MANUEL BODIRSKY
École Polytechnique

Homogeneous Structures and Siggers Terms

In this talk I present classification results for structures B that are first-order definable over a countably infinite homogeneous structure. Those results resemble the situation in the finite: either 3SAT has a primitive positive interpretation in an expansion of (the model-complete core of) B by finitely many constants, or B has a four-ary polymorphism that satisfies the Siggers term identities "up to automorphisms". In the first case, CSP(B) is NP-hard by general principles. In the second case, one would expect that CSP(B) can be solved in polynomial-time, and indeed this has been verified for all structures B that are definable over the order of the rationals, the Rado graph, the universal homogeneous C-relation, or the equivalence relation with infinitely many infinite classes.

ANDREI BULATOV
Simon Fraser University

Conservative dichotomy revisited

The first proof of the Dichotomy conjecture for conservative algebras (2003) introduced colored graphs of relational structures and suggested a very sophisticated solution algorithm for polynomial time problems. The alternative proof of the same result given by Barto in 2011 uses a different approach and presents a significantly simpler algorithm. In this talk we make use of recent developments in the CSP such as the Generalized Majority-Minority algorithm and Maroti's reduction to simplify dramatically the old proof.

VICTOR DALMAU
Universitat Pompeu Fabra

Robust approximation of CSPs

Let Γ be a constraint language and let I be an instance of CSP(Γ). An assignment s (of the variables of I) is a $(1-\epsilon)$-solution, $\epsilon > 0$, of I if it satifies at least a $(1-\epsilon)$-fraction of its constraints. If such an assignment exists we say that I is $(1-\epsilon)$-satisfiable.

Let f be a real function. An f-approximation algorithm (for CSP(Γ)) is an algorithm that receives as input an instance I of CSP(Γ) and returns an assignment s, in such a way that if I is $(1-\epsilon)$-satisfiable then s is a $(1-f(\epsilon))$-solution. If, additionally, $\lim_{\epsilon \to 0} f(\epsilon) = 0$ then we say that CSP(Γ) has a *robust approximation* algorithm. It has been conjectured by Guruswami and Zhou that CSP(Γ) has a robust approximation algorithm if and only if Γ has bounded width. In this talk we shall present some initial results towards proving this conjecture.

MARTIN DYER
University of Leeds

On the complexity of # CSP

This is joint work with David Richerby.

Andrei Bulatov showed that # CSP has a dichotomy, using methods from universal algebra. We will describe a simpler proof of the result, based on succinct representations for a certain class of relations, similar to one described by Bulatov and Dalmau. This leads to a new criterion for the dichotomy, which differs from Bulatov's, but can be shown to be equivalent to it. The advantage of the new criterion is that it leads to a decision procedure for the dichotomy, a major question left open by Bulatov. The talk will conclude with an outline of this decision algorithm.

PETER JEAVONS
University of Oxford

A Galois Connection for Valued Constraints

The valued CSP is a generalization of the CSP which associates a cost with each possible assignment, and seeks to minimise the total cost. It includes the CSP as a special case (where all costs are either 0 or infinite) but also allows many more discrete optimisation problems to be modelled easily. I will introduce the VCSP and discuss various analogs to the notion of polymorphism that have been developed for analysing its complexity. In particular, I will show that the standard Galois connection given by the functions Pol and Inv can be generalized to the valued case.

VLADIMIR KOLMOGOROV
University College London

The complexity of conservative valued CSPs

This is joint work with Stanislav Zivny.

In a valued constraint satisfaction problem (VCSP), a *language* is a fixed set of cost functions over a fixed domain. An *instance* from this language is specified by a sum of cost functions from the language, and the goal is to minimise the sum. Classifying the complexity of such minimisation problems for different languages has been an active research area. While a significant progress has been made for CSPs (i.e. languages containing only $\{0, \infty\}$-valued functions, or relations), fewer results are known for the more general case of VCSPs.

I will talk about our recent work on general-valued languages containing all possible unary cost functions. We call such languages *conservative*. (This definition differs from

the one used in the CSP literature for relations, where a language was called conservative if it contains all possible unary relations.)

We prove a Schaefer-style dichotomy theorem for conservative languages: if all functions in the language satisfy a certain condition then the language can be minimized in polynomial time (via a new algorithm that we developed), otherwise it is NP-hard. This is the first complete classification for general-valued constraint languages over non-Boolean domains. Our work generalises previous results obtained by Cohen et al. [AIJ'06], Takhanov [STACS'10], and Deineko et al. [JACM'08]. Moreover, our proofs do not rely on any computer-assisted search as in Deineko et al. [JACM'08].

MARCIN KOZIK
Jagiellonian University

CSPs with near-unanimity polymorphisms are solvable by linear Datalog

This is joint work with Libor Barto and Ross Willard.

We generalize the result of Dalmau and Krokhin and prove that a finite template with compatible near-unanimity polymorphisms (or equivalently, by Barto's result, Jónsson polymorphisms) generates a CSP with a complement solvable by a program of linear Datalog. This places the CSP in NL.

BENOIT LAROSE
Champlain College

NU operations, graphs and finite duality

This is joint work in progress with T. Feder, P. Hell, C. Loten, M. Siggers, C. Tardif.

We present a unifying framework allowing us to describe graphs admitting a k-NU operation via finite duality. Our results encompass and extend several results from the past 30 years on majority operations.

MIKLÓS MARÓTI
University of Szeged

Beyond bounded width and few subpowers

There are two important algorithms for solving constraint satisfaction problems (CSP) over restricted classes of algebras. One is the $(2,3)$ *consistency algorithm* for solving instances of the CSP where the algebras are in a congruence meet-semidistributive variety. The other is the *generalized Gaussian elimination algorithm* for solving instances of the CSP where the algebras have few subpowers (equivalently have edge terms). We present two new algorithms, one that combines the $(2,3)$ consistency algorithm with the generalized Gaussian elimination algorithm, and the other is a new reduction technique that uses consistent set of unary polynomials.

BARNABY MARTIN
Durham University

A tetrachotomy for positive equality-free logic

We consider the problem of evaluating positive equality-free sentences of first-order logic on a fixed, finite relational structure B. This may be seen as a generalisation of the Quantified Constraint Satisfaction Problem QCSP(B). We argue that our generalisation is not totally arbitrary, and that ours is the only problem in a large class - other than the CSP and QCSP - whose complexity taxonomy was unsolved.

We introduce surjective hyper-endomorphisms in order to give a Galois connection that characterises definability in positive equality-free FO. Through the algebraic method we are able to characterise the complexity of our problem for all finite structures B. Specifically, the problem is either in Logspace, NP-complete, co-NP-complete or Pspace-complete. The problem appears obtuse, but possesses a surprising elegance. Perhaps there are lessons to be learnt in the methodology of the solution of this case, for the continuing program for CSP.

RALPH MCKENZIE
Vanderbilt University

Finitely related clones and algebras with cube-terms; Valeriote's conjecture for finite algebras in congruence modular algebras and its consequence for the CSP dichotomy conjecture

We say that an algebra $\mathbf{A}$ has *sparse relational clone* (or equivalently, has *few subpowers*) if there is a positive constant k so that $|\mathrm{Sub}(\mathbf{A}^n)| \leq 2^{n^k}$ for all integers $n > 1$. It is known that a finite algebra $\mathbf{A}$ has sparse relational clone iff $\mathbf{A}$ has one (or equivalently, all) of

the following: a cube term, an edge term, a parallelogram term. Recently, Aichinger, Mayr and McKenzie have shown that every finite algebra **F** with sparse relational clone is finitely related; that is, there is a finitary relation ρ over F so that the finitary operations that respect ρ constitute precisely the clone of all term operations of **F**.

Let F be any finite set. We work with subclones of the clone $\mathcal{I}$ of all idempotent operations over F. In this talk, we describe finitely many clones $\mathcal{E}_1, \ldots, \mathcal{E}_k$ included in $\mathcal{I}$ which turn out to be precisely the maximal subclones of $\mathcal{I}$ that fail to contain a cube operation. They are also precisely the maximal not finitely related subclones of $\mathcal{I}$. For each of these clones, there is a trivial polynomial-time algorithm to determine if an operation belongs to it. This yields an algorithm to determine, given finitely many operations $f_1, \ldots, f_\ell$ over a finite set F, whether or not the clone generated by these operations contains a cube operation. With F fixed, the algorithm is polynomial-time as a function of its input; but with variable F it becomes exponential-time. There is a problem here, to determine if this computational problem with input $F, f_1, \ldots, f_\ell$ has a polynomial-time algorithm.

Applying the theorem of Aichinger-Mayr-McKenzie, we obtain the corollary that a finite idempotent algebra **F** has sparse relational clone if and only if **F** is inherently finitely related; i.e., every algebra obtained by enriching the set of operations of **F** is finitely related. Using a recent result of Libor Barto, and an application of commutator theory, we find interesting equivalent conditions for the truth of Valeriote's conjecture that a finite algebra in a congruence modular variety is finitely related if and only if it has a cube term.

MICHAEL PINSKER
Université Denis Diderot – Paris 7

Making the infinite finite: Polymorphisms on Ramsey structures

Many natural computational problems can be formulated elegantly as CSPs over a countably infinite structure S (for example, given a digraph, decide whether or not it is acyclic). However, the price to pay for such a formulation is generally high, as there are many operations possibly proving tractability, and many relations possibly proving hardness of such CSPs, and both operations and relations are infinite objects.

We present a general method for obtaining complexity classifications for the CSP of all structures which are first-order definable in an ordered homogeneous Ramsey structure. This method roughly consists in finding patterns of regular behaviour in arbitrary operations on the structure - such patterns allow to relate infinite operations (in particular, polymorphisms) to operations on finite sets, and to translate existing tractability results on finite sets back to the infinite.

In particular, given a relation R in this context, it is possible to calculate a finite number of "minimal" generic polymorphisms that violate R; such generic polymorphisms can be

represented as finite functions, and used by algorithms. This implies, for example, that the following problem becomes decidable: Given a relation R and a relational structure S, both of which have first-order definitions in an ordered homogeneous Ramsey structure which is finitely bounded, does R have a primitive positive definition from S?

JOHAN THAPPER
École Polytechnique

Min CSP on Four Elements: Moving Beyond Submodularity

This is joint work with Peter Jonsson and Fredrik Kuivinen.

We present a computational complexity classification of Min CSP (or, equivalently Max CSP) for constraint languages over a four-element domain. The result shows that every such problem is either tractable or NP-complete. Similar classifications are known for the two- and three-element domain cases. To prove our result, we introduce a new class of tractable problems based on combining submodular- and bisubmodular function minimisation. We further establish some basic techniques which allow us to obtain the result without relying on computer-assisted case analyses (which were heavily used in the classification of Max CSP on a three-element domain). Our result shows, for the first time, that submodularity is not the only source of tractability for Min CSP.