

## **REPORT ON BCTCS 2021**

### **The 37th British Colloquium for Theoretical Computer Science**

**29–31 March 2021, University of Liverpool**

Patrick Totzke and Michele Zito

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 2021 was hosted by the University of Liverpool and held from 29<sup>th</sup> to 31<sup>st</sup> March, 2021. Due to the on-going COVID-19 pandemic, all talks were presented over Zoom, and Gather . town was used to encourage virtual networking. The meeting was well attended, attracting 128 registered participants, and featured an interesting and wide-ranging programme of five invited talks and 25 contributed talks.

The meeting spanned three days, with each morning session and each afternoon session opening with an invited keynote lecture. The meeting opened on the first morning with an invited talk by Eleni Akrida (Durham) on algorithmic properties of temporal graphs, and the afternoon session opened with an invited talk by Herbert Edelsbrunner (Vienna) on the properties of random polytopes inscribed on the 2-sphere. The second day opened with an invited talk by Alessandra Russo (London) on logic-based machine learning, and the afternoon session opened with the LMS Keynote Lecture in Discrete Maths delivered by Professor Jan Krájček (Prague), who gave an engaging presentation on proof complexity. The third and final day of the meeting opened with an invited lecture by Kousha Etessami (Edinburgh) on computing fixed points of monotone functions.

BCTCS 2022 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](http://www.bctcs.ac.uk).

## Invited Talks

**Eleni Akrida (Durham University)**

*Temporal graphs: Algorithms and complexity*

We discuss here the notion of temporal graphs, which are abstract models of networks that change as time progresses. In particular, a temporal graph on an underlying graph  $G$  is such that the edges of  $G$  have labels that are positive integers indicating discrete times (moments) of availability. A basic notion that arises from that definition is that of a journey, which extends the notion of a path in static graphs: it is a path in which subsequent edges appear with increasing labels. In this talk, we present recent work on temporal graphs through the lens of algorithms and complexity. We examine foremost journeys, design issues of temporal networks, as well as issues of traversal including temporal variants of the Travelling salesman Problem.

**Herbert Edelsbrunner (IST Vienna)**

*Random triangles and random inscribed polytopes*

Given three random points on a circle, the triangle they form is acute with probability  $1/4$ . In contrast, the triangle formed by three random points in the 2-sphere is acute with probability  $1/2$ . Both of these claims have short geometric proofs. We use the latter fact to prove that a triangle in the boundary of a random inscribed 3-polytope is acute with probability  $1/2$ .

Picking  $n$  points uniformly at random on the 2-sphere, we take the convex hull, which is an inscribed 3-polytope. The expected mean width, surface area, and volume of this polytope are  $2(n-1)/(n+1)$ ,  $4\pi[(n-1)(n-2)]/[(n+1)(n+2)]$ , and  $(4\pi/3)[(n-1)(n-2)(n-3)]/[(n+1)(n+2)(n+3)]$ . These formulas are not new but our combinatorial proofs are.

This is joint work with Arseniy Akopyan and Anton Nikitenko.

**Kousha Etessami (University of Edinburgh)**

*The complexity of computing a fixed point of a monotone function, and some applications*

The task of computing a fixed point of a monotone function arises in a variety of applications. In this talk I describe recent work in which we have studied the computational complexity of computing a fixed point of a given monotone function that maps a finite  $d$ -dimensional grid lattice with sides of length  $N = 2^n$  to itself, where the monotone function is presented succinctly via a boolean circuit with  $d \cdot n$  input gates and  $d \cdot n$  output gates, and the underlying ordering is the standard coordinate-wise partial order on  $d$ -dimensional vectors in  $[N]^d$ . By Tarski's theorem, any monotone function has a fixed point.

We refer to the search problem of either finding a fixed point or finding a wit-

ness pair for non-monotonicity as the Tarski problem. It turns out that Tarski subsumes a number of important computational problems. In particular, it is essentially computationally equivalent to the task of computing a pure Nash Equilibrium for (succinctly presented) super-modular games, or games with strategic complementarities, which are very widely studied classes of games in economics. Also, computing the value of Condon's turn-based simple stochastic (reachability) games, as well as the more general problem of computing the value of Shapley's original stochastic games to within a given accuracy, is reducible to Tarski. Tarski is contained in both the total search complexity classes PLS and PPAD.

This is joint work with C. Papadimitriou, A. Rubinfeld, and M. Yannakakis.

**Jan Krajíček (Charles University, Prague)**

*LMS Keynote Lecture in Discrete Maths*

*Proof complexity*

Proof complexity is an area connecting mathematical logic and computational complexity theory. It has several facets and I try to explain what some of these are. In particular, I discuss a few ways proof complexity relates to the existence of algorithms of various types.

**Alessandra Russo (Imperial College London)**

*Logic-based learning for interpretable AI*

Learning interpretable programs from data is one of the main challenges of AI. Over the last two decades there has been a growing interest in logic-based learning, a field of machine learning aimed at developing algorithms and systems for learning programs that can explain labelled data in the context of given background knowledge. Contrary to the black-box Deep Learning methods, logic-based learning systems learn programs that can be easily expressed into plain English and explained to human users, so facilitating a closer interaction between humans and machines. In this talk, I present recent advances in the field of logic-based learning. I introduce frameworks and algorithms for learning different classes of programs, ranging from definite programs under the Least Herbrand model semantics, to highly expressive non-monotonic programs under the Answer Set semantics. Such programs include normal rules, non-determinism, choice rules, hard and weak constraints, which are particularly useful for modelling human preferences. I discuss the relationship between these frameworks and illustrate a range of real-world problems that have been addressed using these systems, and open challenges in this area.

## Contributed Talks

**Duncan Adamson (University of Liverpool)**

***Ranking bracelets in polynomial time***

Ranking and unranking combinatorial objects is a well studied problem, with classical results for words, trees and partitions. More recently polynomial time algorithms have been provided for necklaces, also known as cyclic words. Despite this, the problem of ranking and unranking bracelets (also known as turnover necklaces) has remained a noted open problem. In this talk the first polynomial time algorithms for ranking and unranking bracelets are presented.

**Benjamin Bumpus (University of Glasgow)**

***Spined categories: generalising tree-width beyond graphs***

Problems that are NP-hard in general are often tractable on inputs that have a recursive structure. For instance consider classes defined in terms of “graph decompositions” such as of bounded tree- or clique-width graphs. Given the algorithmic success of graph decompositions, it is natural to seek analogues of these notions in other settings. What should a “tree-width- $k$ ” digraph or lattice or temporal graph even look like?

Since most decomposition notions are defined in terms of the internal structure of the decomposed object, generalizing a given notion of decomposition to a larger class of objects tends to be an arduous task. In this talk I show how this difficulty can be reduced significantly by finding a characteristic property formulated purely in terms of the category that the decomposed objects inhabit, which defines the decomposition independently of the internal structure. I introduce an abstract characterisation of tree-width as a vast generalisation of Halin’s definition of tree-width as the maximal graph parameter sharing certain properties with the Hadwiger number and chromatic number. Our uniform construction of tree-width provides a roadmap to the discovery of new tree-width-like parameters simply by collecting the relevant objects into our new notion of a spined category.

This is joint work with Zoltan A. Kocsis.

**Marcel Dall’Agnol (University of Warwick)**

***Quantum proofs of proximity***

With the prevalence of massive datasets and small computing devices, offloading computation becomes increasingly important. One may model this problem as follows: a weak device (the *verifier*) holds an input string and communicates with an all-powerful computer (the *prover*) so as to verify the validity of this string. Since the verifier cannot perform the verification on its own, the prover provides a digest of the computational task, allowing the verifier to check that it was performed correctly. In the extreme setting of property testing, the verifier cannot

even read all of its input, and must decide whether it is valid or far from any valid string; in this case, the the digest is called a proof of proximity. We formalise the notion of quantum proofs of proximity, where prover and verifier are quantum computers, and show that a large class of properties admits quantum speedups. We also begin to chart the complexity landscape of the quantum/classical as well as the interactive/non-interactive variants of these proof systems.

This is joint work with Tom Gur and Subhayan Roy Moulik.

**Bhaskar DasGupta (University of Illinois at Chicago)**

*Maximizing coverage while ensuring fairness: a tale of conflicting objectives*

Ensuring fairness in computational problems has emerged as a key topic during recent years, buoyed by considerations for equitable resource distributions and social justice. It is possible to incorporate fairness in computational problems from several perspectives, such as using optimization, game-theoretic or machine learning frameworks. In this talk we address the problem of incorporating fairness from a combinatorial optimization perspective. We formulate a combinatorial optimization framework, suitable for analysis by researchers in approximation algorithms and related areas, that incorporates fairness in maximum coverage problems as an interplay between two conflicting objectives. Fairness is imposed in coverage by using colouring constraints that minimizes the discrepancies between number of elements of different colors covered by selected sets; this is in contrast to the usual discrepancy minimization problems studied extensively in the literature where (usually two) colors are not given a priori but need to be selected to minimize the maximum color discrepancy of each individual set. Our main results are a set of randomized and deterministic approximation algorithms that attempts to simultaneously approximate both fairness and coverage in this framework.

This is joint work with A Asudeh, T Berger-Wolf and A Sidiropoulos.

**Andrei-Cristian Diaconu (University of Oxford)**

*Controlling control: functional language design and implementation*

The design and implementation of functional languages is a problem that has been tackled from many perspectives. Generally, one starts with a formal specification of the semantics, after which this is used to derive an implementation of the language (e.g., an interpreter). The usual (operational) approaches to giving semantics are small-step, big-step and definitional interpreters, the last one capturing both a formal specification of the language and an actual implementation. However, the presentation of a functional language should aim to use a mixture of these (e.g., provide a small-step semantics, and then implement a definitional interpreter). This becomes a problem when languages with more complex control mechanisms such as exception handling and concurrency are considered, because the mentioned approaches generally do not scale well.

In this talk, we present a unified view of semantics and implementation, that scales well with the complexity of the language, whilst still giving one the ability to reason about it. We first present a simple and powerful way to specify small-step operational semantics using evaluation contexts, which allow us to succinctly capture complex control mechanisms. We then show how the idea of evaluation contexts is tightly linked to delimited continuations. Furthermore, delimited continuations can then be used to implement a definitional interpreter, which is both concise and highly extensible. To achieve this, we make use of various control operators and employ multi-prompt delimited continuations. Finally, we explore how the delimited continuation approach relates to other similar approaches of implementing interpreters, such as effects-and-handlers and monadic reflection.

**Alejandro Flores-Lamas (Royal Holloway, University of London)**

***Solving the distance  $k$ -clique problem in 1-outerplanar graphs***

A clique is a subset of vertices of a simple graph  $G$  in which each vertex in the set is adjacent to all other vertices in it; in other words, any pair of such vertices are connected. Similarly, a distance  $k$ -clique is a subset of vertices in which any pair of vertices in the set are connected by at most  $k$ -edges; thus, a clique is a distance 1-clique set. A common formulation of these sets as a problem asks to find a Maximum Distance  $k$ -Clique, MDkC, from a given graph; i.e., the largest cardinality set. This problem is NP-hard in arbitrary graphs. In this work, we address the MDkC problem in 1-outerplanar graphs; i.e., graphs that admit a drawing on the plane such that its edges only intersect at the vertices, and each vertex lies on the outer face of the drawing. Our approach to solving this problem was first finding an MDkC set in a tree. Then, we adapted this algorithm to run in 1-outerplanar graphs. The proposed algorithm uses dynamic programming, and it goes through two phases. In the first phase, the algorithm follows a postorder traversal on a tree decomposition of  $G$  to compute the size of an MDkC set. Then, during a second traversal (top-bottom) it identifies the vertices that belong to the set. The running time of this algorithm is  $O((k + 1) \cdot 2^{\tau+1} \cdot n)$ , where  $k$  comes from the MDkC problem, while  $\tau$  and  $n$  are the treewidth and order of  $G$ , respectively.

**Arved Friedemann (Swansea University)**

***Propagator networks for unification and proof search***

Logical languages have often shown to be useful as they give compact code for a huge variety of search problems. However, their performance highly depends on the underlying search engine. Propagator networks have been developed that combine the principles of SAT, SMT and CSP solvers all in one formalism. These networks can be used to get a general notion of what unit propagation means for more than just boolean functions and they can be easily parallelised. Here, we present a framework for propagator networks that can be used to easily model

unification and branching and create a logical language to perform proof search for PROLOG-like languages.

This is joint work with Philipp Dargel.

**Cameron Hargreaves (University of Liverpool)**

***The earth movers' distance as a metric for inorganic compositions***

Computational chemists have an intuitive grasp of what makes two compositions “chemically similar”, but this concept is surprisingly difficult to capture using standard numeric techniques. Here we present the earth movers' distance (EMD) as a measure of similarity between two compounds, which operates by optimally pairing elements in a source composition to their most similar elements in a target composition and scoring the resultant similarities. This is akin to the cognitive process by which humans judge chemical similarity and as such, the resultant distances consistently align with human understanding. The chemical formula gives a mathematically abstract representation of the inorganic material which is typically difficult to work with. By contrast, we can use a well-formed distance in a wide range of analytical techniques, allowing us to cement these abstractions into tangible information. This is demonstrated effectively on the 12,567 binary structures of the ICSD, where we use this distance to plot detailed chemical maps which separate compositions into families of clear similarity both on a global and local scale. These maps can be clustered using unsupervised machine learning techniques to automatically partition our compounds into digestible subgroups which enables us to identify and distil critical chemical trends that would otherwise have been overlooked. We can additionally use these distances for the automated retrieval of structures from chemical databases, where an exact formulation may never have been reported but a closely related structure provides the reference needed. This metric is fast to compute between two compositions in practice, making it a strong candidate for many other applications in data driven materials discovery. We discuss how the EMD is calculated between two compositions, and demonstrate strengths against the standard metric.

**Alexandros Hollender (University of Oxford)**

***The complexity of gradient descent:  $CLS = PPAD \cap PLS$***

We study search problems that can be solved by performing gradient descent on a bounded convex polytopal domain and show that this class is equal to the intersection of two well-known classes: PPAD and PLS. As our main underlying technical contribution, we show that computing a Karush-Kuhn-Tucker (KKT) point of a continuously differentiable function over the domain  $[0, 1]^2$  is PPAD  $\cap$  PLS-complete. This is the first natural problem to be shown complete for this class. Our results also imply that the class CLS (Continuous Local Search) – which was defined by Daskalakis and Papadimitriou as a more “natural” counter-

part to  $\text{PPAD} \cap \text{PLS}$  and contains many interesting problems – is itself equal to  $\text{PPAD} \cap \text{PLS}$ .

This is joint work with John Fearnley, Paul Goldberg and Rahul Savani.

### **Finnbar Keating (University of Warwick)**

#### ***Resolving data flow dependencies of concurrent programs with a graded monad***

In concurrent programs that interact with memory, it is important to know that variables are read from and written to in the correct order to prevent race conditions or inconsistent data. To this end, we introduce a graded monad for identifying the exact memory locations read and written by a given function. With this we can replace Haskell’s IO monad with a variant indexed by a representation of the operations performed. This new information allows us to identify data-flow dependencies between functions at the type-level; that is, if one function reads a memory cell that is written to by another function we might wish to evaluate the latter before the former or vice-versa depending on the context.

With these dependencies known at compile-time, we then provide a function that establishes the correct order of execution for these program parts at compile-time with respect to their data-flow dependencies. This compile-time sort even allows us to identify functions that can be run in parallel, and can skip evaluation of code that relies on state that has not changed. We also prevent compilation in the presence of two errors: race conditions via multiple functions writing to the same memory cell, and dependency loops from functions writing to cells that other functions read from.

### **Peter Kiss (University of Warwick)**

#### ***Dynamic matchings***

We present a framework for deterministically rounding a dynamic fractional matching. Applying our framework in a black-box manner on top of existing fractional matching algorithms, we derive the following new results: the first deterministic algorithm for maintaining a  $(2-\delta)$ -approximate maximum matching in a fully dynamic bipartite graph, in arbitrarily small polynomial update time; the first deterministic algorithm for maintaining a  $(1+\delta)$ -approximate maximum matching in a decremental bipartite graph, in polylogarithmic update time; and the first deterministic algorithm for maintaining a  $(2+\delta)$ -approximate maximum matching in a fully dynamic general graph, in *small* polylogarithmic (specifically,  $O(\log^4 n)$ ) update time.

Our rounding scheme works by maintaining a good *matching-sparsifier* with bounded arboricity, and then applying the algorithm of Peleg and Solomon to maintain a near-optimal matching in this low arboricity graph. To the best of our knowledge, this is the first dynamic matching algorithm that works on general graphs by using an algorithm for low-arboricity graphs as a black-box subroutine.



**Nina Klobas (Durham University)**

***Fast recognition of some parametric graph families***

Understanding the cycle (or anticycle) structure of a graph is fundamentally related to graph families such as trees, perfect graphs, bipartite graphs, chordal graphs, pancyclic graphs, and many others. A particularly strong cycle-related property is the notion of cycle-regularity, introduced by Mollard, which has been used to better understand the structure of graph families such as hypercubes or generalized Petersen graphs. In this talk we present three graph families, namely I-graphs, double generalized Petersen graphs and folded cubes and show how their cyclic structure helped us devise linear time recognition algorithms for them.

**Joe Livesey (University of Liverpool)**

***Propositional gossip protocols***

Gossip protocols are programs that can be used by a group of agents to synchronise what they know. Assuming each agent holds a secret, the goal of a protocol is to reach a situation in which all agents know all secrets. Distributed epistemic gossip protocols use epistemic formulas in the component programs for the agents.

We investigate open problems regarding propositional gossip protocols, in which whether an agent wants to make a call depends only on their currently known set of secrets. Specifically, we show that all correct propositional gossip protocols, i.e., ones that always terminate in a situation where all agents know all secrets, require the communication graph to be complete, whilst investigating the minimum number of calls required. We also investigate the complexity of the problem of checking correctness of a given propositional gossip protocol, before discussing implementing such a check with model checker NuSMV.

**Lily Major (Aberyswyth University)**

***Developments with manipulating Lyndon factorizations using evolutionary computation methods***

String factorization is an important tool for partitioning data for parallel processing and other algorithmic techniques often found in the context of big data applications such as bioinformatics or compression. Duval's well-known linear algorithm uniquely factors a string over an ordered alphabet into Lyndon words, patterned strings which are strictly smaller than all of their cyclic rotations. While Duval's algorithm produces a pre-determined factorization, modern applications motivate the demand for factorizations with specific properties, e.g., those that minimize/maximize the number of factors, or consist of factors with similar lengths.

We considered the problem of finding an alphabet ordering that yields a Lyndon factorization with such properties. For problems with no known heuristics, evolutionary computation methods may be useful, as such, we produced a flexible

evolutionary framework and evaluated it on biological sequence data. For the minimization case, we proposed a new problem-specific heuristic, Flexi-Duval, and a problem-specific mutation operator for Lyndon factorization. We have shown that for individual amino acid sequences, very long (often single) Lyndon factors can be produced using Flexi-Duval, and in comparison with Duval's algorithm, a much more balanced length for each Lyndon factor is achievable for specific texts.

We also consider the balancing problem for a much larger range of input texts, including random sequences, repetitions with some introduced noise, and natural language text corpora. We consider improvements to the Flexi-Duval algorithm and evaluate it with the broader input texts, and consider the problem of maximizing the number of Lyndon factors via evolutionary computation methods.

**Diptapriyo Majumdar (Royal Holloway, University of London)**

***Tractability of Konig edge deletion problems***

A graph is said to be a Konig graph if the size of its maximum matching is equal to the size of its minimum vertex cover. The Konig edge deletion problem asks if a graph has a set of at most  $k$  edges whose deletion results in a Konig graph. While the vertex deletion version of this problem was shown to be fixed-parameter tractable (FPT) more than a decade ago, the FPT status of Konig edge deletion has remained open since then. It has been conjectured to be  $W[1]$ -hard in several papers. In this paper, we settle the conjecture by proving it  $W[1]$ -hard. In addition, we prove that a variant of this problem, where we are given a graph  $G$  and a maximum matching  $M$ , and we want a Konig edge deletion set of size at most  $k$  that is disjoint from  $M$ , is fixed-parameter tractable.

This is joint work with Rian Neogi, Venkatesh Raman, and S. Vaishali.

**Michael McKay (University of Glasgow)**

***A non-trivial polynomial time algorithm for a 3D stable roommates problem***

In this talk we consider possible formalisations of the three-dimensional stable roommates problem (3D-SR). Players specify preference lists over their peers, and the task is to partition the players into triples based on their preferences. A number of hardness results exist under various schemes of preference representation. We consider a formalisation of 3D-SR in which agents' preferences are additively separable and binary, named 3D-SR-AS-BIN. The decision problem then asks whether we can partition the players into triples so that no three players would prefer to be in a set together than remain in their current triples. We show that 3D-SR-AS-BIN is NP-complete and consider its restriction in which preferences are symmetric, named 3D-SR-SAS-BIN. We show that every instance of 3D-SR-SAS-BIN contains a stable matching that can be found using a non-trivial algorithm in polynomial time. These results help us explore the boundary between NP-hardness and polynomial-time solvability in problems of coalition formation.

**Jay Morgan (Swansea University)**

*Trustable machine learning systems*

Machine Learning (ML) has had a remarkable impact on society. Everything from the phones in our pockets, to the cars that we drive, are being increasingly outfitted with this progressively sophisticated suite of algorithms. But while many of the most basic and fundamental algorithms from ML can be formally verified and tested for safety without much trouble, the same may not be said for Deep Learning (DL) – a prominent forerunner in the state-of-the-art for ML research. These DL models, while performing simple matrix-to-matrix operations at a micro-level, have evolved in scale far past what is tractable for current formal verification methods – all in the pursuit of improving accuracy and performance. This issue of tractability is unsettling considering that the existence of adversarial examples is well known in the ML community. These adversarial examples occur when very small changes to the input space result in a large change in the output space and cause a misclassification made by the DL model. In the context of self-driving vehicles, small defects and visual artefacts in the sensor input of the DL model, could lead the vehicle to wrongly conclude a stop sign indicates to continue driving where it should have stopped. While the manufacturers will need to put safeguards in place to prevent this from happening, we should formally prove the (non)-existence of these adversarial examples in the DL model itself. In this presentation, I present the foundational knowledge for understanding adversarial examples, how we can use the input space to dictate the search space for the existence of these examples, and demonstrate their presence with the use of SAT-solving. This work, as a free and open-source project, provides a framework for ML practitioners to verify their own architectures.

**Eric Munday (University of Edinburgh)**

*Forcing infinite memory for  $\liminf$  total payoff objectives in countable MDPs*

We look at an example of a countable MDP with integer-valued transition rewards. Every infinite run induces a sequence of total payoffs (the sequence of sums of all rewards seen so far). The objective is to maximise the probability that the  $\liminf$  is non negative. We present a counterexample to show that infinite memory is required in order to satisfy the  $\liminf$  objective for epsilon-optimal strategies. Furthermore, this holds even if the step-counter is implicit in the state and the MDP is finitely branching.

**Kheeran Naidu (University of Bristol)**

*A unifying class of algorithms for semi streaming bipartite maximum matching*

In the semi-streaming model of Feigenbaum et al., a graph with  $n$  vertices is presented to an algorithm as a stream of edges where the storage space of the algo-

rithm is bounded by  $O(n \text{ polylog } n)$ . It provides an efficient model for processing massive graphs which have quickly become widespread. An algorithm in this model typically takes anywhere from one pass up to logarithmically many passes of the stream, in the same order.

A long-standing open problem is to improve upon a  $\frac{1}{2}$ -approximation for maximum cardinality matching (MCM) in one pass of the stream (in adversarial order). Currently, a simple greedy algorithm achieves the state-of-the-art approximation factor. Improving upon a  $\frac{1}{2}$ -approximation in two passes for bipartite MCM, however, was achieved by Konrad et al. with a  $(\frac{1}{2} + 0.0192)$ -approximation. Kale and Tirodkar, and Esfandiari et al. independently improved these bounds to a  $(\frac{1}{2} + 0.0625)$ - and  $(\frac{1}{2} + 0.0833)$ -approximation, respectively. Most recently, Konrad set the state-of-the-art for bipartite MCM at a  $(\frac{1}{2} + 0.0858)$ -approximation in two passes of the stream.

We present a wider class of two-pass bipartite MCM algorithms of which the above algorithms are special cases. We show that there are two optimal algorithms in this class which achieve the state-of-the-art, one of which is a novel algorithm. Finally, we construct a worst-case example which achieves the analytical lower bound, proving that: the analysis is indeed tight; a  $(\frac{1}{2} + 0.0858)$ -approximation is best for this class of algorithms; and a completely different algorithmic approach will be needed to further improve this bound.

This is joint work with Christian Konrad.

**Adam Ó Conghaile (University of Cambridge)**

*Partition games, compositionally: Structure and power of linear algebraic games*

Spoiler-duplicator games have been a crucial tool for proving expressibility upper bounds on logical languages. Recently, a large body of work on game comonads has exposed a fascinating compositional structure lying beneath the surface of many of these games, relating games to each other and to relevant structural parameters such as treewidth. This has helped to better understand the pebble game, modal bisimulation game and games for generalised quantifiers. One class of games whose category theoretic structure is not yet understood is that of partition games. As explored in the thesis of Bjarki Holm, partitions give a rich grammar for constructing spoiler-duplicator games for various logics including the matrix equivalence and invertible maps games for linear algebraic logics. In this talk, I briefly review recent developments in game comonads before introducing partition games and describing where they might fit into this compositional picture.

This is joint work with Anuj Dawar.

**Bruno Pasqualotto Cavalari (University of Warwick)**

*Monotone circuit lower bounds from robust sunflowers*

Monotone Boolean circuits form one of the largest natural circuit classes for which

we are able to prove exponential size lower bounds. Such lower bounds play a pivotal role in complexity theory, being a proxy for lower bounds on communication complexity, proof complexity and optimization. For over 20 years, the best known lower bound on the size of monotone circuits computing an explicit  $n$ -bit monotone Boolean function was  $\exp(n^{1/3-o(1)})$ .

In this talk, we motivate the study of monotone circuits and their applications, and present the first lower bound for monotone circuit size of order  $\exp(n^{1/2-o(1)})$ . The proof employs the classical approximation method of Razborov and recent robust sunflower bounds of Alweiss, Lovett, Wu and Zhang, and of Rao.

This is joint work with Mrinal Kumar and Benjamin Rossman.

### **Nicos Protopapas (University of Liverpool)**

#### ***Impartial selection, additive guarantees, and prior information***

We study the problem of impartial selection, a topic that lies in the intersection of computational social choice and mechanism design. The goal is to select the most popular individual among a set of community members. The input can be modelled as a directed graph, where each node represents an individual, and a directed edge indicates nomination or approval of a community member to another. An impartial mechanism is robust to potential selfish behaviour of the individuals and provides appropriate incentives to voters to report their true preferences by ensuring that the chance of a node to become a winner does not depend on its outgoing edges. Our goal is to design impartial mechanisms that select a node with an in-degree that is as close as possible to the highest in-degree. Recent progress has identified impartial selection rules with optimal approximation ratios. It was noted though that worst-case instances are graphs with few vertices. Motivated by this fact, we propose the study of additive approximation, the difference between the highest number of nominations and the number of nominations of the selected member, as an alternative measure of quality. We present randomized impartial selection mechanisms with additive approximation guarantees of  $o(n)$ , where  $n$  is the number of nodes in the input graph. We furthermore demonstrate that prior information on voters' preferences can be useful in the design of simple (deterministic) impartial selection mechanisms with good additive approximation guarantees. In this direction, we consider different models of prior information and analyze the performance of a natural selection mechanism that we call approval voting with default.

### **Benjamin Smith (University of Liverpool)**

#### ***Fixed-parameter tractability of pinwheel scheduling***

The pinwheel scheduling problem asks if it is possible to construct a perpetual schedule for a set of  $k$  tasks of unit length where the maximum gap between repetitions of the same task is limited by an ordered multiset  $A$  of  $n$  constraints, one per

task. We demonstrate that pinwheel scheduling is fixed-parameter tractable with respect to the parameter  $k$ . We also show that, for any given value of  $n$ , a finite set of schedules can solve *all solvable* pinwheel scheduling instances. We then introduce exhaustive-search algorithms for both pinwheel scheduling instances and partial pinwheel scheduling instances (where only a prefix of  $A$  is known and gaps have to be left in the schedule), and use that to construct the Pareto surfaces over all possible schedules for  $k$  tasks, for  $k \leq 5$ . These results have implications for the bamboo garden trimming problem of Gasieniec et al.

**Siani A. Smith (Durham University)**

***The complexity of acyclic, star and injective colouring for  $H$ -free graphs***

Colouring is among the best known problems in computer science. One direction of study has been to consider the complexity of colouring for  $H$ -free graphs, that is graphs which do not contain  $H$  as an induced subgraph. A full dichotomy is known in the case where the number of available colours,  $k$ , is part of the input whereas infinitely many open cases remain where  $k$  is fixed.

We study three variants of the colouring problem, acyclic, star and injective colouring. An acyclic colouring of a graph  $G$  is a proper colouring in which the union of any two colour classes induces a forest. Similarly a star colouring of  $G$  is a proper colouring in which the union of any two colour classes is  $P_4$ -free, in other words it induces a star forest. Finally, an injective colouring, also known as a distance 2 colouring or an  $L(1, 1)$ -labelling, is a proper colouring in which the union of any two colour classes is  $P_3$ -free. In each case we combine new and known results to obtain a full complexity dichotomy where  $k$  is fixed and classify all but at most finitely many graphs  $H$  where  $k$  is part of the input. In fact, for both star and injective colouring, we show that only one open case remains.

This is joint work with J Bok, N Jedličková, B Martin and D Paulusma.

**Federico Vastarini (University of York)**

***Random graph generation using hyperedge replacement grammars***

This work addresses the problem to satisfy the necessity of graph based applications when requesting random test data. May they involve classical algorithms, software testing, pointer manipulation, pattern recognition or even complex networks, the proposed solution, allows for the specification of any graph domain. Such approach, that, to the best of our knowledge has never been followed by any pre-existing method, uses context-free hyperedge replacement grammars to define graph classes, over which, elements are sampled uniformly at random.

In order to achieve this, a generalised efficient version of Mairson's algorithms for the sampling of strings over non-ambiguous context-free grammars in Chomsky normal form is adapted to the setting of hyperedge replacement. The presented method is correct in that given a non-ambiguous hyperedge replacement grammar

$G$  in Chomsky normal form and a size  $n$ , a  $n$ -hypergraph in the language  $L(G)$  is generated uniformly at random.

**Pavel Vesely (University of Warwick)**

***Relative error streaming quantiles***

Approximating distributions and quantiles is a fundamental task in data analysis. We consider streaming algorithms that make just one pass over a massive dataset and must use a small amount of memory, polylogarithmic in the input size. The problem of approximating the median and other quantiles with an additive error is sufficiently well-understood, however, the case of the stronger relative error is still open. Such an error guarantee provides more accurate estimates for extreme quantile queries (such as the 99.5th percentile), thus helping to understand the tails of the distribution. We describe a new streaming algorithm providing the relative error guarantee with a nearly optimal trade-off between space usage and accuracy.

This is joint work with G Cormode, Z Karnin, E Liberty, and J Thaler.