Report on BCTCS 2004

The 20th British Colloquium for Theoretical Computer Science
5–8 April 2004, Pitlochry, Scotland

Stephan Reiff-Marganiec and Carron Shankland

The 20th British Colloquium for Theoretical Computer Science, BCTCS 2004, was held 5–8 April 2004 in Pitlochry (Scotland). The event was hosted by the Department of Computing Science and Mathematics, University of Stirling and was supported by LMS, EPSRC, BCS-FACS and Microsoft Research.

The colloquium featured an interesting and wide ranging program of five invited talks, two invited tutorials, and 19 contributed talks. The invited talks were presented by Dr Luca Cardelli (Microsoft Research, UK), Dr Sharon Curtis (Oxford Brookes University), Prof José Fiadeiro (University of Leicester), Dr Rob Irving (University of Glasgow) and Prof Ken Turner (University of Stirling); and the invited tutorials were presented by Dr Rachel Norman (University of Stirling) and Prof Rick Thomas (University of Leicester). Two prizes for outstanding talks, sponsored by BCS-FACS, were awarded to Alastair Donaldson (University of Glasgow) and Corrina Elsenbroich (King's College London).

The main purpose of BCTCS is to provide an opportunity for networking in the UK Theoretical Computer Science community. As usual, participants took great advantage of this. A high point of the colloquium was the formal dinner at which the outgoing president (Dr Chris Tofts, HP Labs Bristol) mused on the past and future of BCTCS. He was presented with a small token of our esteem by the incoming president (Prof Faron Moller, University of Wales Swansea), purchased earlier in the day when we all visited Edradour distillery for a whisky tasting.

The 21st British Colloquium for Theoretical Computer Science, BCTCS 2005, will be hosted by the University of Nottingham during 22-25 March 2005. Anyone wishing to contribute talks concerning aspects of Theoretical Computer Science is warmly invited to do so. Details are available from the Colloquium web site at `http://www.cs.nott.ac.uk/~gmh/bctcs2005.html`.

Invited Talks

Luca Cardelli, Microsoft Research, Cambridge, ***Membrane interactions***
Computer simulation of biological systems is computationally intensive. Fortunately, living cells are not just random collections of billions of molecules: they are extremely well-organized (although stochastic) systems. If we can model their organization, we should be able to simulate them abstractly and effectively. A large part of the organization of (eukaryotic) cells depends on hierarchies of nested

membranes, their properties, the proteins bound to membrane surfaces, and the way membranes interact dynamically with each other. I discuss some preliminary efforts in modeling dynamic membrane systems. The long-term goal is to represent the structure and function of biological systems via formal languages, for description, simulation, analysis and (eventually) compilation.

Sharon Curtis, Oxford Brookes University, ***Functional fractal image compression***
Fractal image compression is an ingenious technique that compresses an image by representing it as a partitioned iterated function system. As functions are fundmentally involved in this modeling of images, it is a natural step to express the image compression and decompression algorithms in the world of functional programming. This talk will give an introduction to fractal image compression and an account of recent experiences implementing the compression and decompression in Haskell.

José Fiadeiro, University of Leicester, ***Software architectures in 3D***
The promise of Software Architectures for controlling the complexity of system construction and evolution is based on a conceptually simple separation of concerns: that between the *computational* aspects associated with the functionality of basic services, and the mechanisms through which their interactions are *coordinated*. This model is too "static" for the new generation of systems that will have to operate in dynamic configuration topologies. While we are beginning to understand how to address mobility of computations, the effects of mobility on coordination are only now being recognised as an additional factor of complexity, one for which current architectural concepts and techniques are not prepared. This lecture will address the challenge of making Software Architectures three-dimensional by proposing a mathematical model in which Distribution is recognised as a separate concern from Computation and Coordination.

Rob Irving, University of Glasgow, ***Fifty years of stable marriage***
It is over fifty years since the classical stable matching algorithm was first used in a large scale practical application. Since then, the stable marriage problem has been the focus of extensive interest and research on the part of computer scientists, mathematicians, game theorists, economists, and others. This talk presents an overview of stable matching problems, primarily from the standpoint of the algorithmicist. We emphasise a range of recent developments, and explore their practical implications as well as theoretical significance.

Ken Turner, University of Stirling, ***Test generation for radiotherapy accelerators***
System specification with LOTOS (Language Of Temporal Ordering Specification) will be briefly described. To make test generation practicable, specifications are annotated with event constraints using PCL (Parameter Constraint Language)

for stating test purposes. Automated test generation follows the principle of input-output conformance to check that an implementation agrees with its specification. Test suites are created from a transition tour of the automaton generated from the specification. The approach has been applied to a safety-critical systems case study: radiotherapy accelerators as used in cancer treatment. The automatically generated test suite can be executed manually or automatically. The goal is to discover situations in which an accelerator does not conform to its specification.

Invited Tutorials

Rachel Norman, University of Stirling, ***Modelling of biological systems***
The aim of this tutorial will be to convince you that modeling biological systems is important and exciting. Mathematical models have been used successfully to model many different types of biological problems. However, as the models become more biologically realistic and therefore more mathematically complex, theoretical biologists are turning to theoretical computing science techniques to see what they can bring to the party. The main emphasis of the talk will be to give you some background and examples of the types of problems addressed and types of mathematical models used. I then go on to discuss my ideas about the types of TCS techniques that could be used and how the two could be linked together. Biological emphasis will be on infectious disease systems but the techniques could be applied to any area of biology.

Rick Thomas, University of Leicester, ***Formal languages and word problems of groups***
In this talk we survey some connections between formal language theory and group theory. Groups arise in many ways and are often described by "presentations." A presentation for a group G consists of a set of generators for G and a set of relations between words in the generators sufficient to define the group operation. Given a presentation for G with a finite set of generators and a finite set of relations, we say that G is "finitely presented." A classical result in group theory is the unsolvability of the word problem for finitely presented groups: there are finite presentations such that there is no algorithm to decide whether or not a word in the generators represents the identity element of the group defined by that presentation. An elegant result of Boone and Higman describes which finitely generated groups have a solvable word problem. A natural question arises from this: if we take some restricted model of computation, which groups have a word problem which is decidable within that model? The restriction might be, eg, that the word problem is accepted by a particular type of automaton or generated by a particular type of grammar. We survey some of what is known in this field; we will not assume any knowledge of group theory, and review what we need from formal language theory.

Contributed Talks

Alastair Donaldson, University of Glasgow, ***Investigating structural symmetry in models of concurrent systems***

Symmetry reduction techniques have been successful in combating state space explosion in model checking. Such techniques usually assume that symmetries of a concurrent system are known in advance and are exploited when model checking the original system description. Two obvious questions ask if we can we detect symmetries of a concurrent system from its textual description, and if it is possible to apply symmetry reduction directly to the textual description of a system so that a reduced textual description could be model checked in the usual way. In this talk I describe attempts to answer these questions.

Corinna Elsenbroich, King's College London, ***Goal directed dynamic abduction***

Building on an idea by Batens et al. [Dynamic Dialectical Logics, 1989] we define a goal directed dynamic logic which facilitates reasoning with inconsistencies while giving a larger set of conclusions than a normal paraconsistent logic. The logic is dynamical as conclusions drawn at certain stages of the proof are conditional and can be withdrawn later on if the conditions are no longer met. This dynamic is what is needed for abductive reasoning so that we can define an abduction rule integrated into the above framework.

Fedor Fomenko, Edinburgh University, ***Dual-context multicategories***

Bellantoni and Cook gave a recursion-theoretic description of PTIME computations independent of externally imposed resource bounds, system B. They separated function inputs into *normal* and *safe* inputs; functions can apply any PTIME operation to normal inputs, but on safe inputs allow only operations which increase the length by an additive constant. Their safe recursion scheme defines new functions by recursion on normal inputs, and recursively defined values are processed only through safe inputs. Barber proposed DILL, a sequent style formulation of Intuitionistic Linear Logic with clear introduction-elimination rules for all connectives, which splits contexts into two parts: one for intuitionistic variables and one for linear variables. Weakening and contraction are allowed only for intuitionistic variables. We extend the definition of multicategory [Lambek] to dual-context multicategory with contexts separated into two parts. We define linear-linear (LL), intuitionistic-linear (IL) and intuitionistic-intuitionistic (II) multicategories. IL multicategories provide models for DILL. We consider the strong endofunctor of dual-context multicategories, an extension of the strong endofunctor on usual categories, and define initial algebras for such endofunctors corresponding to the Bellantoni-Cook safe recursion scheme, and show how system B can be modeled in II dual-context multicategories.

Ana Fonseca, University of Leicester, ***The irreducible word problem of groups***

Given a group G with a finite generating set X, the word problem of G with respect to X is the set of words over the elements of X and their inverses which are equal to the identity in G. The irreducible word problem of G is then the subset of the word problem consisting of those nonempty words that have no nonempty proper subword which is equal to the identity. There are many interesting links between word problems and irreducible word problems on the one hand and formal language theory on the other. In particular, groups whose irreducible word problem is either a finite or regular language have been classified. I present some results concerning groups with a context-free irreducible word problem.

David Gabelaia, King's College London, ***Modal logics of submaximal spaces***
The operations of closure and derivation on a topological space can be understood as normal modal operators working on the boolean algebra of all subsets of the underlying set. This gives rise to modal logics associated with topological spaces. It turns out that the formalism with the closure operator is rather weak in feeling the topological structure, whereas the derivation operator offers more expressivity, and allows for differentiating natural topological classes by logical means. In this talk we concentrate on the subclasses of submaximal spaces (among others door spaces and maximal spaces), which have been the subject of increasing interest in the topological community. Such spaces are also of interest in digital topology (eg. Khalimsky line is submaximal). We demonstrate that the equations involving the closure operator alone fail to distinguish these classes. On the other hand, equations built using the derivation operator capture precisely the individual characteristics of each of them. Finally, for each of the topological classes concerned, a simple class of finite Kripke frames generating the same modal validities will be exhibited.

Richard Geary, University of Leicester, ***Representing XML documents compactly***
We consider succinct or space-efficient representations of trees that effectively support a variety of navigation operations. Our representation takes 2n + o(n) bits to represent a tree of n nodes, which is within o(n) bits of the information-theoretic minimum and supports all operations in O(1) time on the RAM model. In addition to the existing motivations for studying such data structures, we are motivated by the problem of representing XML documents compactly so that XPath queries can be supported efficiently.

Neil Ghani, University of Leicester, ***Categories of containers***
Efficient representation and manipulation of data is a fundamental task in constructing large software systems. One successful approach has been Hindley-Milner polymorphism which provides predefined mechanisms for manipulating data structures providing they are parametric in the data. I will talk about work on containers which has lead to some intriguing new insights into polymorphism.

In particular I will classify all the polymorphic programs between containers and show they have a remarkably simple form. I then extend this to cover quotient types where a similar beautiful picture emerges.

Horacio Gonzalez-Velez, University of Edinburgh, ***On the issues of the application of algorithmic skeletons to metacomputing***
Metacomputing refers to the execution of parallel code on multiple supercomputers or distributed clusters. Grid computing encompasses not only metacomputing in the form of computational grids, but also scheduling, enhanced security, data access and migration, peer-to-peer processing, and other technologies. With the advent of computational grids as part of the world initiative on the Grid, the use of proven programming paradigms has become an issue. Current trends in scientific application programming for computational grids include the development of frameworks for distributed programming, Internet-based processing or component-based technology, as well as the use of grid-based portals. Algorithmic skeletons, conceived as higher order functions corresponding to good parallel algorithmic techniques, provide reliable methods for the design of high-level, structured parallel programming solutions. In our approach we do not intend to adapt mainstream, distributed Grid development technology into the main programming arena, but conversely to apply parallel programming techniques to solve scientific or industrial problems in the Grid. Parallel algorithmic solutions often show certain patterns (pipelines, process farms, branch-and-bound) that are naturally addressed using skeletons. Hence, it may be of benefit as a novel way of programming a grid-oriented computational environment.

Will Harwood, University of Wales Swansea, ***Using unique decomposition bases to efficiently compute bisimilarity on normed Basic Parallel Processes***
Jančar and Kot show (AVIS 2004) that the PSPACE BPP-bisimilarity-checking algorithm of Jančar (LICS 2003), which characterises states modulo bisimilarity using sets of NORMS, can be made to run in $O(n^3)$ time on the subclass of *normed* Basic Parallel Processes. They consider this to be a significant improvement over the previous polynomial-time algorithm for bisimilarity on normed BPP, by Moller, Hirshfeld and Jerrum (MSCS 1996), which relied on *unique decomposition bases* and for which no explicit upper bound was given. However, while the latter's related algorithm for normed Basic Process Algebra (TCS 1996) performs (naïvely) in $O(n^{19})$ time, we show here that the decomposition base method for normed Basic Parallel Processes is not so pessimistic.

Roman Kontchakov, King's College London, ***Combining spatial and temporal logics: expressiveness versus complexity***
Modal logic S4u is complete with respect to the topological semantics where the diamond modality is interpreted as the closure operator of topological spaces. On the other hand, RCC-8, a fragment of the region connection calculus RCC with

eight jointly exhaustive and pairwise disjoint predicates, is known to be embeddable into S4u. In this talk we first consider several intermediate fragments of S4u extending RCC-8 whose computational complexity varies from NP to PSPACE. Then we construct and investigate a hierarchy of spatio-temporal formalisms that results from various combinations of the propositional temporal logic PTL with these fragments of S4u. We demonstrate how different 'blending' principles as well as spatial and temporal components give rise to NP-, PSPACE-, EXPSPACE-, 2EXPSPACE-complete and even undecidable spatio-temporal logics.

David Manlove, University of Glasgow, ***Approximability results for induced matchings in graphs***
An induced matching in a graph $G = (V, E)$ is a set of edges M, no pair of which are adjacent or joined by an edge in G. Let MIM be the problem of finding a maximum induced matching in a graph G. MIM has applications in channel assignment, VLSI design and network flow problems. However MIM is known to be NP-hard, even if G is d-regular (i.e., each vertex has degree exactly d). In this talk we consider the approximability of MIM in d-regular graphs. We show that, for each $d \geq 3$, MIM admits an approximation algorithm with performance guarantee $d-1$. On the other hand, we show that, for the same class of graphs, MIM does not admit a polynomial-time approximation scheme unless P=NP. This is joint work with Billy Duckworth and Michele Zito.

Richard McKinley, University of Bath, ***Categorical models of first order classical sequent proofs***
It has been thought for a long time that there were no non-trivial models of classical logic that admit full cut elimination - Fuhrmann and Pym have recently shown models which escape the usual collapse to a boolean algebra, by enriching the proof spaces with a partial order structure - this partial order interprets the structure of cut-elimination. I will briefly recap the notions of sequent calculus and cut-elimination, and the non-collapsing models, and then present the generalisation of this result to the first order case, using an indexed category approach - we exhibit a sound, complete class of models, with examples, and consider what new light the structure of these models might shed on Herbrand's Theorem.

Jose Raymundo Marcial-Romero, Birmingham University, ***Semantics of a sequential language for exact real-number computation***
We present a programming language with a built-in ground type for real numbers. In order for the language to be sufficiently expressive but still sequential, we adopt a construction by Boehm and Cartwright. The non-deterministic nature of the construction suggests the use of powerdomains to obtain a denotational semantics for the language. We show that the construction cannot be modeled by the Plotkin or Smyth powerdomains, but that the Hoare powerdomain gives a computationally adequate semantics. As is well known, Hoare semantics can be

used in order to establish partial correctness only. As computations on the reals are infinite, one cannot decompose total correctness into partial correctness and termination as it is traditionally done. We instead introduce a suitable operational notion of strong convergence and show that total correctness can be proved by establishing partial correctness (by denotational methods) and strong convergence (by operational methods). We illustrate this with a representative example.

Graham Oliver, University of Leicester, ***Automatic presentations and groups***
A structure is said to be computable if there is a coding for its domain and relations that can be read by Turing machines. Restricting the Turing machines in the definition to finite automata then gives us structures which are particularly simple computationally; these structures are said to have an automatic presentation. This talk will review the general area, before presenting a classification in a specific case: a finitely generated group has an automatic presentation if and only if it has an abelian subgroup of finite index.

Igor Potapov, University of Liverpool, ***From discrete games to counter automata and matrix reachability problem***
In recent years, discrete games have attracted more and more attention in computation theory as the dynamics complexity is directly connected to computational power of the game. In models of discrete adaptive games, a finite number of players or agents pick from a finite set of actions, at discrete times, and receive rewards that depend on their own and the other agents' actions. Each agent chooses its actions according to a strategy; if agents have more than one strategy available, they use a second layer strategy to choose among them. For computation power analysis we suggest representing the whole process as a counter automaton with guards. In that case, many questions about the predictability of such games reduce to reachability problems in automata-like systems. In this presentation we show new decidability and undecidability results for counter automata with guards which are connected to standard computational models, reachability problems for piecewise and non-deterministic iterative maps, matrix semigroups etc. This is joint work with Alexei Lisitsa.

Paul Sant, King's College London, ***Graph problems in a data streaming context***
In the Data Streaming model input is presented as a stream of data, for example, a stream of edges of a graph. This model has been studied by, eg, Muthukrishnan and Cormode. An added constraint is that only polylogarithmic space is available for storage. Much of the work in the data streaming model has concentrated on finding most frequent items in a data set, approximating the L_p norm, etc. However, graph problems have not been extensively studied in this model and only recently have any results been published on what is possible (from a graph theoretic perspective) within this model. In this talk I describe some graph problems

and results that are known and explain the main problems that are encountered when looking at graph problems in the data streaming context.

Francois Siewe, De Montfort University, ***An algebra of security policy***
Despite considerable work on access control policies, enforcing multiple policies is still a challenge to achieve the level of security required in many real-world systems. Also, current approaches address security settings independently, and their incorporation into the systems development lifecycle is not well understood. We present an algebra for the specification of access control policies which can handle the enforcement of multiple policies through policies composition. Interval Temporal Logic (ITL) is our underlying framework and policies are modeled as safety properties expressing how authorizations are granted over time. The approach is compositional and can be used to specify system properties such as functional and temporal requirements. The use of a common formalism eases the integration of security requirements into system requirements so that they can be reasoned about uniformly throughout the development lifecycle. Policy specifications are executable in Tempura, a simulation tool for ITL, letting you validate a formal specification against the security requirement before proceeding to design and implementation.

Gabriel Valiente, Technical University of Catalonia, ***Tree matching: metrics and algorithms***
Tree matching is a natural generalisation of string matching to trees, with practical applications in combinatorial pattern matching, pattern recognition, chemical structure search, computational molecular biology, and other areas of engineering and life sciences. In this talk, a gentle introduction to tree matching will be given, recent algorithmic results will be reviewed, and similarity measures based on tree matching will be discussed.

Daniel Wiley, University of Bath, ***A logic for describing operating systems in terms of their resources and processes***
Instances of the Logic of Bunched Implications have been successfully used to describe resource semantics. Judgements of the form a resource model satisfies a bunched logical assertion have been used to describe systems such as heap storage. We hope to enrich this judgment to make assertions about resources relative to processes acting on these resources satisfying bunched assertions. We demonstrate that, through this enrichment, we are able to describe a complex, dynamic system, an operating system, and in doing so, we gain a significant insight into the properties and resource nature of such systems.