Abstracts Collection
25th International Symposium on Theoretical Aspects of Computer Science (STACS 2008)

Susanne Albers[1], Pascal Weil[2]

[1] Universität Freiburg, DE
salbers@informatik.uni-freiburg.de
[2] Université Bordeaux, FR
weil@labri.u-bordeaux.fr

Abstract. The Symposium on Theoretical Aspects of Computer Science (STACS) is held alternately in France and in Germany. The conference of February 21-23, 2008, held in Bordeaux, is the 25th in this series. Previous meetings took place in Paris (1984), Saarbrücken (1985), Orsay (1986), Passau (1987), Bordeaux (1988), Paderborn (1989), Rouen (1990), Hamburg (1991), Cachan (1992), Würzburg 1993), Caen (1994), München (1995), Grenoble (1996), Lübeck (1997), Paris (1998), Trier (1999), Lille (2000), Dresden (2001), Antibes (2002), Berlin (2003), Montpellier (2004), Stuttgart (2005), Marseille (2006) and Aachen (2007).

Keywords. Symposium, theoretical computer science, algorithms and data structures, automata and formal languages, computational and structural complexity, logic in computer science, applications

Preface – 25th International Symposium on Theoretical Aspects of Computer Science

The interest in STACS has remained at a high level over the past years. The STACS 2008 call for papers led to approximately 200 submissions from 38 countries. Each was assigned to at least three program committee members. The program committee held a 2-week long electronic meeting at the end of November, to select 54 papers. As co-chairs of this committee, we would like to sincerely thank its members and the many external referees for the valuable work they put into the reviewing process. The overall very high quality of the papers that were submitted to the conference made this selection a difficult task.

We would like to express our thanks to the three invited speakers, Maxime Crochemore, Thomas Schwentick and Mihalis Yannakakis, for their contributions to the proceedings.

Special thanks are due to A. Voronkov for his EasyChair software (www.easychair.org) which gives the organisers of conferences such as STACS a remarkable level of comfort; to Ralf Klasing for helping us explore the many possibilities of this brilliant software; to Emilka Bojanczyk for the design of the

STACS poster, proceedings and logo; and to the members of the Organizing Committee, chaired by David Janin.

An innovation in this year's STACS is the electronic format of the publication. A printed version was also available at the conference, with ISBN 978-3-939897-06-4. The electronic proceedings are available through several portals, and in particular through HAL and DROPS. HAL is an electronic repository managed by several French research agencies, and DROPS is the Dagstuhl Research Online Publication Server. We want to thank both these servers for hosting the proceedings of STACS and guaranteeing them perennial availability.

Keywords: Symposium, theoretical computer science, algorithms and data structures, automata and formal languages, computational and structural complexity, logic in computer science, applications

Joint work of: Albers, Susanne; Weil, Pascal

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2008/1326

Pushdown Compression

The pressing need for efficient compression schemes for XML documents has recently been focused on stack computation (Hariharan and Shankar 2006, League and Eng 2007), and in particular calls for a formulation of information-lossless stack or pushdown compressors that allows a formal analysis of their performance and a more ambitious use of the stack in XML compression, where so far it is mainly connected to parsing mechanisms. In this paper we introduce the model of pushdown compressor, based on pushdown transducers that compute a single injective function while keeping the widest generality regarding stack computation.

The celebrated Lempel-Ziv algorithm LZ78 was introduced as a general purpose compression algorithm that outperforms finite-state compressors on all sequences. We compare the performance of the Lempel-Ziv algorithm with that of the pushdown compressors, or compression algorithms that can be implemented with a pushdown transducer. This comparison is made without any a priori assumption on the data's source and considering the asymptotic compression ratio for infinite sequences. We prove that Lempel-Ziv is incomparable with pushdown compressors.

Keywords: Finite-state compression, Lempel-Ziv algorithm, pumping-lemma, pushdown compression, XML document

Joint work of: Albert, Pilar; Mayordomo, Elvira; Moser, Philip; Perifel, Sylvain

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1332

Quantum search with variable times

Since Grover's seminal work, quantum search has been studied in great detail. In the usual search problem, we have a collection of n items $x_1, \ldots, x_n$ and we would like to find $i : x_i = 1$.

We consider a new variant of this problem in which evaluating x_i for different i may take a different number of time steps.

Let t_i be the number of time steps required to evaluate x_i.

If the numbers t_i are known in advance, we give an algorithm that solves the problem in $O(\sqrt{t_1^2 + t_2^2 + \ldots + t_n^2})$ steps. This is optimal, as we also show a matching lower bound.

The case, when t_i are not known in advance, can be solved with a polylogarithmic overhead. We also give an application of our new search algorithm to computing read-once functions.

Joint work of: Ambainis, Andris

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1333

Structural aspects of tilings

In this paper, we study the structure of the set of tilings produced by any given tile-set. For better understanding this structure, we address the set of finite patterns that each tiling contains.

This set of patterns can be analyzed in two different contexts: the first one is combinatorial and the other topological. These two approaches have independent merits and, once combined, provide somehow surprising results.

The particular case where the set of produced tilings is countable is deeply investigated while we prove that the uncountable case may have a completely different structure.

We introduce a pattern preorder and also make use of Cantor-Bendixson rank. Our first main result is that a tile-set that produces only periodic tilings produces only a finite number of them. Our second main result exhibits a tiling with exactly one vector of periodicity in the countable case.

Keywords: Tiling, domino, patterns, tiling preorder, tiling structure

Joint work of: Ballier, Alexis; Durand, Bruno; Jeandal, Emmanuel

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1334

Limit complexities revisited

The main goal of this paper is to put some known results in a common perspective and to simplify their proofs.

We start with a simple proof of a result from (Vereshchagin, 2002) saying that $\limsup_{n} \mathrm{KS}(x|n)$ (here $\mathrm{KS}(x|n)$ is conditional (plain) Kolmogorov complexity of x when n is known) equals $\mathrm{KS}^{\mathbf{0}'(\mathbf{x})}$, the plain Kolmogorov complexity with $\mathbf{0}'$-oracle.

Then we use the same argument to prove similar results for prefix complexity (and also improve results of (Muchnik, 1987) about limit frequencies), a priori probability on binary tree and measure of effectively open sets. As a by-product, we get a criterion of $\mathbf{0}'$ Martin-Löf randomness (called also 2-randomness) proved in (Miller, 2004): a sequence ω is 2-random if and only if there exists c such that any prefix x of ω is a prefix of some string y such that $\mathrm{KS}(y) \geq |y| - c$. (In the 1960ies this property was suggested in (Kolmogorov, 1968) as one of possible randomness definitions; its equivalence to 2-randomness was shown in (Miller, 2004) while proving another 2-randomness criterion (see also (Nies et al. 2005)): ω is 2-random if and only if $\mathrm{KS}(x) \geq |x| - c$ for some c and infinitely many prefixes x of ω.

Finally, we show that the low-basis theorem can be used to get alternative proofs for these results and to improve the result about effectively open sets; this stronger version implies the 2-randomness criterion mentioned in the previous sentence.

Keywords: Kolmogorov complexity, limit complexities, limit frequencies, 2-randomness, low basis

Joint work of: Bienvenu, Laurent; Muchnik, Andrej; Shen, Alexander; Veraschagin, Nikolay

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1335

Trimmed Moebius Inversion and Graphs of Bounded Degree

We study ways to expedite Yates's algorithm for computing the zeta and Moebius transforms of a function defined on the subset lattice.

We develop a trimmed variant of Moebius inversion that proceeds point by point, finishing the calculation at a subset before considering its supersets. For an n-element universe U and a family $\mathcal{F}$ of its subsets, trimmed Moebius inversion allows us to compute the number of packings, coverings, and partitions of U with k sets from $\mathcal{F}$ in time within a polynomial factor (in n) of the number of supersets of the members of $\mathcal{F}$.

Relying on an intersection theorem of Chung et al. (1986) to bound the sizes of set families, we apply these ideas to well-studied combinatorial optimisation problems on graphs of maximum degree Δ. In particular, we show how to compute the Domatic Number in time within a polynomial factor of $(2^{\Delta+1-2)^n/(\Delta+1)}$ and the Chromatic Number in time within a polynomial factor of $(2^{\Delta+1-\Delta-1)^n/(\Delta+1)}$. For any constant Δ, these bounds are $O\big((2-\epsilon)^n\big)$ for $\epsilon > 0$ independent of the number of vertices n.

Joint work of: Björklund, Andreas; Husfeldt, Thore; Kaski, Petteri; Koivisto, Mikko

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1336

On the Complexity of the Interlace Polynomial

We consider the two-variable interlace polynomial introduced by Arratia, Bollobàs and Sorkin (2004). We develop two graph transformations which allow us to derive point-to-point reductions for the interlace polynomial. Exploiting these reductions we obtain new results concerning the computational complexity of evaluating the interlace polynomial at a fixed point. Regarding exact evaluation, we prove that the interlace polynomial is #P-hard to evaluate at every point of the plane, except at one line, where it is trivially polynomial time computable, and four lines and two points, where the complexity mostly is still open. This solves a problem posed by Arratia, Bollobàs and Sorkin (2004). In particular, we observe that three specializations of the two-variable interlace polynomial, the vertex-nullity interlace polynomial, the vertex-rank interlace polynomial and the independent set polynomial, are almost everywhere #P-hard to evaluate, too. For the independent set polynomial, our reductions allow us to prove that it is even hard to approximate at every point except at -1 and 0.

Keywords: Computational complexity, approximation, interlace polynomial, independent set polynomial, graph transformation

Joint work of: Bläser, Markus; Hoffmann, Christian

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1337

Minimizing Flow Time in the Wireless Gathering Problem

We address the problem of efficient data gathering in a wireless network through multi-hop communication. We focus on the objective of minimizing the maximum flow time of a data packet. We prove that no polynomial time algorithm for this problem can have approximation ratio less than $\Omega(m^{1/3}$ when m packets have to be transmitted, unless $P = NP$. We then use resource augmentation to assess the performance of a FIFO-like strategy. We prove that this strategy is 5-speed optimal, i.e., its cost remains within the optimal cost if we allow the algorithm to transmit data at a speed 5 times higher than that of the optimal solution we compare to.

Keywords: Wireless networks, data gathering, approximation algorithms, distributed algorithms

Joint work of: Bonifaci, Vincenzo; Korteweg, Peter; Marchetti-Spaccamela, Alberto

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1338

On Termination for Faulty Channel Machines

A channel machine consists of a finite controller together with several fifo channels; the controller can read messages from the head of a channel and write messages to the tail of a channel. In this paper, we focus on channel machines with insertion errors, i.e., machines in whose channels messages can spontaneously appear.

Such devices have been previously introduced in the study of Metric Temporal Logic. We consider the termination problem: are all the computations of a given insertion channel machine finite? We show that this problem has non-elementary, yet primitive recursive complexity.

Keywords: Automated Verification, Computational Complexity

Joint work of: Bouyer, Patricia; Markey, Nicolas; Quaknine, Joël; Schnoebelen, Philippe; Worrell, James

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1339

Stackelberg Network Pricing Games

We study a multi-player one-round game termed Stackelberg Network Pricing Game, in which a leader can set prices for a subset of m priceable edges in a graph. The other edges have a fixed cost.

Based on the leader's decision one or more followers optimize a polynomial-time solvable combinatorial minimization problem and choose a minimum cost solution satisfying their requirements based on the fixed costs and the leader's prices. The leader receives as revenue the total amount of prices paid by the followers for priceable edges in their solutions, and the problem is to find revenue maximizing prices. Our model extends several known pricing problems, including single-minded and unit-demand pricing, as well as Stackelberg pricing for certain follower problems like shortest path or minimum spanning tree. Our first main result is a tight analysis of a single-price algorithm for the single follower game, which provides a $(1+\varepsilon)\log m$-approximation for any $\varepsilon > 0$. This can be extended to provide a $(1+\varepsilon)(\log k + \log m)$-approximation for the general problem and k followers. The latter result is essentially best possible, as the problem is shown to be hard to approximate within $\mathcal{O}(\log^{\varepsilon} k + \log^{\varepsilon} m)$. If followers have demands, the single-price algorithm provides a $(1+\varepsilon)m^2$-approximation, and the problem is hard to approximate within $\mathcal{O}(m^{\varepsilon})$ for some $\varepsilon > 0$. Our second main result is a polynomial time algorithm for revenue maximization in the special case of Stackelberg bipartite vertex cover, which is based on non-trivial max-flow and LP-duality techniques. Our results can be extended to provide constant-factor approximations for any constant number of followers.

Keywords: Stackelberg Games, Algorithmic Pricing, Approximation Algorithms, Inapproximability

Joint work of: Briest, Patrick; Hoefer, Martin; Krysta, Piotr

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1340

Sublinear Communication Protocols for Multi-Party Pointer Jumping and a Related Lower Bound

We study the one-way number-on-the-forehead (NOF) communication complexity of the k-layer pointer jumping problem with n vertices per layer. This classic problem, which has connections to many aspects of complexity theory, has seen a recent burst of research activity, seemingly preparing the ground for an $\Omega(n)$ lower bound, for constant k. Our first result is a surprising sublinear — i.e., $o(n)$ — upper bound for the problem that holds for $k \geq 3$, dashing hopes for such a lower bound.

A closer look at the protocol achieving the upper bound shows that all but one of the players involved are collapsing, i.e., their messages depend only on the composition of the layers ahead of them. We consider protocols for the pointer jumping problem where all players are collapsing. Our second result shows that a strong $n - O(\log n)$ lower bound does hold in this case. Our third result is another upper bound showing that nontrivial protocols for (a non-Boolean version of) pointer jumping are possible even when all players are collapsing.

Our lower bound result uses a novel proof technique, different from those of earlier lower bounds that had an information-theoretic flavor. We hope this is useful in further study of the problem.

Keywords: Communication complexity, pointer jumping, number on the forehead

Joint work of: Brody, Joshua; Chakrabarti, Amit

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1341

Finding Irrefutable Certificates for S_2^p via Arthur and Merlin

We show that $S_2^p \subseteq P^{prAM}$, where S_2^p is the symmetric alternation class and $prAM$ refers to the promise version of the Arthur-Merlin class AM. This is derived as a consequence of our main result that presents an FP^{prAM} algorithm for finding a small set of "collectively irrefutable certificates" of a given S_2-type matrix. The main result also yields some new consequences of the hypothesis that NP has polynomial size circuits. It is known that the above hypothesis implies a collapse of the polynomial time hierarchy (PH) to $S_2^p \subseteq ZPP^{NP}$ (Cai 2007, Köbler and Watanabe 1998).

Under the same hypothesis, we show that PH collapses to P^{prMA}. We also describe an FP^{prMA} algorithm for learning polynomial size circuits for SAT, assuming such circuits exist.

For the same problem, the previously best known result was a ZPP^{NP} algorithm (Bshouty et al. 1996).

Keywords: Symmetric alternation, promise-AM, Karp–Lipton theorem, learning circuits

Joint work of: Chakaravarthy, Venkatesan T.; Roy, Sambuddha

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1342

Quantifying Homology Classes

We develop a method for measuring homology classes. This involves three problems. First, we define the size of a homology class, using ideas from relative homology. Second, we define an optimal basis of a homology group to be the basis whose elements' size have the minimal sum. We provide a greedy algorithm to compute the optimal basis and measure classes in it. The algorithm runs in $O(\beta^4 n^3 \log^2 n)$ time, where n is the size of the simplicial complex and β is the Betti number of the homology group. Third, we discuss different ways of localizing homology classes and prove some hardness results.

Keywords: Computational Topology, Computational Geometry, Homology, Persistent Homology, Localization, Optimization

Joint work of: Chen, Chao; Freedman, Daniel

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1343

Understanding maximal repetitions in strings

The cornerstone of any algorithm computing all repetitions in a string of length n in $\mathcal{O}(n)$ time is the fact that the number of runs (or maximal repetitions) is $\mathcal{O}(n)$.

We give a simple proof of this result. As a consequence of our approach, the stronger result concerning the linearity of the sum of exponents of all runs follows easily.

Keywords: Combinatorics on words, repetitions in strings, runs, maximal repetitions, maximal periodicities, sum of exponents

Joint work of: Crochemore, Maxime; Ilie, Lucian

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1344

Connecting Polygonizations via Stretches and Twangs

We show that the space of polygonizations of a fixed planar point set S of n points is connected by $O(n^2)$ "moves" between simple polygons. Each move is composed of a sequence of atomic moves called "stretches" and "twangs". These atomic moves walk between weakly simple "polygonal wraps" of S. These moves show promise to serve as a basis for generating random polygons.

Keywords: Polygons, polygonization, random polygons, connected configuration space

Joint work of: Damian, Mirela; Flatland, Robin; O'Rourke, Joseph; Ramaswani, Suneeta

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1345

Deterministically Isolating a Perfect Matching in Bipartite Planar Graphs

We present a deterministic way of assigning small (log bit) weights to the edges of a bipartite planar graph so that the minimum weight perfect matching becomes unique. The isolation lemma as described in (Mulmuley et al. 1987) achieves the same for general graphs using a randomized weighting scheme, whereas we can do it deterministically when restricted to bipartite planar graphs. As a consequence, we reduce both decision and construction versions of the matching problem to testing whether a matrix is singular, under the promise that its determinant is 0 or 1, thus obtaining a highly parallel SPL algorithm for bipartite planar graphs. This improves the earlier known bounds of non-uniform SPL by (Allender et al. 1999) and NC^2 by (Miller and Naor 1995, Mahajan and Varadarajan 2000). It also rekindles the hope of obtaining a deterministic parallel algorithm for constructing a perfect matching in non-bipartite planar graphs, which has been open for a long time. Our techniques are elementary and simple.

Joint work of: Datta, Samir; Kulkarni, Raghav; Roy, Sambuddha

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1346

Tight Bounds for Blind Search on the Integers

We analyze a simple random process in which a token is moved in the interval $A = \{0, \ldots, n\}$: Fix a probability distribution μ over $\{1, \ldots, n\}$. Initially, the token is placed in a random position in A. In round t, a random value d is chosen according to μ. If the token is in position $a \geq d$, then it is moved to position $a - d$. Otherwise it stays put. Let T be the number of rounds until the token reaches position 0. We show tight bounds for the expectation of T for the optimal distribution μ. More precisely, we show that $\min_\mu\{E_\mu(T)\} = \Theta\left((\log n)^2\right)$. For the proof, a novel potential function argument is introduced. The research is motivated by the problem of approximating the minimum of a continuous function over $[0, 1]$ with a "blind" optimization strategy.

Joint work of: Dietzfelbinger, Martin; Rowe, Jonathan E.; Wegener, Ingo; Woelfel, Philip

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1348

Discrete Jordan Curve Theorem: A proof formalized in Coq with hypermaps

This paper presents a formalized proof of a discrete form of the Jordan Curve Theorem. It is based on a hypermap model of planar subdivisions, formal specifications and proofs assisted by the Coq system. Fundamental properties are proven by structural or noetherian induction: Genus Theorem, Euler's Formula, constructive planarity criteria. A notion of ring of faces is inductively defined and a Jordan Curve Theorem is stated and proven for any planar hypermap.

Keywords: Formal specifications, Computational topology, Computer-aided proofs, Coq, Planar subdivisions, Hypermaps, Jordan Curve Theorem

Joint work of: Dufourd, Jean-Francois

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1349

Trimming of Graphs, with Application to Point Labeling

For $t, g > 0$, a vertex-weighted graph of total weight W is (t, g)-trimmable if it contains a vertex-induced subgraph of total weight at least $(1 - 1/t)W$ and with no simple path of more than g edges. A family of graphs is trimmable if for each constant $t > 0$, there is a constant $g = g(t)$ such that every vertex-weighted graph in the family is (t, g)-trimmable. We show that every family of graphs of bounded domino treewidth is trimmable. This implies that every family of graphs of bounded degree is trimmable if the graphs in the family have bounded treewidth or are planar. Based on this result, we derive a polynomial-time approximation scheme for the problem of labeling weighted points with nonoverlapping sliding labels of unit height and given lengths so as to maximize the total weight of the labeled points. This settles one of the last major open questions in the theory of map labeling.

Keywords: Trimming weighted graphs, domino treewidth, planar graphs, point-feature label placement, map labeling, polynomial-time approximation schemes

Joint work of: Erlebach, Thomas; Hagerup, Torben; Jansen, Klaus; Minzlaff, Moritz; Wolff, Alexander

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1350

Convergence Thresholds of Newton's Method for Monotone Polynomial Equations

Monotone systems of polynomial equations (MSPEs) are systems of fixed-point equations $X_1 = f_1(X_1, \ldots, X_n)$, $\ldots, X_n = f_n(X_1, \ldots, X_n)$ where each f_i is a polynomial with positive real coefficients. The question of computing the least non-negative solution of a given MSPE $\boldsymbol{X} = \boldsymbol{f}(\boldsymbol{X})$ arises naturally in the analysis of stochastic models such as stochastic context-free grammars, probabilistic pushdown automata, and back-button processes. Etessami and Yannakakis have recently adapted Newton's iterative method to MSPEs. In a previous paper we have proved the existence of a threshold $k_{\boldsymbol{f}}$ for strongly connected MSPEs, such that after $k_{\boldsymbol{f}}$ iterations of Newton's method each new iteration computes at least 1 new bit of the solution. However, the proof was purely existential. In this paper we give an upper bound for $k_{\boldsymbol{f}}$ as a function of the minimal component of the least fixed-point $\mu\boldsymbol{f}$ of $\boldsymbol{f}(\boldsymbol{X})$. Using this result we show that $k_{\boldsymbol{f}}$ is at most single exponential resp. linear for strongly connected MSPEs derived from probabilistic pushdown automata resp. from back-button processes. Further, we prove the existence of a threshold for arbitrary MSPEs after which each new iteration computes at least $1/w2^h$ new bits of the solution, where w and h are the width and height of the DAG of strongly connected components.

Keywords: Newton's Method, Fixed-Point Equations, Formal Verification of Software, Probabilistic Pushdown Systems

Joint work of: Esparza, Javier; Kiefer, Stefan; Luttenberger, Michael

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1351

Model Checking Games for the Quantitative μ-Calculus

We investigate quantitative extensions of modal logic and the modal μ-calculus, and study the question whether the tight connection between logic and games can be lifted from the qualitative logics to their quantitative counterparts. It turns out that, if the quantitative μ-calculus is defined in an appropriate way respecting the duality properties between the logical operators, then its model checking problem can indeed be characterised by a quantitative variant of parity games. However, these quantitative games have quite different properties than their classical counterparts, in particular they are, in general, not positionally determined. The correspondence between the logic and the games goes both ways: the value of a formula on a quantitative transition system coincides with the value of the associated quantitative game, and conversely, the values of quantitative parity games are definable in the quantitative μ-calculus.

Keywords: Games, logic, model checking, quantitative logics

Joint work of: Fischer, Diana; Grädel, Erich; Kaiser, Lukasz

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1352

Order-Invariant MSO is Stronger than Counting MSO in the Finite

We compare the expressiveness of two extensions of monadic second-order logic (MSO) over the class of finite structures. The first, counting monadic second-order logic (CMSO), extends MSO with first-order modulo-counting quantifiers, allowing the expression of queries like "the number of elements in the structure is even".

The second extension allows the use of an additional binary predicate, not contained in the signature of the queried structure, that must be interpreted as an arbitrary linear order on its universe, obtaining order-invariant MSO.

While it is straightforward that every CMSO formula can be translated into an equivalent order-invariant MSO formula, the converse had not yet been settled. Courcelle showed that for restricted classes of structures both order-invariant MSO and CMSO are equally expressive, but conjectured that, in general, order-invariant MSO is stronger than CMSO.

We affirm this conjecture by presenting a class of structures that is order-invariantly definable in MSO but not definable in CMSO.

Keywords: MSO, Counting MSO, order-invariance, expressiveness, Ehrenfeucht-Fraïssé game

Joint work of: Ganzow, Tobias; Rubin, Sasha

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1353

Succinctness of the Complement and Intersection of Regular Expressions

We study the succinctness of the complement and intersection of regular expressions. In particular, we show that when constructing a regular expression defining the complement of a given regular expression, a double exponential size increase cannot be avoided.

Similarly, when constructing a regular expression defining the intersection of a fixed and an arbitrary number of regular expressions, an exponential and double exponential size increase, respectively, can in worst-case not be avoided. All mentioned lower bounds improve the existing ones by one exponential and are tight in the sense that the target expression can be constructed in the corresponding time class, i.e., exponential or double exponential time. As a by-product, we generalize a theorem by Ehrenfeucht and Zeiger stating that there is a class of DFAs which are exponentially more succinct than regular expressions, to a fixed four-letter alphabet. When the given regular expressions are one-unambiguous, as for instance required by the XML Schema specification, the complement can be computed in polynomial time whereas the bounds concerning intersection continue to hold. For the subclass of single-occurrence regular expressions, we prove a tight exponential lower bound for intersection.

Joint work of: Gelade, Wouter; Neven, Frank

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1354

Efficient Algorithms for Membership in Boolean Hierarchies of Regular Languages

The purpose of this paper is to provide efficient algorithms that decide membership for classes of several Boolean hierarchies for which efficiency (or even decidability) were previously not known.

We develop new forbidden-chain characterizations for the single levels of these hierarchies and obtain the following results:

- The classes of the Boolean hierarchy over level Σ_1 of the dot-depth hierarchy are decidable in NL (previously only the decidability was known). The same remains true if predicates mod d for fixed d are allowed.
- If modular predicates for arbitrary d are allowed, then the classes of the Boolean hierarchy over level Σ_1 are decidable.
- For the restricted case of a two-letter alphabet, the classes of the Boolean hierarchy over level Σ_2 of the Straubing-Thérien hierarchy are decidable in NL. This is the first decidability result for this hierarchy.
- The membership problems for all mentioned Boolean-hierarchy classes are logspace many-one hard for NL.
- The membership problems for quasi-aperiodic languages and for d-quasi-aperiodic languages are logspace many-one complete for $PSPACE$.

Keywords: Automata and formal languages, computational complexity, dot-depth hierarchy, Boolean hierarchy, decidability, efficient algorithms

Joint work of: Glasser, Christian; Schmitz, Heinz; Selivanov, Victor

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1355

On the Complexity of Elementary Modal Logics

Modal logics are widely used in computer science. The complexity of modal satisfiability problems has been investigated since the 1970s, usually proving results on a case-by-case basis. We prove a very general classification for a wide class of relevant logics:

Many important subclasses of modal logics can be obtained by restricting the allowed models with first-order Horn formulas. We show that the satisfiability problem for each of these logics is either NP-complete or PSPACE-hard, and exhibit a simple classification criterion. Further, we prove matching PSPACE upper bounds for many of the PSPACE-hard logics.

Joint work of: Hemaspaandra, Edith; Schnoor, Henning

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1356

Fixed Parameter Polynomial Time Algorithms for Maximum Agreement and Compatible Supertrees

Consider a set of labels L and a set of trees $\mathcal{T} = \{\mathcal{T}^{(1)}, \mathcal{T}^{(2)}, \ldots, \mathcal{T}^{(k)}\}$ where each tree $\mathcal{T}^{(i)}$ is distinctly leaf-labeled by some subset of L. One fundamental problem is to find the biggest tree (denoted as supertree) to represent $\mathcal{T}$ which minimizes the disagreements with the trees in $\mathcal{T}$ under certain criteria. This problem finds applications in phylogenetics, database, and data mining. In this paper, we focus on two particular supertree problems, namely, the maximum agreement supertree problem (MASP) and the maximum compatible supertree problem (MCSP). These two problems are known to be NP-hard for $k \geq 3$. This paper gives the first polynomial time algorithms for both MASP and MCSP when both k and the maximum degree D of the trees are constant.

Keywords: Maximum agreement supertree, maximum compatible supertree

Joint work of: Hoang, Viet Tung; Sung, Wing-Kin

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1357

Computing Minimum Spanning Trees with Uncertainty

We consider the minimum spanning tree problem in a setting where information about the edge weights of the given graph is uncertain.

Initially, for each edge e of the graph only a set A_e, called an uncertainty area, that contains the actual edge weight w_e is known. The algorithm can 'update' e to obtain the edge weight $w_e \in A_e$. The task is to output the edge set of a minimum spanning tree after a minimum number of updates. An algorithm is k-update competitive if it makes at most k times as many updates as the optimum. We present a 2-update competitive algorithm if all areas A_e are open or trivial, which is the best possible among deterministic algorithms. The condition on the areas A_e is to exclude degenerate inputs for which no constant update competitive algorithm can exist.

Next, we consider a setting where the vertices of the graph correspond to points in Euclidean space and the weight of an edge is equal to the distance of its endpoints. The location of each point is initially given as an uncertainty area, and an update reveals the exact location of the point. We give a general relation between the edge uncertainty and the vertex uncertainty versions of a problem and use it to derive a 4-update competitive algorithm for the minimum spanning tree problem in the vertex uncertainty model. Again, we show that this is best possible among deterministic algorithms.

Keywords: Algorithms and data structures, Current challenges, mobile and net computing

Joint work of: Hoffmann, Michael; Erlebach, Thomas; Krizanc, Danny; Mihal'ák, Matús; Raman, Rajeev

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1358

Improved Algorithms for the Range Next Value Problem and Applications

The Range Next Value problem (Problem RNV) is a recent interesting variant of the range search problems, where the query is for the immediate next (or equal) value of a given number within a given interval of an array. Problem RNV was introduced and studied very recently by Crochemore et. al [Finding Patterns In Given Intervals, MFCS 2007]. In this paper, we present improved algorithms for Problem RNV. We also show how this problem can be used to achieve optimal query time for a number of interesting variants of the classic pattern matching problems.

Keywords: Algorithms, Data structures

Joint work of: Iliopoulos, Costas S.; Crochemore, Maxime; Kubica, Marcin; Rahman, M. Sohel; Walen, Tomasz

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1359

Cardinality and counting quantifiers on omega-automatic structures

We investigate structures that can be represented by omega-automata, so called omega-automatic structures, and prove that relations defined over such structures in first-order logic expanded by the first-order quantifiers 'there exist at most $\aleph_0$ many', 'there exist finitely many' and 'there exist k modulo m many' are omega-regular. The proof identifies certain algebraic properties of omega-semigroups.

As a consequence an omega-regular equivalence relation of countable index has an omega-regular set of representatives. This implies Blumensath's conjecture that a countable structure with an ω-automatic presentation can be represented using automata on finite words. This also complements a very recent result of Hjörth, Khoussainov, Montalban and Nies showing that there is an omega-automatic structure which has no injective presentation.

Keywords: ω-automatic presentations, ω-semigroups, ω-automata

Joint work of: Kaiser, Lukasz; Rubin, Sascha; Bárány, Vince

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1360

On the Induced Matching Problem

We study extremal questions on induced matchings in several natural graph classes. We argue that these questions should be asked for twinless graphs, that is graphs not containing two vertices with the same neighborhood. We show that planar twinless graphs always contain an induced matching of size at least $n/40$ while there are planar twinless graphs that do not contain an induced matching of size $(n+10)/27$. We derive similar results for outerplanar graphs and graphs of bounded genus. These extremal results can be applied to the area of parameterized computation. For example, we show that the induced matching problem on planar graphs has a kernel of size at most $40k$ that is computable in linear time; this significantly improves the results of Moser and Sikdar (2007). We also show that we can decide in time $O(91^k + n)$ whether a planar graph contains an induced matching of size at least k.

Keywords: Induced matching, bounded genus graphs, parameterized algorithms, kernel

Joint work of: Kanj, Iyad A.; Pelsmajer, Michael J., Schaefer, Marcus; Xia, Ge

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1361

The Frobenius Problem in a Free Monoid

The classical Frobenius problem over $\mathbb{N}$ is to compute the largest integer g not representable as a non-negative integer linear combination of non-negative integers $x_1, x_2, \ldots, x_k$, where $\gcd(x_1, x_2, \ldots, x_k) = 1$. In this paper we consider novel generalizations of the Frobenius problem to the noncommutative setting of a free monoid. Unlike the commutative case, where the bound on g is quadratic, we are able to show exponential or subexponential behavior for several analogues of g, with the precise bound depending on the particular measure chosen.

Keywords: Combinatorics on words, Frobenius problem, free monoid

Joint work of: Kao, Jui-Yi; Shallit, Jeffrey; Xu, Zhi

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1362

Space Hierarchy Results for Randomized Models

We prove space hierarchy and separation results for randomized and other semantic models of computation with advice. Previous works on hierarchy and separation theorems for such models focused on time as the resource. We obtain tighter results with space as the resource. Our main theorems are the following. Let $s(n)$ be any space-constructible function that is $\Omega(\log n)$ and such that $s(an) = O(s(n))$ for all constants a, and let $s'(n)$ be any function that is $\omega(s(n))$.

- There exists a language computable by two-sided error randomized machines using $s'(n)$ space and one bit of advice that is not computable by two-sided error randomized machines using $s(n)$ space and $\min(s(n), n)$ bits of advice.
- There exists a language computable by zero-sided error randomized machines in space $s'(n)$ with one bit of advice that is not computable by one-sided error randomized machines using $s(n)$ space and $\min(s(n), n)$ bits of advice.

The condition that $s(an) = O(s(n))$ is a technical condition satisfied by typical space bounds that are at most linear. We also obtain weaker results that apply to generic semantic models of computation.

Keywords: Computations with Advice, Space Hierarchy, Randomized Machine, Promise Classes, Semantic Models

Joint work of: Kinne, Jeff; van Melkebeek, Dieter

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1363

Ehrenfeucht-Fraïssé Goes Automatic for Real Addition

Various logical theories can be decided by automata-theoretic methods. Notable examples are Presburger arithmetic $\mathrm{FO}(Z, +, <)$ and the linear arithmetic over the reals $\mathrm{FO}(R, +, <)$, for which effective decision procedures can be built using automata. Despite the practical use of automata to decide logical theories, many research questions are still only partly answered in this area.

One of these questions is the complexity of such decision procedures and the related question about the minimal size of the automata of the languages that can be described by formulas in the respective logic. In this paper, we establish a double exponential upper bound on the automata size for $\mathrm{FO}(R, +, <)$ and an exponential upper bound for the discrete order over the integers $\mathrm{FO}(Z, <)$.

The proofs of these upper bounds are based on Ehrenfeucht-Fraïssé games. The application of this mathematical tool has a similar flavor as in computational complexity theory, where it can often be used to establish tight upper bounds of the decision problem for logical theories.

Keywords: Automata theory, automata-based decision procedures for logical theories, upper bounds, minimal sizes of automata, linear arithmetic over the reals, first-order equivalence, complexity

Joint work of: Klaedtke, Felix

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1364

New Combinatorial Complete One-Way Functions

In 2003, Leonid A. Levin presented the idea of a combinatorial complete one-way function and a sketch of the proof that Tiling represents such a function. In this paper, we present two new one-way functions based on semi-Thue string

rewriting systems and a version of the Post Correspondence Problem and prove their completeness. Besides, we present an alternative proof of Levin's result. We also discuss the properties a combinatorial problem should have in order to hold a complete one-way function.

Joint work of: Kojevnikov, Arist ; Nikolenko, Sergey I.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1365

Compatibility of Shelah and Stupp's and Muchnik's iteration with fragments of monadic second order logic

We investigate the relation between the theory of the iterations in the sense of Shelah-Stupp and of Muchnik, resp., and the theory of the base structure for several logics. These logics are obtained from the restriction of set quantification in monadic second order logic to certain subsets like, e.g., finite sets, chains, and finite unions of chains. We show that these theories of the Shelah-Stupp iteration can be reduced to corresponding theories of the base structure. This fails for Muchnik's iteration.

Keywords: Logic in computer science, Rabin's tree theorem

Joint work of: Kuske, Dietrich

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1366

Geometric Set Cover and Hitting Sets for Polytopes in R^3

Suppose we are given a finite set of points P in IR^3 and a collection of polytopes $\mathcal{T}$ that are all translates of the same polytope T. We consider two problems in this paper.

The first is the set cover problem where we want to select a minimal number of polytopes from the collection $\mathcal{T}$ such that their union covers all input points P. The second problem that we consider is finding a hitting set for the set of polytopes $\mathcal{T}$, that is, we want to select a minimal number of points from the input points P such that every given polytope is hit by at least one point.

We give the first constant-factor approximation algorithms for both problems. We achieve this by providing an epsilon-net for translates of a polytope in R^3 of size $O(\frac{1}{\epsilon})$.

Keywords: Computational Geometry, Epsilon-Nets, Set Cover, Hitting Sets

Joint work of: Laue, Sören

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1367

A Theory for Valiant's Matchcircuits (Extended Abstract)

The computational function of a matchgate is represented by its character matrix. In this article, we show that all nonsingular character matrices are closed under matrix inverse operation, so that for every k, the nonsingular character matrices of k-bit matchgates form a group, extending the recent work of Cai and Choudhary (2006) of the same result for the case of $k = 2$, and that the single and the two-bit matchgates are universal for matchcircuits, answering a question of Valiant (2002).

Keywords: Pfaffian, Matchgate, Matchcircuit

Joint work of: Li, Angsheng; Xia, Mingji

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1368

Rent, Lease or Buy: Randomized Algorithms for Multislope Ski Rental

In the Multislope Ski Rental problem, the user needs a certain resource for some unknown period of time. To use the resource, the user must subscribe to one of several options, each of which consists of a one-time setup cost ("buying price"), and cost proportional to the duration of the usage ("rental rate"). The larger the price, the smaller the rent. The actual usage time is determined by an adversary, and the goal of an algorithm is to minimize the cost by choosing the best option at any point in time.

Multislope Ski Rental is a natural generalization of the classical Ski Rental problem (where the only options are pure rent and pure buy), which is one of the fundamental problems of online computation. The Multislope Ski Rental problem is an abstraction of many problems where online decisions cannot be modeled by just two options, e.g., power management in systems which can be shut down in parts. In this paper we study randomized algorithms for Multislope Ski Rental. Our results include the best possible online randomized strategy for any additive instance, where the cost of switching from one option to another is the difference in their buying prices; and an algorithm that produces an e-competitive randomized strategy for any (non-additive) instance.

Keywords: Competitive analysis, ski rental, randomized algorithms

Joint work of: Lotker, Zvi; Patt-Shamir, Boaz; Rawitz, Dror

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1331

Lower bounds for adaptive linearity tests

Linearity tests are randomized algorithms which have oracle access to the truth table of some function f, and are supposed to distinguish between linear functions and functions which are far from linear. Linearity tests were first introduced by (Blum, Luby and Rubenfeld, 1993), and were later used in the PCP theorem, among other applications. The quality of a linearity test is described by its correctness c - the probability it accepts linear functions, its soundness s - the probability it accepts functions far from linear, and its query complexity q - the number of queries it makes.

Linearity tests were studied in order to decrease the soundness of linearity tests, while keeping the query complexity small (for one reason, to improve PCP constructions). Samorodnitsky and Trevisan (Samorodnitsky and Trevisan 2000) constructed the Complete Graph Test, and prove that no Hyper Graph Test can perform better than the Complete Graph Test. Later in (Samorodnitsky and Trevisan 2006) they prove, among other results, that no non-adaptive linearity test can perform better than the Complete Graph Test.

Their proof uses the algebraic machinery of the Gowers Norm. A result by (Ben-Sasson, Harsha and Raskhodnikova 2005) allows to generalize this lower bound also to adaptive linearity tests.

We also prove the same optimal lower bound for adaptive linearity test, but our proof technique is arguably simpler and more direct than the one used in (Samorodnitsky and Trevisan 2006). We also study, like (Samorodnitsky and Trevisan 2006), the behavior of linearity tests on quadratic functions. However, instead of analyzing the Gowers Norm of certain functions, we provide a more direct combinatorial proof, studying the behavior of linearity tests on random quadratic functions. This proof technique also lets us prove directly the lower bound also for adaptive linearity tests.

Keywords: Property testing, Linearity testing, Adaptive tests, Lower Property testing, Linearity testing, Adaptive tests, Lower

Joint work of: Lovett, Shachar

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1313

An Improved Randomized Truthful Mechanism for Scheduling Unrelated Machines

We study the scheduling problem on unrelated machines in the mechanism design setting. This problem was proposed and studied in the seminal paper (Nisan and Ronen 1999), where they gave a 1.75-approximation randomized truthful mechanism for the case of two machines. We improve this result by a 1.6737-approximation randomized truthful mechanism. We also generalize our result to a $0.8368m$-approximation mechanism for task scheduling with m machines, which improve the previous best upper bound of $0.875m$ (Mu'alem and Schapira 2007)

Keywords: Truthful mechanism, scheduling

Joint work of: Lu, Pinyan; Yu, Changyuan

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1314

Lagrangian Relaxation and Partial Cover (Extended Abstract)

Julián Mestre (Max-Planck-Institut für Informatik, D)

Lagrangian relaxation has been used extensively in the design of approximation algorithms. This paper studies its strengths and limitations when applied to Partial Cover.

We show that for Partial Cover in general no algorithm that uses Lagrangian relaxation and a Lagrangian Multiplier Preserving (LMP) α-approximation as a black box can yield an approximation factor better than $\frac{4}{3}\alpha$. This matches the upper bound given by Könemann et al. (ESA 2006, pages 468–479).

Faced with this limitation we study a specific, yet broad class of covering problems: Partial Totally Balanced Cover. By carefully analyzing the inner workings of the LMP algorithm we are able to give an almost tight characterization of the integrality gap of the standard linear relaxation of the problem. As a consequence we obtain improved approximations for the Partial version of Multicut and Path Hitting on Trees, Rectangle Stabbing, and Set Cover with ρ-Blocks.

Keywords: Lagrangian Relaxation, Partial Cover, Primal-Dual Algorithms

Joint work of: Mestre, Julián

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1315

On Dynamic Breadth-First Search in External-Memory

We provide the first non-trivial result on dynamic breadth-first search (BFS) in external-memory: For general sparse undirected graphs of initially n nodes and $O(n)$ edges and monotone update sequences of either $\Theta(n)$ edge insertions or $\Theta(n)$ edge deletions, we prove an amortized high-probability bound of $O(n/B^{2/3} + \mathrm{sort}(n) \cdot \log B)$ I/Os per update. In contrast, the currently best approach for static BFS on sparse undirected graphs requires $\Omega(n/B^{1/2} + \mathrm{sort}(n))$ I/Os.

Keywords: External Memory, Dynamic Graph Algorithms, BFS, Randomization

Joint work of: Meyer, Ulrich

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1316

Analytic aspects of the shuffle product

There exist very lucid explanations of the combinatorial origins of rational and algebraic functions, in particular with respect to regular and context free languages. In the search to understand how to extend these natural correspondences, we find that the shuffle product models many key aspects of D-finite generating functions, a class which contains algebraic. We consider several different takes on the shuffle product, shuffle closure, and shuffle grammars, and give explicit generating function consequences. In the process, we define a grammar class that models D-finite generating functions.

Keywords: Generating functions, formal languages, shuffle product

Joint work of: Mishna, Marni; Zabrocki, Mike

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1317

Weak index versus Borel rank

We investigate weak recognizability of deterministic languages of infinite trees. We prove that for deterministic languages the Borel hierarchy and the weak index hierarchy coincide.

Furthermore, we propose a procedure computing for a deterministic automaton an equivalent minimal index weak automaton with a quadratic number of states. The algorithm works within the time of solving the emptiness problem.

Keywords: Weak index, Borel rank, deterministic tree automata

Joint work of: Murlak, Filip

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1318

Complexity of solutions of equations over sets of natural numbers

Systems of equations over sets of natural numbers (or, equivalently, language equations over a one-letter alphabet) of the form $X_i = \varphi_i(X_1, \ldots, X_n)$ $(1 \leqslant i \leqslant n)$ are considered. Expressions φ_i may contain the operations of union, intersection and pairwise sum $A \oplus B = \{x + y \mid x \in A, y \in B\}$. A system with an EXPTIME-complete least solution is constructed, and it is established that least solutions of all such systems are in EXPTIME. The general membership problem for these equations is proved to be EXPTIME-complete.

Joint work of: Okhotin, Alexander; Jez, Artur

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1319

On Geometric Spanners of Euclidean and Unit Disk Graphs

We consider the problem of constructing bounded-degree planar geometric spanners of Euclidean and unit-disk graphs. It is well known that the Delaunay subgraph is a planar geometric spanner with stretch factor $C_{del} \approx 2.42$; however, its degree may not be bounded. Our first result is a very simple linear time algorithm for constructing a subgraph of the Delaunay graph with stretch factor $\rho = 1 + 2\pi(k \cos \frac{\pi}{k})^{-1}$ and degree bounded by k, for any integer parameter $k \geq 14$. This result immediately implies an algorithm for constructing a planar geometric spanner of a Euclidean graph with stretch factor $\rho \cdot C_{del}$ and degree bounded by k, for any integer parameter $k \geq 14$.

Moreover, the resulting spanner contains a Euclidean Minimum Spanning Tree (EMST) as a subgraph. Our second contribution lies in developing the structural results necessary to transfer our analysis and algorithm from Euclidean graphs to unit disk graphs, the usual model for wireless ad-hoc networks. We obtain a very simple distributed, *strictly-localized* algorithm that, given a unit disk graph embedded in the plane, constructs a geometric spanner with the above stretch factor and degree bound, and also containing an EMST as a subgraph. The obtained results dramatically improve the previous results in all aspects, as shown in the paper.

Keywords: Geometric spanner, euclidean graph, unit disk graph, wireless ad-hoc networks

Joint work of: Perkovic, Ljubomir; Kanj, Iyad A.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1320

A Mahler's theorem for functions from words to integers

In this paper, we prove an extension of Mahler's theorem, a celebrated result of p-adic analysis. Mahler's original result states that a function from N to Z is uniformly continuous for the p-adic metric d_p if and only if it can be uniformly approximated by polynomial functions. We prove the same result for functions from A^* to Z, where d_p is now the profinite metric defined by p-groups (pro-p metric).

Keywords: p-adic topology, binomial coefficients, Mahler's theorem, p-group languages

Joint work of: Pin, Jean-Éric; Silva. Pedro V.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1321

Distinguishing Short Quantum Computations

Distinguishing logarithmic depth quantum circuits on mixed states is shown to be complete for QIP, the class of problems having quantum interactive proof systems. Circuits in this model can represent arbitrary quantum processes, and thus this result has implications for the verification of implementations of quantum algorithms. The distinguishability problem is also complete for QIP on constant depth circuits containing the unbounded fan-out gate. These results are shown by reducing a QIP-complete problem to a logarithmic depth version of itself using a parallelization technique.

Keywords: Quantum information, computational complexity, quantum circuits, quantum interactive proof systems

Joint work of: Rosgen, Bill

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1322

Factoring Polynomials over Finite Fields using Balance Test

We study the problem of factoring univariate polynomials over finite fields. Under the assumption of the Extended Riemann Hypothesis (ERH), (Gao, 2001) designed a polynomial time algorithm that fails to factor only if the input polynomial satisfies a strong symmetry property, namely square balance. In this paper, we propose an extension of Gao's algorithm that fails only under an even stronger symmetry property. We also show that our property can be used to improve the time complexity of best deterministic algorithms on most input polynomials. The property also yields a new randomized polynomial time algorithm.

Keywords: Algebraic Algorithms, polynomial factorization, finite fields

Joint work of: Saha, Chandan

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1323

On the decomposition of k-valued rational relations

We give a new, and hopefully more easily understandable, structural proof of the decomposition of a k-valued transducer into k unambiguous functional ones, a result established by A. Weber in 1996. Our construction is based on a lexicographic ordering of computations of automata and on two coverings that can be build by means of this ordering. The complexity of the construction, measured as the number of states of the transducers involved in the decomposition, improves the original one by one exponential.

Moreover, this method allows further generalisation that solves the problem of decomposition of rational relations with bounded length-degree, which was left open in Weber's paper.

Keywords: Rational relation, k-valued transducer, unambiguous transducer, covering of automata

Joint work of: Sakarovitch, Jacques; de Souza, Rodrigo

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1324

A little bit infinite? On adding data to finitely labelled structures (Abstract)

Finite or infinite strings or trees with labels from a finite alphabet play an important role in computer science. They can be used to model many interesting objects including system runs in Automated Verification and XML documents in Database Theory. They allow the application of formal tools like logical formulas to specify properties and automata for their implementation. In this framework, many reasoning tasks that are undecidable for general computational models can be solved algorithmically, sometimes even eficiently.

Joint work of: Schwentick, Thomas

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1325

The Isomorphism Problem for Planar 3-Connected Graphs is in Unambiguous Logspace

The isomorphism problem for planar graphs is known to be efficiently solvable. For planar 3-connected graphs, the isomorphism problem can be solved by efficient parallel algorithms, it is in the class AC^1.

In this paper we improve the upper bound for planar 3-connected graphs to unambiguous logspace, in fact to $UL \cap coUL$. As a consequence of our method we get that the isomorphism problem for oriented graphs is in NL. We also show that the problems are hard for L.

Joint work of: Thierauf, Thomas; Wagner, Fabian

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1327

Efficient Minimization of DFAs with Partial Transition

Let PT-DFA mean a deterministic finite automaton whose transition relation is a partial function. We present an algorithm for minimizing a PT-DFA in $O(m \lg n)$ time and $O(m+n+\alpha)$ memory, where n is the number of states, m is the number of defined transitions, and α is the size of the alphabet.

Time consumption does not depend on α, because the α term arises from an array that is accessed at random and never initialized. It is not needed, if

transitions are in a suitable order in the input. The algorithm uses two instances of an array-based data structure for maintaining a refinable partition.

Its operations are all amortized constant time. One instance represents the classical blocks and the other a partition of transitions. Our measurements demonstrate the speed advantage of our algorithm on PT-DFAs over an $O(\alpha n \lg n)$ time, $O(\alpha n)$ memory algorithm.

Keywords: Deterministic finite automaton, sparse adjacency matrix, partition refinement

Joint work of: Valmari, Antti ; Lehtinen, Petri

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1328

Geodesic Fréchet Distance Inside a Simple Polygon

We unveil an alluring alternative to parametric search that applies to both the non-geodesic and geodesic Fréchet optimization problems. This randomized approach is based on a variant of red-blue intersections and is appealing due to its elegance and practical efficiency when compared to parametric search.

We present the first algorithm for the geodesic Fréchet distance between two polygonal curves A and B inside a simple bounding polygon P. The geodesic Fréchet decision problem is solved almost as fast as its non-geodesic sibling and requires $O(N^{2 \log k})$ time and $O(k+N)$ space after $O(k)$ preprocessing, where N is the larger of the complexities of A and B and k is the complexity of P. The geodesic Fréchet optimization problem is solved by a randomized approach in $O(k+N^{2 \log kN \log N})$ expected time and $O(k+N^2)$ space. This runtime is only a logarithmic factor larger than the standard non-geodesic Fréchet algorithm (Alt and Godau 1995). Results are also presented for the geodesic Fréchet distance in a polygonal domain with obstacles and the geodesic Hausdorff distance for sets of points or sets of line segments inside a simple polygon P.

Keywords: Fréchet Distance, Geodesic, Parametric Search, Simple Polygon

Joint work of: Wenk, Carola; Cook, Atlas F.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1330

Equilibria, Fixed Points, and Complexity Classes

Many models from a variety of areas involve the computation of an equilibrium or fixed point of some kind. Examples include Nash equilibria in games; market equilibria; computing optimal strategies and the values of competitive games (stochastic and other games); stable configurations of neural networks; analysing basic stochastic models for evolution like branching processes and for language like stochastic context-free grammars; and models that incorporate the basic primitives of probability and recursion like recursive Markov chains. It is not

known whether these problems can be solved in polynomial time. There are certain common computational principles underlying different types of equilibria, which are captured by the complexity classes PLS, PPAD, and FIXP. Representative complete problems for these classes are respectively, pure Nash equilibria in games where they are guaranteed to exist, (mixed) Nash equilibria in 2-player normal form games, and (mixed) Nash equilibria in normal form games with 3 (or more) players. This paper reviews the underlying computational principles and the corresponding classes.

Keywords: Equilibria, Fixed points, Computational Complexity, Game Theory

Joint work of: Yannakakis, Mihalis

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1311

Weighted Matching in the Semi-Streaming Model

We reduce the best known approximation ratio for finding a weighted matching of a graph using a one-pass semi-streaming algorithm from 5.828 to 5.585. The semi-streaming model forbids random access to the input and restricts the memory to $\mathcal{O}(n \cdot \text{polylog}\, n)$ bits. It was introduced by Muthukrishnan in 2003 and is appropriate when dealing with massive graphs.

Keywords: Semi-streaming algorithm, matching, approximation algorithm, graph algorithm

Joint work of: Zelke, Mariano

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1312

Shortest Vertex-Disjoint Two-Face Paths in Planar Graphs

Let G be a directed planar graph of complexity n, each arc having a nonnegative length. Let s and t be two distinct faces of G; let $s_1, \ldots, s_k$ be vertices incident with s; let $t_1, \ldots, t_k$ be vertices incident with t. We give an algorithm to compute k pairwise vertex-disjoint paths connecting the pairs (s_i, t_i) in G, with minimal total length, in $O(kn \log n)$ time.

Keywords: Algorithm, planar graph, disjoint paths, shortest path

Joint work of: de Verdiére, Éric; Schrijver

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1347

Design by Measure and Conquer, A Faster Exact Algorithm for Dominating Set

The measure and conquer approach has proven to be a powerful tool to analyse exact algorithms for combinatorial problems, like Dominating Set and Independent Set. In this paper, we propose to use measure and conquer also as a tool in the design of algorithms.

In an iterative process, we can obtain a series of branch and reduce algorithms. A mathematical analysis of an algorithm in the series with measure and conquer results in a quasiconvex programming problem. The solution by computer to this problem not only gives a bound on the running time, but also can give a new reduction rule, thus giving a new, possibly faster algorithm. This makes design by measure and conquer a form of computer aided algorithm design.

When we apply the methodology to a Set Cover modelling of the Dominating Set problem, we obtain the currently fastest known exact algorithms for Dominating Set: an algorithm that uses $O(1.5134^n)$ time and polynomial space, and an algorithm that uses $O(1.5063^n)$ time.

Keywords: Exact algorithms, exponential time algorithms, branch and reduce, measure and conquer, dominating set, computer aided algorithm design

Joint work of: van Rooij, Johan M. M.; Bodlaender, Hans L.

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2008/1329