

REPORT ON BCTCS 2014

The 30th British Colloquium for Theoretical Computer Science 9–11 April 2014, Loughborough University

Paul Bell, Daniel Reidenbach

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

BCTCS 2014 was hosted by the Department of Computer Science at Loughborough University, and held from 9th – 11th April 2014. The event attracted over 40 participants from sixteen universities, and featured an engaging and wide-ranging programme of four invited talks and 25 contributed talks. These were in large part from PhD students and covered the full gamut of topics in Theoretical Computer Science. Abstracts of the talks are provided below.

The conference began with an invited talk by Leszek Gąsieniec, University of Liverpool, entitled “Distributed maintenance of mobile entities”. Other invited talks were given by Timo Kötzing, Friedrich-Schiller-Universität Jena/Germany (“Recent advances in inductive inference”) and Achim Jung, University of Birmingham (“A modal Belnap logic”). As in previous years, the London Mathematical Society (LMS) sponsored a keynote talk in Discrete Mathematics, which this year was given by Jeffrey Shallit, University of Waterloo/Canada, on “Open problems in automata theory”. The financial support of this lecture by the LMS is gratefully acknowledged. We also acknowledge the generous financial support of the Heilbronn Institute for Mathematical Research, which provided 24 bursaries to cover the full costs of attendance for research students.

Invited Talks at BCTCS 2014

Leszek Gąsieniec, University of Liverpool

Distributed maintenance of mobile entities

With the recent advent of ad-hoc, not well-structured, large, and (very often) dynamic network environments there is a strong need for more robust, universal, and inexpensive distributed network protocols. The purpose of these protocols is to support basic network formation and integrity mechanisms as well as more dedicated tasks such as information dissemination, network search and exploration, network monitoring and others.

One of the novel and promising alternatives in supporting such network protocols are dedicated teams of mobile entities (MEs) that can work independently on top of basic network system routines. The MEs' ability to communicate and to move within the environment impels the design and implementation of efficient formation, communication and navigation mechanisms including motion control and coordination mechanisms that allow MEs to perform dedicated tasks collectively.

We will provide an introduction to the field and will discuss several extensively studied algorithmic problems as well as those just touched upon in the recent years. The talk will be concluded with open problems.

Achim Jung, University of Birmingham

A modal Belnap logic

Four valued logic was introduced by Nuel Belnap in the 70s. It is very easy to motivate and seems to be central to Computer Science; in fact, one of his papers on the subject was called "How a computer should think". Adding mathematical structure to his basic ideas turned out not to be so easy, however. Much work was done by Arieli and Avron in the 90s, and more recently, by Umberto Rivieccio, a collaborator on the work to be presented, which concerns a modal extension of Belnap's work. The topic is also related to my longstanding interest in using Stone Duality to link semantics and logic for computer science.

Timo Kötzing, Friedrich-Schiller-Universität Jena/Germany

Recent advances in inductive inference

3,5,7,11,13 – what's next? What general rule (apparently) produces this sequence? Maybe the sequence lists all the odd primes, but what if the next datum is 15? Maybe all odd numbers that are not squares? Since the 1960's there are formal models defining what it means to learn or predict such sequences; this area of research is called inductive inference. In this talk I will briefly review the main classical results and then focus on recent advances in inductive inference, espe-

cially concerning the development of general techniques. Applications of these techniques include, for example, questions regarding avoidance of seemingly inefficient learning behavior.

Jeffrey Shallit, University of Waterloo/Canada, the LMS-sponsored keynote speaker in Discrete Mathematics

Open problems in automata theory

In this talk I will survey some of my favorite open problems from automata theory, including the separating words problem, decidability problems related to number theory and the Endrullis-Hendriks problem on transducers.

Contributed Talks at BCTCS 2014

Eleni Akrida, University of Liverpool

Ephemeral networks with random availability of links: diameter and connectivity

In this work we consider temporal networks, the links of which are available only at random times (randomly available temporal networks). Our networks are ephemeral in the sense that their links appear sporadically, only at certain times, within a given maximum time (called lifetime of the network). More specifically, our temporal networks notion concerns networks, whose edges are assigned one or more random discrete-time labels drawn from a set of natural numbers. The labels of an edge indicate the discrete moments in time at which the edge is available. In such networks, information (e.g., messages) have to follow temporal paths, i.e., paths, the edges of which are assigned a strictly increasing sequence of labels. We first examine a very hostile network: a clique, each edge of which is known to be available only one random time in the time period $\{1, 2, \dots, n\}$ (where n is the number of vertices). How fast can a vertex send a message to all other vertices in such a network? To answer this, we define the notion of the Temporal Diameter for the random temporal clique and prove that it is $\Theta(\log n)$ with high probability and in expectation. In fact, we show that information dissemination is very fast with high probability even in this hostile network with regard to availability. This result is similar to the results for the random phone-call model. Our model, though, is weaker. Our availability assumptions are different and randomness is provided only by the input. We show here that the temporal diameter of the clique is crucially affected by the clique's lifetime, a , e.g., when a is asymptotically larger than the number of vertices, n , then the temporal diameter must be $\Omega((a/n) * \log n)$. We, then, consider the least number, r , of random instances at which an edge is available, in order to guarantee at least a temporal path between any pair of vertices of the network (notice that the clique is the only network for

which just one instance of availability per edge, even non-random, suffices for this). We show that r is $\Omega(\log n)$ even for some networks of diameter 2. Finally, we compare this cost to an (optimal) deterministic allocation of labels of availability that guarantees a temporal path between any pair of vertices. For this reason, we introduce the notion of the Price of Randomness and we show an upper bound for general networks.

Theofanis Apostolopoulos, King's College London

Sparse signal recovery as a non-linear problem

My contributed talk will focus on a novel research field, called Compressed Sensing (CS) method, which has attracted considerable research with several new application areas, mainly signal and image compression. It was introduced recently for simultaneously sampling and compressing signals and enabling new reconstruction techniques for applications where the standard sampling process is not feasible or very expensive. In fact, CS adopts a new sampling scheme that does not follow the principle of conventional approach depicted by the sampling theorem of Nyquist-Shannon. The goal is to efficiently recover any type of signal, such as speech and image data, using what was previously considered as highly incomplete and inaccurate (under-sampled) measurements. This is an ill-posed inverse problem, which can be solved as an l_0 norm based optimisation problem, with the aim to find the best fit which minimises the difference between the solution and the observations while satisfying all the given constraints. In this talk, I will also introduce a new swarm based heuristic for efficiently recovering signals, with high probability. It is an iterative process which finds an approximation of the l_0 -norm based problem viewed as a combinatorial optimization problem. In each iteration every agent calculates and carries a slightly different feasible solution based on the current best (optimal) solution, which is necessary so as to avoid being trapped to one of the numerous local minima. This method is very efficient and quick compared to other conventional methods, such as the classical log-barrier and Least squares methods, even under the presence of noise, based on experimental results. In particular, the heuristic is compared with other alternative sparse recover methods in terms of complexity, computational time, samples size, and recovery error. Possible improvement for enhancing the performance of the heuristic could be to re-weight the approximate l_0 norm, by using coefficients at every iteration; an approach that has been applied successfully to similar l_0 , l_1 and l_2 norm based CS problems.

Christopher Bak, University of York

Towards an implementation of rooted graph programs

Rooted Graphs are used to improve the efficiency of graph matching when applying graph rewriting rules. The basic idea is to automatically match a node in the

rule to a node in the host graph, restricting the search space to a small area of the host graph. The increase in efficiency comes at the cost of flexibility: graph programs consisting of rooted rules are significantly more complicated than equivalent programs with standard rules. I present a Topological Sorting graph program which strikes a balance by using both rooted and non-rooted rules to solve the problem, and discuss some issues in the ongoing implementation of rooted graph programs in the graph programming language GP.

Jannis Bulian, University of Cambridge

Graph isomorphism parameterized by elimination distance to bounded degree

A commonly studied means of parameterizing graph problems is the deletion distance from triviality, which measures the number of vertices that need to be deleted from a graph to place it in some class for which efficient algorithms are known. In the context of graph isomorphism, we define triviality to mean a graph with maximum degree bounded by a constant, as such graph classes admit polynomial-time isomorphism tests. We generalise deletion distance to a measure we call elimination distance to triviality, based on elimination trees or tree-depth decompositions. We establish that graph isomorphism is FPT when parameterized by elimination distance to bounded degree, generalising results of Bouland et al. on isomorphism parameterized by tree-depth.

Leroy Chew, University of Leeds

The complexity of theorem proving in circumscription and minimal entailment

Circumscription is one of the main formalisms for non-monotonic reasoning. It uses reasoning with minimal models, the key idea being that minimal models have as few exceptions as possible. In this contribution we provide the first comprehensive proof-complexity analysis of different proof systems for propositional circumscription. In particular, we investigate two sequent-style calculi: MLK defined by Olivetti (J. Autom. Reasoning, 1992) and CIRC introduced by Bonatti and Olivetti (ACM ToCL, 2002), and the tableaux calculus NTAB suggested by Niemelä (TABLEAUX, 1996). In our analysis we obtain exponential lower bounds for the proof size in NTAB and CIRC and show a polynomial simulation of CIRC by MLK. This yields a chain $\text{NTAB} < \text{CIRC} < \text{MLK}$ of proof systems for circumscription of strictly increasing strength with respect to lengths of proofs.

Michalis Christofi, King's College London

Worst-case behavior of distributed algorithms for the maximum concurrent flow problem

A Multicommodity Flow Problem is a problem of designing flows of commodities in a common network. The flows must be feasible, that is they cannot exceed the edge capacities, and they must satisfy the demand of each commodity. Multi-

commodity flow problems have a wide variety of important applications in areas such as VLSI circuit design, network design, production and distribution of goods, transportation systems and communication systems. We consider the multicommodity flow problem which is called the Maximum Concurrent Flow problem. The objective is to minimise the maximum edge congestion, where the congestion of an edge is defined as the ratio of the flow to the capacity. In this talk we discuss algorithms which solve this problem in the following distributed manner: one agent controls one commodity, and the agents communicate at the end of each computation round via a billboard. Algorithms of this type were proposed by Awerbuch, Khandekar and Rao [SODA 2007] and Awerbuch and Khandekar [PODC 2007], who showed that an approximate solution can be reached in the number of rounds which is linear in the maximum length L of a path followed by any flow. We show that this running-time bound is asymptotically tight by constructing a worst-case input network and analyzing the performance of the algorithms on this network. We also propose a heuristic improvement of these algorithms, analyze its performance on our worst-case input, and indicate why we should expect that it improves running times on general networks.

Alejandro Erickson, Durham University

Computer science takes back data centre networks from engineering

Companies like Google, Amazon, and Microsoft house massive warehouses full of interconnected computers which provide services to the whole world. The demand for such services is pushing the limits of traditional data centre designs, and this research area, long dominated by engineers, is becoming a hot topic in theoretical computer science. How can currently available equipment be interconnected in order to increase the size and performance of data centres while reducing the relative cost?

I give an overview of some recent "computer science-y" developments in the world of data centres, and I discuss a novel approach for converting an arbitrary graph into a dual-port server-centric data centre network.

This work is supported by the EPSRC grant "INPUT: Interconnection Networks, Practice Unites with Theory".

Carl Feghali, Durham University

On the complexity of partitioning graphs into disjoint cliques and a triangle-free subgraph

We investigate the computational complexity of deciding whether the vertices of a graph can be partitioned into a disjoint union of cliques and a triangle-free subgraph. This problem is known to be NP-complete on arbitrary graphs. Our hardness results are on planar graphs and perfect graphs. In contrast, we provide a finite list of forbidden induced subgraphs for cographs with such a partition, thus

yielding a linear time recognition algorithm.

(Joint work with Faisal N. Abu-khazam and Haiko Müller.)

Michael Gale, University of Cambridge

Solving an existential crisis in Haskell

Haskell's type system provides mechanisms for type refinement within the scope of certain value expressions if GADTs or type classes are used. The type system propagates sufficient information to ensure that nothing can go wrong even if types are erased from the run-time representation of a program. This is not the case when we are using existential types, where we deliberately hide concrete types from the type system. Nevertheless, we may desire to eliminate existential types in a different part of a program in order to restore the original types. For this purpose, we propose an extension to Haskell which allows programmers to restrict existential types within individual data constructors to finite, but open, domains of types. Each type in such a domain must be associated with a value tag that is then stored at run time to allow it to serve as witness in a case expression.

Thomas Gorry, University of Liverpool

The evacuation problem: group search on the line

This talk will consider the Group Search Problem, or Evacuation Problem, in which K mobile entities located on the line perform a search for a specific destination. The mobile entities are initially placed at the same point (origin) on the line and the target is located at some unknown distance (d) either to the left or to the right of the origin. All mobile entities must simultaneously occupy the destination, and the goal is to minimize the time necessary for this to happen. The problem where $K = 1$ is called the cow-path problem, and the complexity of this is known to be $9d$ in the worst case (when the cow moves at unit speed), it is also known that this is the case for $K \geq 1$ mobile entities travelling at unit speed. This talk presents a clear argument for this claim as well examining the case when $K = 2$ mobile entities with different speeds, showing a surprising result that the bound of $9d$ can still be achieved when 1 mobile entity has unit speed and the other moves with speed at least $1/3$.

Ivaylo Hristakiev, University of York

Analysing graph programs for confluence

The graph programming language GP, developed at York, is an experimental domain-specific language for high-level problem solving on graphs and graph-like structures. In general, graph programs are highly nondeterministic because of their rule-based nature. However, a special case on nondeterminism called confluence ensures the functional behaviour of the execution. Confluence detection is done through construction of critical pairs, which represent conflicts in minimal

context. This technique has been extended to several variations of graph transformation, but not to GP. In this talk, I will present what these extensions are together with their associated issues and also report on ongoing work on static confluence checking of GP programs.

Augustine Kwanashie, University of Glasgow

Profile-based optimal matchings in the student/project allocation problem

In the *Student/Project Allocation problem (SPA)* we seek to assign students to group or individual projects offered by lecturers. Students are required to provide a list of projects they find acceptable in order of preference. Each student can be assigned to at most one project and there are constraints on the maximum number of students that can be assigned to each project and lecturer. A matching in this context is a set of student-project pairs that satisfies these constraints.

We seek to find matchings that satisfy optimality criteria based on the *profile* of a matching. This is a vector whose i th component indicates the number of students obtaining their i th-choice project. Various profile-based optimality criteria have been studied. For example, one matching M_1 may be preferred to another matching M_2 if M_1 has more students with first-choice projects than M_2 .

In this talk we present an efficient algorithm for finding optimal matchings to SPA problems based on various well known profile-based optimality criteria. We model SPA as a network flow problem and describe a modified augmenting path algorithm for finding a maximum flow which can then be transformed to an optimal SPA matching. This approach allows for additional constraints, such as project and lecturer lower quotas, to be handled flexibly without modifying the original algorithm.

Karoliina Lehtinen, University of Edinburgh

Syntactic and semantic complexity in modal μ

The modal μ calculus is a temporal logic evaluated on labelled transition systems. It combines next-state modalities with greatest and least fixpoint operators, resulting in a logic capable of expressing both finite and infinite behaviour such as reachability, safety, eventual safety and much more. In particular, it subsumes many temporal logics such as LTL and CTL. Despite its high expressiveness, the core algorithmic problems around modal μ remain decidable: the model-checking problem is widely conjectured to be in P and satisfiability is EXPTIME-complete. This makes modal μ a widely studied formalism for program verification.

Modal μ 's expressiveness is based on a simple but productive syntax: by increasing the number of alternations between greatest and least fixpoint operators, modal μ can express properties of increasing complexity. However, formulas of large alternation depth can also express much simpler properties and currently we lack the tools to differentiate between inherent and accidental complexity. Given

that the current best model checking algorithms are exponential in a function of the alternation depth of a formula, deciding whether a formula can be expressed with fewer alternations remains one of the main open problems surrounding modal μ .

This talk presents work on identifying and simplifying non-strict formulas, that is to say formulas that are equivalent to a formula with fewer alternations. The strictness of a formula implies the satisfiability of a set of derived formulas describing systems that witness the necessity of each alternation. If these witnesses do not exist for some formula, there are syntactic transformations which yield a formula of lower alternation depth.

Hsiang-Hsuan Liu, University of Liverpool

Scheduling for electricity cost in smart grid

We study an online scheduling problem arising in demand response management in smart grid. Consumers send in power requests with a flexible set of timeslots during which their requests can be served. For example, a consumer may request the dishwasher to operate for one hour during the periods 8am to 11am or 2pm to 4pm. The grid controller, upon receiving power requests, schedules each request within the specified duration. The electricity cost is measured by a convex function of the load in each timeslot. The objective of the problem is to schedule all requests with the minimum total electricity cost. As a first attempt, we consider a special case in which the power requirement and the duration a request needs service are both unit-size. For this problem, we present a polynomial time online algorithm that gives an optimal solution and show that the time complexity can be further improved if the given set of timeslots is a contiguous interval.

Iain McBride, University of Glasgow

Modelling practical placement of trainee teachers to schools

Several countries successfully use centralized matching schemes to assign students to study places or recent graduates to their first positions in a labour market. In this work we describe a model motivated by specific features of the Slovak and Czech education systems where each recently graduated trainee teacher specializes in a small number of subjects, each school has an overall capacity and further each school has partial capacities with respect to each of the available subjects. We show that the problem is unlikely to be efficiently solvable even under severe restrictions on the total number of subjects available, the partial capacities of schools for the available subjects and the number of acceptable schools each trainee teacher may list. Since these results suggest an efficient method of producing optimal solutions is unlikely, we present an integer programming model for finding a maximum cardinality matching in an instance of the teachers assignment problem and we present the results of the application of this IP model to real data

from the allocation process for allocating trainee teachers in Slovakia.

This is joint work with Tamás Fleiner (Budapest University of Technology and Economics), Katarína Cechlárová (P.J. Šafárik University, Košice, Slovakia), David Manlove (University of Glasgow) and Eva Potpinková (P.J. Šafárik University, Košice, Slovakia).

Markus Pfeiffer, University of St Andrews

The rational hierarchy of semigroups

The word problem for semigroups is known to be undecidable in general. On the other hand, deciding the word problem of the natural numbers or the integers is simple. My research focuses on finding classes of semigroups with word problem decidable by different types of automata. In this talk I will introduce what I call the rational hierarchy of semigroups, semigroups that have word problem decidable by asynchronous, two-tape, finite state automata, and conjunctions, Boolean combinations of such automata.

Jean Jose Razafindrakoto, Swansea University

Provably total search problems in fragments of bounded arithmetic below polynomial-time

In bounded arithmetic, a host of theories have been developed and which correspond to complexity classes within the polynomial hierarchy and below polynomial-time (see Cook and Nguyen's monograph "Logical Foundations of Proof Complexity, Cambridge University Press, New York, NY, USA, first ed., 2010", for an overview). Recent research tries to characterize the provably total NP search problems in such theories, where a total NP search problem is provably total in a theory \mathcal{T} if it can be formalized in the language of \mathcal{T} and \mathcal{T} can prove that for each instance, there exists a solution to the search problem.

Given a class S of provably total NP search problems for some theory, the general aim of our research project is to identify some specific provably total NP search problem class (usually defined via some specific combinatorial principle) which is complete within S under AC^0 -many-one reduction; completeness should be proven using AC^0 -reasoning only. For the theory related to polynomial-time, we identify the search problem class Inflationary Iteration (*IITER*) which serves our above described aim. A function F (defined on finite strings) is inflationary if X is a subset of $F(X)$ (under the natural identification of strings with finite sets). An *IITER* principle is defined as a special case of the iteration principle, in which the iterated function has to be AC^0 -computable and inflationary.

Cook and Nguyen have a generic way of defining a bounded arithmetic theory VC for complexity classes C below polynomial-time. For such a theory VC , we define a search problem class $KPT[C]$ which serves our above described aim. These problems are based on a version of Herbrand's theorem, proven by Kra-

jíček, Pudlak and Takeuti in “Bounded arithmetic and the polynomial hierarchy, Ann. Pure Appl. Logic, 52(1-2):143-153, 1991, International Symposium on Mathematical Logic and its Applications (Nagoya, 1988)”.

This is joint work with Arnold Beckmann.

Paolo Serafino, Teesside University

Heterogeneous facility location without money

Mechanism Design is a novel research field mainly concerned with optimization problems that have to operate under the assumption that their input is distributed across *selfish* agents. In this setting, *mechanisms* (i.e., typically *allocation algorithms*) have to elicit their input from the agents and have to ensure (usually via suitable payment functions) that agents report *truthfully* the part of input they possess. The challenge faced in this setting is that agents are not reliable, in the sense that they can misreport their private information. Alas, it is often the case that monetary transfers between the mechanism and the agents cannot be performed. Motivated by this kind of considerations, Procaccia and Tennenholz (*Approximate Mechanism Design Without Money*, EC09) proposed the research agenda of approximate mechanism design without money, which aims at leveraging approximation, instead of payments, as a means to enforce truthfulness. In this line of work, the simple yet general and elegant problem of facility location has attracted much interest. The model which is typically considered therein features single-parameter agents (i.e., agents whose type is a single number encoding their position on a real line). In the wake of this line of research, we formulate and initiate the study of *heterogeneous facility location without money*, a problem akin to the traditional facility location problem but featuring multi-parameter agents. More specifically, we study truthful mechanisms without money for the problem in which *heterogeneous* facilities (facilities serving different purposes) have to be located and agents are only interested in some of them. We study the approximation ratio that can be achieved by truthful mechanisms in this setting, deriving some approximation bounds which make a surprising parallel with our knowledge of truthfulness for the classical single-dimensional facility location problem.

Yiannis Siantos, King’s College London

Inferring network properties and embedded structure using random walks

We study the use of random walks to estimate global properties of graphs, for example the number of edges, vertices, triangles, and generally, the number of small fixed subgraphs. We consider two methods for this based on first returns of random walks: the cycle formula of regenerative processes and random walks with weights based on the property under consideration. In addition we use these methods to infer the embedded structure of graphs, such as whether a pre-defined subset of vertices is better connected internally than the rest of the graph. We dis-

cuss theoretical foundations for both methods and present experimental results on the rate of convergence of all the estimates. Both theory and experiments highlight the importance of high-weight vertices for the efficiency of either method.

Rob van Stee, University of Leicester

An optimal online bin packing algorithm

In the online bin packing problem, items of size at most 1 arrive one by one and need to be packed into bins of size 1 without knowledge of future items. We measure the performance of algorithms for this problem by comparing the number of bins used to the optimal number of bins. The competitive ratio of an algorithm is the highest possible ratio between these numbers (i.e., for all possible inputs).

We present an online bin packing algorithm with absolute competitive ratio $5/3$, which is best possible. The previous best known algorithms for this problem were Best Fit and First Fit, which were only recently shown to be exactly 1.7-competitive.

Anthony Stewart, Durham University

Parallel knock-out schemes for special graph classes

We consider parallel knock-out schemes for graphs. These schemes proceed in rounds. In the first round each vertex in the graph selects exactly one of its neighbours, and then all the selected vertices are eliminated simultaneously. In subsequent rounds this procedure is repeated in the subgraph induced by those vertices not yet eliminated. The scheme continues until there are no vertices left, or until an isolated vertex is obtained (since an isolated vertex will never be eliminated). A graph is reducible if there exists a parallel knock-out scheme that eliminates every vertex in the graph (for instance a graph that has a Hamilton cycle is reducible within one round). The Parallel Knock-Out problem is that of deciding whether a graph is reducible. This problem is known to be NP-complete. We discuss known results for this problem together with a number of new results for special graph classes (such as split graphs).

Nihan Tokac, Durham University

Fixed parameter tractability of hybridization number and rooted subtree prune and regraft distance

The decision problems computing hybridization number and rooted subtree prune and regraft distance are important to understand and model reticulation events in evolutionary biology. In this paper, we show that computing hybridization number is fixed parameter tractable when the parameter is the minimum level of network on T and T' . As well as, computing rooted subtree prune and regraft distance between two rooted binary phylogenetic trees on the same label set is fixed parameter tractable when the parameter is the minimum rSPR-level of network on T

and T' .

Chalita Toopsuwan, King's College London

Maximal anti-exponent of gapped palindromes

A palindrome is a string that reads the same backward and forward. We consider gapped palindromes which are words of the form $uv\tilde{u}$ for some words u, v with $|v| \geq 2$ where \tilde{u} denotes the reversal of u . Mimicking the standard notion of string exponent, we define the antiexponent of a gapped palindrome $uv\tilde{u}$ as the quotient of $|uv\tilde{u}|$ over $|uv|$. We apply techniques based on the use of a suffix automaton and on the reversed Lempel-Ziv factorisation to an input string y containing no ordinary palindrome, and design an algorithm to compute the maximal anti-exponent of gapped palindromes of the string. Our algorithm runs in linear-time on a fixed-size alphabet in contrast to a naive cubic time solution.

William Whistler, Durham University

The counting complexity of planar graph homomorphism problems

In this talk I present my current progress on classifying the counting complexity of graph homomorphism problems with inputs restricted to planar graphs.

Michele Zito, University of Liverpool

Relaxation and rounding for appliance allocation in the smart grid

We introduce a scheduling algorithm for a set of air-conditioners deployed in a building whose electricity comes from the grid as well as from renewable sources. The main objective of this study is to introduce a heuristic algorithm that is able to reduce electricity bills, keeping the temperature within comfort levels in the building, and maximizing the utilization of domestic renewable power. The algorithm uses relaxation in order to convert a Mixed Integer Linear Program into an LP problem and then a rounding mechanism to increase the utilization of domestic renewable power by scheduling the load on times where there is enough renewable power even this period is peak hours in terms of cost.

(This is joint work with M. Arikiez, A. Fernandez Anta, F. Grasso, and D. Kowalski.)