

REPORT ON BCTCS 2018

The 34th British Colloquium for Theoretical Computer Science 26–28 March 2018, Royal Holloway, University of London

Matthew Hague

The British Colloquium for Theoretical Computer Science (BCTCS) is an annual forum in which researchers in Theoretical Computer Science can 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 to benefit from contact with established researchers.

BCTCS 2018 was hosted by Royal Holloway, University of London, and held from 26th to 28th March, 2018. The event attracted over 40 participants (approx. 25% female), and featured an interesting and wide-ranging programme of five invited talks and 27 contributed talks. Abstracts for all of the talks from BCTCS 2018 are provided below.

We are grateful to the *Heilbronn Institute for Mathematical Research* who provided funds covering 7 PhD students. We are also thankful to the *London Mathematical Society* for their annual sponsorship of the *LMS Keynote Speaker*. This year, the LMS speaker was Prof. John E. Hopcroft (Cornell University).

The opening keynote lecture was given by Prof. Marta Kwiatkowska, covering the highly urgent topic of “Safety Verification for Deep Neural Networks”. It was followed by a session of submitted talks on topics related to probabilistic verification. The afternoon consisted of two submitted sessions on matrices and machine learning respectively.

The day ended with the LMS keynote talk given by Prof. John E. Hopcroft on “Research in Deep Learning”. Hopcroft is famous for his foundational work and is the recipient of many prestigious awards, including a Turing award. The event was opened to the public as a *Royal Holloway Distinguished Lecture* and attracted over 80 attendees. The talk was followed by a drinks reception funded by the Royal Holloway Department of Computer Science.

The conference resumed on Tuesday morning with an invited lecture by Prof. Agata Ciabattoni about “Intermediate Logics: From Hypersequents to Parallel Computation”. The next submitted talks session continued the theme, covering topics in logic and verification.

In the afternoon the conference focussed on distributed systems and biological systems. The keynote was given by Dr. Thomas Sauerwald and showed how randomised algorithms can be used to, for example, implement load balancing in a network.

The social event on Tuesday was a banquet held in the artistic surroundings of the Royal Holloway picture gallery.

Wednesday was the final day and was kicked off by Prof. Alexandra Silva giving a presentation of the successful NetKAT programme, showing how Kleene Algebra can be used to program and reason about packet-switched networks. This included recent extensions to include probabilistic properties. The topic of algebra was continued in the submitted talks until lunch.

After lunch the conference closed with two submitted sessions covering matching problems, and graphs and constraint problems respectively.

BCTCS 2019 will be hosted by Durham University. Researchers and PhD students wishing to contribute talks concerning any aspect of Theoretical Computer Science are cordially invited to do so. Further details are available from the BCTCS website at www.bctcs.ac.uk.

Invited Talks at BCTCS 2018

Agata Ciabattoni (TU Wien)

Intermediate logics: from hypersequents to parallel computation

We provide a general proof-theoretic framework connecting logic and concurrent computation. We describe an algorithm for introducing analytic calculi for a large class of logics in a uniform and systematic way. The introduced calculi are used to provide a first concurrent computational interpretation for many intermediate logics, classical logic included. We use the Curry-Howard correspondence to obtain new typed concurrent λ -calculi, each of which features a specific communication mechanism and implements forms of code mobility.

John E. Hopcroft (Cornell University)

LMS Keynote Lecture in Discrete Maths

Research in Deep Learning

A major advance in AI occurred in 2012 when AlexNet won the ImageNet competition with a deep network. The success was sufficiently better than previous years that deep networks were applied in many applications with great success. However, there is little understanding of why deep learning works. This talk will illustrate current research directions in the area at a level for a general scientific audience.

Marta Kwiatkowska (University of Oxford)

Safety Verification for Deep Neural Networks

Deep neural networks have achieved impressive experimental results in image classification, but can surprisingly be unstable with respect to adversarial perturbations, that is, minimal changes to the input image that cause the network

to misclassify it. With potential applications including perception modules and end-to-end controllers for self-driving cars, this raises concerns about their safety. This lecture will describe progress with developing a novel automated verification framework for deep neural networks to ensure safety of their classification decisions with respect to image manipulations, for example scratches or changes to camera angle or lighting conditions, that should not affect the classification. The techniques work directly with the network code and, in contrast to existing methods, can offer guarantees that adversarial examples are found if they exist. We implement the techniques using Z3 and evaluate them on state-of-the-art networks, including regularised and deep learning networks. We also compare against existing techniques to search for adversarial examples.

Thomas Sauerwald (University of Cambridge)

Randomised Distributed Algorithms

Randomised Algorithms are algorithms which are able to use random bits in order to determine the next step. Such algorithms tend to be more elegant and easier to implement, and in some cases they are even provably superior to any deterministic solution. In this talk we will start by giving a brief introduction into the field of randomised algorithms. We will then study two more specific examples of randomised algorithms in the context of large distributed networks, including (i) an algorithm for load balancing based on randomised rounding and (ii) an algorithm for reaching consensus based on random sampling.

Alexandra Silva (University College London)

Probabilistic Program Equivalence for NetKAT

We tackle the problem of deciding whether two probabilistic programs are equivalent in the context of Probabilistic NetKAT, a formal language for reasoning about the behaviour of packet-switched networks. The main challenge lies in reasoning about iteration, which we address by a reduction to finite-state absorbing Markov chains. Building on this approach, we develop an effective decision procedure based on stochastic matrices. Through an extended case study with a real-world data centre network, we show how to use these techniques to automatically verify various properties of interests, including resilience in the presence of failures. (This is joint work with Steffen Smolka, David Khan, Praveen Kumar, Nate Foster, Justin Hsu, and Dexter Kozen.)

Contributed Talks at BCTCS 2018

Hoda Abbasizanjani (Swansea University)

Understanding Minimal Unsatisfiability

A wide range of computational problems can be represented as propositional sat-

isfiability problems (SATs), where SAT solvers can decide whether the problem is satisfiable or not. However, in many cases knowing that a problem or a system is unsatisfiable (inconsistent) is not enough and further information on the reasons for unsatisfiability is required. SAT problems are usually represented by clause-sets (conjunctive normal forms as set-systems). Every unsatisfiable clause-set can have many minimally unsatisfiable sub-clause-sets (unsatisfiable formulas without redundancy) which are considered as the reasons for unsatisfiability. We consider minimal unsatisfiability problems and we use the deficiency of a clause-set (the difference between the number of clauses and the number of variables) as a measure for classifying minimally unsatisfiable clause-sets (MUs). It is conjectured that for a fixed deficiency, MUs can be reduced to finitely many patterns. In this talk, we investigate the basic patterns for MUs with small deficiencies.

Hanadi Alkhudhayr (Newcastle University)

A Formal Framework for Composing Boolean Networks

Boolean networks are a widely used qualitative modelling approach which allows the abstract description of a biological system. One issue with the application of Boolean networks is the state space explosion problem which limits the applicability of the approach to large realistic systems. In this paper we investigate developing a compositional framework for Boolean networks to facilitate the construction and analysis of large scale models. The compositional approach we present is based on merging entities between Boolean networks using conjunction, and we introduce the notion of compatibility which formalises the preservation of behaviour under composition. We investigate characterising compatibility and develop a notion of trace alignment which is sufficient to ensure compatibility. The compositional framework developed is supported by a prototype tool that automates composition and analysis.

Rayan Alnamnakany (Kings College University)

Analysis of the Post-transcriptional Control of Gene Expression by MicroRNAs and RNA-binding Proteins by Using Machine Learning Methods

An organism is made of a number of components, including RNA and protein. Recent studies show that the post-transcriptional level is critically regulated by a number of transcript-factors, including miRNAs and RBPs. The expression of a gene mainly depends on the interaction between RBPs and miRNAs; RNA-binding proteins and micro-RNA are considered as two of the important factors in regulating the expression of genes. RBPs and MicroRNAs have been involved in a number of human diseases, such as cancers. Identifying RNA-protein interactions has become one of the central questions in biomedical research as understanding this mechanism may give the chance to design new medicines that cure diseases. These two factors can work competitively or cooperatively between each other,

and either indirectly or directly in order to adjust the expression of their target mRNAs. The new topic emerges of designing algorithms and software tools that account for the interplay between microRNA and RBPs in post-transcriptional regulation. At present, there is only one Bioinformatics tool, namely SimiRa that addresses the issue of microRNA RBP cooperation. While SimiRa relies on gene-expression data analysis and pathway features, the proposed research aims at the performance analysis of different machine learning methods in regard to the prediction of mRNA regulation by microRNAs and RBPs. The proposed research will focus on structural elements, i.e., features of microRNA binding sites as well as RBP binding sites in secondary RNA structures, including meta-stable RNA secondary structures.

Gary Bennett (Loughborough University)

Leader Election, a Brief Encounter

In a distributed system there are times when a single entity acts as a coordinator, controlling all of the other entities in the system during the execution of a task. The need for a coordinator arises if having one simplifies the solution or because it is intrinsic to the problem. The classic problem of choosing such a coordinator is known as Leader Election. In this talk, we present a brief historical overview of Leader Election as well as some insights from our ongoing research investigating asynchronous (and possibly self-stabilising) Leader Election algorithms.

Joshua Blinkhorn (University of Leeds)

Size, Cost and Capacity: A Semantic Technique for Hard Random QBFs

The central question in proof complexity can be stated as follows: Given a logical theory and a provable theorem, what is the size of the shortest proof? Proof complexity is intrinsically linked to recent noteworthy innovations in solving, owing to the fact that any decision procedure implicitly defines a proof system for the underlying language. Moreover, proof-size lower bounds correspond to best case scenarios for the solver's consumption of computational resources, namely time and memory. Arguably more important than the lower bounds themselves is the introduction of general techniques for showing them. In this talk, we introduce a new technique for proof-size lower bounds in various proof systems for quantified Boolean formulas (QBF), the prototypical PSPACE-complete language. Our technique is based on two semantic measures: the "cost" of a QBF, and the "capacity" of a proof. Relating these two measures in the context of strategy extraction, our central theorem provides absolute lower bounds on proof size over a family of natural QBF proof systems. We exemplify the technique by proving the hardness of a new family of QBFs representing equality. Our main application provides the first "genuine" proof-size lower bounds for random QBFs, which apply to the QBF analogues of Resolution, Cutting Planes, and Polynomial Calculus. (This is

joint work with Olaf Beyersdorff and Luke Hinde.)

Brendan Case (University of Birmingham)

Population Dynamics and Runtime Analysis of Self-Adapting, Non-Elitist Evolutionary Algorithms

A common problem when designing Evolutionary Algorithms, in which a population of solutions undergo random mutation and natural selection to solve optimisation problems, is choosing the correct global parameters such as mutation rate and population size to achieve optimal behaviour. Further, there are numerous settings where the best choice of a fixed parameter is either unknown or cannot exist. How to apply such tuning strategies, and in which contexts they are actually effective, is not well understood particularly from a theoretical perspective. Tuning an individual's mutation rate is one such strategy for which a number of different approaches have been explored. These include self-adjustment, where the mutation rate changes over time according to a global update scheme, and self-adaptation, in which the mutation rate is encoded into the gene of the individual itself. Using results from probability and runtime analysis, we explore the latter, focusing on how self-adaptation of mutation rates can be shown to solve simple variable-length problem instances faster than their fixed-mutation counterparts. We also attempt to gain better insight into whether a self-adapting population's affinity for conservation using very low mutation rates can overtake the incentive of making progress using higher mutation rates.

Frances Cooper (University of Glasgow)

A $3/2$ -approximation algorithm for the student-project allocation problem with ties

In universities all over the world, students must be assigned to dissertation projects as part of their degree programmes. In many cases students are required to rank a subset of the available projects in order of preference, and likewise, lecturers rank in order of preference those students who have applied for their projects. A centralised allocation is then conducted which gives a matching of students to projects. A key consideration is that the matching should be stable, which ensures that no student and lecturer who are not matched together have an incentive to form an assignment with one another after the matching has been announced. In this talk we consider the case where preference lists need not be strictly ordered, and may contain ties. In this scenario stable matchings can be of different sizes and it is known that the problem of finding a maximum-sized stable matching is NP-hard. We present an approximation algorithm for this problem that has a performance guarantee of $3/2$.

Charlie Dickens (Turing Institute)

Leveraging Well-Conditioned Bases: Streaming and Distributed Summaries in Minkowski p -Norms

Work on approximate linear algebra has led to efficient distributed and streaming algorithms for problems such as approximate matrix multiplication, low rank approximation, and regression, primarily for the Euclidean norm l_2 . We study other l_p norms, which are more robust for $p < 2$, and can be used to find outliers for $p > 2$. Unlike previous algorithms for such norms, we give algorithms that are (1) deterministic, (2) work simultaneously for every $p \geq 1$ (including $p = \omega$), and (3) can be implemented in both distributed and streaming environments. We study l_p -regression, entrywise l_p -low rank approximation, and versions of approximate matrix multiplication.

Simon Docherty (University College London)

Modular Tableaux Calculi for Bunched Logics and Separation Theories

In recent years the key principles behind Separation Logic have been generalized to generate formalisms for a number of verification tasks in program analysis via the formulation of “non-standard” models utilizing notions of separation distinct from heap disjointness. These models can typically be characterized by a separation theory, a collection of first-order axioms in the signature of the model’s underlying ordered monoid. While all separation theories are interpreted by models that instantiate a common mathematical structure, many are undefinable in Separation Logic itself and determine different classes of valid formulae, leading to incompleteness for existing proof systems. Generalizing systems utilized in the proof theory of bunched logics (the underlying propositional basis of Separation Logic), we propose a framework of tableaux calculi that are generically extendable by rules corresponding to coherent formulae. All separation theories in the literature – as well as axioms for a number of related formalisms appropriate for reasoning about complex systems, security, and concurrency – can be presented as coherent formulae, and so this framework yields proof systems suitable for the zoo of separation logics and a range of other logics with applications across computer science. Parametric soundness and completeness of the framework is proved by a novel representation of tableaux systems as coherent theories, suggesting a strategy for implementation and a tentative first step towards a new logical framework for non-classical logics. (This is joint work with David Pym.)

Raisa Dzhamtyrova (Royal Holloway, University of London)

Aggregating Algorithm for Prediction of Packs

We formulate a protocol for prediction of packs, which is a special case of prediction under delayed feedback. Under this protocol, the learner must make a few predictions without seeing the outcomes and then the outcomes are revealed.

We develop the theory of prediction with expert advice for packs. By extending Vovk's Aggregating Algorithm to this problem we obtain a number of algorithms with tight upper bounds. We carry out empirical experiments on housing data.

Alexander P. Jeffrey (Sussex University)

Asynchronous Sessions with Implicit Messages and Functions

Runtime safety of Lambda And Session Types (LAST) with asynchronous session types is well established. Implicit function types have recently arisen as a powerful abstraction mechanism for lambda calculus. We introduce implicit function types to LAST, and generalise the concept of an implicit function to an implicit message, a concurrent analogue of implicit functions. We argue that the resulting system provides useful abstractions for programming languages with session types. We formalise the system and prove type soundness.

Tobias Kappé (University College London)

Concurrent Kleene Algebra

Kleene Algebra (KA) is a successful framework for studying program flow; it sports a simple yet expressive syntax, as well as efficient algorithms for model checking. Recently, Concurrent Kleene Algebra (CKA) has been proposed by Hoare et al. as an extension of KA, aimed at incorporating concurrent composition. We shall review some hurdles recently taken in the process of extending the toolkit of KA to one for CKA, as well as some problems yet to tackle.

Andrew Lewis-Pye (London School of Economics)

Finding short paths in small worlds

While the problem of finding short paths in graphs has been extensively studied, the situation changes dramatically when one restricts to the case of graphs satisfying conditions like the small world property, common to most networks seen in nature. In this talk we discuss the idemetric property, which formalises the idea that most nodes in a graph have similar distances between them, and which we suggest is likely to be satisfied by many small-world network models. As evidence for this claim, we have shown, for example, that the Watts-Strogatz model is idemetric for a wide range of parameters. For graphs with the idemetric property, it is easily observed that the all-pairs shortest path problem can be reduced to the single-source shortest path problem, so long as one is prepared to accept solutions which are of stretch close to 2 with high probability. So the suggestion is that this is a common property, giving rise to very efficient short path finding algorithms in contexts where $O(n)$ preprocessing is permitted.

Andrew Lewis-Smith (Queen Mary, University of London)

Kripke semantics for sub-structural logics (basic logic and GBL logic)

We investigate intermediate logics that retain a weak form of contraction. Whilst intermediate logics are generally constructive with a well understood proof theory, the same cannot be said for logics with restricted contraction. This is partly because such systems have a rich semantic motivation, being many-valued or “fuzzy.” The result is that the majority of work in such logics focus on algebraic and semantic aspects, downplaying questions of proof. Indeed, the lack of a sufficiently worked-out proof theory is even worse in the case of so-called intermediate logics with fuzzy semantics. Generalized Basic Logic (GBL) is one such logic, extending the Basic Logic (BL) of Hajek by adding pre-linearity to the axioms. We have succeeded in extending an algebraic semantics of Urquhart to GBL (from Hajek’s BL), have proven soundness for BL under this semantics, and are currently working on the completeness result. We have identified a connection with Kripke frames in the work of Bova-Montagna which could help simplify the existent approaches to fuzzy logic as it extends Kripke frames to the case of unit-interval. In the fullness of time we intend to algebraically characterize the relation between BM-frames, Kripke frames, and Totally-ordered commutative monoids.

Ciaran McCreesh (University of Glasgow)

Subgraph Isomorphism in Practice

The subgraph isomorphism problem is to find a copy of a little “pattern” graph inside a larger “target” graph. Despite being NP-complete, practical algorithms can often solve instances with graphs involving many thousands of vertices. I will give an overview of applied subgraph isomorphism solving, and present a rough outline of the state of the art algorithm that has been developed at Glasgow over the past several years. This algorithm does nothing to improve theoretical worst-case complexity, but in practice it is many orders of magnitude better than earlier algorithms. To justify this claim, I will discuss how we evaluate subgraph isomorphism algorithms empirically, with a particular focus on what makes the problem hard in practice. For the non-induced variant of the problem on randomly generated pattern and target graphs, we see a satisfiable / unsatisfiable phase transition and a corresponding complexity peak which is similar to what we see with random boolean satisfiability instances, but with induced or labelled variants of the problem, much richer behaviour arises. Finally, I’ll explain the connection between “really hard” instances and the design of search order strategies. (This is joint work with Patrick Prosser and James Trimble.)

Grzegorz Muszynski (University of Liverpool)

Using Disjoint-Set Data Structure In The Study Of Extreme Weather Events

We address the problem of algorithm tracking changes in the topology of connected components of scalar fields on geometric objects such as grids. In particular, an algorithm will be discussed based on the Union-Find Disjoint-Set data

structure whose time complexity was bounded by $O(\log^*(n))$ by John Hopcroft in 1973. In this talk, I will also give some examples from our ongoing study of extreme weather events.

Sofiat Olaosebikan (University of Glasgow)

An Integer Programming formulation for a matching problem

Matching problems generally involve assigning a set of agents to another set of agents based on the preferences of the agents and some problem-specific constraints. This class of problems was first studied by Gale and Shapley. We present an example of a matching problem, the Student-Project Allocation problem with preferences over Projects (SPA-P), which involves a set of students, projects and lecturers. Each project is offered by one lecturer, and both projects and lecturers have capacity constraints. Each lecturer and student has preferences over acceptable projects. We seek a matching which is an allocation of students to projects that respects these preferences and capacities. We also seek to ensure that no student and lecturer could improve relative to their current assignment by forming an arrangement outside the matching involving some project, a concept termed stability. Given an instance of SPA-P, stable matchings can have different sizes. This naturally leads to the problem of finding a maximum cardinality stable matching (MAX-SPA-P), which is NP-hard. Manlove and O'Malley gave an approximation algorithm to solve MAX-SPA-P with a performance guarantee of 2. Subsequently, Iwama et al. described an improved approximation algorithm with a performance guarantee of $3/2$. We describe an Integer Programming (IP) formulation to enable MAX-SPA-P to be solved optimally. Finally, we present some experimental results arising from an empirical analysis that compares the approximation algorithms and the IP model for randomly generated SPA-P instances. (This is joint work with David Manlove, Duncan Milne.)

William Pettersson (University of Glasgow)

Stable matching problem for adoptions

The stable matching problem matches pairs of agents from two distinct sets to optimise some particular objective. In this talk, I will show how the stable matching problem can be applied to find adoptive families for children in need of loving, permanent homes. In such a problem, each family-child pair is given a score, based on certain attributes of the child and family. These can easily be converted into instances of a stable matching problem, and if these instances are small enough they can be directly solved, though larger instances may be intractable. We demonstrate two methods of adapting this conversion process to produce limited instances which still produce stable matchings with suitable sizes and total scores. This particular problem also has a second interesting aspect. The aforementioned scores are not the ultimate decider of a pairing. Some of the factors relevant to a

child and family are not quantifiable, and as such are not reflected in the scores attributed to the instance. Instead, the expectation is that a social worker has the final ability to select the correct matching for each child. With this in mind, we also discuss ways of adapting the problem to present, for each child, a limited selection of families such that the final solution still satisfies some set of expected outcomes, such as stability.

David Purser (University of Warwick)

Adapting the Kantorovich Metric to ϵ, δ -Differential Privacy

Differential privacy often relies upon manually crafted mathematical proofs. In this talk, we consider automata-based modelling to automatically verify differential privacy properties. The Kantorovich metric has been used by Desharnais et al. to analyse bisimilarity of Markov chains; by van Breugel and Worrell to analyse probabilistic automata; and by Xu, Chatzikokolakis and Lin to analyse ϵ -differential privacy. We present work in progress to further generalise the Kantorovich metric to analyse ϵ, δ -differential privacy, using a distance introduced by Barthe et al. We consider its properties, how it can be calculated and the problems to which it can be applied.

Tobias Rosenberger (Swansea University)

Verified Safe and Secure Software Services

Software services – deployed on various machines, in the cloud, Internet of Things devices, etc, and cooperating to perform their various tasks – play an increasingly important role in modern economy. In many cases, incorrect behaviour of these services would put lives, health, property or privacy at risk, be it because they handle money, handle confidential data, control vehicles, exchange records of prescribed or performed medical treatments, etc. It is obviously desirable to ensure certain safety and security properties of these software services in a way suitable to, and taking advantage of, their distinct characteristics. In this talk, we explore what the characteristic properties of services are, and sketch how these inform a tool-supported development process for formally verified software that is tailored to systems composed of services.

Mehrnoosh Sadrzadeh (Queen Mary, University of London)

Fuzzy Generalised Quantifiers for Natural Language in Compositional Distributional Semantics

Recent work on compositional distributional models of natural language shows that bialgebras over finite dimensional vector spaces can model generalised quantifiers. This technique requires one to construct the vector space over powersets, and therefore is computationally costly. In this paper, we overcome this problem by considering fuzzy versions of quantifiers along the lines of Zadeh, within the

category of many valued relations. We show that this category is a concrete instantiation of the compositional distributional model. We show that the semantics obtained in this model are equivalent to the semantics of the fuzzy quantifiers of Zadeh. As a result, we are now able to treat fuzzy quantification without requiring a powerset construction. (This is joint work with Matej Dostal, Mehrnoosh Sadrzadeh and Gijs Wijnholds.)

Radu Ștefan Mincu (University of Bucharest)

Local search algorithms for wireless mesh networks

In a multi-channel Wireless Mesh Networks (WMN), each node is able to use multiple non-overlapping frequency channels. Raniwala et al. propose and study several such architectures in which a computer can have multiple network interface cards. These architectures are modeled as the graph problem *maximum edge q -coloring* and studied by Feng et. al, by Adamazek and Popa. Larjomaa and Popa define and study an alternative variant problem, *min-max edge q -coloring*. These two graph problems are studied mainly from the theoretical perspective. In this talk, we consider *min-max edge q -coloring* from a practical perspective, and describe three heuristic approximation algorithms. These algorithms are based on local search methods: hill climbing, simulated annealing and tabu search. Algorithms for particular graph classes (e.g., trees, planar graphs, cliques, bi-cliques, hypergraphs), were proposed by Larjomaa and Popa; our algorithms are applicable to general graphs. We study and compare the running data for all three algorithms on Unit Disk Graphs, as well as graphs from the DIMACS vertex coloring benchmark dataset. (This is joint work with Alexandru Popa.)

Zak Tonks (University of Bath)

Fast Matrix Inversion in Computer Algebra

James R. Bunch and John E. Hopcroft improved upon Strassen's (1969) method for matrix inversion via fast matrix multiplication in 1974. Bunch-Hopcroft handled the case in which principal submatrices are singular, and presented a modification for providing LU factorisation via this scheme. Such fast matrix multiplication techniques recurse via 7 multiplications on submatrices instead of the naïve 8, and achieve a worst case complexity of $O(n^\omega)$ where $\omega = \log_2 7$. However, Bunch-Hopcroft's method assumes that the input matrix is over a field – in particular the recursive nature of the algorithm requires that certain elements and sub-determinants are invertible. But this is not always true of a ring, and in doing Computer Algebra we are most interested in rings of polynomials. In this talk, we present a fraction-free formulation of the algorithm suited to (dense) matrices of sparse polynomials, where the intention is that such a method should be more efficient than interpolation methods. In such a way, it is attempted to provide for these matrix inversion methods what Bareiss-Dodgson did for Gaussian Elimination.

Mariia Vasileva (Newcastle University)

An evaluation of estimation techniques for probabilistic reachability

We evaluate numerically-precise Monte Carlo (MC), Quasi-Monte Carlo (QMC) and Randomized Monte Carlo (RQMC) methods for computing probabilistic reachability in hybrid systems with random parameters. Computing reachability probability amounts to computing (multi-dimensional) integrals. In particular, we pay attention to the QMC and RQMC methods due to their theoretical benefits in convergence speed with respect to the MC method. The Koksma-Hlawka inequality is a standard result that bounds the approximation of an integral by QMC techniques. However, it is not useful in practice because it depends on the variation of the integrand function, which is in general difficult to compute. The question arises whether it is possible to apply statistical or empirical methods for estimating the approximation error. In this paper we compare a number of interval estimation techniques: Central Limit Theorem (CLT), Wilson, Agresti-Coul, Logit, Anscombe, Arcsine, Bayesian and Quint methods. We also introduce a new approach based on the CLT for computing confidence intervals for probabilities near the borders of the $[0,1]$ interval. Based on our analysis, we provide justification for the use of the developed approach and suggest usage guidelines for probability estimation techniques. (This is joint work with Paolo Zuliani.)

Gleifer Vaz Alvez (Federal University of Technology, Parana, Brazil)

Formal Specification of Autonomous Systems Properties by Using a Temporal-Deontic based Logic

Autonomous Systems (AS) are largely applied in critical and reliable systems, like Autonomous Vehicles (AV), Aircrafts, Robots, and so on, and one should thus be concerned with formally verifying that the AS is working as designed. Model Checking can be used, though the properties and behaviour of an AS are constantly changing. An AS can be represented by means of an intelligent agent (or multiagents), where we describe its behaviour of through agent beliefs, plans and goals (within the so-called BDI agent architecture). Formally verify our intelligent agent can then be done with the MCAPL (Model Checking Agent Programming Language) framework, in which agent code is written in an agent programming language, for instance, Gwendolen, and then formally verified with AJPF (Agent Java Path Finder) where formal specifications are written in temporal logic. Previously, we established the SAE (Simulate Automotive Environment) to formally verify a simple intelligent agent controlling the basic behaviour of an autonomous car, and wrote temporal logic specifications with PSL (Property Specification Language). However, in order to properly specify the behaviour of an AS, we extend PSL language with deontic logic operators; this makes it possible to formally specify scenarios where the behaviour of our intelligent agent can be allowed or

forbidden. For instance, an autonomous car may be forbidden to cross a red signal light unless the agent believes there is an emergency situation making it necessary to change the agent behaviour from forbidden to allowed.

Thomas Wright (University of Edinburgh)

The bond-calculus for biochemical modelling

In this talk, we present the bond-calculus, a process algebra for modelling biological and chemical systems featuring nonlinear dynamics, multiway interactions, and dynamic bonding of agents. Mathematical models based on differential equations have been instrumental in modelling and understanding the dynamics of biological systems. Quantitative process algebras aim to build higher level descriptions of biological systems, capturing the agents and interactions underlying their behaviour, and can be compiled down to a range of lower level mathematical models. The bond-calculus builds on Kwiatkowski, Banks, and Stark's continuous pi-calculus by adding a flexible multiway communication operation based on pattern matching and general kinetic laws. The language has a compositional semantics based on vector fields and linear operators, allowing simulation and analysis through extraction of differential equations. We apply our framework to V. A. Kuznetsov's classic model of immune response to tumour growth, demonstrating the key features of our language, whilst validating the link between the dynamics of the model and our conceptual understanding of immune action and capturing the behaviour of the agents in a modular and extensible manner.

Elena Zamaraeva (University of Warwick)

On the specification number of threshold functions

We say that a set S of points completely specifies a function f within some class of Boolean functions if for any other function g from this class there exists a point x in S such that $f(x)$ and $g(x)$ differ. The minimum number of points in S is the specification number of f . Interest in the specification number of f comes from its connection to the complexity of learning with membership queries: the specification number of f is a lower bound of the complexity of learning with membership queries for f . Hu (1965) showed that the specification number of a threshold function of n variables is at least $n+1$; Anthony et al. (1995) showed that this bound is attained for nested functions, and conjectured that for all other threshold functions the specification number is strictly greater. We resolve this conjecture negatively by exhibiting threshold Boolean functions of n variables which are non-nested and for which the specification number is $n + 1$. Our next goal is to characterize all threshold functions with the minimum specification number. We will present some progress towards this goal.