

REPORT ON BCTCS 2024

The 40th British Colloquium for Theoretical Computer Science 4–5 April 2024, University of Bath

Thomas Powell

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.

BCTCS 2024 was hosted by the University of Bath and held from 4th to 5th April 2024, where it was co-located with the 5th meeting of the Southern and Midlands Logic Seminar. The event attracted 52 registered participants, and featured an interesting and wide-ranging programme. A total of 23 contributed talks – predominantly by PhD students – were presented at the meeting alongside the four keynote speakers. We are grateful for the sponsorship of the British Logic Colloquium which provided travel bursaries to a number of PhD students.

BCTCS 2024 was dedicated to the memory of Professor Alan Gibbons (1942-2024) who served as the second Chairman of the BCTCS from 1992 to 1998.

BCTCS 2025 will be hosted by the University of Strathclyde from 14rd–16th April 2025. 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

Anupam Das (University of Birmingham)

Proof theoretic approaches for regular languages (and beyond)

In the second half of the 20th century various theories of regular expressions were proposed, eventually giving rise to the notion of a Kleene Algebra (KA). Kozen and Krob independently proved the completeness of KA for the model of regular languages, a now celebrated result that has been refined and generalised over the years. In recent years proof theoretic approaches to regular languages have been studied, providing alternative routes to metalogical results like completeness and decidability. In this talk, I will present a new such approach from a different starting point: right-linear (or left-linear) grammars. The resulting ‘right-linear

algebras' are more general than KAs, e.g. admitting ω -regular languages as a model too, and enjoy a simpler proof theory. This allows us to recover (more general) metalogical results in a robust way, combining techniques from cyclic proofs, games and automata.

Alex Kavvos (University of Bristol)

Two-dimensional Kripke Semantics

We revisit the duality between Kripke and algebraic semantics of intuitionistic and modal logic. We show that it is possible to raise this to the level of proofs, bridging the gap between Kripke semantics and more recent developments in modal type theory and the categorical semantics of modal logic.

Stuart Matthews (Capgemini Engineering)

Software engineering with formal methods: Quality, Time & Cost

Formal Methods – applying mathematics to software engineering – have seen a gradual industrial take-up since the early adopters in the 1980s. Early motivations were largely around questions of quality (Are we building the right system? Have we built the system right?) Until relatively recently, our assessment of the success of formal methods in becoming a mainstream industrial technique was relatively conservative, leading us to ask in a 2017 Royal Society paper, “Will the seedling ever flower?” In this talk I will give an update on this question, the answer now being rather more up-beat than it was seven years ago. We will look at the reasons for the increased penetration of formal methods, some of the companies that are leading the charge, and recent tool developments at Capgemini that are driving greater automation in the software development pipeline.

Monika Seisenberger (Swansea University)

Logic in Railway Verification

In this talk we present various logical methods we used for the verification of discrete and continuous Railway control systems. We start with an explanation of Ladder Logic, which is used in industrial control applications, present a semantics and show how a problem formulated in ladder logic amounts to a SAT solving problem. Then we talk about the modelling and verification of ERTMS (via modelchecking), and finally present more recent applications of SMT solving to the verification of Railway scheme plans. We conclude with a number of current challenges of formal methods in Railway.

Contributed Talks

Cezar-Mihail Alexandru (University of Bristol)

Interval Selection in Sliding Windows

We initiate the study of the Interval Selection problem in the (streaming) sliding window model of computation. We also establish a simple theoretical framework for proving memory lower bounds in this model which might be used beyond the Interval Selection problem. In this problem, an algorithm receives a potentially infinite stream of intervals on the line, and the objective is to maintain at every moment an approximation to a largest possible subset of disjoint intervals among the L most recent intervals, for some integer L . When processing streams, it is reasonable to focus on the most recent data items as it is modelled in the sliding window model since the near past usually affects the present more strongly than older data. In this talk, we will present the following results: In the unit-length intervals case, we give an optimal 2-approximation sliding window algorithm with space $O(|OPT|)$, and we briefly show that any sliding window algorithm that computes a $(2 - \varepsilon)$ -approximation must use space $\Omega(L)$ for any $\varepsilon > 0$. In the arbitrary-length case, we give a $(11/3 + \varepsilon)$ -approximation sliding window algorithm with space $O(|OPT|)$, for any constant $\varepsilon > 0$. We also briefly show that space $\Omega(L)$ is needed for algorithms that compute a $(2.5 - \varepsilon)$ -approximation, for any $\varepsilon > 0$. Our main technical contribution is an improvement over the smooth histogram technique, which consists of running independent copies of a traditional streaming algorithm with different start times. By employing the one-pass 2-approximation streaming algorithm by Cabello and Pérez-Lantero for Interval Selection on arbitrary-length intervals as the underlying algorithm, the smooth histogram technique immediately yields a $(4 + \varepsilon)$ -approximation in this setting. Our improvement is obtained by forwarding the structure of the intervals identified in a run to the subsequent run, which constrains the “shape” of an optimal solution and allows us to target optimal intervals differently. This is based on joint work with Dr. Christian Konrad.

Pete Austin (University of Liverpool)

Parity Games on Temporal Graphs

Temporal graphs are a popular modelling mechanism for dynamic complex systems that extend ordinary graphs with discrete time. Simply put, time progresses one unit per step and the availability of edges can change with time. We consider the complexity of solving ω -regular games played on temporal graphs where the edge availability is ultimately periodic and fixed a priori. We show that solving parity games on temporal graphs is decidable in PSPACE, only assuming the edge predicate itself is in PSPACE. A matching lower bound already holds for what we call punctual reachability games on static graphs, where one player wants to reach the target at a given, binary encoded, point in time. We further study syntactic restrictions that imply more efficient procedures. In particular, if the edge predicate is in P and is monotonically increasing for one player and decreasing for the other, then the complexity of solving games is only polynomially increased compared to

static graphs.

Hollie Baker (University of Bath)

Smooth stratification: an efficient implementation using SACLIB

A smooth stratification of an algebraic variety is a partition of that variety into effectively nonsingular subsets. Gabrielov and Vorobjov presented an algorithm for computing a basic weak stratification of a semi-Pfaffian set. I will present this algorithm along with a worked example and then describe its implementation on top of the computer algebra library SACLIB. Since SACLIB is a polynomial library, we will work in the category Semialgebraic Set, to which Gabrielov and Vorobjov's algorithm can be applied directly.

Karl Boddy (Newcastle University)

Bounded clique-width versus well-quasi-orderability by induced subgraphs

Boundedness of clique-width is a graph class property that has been extensively studied due its useful properties with respect to designing algorithms, with many NP problems tractable on graph classes with bounded clique-width. The graph class property of being well-quasi-ordered by the induced subgraph relation, at first seems completely unrelated when studying their definitions. However it has been conjectured by Lozin, Razgon and Zamarae that well-quasi-orderability implies boundedness of clique-width on finitely defined graph classes; that is classes which are defined by a finite number of forbidden induced subgraphs, which includes a large number of useful classes. As calculating clique-width in general is NP-hard, it would be useful to unite the two notions, expanding the number of classes we know to have bounded clique-width. We give an overview of the progress on this conjecture, including graph theoretical tools that are shared by the two properties. We focus on classes that are defined using only one or two forbidden graphs, where we have closed all but one case.

Adithya Diddapur (University of Bristol)

Interval Selection in Data Streams: Weighted Intervals and the Insertion-Deletion Setting

We study the problem in data streams: Given a stream of n intervals on the line, the objective is to compute a largest possible subset of non-overlapping intervals using $\tilde{O}(|OPT|)$ space, where $|OPT|$ is the size of an optimal solution. Previous work gave a $3/2$ -approximation for unit-length and a 2 -approximation for variable-length intervals, along with corresponding (tight) lower bounds. We extend this line of work to weighted intervals as well as to insertion-deletion streams. Our results include: When considering weighted intervals, a $(3/2 + \epsilon)$ -approximation can be achieved for unit intervals, but any constant factor approximation for variable length intervals requires space $\Omega(n)$. In the insertion-deletion setting where

intervals can both be added and deleted, we prove that, even without weights, computing a constant factor approximation for variable length intervals requires space $\Omega(n)$, whereas in the weighted unit-length intervals case a $(2 + \varepsilon)$ -approximation can be obtained. Our lower bound results are obtained via reductions to the recently introduced communication problem, further demonstrating the strength of this problem in the context of streaming geometric independent set problems. This talk is based on joint work with Jacques Dark and Christian Konrad.

Tala Eagling-Vose (Durham University)

Graph Homomorphisms, Colouring Games, and Forbidden Subgraphs

Recently, a framework has been proposed to investigate monotone graph classes, where a finite set of graphs is prohibited as subgraphs. In this talk, we explore both the framework and its impact through the lens of the graph homomorphism problem alongside some of its variants. A problem falls within our framework if it is efficiently solvable for graph classes of bounded treewidth, hard for the class of subcubic graphs, and remains hard when every edge is repeatedly subdivided. Under these conditions, a complete dichotomy is established for H-subgraph-free graphs, where H is any finite set of graphs. Graph homomorphism receives significant attention as it generalises the well-known graph colouring problem. In this talk, we consider a colouring game in which two players take turns colouring vertices of a graph. The first player wins if a proper colouring is obtained. Through this game, we demonstrate the presence of problems within the framework whose hardness is in the polynomial hierarchy above NP, in particular, a dichotomy exists between cases within P and those that are $\Pi P2k$ -complete.

Nathan Flaherty (University of Liverpool)

Collision-Free Robot Scheduling

Mobile robots are becoming an increasingly common part of scientific work within laboratory environments and therefore effectively coordinating them to complete necessary tasks around the laboratory whilst preventing collisions is of great practical interest. We model this theoretically by having a graph $G = (V, E)$ along with some tasks and robots, each task is on a vertex and has a certain duration, the time taken to complete the task, we then have robots which are agents that can move between adjacent vertices on the graph. The problem being to find schedules for the robots which complete all the tasks in G without any two robots sharing the same vertex or edge at any given time. We have shown that this is problem is NP-Complete even for some quite simple graph classes such as stars and cliques. This is work done jointly with Duncan Adamson, Igor Potapov and Paul Spirakis.

Iris van der Giessen (University of Birmingham)

Intuitionistic Gödel-Löb Logic, à la Simpson: Proof Theory and Birelational

Semantics

The Gödel-Löb's axiom is a modal axiom famous from the provability logic GL of Peano Arithmetic. The axiom represents a form of induction and is used in different constructive/intuitionistic settings to capture forms of recursion. We introduce a new Intuitionistic version of Gödel-Löb modal logic GL towards a computational interpretation of GL. While existing intuitionistic versions of GL are typically defined over only the box and not the diamond, we include both modalities. In the style of Alex Simpson, we develop a non-wellfounded labelled proof theory and coinciding birelational semantics and call the resulting logic IGL. This is joint work with Anupam Das and Sonia Marin.

Giulio Guerrieri (University of Sussex)

The theory of meaningfulness in the call-by-value lambda-calculus

The untyped lambda-calculus is a simple and Turing-complete model of computation that represents the kernel of any functional programming language. The semantics of the untyped call-by-name lambda-calculus is a well-developed field built around the concept of solvable terms, which are elegantly characterized in many different ways. In particular, unsolvable terms provide a consistent notion of meaningless terms, and meaningful terms can be identified with the solvable ones. The semantics of the untyped call-by-value lambda-calculus (CbV, which is closer to the real implementations of programming languages) is instead still in its early stages, because of some inherent difficulties but also because CbV solvable terms are less studied and understood than in call-by-name. On the one hand, we show that a carefully crafted presentation of CbV allows us to recover many of the properties that solvability has in call-by-name, in particular qualitative and quantitative characterizations via multi types. On the other hand, we stress that, in CbV, solvability plays a different role: identifying unsolvable terms as meaningless induces an inconsistent theory. We argue that in CbV, the correct notion of meaningful terms is captured by the concept of potential valuability. In particular, terms that are not potentially valuable provide a consistent notion of meaningless terms in CbV.

Thomas Karam (University of Oxford)

Fourier analysis modulo p on the Boolean cube

Fourier analysis in theoretical computer science is most commonly defined on the Boolean cube $\{0, 1\}^n$ identified with \mathbb{Z}_2^n . Recently, generalisations of that setting have been studied, involving restrictions to $\{0, 1\}^n$ of characters modulo p for some arbitrary prime p . Unlike in the case of mod-2 characters, different mod- p characters are no longer orthogonal when restricted to the Boolean cube. Equivalently, sets of mod- p linear forms which for $p = 2$ would be independent in a probabilis-

tic sense provided that they are linearly independent, exhibit significantly more complicated behaviour for general p . We will begin by discussing some of the basic phenomena involved, and then briefly mention some recent progress and remaining difficulties.

David Kutner (Durham University)

Reconfigurable routing in data center networks

The Reconfigurable Routing Problem (RRP) in hybrid networks is, in short, the problem of finding settings for optical switches augmenting a static network so as to achieve optimal delivery of some given workload. The problem has previously been studied in various scenarios with both tractability and NP-hardness results obtained. However, the data center and interconnection networks to which the problem is most relevant are almost always such that the static network is highly structured, whereas all previous results assume that the static network can be arbitrary (which makes existing computational hardness results less technologically relevant and also easier to obtain). In this talk, we present results for restrictions of RRP where the underlying static network is highly structured, for example consisting of a hypercube.

Lukas Holter Melgaard (University of Birmingham)

Cyclic proofs for (arithmetical) inductive definitions

We investigate the cyclic proof theory of extensions of Peano Arithmetic by (finitely iterated) inductive definitions. Such theories are essential to proof theoretic analyses of certain ‘impredicative’ theories; moreover, our cyclic systems naturally subsume Simpson’s Cyclic Arithmetic. Our main result is that cyclic and inductive systems for arithmetical inductive definitions are equally powerful. We conduct a metamathematical argument, formalising the soundness of cyclic proofs within second-order arithmetic and appealing to conservativity. This approach is inspired by those of Simpson and Das for Cyclic Arithmetic, however we must further address a difficulty that the closure ordinals of our inductive definitions (around Church-Kleene) far exceed the proof theoretic ordinal of the appropriate metatheory (around Bachmann-Howard or Bucholz), and so explicit induction on their notations is not possible. For this reason, we rather rely on a formalisation of the theory of (recursive) ordinals within second-order arithmetic.

Kheeran Naidu (University of Bristol)

Multi-Pass Streaming Lower Bounds for Graph Problems

Many one-pass graph streaming lower bounds (e.g., connectivity, maximum matching, maximal independent set) are proven by a reduction to the classic one-round one-way two-party communication problem INDEX, which is hard in one round. In two rounds, however, it becomes exponentially easier and is thus not suitable

for proving non-trivial multi-pass lower bounds. In this talk, we look at the communication game HiddenStrings introduced by Konrad and Naidu, which can be seen as a generalisation of INDEX, and show that it is suitably hard in multiple rounds/passes. In particular, we can prove multi-pass semi-streaming lower bounds for Approximate Maximum Matching and Maximal Independent Set (by Assadi et al.). This is based on joint works with Sepehr Assadi, Christian Konrad, and Janani Sundaresan.

Ian Price (Swansea University)

Two-Way Reversible Transducers as Functors

Hines suggested a notion of planarity for two-way deterministic finite automata, conjecturing links between automata theory and Temperley-Lieb algebras familiar from Physics and Knot theory. Taking inspiration from this, Pradic and Nguyễn introduced a notion of planarity for two-way finite reversible automata (2RFAs) and transducers (2RFTs), then showed that these compute star-free languages and first-order transductions, respectively. Building on this work, I will explain how to view 2RFTs as a functor from a particular “automata shape” category, due to Colcombet & Petrisan, to a category of planar transition diagrams. This category is autonomous (like compact-closed, but not symmetric) and can be viewed as a target for interpreting non-commutative lambda calculi. I will work through a particular example and hint at future applications to linear logic.

Tymofii Prokopenko (University of Liverpool)

Capturing an invisible robber on a planar graph

“Cops and robbers” is a game on a graph where a team of cops starting at a given vertex on a given graph (known as an arena) aims to capture a robber starting at another vertex (i.e. at least one of the team of cops must be in the same vertex as the robber). Moves alternate between a team of cops and a robber and in classical settings all players move along the edges of a given graph. Two main variants of a game with respect to the visibility of a robber by a team of cops: full visibility and zero visibility. In full visibility case the cops know the location of the robber on the graph and thus have full information at any time. In the zero visibility case, the robber is invisible to the cops throughout the game and only at the moment of capturing the robber the location becomes known. We present an algorithm for capturing the robber in the case of zero-visibility by a team of robots on a planar graph which extends the known results for simpler structures like trees. Following the formal analysis, we estimate the number of cops and the time/space complexity required for capturing the robber. Moreover, we define the new problem for a game with a mixed setting of full and zero visibility of a robber.

Chris Purdy (Royal Holloway, University of London)

A cyclical proof system for higher-order fixed point logic

Non-well-founded proof systems, allowing infinitely deep derivations, facilitate implicit forms of (co)inductive reasoning, validated by a so-called global trace condition requiring each infinite path in the derivation to witness some infinite descent. Reasoning using cyclic proofs, the finite representations of regular such non-well-founded derivations, is often more natural than using induction rules, since explicit induction invariants (hypotheses) are not required. We present work-in-progress on a non-well-founded (cyclic) natural deduction system for higher order fixed point logic with explicit approximations μHOLex . To our knowledge, ours is the first cyclic proof system to combine fixed point operators and negation in a higher order setting. We also discuss the development of a suitable global trace condition in the absence of explicit approximations, as well as potential propositions-as-types interpretations for both of these deduction systems.

Andrew Ryzhikov (University of Oxford)

Fixed Vector Addition Systems and Bounded Arithmetic

Recently, there has been significant success in understanding the computational complexity of the reachability problems in vector addition systems with states (VASS). We consider a variant of this problem where the VASS is fixed, and only its initial and target configurations are a part of the input. If the input is given in the binary encoding, we show that there exists a VASS such that this problem is PSPACE-hard. If the input is given in the unary encoding, the problem becomes more difficult to analyse, since in this case the input consists of a constant-length sequence of unary numbers. Despite that, VASS remain quite expressive: in particular, for every fixed formula of bounded arithmetic with counting there exists a fixed VASS simulating this formula in some natural way. This is a joint work with Christoph Haase and Andrei Draghici.

James Swire (Swansea University)

An empirical evaluation of a quality assurance process: A case study in the railway sector

Critical systems, such as railways, require multiple levels of testing and verification to be fit for use. The objective of this research is to determine the quality of a verification approach for the control software of railway interlocking computers. As such, we will utilise statistics to determine the likelihood of a fault being detected on any given track plan using its respective safety properties. To accomplish this, we will inject faults into an otherwise correct control software and see if the verification approach detects them. A secondary aim of this research is to automate this fault injection process. Our research aims to establish a fault probability space, determined by these statistics, of control software after passing the

verification process.

Marc Thatcher (University of Sussex)

Parallel Functional Programming with Interaction Nets

Interaction nets are a graph-rewriting system based on Girard’s linear logic. Suitable as both a computational model and a simple programming language, a most interesting aspect of the latter case is the ease with which they can move from sequential to parallel execution. The combination of locality (graph rewrites are defined on pairs of nodes which do not affect the greater graph) and linearity (any sub-graph’s interface with the greater graph must be unchanged by the reduction rules) means that nets can automatically take advantage of any parallelism in the algorithm, without any extra work from the programmer. Interaction net programming languages proposed to date have used various forms of explicit syntax, which is easy to reason about but hard to use. We propose a functional-like programming language for defining interaction nets with a bijective mapping from the language to nets and show how this can be used to provide parallel programming “for free”. An alpha-version of an implementation of this language is available.

Riccardo Treglia (Kings College London)

Monadic Intersection Types, Relationally

We consider an extension of intersection types to a computational lambda-calculus with algebraic operations à la Plotkin and Power. We achieve this by considering monadic intersections – whereby computational effects appear not only in the operational semantics, but also in the type system. We will focus on why in the effectful setting termination is not anymore the only property of interest, and consequently how to analyze the interactive behaviour of typed programs with the environment is of interest. As the main result, our type system is able to characterize the natural notion of observation, both in the finite and in the infinitary setting, and for a wide class of effects. This is achieved via a novel combination of syntactic techniques with abstract relational reasoning, which we will briefly introduce as a tool to lift all the required notions to the monadic setting. This is joint work with Francesco Gavazzo and Gabriele Vanoni.

Sean Watters (University of Strathclyde)

Extensional Finite Sets and Multisets in Agda

Lists and list-like types are easy to represent in type theory. However, sets and multisets are more troublesome; without advanced features such as quotient types, representations of these typically lack either decidable equality or the correct equational theory. In this talk, I will discuss how to realise finite sets and multisets in Agda, while maintaining the above desiderata. We achieve this using a generalisation of Catarina Coquand’s data type of fresh lists. Our constructions

satisfy properties akin to (multi)set extensionality, which allows us to prove that they satisfy the universal properties of the free idempotent commutative monoid and free commutative monoid, respectively. I will also show how fresh lists realise many other free algebraic structures. This is joint work with Clemens Kupke and Fredrik Nordvall Forsberg.

Xin Ye (Durham University)

Computing Balanced Solutions for Large International Kidney Exchange Schemes When Cycle Length Is Unbounded

In kidney exchange programmes (KEP) patients may swap their incompatible donors leading to cycles of kidney transplants. Nowadays, countries try to merge their national patient-donor pools leading to international KEPs (IKEPs). As shown in the literature, long-term stability of an IKEP can be achieved through a credit-based system. In each round, every country is prescribed a “fair” initial allocation of kidney transplants. The initial allocation, which we obtain by using solution concepts from cooperative game theory, is adjusted by incorporating credits from the previous round, yielding the target allocation. The goal is to find, in each round, an optimal solution that closely approximates this target allocation. There is a known polynomial-time algorithm for finding an optimal solution that lexicographically minimizes the country deviations from the target allocation if only 2-cycles (matchings) are permitted. In practice, kidney swaps along longer cycles may be performed. However, the problem of computing optimal solutions for maximum cycle length ℓ is NP-hard for every $\ell \geq 3$. This situation changes back to polynomial time once we allow unbounded cycle length. However, in contrast to the case where $\ell = 2$, we show that for $\ell = \infty$, lexicographical minimization is only polynomial-time solvable under additional conditions (assuming $\text{sP} \neq \text{NP}$). Nevertheless, the fact that the optimal solutions themselves can be computed in polynomial time if $\ell = \infty$ still enables us to perform a large scale experimental study for showing how stability and total social welfare are affected when we set $\ell = \infty$ instead of $\ell = 2$.

Tansholpan Zhanabekova (University of Liverpool)

Semantic Flowers for Good-for-Games and Deterministic Automata

We present an innovative approach for capturing the complexity of ω -regular languages using the concept of flowers. This semantic tool combines two syntax-based definitions, namely the Mostowski hierarchy of word languages and syntactic flowers. The former is based on deterministic parity automata with a limited number of priorities, while the latter simplifies deterministic parity automata by reducing the number of priorities used, without altering their structure. Synthesising these two approaches yields a semantic concept of flowers, which offers a more effective way of dealing with the complexity of ω -regular languages. This

BEATCS no 143

latter provides a comprehensive definition of semantic flowers and shows that it captures the complexity of ω -regular languages. We also show that this natural concept yields simple proofs of the expressive power of good-for-games automata.