

REPORT ON BCTCS & ALGOUK 2019

**The 35th British Colloquium for Theoretical Computer Science,
colocated with the 3rd AlgoUK workshop**

15–17 April 2019, Durham University

Matthew Johnson, Barnaby Martin, George B. Mertzios, Daniël Paulusma

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 a welcoming environment for PhD students to gain experience in presenting their work to a broader audience, and to benefit from contact with established researchers.

AlgoUK is a research network for Algorithms and Complexity in the UK, which is funded by an EPSRC network grant. The key aim of the AlgoUK network is to facilitate multi-faceted interactions within the UK's algorithmic research community and also to promote inter-disciplinary cooperation between these researchers and those in industry and other subjects such as Mathematics, Biology, and Chemistry. The founding research groups are at Durham, Kings College London, Leicester, Liverpool, Royal Holloway, and Warwick, but the network is open to all who wish to engage with it.

The joint event BCTCS & AlgoUK 2019 was hosted by Durham University, and held between 15 – 17 April 2019. The event was very successful, as it attracted 83 participants (approx. 25% female) and featured an interesting and wide-ranging programme of 10 invited talks and 22 contributed talks. Within AlgoUK there were four invited talks on Non-Standard Computing and four invited talks on Algorithmics. Within BCTCS there were two invited talks, one on graph theory and perfect graphs and one on logic. Abstracts for all of the talks are provided below.

We are grateful to the *Heilbronn Institute for Mathematical Research* who provided funds covering 7 PhD students (3 female and 4 male, from 7 different universities). We are also thankful to the *London Mathematical Society* for their annual sponsorship of the *LMS Keynote Speaker*. This year, the LMS speaker was Professor Maria Chudnovsky (Princeton University).

The joint event started on Monday afternoon, 15 April, with the four invited talks on Non-Standard Computing. The opening invited lecture was given by Professor Andrew Adamatzky, giving an overview of liquid based computing. The session also included talks from Professor Susan Stepney on unconventional computing, and from Dr Vivien Kendon and Professor Richard Jozsa on aspects of quantum computing. The evening finished in a celebratory atmosphere with

a welcoming drinks reception at the Common Room of the Department of Computer Science at Durham. The reception was kindly funded by the Department of Computer Science.

The conference resumed on Tuesday morning, 16 April, with the four invited talks on Algorithmics. Professor James Worrell told us about computing polynomial invariants, Dr Bahar Rastegari introduced a new variant of the stable marriage problem, Dr Kitty Meeks discussed interesting connections between decision and approximate counting problems, and Professor Bernhard von Stengel spoke about algorithms for bimatrix games.

The joint event continued after lunch with the LMS invited speaker Professor Maria Chudnovsky who presented older and recent work associated with the recognition of perfect graphs. On Tuesday afternoon there were two parallel tracks of contributed talks. The evening closed with a convivial dinner at the Durham University College, located inside the famous and imposing Castle of Durham.

Wednesday was the final day of the joint event. The day started with two parallel sessions of contributed talks, followed by the invited talk of Dr Ulrich Berger, on intuitionistic proofs and program synthesis, and closed with two final parallel sessions of contributed talks.

BCTCS 2019 will be hosted by Swansea 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.

BCTCS Invited Talks

Ulrich Berger

IFP - A Logic for Program Extraction

That proofs correspond to computation is an old idea of logic going back to Kolmogorov, Brouwer, Boole and many others. In computer science this is known as the Curry-Howard correspondence which identifies the typed lambda-calculus both as a notation system for proofs and a programming language.

A related - and regarding specification more expressive - correspondence between proofs and programs is provided by Kleene's realizability which can be viewed as a computationally enriched semantics of logic and arithmetic. Kleene proved the fundamental Soundness Theorem which shows that from a constructive proof of a formula one can extract an algorithm that realizes the proven formula, that is, solves the computational problem it expresses.

Realizability offers an attractive alternative to program verification: Instead of first writing a program and then trying to prove that it meets its specification, one first proves that the specification is satisfiable and then extracts a program

(provably) satisfying it.

The advantages of program extraction over verification are obvious: While it is in general undecidable whether or not a program meets its specification, the correctness of a proof can be easily checked. Hence program extraction not only yields correct programs but also provides formal correctness proofs for free.

Despite the virtues of program extraction, currently dominating technologies for producing formally certified software are based on traditional 'programming-first' methodologies such as model-checking and formal verification. The main reasons are that program extraction is considered too restricted, since it covers only functional programming but misses out on many other computational paradigms, and too complicated and demanding for software developers to be of practical relevance. This talk will present a logical system IFP (Intuitionistic Fixed point Logic) as a basis for program extraction that addresses these issues.

After a brief introduction to IFP an overview of the main case studies - most of them implemented in the Minlog proof system - is given. This is followed by an outline of new developments aiming to extend program extraction to capture non-functional computational paradigms such as concurrency and state based computation.

Maria Chudnovsky (LMS Keynote Lecturer in Discrete Mathematics)

Detecting odd holes

A hole in a graph is an induced cycle of length at least four; and a hole is odd if it has an odd number of vertices. In 2003 a polynomial-time algorithm was found to test whether a graph or its complement contains an odd hole, thus providing a polynomial-time algorithm to test if a graph is perfect. However, the complexity of testing for odd holes (without accepting the complement outcome) remained unknown. This question was made even more tantalizing by a theorem of D. Bienstock that states that testing for the existence of an odd hole through a given vertex is NP-complete. Recently, in joint work with Alex Scott, Paul Seymour and Sophie Spirkl, we were able to design a polynomial time algorithm to test for odd holes.

AlgoUK Invited Talks on Non-Standard Computing

Andrew Adamatzky

Liquid computers

We show that a liquid is a universal computing substrate. A liquid can be used to represent signals, actuate mechanical computing devices and to modify signals via chemical reactions. A liquid per se can host thousands of chemical micro-processors. We give an overview of liquid based computing devices developed over hundreds of years. In these devices liquid is used in various roles. Mass

transfer analogies are employed in hydraulic mathematical machines and integrators. Flows of fluid explore and map templates, and solve mazes, in fluid mappers. Interacting fluid jets realise logical gates in fluidic logic devices, where Boolean values are represented by a flow pressure in the channels. The liquid can be discretised in droplets and liquid marbles and computation is implemented via marbles/droplets colliding with each other or actuating mechanical devices. We give examples of liquid reaction-diffusion chemical computers, where oxidation wave fronts or diffusing fronts of precipitation solve logical and geometrical problems. We will also present advances in liquid electronics, liquid switches and transistors. If time allows we will talk about colloidal computers, including sol, emulsion and foam.

Richard Jozsa

Clifford operations and the verification of quantum computations

Quantum computation has a remarkably rich diversity of models, each offering different perspectives on its relation to classical computing, and also on implementational challenges. We will introduce and discuss the model of quantum computing based on Clifford gates and so-called magic states.

In this model, efficiently classically simulatable quantum processes are elevated to full universal quantum computing power by inclusion of an additional resource that can be argued to be entirely classical. We will also outline how this feature suggests a novel and simple approach to the issue of verifying the correctness of a quantum computer that is operating in the high complexity regime, beyond the remit of any foreseeable classical computing power.

Vivian Kendon

Continuous-time quantum computing

We present a generalising, unifying view of quantum computing with continuous-time quantum walks, adiabatic quantum computing, quantum annealing, and a range of special purpose quantum simulators. The key feature they all share is evolving the initial quantum state to the final quantum state using a continuous-time process. Looking more closely, the similarities go further, with the same Hamiltonians used to encode the problems across adiabatic quantum computing, quantum annealing, and quantum walks for implementing the same algorithm. The differences are in the time-dependence of the Hamiltonian, and how long it is applied for. This invites the question of which is best, especially in a practical setting where resources such as quantum coherence are severely limited. The shorter run times of quantum walks and quantum annealing may be favoured over slower adiabatic strategies, and hybrid algorithms exploiting the strengths of each strategy allow the best performance to be extracted from real hardware [Morley et al., Arxiv 2017].

This unified view also has implications for the design of quantum computing hardware, which can potentially be used to implement any or all types of continuous time evolution, as required for the particular problem to be solved. We therefore identify *continuous-time quantum computing* (CTQC) as a unifying model of quantum computation with a continuous-time evolution applied to discrete data. It consists of a register of qudits – with qubits as a special case – to which Hamiltonian time evolution is applied that can vary continuously with time, along with coupling to an engineered environment that contributes decoherence, energy exchange and other open quantum system effects. As with the circuit model, the output can be the quantum state in the register at the end of the process, or the classical outcomes of measurements in the computational basis on the quantum register. This contrasts with the circuit model, in which discrete unitary quantum gates are applied to the register, with error correction to ensure the computation remains coherent throughout.

Continuous-time quantum computing is complementary to the circuit model, and it provides a computational paradigm for quantum computing in which techniques can be adopted across what were previously seen as separate approaches. In particular, it suggests that hardware platforms suitable for one should be suitable for all. This provides a focus for hardware design and development of continuous-time quantum computers, alongside digital quantum computers.

Susan Stepney

Unconventional design of unconventional computers

There are many proposed unconventional computational models: reservoir computing, general purpose analogue computing, membrane computing, quantum computing, morphogenetic computing, and more. There are just as many proposed unconventional computing substrates: slime moulds, carbon nanotubes, gene engineered bacteria, gold nanoparticle networks, memristors, optical relays, and more. But how to match model and substrate to get an effective unconventional computer?

We have developed CHARC, a framework to characterise how well some material substrate can instantiate and implement a given computational model. We have applied this to a reservoir computing model instantiated in carbon nanotubes, and also in simulations. This works by comparing candidate instantiations with a baseline system, and using novelty search over computational behaviour space to quantify the achievable scope of the instantiation.

I will describe this framework, and how we are developing it further: to characterise how faithfully some abstract model represents the computing properties of a given substrate; and to co-design unconventional models and substrates that are natural computational partners.

AlgoUK Invited Talks on Algorithmics

Kitty Meeks

From decision to (approximate) counting

Decision problems – those in which the goal is to return an answer of “YES” or “NO” – are at the heart of the theory of computational complexity, and remain the most studied problems in the theoretical algorithms community. However, in many real-world applications it is not enough just to decide whether a the problem under consideration admits a solution: we often want to find all solutions, or at least count (either exactly or approximately) their total number. It is clear that finding or counting all solutions is at least as computationally difficult as deciding whether there exists a single, and indeed in many cases it is strictly harder (assuming $P \neq NP$) even to count approximately the number of solutions than it is to decide whether there exists at least one.

In this talk I will discuss a restricted family of problems, in which we are interested in solutions of a specified size: for example, solutions could be copies of a given k -vertex graph H in a large host graph G , or more generally k -vertex subgraphs of G that have some specified property (e.g. k -vertex subgraphs that are connected). In this setting, although *exact* counting is strictly harder than decision (assuming standard assumptions in parameterised complexity), the methods typically used to separate *approximate* counting from decision break down. Indeed, I will demonstrate a method that, subject to certain additional assumptions, allows us to transform an efficient decision algorithm for a problem of this form into an approximate counting algorithm, with only a small increase in running time.

Bahar Rastegari

Stable marriage with groups of similar agents

Matching problems occur in various applications and scenarios such as the assignment of children to schools and junior doctors to hospitals. In such problems, it is often understood the participants (which we will refer to as agents) have ordinal preferences over a subsets of agents. It is also widely accepted that a “good” solution must be *stable*. Intuitively speaking, a stable solution guarantees that no subset of agents find it in their best interest to leave the prescribed solution and seek an assignment amongst themselves.

Many important stable matching problems are known to be NP-hard, even when strong restrictions are placed on the input. In this talk, I’ll identify structural properties of instances of stable matching problems which will allow us to design efficient algorithms using elementary techniques. I focus on the setting in which all agents involved in some matching problem can be partitioned into k different *types*, where the type of an agent determines his or her preferences, and agents have preferences over types (which may be refined by more detailed preferences

within a single type). I also consider a generalisation in which each agent may consider some small collection of other agents to be exceptional, and rank these in a way that is not consistent with their types.

I will show that (for the case without exceptions), the well-known NP-hard matching problem MAX SMTI (that of finding the maximum cardinality stable matching in an instance of stable marriage with ties and incomplete lists) belongs to the parameterised complexity class FPT when parameterised by the number of different types of agents needed to describe the instance. I will demonstrate how this tractability result can be extended to the setting in which each agent promotes at most one “exceptional” candidate to the top of his/her list (when preferences within types are not refined), but that the problem remains NP-hard if preference lists can contain two or more exceptions and the exceptional candidates can be placed anywhere in the preference lists.

Bernhard von Stengel

Fast Algorithms for Rank-1 Bimatrix Games

The rank of a bimatrix game is the matrix rank of the sum of the two payoff matrices. For a game of rank k , the set of its Nash equilibria is the intersection of a one-dimensional set of equilibria of parameterized games of rank $k - 1$ with a hyperplane. For a game of rank 1, this is a monotonic path of solutions to a parameterized zero-sum game and hence LP. One intersection of this path with the hyperplane can be found in polynomial time by binary search. This gives a polynomial-time algorithm to find an equilibrium of a game of rank 1. On the other hand, there are rank-1 games with exponentially many equilibria, which answers an open problem by Kannan and Theobald [SODA 2017; Economic Theory 2010]. Hence, the exponential multiplicity of equilibria does not prevent finding one equilibrium in polynomial time. In this spirit, we conjecture that the *uniqueness* of an equilibrium of a rank-1 game is co-NP-hard. For games of higher rank, the parameterized equilibria of lower rank correspond to the computation of the global Newton method by Govindan and Wilson [Journal of Economic Theory 2003], which for two-player games is a special case of Lemke’s algorithm.

James Worrell

Synthesising polynomial program invariants

Automatically generating invariants is a fundamental challenge in program analysis and verification. In this talk we give a select overview of previous work on this problem, starting with Michael Karr’s 1976 algorithm for computing affine invariants in affine programs, i.e., programs having only non-deterministic (as opposed to conditional) branching and all of whose assignments are given by affine expressions. We then present the main result of the talk—an algorithm to compute all polynomial invariants that hold at each location of a given affine program. Our

main tool is an algebraic result of independent interest: given a finite set of rational square matrices of the same dimension, we show how to compute the Zariski closure of the semigroup that they generate.

Contributed BCTCS & AlgoUK Talks

Hanin Abdulrahman

A Formal Framework for the Decomposition and Analysis of Boolean Networks

Qualitative modelling techniques are a key tool in modelling and analysing biological systems since they avoid the need for problematic quantitative data about reaction rates. They are seen as an important approach to enabling development methods and tools for analysing and synthesising large biological systems. *Boolean Networks* (BNs) are a well-known qualitative modelling approach that has been successfully used to model and analyse complex biological networks such as *gene regulatory networks* and *signal transduction networks*. However, modelling and analysing large-scale regulatory systems using BNs is challenging due to the well-known *state space explosion problem*.

In this PhD research, we are investigating the development of a formal framework for decomposing BNs to allow the analysis of large-scale models. The aim is to develop a theoretical framework and practical techniques with tool support to facilitate the automatic decomposition and analysis of large BNs. The key idea is to decompose a BN into sub-models by splitting and duplicating entities. This work builds on the existing compositional framework (developed at Newcastle University) that allows entities to be composed using Boolean operators. This presentation will cover our initial work in:

- Decomposing an interaction graph
- Decomposition quality metrics
- Boolean function decomposition

Duncan Adamson

Crystal Structure Prediction by Vertex Removal in Euclidean Space

Crystal structure prediction remains a major problem in computational chemistry, creating an essential barrier in the development of new materials. The problem is to find the lowest-energy crystal structures which are expected by chemists as potentially "good" ones.

The approach to crystal structure prediction that we are focusing on is to take a unit cell that has been densely packed with ions, then to remove ions from it. Our goal in the approach is to minimise the total energy with the constraint that the

final charge is neutral, leaving us with a structure that may be close to a realistic structure. The motivation for this approach from a computational perspective is twofold, first it is comprised of a simple base operation - that being the removal of an ion - which we may reason about. Secondly we can combine this approach with the opposite operation of inserting ions to reach any potential structure, allowing us to reason about the optimal arrangement for the crystal.

In this talk we will discuss this problem through the lens of graph theory, representing the structure as a set of vertices with a positive or negative charge embedded into 3-dimensional Euclidean space forming a complete graph with edges having weights corresponding to the energy between the ions in the structure. We will present a formal definition of our problem based on this input, as well as a class of functions which we may use to define the energy between ions. Finally we will conclude by showing that this problem is *NP-Complete* by a reduction from the *Clique Problem* for both the general case, and when we restrict it to the *Buckingham-Coulomb Potential* function used in chemistry.

Hanadi Alkudhayr

Boolean Networks Composition: Multiple Entities and Multiple Models

In order to study and analyse complex biological systems, a range of *qualitative modelling techniques* have emerged such as *Boolean Networks*, *Multi-Valued Networks* and *Petri Nets*. Boolean networks are a widely used qualitative modelling approach which can represent the state of each biological component (e. g. genes, protein and other chemical signals) abstractly as a Boolean value, where 1 represent the entity is active (e. g. gene is expressed or a protein is present) and 0 that it is inactive (e. g. gene is inhibited). A key challenge in analysing Boolean networks is the well-known state space explosion problem which constrains the applicability of the approach to large realistic systems.

We consider developing compositional techniques to address this limitation by extending the basic compositional theory developed previously by the authors. The compositional approach we present is based on composing Boolean networks by merging entities using a binary Boolean operator (there are 16 such binary Boolean operators to choose), and we formalise the behaviour preservation under composition using a notion of *compatibility*. We consider what it means to preserve the underlying behaviour of Boolean networks in a composed model when the composition involves multiple entities and multiple component Boolean networks.

It turns out that the compatibility property is problematic to check as it references the full behaviour of the composed model, and so we consider characterising the preservation of component Boolean network behaviour when the composition involves multiple entities and multiple component Boolean networks. We extend the definition of an *interference* graph to encompass the 16 Binary Boolean opera-

tors. We formalise this interference using a labelled state graph approach and then use this to define an extended alignment property called *weak alignment*. We show formally that weak alignment completely characterises compatibility when the underlying Boolean operator being used is idempotent. Therefore, weak alignment property provides an important scalable means of checking behaviour preservation in composed components without the need to reference the composed model. The compositional framework presented here is supported by a prototype tool that automates the composition process and associated behaviour preservation analysis.

Benjamin Merlin Bumpus

The Width of Minimum Cost Tree Decomposition

Tree decompositions have been very successfully employed in the design of parameterized graph algorithms. Typically the running time of such algorithms depends on the width of the decomposition provided: i.e the size of its largest bag. For this reason much effort has been directed towards finding tree decompositions with minimum width. However, this is not the right approach: the width of a tree decomposition which minimizes the the running time of some algorithm is not always minimum. This talk will will address progress related to the question:

“is the width of an ‘algorithmically best’ tree decomposition bounded with respect to treewidth?”

Nicholas Chancellor

Trading off Optimality and Robustness using Quantum Annealing

Quantum annealing is an analog of simulated annealing but where fluctuations are quantum as well as thermal, the most well know examples of quantum annealing devices are those produced by D-Wave Systems Inc, which have roughly 2000 qubits. While traditionally, the measure of success of a quantum annealing protocol is the ability to find optimal solutions, I argue that one could also consider the ability to find solutions which are both optimal and robust, in the sense that a small change to a solution is likely to still lead to a good solution. It is well known that quantum annealing naturally tends to preferentially find solutions which are nearby (in the Hamming distance sense) other solutions which are similarly optimal, while this is traditionally viewed as a weakness in that it can lead to suboptimal solutions, I demonstrate that it also can be used as a strength.

I will show experimental results which demonstrate how advanced features of the D-Wave device can be used to convert an optimal solution which is not very robust to a more robust solution in a controllable way for both binary and mixed binary/integer synthetic optimization problems. These results suggest a new use of quantum annealers as a tool for trading off between optimality and robustness

in hybrid quantum/classical algorithms. I will also discuss the advanced features used in this protocol, in particular the broader context of how the reverse annealing feature may be used in hybrid quantum/classical algorithms and give some experimental results in this direction.

Rajesh Chitnis

Towards a Theory of Parameterized Streaming Algorithms

Parameterized complexity attempts to give a more fine-grained analysis of the complexity of problems: instead of measuring the running time as a function of only the input size, we analyze the running time with respect to additional parameters. This approach has proven to be highly successful in delineating our understanding of NP-hard problems. Given this success with the TIME resource, it seems but natural to use this approach for dealing with the SPACE resource. First attempts in this direction have considered a few individual problems, with some success: Fafianie and Kratsch [MFCS'14] and Chitnis et al. [SODA'15] introduced the notions of streaming kernels and parameterized streaming algorithms respectively. For example, the latter shows how to refine the $\Omega(n^2)$ bit lower bound for finding a minimum Vertex Cover (VC) in the streaming setting by designing an algorithm for the parameterized k -VC problem which uses $O(k^2 \log n)$ bits.

We initiate a systematic study of graph problems from the paradigm of parameterized streaming algorithms. We first define a natural hierarchy of space complexity classes of FPS, SubPS, SemiPS, SupPS and BrutePS, and then obtain tight classifications for several well-studied graph problems such as Longest Path, Feedback Vertex Set, Dominating Set, Girth, Clique, Treewidth, etc. into this hierarchy. On the algorithmic side, our parameterized streaming algorithms use techniques from the FPT world such as bidimensionality, iterative compression and bounded-depth search trees. On the hardness side, we obtain lower bounds for the parameterized streaming complexity of various problems via novel reductions from problems in communication complexity. We also show how (conditional) lower bounds for kernels and W-hard problems translate to lower bounds for parameterized streaming algorithms.

Parameterized algorithms and streaming algorithms are approaches to cope with TIME and SPACE intractability respectively. It is our hope that this work on parameterized streaming algorithms leads to two-way flow of ideas between these two previously separated areas of theoretical computer science.

Christian Coester

The online k -taxi problem

In the k -taxi problem, a generalization of the k -server problem, an algorithm controls k taxis to serve a sequence of requests in a metric space. A request consists of two points s and t , representing a passenger that wants to be carried by a taxi from

s to t . For each request and without knowledge of future requests, the algorithm has to select a taxi to transport the passenger. The goal is to minimize the total distance traveled by all taxis. The problem comes in two flavors, called the easy and the hard k -taxi problem: In the easy k -taxi problem, the cost is defined as the total distance traveled by the taxis; in the hard k -taxi problem, the cost is only the distance of empty runs.

The easy k -taxi problem is exactly equivalent to the k -server problem. The talk will focus mostly on the hard version, which is substantially more difficult. For hierarchically separated trees, I will present a memoryless randomized algorithm with optimal competitive ratio $2^k - 1$. This implies an $O(2^k \log n)$ -competitive algorithm for arbitrary n -point metrics, which is the first competitive algorithm for the hard k -taxi problem for general finite metric spaces and general k . I will also describe the main ideas of an algorithm based on growing, shrinking and shifting regions which achieves a constant competitive ratio for three taxis on the line (abstracting the scheduling of three elevators).

Frances Cooper

Rank-maximal stable matchings in the Stable Marriage problem

In 1962, Gale and Shapley described the Stable Marriage problem (SM) in their seminal paper "College Admissions and the Stability of Marriage". In this setting we have two sets of agents: men and women, where each man has a strictly-ordered preference list over a subset of women and vice versa. We seek a stable matching, that is, an assignment of men to women such that no man-woman pair (who are not currently assigned each other) would prefer to be assigned to each other than their current partners (if any).

Each instance of SM has at least one stable matching, and may indeed have many. Although all stable matchings for an instance have the same size, they may differ greatly in terms of the happiness of individual agents. This motivates us to find a stable matching with additional properties. One such example is a rank-maximal stable matching, which is a stable matching such that the greatest number of men and women gain their first-choice partner, and subject to that, their second choice, and so on. A polynomial-time algorithm for computing such a matching was described by Gusfield and Irving in 1989. However, this algorithm relies upon exponential weights which, when used in practice, may result in overflows or memory issues. In this talk I describe a new polynomial-time algorithm to compute a rank-maximal stable matching using a combinatorial approach, without the need to revert to exponential weights.

Brian Courtehoue and Detlef Plump

Rule-Based Graph Programming

GP 2 is a computationally complete, non-deterministic language based on rules.

In a given host graph, one looks for a match of a specific graph, and replaces that match with another graph. The existence and uniqueness of these constructions is based on double-pushouts from category theory. It has relatively simple syntax and semantics, which helps making complexity or verification arguments, and comes with a compiler creating C code.

In general, there is a significant cost to matching a graph in the host graph. To lessen that cost, the language has been expanded to rooted graphs. Some nodes are distinguished as root nodes that can be accessed in constant time. They can be thought of as the nodes that are marked as the “current node” in illustrations of graph algorithms.

In this talk, we discuss some examples of rooted GP 2 programs. In particular, we will focus on identifying connected components, topological sorting and 2-colouring.

Matthew England

Cylindrical Algebraic Decomposition with Equational Constraints

Cylindrical Algebraic Decomposition (CAD) has long been one of the most important algorithms within Symbolic Computation, as a tool to perform quantifier elimination in first order logic over the reals. Its worst case complexity bound is doubly exponential in the number of variables which limits its use in practice, however, it has found application throughout science and engineering, with published applications as diverse as biological network analysis and economic hypothesis refutation. In particular, it is finding prominence in the Satisfiability Checking community as a subtool to check the satisfiability of formula in nonlinear real arithmetic in larger SMT-solver problem instances.

The original CAD algorithm produces decompositions according to the signs of polynomials, when what is usually required is a decomposition according to the truth of a formula containing those polynomials. One approach to achieve that coarser (but hopefully cheaper) decomposition is to reduce the polynomials identified in the CAD to reflect a logical structure which reduces the solution space dimension: the presence of Equational Constraints (ECs). This talk will summarise recent progress on the theory of CAD with ECs. In particular, we analyse the effect of ECs on the complexity of CAD and demonstrate that the double exponent in the worst case complexity bound reduces in line with the number of ECs.

Arved Friedemann, Jay Morgan

Verifiable Models to Machine Learning Tasks using SAT

Machine learning (ML) provides state-of-the-art performance in many tasks where the rules and representations required for classification cannot be described manually. However, due to the complexities in describing all potential behaviour for

an input distribution, their use in safety critical systems is limited. For ML to be applicable for such systems, the robustness against noise, reachability, and similar must be verified. On the other hand, Logical formulas, like circuit representations, are verifiable and can be inferred using synthesis techniques. We propose the usage of a Boolean Satisfiability (SAT) solver that uses unit propagation and clause learning to create a logical model that fits given observations and verifies safety criteria. In doing so, we show interesting comparisons between concepts implemented in SAT and ML that allows for further research. Our implementation is tested on a classic MNIST dataset where evaluations of training time, accuracy, and verifiability are made.

Abdul Ghani

The Sum-Of-Squares Hierarchy and Proof Complexity

The Sum-Of-Squares hierarchy – also known as the Lasserre or Positivstellensatz hierarchies – is a method of generating semidefinite relaxations of polynomial optimization problems. It has found applications in a wide range of diverse fields, from robotics and control to the theory of constraint satisfaction.

In this talk we will view the Sum-Of-Squares hierarchy as a propositional proof system - i.e., as a method of demonstrating the unsatisfiability (or the tautology-hood) of propositional statements. In particular we will be concerned with the complexity of refuting translations of unsatisfiable (in finite domains) first-order combinatorial statements. This complexity is very intimately related with the degree of the relaxations.

This introductory talk will begin with relevant background and definitions. This will be followed by some key results and applications of Sum-Of-Squares throughout theoretical computer science. We will finish by discussing its relationship to proof complexity, including some recent results and open questions.

Tyson Jones

Variational algorithms for noisy quantum computers

With large, error-corrected quantum computers still a distant dream, what fun can be had with the smaller, error-prone, near-future machines? Even noisy classical-quantum hybrids can be used to study quantum systems too large to be simulated on a classical computer. Variational algorithms, whereby a classical processor assists a quantum coprocessor to optimise a set of parameters, show promising noise resilience. While challenging to scale, they enable us to simulate quantum dynamics and even solve chemistry problems of intermediate scale. I'll introduce solving optimisation problems with variational algorithms, discuss the edge they give us in simulating dynamics, and how we've re-purposed them to study energy spectra and perform quantum circuit recompilation.

Elektra Kypridemou

Correlation of metrics in random graphs

In this work, we are interested in the study of the pair-wise correlations between several metrics of random graphs of small order, which are embedded in a two-dimensional Euclidean plane. We present a preliminary investigation of 11 measures in two well-known models of random graphs: the *Erdős-Rényi graphs* (which we assume randomly embedded in a circle $C(O, R)$ of centre $O(0, 0)$ and radius R) and *random geometric graphs* (RGG) over the same $C(O, R)$. Relevant descriptive measures include the size of the smallest dominating sets, the clique and independence numbers, the number of connected components, and measures of the graph clustering properties.

Studying such correlations analytically is very difficult. We review the existing literature on correlation analysis inspired by real-world complex networks, such as the Internet and social networks. This work, however, is motivated by applications in Psychology, where one is interested in the perceived numerosity of configurations of dots spread in a two-dimensional surface. In applications like these, correlations could also be useful to answer questions about one measure based on the properties of another, related one.

After analysing the pair-wise Spearman's rank correlation coefficients of the measures mentioned above, we find that strong relationships exist between groups of measures, both in the Erdős-Rényi random graphs and in the RGGs. Not surprisingly, however, these correlations highly depend on the main graph parameters of the two models (the p value in $G_{n,p}$ and the d value in G_d).

Sofiat Olaosebikan

Strongly Stable Matchings in the Student Project Allocation Problem with Ties

The *Student-Project Allocation problem with lecturer preferences over Students* (SPA-S) involves assigning students to projects based on student preferences over projects, lecturer preferences over students, and the maximum number of students that each project and lecturer can accommodate. This classical model assumes that preference lists are strictly ordered. In this work, we study a variant of SPA-S where ties are allowed in the preference lists of students and lecturers, which we refer to as the *Student-Project Allocation problem with lecturer preferences over Students with Ties* (SPA-ST). In this context, what we seek is a *stable matching* of students to projects. However, when the preference lists involve ties, three different stability concepts naturally arise; namely, *weak stability*, *strong stability* and *super-stability*.

In this talk, I will focus on the concept of strong stability in SPA-ST. Further, I will describe a polynomial-time algorithm to find a strongly stable matching or to report that no such matching exists, given an instance of SPA-ST. Our algorithm

runs in $O(m^2)$ time, where m is the total length of the students' preference lists.

Giacomo Paesani

On the Price of Independence for Graph Transversals

We compare the minimum value of classical graph transversals vertex cover, feedback vertex set and odd cycle transversal with the minimum value of the corresponding variants, in which we require the transversal to be an independent set. That is, let $vc(G)$, $fvs(G)$ and $oct(G)$ denote, respectively, the size of a minimum vertex cover, minimum feedback vertex set and minimum odd cycle transversal in a graph G . One can then ask, when looking for these sets in a graph, how much bigger might they be if we require that they are independent; that is, what is the *price of independence*? If G has a vertex cover, feedback vertex set or odd cycle transversal that is an independent set, then we let, respectively, $ivc(G)$, $ifvs(G)$ or $ioct(G)$ denote the minimum size of such a set.

First we investigate for which graphs H the values of $ivc(G)$, $ifvs(G)$ and $ioct(G)$ are *bounded* in terms of $vc(G)$, $fvs(G)$ and $oct(G)$, respectively, when the graph G belongs to the class of H -free graphs. We find complete classifications for vertex cover and feedback vertex set and an almost complete classification for odd cycle transversal (subject to three open cases). Second we investigate for which graphs H the values of $ivc(G)$, $ifvs(G)$ and $ioct(G)$ are *equal* to $vc(G)$, $fvs(G)$ and $oct(G)$, respectively, when the graph G belongs to the class of H -free graphs. We find complete classifications for vertex cover and almost complete classifications for feedback vertex set (subject to one open case) and odd cycle transversal (subject to three open cases).

Filippos Pantekis

Graph Layout Algorithms for Railways

The development and application of formal methods is a long standing research topic within the field of Computer Science. One particular challenge that remains is the widespread uptake of formal methods in to industrial practices. This is an area that the Swansea Railway Verification Group (SRVG) has been addressing in collaboration with Siemens Rail Automation UK.

The OnTrack tool-set developed by the SRVG explores a methodology for developing graphical domain specific languages for modelling and verification to aid in the uptake of formal methods within industry. The OnTrack tool-set illustrates the successful application of this methodology within the railway domain, successfully addressing issues surrounding faithful modelling of the domain, scalability of verification and accessibility to modelling and verification processes for practitioners within the domain.

However, there are still a number of unresolved issues, including that of importing standard railway designs from existing design tools within the railway

domain. Here, the importation of data sets has been solved; however, the efficient layout of this data within the graphical domain specific language of OnTrack has not. In this talk, we explore how graph layout algorithms, such as simulated annealing, can be adopted to provide a solution.

Pavel Veselý

A ϕ -Competitive Algorithm for Scheduling Packets with Deadlines

In the online packet scheduling problem with deadlines, the goal is to schedule transmissions of packets that arrive over time in a network switch and need to be sent across a link. Each packet has a deadline, representing its urgency, and a non-negative weight, that represents its priority. Only one packet can be transmitted in any time slot, so, if the system is overloaded, some packets will inevitably miss their deadlines and be dropped. In this scenario, the natural objective is to compute a transmission schedule that maximizes the total weight of packets which are successfully transmitted. The problem is inherently online, with the scheduling decisions made without the knowledge of future packet arrivals. The goal is to develop an algorithm with the lowest possible competitive ratio, which is the worst-case ratio between the optimum total weight of a schedule (computed by an offline algorithm) and the weight of a schedule computed by a (deterministic) online algorithm.

In this talk, we present a ϕ -competitive online algorithm (where $\phi \approx 1.618$ is the golden ratio), matching the previously established lower bound. We then outline main ideas of its analysis and discuss possible extensions of this work.

Aled Walters

Specification Based Testing in the context of Railway Signalling Systems

The railway domain is safety-critical: technical failure can lead to financial loss or the loss of human life. Within railway systems, signalling is an important safety measure. Its objectives include separation of trains to avoid train collisions, setting speed restrictions to avoid train derailment, and the coordination of train and point movement.

The European Rail Traffic Management System (ERTMS) is a modern railway control system which is based on the communication between trains and interlocking via a so-called radio block centre (RBC). In our PhD project, we are investigating how to provide quality control of RBC implementations through specification based testing. To this end, we formally model and verify the ERTMS. This yields a specification which has been proven to be correct with regards to selected safety properties. By the systematic testing of real world implementations of a RBC against our specification, we then establish that these implementations conform to our specification. In particular we aim to apply Gaudel's testing framework that allows the argument that testing is complete, i.e., testing is equivalent to perform-

ing a formal proof.

In this talk we will discuss a concrete, simple, first model of ERTMS in UP-PAAL. Furthermore, we will give an outlook for our overall project along with our aims and methodologies moving forward. The research is done in collaboration with our Industry partner Siemens Rail Automation.

Elena Zamaraeva

On the number of 2-threshold functions

We consider 2-threshold functions over a 2-dimensional integer grid of a fixed size $M \times N$, that is the functions which can be represented as the conjunction of two threshold functions. The asymptotic on the number of threshold functions is known to be $\frac{6M^2N^2}{\pi^2} + O(MN^2)$ [Haukkanen and Merikoski, Discrete Applied Mathematics 2013]. It immediately gives an upper bound on the number of 2-threshold functions which is $\frac{18M^4N^4}{\pi^4} + O(M^3N^4)$. In this work we provide an asymptotic formula for the number of 2-threshold functions. To achieve this goal we establish a one-to-one correspondence between almost all 2-threshold functions and pairs of integer segments with specific properties which we call *proper* pairs of segments. We expect this bijection to be useful in algorithmic studies of 2-threshold functions.

Yong Zhang

Hybrid Online Safety Observer for Train Control System On-board Equipment

The on-board equipment of train control plays a key role in protecting trains from over-speeding. The demand of fully automatic operation puts forward higher requirements for the precise control of railway signalling systems. Conventional over-speed protection methods monitor the speed at discrete time instants. However, the over-speed behavior between discrete time instants cannot be detected, which may cause potential risks.

To address this problem, we propose a hybrid safety observation method to monitor train speed throughout the operation. In the proposed method, train behavior is modeled with a hybrid automata, which takes parameters that affect train movement into consideration. The parameters can be updated along with the movement of the train. Consequently, the train behavior in the next control cycles can be captured continuously by computing the reachable set of the model. The safety properties are also classified as a set of states. This set is obtained according to the dynamical and static speed limits regarding various over-speed protection principles. The intersection check between the reachable set and safety property set is performed based on the translated space representations. Finally, we make a judgement on whether the train speed could exceed the speed limit in a short future period.

In order to illustrate our proposed method, we present a case study based on

the real engineering data from the Chinese railway industry. A plant train model is constructed according to the specific train dynamics. A simplified over-speeding control logic is adopted in the case of the predefined scenario. The results show that the observation method is feasible in verifying the violations of CTCS-3 on-board equipment safety properties.

Michele Zito

Boolean network tomography for node failure identification in line of sight networks

A central issue in the study of communication networks is that of constantly ensuring that the structure works reliably. Network Tomography is a family of distributed failure detection algorithms based on the spreading of end-to-end measurements along the network edges. In particular, Boolean Network Tomography, (BNT for short), uses boolean measurements. Introduced in [Duffield, IEEE Trans. on Inf. Th. 2006; Ghita et al., CoNEXT 2011], BNT has recently attracted a lot of interest given the importance for example of recognizing failing nodes or links in a network, and given the simple measurements (1 failing/ 0 working) we runs through the network.

In this work we use BNT to identify corrupted/failing *nodes*. We study the identifiability of simultaneous node failures in Line-of-Sight (LoS for short) networks. LoS networks were introduced by Frieze *et al.* [SODA 2007] and have been widely studied ever since as they can be used to model communication patterns in a geometric environment containing obstacles. Like grids, LoS networks can be embedded in a finite cube of \mathbb{Z}^d , for some positive integer d . But LoS networks generalize grids in that edges are allowed between nodes that are not necessarily next to each other in the network embedding. The main contribution of our work is the analysis of a strategy to identify large sets of simultaneously failing nodes in Line-of Sight networks. Using the network vertex-connectivity, $\kappa(G)$, (i.e. the size of the minimal set of node disconnecting the graph) we are able to prove that the maximum number of simultaneously failing nodes that can be detected in any LoS network G using BNT is at least $\kappa(G) - 1$.