Vorläufiges Programm & Abstracts zum 65. Theorietag in Paderborn

Englische Vorträge sind blau hinterlegt.
English talks are indicated with a blue background.

Donnerstag, 14. Februar 2013

ab 10:00	Ankunft + Registrierung + Kaffee	
11:00	Willkommensgruß	Friedhelm Meyer auf der Heide (Paderborn)
Session 1	Chair: TBA	
11:15	CONE-Hashing: A distributed self-stabilizing algorithm for a heterogeneous storage system	Sebastian Kniesburges (Paderborn)
11:45	Combining Network Creation Games and Self-Optimizing Binary Search Trees – Results and Struggles	Thim Strothmann (Paderborn)
12:15	Closed Sets and Operators thereon: Representations, Computability and Complexity	Carsten Rösnick (TU Darmstadt)
13:00	60 min Mittagspause	
Keynote 1	Chair: TBA	
14:00 (english)	Arithmetic circuits: challenges and new developments	Peter Bürgisser (Paderborn)
15:00	15 min Pause	
Session 2	Chair: TBA	
15:15 (english)	The Delta-ball game and the diameter of polyhedra	Nicolai Hähnle (TU Berlin)
15:45	How applying Myhill-Nerode methods to hypergraphs helps mastering the Art of Trellis Decoding	René van Bevern (TU Berlin)
16:15 (english)	Greedy Selfish Network Creation	Pascal Lenzner (HU Berlin)

16:45	30 min Pause	
Session 3	Chair: TBA	
17:15	Analyzing Randomly Placed Multiple Antennas for MIMO Wireless Communication	Thomas Janson (Freiburg)
17:45	Partikelschwarmoptimierung findet fast sicher ein lokales Optimum	Manuel Schmitt (Erlangen-Nürnberg)
18:15 (english)	Advances in Localized Topology Control Algorithms for Wireless Networks	Florentin Neumann (Koblenz-Landau)
20:00	Socializing (Paderborner Brauhaus)	

Freitag, 15. Februar 2013

Keynote 2	Chair: Christian Scheideler	
09:15	An Overview about Graph Sparsification	Harald Räcke (TU München)
10:15	30 min Pause + Gruppenfoto	
Session 4	Chair: TBA	
10:45 (english)	IRIS: A Robust Information System Against Insider DoS-Attacks	Martina Eikel (Paderborn)
11:15	Combinatorial Feature Selection Problems: Parameterized Algorithms and Complexity	Vincent Froese (TU Berlin)
11:45	On the Parameterized and Approximation Hardness of Metric Dimension	André Nichterlein (TU Berlin)
12:15 (english)	Profitable Scheduling via Speed Scaling	Peter Kling (Paderborn)
12:45	Abschlussworte	Friedhelm Meyer auf der Heide (Paderborn)

Abstracts

How applying Myhill-Nerode methods to hypergraphs helps mastering the Art of Trellis Decoding

Speaker: René van Bevern, Technische Universität Berlin

Abstract: A trellis is a graph associated with a linear code that is used for maximum-likelihood decoding. The decoding complexity of a linear code is strongly influenced by the state complexity of the trellis, which highly depends on the coordinate permutation of the linear code. The problem of finding the coordinate permutation of a linear code such that the state-complexity of the associated trellis is at most k has been referred to as the Art of Trellis Decoding and is NP-hard. We show that it is linear-time solvable if k is constant. To this end, we reduce the problem to Hypergraph Cutwidth, which we show to be fixed-parameter linear by, for the first time, applying an analog of the Myhill-Nerode theorem from formal language theory to a hypergraph problem. In contrast, we also exploit Myhill-Nerode methods to show that for instances whose incidence graphs have bounded treewidth, Hypertree Width and variants cannot be solved by simple finite tree automata.

Arithmetic circuits: challenges and new developments

Speaker: Peter Bürgisser, Universität Paderborn

Abstract: The complexity of a polynomial is defined as the minimal number of arithmetic operations required to build it up from the variables and constants. Proving lower bounds on the complexity of specific polynomials is very challenging: the most famous question being the permanent versus determinant problem. It turned out that for its study, one may focus on arithmetic circuits of depth four. In the talk we shall survey two recent approaches on this that lead to new and fascinating mathematical questions.

For univariate polynomials, depth four circuits express a polynomial as a sum of products of sparse polynomials of small size. How many real zeros may such a polynomial have? Pascal Koiran discovered a surprising connection: if the number of real zeros must be necessary small (at most polynomial in the circuit size), then the permanent is hard (its complexity is superpolynomial in the number of variables). The corresponding question of real algebra is wide open.

In another direction, Gupta, Kamath, Kayal and Saptharishi recently proved an exponential lower bound on depth four circuits computing the permanent or determinant. This is based on quantifying the amount of algebraic relations between the higher order partial derivatives of the polynomial under question. While the bound obtained for the determinant seems to close to the optimum, there seems to be room for improvement for the permanent. For this, understanding the Hilbert function of the ideal of permanental minors is required. (One has a good understanding of the Hilbert function of determinantal minors.)

IRIS: A Robust Information System Against Insider DoS-Attacks

Speaker: Martina Eikel, Universität Paderborn

Abstract: In this work we present the first scalable distributed information system, called IRIS, that is provably robust against Denial-of-Service (DoS) attacks by a \emph{current insider}. We allow a current insider to have \emph{complete} knowledge about the information system and to have the power to block \emph{any ε-fraction} of its servers by a DoS attack, where ε can be chosen up to a constant.

Previously, only solutions against DoS-attacks of past insiders were known, where a past insider only has complete knowledge about some {\em past} time point t_0 of the information system. Scheideler et al. (DISC 2006, SPAA2009) showed that in this case it is possible to design an information system so that any information that was last updated {\em after} t_0 is safe against a DoS-attack. But their constructions would not work at all for a current insider.

The key idea behind our system is to make extensive use of coding. More precisely, we present two alternative distributed coding strategies with a low redundancy that make it impossible even for a current insider to prevent access to any data item even when blocking an ε-fraction of servers of its choice.

Combinatorial Feature Selection Problems: Parameterized Algorithms and Complexity

Speaker: Vincent Froese, Technische Universität Berlin

Abstract: The term combinatorial feature selection refers to a general class of problems that, given a set of high-dimensional objects, ask for selecting a subset of dimensions (features) such that some desired property holds for the dataset restricted to the selected dimensions. Problems of this kind arise as data preprocessing tasks in areas such as machine learning. Charikar et al. defined a general framework for combinatorial feature selection problems comprising cluster analysis as well as dimension reduction. We consider the Hidden Cluster problem, where the goal is to eliminate noisy dimensions such that in the remaining dimensions the data can be divided into a given number of clusters, and the Distinct Vectors problem, which aims at finding a minimum number of dimensions such that all given points are still distinct. Moreover, we introduce the Hidden Cluster Graph problem, where the number of cluster centers is not known. We analyze the above problems in terms of their parameterized complexity and provide parameterized hardness results as well as fixed-parameter algorithms. For example, we show that they are hard to solve with respect to some natural parameters such as the number of dimensions to select or the number of dimensions to delete. In order to obtain fixed-parameter tractability, we focus on parameters such as the number of cluster centers, the radius of a cluster, the size of the alphabet, or the maximum pairwise Hamming distance of the data points. We show that the problems are fixed-parameter tractable for some combinations of the mentioned parameters by providing problem kernels and fixed-parameter tractable algorithms.

The Delta-ball game and the diameter of polyhedra

Speaker: Nicolai Hähnle, Technische Universität Berlin

Abstract: The polynomial Hirsch conjecture, motivated by empirical observations in linear programming, is one of the fundamental open problems in the study of polyhedra. It states that the diameter of the (vertex-edge)-graph of any d-dimensional polyhedron with n facets is bounded by a polynomial in d and n. The best known upper bound is n^{1 + log d}, while the best known lower bound construction has a diameter of (1 + epsilon) n for a fixed epsilon > 0.

We (very) briefly outline some approaches to the problem and state some recent results (joint work with Bonifas, Di Summa, Eisenbrand, and Niemeier and with Klee and Pilaud).

Analyzing Randomly Placed Multiple Antennas for MIMO Wireless Communication

Speaker: Thomas Janson, Universität Freiburg

Abstract: We present an analytical approach for determining the signal-to-noise-ratio (SNR) of m multiple antennas in the line-of-sight case. The antennas are placed at random positions within a disc of given diameter d. We characterize the angular signal strength with three sectors: the main beam, the side beams and an area of white Gaussian noise. The SNR and the sector angles depend on d, m, and the wavelength λ. It turns out that for randomized antenna positions the analysis can be reduced to the analysis of a random geometric walk in two dimensions. The angle of the main beam is approximately λ/d with a SNR proportional to $\sqrt{m}$. For the side beams the SNR is proportional to $\hbox{sinc}(2\alpha d/\lambda)$ where α denotes the angle deviating from the target. The range of the side beams is limited to an approximate angle of $\lambda/d\sqrt{m}$. Beyond this angle we observe average white Gaussian noise.

Profitable Scheduling via Speed Scaling

Speaker: Peter Kling, Universität Paderborn

Abstract: We consider a scheduling problem where jobs arrive over time and are preemptable. They have different workloads, values, and deadlines. To process a job's workload until its deadline the scheduler must invest a certain amount of energy. However, the scheduler may also decide not to finish a job but instead to suffer a loss equaling the job's value. The cost of a schedule is the sum of invested energy and lost values. In order to finish a job the scheduler has to determine which processors to use and set their speeds accordingly. A processor's energy consumption is power Power(s) integrated over time, where Power(s)=s^α is the power consumption when running at speed s. Since we consider the online variant of the problem, the scheduler has no knowledge about future jobs. This problem was introduced by Chan et al. [WAOA, 2010] for the case of a single processor. They presented an online algorithm which is $\alpha^\alpha+2e\alpha$-competitive. We provide an online algorithm for the case of multiple processors with an improved competitive ratio of α^α.

Our analysis is significantly different from the typical potential function argument which is dominant in the analysis of online algorithms in this research area. Instead, we make use of a framework recently suggested by Gupta et al. [WAOA, 2012]. It utilizes well-known tools from convex optimization, especially duality theory and primal-dual algorithms. We develop a convex programming formulation and design a greedy primal-dual online algorithm for the problem at hand. Our result shows that this new technique has the potential to yield a cleaner and tighter analysis than classical methods.

CONE-Hashing: A distributed self-stabilizing algorithm for a heterogeneous storage system

Speaker: Sebastian Kniesburges, Universität Paderborn

Abstract: A heterogeneous storage system consists of several hosts having different storage capacities, and the assignment of data to hosts must be done in a way such that the number of data stored at a host is related to its capacity. We study the problem of maintaining a distributed heterogeneous storage system, consisting of heterogeneous nodes in a self-stabilizing manner (i.e. the network has the ability to recover out of any weakly-connected state). In this talk we will present a general model for topological self-stabilization and apply this model on the specific CONE-DHT problem. For the CONE-DHT with n nodes we will show fair load balancing, storage efficiency and further network properties like a low degree (O(log n)) and a low routing distance (O(log n) hops). We will then present a distributed self-stabilizing algorithm that transforms any weakly connected network into a CONE-DHT and show that the amount of structural changes in a stable state due to external dynamics like the joining or leaving or capacity change of a single node is bounded by O(log^2 n).

Greedy Selfish Network Creation

Speaker: Pascal Lenzner, Humboldt-Universität Berlin

Abstract: Network Creation Games attempt to model the creation and evolution of complex networks, which are build in a decentralized way by selfish agents. Each selfish agent tries to balance the two contradicting objectives of maximizing the quality of network usage and minimizing the number of (costly) links that have to be created and maintained to connect to the network. Clearly, the Internet is a prominent example of such a network and it can be considered as an equilibrium state of a Network Creation Game. Here, the agents are Internet Service Providers and links connect different autonomous systems.

I will present insights on the impact of greediness in the Network Creation Games introduced by Fabrikant et al. [PODC'03]. Based on the ideas that ISPs prefer smooth adaptations over radically re-designing their infrastructure and that computing a best possible strategy is NP-hard, I will introduce and analyze a very natural solution concept, called the Greedy Equilibrium.

It turns out that naive greedy play leads to remarkably stable networks in the SUM-version, where agents attempt to minimize their average distance to all other agents. For the MAX-version, where agents attempt to minimize their maximum distance, I will present positive results for tree-networks and a strong negative result for general networks.

Advances in Localized Topology Control Algorithms for Wireless Networks

Speaker: Florentin Neumann, Universität Koblenz-Landau

Abstract: In order to guarantee message delivery, localized geographic routing algorithms for wireless ad hoc networks require the underlying communication graph to be planar. As such network graphs typically are not planar, prior to a routing step, a node needs to locally compute a view on a planar subgraph with small stretch factor to improve the performance of the routing algorithm.

In my talk, recent advances in localized topology control are presented: Under the unit disk graph model, the partial Delaunay triangulation (PDT) is a planar spanner with constant Euclidean stretch factor and it can be computed locally in a reactive manner (i.e., without full knowledge on the node's neighborhood) using an optimal number of messages. Moreover, it is sketched that PDT is also useful for Euclidean spanner construction under the quasi unit disk graph model.

On the Parameterized and Approximation Hardness of Metric Dimension

Speaker: André Nichterlein, Technische Universität Berlin

Abstract: The NP-hard Metric Dimension problem is to decide for a given graph G and a positive integer k whether there is a vertex subset of size at most k that separates all vertex pairs in G. Herein, a vertex v separates a pair {u, w} if the distance (length of a shortest path) between v and u is different from the distance of v and w. We give a polynomial-time computable reduction from the Bipartite Dominating Set problem to Metric Dimension on maximum degree three graphs such that there is a one-to-one correspondence between the solution sets of both problems. There are two main consequences of this:

First, it proves that Metric Dimension on maximum degree three graphs is W[2]-hard with respect to the parameter k. This answers an open question concerning the parameterized complexity of Metric Dimension posed by Lokshtanov [Dagstuhl seminar, 2009] and also by Diaz et al. [ESA'12]. Additionally, it implies that a trivial n^O(k)-time algorithm cannot be improved to an n^o(k)-time algorithm, unless the assumption FPT = W[1] fails.

Second, as Bipartite Dominating Set is inapproximable within o(log n), it follows that Metric Dimension on maximum degree three graphs is also inapproximable by a factor of o(log n), unless NP = P. This strengthens the result of Hauptmann et al. [JDA'12] who proved APX-hardness on bounded-degree graphs.

This is joint work with Sepp Hartung.

An Overview about Graph Sparsification

Speaker: Harald Räcke, Technische Universität München

Abstract: Given an undirected graph $G=(V,E)$ \emph{edge sparsification methods} try to compute a sparse graph $H=(V,E')$ such that H approximates the cuts in G up to a small error term. The existence of fast edge sparsification methods are fundamental for obtaining good running times for many optimization problems on graphs.

Given a graph $G=(V,E)$ together with a subset $T\subseteq V$ of terminal nodes, \emph{vertex sparsifiers} produce a graph $H'=(V',E')$ with $T\subseteq V'$ and $|V'|\approx |V|$ such that H' approximates the cut-structure among the terminals in G up to a small error. These sparsification methods are mainly used for improving the approximation guarantees for graph problems.
In this talk I give an overview of some of the recent advances in the area of graph sparsification.

Closed Sets and Operators thereon: Representations, Computability and Complexity

Speaker: Carsten Rösnick, Technische Universität Darmstadt

Abstract: In Computable Analysis (which goes way back to Turing, and got a thorough mathematical foundation in the subsequent decades), not only encodings of real numbers and functions are considered (cf. the equivalently-named monograph (2000) by K.Weihrauch), but also representations of subsets of, say, IR^d, for fixed dimension d. In this talk will see four (of many) seemingly different ways to encode closed/compact/convex/regular sets via infinite binary strings, how to feed them to an OTM, and how to compute with them operators like union and intersection, point finding and projection, or even the image and the inverse --- always uniform (!) in the input (continuous functions). In 2010, A.Kawamura and S.Cook introduced the notion of second-order polytime (based on ideas going back to K.Mehlhorn ('76) and B.Kapron/S.Cook ('96)). Based on this notion, we will see how to measure the (second--order, parameterized) complexity of the aforementioned operations (and operators), and whether (or not) these representations are (polytime) equivalent.

Partikelschwarmoptimierung findet fast sicher ein lokales Optimum

Speaker: Manuel Schmitt, Friedrich-Alexander-Universität Erlangen-Nürnberg

Abstract: Partikelschwarmoptimierung ist eine von mehreren sehr bekannten, durch Naturphänomene inspirierten Meta-Heuristiken für Black-Box-Optimierung von Funk- tionen auf kontinuierlichem Definitionsbereich. Im Unterschied zu genetischen oder evolutionsbasierten Optimierungsverfahren kooperieren die als Partikel bezeichneten Einheiten miteinander und teilen Informationen über den Suchraum. Obwohl diese Optimierungsmethode weite Anwendungsfelder hat, sind bisher nur Teilaspekte ihrer Funktionsweise formal nachgewiesen. Beispielsweise wurde nachgewiesen, dass unter gewissen Voraussetzungen die Partikel ein kon- vergentes Verhalten aufweisen. Es wird untersucht, ob und unter welchen Bedin- gungen der Partikelschwarm ein lokales Optimum finden kann. Dabei stellt sich heraus, dass für eine allgemeine Klasse von Zielfunktionen eine leichte Modifika- tion des Algorithmus ausreicht, um zu gewährleisten, dass der Partikelschwarm fast sicher ein lokales Optimum findet.

Combining Network Creation Games and Self-Optimizing Binary Search Trees – Results and Struggles

Speaker: Thim Strothmann, Universität Paderborn

Abstract: We propose a new network creation game which utilizes a binary search tree as the network structure and in which selfish nodes want to communicate with only a certain subset of all nodes. Each player of the game can change the structure of the network by using mechanics of self-optimizing binary search trees.

We prove that in many cases the game does not converge and relate convergence to certain structures of the communication interests. Furthermore, we present how good the selfish nodes perform compared to a social optimum in convergent situations by computing the Price of Anarchy.

We use the solution concept of sink equilibria to analyze the game in non-convergent situations. We compute the Price of Sinking for different non-convergent structures of the communication interests and deduce that selfishness can lead to quadratic social cost, although an optimum only has cost of $O(n \cdot \log(n))$. This leads to a Price of Sinking of $O(n/\log(n))$. We complete the talk by discussing variations of the original game.