Program Tuesday, March 7

13:00–13:25	Julian Bitterlich TU Darmstadt	Acyclicity, Simple Connectivity, and Unbranched Covers of Hypergraphs
13:25–13:50	Kord Eickmeyer TU Darmstadt	FO model checking on map graphs
13:50–14:15	Nicole Schweikardt Humboldt-Universität zu Berlin	First-order logic with counting: at least, *weak* Hanf normal forms always exist and can be computed!
14:15–14:45	Coffee Break	
14:45–15:10	Florian Bruse Universität Kassel	The State of HFL Model-Checking
15:10–15:35	Felix Canavoi TU Darmstadt	A Modal Characterisation Theorem for Common Knowledge Logic
15:35–16:00	Marco Voigt Max-Planck-Institut für Informatik	A fine-grained hierarchy of complete problems in the separated fragment of first-order logic
16:00–16:15	Coffee Break	
16:15–16:40	Berit Grußien Humboldt-Universität zu Berlin	Capturing Polynomial Time using Modular Decomposition
16:40–17:05	Roman Rabinovich TU Berlin	Model-Checking for Successor-Invariant First-Order Formulas on Graph Classes of Bounded Expansion
17:05–17:30	Sebastian Siebertz TU Berlin	Polynomial Kernels and Wideness Properties of Nowhere Dense Graph Classes
17:30–17:45	Coffee Break	
17:45-18:10	Isolde Adler University of Leeds	Property Testing for structures of bounded degree
18:10–18:35	Martin Ritzert RWTH Aachen University	Learning first-order definable concepts over structures of small degree
20:00	Dinner	

Program Wednesday, March 8

9:00–9:25	Andre Frochaux Humboldt-Universität zu Berlin	Static Analysis of Monadic Datalog on Finite Labeled Trees
9:25–9:50	Erich Grädel RWTH Aachen University	Provenance Analysis in Logic and Games
9:50–10:15	Matthias Niewerth Universität Bayreuth	Enumeration of MSO Queries on Strings with Constant Delay and Logarithmic Updates
10:15–10:45	Coffee Break	
10:45–11:10	Christoph Berkholz Humboldt-Universität zu Berlin	Answering Conjunctive Queries under Updates
11:10–11:35	Jens Keppeler Humboldt-Universität zu Berlin	Answering FO+MOD queries under updates on bounded degree databases
11:35–12:00	Nils Vortmeier TU Dortmund	Connecting AC^1 and Dynamic Descriptive Complexity
12:00–12:15	Coffee Break	
12:15–12:40	Martin Grohe RWTH Aachen University	The Descriptive Complexity of Solving Linear Equation Systems and its Applications
12:40–13:05	Daniel Neuen RWTH Aachen University	An exponential lower bound for Individualization-Refinement algorithms for Graph Isomorphism

Abstracts

Acyclicity, Simple Connectivity, and Unbranched Covers of Hypergraphs

Tuesday, 13:00–13:25
by *Julian Bitterlich*, *TU Darmstadt*

Constructing finite, highly acyclic covers of hypergraphs is hard. Classical methods for graph covers do not work in this setting and new sophisticated techniques have to be used. In general, these covers have to be branched.

Our research aims at characterising the class of hypergraphs that admit finite, highly acyclic unbranched covers. To this end we apply techniques that are similar to the ones used for constructions of unbranched graph covers. Perhaps unsurprisingly, this question leads to a connection with basic algebraic topology.

In this talk we summarise the main points of our research:

1. We find a new, intuitive characterisation of acyclicity in hypergraphs.
2. We show that simple connectivity and acyclicity of hypergraphs do agree on the class of locally acyclic hypergraphs.
3. We find that the class of locally acyclic hypergraphs can also be characterised as those hypergraphs that admit unbranched acyclic covers. Furthermore we show that finite, locally acyclic hypergraphs admit finite, highly acyclic, unbranched covers.
4. At the heart of the proof of the aforementioned results lies the theory of what we call 'granulated covers', that is, graph covers that do not unravel certain cycles. We show that granulated covers provide universal elements and finite approximations of these. This result is interesting from a purely graph theoretic point of view.

Potentially, these results can lead to some van Benthem-style characterisation theorem for guarded logic with counting on suitable "frames", but this is still work in progress.

FO model checking on map graphs

Tuesday, 13:25–13:50
by *Kord Eickmeyer*, *TU Darmstadt*

Starting with Courcelle's Theorem in the 1990's, a range of efficient algorithms for model checking on restricted classes of relational structures have been obtained. For first-order logic in particular, Grohe et al. gave a near-linear time fpt algorithm for model checking on nowhere dense graphs. For monotone graph classes (i.e. those closed under taking subgraphs) this result is tight in the sense that model checking on effectively somewhere dense graphs is not fixed-parameter tractable (modulo the usual complexity theoretic assumptions).

For non-sparse graph classes few results are known, notably MSO model checking for graphs of bounded clique-width and FO model checking on certain interval graphs and on posets of bounded width. We show that FO model checking on map graphs is tractable as well.

Map graphs are a generalisation of planar graphs; a graph is a map graph if it can be drawn in the plane in such a way that to every vertex there corresponds a simply connected region in the plane, and such that two vertices are adjacent if, and only if, their corresponding regions touch. Since the number of regions that touch in a single point is not bounded, map graphs may contain arbitrarily large cliques. The class of map graphs is closed under taking induced subgraphs, but not under edge removal.

This is joint work with Ken-ichi Kawarabayashi.

First-order logic with counting: at least, *weak* Hanf normal forms always exist and can be computed!

Tuesday, 13:50–14:15

by *Nicole Schweikardt*, *Humboldt-Universität zu Berlin*

This talk presents recent joint work with Dietrich Kuske. We introduce the logic FOCN(P) which extends first-order logic by counting and by numerical predicates from a set P, and which can be viewed as a natural generalization of various counting logics that have been studied in the literature.

We obtain a locality result showing that every FOCN(P)-formula can be transformed into a formula in Hanf normal form that is equivalent on all finite structures of degree at most d. A formula is in Hanf normal form if it is a Boolean combination of formulas describing the neighbourhood around its tuple of free variables and arithmetic sentences with predicates from P over atomic statements describing the number of realizations of a type with a single center. The transformation into Hanf normal form can be achieved in time elementary in d and the size of the input formula.

From this locality result, we infer the following applications:

(1) The Hanf-locality rank of first-order formulas of bounded quantifier alternation depth only grows polynomially with the formula size.

(2) The model checking problem for the fragment FOC(P) of FOCN(P) on structures of bounded degree is fixed-parameter tractable (with elementary parameter dependence).

(3) The query evaluation problem for fixed queries from FOC(P) over fully dynamic databases of degree at most d can be solved efficiently: there is a dynamic algorithm that can enumerate the tuples in the query result with constant delay, and that allows to compute the size of the query result and to test if a given tuple belongs to the query result within constant time after every database update.

The State of HFL Model-Checking

Tuesday, 14:45–15:10

by *Florian Bruse*, *Universität Kassel*

Higher-Order Modal Fixpoint Logic (HFL) is an extension of the modal mu-calculus. We give an overview over the state of the HFL model-checking problem, with focus on operational semantics for HFL. We also briefly sketch translations between HFL model-checking and the model-checking problem for higher-order recursion schemes. This summarizes results presented at GandALF 2016 and POPL 2017.

A Modal Characterisation Theorem for Common Knowledge Logic

Tuesday, 15:10–15:35
by *Felix Canavoi*, *TU Darmstadt*

We investigate epistemic modal logic with common knowledge modalities for groups of agents and obtain two van Benthem style model-theoretic characterisation theorems, a classical version and one in the sense of finite model theory. In order to do that we show that basic modal logic is expressively equivalent to the bisimulation invariant fragment of first-order logic over the non-elementary class of so-called (finite) common knowledge Kripke frames. The transitive nature of common knowledge modalities poses a new challenge that cannot be tackled by the usual locality based techniques. For (finite) common knowledge Kripke frames we construct (finite) bisimilar companions that are based on Cayley graphs of suitable groups with specific acyclicity properties. These Cayley graphs possess a very intricate yet highly regular structure that allows for a locality analysis on multiple scales that are induced by distance measures w.r.t. various coalitions of agents.

A fine-grained hierarchy of complete problems in the separated fragment of first-order logic

Tuesday, 15:35–16:00
by *Marco Voigt*, *Max-Planck-Institut für Informatik*

Recently, the separated fragment (SF) has been introduced and proved to be decidable [1]. Its defining principle is that universally and existentially quantified variables may not occur together in atoms. Topmost existential quantifier blocks are exempt from this rule. SF contains every Bernays–Schönfinkel–Ramsey sentence and every monadic first-order sentence. The known upper bound on the time required to decide SF's satisfiability problem is formulated in terms of quantifier alternations: Given an SF sentence $\exists\vec{\mathbf{z}}\forall\vec{\mathbf{x}}_1\exists\vec{\mathbf{y}}_1\ldots\forall\vec{\mathbf{x}}_n\exists\vec{\mathbf{y}}_n.\psi$ in which ψ is quantifier free, satisfiability can be decided in nondeterministic n-fold exponential time. In this talk, a more fine-grained analysis of the complexity of SF-satisfiability shall be presented. One can derive an upper and a lower bound in terms of *the degree ∂ of interaction of existential variables (short: degree)*—a measure of how many separate existential quantifier blocks in a sentence are connected via joint occurrences of variables in atoms. The satisfiability problem for the set $\mathrm{SF}_{\partial\leq k}$ of all SF sentences that have degree k or smaller is complete for k-NExpTime. Consequently, SF-satisfiability is non-elementary in general, since SF is defined without restrictions on the degree. This refined upper bound yields the well-known NExpTime-completeness result for the satisfiability for monadic first-order sentences, since monadic sentences exhibit a degree of $\partial = 1$, in spite of the quantifier alternations that may occur.

The old analysis results in an upper bound that becomes significantly more inaccurate with every additional $\forall\exists$-alternation.

Capturing Polynomial Time using Modular Decomposition

Tuesday, 16:15–16:40

by *Berit Grußien*, *Humboldt-Universität zu Berlin*

The question of whether there is a logic that captures polynomial time is one of the main open problems in descriptive complexity theory and database theory. In 2010 Grohe showed that fixed point logic with counting captures polynomial time on all classes of graphs with excluded minors. We now consider classes of graphs with excluded induced subgraphs. For such graph classes, an effective graph decomposition, called modular decomposition, was introduced by Gallai in 1976. The graphs that are non-decomposable with respect to modular decomposition are called prime. We present a tool, the Modular Decomposition Theorem, that reduces (definable) canonization of a graph class C to (definable) canonization of the class of prime graphs of C that are colored with binary relations on a linearly ordered set. By an application of the Modular Decomposition Theorem, we show that fixed point logic with counting also captures polynomial time on the class of permutation graphs. As a side effect of the Modular Decomposition Theorem, we further obtain that the modular decomposition tree is computable in logarithmic space. It follows that cograph recognition and cograph canonization is computable in logarithmic space.

Model-Checking for Successor-Invariant First-Order Formulas on Graph Classes of Bounded Expansion

Tuesday, 16:40–17:05

by *Roman Rabinovich*, *TU Berlin*

A successor-invariant first-order formula is a formula that has access to an auxiliary successor relation on a structure's universe, but the model relation is independent of the particular interpretation of this relation. It is well known that successor-invariant formulas are more expressive on finite structures than plain first-order formulas without a successor relation. This naturally raises the question whether this increase in expressive power comes at an extra cost to solve the model-checking problem, that is, the problem to decide whether a given structure together with some (and hence every) successor relation is a model of a given formula.

It was shown earlier that adding successor-invariance to first-order logic essentially comes at no extra cost for the model-checking problem on classes of finite structures whose underlying Gaifman graph is planar, excludes a fixed minor or a fixed topological minor (presented at AlmoTh 2016 by Kord Eickmeyer). In this work we show that the model-checking problem for successor-invariant formulas is fixed-parameter tractable on any class of finite structures whose underlying Gaifman graphs form a class of bounded expansion. Our result generalizes all earlier results and comes close to the best tractability results on nowhere dense classes of graphs currently known for plain first-order logic.

Polynomial Kernels and Wideness Properties of Nowhere Dense Graph Classes

Tuesday, 17:05–17:30

by *Sebastian Siebertz*, *TU Berlin*

Nowhere dense classes of graphs are very general classes of uniformly sparse graphs with several seemingly unrelated characterisations. From an algorithmic perspective, a characterisation of these classes in terms of uniform quasi-wideness , a concept originating in finite model theory, has proved to be particularly useful. Uniform quasi-wideness is used in many fpt-algorithms on nowhere dense classes. However, the existing constructions showing the equivalence of nowhere denseness and uniform quasi-wideness imply a non-elementary blow up in the parameter dependence of the fpt-algorithms, making them infeasible in practice. As a first main result of this paper, we use tools from logic, in particular from a sub-field of model theory known as stability theory, to establish polynomial bounds for the equivalence of nowhere denseness and uniform quasi-wideness. A powerful method in parameterized complexity theory is to compute a problem kernel in a pre-computation step, that is, to reduce the input instance in polynomial time to a sub-instance of size bounded in the parameter only (independently of the input graph size). Our new tools allow us to obtain for every fixed value of r a polynomial kernel for the distance-r dominating set problem on nowhere dense classes of graphs. This result is particularly interesting, as it implies that for every class C of graphs which is closed under subgraphs, the distance-r dominating set problem admits a kernel on C for every value of r if, and only if, it admits a polynomial kernel for every value of r (under the standard assumption of parameterized complexity theory that FPT is distinct from W[2]).

Property Testing for structures of bounded degree

Tuesday, 17:45-18:10

by *Isolde Adler*, *University of Leeds*

Property testing (for a property P) asks for a given input, whether it has property P, or is ”far” from having that property. A ”testing algorithm” is a probabilistic algorithm that answers this question with high probability correctly, by only looking at small parts of the input. Testing algorithms are thought of as ”extremely efficient”, making them relevant in the context of big data.

We extend the bounded degree model of property testing from graphs to relational structures, and we discuss the model. We show that every property definable in first-order logic is testable with a constant number of queries in polylogarithmic time. On structures of bounded tree-width, a similar statement holds for monadic second-order logic.

This is joint work with Frederik Harwath.

Learning first-order definable concepts over structures of small degree

Tuesday, 18:10–18:35

by *Martin Ritzert*, *RWTH Aachen University*

We consider a declarative framework for machine learning where concepts and hypotheses are defined by formulas of a logic over some "background structure". We show that within this framework, concepts defined by first-order formulas over a background structure of at most polylogarithmic degree can be learned in polylogarithmic time in the "probably approximately correct" learning sense.

Static Analysis of Monadic Datalog on Finite Labeled Trees

Wednesday, 9:00–9:25

by *Andre Frochaux*, *Humboldt-Universität zu Berlin*

We investigate into the decidability and complexity of the fundamental problems entailed by static analysis of monadic datalog on finite labeled trees using various schemas.

Static analysis is used for optimizing queries without considering concrete database instances but exploiting information about the represented structure. Static analysis relies on three basic decision problems. First, the emptiness problem, whose task is to decide whether a query returns the empty result on every database. Second, the equivalence problem, asking if the result of two given queries always coincides on every database. And finally, the query containment problem, where it is to decide whether on every database a given query produces a subset of the results of a second given query. We are interested in finding out whether these problems are decidable and, if so, what their complexity is.

We show that, for monadic datalog on finite (un)ranked labeled trees, the mentoined problems are complete for 2EXPTIME when (a) considering unordered trees using the axes child and descendant, and when (b) considering ordered trees using the axes firstchild, nextsibling, child, and descendant. When omitting the descendant-axis, we obtain that in both cases the problems are EXPTIME-complete.

Finally, we briefly investigate into the aforementioned problems of monadic datalog under bag-(set-)semantics.

Provenance Analysis in Logic and Games

Wednesday, 9:25–9:50

by *Erich Grädel*, *RWTH Aachen University*

Provenance analysis for database transformations is used to track the provenance or dependence of computed facts from different input items. For positive query languages, such as (unions of) conjunctive queries or datalog, it has been shown that provenance analysis can been done via interpretations in commutative semirings, to answer questions about the trust in, the cost of, or the required clearance level for derived facts, or the number of derivation trees that are available for establishing such a fact.

We generalize this analysis to logics that include negation, such as full FO and LFP, exploiting connections between logic and games. Beyond the already familiar applications to query evaluation (and hence logic), provenance analysis also has interesting interpretations in finite and infinite games to answer more subtle questions than just who wins the game, such the number or costs of winning strategies, or issues such as confidence and trust in game-theoretic settings. The mathematical basis of this approach are interpretations of logics and games in semirings of polynomials (or power series).

This is joint work with Val Tannen.

Enumeration of MSO Queries on Strings with Constant Delay and Logarithmic Updates

Wednesday, 9:50–10:15

by *Matthias Niewerth*, *Universität Bayreuth*

We consider the enumeration of MSO queries over strings under updates. For each MSO query we build an index structure enjoying the following properties: The index structure can be constructed in linear time, it can be updated in logarithmic time and it allows for constant delay time enumeration.

This improves from the previous known index structures allowing for constant delay enumeration that would need to be reconstructed from scratch, hence in linear time, in the presence of updates.

We allow relabeling updates, insertion of individual labels and removal of individual labels.

Answering Conjunctive Queries under Updates

Wednesday, 10:45–11:10

by *Christoph Berkholz*, *Humboldt-Universität zu Berlin*

In this talk I present a recent joint work with Jens Keppeler and Nicole Schweikardt on the algorithmic task of answering conjunctive queries on dynamic databases. Our main result is a dichotomy theorem that precisely characterises those conjunctive queries that can be evaluated efficiently.

This work will be appear at PODS 2017 and the more detailed abstract of the paper can be found below.

We consider the task of enumerating and counting answers to k-ary conjunctive queries against relational databases that may be updated by inserting or deleting tuples. We exhibit a new notion of q-hierarchical conjunctive queries and show that these can be maintained efficiently in the following sense. During a linear time preprocessing phase, we can build a data structure that enables constant delay enumeration of the query results; and when the database is updated, we can update the data structure and restart the enumeration phase within constant time. For the special case of self-join free conjunctive queries we obtain a dichotomy: if a query is not q-hierarchical, then query enumeration with sublinear delay and sublinear update time (and arbitrary preprocessing time) is impossible. For Boolean conjunctive queries and the more general problem of counting the number of solutions of a k-ary query we obtain complete dichotomies: if the query's homomorphic core is q-hierarchical, then size of the the query result can be computed in linear time and maintained with constant update time. Otherwise, the size of the query result cannot be maintained with sublinear update time. All our lower bounds rely on the OMv-conjecture, a conjecture on the hardness of online matrix-vector multiplication that has recently emerged in the field of fine-grained complexity to characterise the hardness of dynamic problems. The lower bound for the counting problem additionally relies on the orthogonal vectors conjecture, which in turn is implied by the strong exponential time hypothesis.

Answering FO+MOD queries under updates on bounded degree databases

Wednesday, 11:10–11:35

by *Jens Keppeler*, *Humboldt-Universität zu Berlin*

In this talk I will present joint work with Christoph Berkholz and Nicole Schweikardt about answering FO+MOD queries under updates on bounded degree databases (to appear in Proc. ICDT 2017).

We investigate the query evaluation problem for fixed queries over fully dynamic databases, where tuples can be inserted or deleted. The task is to design a dynamic algorithm that immediately reports the new result of a fixed query after every database update.

We consider queries in first-order logic (FO) and its extension with modulo-counting quantifiers (FO+MOD), and show that they can be efficiently evaluated under updates, provided that the dynamic database does not exceed a certain degree bound.

In particular, we construct a data structure that allows to answer a Boolean FO+MOD query and to compute the size of the result of a non-Boolean query within constant time after every database update. Furthermore, after every update we are able to immediately enumerate the new query result with constant delay between the output tuples. The time needed to build the data structure is linear in the size of the database.

Our results extend earlier work on the evaluation of first-order queries on static databases of bounded degree and rely on an effective Hanf normal form for FO+MOD recently obtained by Heimberg, Kuske, and Schweikardt (LICS 2016).

Connecting AC^1 and Dynamic Descriptive Complexity

Wednesday, 11:35–12:00

by *Nils Vortmeier*, *TU Dortmund*

A dynamic program, as introduced by Dong, Su and Topor and Patnaik and Immerman, maintains the result of a fixed query for an input database which is subject to changes. It can use an auxiliary database whose relations are updated via first-order formulas upon modifications of the input database.

Although the update mechanism is based on first-order logic, which corresponds to the circuit complexity class AC^0, in this talk we discuss two recent approaches to utilizing AC^1 algorithms. At first we explain how all AC^1 queries can be maintained, if changes to the database are specified by first-order queries without parameters. For this we show that a dynamic program can prepare for all possible future changes. Then we show how a dynamic program may use auxiliary data constructed by AC^1 algorithms. For this, a dynamic program has to "catch up" with changes that have already happened, but have not yet been processed. As an example, with this approach we are able to show a dynamic version of Courcelle's Theorem: all MSO definable properties can be maintained for graphs with bounded treewidth. This talk presents joint work with Thomas Schwentick, Thomas Zeume, Samir Datta, and Anish Mukherjee.

The Descriptive Complexity of Solving Linear Equation Systems and its Applications

Wednesday, 12:15–12:40

by *Martin Grohe*, *RWTH Aachen University*

We prove that the solvability of systems of linear equations and related linear algebraic properties are definable in a fragment of fixed-point logic with counting that only allows polylogarithmically many iterations of the fixed-point operators. This allows us to separate the descriptive complexity of solving linear equations from full fixed-point logic with counting by logical means. As an application of these results, we separate an extension of first order logic with a rank operator from fixed-point logic with counting, solving an open problem due to Holm [PhD thesis, Cambridge University, 2010].

We then draw a connection between this work in descriptive complexity theory to graph isomorphism testing and propositional proof complexity. Answering an open question from [Berkholz and Grohe, ICALP 2015], we separate the strength of certain algebraic graph-isomorphism tests. This result can also be phrased as a separation of the algebraic propositional proof systems "Nullstellensatz" and "monomial PC".

This is joint work with Wied Pakusa.

An exponential lower bound for Individualization-Refinement algorithms for Graph Isomorphism

Wednesday, 12:40–13:05

by *Daniel Neuen*, *RWTH Aachen University*

The individualization-refinement paradigm provides a strong toolbox for testing isomorphism of two graphs and indeed, the currently fastest implementations of isomorphism solvers all follow this approach. While these solvers are fast in practice, from the theoretical point of view, no general lower bounds concerning the worst case complexity of these tools are known. In fact, it is an open question whether individualization-refinement algorithms can achieve upper bounds on the running time similar to the more theoretical techniques based on a group theoretic approach. In this work we give a negative answer to this question and construct a family of graphs on which algorithms based on the individualization-refinement paradigm require exponential time. Contrary to a previous construction of Miyazaki, that only applies to a specific algorithm within the individualization-refinement framework, our construction is immune to changing the cell selector, or adding various heuristic invariants to the algorithm. Furthermore, our graphs also provide exponential lower bounds in the case when the k-dimensional Weisfeiler Leman algorithm is used to replace the standard color refinement operator and the arguments even work when the entire automorphism group of the inputs is initially provided to the algorithm.

References

[1] T. Sturm, M. Voigt, and C. Weidenbach. Deciding first-order satisfiability when universal and existential variables are separated. *Logic in Computer Science (LICS'16)*, page 86–95, 2016.