

REPORT ON BCTCS 2016

The 32nd British Colloquium for Theoretical Computer Science 22–24 March 2016, Queen’s University Belfast

Amitabh Trehan

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 2016 was hosted by Queen’s University Belfast (*QUB*), and held from 22nd to 24th March, 2016. The event attracted over 30 participants, and featured an interesting and wide-ranging programme of six invited talks and 20 contributed talks. Abstracts for all of the talks from BCTCS 2015 are provided below. We are particularly thankful to the *Heilbronn Institute for Mathematical Research* who provided bursaries for 7 PhD students (including 2 female students). The many (sponsored and non-sponsored) PhD students ensured enthusiastic participation and talks over a wide range of research areas. QUB generously provided free use of the venue, the newly opened and redesigned *Graduate School*, to promote learning among graduate students. We are also thankful to the *London Maths Society (LMS)* for their annual sponsorship of the *LMS keynote speaker in Discrete Maths* - Prof. Valerie King (University of Victoria, Canada).

The talks covered a wide range of topics in TCS and there was global participation ranging from Canada, USA, Turkey, Iceland, the UK, of course, and the republic of Ireland. The opening talk was given by Matthew Hennessy (Trinity College Dublin). Matthew gave an overview of recent research in transactional distributed systems outlining semantic theories for process calculus for co-operating transactions. Michael Butler (Southampton) continued the theme in the afternoon with a talk on Event Refinement Structures (ERS) for the Event-B formal method.

Wednesday began with an interesting talk by Magnus M Halldórsson (Reykjavik University) on the philosophical and practical questions involving choosing the right problems to work on continuing with his recent results involving scheduling problems arising from wireless networks. The morning continued with Gregory Chockler (Royal Holloway) talking about using erasure codes with replication for reliable storage in asynchronous distributed systems. In the pre-lunch session, Bhaskar DasGupta (University of Illinois, Chicago, USA) gave a talk featuring non-trivial bounds on node expansions and cut-sizes for Gromov-hyperbolic graphs (or, hyperbolic graphs for short).

The session continued with a talk from our LMS keynote speaker, Valerie King. Valerie is an ACM fellow who is well known for her work on algorithms (in particular, dynamic and distributed algorithms). Valerie with Jared Saia (University of New Mexico) and others have made extensive and fundamental advances in solving the very important problem of Byzantine agreement. They have used a number of innovative techniques including spectral methods for achieving their results. In her talk, she described work on efficient algorithms for Byzantine agreement without secrecy by making connections with a new collective coin flipping problem. The afternoon also featured a stroll to the nearby Botanic gardens and the Ulster museum.

The final day began with another interesting talk- by Bruce Kapron (University of Victoria, Canada) on Gambling, information and encryption security. The talk featured many animated race horses, Yao's question (1982)- can encryption security may be characterized using computational information? His talk would have probably benefited many bookies before the Grand National. Cassio D. Campos (QUB) gave a talk on Inferential Complexity in Probabilistic Graphical Models which should probably have featured on day one considering the excellent introduction to Belfast that it gave! The final invited speaker was Robert Giles (Economics, QUB) who gave a talk on the underlying game theoretic concepts behind consent in network formation, an area in which he has extensive experience.

BCTCS 2017 will be hosted by St Andrews University, Researchers and PhD students wishing to contribute talks concerning any aspect of Theoretical Computer Science are cordially invited to do so. Further details are available from the BCTCS website at www.bctcs.ac.uk.

Invited Talks at BCTCS 2016

Michael J Butler (Southampton University)

Verification Patterns for Refinement

Event-B is a general purpose refinement-oriented formal method. Event Refinement Structures (ERS) provide additional structure for refining Event-B models, in particular, for refining course-grained atomicity to fine-grained atomicity when reasoning about concurrent and distributed systems. This talk presents specification-oriented patterns for ERS refinement and verification. A specification-oriented pattern is determined by the shape of the problem description rather than the solution description.

Robert Giles (Queen's University Belfast)

Consent in Network Formation: Game Theoretic Solutions

We consider the formation of networks under the principle of mutual consent and costly link formation. This problem has attracted considerable game theoretic

analysis that we survey in this presentation. First, we consider link-based stability concepts founded on the seminal work by Jackson and Wolinsky (1996). In particular, we survey some refinements of pairwise stability notions in the context of costly link formation. Second, we look at non-cooperative game theoretic analysis of consent in link formation known as the Myerson network formation game. The deficiency of the Nash equilibrium concept is shown. Instead a belief-based equilibrium concept, known as a self-confirming equilibrium, is explored to model meaningful network formation in this context. An equivalence with strictly pairwise stable networks is shown. Finally, we explore the role of correlations of payoff functions through the notion of a potential. If network payoffs admit a potential function, dynamic network formation algorithms can be devised that converge to strictly pairwise stable networks that are supported as a self-confirming monadic equilibrium in the standard Myerson network formation game.

Magnus M Hallsdórrson (Reykjavik University)

“What problem should I solve?” and Efficiency in Wireless Networks

The selection of topic to work on is perhaps the most important aspect of research, and one fraught with pitfalls. When we determine significance of a problem from the “importance in applications”, it is important to understand the assumptions underlying the formulations, since all abstractions leak somewhere. We ponder these issues while examining recent progress in scheduling problems underlying wireless networking. The modeling of reality, in this case “interference”, is crucial. Simple abstractions can be valuable, as long as we are aware of their limitations. This brings up the issue of wider interest: how well can simpler abstractions approximate more detailed/complex ones?

Matthew Hennessy (Trinity College Dublin)

Behavioural Theories for Co-operating Transactions

Relaxing the isolation requirements on transactions leads to systems in which transactions can co-operate to achieve distributed goals. However in the absence of isolation it is not easy to understand the desired behaviour of transactional systems, or the extent to which the other standard ACID properties of transactions can be maintained: atomicity, consistency and durability. In this talk I give an overview of some recent research in this area, outlining semantic theories for a process calculus augmented by a new construct for co-operating transactions. In particular I focus on property logics which can be used to distinguish behaviourally between such transactions.

Bruce Kapron (University of Victoria, Canada)

Gambling, Computational Information and Encryption Security

We revisit the question, originally posed by Yao (1982), of whether encryption

security may be characterized using computational information. Yao provided an affirmative answer, using a compression-based notion of computational information to give a characterization equivalent to the standard computational notion of semantic security. We give two other equivalent characterizations. The first uses a computational formulation of Kelly's (1957) model for "gambling with inside information", leading to an encryption notion which is similar to Yao's but where encrypted data is used by an adversary to place bets maximizing the rate of growth of his wealth over a sequence of independent, identically distributed events. The difficulty of this gambling task is closely related to Vadhan and Zheng's (2011) notion of KL-hardness, a form of which is equivalent to a conditional form of the pseudoentropy introduced by Håstad et. al. (1999). Using techniques introduced to prove this equivalence, we also give a characterization of encryption security in terms of conditional pseudoentropy. Finally, we reconsider the gambling model with respect to adversaries with linear utility in an attempt to understand whether assumptions about the rationality of adversaries may impact the level of security achieved by an encryption scheme. (Joint work with Mohammad Hajiabadi.)

Valerie King (University of Victoria, Canada)

Tossing a Collective Coin and Coming to Agreement

Over 35 years ago, Leslie Lamport formulated a fundamental problem of coordination in a distributed network. He asked us to imagine an army led by generals, who send messages to each other with the goal of coming to agreement on a strategy. Planted among those generals are spies who seek to thwart this goal. Not long after this Byzantine agreement problem was presented, there were a few developments: an impossibility for any deterministic scheme, a randomized exponential time algorithm, and a demonstration that one globally known coin toss could solve the problem in constant expected time.

Some researchers turned to the use of committed secret coinflips via cryptography, while others turned to the study of collective coin flipping with full information. Recently, the need for Byzantine agreement without the overhead of cryptographic techniques has arisen in decentralized digital currency systems.

I will describe joint work with Jared Saia and others on efficient algorithms for Byzantine agreement without secrecy, including the first polynomial time algorithm for this problem in a fully asynchronous model, which is obtained by solving a new collective coin flipping problem.

Contributed Talks at BCTCS 2016

Athraa Al-Krizi (University of Liverpool)

Probabilistic Model Checking of One-Dimensional Nano Communication System

Molecular communication is a bio-inspired paradigm in which molecules are transmitted, propagated and received between nanoscale machines. Establishing controlled molecular transmissions between these nanomachines represents a major challenge. Many studies have aimed to model the physical medium (channel) of molecular communication, primarily from a communication or information-theoretical perspective.

In this talk we model a simple time-slotted communication system between nanoscale machines in a one-dimensional environment. This communication system employs some bio-inspired rules that can be checked at each interval. The system model has been verified using the probabilistic model checking tool PRISM on different sized networks. We were able to verify that acknowledgement has been obtained, and thus, communication between these nanonodes has been ascertained. The results were promising for further study of more complex scenarios such as multi-access channels.

Thomas van Binsbergen (Royal Holloway, University of London)

Executable Component-Based Semantics

To improve the practicality of formal semantic definitions, the P_LanCompS project has developed a component-based approach. In this approach, the semantics of a language is defined by translating its constructs to combinations of so-called fundamental constructs, or ‘funcons’. Each funcon is defined using a modular variant of structural operational semantics, and forms a language-independent component that can be reused in definitions of different languages.

For specifying component-based semantics, we have designed and implemented a meta-language called CBS. CBS includes specification of abstract syntax, of its translation to funcons, and of the funcons themselves. In this talk we discuss the compilation of CBS funcon specifications to Haskell code. In particular, we shall discuss how modularity is obtained in Haskell definitions of funcons.

Joshua Blinkhorn (University of Leeds)

Dependency Schemes: Semantics and Soundness in QBF Calculi

The tremendous success of SAT solvers in recent years has led to a natural extension to quantified Boolean formula (QBF) solving. Whereas SAT is the canonical NP-complete decision problem, deciding QBF is PSPACE-complete, and captures the problem of determining winning strategies in two-player games with perfect information. This semantic interpretation is central to the recent methods for proving lower bounds via the strategy extraction paradigm.

The linear ordering of a QBF’s quantifier prefix imposes a trivial “dependency structure” upon its variables which, unfortunately, identifies dependencies which are not essential for particular instances. A dependency scheme is an algorithm which attempts to identify such spurious dependencies directly from the syntactic

form of an instance; the results are used in state-of-the-art QBF solving to optimise performance. Since every unsuccessful run of a solver provides a proof of falsity, this raises questions about the underlying dependency proof systems: What is their proof complexity relative to other QBF calculi? For which dependency schemes are the underlying proof systems sound? Is strategy extraction possible? The suggestion also arises to implement dependency schemes in stronger calculi, for example in the QBF analogue of propositional Frege systems. Employing a semantic framework, the talk will present some new results, targeting an improved understanding of variable dependency in QBF calculi.

Michele Bottone (University of Middlesex)

The Agoric Process

Many phenomena can be viewed as systems arising from the interactions of agents, where an important aspect is computation, defined as the abstract representation of a process in terms of states and transitions. Such computational systems have the ability to share information and allocate resources and to parcel the computation in efficient ways through some form of signalling; they also occur in dynamically changing environments with asynchronous and unpredictable changes, including the possibility of new agents entering the system or leaving it. This basic infrastructure of open systems shares many similarities with a marketplace, where a collection of agents meet to exchange things of value to its participants. We use the term “agoric processes” to capture the mathematical essence of such information-processing mechanisms.

In this talk, I argue that the rise of connected and autonomous systems provides a useful backdrop both for experimentation and attaining a mathematical theory of distributed computation, and present a formalisation of the intuitive notion of an agoric process as a computational object living on some graph structure together with information flows. In this setting, information is an equivalence class of all ways of describing the same information possessed by agents, and one can use the Rota-Wallstrom theory of integration of functions indexed by set partitions to represent higher-level concepts.

Cassio P De Campos (Queen’s University Belfast)

Inferential Complexity in Probabilistic Graphical Models

Computations such as evaluating posterior probability distributions and finding joint value assignments with maximum posterior probability are of great importance in practical applications of probabilistic graphical models. These computations, however, are intractable in general, both when the results are computed exactly and when they are approximated. In order to successfully apply probabilistic graphical models in practical situations, it is crucial to understand what does and what does not make such computations hard. In this talk we guide the

audience through some of the most important computational complexity proofs and give insights about the boundary between tractable and intractable.

Gregory Chockler (Royal Holloway, University of London)

Space Bounds for Reliable Storage: Fundamental Limits of Coding

We study the space requirements of reliable storage algorithms in asynchronous distributed systems. A series of recent works have advocated using coding-based techniques (and in particular, erasure codes) as a way of reducing space overheads incurred by the standard replication approaches. However, a closer look reveals that they incur extra costs in certain scenarios. Specifically, if multiple clients access the storage concurrently, then existing asynchronous code-based algorithms may store a number of copies of the data that grows linearly with the number of concurrent clients. We establish a lower bound showing that this limitation is indeed inherent. We also present a reliable storage algorithm with matching space complexity thus proving that our bound is tight. Our algorithm is based on a new technique combining erasure codes with replication so as to obtain the best of both. I will start by introducing and motivating the problem, followed by an overview of the key ideas and techniques behind our results. (Joint work with Yuval Cassuto, Idit Keidar and Alexander Spiegelman.)

Bhaskar DasGupta (University of Illinois, USA)

Node Expansions and Cuts in Gromov-hyperbolic Graphs

Gromov-hyperbolic graphs (or, hyperbolic graphs for short) represent an interesting class of “non-expander” graphs. Originally conceived by Gromov in 1987 in a different context while studying fundamental groups of a Riemann surface, the hyperbolicity measure for graphs has recently been a quite popular measure in the network science community in quantifying “curvature” and “closeness to a tree topology” for a given network, and many real-world networks have been empirically observed to be hyperbolic.

In this talk, we give constructive non-trivial bounds on node expansions and cut-sizes for hyperbolic graphs, and show that witnesses for such non-expansion or cut-size can in fact be computed in polynomial time. We also provide some algorithmic consequences of these bounds and their related proof techniques for a few problems related to cuts and paths for hyperbolic graphs, such as the existence of a large family of s-t cuts with small number of cut-edges when s and t are at least logarithmically far apart, efficient approximation of hitting sets of size-constrained cuts, and a polynomial-time solution for a type of small-set expansion problem originally proposed by Arora, Barak and Steurer.

Colm Ó Dúnlain (Trinity College Dublin)

An almost-confluent congruential language which is not Church-Rosser con-

gruential

It is fairly easy to show that every regular set is an almost-confluent congruential language (ACCL), and Diekert et al (2015) showed that every regular set is a Church-Rosser congruential language (CRCL). The existence of an ACCL which is not a CRCL remained an open question. In this talk we present one such ACCL.

Marie Farrel (Maynooth University, Ireland)

A Logical Framework for Integrating Software Models via Refinement

Modern software development focuses on model-driven engineering: the construction, maintenance and integration of software models, ranging from formal design documents through to program code. We frequently model software at different levels of abstraction, starting with a very high level abstract specification and finishing with a detailed concrete implementation. In formal software engineering we can map between these levels of abstraction in a verifiable way through a process known as refinement. The question is how to combine information from models which focus on different aspects of the software system.

The theory of institutions observes that once the syntax and semantics of a formal system have been defined in a uniform way, a set of specification building operators can be defined that allow you to write, modularise and build up specifications that can be defined in a formalism-independent manner. Event-B is a formal specification language that enables the user to prove safety properties of a specification. It facilitates the modelling of systems at different levels of abstraction through the verifiable process of refinement. Our goal is to represent the Event-B formalism in terms of institutions and provide modularisation constructs which increase the scalability of Event-B for use in larger projects. A benefit of this approach is the increased interoperability of Event-B via institution comorphisms to allow aspects of the system to be specified in different formalisms and included in the final Event-B specification.

Andrew Healy (Maynooth University, Ireland)

Evaluating SMT solvers for software verification

SMT (Satisfiability Modulo Theories) solvers are an important component in deductive software verification systems. Such systems usually differ greatly in terms of specification language and approach to intermediate verification condition generation. This variety hinders the development of a common benchmark suite and makes the comparative evaluation of verification systems difficult. By using a large benchmark suite designed to test the capabilities of SMT solvers beyond software verification (in problem domains such as operations research and cryptography, for example), we aim to identify a subset of the large benchmark suite that can be used as a surrogate suite for software verification projects. We use

dynamic profiling to obtain a feature vector that characterises the workload of the solver before using techniques from cluster analysis to form new suites based on the observed behaviour of the solver under a verification workload.

We present the workflow and results of this process. A useful outcome will be a prediction of the most appropriate solver or combination of solvers to choose for a given verification problem. Such a model will show the correspondence of input to result in a way that would be useful to verification system and SMT tool developers as well as end users.

Ruth Hoffman (University of Glasgow)

Autonomous Agent Behaviour Modelled In PRISM

With the rising popularity of autonomous systems and their increased deployment within the public domain, ensuring the safety of these systems is crucial. Although testing is a necessary part in the process of deploying such systems, simulation and formal verification are key tools, especially at the early stages of design. Simulation allows us to view the continuous dynamics and monitor the behaviour of a system. On the other hand, formal verification of autonomous systems allows for a cheap, fast, and extensive way to check for safety and correct functionality of autonomous systems that is not possible using simulations alone. In this talk I demonstrate a simulation and the corresponding probabilistic model of an unmanned aerial vehicle (UAV) in an exemplary autonomous scenario and present results of both models. Further, I propose a definition of autonomy which can be used to model autonomous systems, followed by a discussion on how simulations inform probabilistic models.

Alison Jones (Swansea University)

Extracting Monadic Parsers from Proofs

My talk outlines a proof-theoretic approach to developing correct and terminating monadic parsers. Using modified realizability, we extract formally verified and terminating programs from formal proofs. By extracting both primitive parsers and parser combinators, it is ensured that all complex parsers built from these are also correct, complete and terminating for any input. We demonstrate the viability of our approach by means of two case studies: we extract (1) a small arithmetic calculator and (2) a non-deterministic natural language parser. The work is being carried out in the interactive proof system Minlog. (Joint work with Ulrich Berger and Monika Seisenberger)

Josh Lockhart (University College London)

An entanglement detection problem in a faulty quantum computer.

Quantum computers and algorithms promise to be a very disruptive technology. A key ingredient in algorithms that run on a quantum device is the ability for the

hardware components to be entangled with one another. If two quantum objects become entangled, then the simple act of checking the state of one object can instantaneously affect the state of the other. A great deal of effort has been expended studying the questions of what quantum entanglement really is, how to create it, and how to keep it around long enough to perform useful work. In this talk we outline new combinatorial techniques for verifying the existence of entanglement in a quantum computer. We show how graph theory can be used to think about a particular model of a faulty quantum computer, in an attempt to gain complexity theoretic insight into the problem of checking for the existence of entanglement. Specifically, we consider a restricted class of quantum states that can be represented by the combinatorial Laplacian matrix of a graph. This object encodes the states of a quantum computer that has suffered from a particular kind of error in its operation: the computer promised to construct a certain state but instead it has erroneously constructed one of a number of alternative states. We prove that our graph representation allows for certain well known entanglement criteria to be re-expressed in terms of the structure of the graph corresponding to the quantum state. Hence, the entanglement, or lack thereof, in the quantum computer after the fault can be tested via this purely combinatorial method.

David Kohan Marzagao (King's College London)

Flag Coordination Games and Random Walks

The main goal of this talk is to establish and explore a connection between flag coordination games and random walks. A Flag Coordination Game can be seen as graph colouring game where, in each round, each node colours itself based on an algorithm and on the range of visibility this node has. For example, assume you have 20 people in a circle playing a flag coordination game, each of them holding a red flag and a blue flag. Their goal is to reach a configuration where each of them is raising a different flag from both neighbours, i.e., to achieve a proper colouring of the graph. Their visibility is limited in that they can only see the flags raised by their immediate neighbours. The game starts with random flags and players follow an algorithm to decide which flag to raise at each new round. More precisely, if their neighbours are raising the same colour, choose the other colour, else, randomize between red and blue. What is the probability that they eventually reach their goal? What is the expected number of rounds for the game to finish? By proving an equivalence between such a game and a game of annihilating random walking particles, we can prove a theorem regarding the probabilities involved in the game, as well as an upper bound for the expected time it takes for the game to end. We can, then, generalize some of the results for graphs such that every vertex has degree 2.

Brett McLean (University College London)

The Finite Representation Property for Composition, Intersection, Domain and Range.

Motivated by modelling collections of deterministic programs, one might be interested in algebras isomorphic to a set of partial functions equipped with some set-theoretically-defined operations. We call such algebras representable and call the isomorphisms representations. The finite representation property holds for a signature of operations if any finite representable algebra can be represented using only a finite set as a base for the partial functions. I will describe a proof that the finite representation property holds for some of the most expressive signatures considered, containing the composition, intersection, domain and range operations. (Joint work with Szabolcs Mikulás.)

Reino Niskanen (University of Liverpool)

Undecidability of 2-dimensional Robot Games and Other Attacker-Defender Games

In this talk we consider simple two-player vector addition games, called robot games. In robot games, two players, Adam and Eve, are given sets of moves which they use to push a token on the integer lattice \mathbb{Z}^n starting from some initial point; Eve tries to push the token to the origin, whilst Adam tries to prevent this. The decision problem of the game is to determine which of the players has a strategy that guarantees victory. By constructing a game that simulates a 2-counter Minsky machine, we show that it is undecidable whether Eve has a winning strategy already when the game is played on a plane \mathbb{Z}^2 .

We also consider other two-player games, called Attacker-Defender games, that generalize robot games by having more complex state spaces and sets of moves, allowing or disallowing certain moves depending on the internal states of players. We consider games played on topological braids where Eve tries to unbraid a braid, or on vectors, where the moves are linear transformations rather than addition of vectors in robot games. We show that it is also undecidable whether Eve has a winning strategy in these games by constructing games that simulate one-counter automata on infinite words. (Joint work with Vesa Halava, Tero Harju, Igor Potapov and Julien Reichert).

Daniel Playfair (Queen's University Belfast)

Selection of optimal checkpoints from query plan graphs

Database fault tolerance is important for supporting critical business operations. The existing approach is to provide replicated fault tolerant data stores. Such solutions protect data and can offer availability. However, in the event of failures, work being performed at the time of execution is lost. Existing research into solving this problem and providing intra-query fault tolerance focuses on distributed or

row-oriented databases. Such solutions are not suitable for use with the column-oriented in-memory databases increasingly used for high-performance workloads.

A key requirement to providing intra-query fault tolerance is the ability to devise an efficient checkpoint plan for a given query plan. We introduce and discuss a general model for reasoning about query plans. We make observations about the location of worst case queries in such plans, and introduce algorithms for calculating the worst case execution time of checkpointed queries. We also discuss aspects of algorithms for producing checkpoint plans within such a model.

Craig Reilly (University of Glasgow)

Enumeration of knots using constraint programming

This presentation details a novel approach, using constraint programming, to the problem of enumerating knot diagrams. Our enumeration makes use of Gauss code representations of knot diagrams, which leads to an obvious modelling as a constraint satisfaction problem. However, a Gauss code may not always represent a knot diagram (they might represent a virtual knot) and we show that we can add constraints to our model to disallow virtual knots. Further, each knot diagram can be represented by many Gauss codes, and we explain new ways of modelling the enumeration with a view towards symmetry breaking.

Latif Salum (Dokuz Eylul University, Turkey)

The Tractability of Un/Satisfiability

A safe acyclic Petri net (PN) is associated with some Exactly-1 3SAT formula $\phi = c_1 \wedge c_2 \wedge \dots \wedge c_m$, in which a clause $c_k = (z_i \dot{\vee} z_j \dot{\vee} z_u)$ is an exactly-1 disjunction $\dot{\vee}$, rather than disjunction \vee , of three literals: c_k is true exactly when only one of z_i or z_j or z_u is true. Some 2SAT/XOR-SAT formula arising in the *inversed* PN checks if the truth assignment of a literal (a transition firing) z_v is “incompatible” for the satisfiability of the 3SAT formula (the reachability of the target state in the *inversed* PN). If z_v is incompatible, then z_v is discarded and \bar{z}_v becomes true. Therefore, a clause $(\bar{z}_v \dot{\vee} z_i \dot{\vee} z_j)$ *reduces* to the conjunction $(\bar{z}_v \wedge \bar{z}_i \wedge \bar{z}_j)$, and a 3-literal clause $(z_v \dot{\vee} z_u \dot{\vee} z_x)$ *reduces* to the 2-literal clause $(z_u \oplus z_x)$. This reduction facilitates checking (un)satisfiability; the 3SAT formula is (un)satisfiable iff the target state of the *inversed* PN is (un)reachable. The solution complexity is $O(n^5)$. Therefore, it is the case that $\mathcal{P} = \mathcal{NP} = \text{co-}\mathcal{NP}$.

Chhaya Trehan (Queen’s University Belfast)

Energy consumption of parallel workloads: convexity, parallelism and memory intensity

Performance and energy consumption are important and contradicting design criteria for modern multicore processors. Optimizing one whilst imposing a threshold on the other leads to two flavors of optimization: in the laptop problem, the

goal is to maximize performance given a fixed energy budget; in the server problem, the goal is to minimize energy consumption given a fixed performance budget. We deal with the server problem in this talk. Energy and performance of a parallel application running on a multicore chip are convex functions of its varying operating frequency. Analytical frameworks have been built that exploit this in order to tune the operating frequency of the processor. However, existing theoretical work on energy minimization using frequency tuning ignores the time and energy consumed by the processor whilst it waits to access the data it requires from the main memory. We present a new energy-performance model which accounts for the time and energy consumed by a processor on memory accesses in addition to the time and energy consumed on actual CPU instructions. We show that the problem of energy minimization under a performance budget remains a convex optimization problem in the new model. We also investigate how the optimal operating points (frequencies) in our model differ from the operating points in the model that does not account for the energy performance overhead of memory accesses. Finally, we show the relationship between the optimal frequencies and energy-aware scheduling of an application.