Program & Abstracts

73. Theorietag in Hamburg

Thursday, 18 May 2017

9:25 – 9:30	*Welcome*	
Session 1		
09:30 – 10:00	Improved Approximation for Two Dimensional Strip Packing with Polynomial Bounded Width	Rau Malin (Kiel)
10:00 – 10:30	Online Resource Minimization	Nicole Megow (Bremen)
10:30 – 11:00	New Structural Results for Bin Packing with a Constant Number of Item Types	Kim-Manuel Klein (Kiel)
Coffee Break (30 min)		
Session 2		
11:30 – 12:00	Efficient Best-Response Computation for Strategic Network Formation under Attack	Pascal Lenzner (Potsdam)
12:00 – 12:30	Approximation Schemes for Scheduling Jobs on Identical and Uniform Machines	Klaus Jansen (Kiel)
12:30 – 13:00	On minimizing makespan when some jobs cannot be assigned on the same machine	Das Syamantak (Bremen)
Joint Lunch (60 min)		
Keynote 1		
14:00 – 15:00	Relays: A new interconnection model for overlay networks	Christian Scheideler (Paderborn)
Coffee Break (30 min)		
Session 3		

15:30 – 16:00	A Distributed Algorithm for Finding Hamiltonian Cycles in Random Graphs	Volker Turau (Hamburg)
16:00 – 16:30	Parameterized Approximation Algorithms for some Location Problems in Graphs	Arne Leitert (Kent, Ohio)
Social Program & Dinner		

Friday, 19 Mai 2017

Keynote 2		
09:00 – 10:00	Attribute-Based Encryption and Information-Theoretic Crypto	Hoeteck Wee (Paris)
Coffee Break (30 min)		
Session 4		
10:30 – 11:00	Revenue maximization in Stackelberg Pricing	Toni Böhnlein (Köln)
11:00 – 11:30	Continuous Speed Scaling with Variability	Sören Riechers (Paderborn)
11:30 – 12:00	Recognising Multidimensional Euclidean Preferences	Dominik Peters (Oxford)
Joint Lunch + Group Photo (60 min)		
Session 5		
13:00 – 13:30	TeamLTL	Arne Meier (Hannover)
13:30 – 14:00	Forming Tile Shapes with a Single Robot	Robert Gmyr (Paderborn)
14:00 – 14:30	Distributed Monitoring of Network Properties: The Power of Hybrid Networks	Kristian Hinnenthal (Paderborn)
Coffee Break (30 min)		
Session 6		
15:00 – 15:30	Completion Probabilities and Parallel Restart Strategies under an Imposed Deadline	Jan-Hendrik Lorenz (Ulm)

15:30 – 16:00	Routing in Hybrid Communication Networks with Holes - The Distributed Convex Hull Algorithm	Christina Kolb (Paderborn)
16:00 – 16:30	On the computational complexity of super-resolution imaging in discrete tomography	Andreas Alpers (München)
16:30 – 16:35	*Closing*	

Abstracts Keynote Talks

Relays: A new interconnection model for overlay networks

Speaker: Christian Scheideler, Institut für Informatik, Universität Paderborn
Abstract: Research on overlay networks has used a very simple interconnection model in the past: any node can send messages to any other node it knows about, and that knowledge can be freely transferred from node to node. However, as a consequence, problems like safely leaving an overlay network or protecting an overlay network against adversarial behavior cannot be solved without various restrictions on the network state or the power of adversarial nodes.
In order to avoid these restrictions, I propose a new interconnection model based on relays. The basic idea of that model is to give the nodes much more control over their connections. Though the relay concept appears to be quite restrictive, I will show that it preserves universality in a sense that, in principle, one can get from any weakly connected topology to any other weakly connected topology using that concept. I will also demonstrate that the relay concept allows nodes to leave the network without risking disconnectivity while this guarantee cannot be provided with the commonly used interconnection model.

Attribute-Based Encryption and Information-Theoretic Crypto

Speaker: Hoeteck Wee, Computer Science Department, École normale supérieure
Abstract: Can we encrypt data while enabling fine-grained access control and selective computation? In this talk, we will survey how addressing this question led to new connections and questions in information-theoretic cryptography.

Abstracts Contributed Talks

On the computational complexity of super-resolution imaging in discrete tomography

Speaker: Andreas Alpers, Zentrum Mathematik, Technische Universität München
Abstract: Super-resolution imaging aims at improving the resolution of an image by enhancing it with other images or data that might have been acquired using different imaging techniques or modalities. Motivated by applications in plasma physics, we consider the task of doubling the resolution of tomographic grayscale images of binary objects by fusion with double-resolution tomographic data that has been acquired from two viewing angles. We show that this task is polynomial-time solvable if the gray levels have been reliably determined. The task becomes NP-hard if the gray levels of some pixels come with an error of ±1 or larger. The NP-hardness persists for any larger resolution enhancement factor. This means that noise does not only affect the quality of a reconstructed image but, less expectedly, also the algorithmic tractability of the inverse problem itself.

(This is joint work with Peter Gritzmann.)

Revenue maximization in Stackelberg Pricing

Speaker: Toni Böhnlein, Institut für Informatik, Universität zu Köln
Abstract: In a Stackelberg Pricing Game a distinguished player, the leader, chooses prices for a set of items, and the other players, the followers, each seeks to buy a minimum cost feasible subset of the items. The goal of the leader is to maximize her revenue, which is determined by the sold items and their prices. Most previously studied cases of such games can be captured by a combinatorial model where we have a base set of items, some with fixed prices, some priceable, and constraints on the subsets that are feasible for each follower. In this combinatorial setting, Briest et al. and Balcan et al. independently showed that the maximum revenue can be approximated to a factor of H k ~ log k, where k is the number of priceable items.

Our results are twofold. First, we strongly generalize the model by letting the follower minimize any continuous function plus a linear term over any compact subset of R n ≥0 ; the coefficients (or prices) in the linear term are chosen by the leader and determine her revenue. In particular, this includes the fundamental case of linear programs. We give a tight lower bound on the rev- enue of the leader, generalizing the results of Briest et al. and Balcan et al. Besides, we prove that it is strongly NP-hard to decide whether the optimum revenue exceeds the lower bound by an arbitrarily small factor. Second, we study the parameterized complexity of computing the optimal revenue with respect to the number k of priceable items. In the combinatorial setting, given an efficient algorithm for optimal follower solutions, the maximum revenue can be found by enumerating the 2 k subsets of priceable items and computing optimal prices via a result of Briest et al., giving time O(2 k |I|

c) where |I| is the input size. Our main result here is a W[1]-hardness proof for the case where the followers minimize a linear program, ruling out running time f (k)|I| c unless FPT = W[1] and ruling out time |I| o(k) under the Exponential-Time Hypothesis.

On minimizing makespan when some jobs cannot be assigned on the same machine

Speaker: Syamantak Das, Universität Bremen
Abstract: We study the classical scheduling problem of assigning jobs to machines in order to minimize the makespan. It is well-studied and admits an EPTAS on identical machines and a (2-1/m)-approximation algorithm on unrelated machines. In this paper we study a variation in which the input jobs are partitioned into bags and no two jobs from the same bag are allowed to be assigned on the same machine. Such a constraint can easily arise, e.g., due to system stability and redundancy considerations. Unfortunately, as we demonstrate in this paper, the techniques of the above results break down in the presence of these additional constraints. Our first result is a PTAS for the case of identical machines. It enhances the methods from the known (E)PTASs by a finer classification of the input jobs and careful argumentations why a good schedule exists after enumerating over the large jobs. For unrelated machines, we prove that there can be no $(\log n)^{1/4-\epsilon}$-approximation algorithm for the problem for any \epsilon>0, assuming that NP is a subset of ZPTIME}(2^{(\log n)^{O(1)}})$. This holds even in the restricted assignment setting. However, we identify a special case of the latter in which we can do better: if for each bag all its jobs can be assigned to the same set of machines we give an 8-approximation algorithm. It is based on rounding the LP-relaxation of the problem in phases and adjusting the residual fractional solution after each phase in to order to respect the bag constraints.

Forming Tile Shapes with a Single Robot

Speaker: Robert Gmyr, Institut für Informatik, Universität Paderborn
Abstract: We investigate the problem of shape formation with robots on tiles in which a collection of robots has to rearrange a set of movable tiles to form a desired shape. In this preliminary work we consider the case of a single robot operating on an arbitrary number of tiles and present first results towards the formation of simple shapes. Our ultimate goal is to investigate how multiple robots can cooperate to speed up the process of shape formation.

Distributed Monitoring of Network Properties: The Power of Hybrid Networks

Speaker: Kristian Hinnenthal, Institut für Informatik, Universität Paderborn
Abstract: We initiate the study of network monitoring algorithms in a class of hybrid networks in which the nodes are connected by an external network and an internal network

(as a short form for externally and internally controlled network). While the external network lies outside of the control of the nodes (or in our case, the monitoring protocol running in them) and might be exposed to continuous changes, the internal network is fully under the control of the nodes. As an example, consider a group of users with mobile devices having access to the cell phone infrastructure. While the network formed by the WiFi connections of the devices is an external network (as its structure is not necessarily under the control of the monitoring protocol), the connections between the devices via the cell phone infrastructure represent an internal network (as it can be controlled by the monitoring protocol). Our goal is to continuously monitor properties of the external network with the help of the internal network. We present scalable distributed algorithms that efficiently monitor the number of edges, the average node degree, the clustering coefficient, the bipartiteness, and the weight of a minimum spanning tree. Their performance bounds demonstrate that monitoring the external network state with the help of an internal network can be done much more efficiently than just using the external network, as is usually done in the literature.

Approximation Schemes for Scheduling Jobs on Identical and Uniform Machines

Speaker: Klaus Jansen, Institut für Informatik, Universität Kiel
Abstract: Makespan scheduling on identical (and uniform) machines is one of the most basic and fundamental scheduling problem. It asks for an assignment of n jobs to a set of m identical (or uniform) machines that minimizes the makespan (the length of the schedule). The problem is strongly NP-hard even for identical machines, and thus we do not expect a $(1+\epsilon)$-approximation algorithm running in time that depends polynomially on $1/\epsilon$. Furthermore, Chen, Jansen and Zhang recently showed that a running time of $2^{(1/\epsilon)^{1-\delta}} + poly(n)$ of an approximation scheme for identical machines for any $\delta>0$ would imply that the Exponential Time Hypothesis (ETH) for 3-SAT fails. A long sequence of algorithms have been developed to obtain lower dependencies on $1/\epsilon$; the better of which achieves a running time of $2^{\tilde{O}(1/\epsilon^2)}+O(n\log n)$. In this talk we present an algorithm with an improved running time of $2^{\tilde{O}(1/\epsilon)}+O(n\log n)$, which is tight under ETH up to logarithmic factors on the exponent. This is joint work with Kim-Manuel Klein (Kiel) und Jose Verschae (Santiago).

New Structural Results for Bin Packing with a Constant Number of Item Types

Speaker: Kim-Manuel Klein, Institut für Informatik, Universität Kiel
Abstract: We consider the bin packing problem with d different item sizes and revisit the structure theorem given by Goemans and Rothvoß [GR2015] about solutions of the integer cone. We present new techniques on how solutions can be modified and give a new structure theorem that relies on the set of vertices of the underlying integer polytope. As a result of our new structure theorem, we obtain an algorithm for the bin packing problem with

running time |V|^{2^{O(d)}} \cdot enc(I)^{O(1)}, where V is the set of vertices of the integer knapsack polytope and enc(I) is the encoding length of the bin packing instance. The algorithm is fixed parameter tractable, parameterized by the number of vertices of the integer knapsack polytope |V|. This shows that the bin packing problem can be solved efficiently when the underlying integer knapsack polytope has an easy structure, i.e. has a small number of vertices.

Furthermore, we show that the presented bounds of the structure theorem are asymptotically tight. We give a construction of bin packing instances using new structural insights and classical number theoretical theorems which yield the desired lower bound.

Routing in Hybrid Communication Networks with Holes – The Distributed Convex Hull Algorithm

Speaker: Christina Kolb, Institut für Informatik, Universität Paderborn
Abstract: A hybrid communication model consists of a basic ad hoc network of smartphones and a cellular infrastructure, which is a special subset of those smartphones. Whereas the basis ad hoc network is fundamental to enable WLAN connections between smartphones, the celluar infrastructure grants control information over the basic ad hoc network. This kind of information is for example the knowledge of the geographic position and the structure of radio holes. Since one of the aims respective to routing in hybrid communication networks is to avoid routing into radio holes, we are interested in calculating the convex hulls of them.

In this talk, we present the main idea of our distributed convex hull algorithm. This algorithm is faster than related distributed convex hull algorithms.

Parameterized Approximation Algorithms for some Location Problems in Graphs

Speaker: Arne Leitert, Department of Computer Science, Kent State University
Abstract: We develop efficient parameterized, with additive error, approximation algorithms for the (Connected) r-Domination problem and the (Connected) p-Center problem for unweighted and undirected graphs. Given a graph~G, we show how to construct a (connected) $\big(r + \mathcal{O}(\mu) \big)$-dominating set~$D$ with $|D| \leq |D^*|$ efficiently. Here, D^* is a minimum (connected) r-dominating set of~G and μ is our graph parameter, which is the \emph{tree-breadth} or the \emph{cluster diameter in a layering partition} of~G. Additionally, we show that a $+\mathcal{O}(\mu)$-approximation for the (Connected) p-Center problem on~G can be computed in polynomial time. Our interest in these parameters stems from the fact that in many real-world networks, including Internet application networks, web networks, collaboration networks, social networks, biological networks, and others, and in many structured classes of graphs these parameters are small constants.

Efficient Best-Response Computation for Strategic Network Formation under Attack

Speaker: Pascal Lenzner, Hasso-Plattner-Institut, Universität Potsdam
Abstract: Strategic network formation models the uncoordinated creation of a network by selfish agents. Inspired by real world examples, e.g. the Internet, researchers have introduced an abundance of strategic games to study natural phenomena in networks. Most of these games have the conceptual drawback of being computationally intractable. For example, computing a best response strategy or checking whether an equilibrium is reached is NP-hard. Thus, a main challenge in the field is to find models which incorporate many interesting features and to devise efficient algorithms for solving the entailed computational tasks.

In my talk I'll address this challenge by providing an efficient algorithm for computing a Best-Response strategy for the recently introduced model by Goyal et al. [WINE'16]. Their promising model focuses on network robustness by considering an adversary who attacks (and kills) nodes in the network and lets this attack spread virus-like to neighboring nodes.

Completion Probabilities and Parallel Restart Strategies under an Imposed Deadline

Speaker: Jan-Hendrik Lorenz, Institut für Theoretische Informatik, Universität Ulm
Abstract: Randomised algorithms sometimes utilise a restart strategy. This means that a computation is aborted after a certain number of steps and restarted with a new, independent random seed. In many cases, this results in an improved expected runtime. However, we consider a model where a deadline D is present. In this context, the algorithm terminates without finding a solution if its runtime exceeds D. This is, for example, a common scenario in competitions or when renting time on a computer grid.

This scenario provides no guarantee for finding a solution, therefore the expected runtime is not necessarily the best measure for the properties of the algorithm under the imposed deadline D. In this case, the probability of finding a solution before D becomes a measure for the quality of the process. We address this issue and provide upper and lower bounds for this probability under varying assumptions.

It is known that the optimal restart times (regarding the expected runtime) for fixed cut-off algorithm running in parallel are not identical to the optimal restart times for the algorithm running on a single processor. In contrast, we show that the optimal restart times (regarding solution probability) for an algorithm running in parallel are identical to the optimal restart times for the algorithm running on a single processor. Finally, we conclude that the odds of finding a solution scale superlinearly in the number of processors.

Partially Ordered Nondeterministic Finite Automata – Expressivity and Complexity

Speaker: Tomáš Masopust, Institute of Theoretical Computer Science, TU Dresden
Abstract: A finite automaton is partially ordered if the only cycles in its transition diagram are self-loops. In this talk, we first discuss the expressive power of partially ordered NFAs and their relationship to the Straubing-Therien hierarchy. Then we discuss the complexity questions, such as the complexity of deciding universality, inclusion and equivalence. As a result, we show that two well-known classes of regular languages possess both nondeterministic and deterministic automata models, which is not a common situation in automata theory. Finally, we will discuss some applications of the results.

Online Resource Minimization

Speaker: Nicole Megow, Universität Bremen
Abstract: Minimizing resource usage is a key to achieving economic, environmental, or societal goals. We consider the fundamental problem of minimizing the number of machines that is necessary for feasibly scheduling preemptive jobs with release dates and hard deadlines. We consider the online variant of this problem in which every job becomes known to the online algorithm only at its release date. We present new competitive algorithms for this problem. We also discuss the power of job migration and show somewhat surprising strong lower bounds on the best possible performance guarantee.

Joint work with Lin Chen and Kevin Schewior.

TeamLTL

Speaker: Arne Meier, Institut für Theoretische Informatik, Leibniz Universität Hannover
Abstract: In this talk we transfer the notion of team semantics to the Linear Temporal Logic (LTL). We investigate basic properties of this logic and distinguish between a synchronised and an asynchronous version of the semantics. Further, we classify the computational complexity of satisfiability and model checking in this logic. Finally, we discuss the relation of TeamLTL to HyperLTL, the Linear Temporal Logic of Hyperproperties.

This talk presents joint work of Arne Meier and Martin Zimmermann.

Recognising Multidimensional Euclidean Preferences

Speaker: Dominik Peters, University of Oxford
Abstract: Euclidean preferences are a widely studied preference model, in which decision makers and alternatives are embedded in d-dimensional Euclidean space. Decision makers

prefer those alternatives closer to them. This model, also known as multidimensional unfolding, has applications in economics, psychometrics, marketing, and many other fields. We study the problem of deciding whether a given preference profile is d-Euclidean. For the one-dimensional case, polynomial-time algorithms are known. We show that, in contrast, for every other fixed dimension d > 1, the recognition problem is equivalent to the existential theory of the reals (ETR), and so in particular NP-hard. We further show that some Euclidean preference profiles require exponentially many bits in order to specify any Euclidean embedding, and prove that the domain of d-Euclidean preferences does not admit a finite forbidden minor characterisation for any d > 1. We also study dichotomous preferences and the behaviour of other metrics, and survey a variety of related work.

Published in AAAI 2017, https://arxiv.org/abs/1602.08109.

Improved Approximation for Two Dimensional Strip Packing with Polynomial Bounded Width

Speaker: Malin Rau, Institut für Informatik, Universität Kiel
Abstract: We study the well-known two-dimensional strip packing problem. Given is a set of rectangular axis-parallel items and a strip of width W with infinite height. The objective is to find a packing of these items into the strip, which minimizes the packing height. Lately, it has been shown that the lower bound of $3/2$ of the absolute approximation ratio can be beaten when we allow a pseudo-polynomial running-time of type $(n W)^{f(1/\varepsilon)}$, for any function f. If W is polynomially bounded by the number of items, this is a polynomial running-time.

We present a pseudo-polynomial algorithm with approximation ratio $4/3 + \varepsilon$ and running time $(n W)^{1 / \varepsilon^{\mathcal{O}(2^{1 / \varepsilon})}}$.

Continuous Speed Scaling with Variability

Speaker: Sören Riechers, Institut für Informatik, Universität Paderborn
Abstract: In my talk, I introduce an extension of the dynamic speed scaling scheduling model introduced by Yao et al. (1995): A set of jobs, each with a release time, deadline, and workload, has to be scheduled on a single, speed-scalable processor. Both the maximum allowed speed of the processor and the energy costs may vary continuously over time. The objective is to find a feasible schedule that minimizes the total energy costs.

Theoretical algorithm design for speed scaling problems often tends to discretize problems, as our tools in the discrete realm are often better developed or understood. Using the above speed scaling variant with variable, continuous maximal processor speeds and energy prices as an example, I demonstrate that a more direct approach via tools from variational calculus can

not only lead to a very concise and elegant formulation and analysis, but also avoids the “explosion of variables/constraints” that often comes with discretizing.

Nevertheless, I will start by introducing an optimal polynomial algorithm for piecewise constant speed limit. The optimality of the algorithm can be shown by using the well-known KKT conditions from combinatorial optimization. I will then discuss how to extend the algorithm to piecewise continuously differentiable speed limit functions and how to incorporate varying energy costs. Using well-known tools from calculus of variations, I derive combinatorial optimality characteristics for our continuous problem to prove the correctness of our algorithm in the continuous setting.

A Distributed Algorithm for Finding Hamiltonian Cycles in Random Graphs

Speaker: Volker Turau, Technische Universität Hamburg
Abstract: The threshold p_{crit} for the existence of a Hamiltonian cycle in a random graph $G(n,p)$ has been known precisely for some time. We present an asynchronous distributed algorithm that finds with high probability a Hamiltonian cycle in a graph $G(n,p)$ provided $p\ge (\epsilon\log n + \log \log n)/n$ with any $\epsilon>1$. The algorithm works in the asynchronous model and uses messages of size $O(\log n)$. For any $k \in \mathbb{N}$, $0< \epsilon \le 1/k$, and $p=(\log n/n)^{1/k-\epsilon}$ the algorithm terminates in $O(n + n^{2/k-\epsilon}\log^{1-1/k+\epsilon} n)$ rounds. In particular for $p \ge (\log n/n)^{1-\epsilon}$ and any $\epsilon < 1$ it terminates in subquadratic time and for $p \ge (\log n /n)^{1/2-\epsilon}$ and any $\epsilon < 1/2$ in linear time.