

REPORT ON BCTCS 2008

The 24th British Colloquium for Theoretical Computer Science

7-10 April 2008, Grey College, Durham University

Hajo Broersma, Tom Friedetzky, Daniël Paulusma

The British Colloquium for Theoretical Computer Science (BCTCS) is an annual forum for researchers in theoretical computer science to 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 benefit from contact with established researchers.

BCTCS 2008 was hosted by Durham University, and held at Grey College, one of the colleges of Durham University, during 7–10 April 2008. The event attracted 84 participants, and featured an interesting and wide-ranging programme of 7 invited talks and 36 contributed talks; the majority of the contributed talks were given by PhD students. Abstracts for the talks are provided below; further details are available via the BCTCS 2008 website at www.durham.ac.uk/bctcs.2008. The financial support of the Engineering and Physical Sciences Research Council (EPSRC) and the London Mathematical Society (LMS) is gratefully acknowledged.

BCTCS 2009 will be hosted by the University of Warwick from 6–9 April 2009. Researchers and PhD students wishing to contribute talks concerning any aspect of theoretical computer science are warmly welcomed to do so. Further details are available from the BCTCS website at www.bctcs.ac.uk.

Invited Talks at BCTCS 2008

Martín Abadi, Microsoft Research Silicon Valley, USA and University of California, Santa Cruz, U.S.A.

Towards Correct Programming with Transactional Memory

Several recent designs and systems based on transactions aim to facilitate the writing of concurrent programs. In particular, software transactional memory (STM) appears as an intriguing alternative to locks and related constructs for shared-memory concurrency. In this context, we are exploring a model that we call Automatic Mutual Exclusion (AME), in which transactional execution is the default. In this talk we present AME and the underlying semantics and implementation strategies. We also discuss the programming disciplines (various forms of separation) that are required for ensuring the correctness of those implementation strategies.

José Félix Costa, Technical University of Lisbon (IST) and Swansea University

Physics and Computation: An Essay on the Unity of Science through Computability

In the thirties, Alan Turing introduced the Turing Machine to underly a proof of existing bounds on Hilbert's Research Programme. Although considered a foundation of Computer Science, the Turing Machine and other equivalent abstract devices (invented since the time of Turing, Post and Markov) are far from physically realizable. To make things worse, the Church-Turing Thesis – a statement with metamathematical content associated with the Turing Machine – has acquired physical meaning in the last two decades, due mainly to Kreisel. In this lecture, I will address the problem of the physical implementation of a Turing Machine and some deep consequences to Physics and other Sciences (eg, Economics, Chemistry), as well as the problem as to whether or not the Church-Turing Thesis can be considered a refutation tool for physical theories in the sense of Popper.

Artur Czumaj, University of Warwick

Sublinear-time Algorithms

We are often confronted with a huge amount of information that is generated by large scale complex information systems that cannot even be stored in their entirety. Examples include the World Wide Web as a source of information, and performance measurements in networks. In many of these applications, polynomial-time algorithms that are efficient for relatively small inputs may become impractical for input sizes of several gigabytes. Managing and analyzing such data sets forces us to revisit the traditional notions of efficient algorithms. For example, when we consider approximation algorithms for clustering problems in metric spaces with n points, since their input size is $\Theta(n^2)$, they typically have $\Omega(n^2)$ running time. Clearly, such a running time is not feasible for massive data sets. Constructing a sublinear-time algorithm may seem to be an impossible task since it allows one to read only a small fraction of the input. However, in recent years, we have seen the development of sublinear-time algorithms for optimization problems arising in such diverse areas as graph theory, geometry, algebraic computations, and computer graphics.

In this talk, we present a few examples of sublinear-time algorithms. Our main focus will be on techniques used to design sublinear-time algorithms to estimate basic graph parameters (e.g., the average degree of the graph or the cost of its minimum spanning tree). We also discuss the framework of *property testing*, an alternative notion of approximation for decision problems, which has been applied to give sublinear algorithms for a wide variety of problems.

Martin C. Henson, University of Essex

The Bulletin of the EATCS

Varieties of Schema Calculus

Z is a widely-used industrial-strength specification language whose schema calculus comprises an *equational logic* and schema operators that are *non-monotonic* with respect to refinement. In this talk we demonstrate that this is not the only way (we will call it the *first embedding*) in which Z might have been defined. The equational logic, and the lack of monotonicity, result from a choice between two alternative mathematical descriptions of the schema calculus that differ only in the permutation of two semantical aspects: one concerning the schema operators themselves, and the other concerning refinement. The alternative permutation (the *second embedding*) leads to an *inequational logic* and schema operators that are *monotonic* with respect to refinement. The second embedding is equivalent to (a restriction of) the specification logic νZ and the first embedding is a retract of νZ .

In this talk we will consider the following questions. First, why is the Z schema calculus non-monotonic, whereas the schema calculi of νZ , Hoare's and He's UTP designs, and Henson's and Reeves's *sets-of-implementations* model *all* monotonic? Second, why is there *no* refinement semantics based on Z's usual partial relation interpretation of the schema calculus that leads to a monotonic refinement calculus?

Leonid Libkin, University of Edinburgh

Databases Meet Verification, or Nested Words and XML Documents

The main computational tasks in the fields of databases and verification can be summarized as the evaluation of logical formulae on finite (or sometimes infinite but finitely-presented) structures. And yet the differences in logics and types of structures have kept these fields further apart than they should be. This is changing now, partly due to more flexible data formats, such as XML, taking the central role in database research. Navigation over XML documents closely resembles temporal properties used in software verification, which bridges the gap between tools and techniques used in the two fields.

In this talk, I explain how (and why) the fields of databases and verification developed largely independent of each other, and give examples of data processing problems where model-checking techniques could be used. I then concentrate on a particular case of this, related to nested words, which present a nice abstraction for reasoning about programs with recursive procedure calls. I show how results about the XML navigation language XPath can be applied in this context to get expressively complete logics for nested words with good model-checking properties.

The second part of the talk is based on a joint paper with Alur, Arenas, Barceló, Etessami, and Immerman from LICS 2007.

Rolf Niedermeier, Friedrich-Schiller-Universität, Jena, Germany

Trends in Parameterized Algorithmics

Parameterized algorithmics is a new and compelling methodology for the design and analysis of algorithms for computationally hard (typically NP-hard) problems. The mission is to lift the typically “one-dimensional” view on problems usually taken in classical algorithmics to a “multi-dimensional” (or multivariate) one. To understand the intrinsic difficulties of hard computational problems, it seems to be adequate and fruitful to analyze problems under various perspectives, that is, different parameterizations. This is what the field is heading for: a better understanding of the nature of computational complexity and, correspondingly, the detection of new ways and tools for the development of practically useful algorithms by employing specific properties of problem instances. These properties can best be detected using the mathematically rigorous machinery of parameterized complexity analysis. It is our ultimate goal to establish parameterized algorithmics as a practically relevant and scientifically versatile “competitor” of approximation algorithmics. In this line, we here provide a survey on some recent research trends in parameterized algorithmics with some theoretical breakthroughs as well as concrete practical applications of the developed algorithms, having a particular focus on results achieved in the Jena research group.

Gerhard J. Woeginger, Technical University Eindhoven, The Netherlands

Three Assignment Problems and One Theorem

(LMS Keynote Lecture in Discrete Mathematics)

In this talk we discuss three special cases of the quadratic assignment problem. Firstly, a data management problem on the arrangement of data records with given access probabilities in a linear storage medium. Secondly, a sequencing problem on circular arrangements of disks and other objects. Thirdly, a weight balancing problem that deals with the assignment of masses to the vertices of a regular polygon. We analyze the behavior of these three problems, derive some of their combinatorial properties, and formulate one unifying theorem.

Contributed Talks at BCTCS 2008

Timos Antonopoulos, University of Cambridge

On the Expressive Power of Graph Logic

We present results on the expressive power of Graph Logic (GL), a spatial logic for querying graphs introduced by Cardelli et al. (2002). GL is an extension of First Order Logic with a second order quantifier over edges of restricted form, and a first order quantifier over labels of edges. In particular, if G is some graph, all one can do with the second order quantifier is express that there exists a set of

edges X of the graph G such that some formula ϕ is satisfied by the subgraph of G containing exactly the edges in X , and some formula ψ is satisfied by the subgraph of G with exactly the remaining edges not in X .

It has been observed that GL is a sublogic of Monadic Second Order Logic with quantification over edges (MSO for short), and although it seems that GL is strictly less expressive than MSO, many interesting properties have been shown to be expressible in GL. Furthermore it was shown by Dawar et al. (2007) that GL is able to express complete problems on any level of the Polynomial Hierarchy and that GL and MSO are equi-expressive when restricted to words. Marcinkowski (2006) showed that this richer form of MSO is more expressive than GL.

As this richer form of MSO is not the one we usually deal with, the case where we omit the first order quantification over edges from both logics is a more interesting one. We show that this restriction of GL is indeed strictly less expressive than the one of MSO. Moreover we show that this is the case even when restricted to the class of forests.

This is joint work with Anuj Dawar.

Haris Aziz, University of Warwick

Variations in weighted voting games

Weighted voting games model situations where agents with variable voting weight vote in favour of or against a decision. A coalition of agents is winning if and only if the sum of weights of the coalition exceeds or equals a specified quota. *Tolerance* and *amplitude* of a weighted voting game signify the possible variations in a weighted voting game which still keep the game unchanged. We characterize the complexity of computing the tolerance and amplitude of weighted voting games. We give tighter bounds and results for the tolerance and amplitude of key weighted voting games. We then provide limits to how much the Banzhaf index of a player increases or decreases if it splits up into sub-players.

This is joint work with Mike Paterson.

Dénes Bisztray, University of Leicester

Verification of Architectural Refactoring Steps by Rule Extraction

With the success of model-driven development as well as component-based and service-oriented systems, models of software architecture are key artefacts in the software development environments and requirements and improve internal software quality process. To adapt to changes such models have to evolve while preserving aspects of their behaviour. Such behaviour-preserving evolution steps, called refactorings, concern both the structure of the system (its graph of components and connectors) as well as the behaviour of the components themselves.

To avoid the resource-consuming verification of the complete system, we extract refactoring rules from the transformations performed and verify these rules.

Our main result shows that, under certain assumptions, the verification of these rules is sufficient to guarantee the preservation of the behaviour of the source in the target model. In this way, we only verify a small subgraph of the model.

We apply the approach to architectural models using UML component, structure, and activity diagrams, using CSP as a semantic domain.

Clive Blackwell, Royal Holloway, University of London
Reasoning about Cryptographic Protection with Spygraphs

Bigraphs define a process algebra based on category theory. A bigraph is composed of a link and place graph with the same nodes but different edges. The link graph represents logical communication (as in the pi calculus) and the place graph models physical locality (as in the ambient calculus). Spygraphs are a novel extension of bigraphs with additional reaction rules to model cryptographic and other security mechanisms, which is analogous to the extension of the pi-calculus to the spi-calculus.

A system and its users are modelled as spygraphs, where spygraph reaction or rewriting rules private to the defender model the possible security mechanisms, and inverse rules model the access rights of users. They can model attackers with varying rights and abilities including insiders, and multilayer attacks using both physical and logical actions. The defender's objectives can be defined by invariants of the spygraph that represent secure states of the system that only allow authorised actions by legitimate users.

In this talk, we propose operational semantics for the basic cryptographic operations and discuss useful equivalences to help in reasoning about cryptographic protection. The advantages of spygraphs compared to most other formal methods are the explicit representation of both the system structure and physical locality that is fundamental to modelling system security realistically. Spygraphs model the requirements that cryptographic keys must be stored in physically inaccessible locations, which control access to protected resources using private communication channels. We demonstrate the use of spygraphs to model boundary protection by firewalls to control remote access and mobile code. We are investigating the extension of spygraphs to model other security mechanisms so that they can represent complete systems rather than just single components like cryptographic protocols handled by most current formal methods.

William Blum, Oxford University
A Concrete Presentation of Game Semantics

We briefly present a new representation theory for game semantics which is very concrete: instead of playing in an arena game in which P plays the innocent strategy given by a term, the same game is played out over (a souped up version of) the abstract syntax tree of the term itself. The plays that are thus traced out are called

The Bulletin of the EATCS

traversals. More abstractly, traversals are the justified sequences that are obtained by performing parallel-composition *less* the hiding. After stating and explaining a number of Path-Traversal Correspondence Theorems, we present a tool for game semantics based on the new representation.

This is joint work with Luke Ong.

Magnus Bordewich, Durham University

Path Coupling Without Contraction

Path coupling is a useful technique for simplifying the analysis of a coupling of a Markov chain. Rather than defining and analysing the coupling on every pair in $\Omega \times \Omega$, where Ω is the state space of the Markov chain, analysis is done on a subset $S \subseteq \Omega \times \Omega$. If the coefficient of contraction β is strictly less than one, no further analysis is needed in order to show rapid mixing. However, if $\beta=1$ then analysis (of the variance) is still required for all pairs in $\Omega \times \Omega$. We present a new approach which shows rapid mixing in the case $\beta=1$ with a further condition which only needs to be checked for pairs in S , greatly simplifying the work involved. We also present a technique applicable when $\beta=1$ and our condition is not met.

This is joint work with Martin Dyer.

Diana Fulger, University of Nottingham

On Reasoning about Effects: Seeing the Wood through the Trees

Monads are widely used in the writing of computations with side effects, as a powerful tool allowing for more concise and intuitive expressions. One might expect reasoning about them to be just as straightforward, but the evidence in support of this claim is meagre so far. We try to address the issue by studying a plain example: the labelling of a tree with different labels. We take a simple algorithm using a modifiable state, and prove it correct using elementary techniques in a first go, and monads/applicative functors the second time around. We thus obtain a factorisation of the labelling routine, the proof of which suggests taking some liberty on the type system (in fact justified by notational abuse), so as to describe some quite intriguing properties of the state applicative functor in interaction with function composition.

This is joint work with Graham Hutton.

Vashti Galpin, University of Edinburgh

Equivalences for Hybrid Systems

HYPE is a process algebra for hybrid systems. It models the flows that affect the continuous aspects of the system and allows for discrete events to occur which may change the flows. The operational semantics define a transition system labelled with events, and the states of this transition system can be used to obtain the ordinary differential equations (ODEs) which describe the continuous change

of the hybrid system. For well-defined HYPE models, system-bisimilar models have the same ODEs. Semantic equivalences such as bisimilarity have also been defined for other process algebras for hybrid systems and these are considered in the context of HYPE together with new equivalences for HYPE.

This is joint work with Jane Hillston and Luca Bortolussi.

Manish Gaur, University of Sussex

A Routing Calculus for Distributed Computing

We study the modeling of a distributed network with *routers* acting as an active component in determining the quality of service of the network. Our model may be considered as an extension of Distributed Pi-Calculus (Dpi). We provide two models of distributed networks. In both models we describe distributed computations explicitly in the presence of routers in an Internet-like network. A site for computational activity consists of named routers which host computational entities called *nodes*. All the nodes are directly connected to some specific router. These sites run in parallel to form a large distributed network called a *system*. The communication between processes of any two nodes in this network is possible only through their respective routers. However these models differ in method of updating the knowledge of routers about the newly created computing agents. In the first model, knowledge of only those routers are updated which are on the path of communication of a newly created agent whereas in the second we update the knowledge of all the routers in the network. We describe these models in two separate languages. Further in each language we define a configuration consisting of router connectivity and a system. For the consistent behaviour of a configuration we define a different set of conditions on well formed configurations in each language. We prove that well formedness of configurations in both languages is preserved under reductions. We then describe a specification language similar to Dpi and show that, after abstracting away the details of routers and paths, both routing languages are reduction equivalent to Dpi. We also define a notion of precongruence based on the reduction semantics of well formed configurations so that we can relate them with a preorder relation. Finally, with an intention to define a pre-bisimulation between the well formed configurations, we define a labeled transition system for both the languages, and aim to recover the preorder relation based on reduction semantics from the pre-bisimulation defined using labeled transition systems and vice versa.

This is joint work with Matthew Hennessy.

Matthew Hague, Oxford University

Winning Regions of Pushdown Parity Games: A Saturation Method

After surveying the use of pushdown systems in software model-checking, we present a new algorithm for computing the winning regions of a parity game

The Bulletin of the EATCS

played over a pushdown system. This method extends the saturation techniques introduced by Bouajjani, Esparza and Maler and independently by Finkel, Willems and Wolper.

Previous solutions to this problem, due to Serre and independently to Cachat, use an exponential reduction to a finite-state game. Our algorithm provides the first extension of saturation techniques to parity games. We use finite word automata to represent sets of pushdown configurations, which consist of a control state and a stack of characters from a finite alphabet. We begin with a small automaton, and, using a fixed-point characterisation of the winning regions, expand the automaton until the algorithm is complete. Hence, in some fortunate cases, the exponential blow-up may be avoided.

Using Hague and Ong's saturation-based algorithm for backwards reachability analysis of higher-order pushdown systems, this technique generalises naturally to the higher-order case.

This is joint work with C.-H. Luke Ong.

Pim van 't Hof, Durham University

A new characterization of P_6 -free graphs

We study P_6 -free graphs, i.e., graphs that do not contain an induced subgraph isomorphic to a path on six vertices. Our main result is a new characterization of this graph class: a non-trivial graph G is P_6 -free if and only if each connected induced subgraph of G has a dominating induced cycle on six vertices or a dominating (not necessarily induced) complete bipartite subgraph. Our characterization of P_6 -free graphs strengthens results of Liu and Zhou, and of Liu, Peng and Zhao. Our proof has the extra advantage of being constructive: we present an algorithm that finds such a dominating subgraph of a connected P_6 -free graph in polynomial time. We present some applications of our results to special cases of the Hamiltonian cycle problem and the Dominating Cycle problem.

This is joint work with Daniël Paulusma.

Temesghen Kahsai, Swansea University

Theory and Application of Testing from CSP-CASL

At BCTCS 2007 we investigated how to develop a theory for the evaluation of test cases with respect to formal specification. We used the specification language C^2 to define and evaluate black-box tests for reactive systems. We can now present a consolidated theory and a framework for which we also have developed a tool support. One of the major innovation of this approach is the separation of the test oracle and the test evaluation by defining the expected result in terms of *test colouring* and the verdict of a coloured test case in terms of *test execution*. Based on characterization theorems, we can use C^2 -prover in order to color test cases. Furthermore, we have implemented a test environment which evalu-

ates tests on a system under test. As work in progress we present a new testing framework based on C -C for a product line (*horizontal refinement*).

Karim Kanso, Swansea University

Formal Verification of Safety Properties in Systems Defined in Ladder Logic

Ladder logic is a representation of propositional logic that is used for programming *programmable logic controllers*. The ladder can then be interpreted as a Boolean program such that we can reason about various properties of the program such as safety and liveness. The Boolean programs that are of interest are ones which are repeated indefinitely such as those running on a railway interlocking or other critical control system. To verify the safety conditions for this program two steps are needed, firstly from the initial state, the conditions must be satisfied. Secondly, from an arbitrary state where the safety conditions hold, then one iteration of the program the conditions should still hold. Hoare Logic is suited to this purpose as repeated Boolean programs are reasoned about. Hoare logic demands that there are pre-conditions, loop invariants and post-conditions.

Matthew Lakin, University of Cambridge

Animating Inductive Definitions with Binders

This work concerns the automatic derivation of executable prototype systems for the animation of programming languages and calculi, given just their specification in terms of inference rules. By “animating” inductive definitions, we mean performing proof search over the relation defined by the rules in such a way that the binding structure of terms in the object-language is respected. Relations of this form are usually defined as the *least set* which *contains all of the axioms* and is *closed under the rules*. The rules involve schematic patterns which we somehow *instantiate* to produce the underlying set of mathematical objects.

In the straightforward setting without binders, the notion of instantiation is trivial. However, when our schematic patterns contain terms involving *binders*, it is by no means obvious what the “correct” mode of instantiation should be. For example, consider the schematic pattern $\lambda x.\lambda y.(Var\ x)$, where x and y are schematic pattern variables. Given *distinct, concrete names* a and b , are either of $\lambda a.\lambda b.(Var\ a)$ or $\lambda a.\lambda a.(Var\ a)$ valid instantiations of this pattern? This is largely a matter of personal taste—almost everyone would allow the first, but some would also permit the second.

We express the implementations of our schematic inductive definitions in the MLSOS metalanguage, an extension of the FreshML functional programming language. Object-language binding is handled using the nominal techniques developed by Pitts, which provide an elegant treatment of bound names while staying close to informal pen-and-paper syntax. The language features are specifically tailored to this application domain, and include constructs for generating fresh *atoms*

The Bulletin of the EATCS

(which represent bound names in the object-language) and *unification variables* (which stand for unknown object-language terms). There is also a sublanguage of constraints and a branching operator which can be used to simulate proof-search behaviour.

In this talk I discuss the language of schematic inductive definitions and the issues surrounding their implementation in MLSOS. I will also describe work on correctness proofs for this translation.

This is joint work with Andrew Pitts.

Ioannis Lignos, Durham University

Time- and Space-Efficient Periodic Exploration of Undirected Graphs

Efficient graph exploration is one of the core problems in algorithmic graph theory. Here we consider construction of a tour that visits all n nodes of an arbitrary undirected graph. The tour is traversed by a mobile entity in a periodic manner. We assume that the nodes in the graph are unlabeled. However, for each node v the endpoints of the edges incident to v are uniquely identified by labels called *port numbers*. In this setting, the consecutive edges of the tour are determined with the help of appropriate arrangements of port numbers at the nodes of the graph supported by a traversal mechanism stored in the agent's limited memory. We present a new construction of a tour of length $3.5n$ which subsumes the bound $3.75n$ established recently. The improvement is possible due to a novel concept of three-layer decomposition of undirected graphs. We believe that this new combinatorial construction may well be of independent interest.

This work is part of a larger project, which includes a no-memory solution, involving J. Czyzowicz, S. Dobrev, L. Gasieniec, D. Ilcinkas, J. Jansson, R. Klasing, I. Lignos, R. Martin, K. Sadakane and W.-K. Sung.

Gerald Luetzgen, University of York

Safe Reasoning with Logic LTS

This talk presents the concurrency-theoretic framework of *Logic LTS* which allows one to truly mix operational and logic styles of specification. Heterogeneous methodologies supporting multi-paradigm specifications have strong practical motivations, and design notations such as UML already provide superficial support for them.

Logic LTS adds *conjunction* as well as disjunction operators to the standard setting of labelled transition systems with CSP-style parallel composition. Within this framework, *ready simulation* is shown to be fully abstract with respect to failures inclusion. Ready simulation also satisfies standard logic properties, facilitates the inclusion of temporal logic operators and thus lends itself to studying mixed operational and logic languages.

This is joint work with Walter Vogler.

David Manlove, University of Glasgow

Size Versus Stability in the Marriage Problem

Given an instance I of the classical Stable Marriage problem with Incomplete preference lists (), a maximum cardinality matching can be larger than a stable matching. In many large-scale applications of , we seek to match as many agents as possible. This motivates the problem of finding a maximum cardinality matching in I that admits the smallest number of blocking pairs (so is “as stable as possible”). We show that this problem is NP-hard and not approximable within $n^{1-\varepsilon}$, for any $\varepsilon > 0$, unless $P=NP$, where n is the number of men in I . Further, even if all preference lists are of length at most 3, we show that the problem remains NP-hard and not approximable within δ , for some $\delta > 0$. By contrast, we give a polynomial-time algorithm for the case where the preference lists on one side are of length at most 2.

This is joint work with Péter Biró and Shubham Mittal.

Luke Mathieson, Durham University

New Results on The Complexity of Regular Subgraph Editing Problems

Deciding whether a given graph has a regular subgraph (of degree at least 3) is a well-studied problem in the field of classical computational complexity. The problem is in general NP-complete, and remains so even when the input is restricted to bipartite graphs of degree at most $r+1$, where r is the desired regularity of the subgraph. We study the problem in the framework of parameterized complexity, and ask if we can obtain a regular subgraph of a given graph by means of at most k editing steps (vertex deletion, edge deletion, edge addition). Moser and Thilikos have shown that the problem is fixed-parameter tractable when only vertex deletion is available, parameterized by both the number k of deletions and the desired regularity r .

In previous work we showed that we can maintain fixed-parameter tractability if both vertex and edge deletions are available operations. By using a generalisation in the form of annotations (each vertex gets annotated with its own desired degree), we obtain better algorithms than without. We also have shown that, when parameterized by the number k of operations alone (and not the regularity r), the problems are $W[1]$ -hard and thus unlikely to be fixed-parameter tractable. Here, we show that when the further editing operation of edge addition is available, the problems remain fixed-parameter tractable (again, parameterized by the number of editing steps and the regularity), even for the annotated version. We also demonstrate that the problem is solvable in polynomial time if only edge addition is available.

This is joint work with Stefan Szeider.

Divina A. Melomey, University of East London

The Bulletin of the EATCS

Mobility Challenges in Online Application Development

The need for mobilization and mobility in online applications has become paramount in ensuring that client and customers' needs are met at the point of service. This work validates detail mobility requirements for such applications. This validation covers areas such as e-health, online banking and virtual learning environment platforms. Further investigations with regards to these requirements explored mobility in games. Initial work to date has highlighted generic mobility requirements derived from distributed platforms. These requirements implementation validates parameters such as persistency, concurrency, message passing and calling, synchronisation, transparency, serialization and remote invocation in the application environment mentioned. This talk explores requirements needed for development of online applications and how mobility requirements are met. These requirements are then validated against existing online applications with case analysis. This validation process is crucial for assessing mobility of online applications for clients or customers.

Faron Moller, Swansea University *On the Complexity of Parity Games*

Parity Games underlie the model-checking problem for the modal μ -calculus, the complexity of which remains unresolved after more than two decades of intensive research. The community is split into those who believe that this problem – which is known to be both in NP and co-NP – has a polynomial-time solution (without the assumption that $P=NP$) and those who believe that it does not. (A third, pessimistic, faction believes that the answer to this question, along with that of $P=NP$, will remain unknown in our lifetime.)

In this talk we explore the possibility of employing Bounded Arithmetic to resolve this question, motivated by the fact that problems which are both NP and co-NP and which can be formulated within a certain fragment of Bounded Arithmetic necessarily admit a polynomial-time solution. While the problem remains unresolved in this talk, we do proposed this line of attack, and at the very least provide a modest refinement to the complexity of parity games (and in turn the μ -calculus model checking problem): that they lie in the class PLS of Polynomial-time Local Search problems. This result is based on a new proof of memoryless determinacy which can be formalised in Bounded Arithmetic.

The approach we propose may offer a route to a polynomial-time solution. Alternatively, there may be scope in devising a reduction of the problem to some other problem which is hard wrt PLS, thus making the discovery of a polynomial-time solution unlikely according to current wisdom.

This is joint work with Arnold Beckmann.

Mark New, Swansea University

Independent Typing Rules for Basic Programming Constructs

When giving typing rules for a programming language, the formulation of the rules for each construct usually depends on which other constructs are included in the described language. This dependency hinders reuse of typing rules for common constructs in descriptions of different languages.

This talk will show that by using MSOS (a recently-developed modular variant of structural operational semantics), typing rules for individual constructs can be given completely independently. Our illustrative examples include basic constructs involving functions, records, and variants, and deal with both types and subtypes. Apart from being independent and potentially reusable, our typing rules are often notationally simpler than in previous approaches.

This is joint work with Peter Mosses.

Liam O'Reilly, Swansea University

Algorithmic Proof Generation for CSP-CASL-Prover

At BCTCS 2007 we suggested an architecture for a theorem prover for the language $C-C$ based on Isabelle/HOL. $C-C$ integrates the process algebra C with the algebraic specification language C , forming a specification language tailored to the description of distributed systems. We have since realised this architecture to the point that we have algorithms which produce the theorems as well as the proof scripts needed in $C-C$ -Prover. Tests in various scenarios including a proper industrial setting demonstrate that our architecture is feasible.

In this talk we will report these results and discuss alternatives to our original implementation strategy, namely instead of starting from a shallow encoding of CASL in Isabelle/HOL to start from a semi-deep encoding, or even a deep encoding of C .

Grant Olney Passmore, University of Edinburgh

Open Euclidean Relations and Decidable Fragments of the True Existential Theory of the Rational Number Field

Hilbert's Tenth Problem over \mathbb{Q} , the question as to the decidability of the true existential first-order theory of the rational number field, is an important open problem in arithmetic algebraic geometry. Its solution, be it negative (as is generally expected) or positive, will have important implications for many areas of mathematics and theoretical computer science. As such, it is receiving much attention: Shapiro and Shlapentokh have recently obtained the striking result that the true existential first-order theory of the ring of integers of any algebraic number field whose Galois group over the rationals is abelian is undecidable. In addition, Poonen has recently obtained a clever simplification of Julia Robinson's $\forall\exists$ definition \mathbb{Z} over \mathbb{Q} to a $\forall\exists$ definition, and there is some hope that eventually

The Bulletin of the EATCS

a purely existential definition may be found.

Our work concerns research in the other direction: We hope to isolate decidable nonlinear fragments of the true existential first-order theory of the rational number field. We have obtained a number of elementary positive results that depend upon two crucial notions: First, we call a continuous map $f : \mathbb{R}^n \rightarrow \mathbb{R}$ a “real-extended rational endomap” (*rcr*-map) if $\{f(\vec{q}) \mid \vec{q} \in \mathbb{Q}^n\} \subseteq \mathbb{Q}$. Second, we call a binary relation $R : \mathbb{R}^2 \rightarrow 2$ an “open Euclidean relation” (*oe*-relation) if for all *rcr*-maps $f, g : \mathbb{R}^n \rightarrow \mathbb{R}$, $\{\vec{r} \in \mathbb{R}^n \mid R(f(\vec{r}), g(\vec{r}))\}$ is open in \mathbb{R}^n under the Euclidean topology. We have shown that over any first-order language \mathcal{L} with relation symbols $(R_i)_{i \in I}$, function symbols $(f_j)_{j \in J}$, logical symbols $\{\exists, \vee, \wedge\}$, and constant symbols $(c_k)_{k \in K}$, $\text{Th}_{\mathcal{L}}(\langle \mathbb{Q}, (f_j)_{j \in J}, (R_i)_{i \in I}, (c_k)_{k \in K} \rangle)$ is decidable provided that each c_k denotes a rational number and each R_i (resp. f_j) denotes an *oe*-relation (resp. *rcr*-map) that is first-order definable in the full language of ordered rings.

In this talk we present background motivation on Hilbert’s Tenth Problem over \mathbb{Q} , present our decidability results and examples of their use with interesting *rcr*-maps and *oe*-relations, and discuss future directions for our research.

Ondrej Rypacek, University of Nottingham

The Expression Lemma

Algebraic data types and catamorphisms (folds) play a central role in functional programming as they allow programmers to define recursive data structures and operations on them uniformly by structural recursion. Likewise, in object-oriented (OO) programming, recursive hierarchies of object types with virtual methods play a central role for the same reason. There is a semantical correspondence between these two situations which we reveal and formalize categorically. To this end, we assume a coalgebraic model of OO programming with functional objects. In practical terms, the development prepares for refactorings that turn sufficiently disciplined functional folds into OO programs of a designated shape (and v.v.).

This is joint work with Ralf Lämmel.

András Z. Salamon, Oxford University

Complete Constraint Satisfaction Problems

Constraint satisfaction problems (CSPs) form a large subclass of NP, capturing decision problems that ask whether there exists a homomorphism between a source and a target relational structure. GRAPH COLOURING is an NP-complete CSP, 2-SAT is a CSP that can be solved in polynomial time, and GRAPH ISOMORPHISM is a CSP that is commonly believed to have intermediate complexity. A complete structure of arity r is a relational structure that can be thought of as the non-redundant part of an r -dimensional adjacency matrix. If all relations in a CSP have the same arity, then the CSP can be transformed to a representation with a complete source structure. This talk deals with such representations of CSPs, and

how they relate to a conjecture of Grohe.

This is joint work with Peter Jeavons.

Rafiq Saleh, University of Liverpool

Automata on Gauss Words

Knots which are defined as embeddings of a circle in 3-dimensional Euclidean space can be faithfully represented by finite structures, such as graphs or words. One such discrete representations is a Gauss code, which is a word of crossing labels O ("over") and U ("under") with appropriate indices. The computational problems of recognising correct Gauss codes, as well as recognising many properties of knots presented by Gauss codes such as virtuality and unknottedness, are known to be decidable, with different time complexity.

In this talk I will discuss and present lower and upper bounds on the computational power needed to recognize the above properties. Due to the fact that the number of crossing in knots is unbounded, the Gauss words can be seen as strings over an infinite alphabet. In this context we are analysing the computational power in terms of a model of a finite state register automata over an infinite alphabet, introduced and studied by Kaminski et al. and Schwentick et al. Our results are mainly related to several variants of such automata including deterministic, non-deterministic, one-way or two-way register automata.

This is joint work with Alexei Lisitsa and Igor Potapov.

Stephan Scheele, University of Bamberg, Germany

cALC: Towards Constructive DL for Abstraction and Refinement

The success of description logics (DLs) in the many domains of semantic information processing is based on their flexibility to strike a carefully crafted trade-off between expressiveness and implementation efficiency. DLs have their origin in knowledge representation formalisms. As such they aim to hide semantical complexity in compact notation which is domain-specific rather than being general purpose. This leverages syntax to make the handling of logical specifications both by humans as well as reasoning engines run in a much higher gear ("application-level") compared to, say, plain vanilla first-order logic, in which all quantification structure is made explicit ("representation-level"). This work explores some aspects of yet another semantical dimension that can be accommodated within the syntax of DL which opens up when passing from the classical truth-value interpretation to a *constructive interpretation* of DL. We argue that such a refined interpretation is essential to represent applications with partial information adequately and to achieve both consistency under abstraction as well as robustness under refinement. An application area where this aspect is particularly prominent and which motivates our work, is auditing of business transactions. Audit statements about the validity of accounting data, absence of fraud or conformance to financial pro-

The Bulletin of the EATCS

ness standards must constructively take account of many dimensions of abstraction and refinement. We introduce a constructive version of \mathcal{ALC} , called $c\mathcal{ALC}$, and give a sound and complete Hilbert axiomatisation and a Gentzen tableau calculus showing finite model property and decidability for $c\mathcal{ALC}$.

This is joint work with Michael Mendler.

Tomoyuki Suzuki, University of Leicester

Algebraic and Relational Semantics for Distributive Substructural Logic

Full Lambek calculus FL is a logical system obtained by deleting the contraction, exchange and weakening rules from Gentzen's sequent calculus of intuitionistic logic. Extensions of FL are called *substructural logics*, and play an important role in computer science, from the study of resource sensitive computations to artificial intelligence. In this talk, we develop a relational semantics for some substructural logics. We investigate *distributive substructural logics* (*DFL logics* for short)—substructural logics satisfying the distributive law; examples include relevance logics, many-valued logics and superintuitionistic logics. The algebraic models of DFL logics are *DFL-algebras*—distributive residuated lattices with a constant 0. Using the standard Lindenbaum-Tarski method, one can show that every DFL logic is complete with respect to its DFL-algebras.

In modal logic, a connection between modal algebras and transition systems is obtained through Stone duality. Similar to modal logic, we introduce a relational semantics for DFL logics via Stone duality. However, unlike for modal logic, there are two ways of defining descriptive frames as duals of DFL-algebras. We define *quasi-descriptive frames* as tight and compact general frames and we define *descriptive frames* as anti-symmetric quasi-descriptive frames. We compare these two types of descriptive frames: let \mathcal{A} be the category of DFL-algebras; \mathcal{D} be the category of descriptive frames; and \mathcal{D}^q be the category of quasi-descriptive frames. We prove that \mathcal{A} is dually equivalent to \mathcal{D} ; and \mathcal{D} is a full reflexive subcategory of \mathcal{D}^q . In fact, both quasi-descriptive frames and descriptive frames provide us with completeness, and therefore with a relational semantics, for DFL logics. We then prove that every DFL logic is complete with respect to its quasi-descriptive and descriptive frames.

Richard M. Thomas, University of Leicester

Hyperbolic Monoids

The notion of hyperbolic groups has played a fundamental role in computational group theory. There were several equivalent definitions of the notion of hyperbolicity in groups but none of these generalized to monoids. This changed with Gilman's elegant characterization of hyperbolic groups (2002); Duncan and Gilman then suggested in 2004 that the formulation from Gilman's characterization could be taken as the definition of a hyperbolic monoid.

Their definition is entirely natural but we do not have efficient algorithms for dealing with hyperbolic monoids. It is known that the word problem for hyperbolic groups can be solved in linear time but the best known algorithm for the word problem in a hyperbolic monoid is exponential (Hoffmann et al. (2002)). Other questions (such as the conjugacy problem, which can be solved efficiently in hyperbolic groups) are still open as far as hyperbolic monoids are concerned, even as regards decidability. The purpose of this talk is to explain how restricting the definition used by Duncan and Gilman leads to efficient algorithms.

In general, the algorithmic properties of this new class of monoids suggest that they are worthy of further study. The new definitions are equivalent to the previous ones in the group setting and they allow us to develop efficient algorithms for monoids; for example, the word problem can be solved in time $O(n \log n)$. The conjugacy problem is also decidable in such monoids.

This is joint work with Michael Hoffmann.

Chris Tofts, Hewlett-Packard Laboratories, Bristol
Services Sciences Management and Engineering (SSME) – a Theoretical Computer Science Perspective.

The Communications of the ACM dedicated its July 2006 issue to Services Sciences, and some traditional computing companies – IBM (www.research.ibm.com/ssme) and HP (www.services-sciences.org) along with BT – have been expressing an interest in the area. The Cambridge Business School and Institute for management recently held a meeting Succeeding through Service Innovation (www.ifm.eng.cam.ac.uk/ssme) along with the emergence of a knowledge sharing network for the activity within UK universities (www.ssmenetuk.org). Is there anything here that should interest the Working Theoretical Computer Scientist?

This is joint work with Richard Taylor.

A. Dayem Ullah, King's College
A Local Move Set for Protein Folding in Triangular Lattice Models

The problem of predicting the folded structure of a protein based on its amino-acid sequence is one of the central problems in Structural Biology. The 3D structure of a protein determines its function and interaction properties, which is vital in all organisms. Following Anfinsen's thermodynamic hypothesis, we can formulate the protein folding problem as a combinatorial optimization problem, i.e., as finding a conformation with minimum energy state. Here a conformation is a self avoiding path in a specified grid such that grid points represent amino-acids, and the path assembles the protein specific amino-acid sequence. The number and kind of contacts between amino-acids, neighbouring elements in the grid not directly connected, are used to determine the energy of a conformation.

The Bulletin of the EATCS

Various simplified models in terms of lattice and energy function have been introduced to analyze the computational complexity of the problem. One of the most studied models is the HP model introduced by Dill et al. in 1989. In this model, the 20 letter alphabet of amino-acids is replaced by H (hydrophobic) and P (polar), and only H to H contacts contribute to the energy value. Protein folding has been shown to be NP-hard for various lattice models and contact-based energy functions (Hart and Istrail 1997, Crescenzi et al. 1998, Atkins and Hart 1999) and is NP-complete in the HP model (Berger and Leighton 1998). A variety of search-based algorithms has been proposed to solve the problem, mostly based on the HP model (Unger and Moulton 1993; Krasnogor et al. 1999; Blazewicz et al. 2005; Steinhöfel et al. 2007). For local search-based algorithms and folding in rectangular grids, Lesh et al. (2003) introduced a set of local moves, the so called *pull moves*, to be used as neighbourhood function. They showed that the pull move set is reversible and complete which are essential properties for the performance of local search algorithms. Although rectangular grids have been in the focus of the research, other grid models have been introduced that claim to be more realistic. Agarwala et al. (1997) introduced the 2D and 3D triangular lattices allowing more dense and parity-constraint free folded structures. Later, Jiang and Zhu (2005) proposed the concept of 2D hexagonal lattice and adapted the pull move set to the hexagonal lattice. Jiang et al. (2007) provided a tool set for structural analysis of RNA sequence and gave a brief introduction of pull moves on triangular lattice.

In this talk, we first formally introduce a complete *pull move set*, preserving improved locality, on 2D triangular lattice in the HP model. Secondly, we extend it on Face-centred-cubic (FCC) lattice using the fact that the FCC lattice is equivalent to the 3D triangular one. We prove the *reversibility* and *completeness* property for our pull move set in both models, 2D triangular and FCC. Finally, we present experimental results for several benchmark problems in the 2D triangular HP model using a simple tabu-search strategy. We show that, based on the pull move set, the algorithm exhibits a good performance and finds optimum energy configurations. Future research will focus on extending the experiments to the 3D triangular model and formally analyzing the pull moves with respect to different local search strategies such as Simulated Annealing.

This is joint work with H.-J. Böckenhauer, L. Kapsokalivas and K. Steinhöfel.

Meng Wang, Oxford University

Translucent Abstraction of Algebraic Datatypes with Safe Views

Algebraic datatypes and pattern matching are probably two of the best loved features of functional programming languages, a direct result of their elegant and convenient syntax for both program development and reasoning. At the same time, they are widely criticized for breaking abstraction and encapsulation. Much

effort has been put into tackling this conflict, notably Wadler's views and their variants. However, the original design of views assumes non-machine checkable conditions, which prevented them from gaining wider acceptance. In this work, we show how bi-directional transformation techniques can be used for the design of practical views and thus achieve abstraction and encapsulation of algebraic datatypes.

This is joint work with Jeremy Gibbons.

Dominik Wojtczak, University of Edinburgh

Recursive Stochastic Games with Positive Rewards

Motivated by the goal of analyzing the optimal/pessimal expected running time of probabilistic procedural programs, we study 1-exit Recursive Markov Decision Processes (1-RMDPs) and Recursive Simple Stochastic Games (1-RSSGs) with strictly positive rewards on all transitions. These models are equivalent to a game version of stochastic context-free grammars with positive rewards (where the objective of the two players is to minimize/maximize the total expected reward).

We describe systems of linear min-max optimality equations for these games, whose least fixed point solution yields the game's value vector. We show that in such games both players have optimal deterministic stackless and memoryless optimal strategies, and we provide polynomial-time algorithms for computing optimal expected rewards, and optimal strategies, for both maximizing and minimizing 1-RMDPs. It follows that the quantitative decision problem for positive expected reward 1-RSSGs is in $NP \cap co-NP$. We also show there is a simultaneous strategy improvement algorithm that converges in a finite number of steps (each step computable in P-time) to the optimal values and strategies of a 1-RSSG game with positive rewards. Just as in the case of Condon's finite-state SSGs, the worst-case complexity of this algorithm is open. These games are "harder" in the following sense: we show that even the problem of determining whether the value of a 1-RSSG strictly positive reward game is ∞ is already as hard as Condon's quantitative termination problem for finite-state SSGs, whereas for finite-state SSGs with strictly positive rewards this ∞ problem is solvable in polynomial time.

Optimized algorithms, based on both value iteration and strategy improvement, have been implemented in PReMo (a tool for analysis of probabilistic recursive models) for solving the above recursive reward games as well as other classes of recursive stochastic games. We describe some interesting comparative experimental results conducted with PReMo.

This is joint work with Kousha Etessami and Mihalis Yannakakis.

Peter Y. H. Wong, Oxford University

A Relative Timed Semantics for BPMN

Modelling of business processes and workflows is an important area in software

engineering. Business Process Modelling Notation (BPMN) allows developers to take a process-oriented approach to modelling of systems. In our previous work we have given an untimed process semantics to a subset of BPMN in the language of Communicating Sequential Processes (CSP). However due to the lack of a notion of time, it is not possible to model precisely concurrent activities with timing information; this is particularly important when specifying collaboration where the coordination of one participant depends on the execution order of another participant's activities.

In this talk we describe a relative-time semantic model for BPMN. We define the semantics in the language of CSP. This model augments the untimed process semantics by introducing the notion of relative-time in the form of delays over non-deterministic bounded ranges. By using CSP as the semantic domain, we show some properties relating the timed semantics and BPMN's untimed process semantics based on existing CSP's refinement orderings, and illustrate the application of the timed model via a simple business process example. Our timed model allows behavioural properties of BPMN diagrams to be mechanically verified via automatic model-checking provided by the FDR tool.

This is joint work with Jeremy Gibbons.

Yonghong Xiang, Durham University

One-to-Many Node-Disjoint Paths in (n, k) -Star Graphs

We present an algorithm which, given a source node and a set of $n-1$ target nodes in the (n, k) -star graph $S_{n,k}$, where all nodes are distinct, builds a collection of $n-1$ node-disjoint paths, one from each target node to the source. The collection of paths output from the algorithm is such that each path has length at most $6k-7$, and the algorithm has time complexity $O(k^3n^4)$.

This is joint work with Iain Stewart.

Stanislav Živný, Oxford University

Valued Constraints – Modelling, Expressive Power and Algebraic Properties

In this talk I present valued constraints and discuss the usefulness of the VCSP framework. I first discuss problems I am currently working on: an algebraic characterisation of the expressive power of valued constraints; and the connection between valued constraints and the problem of SFM. I then mention results which have been obtained so far. Firstly, I discuss ways in which a fixed collection of valued constraints can be combined to express other valued constraints. I show that in some cases a large class of valued constraints, of all possible arities, can be expressed by using valued constraints of a fixed finite arity. I also show that some simple classes of valued constraints, including the set of all monotonic valued constraints with finite cost values, cannot be expressed by a subset of any fixed

finite arity, and hence form an infinite hierarchy.

Finally, I consider the problem of modelling combinatorial optimisation problems in the VCSP framework. I show that some problems, including the (s, t) - M - C problem and the problem of S F M can be naturally modelled in the VCSP framework. However, for other, apparently similar, problems such as the M - C problem and the problem of S S F M , which also have polynomial-time algorithms, we show that they can only be naturally modelled in the VCSP framework by using valued constraints which are powerful enough to represent NP-complete problems.

This is joint work with David Cohen and Peter Jeavons.