models of concurrency.

Matthew Gwynne, Swansea University
On the hardness of (satisfiable) conjunctive normal forms

A new measure $h(F)$ of "hardness" for formulae F in conjunctive normal form (i.e., clause-sets) is introduced. $h(F)$ for unsatisfiable clause-sets has been studied extensively by Kullmann. However, now we treat satisfiable clause-sets F differently. $h(F)$ for satisfiable F as defined in earlier work by Kullmann has a specific algorithmic viewpoint. It measures resources based on forced assignments and probing, progressing in breadth-first manner. Now we consider boolean functions and their representations, and our new measure looks at "complete" representations of the underlying boolean function of F. We present the SAT representation hypothesis: The task of solving or refuting a SAT problem efficiently is captured by constructing a representation of the underlying boolean function which is of low hardness. For large boolean functions, the task becomes then to find good decompositions such that the component functions have low hardness. What it means to be a good decomposition is the next fundamental question. We consider translations of the Advanced Encryption Standard (AES) and Data Encryption Standard (DES) as examples. We provide translations which have low hardness for all subfunctions in natural decompositions. The performance of SAT solvers on these instances is experimentally investigated. This is joint work with Oliver Kullmann.

Peng He, Imperial College London
Visual hulls from imprecise polyhedral scenes

The visual hull is the best 3D shape of an object one can retrieve from its 2D silhouettes. Classical visual hull methods fail to maintain the exactness and robustness when the input is imprecise. In the solid domain, geometric shapes with imprecision are well modelled and carefully studied. We present a domain theoretic framework to compute the visual hull of a polyhedral scene, in which the vertices of the polyhedra are given with some imprecision. The overall algorithm maintains the same computational complexity as the classical method and generates a partial visual hull which converges to the classical visual hull as the input converges to an exact value.

Phillip James, Swansea University
Towards domain specific languages for verification

In this talk, we present a first study into a methodology for the design of domain specific languages for verification. Often within verification, concrete knowledge of the domain is limited and hence automatic verification is difficult. We aim to show that using an ontology language to capture domain knowledge and then formally mapping this domain knowledge to the CASL specification language,

can aid with automatic verification. As part of our design framework, we show how a graphical domain specific language for specification can be created for a given domain. To illustrate these concepts, we use the railway domain as a first example.

Karim Kanso, Swansea University
A light-weight approach to integrate automated and interactive theorem proving

In this talk I will present a novel integration of automated and interactive theorem proving, particularly suited to industrial verification applications. The integration has been implemented for the dependently-typed proof assistant Agda. Briefly, this embedding is based on the use of external tools and the notion of reflection. The talk will conclude with a simple demo, where a simple control system is verified.

Michal Konečný, Aston University
Numerical proofs using function interval arithmetic

Proving theorems that involve non-linear algebraic expressions over the real numbers is a problem that is hard to address using automated symbolic reasoning. Brute-force evaluation using interval arithmetic and adaptive splitting of the real domain provides a partial solution by which some theorems can be proved and with which counter-examples to a large class of non-theorems can be found. We show how several concepts of the AERN exact real arithmetic library improve the scalability and applicability of this approach. Namely, arithmetic of polynomial intervals reduces the need for excessive domain splitting and makes it feasible to prove theorems that involve the integration operator. Also, support for arithmetic over the lattice of generalised intervals instead of the interval domain allows approximating expressions that explicitly involve (exact real) intervals and the set inclusion relation. Such theorems arise naturally in the verification of floating point programs.

Andrew Lawrence, Swansea University
Modelling and verifying railway control systems using Scade

In this talk I will present research into the use of the Scade Suite from Esterel Technologies for the verification of railway interlockings. This is a feasibility study done in co-operation with Invensys Rail, a leading international company for the design, construction, and validation of railway control systems. We concentrate on the application of modelling and model-checking; specifically we present the development of two different modelling approaches. The first approach takes an industrial specification, translates it into Scade language and then applies model-checking to verify safety properties for the specification. The second is a new approach which attempts to model components of the railway in a

modular fashion and aims to capture the topology which is absent from the industrial approach. We then present an example demonstrating how these components can be combined to form a model of a railway. This railway example then has the pre-existing industrial safety properties applied to it and several new safety properties which speak about the topology. We conclude the talk with a comparison between the two approaches and a discussion of a recently commenced PhD project.

Ioannis Lignos, Durham University, UK
Reconfiguration of Hamiltonian cycles

For any instance of a combinatorial problem, the reconfiguration graph has as vertices the set of feasible solutions, pairs of which are joined by an edge if their difference is minimal. For example, for a graph G, the reconfiguration graphs for the k-colouring problem has as vertices all possible k-colourings of G and pairs of colourings are joined by an edge if the colourings differ on only a single vertex. There has been much interest recently in the study of reconfiguration graphs, particularly in the computational complexity of deciding whether reconfiguration graphs are connected, or whether a pair of given solutions belong to the component. We will review a number of recent results focusing particularly on the Hamiltonian cycle problem. Two Hamilton cycles are adjacent in the reconfiguration graph if they differ only in that a pair of adjacent vertices are "switched". We ask given a graph G and two Hamiltonian cycles C_1 and C_2, is there a path between C_1 and C_2 in the reconfiguration graph; that is, is it possible to transform C_1 into C_2 via a sequence of Hamiltonian cycles using the switching operation? We are interested in the complexity of this decision problem, and present a number of results on graphs of bounded degree.

Loretta Mancini, University of Birmingham
Analysing some 3G mobile protocols

The 3G communication system was introduced in 1999 with the first release of UMTS. It aims at overcoming GSM limitations and weaknesses, while maintaining full interoperability between the two systems. In particular, it offers a better support for mobile data applications and an improved security architecture. In this talk, we present the UMTS architecture and formally analyze some protocols of the mobility management layer. More specifically, we study the anonymity and unlinkability properties of the UMTS authentication and key agreement protocol, using the ProVerif tool to verify them. Furthermore, we show a linkability attack, which exploits the use of error messages and makes it possible to link executions of the protocol originated by the same user. In particular, this attack gives an adversary a way of detecting the presence of a UMTS user in an area and in general of tracing UMTS users' movements across different areas.

David Manlove, University of Glasgow
"Almost stable" matchings in the Roommates problem

An instance of the classical Stable Roommates problem need not admit a stable matching. This motivates the problem of finding a matching that is "as stable as possible", i.e., admits the fewest number of blocking pairs. It is known that this problem is NP-hard and not approximable within $n^{\frac{1}{2}-\varepsilon}$, for any $\varepsilon > 0$, unless P=NP, where n is the number of agents in a given SR instance. We extend the study to the Stable Roommates problem with Incomplete lists. In particular, we consider the case that the lengths of the lists are bounded by some integer d. We show that there is some $\delta > 1$ such that the problem of finding a matching with the minimum number of blocking pairs is not approximable within δ, even if $d = 3$. On the other hand we show that the problem is solvable in polynomial time for $d \le 2$, and we give a $(2d-3)$-approximation algorithm for fixed $d \ge 3$. If the given lists satisfy an additional condition (namely the absence of a so-called elitist odd party), the performance guarantee improves to $2d - 4$. This is joint work with Péter Biró, Hungarian Academy of Sciences and Eric McDermid, University of Wisconsin-Milwaukee.

Antony McCabe, University of Liverpool
Calculating the NTUmin value for set-systems auctions

When designing methods of running auctions for making combinations of purchases, we often wish to examine the performance of these "mechanisms" by looking at how much they might have to pay. In order to do so meaningfully, we would like some benchmark figure to compare against; something that would seem to represent a reasonable "fair" price for the auction. One method that has been adopted is calculating the minimum solution of a particular system of inequalities that represent "fairness". (This is called NTUmin or nu in the literature). One of the problems is that this value has been shown to be NP-complete to calculate. In this talk, we take the definition of this value and see a more intuitive way of looking at the structures it creates, before finally showing that it is also NP-hard to approximate.

Fredrik Nordvall Forsberg, Swansea University
A categorical semantics for inductive-inductive definitions

There are two equivalent well-established approaches to modelling the semantics of datatypes: in type theory, each set A comes equipped with an eliminator which at the same time represents reasoning by induction over A and the definition of recursive functions out of A. A more categorical approach models datatypes as initial T-algebras for a suitable endofunctor T. We now extend the categorical semantics and the equivalence to inductive-inductive definitions, where a set A

is defined together with an *A*-indexed family $B : A \rightarrow$ Set. Both *A* and *B* are inductively-defined in such a way that the constructors for *A* can refer to *B* and vice versa. In addition, the constructors for *B* can refer to the constructors for *A*. This complicates matters and makes us replace the usual algebras with dialgebras.

Ian Pratt-Hartmann, University of Manchester
Topological logics of Euclidean spaces

The field of Artificial Intelligence known as Qualitative Spatial Reasoning is concerned with the problem of representing and manipulating spatial information about everyday objects. In recent decades, much activity in this field has centred on spatial logics: formal languages whose variables range over regions of space, and whose non-logical primitives represent geometrical relations and operations involving those regions. The central problem is to determine whether the configuration described by a given formula is geometrically realizable in 2D or 3D Euclidean space. When the geometrical relations and operations are all topological in character, we speak of a topological logic. Topological logics have been intensively studied in Artificial Intelligence over the last two decades. The best-known of these, RCC8 and RCC5, employ variables ranging over regular closed sets, and a collection of eight (respectively, five) binary predicates standing for some basic topological relations between these sets. An important extension of RCC8, known as BRCC8, additionally features functions denoting certain operations on regular closed sets, such as complementation, agglomeration and taking common parts. None of these languages, however, is able to express the property of connectedness, which is a serious limitation in practical contexts. In this talk we present new results on topological logics in which this limitation does not apply. In particular, we show that, for any logic featuring the BRCC8-operations and a predicate representing the property of being connected, the realizability problem over the Euclidean plane is undecidable.

Stephan Reiff-Marganiec, University of Leicester
Modelling virtual organisations: structure and reconfigurations

Organisations have to adapt rapidly to survive in today's diverse and rapidly-changing environments. The idea of virtual organisations (VOs) emerged as an answer. There is a strong need to understand VOs in a formal way: changes can have side effects and hence one might wish to understand precisely what consequences a change might have. The Virtual Organisation Modelling Language (VOML) consists of sub-languages to model different aspects of VOs such as their structure or operational models: VO-S deals with structural aspects while VO-R addresses reconfigurations. The concepts are exemplified through a travel booking VO that needs to cope with extra demands imposed by the upcoming Olympic games.

Gurchetan Singh, University Of Birmingham
Improving verifiability in electronic voting

Electronic voting is being trialed in many countries, but existing systems have poor properties of security, privacy and verifiability. Many voting systems provide the option of verifiability but the evidence is not intuitive and needs machines to be checked, as it relies on cryptography. In this talk we will discuss how verifiability can be improved by introducing trial votes without compromising incoercibility and its implementation on remote voting system JCJ/Civitas.

Dirk Sudholt, University of Birmingham
Analysis of an iterated local search algorithm for vertex colouring

Hybridizations of evolutionary algorithms and local search are among the best-performing algorithms for vertex colouring. However, the theoretical knowledge about these algorithms is very limited and it is agreed that a solid theoretical foundation is needed. We consider an iterated local search algorithm that iteratively tries to improve a colouring by applying mutation followed by local search. We investigate the capabilities and the limitations of this approach using bounds on the expected number of iterations until an optimal or near-optimal colouring is found. This is done for two different mutation operators and for different graph classes: bipartite graphs, sparse random graphs, and planar graphs.