

09171 Abstracts Collection
Adaptive, Output Sensitive, Online and
Parameterized Algorithms
— Dagstuhl Seminar —

Jeremy Barbay¹, Rolf Klein², Alejandro Lopez-Ortiz³ and Rolf Niedermeier⁴

¹ University of Santiago CL

jeremy@dcc.uchile.cl

² Universität Bonn, D

rolf.klein@uni-bonn.de

³ University of Waterloo, CA

⁴ Universität Jena, D

rolf.niedermeier@uni-jena.de

Abstract. From 19.01. to 24.04.2009, the Dagstuhl Seminar 09171 “Adaptive, Output Sensitive, Online and Parameterized Algorithms ” was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Adaptive analysis, instance optimal algorithms, fixed parameter tractable, output sensitive algorithms

**09171 Executive Summary – Adaptive, Output Sensitive,
Online and Parameterized Algorithms**

Traditionally the analysis of algorithms measures the complexity of a problem or algorithm in terms of the worst-case behavior over all inputs of a given size. However, in certain cases an improved algorithm can be obtained by considering a finer partition of the input space. As this idea has been independently rediscovered in many areas, the workshop gathered participants from different fields in order to explore the impact and the limits of this technique, in the hope to spring new collaboration and to seed the unification of the technique.

Keywords: Adaptive analysis, instance optimal algorithms, fixed parameter tractable, output sensitive algorithms

Joint work of: Barbay, Jérémy; Klein, Rolf; López-Ortiz, Alejandro; Niedermeier, Rolf

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2009/2120>

Parameterized/Adaptive Analysis of Online Steiner Tree Problems

Spyros Angelopoulos (MPI für Informatik - Saarbrücken, DE)

Steiner tree problems occupy a central place in both areas of approximation and on-line algorithms. Many variants have been studied from the point of view of competitive analysis, and for several of these variants tight bounds are known. However, in several cases, worst-case analysis is overly pessimistic, which fails to explain the relative performance of algorithms. We show how adaptive analysis can help resolve this problem. As case studies, we consider the Steiner tree problem in directed graphs, and the Priority Steiner tree problem.

Keywords: Online algorithms, Steiner tree problems, adaptive and parameterized analysis

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2009/2121>

The relation between Adaptive Algorithms and Compressed Succinct Data Structures

Jeremy Barbay (DCC - Universidad de Chile, CL)

Restricted classes of permutations can be sorted in time $o(n \lg n)$, faster than $O(n \lg n)$, through adaptive sorting algorithms. We show that each of those results imply a compression scheme for permutation, and that some of those result even in a compressed succinct data structure, which can be manipulated without uncompressing the whole permutation.

Those results are important in the case of self indices of text, which are strongly based on permutations, but application of the relation between fast algorithms and small data structure is not limited to permutations: it yields also compression of binary relations, integers, etcetera.

Keywords: Adaptive algorithms succinct data structure

Alternative Measures for Computational Complexity with applications to Machine Learning

Shai Ben-David (University of Waterloo, CA)

We address the issue of problems that are hard from a worst-case computational complexity view, but turn to be efficiently solvable in practice.

Having in mind a variety of algorithmic problems in machine learning in that category, we propose an alternative measure of computational complexity that takes into account not just the input size but also its "robustness" for the given task. The new measure coincides with the common worst-case complexity for "nicely behaving" data, but becomes less demanding for very non-robust (or, if you wish, non-informative, or irrelevant) data.

We argue that most naturally arising inputs are "nicely behaving" thus providing a possible explanation to the theory-practice complexity gap in the machine learning domain.

Using Property Testing for Efficient Sort of Relations

Sagi Ben-Moshe (Technion - Haifa, IL)

Sometimes, relational operations need to be performed on nearly-sorted relations, in which most tuples appear close to their place in the order (but not always exactly in place), while a small number of tuples may be completely out of their place in the order.

Such relations can result from updates performed on relations that were originally sorted, or as interim results when evaluating a complex query.

Currently, in such cases nearly-sorted relations are treated the same as unsorted relations, and when relational operations are evaluated for them, a generic algorithm is used. However, many operations can be computed more efficiently by an algorithm that exploits this near-ordering, yet to benefit from using such an algorithm, we also need to refrain from using the wrong algorithm for a relation which is not nearly-sorted. For this, an efficient test is required, i.e., a very fast approximation algorithm for establishing whether a given relation is sufficiently nearly-sorted.

In this paper, we provide the theoretical foundations for the problem of query evaluation over nearly-sorted relations. We formally define what it means for a relation to be nearly-sorted and show how operations over such relations can be executed efficiently. We provide efficient probabilistic tests for some cases of nearly-sorted relations (including the most general one) and illustrate how to incorporate them into a database management system. Proving the efficiency of these tests constitutes a substantial part of this work.

Finally, we provide experimental results, which show that our approach outperforms standard evaluation methods.

Parameterized Complexity of Kemeny Ranking

Nadja Betzler (Universität Jena, DE)

The computation of Kemeny rankings is central to many applications in the context of rank aggregation. Unfortunately, the problem is NP-hard. We consider various natural parameterizations resulting in efficient parameterized algorithms for relevant scenarios. In particular, we present a search tree for the parameter "Kemeny score" and give dynamic programming algorithms for the parameters "number of candidates", "maximum candidate range", and "average KT-distance between pairs of votes."

Joint work with Michael R. Fellows, Jiong Guo, Rolf Niedermeier, Frances A. Rosamond

Based on "Fixed-parameter algorithms for Kemeny rankings" (AAIM 2009) and "How similarity helps to efficiently compute Kemeny rankings" (AAMAS 2009)

Keywords: Voting systems, consensus, fixed-parameter tractability

Five analysis techniques for on-line algorithms

Joan Boyar (Univ. of Southern Denmark - Odense, DK)

Five techniques for analyzing the performance of on-line algorithms are compared on algorithms for classical and dual bin packing. The five techniques are competitive analysis, the accommodating function, resource augmentation, relative worst order analysis, and bijective analysis. For classical bin packing First-Fit is compared to Worst-Fit and Worst-Fit is compared to Next-Fit. For Dual bin packing, First-Fit is compared to Worst-Fit. Even with these simple algorithms, there are some open problems.

Alternative Measures for Analysis of Online Algorithms

Reza Dorrigiv (University of Waterloo, CA)

In this talk I show that there is a large gap between the theory and practice of paging as well as other important on-line algorithms. In particular, the competitive ratio which is the main technique for analysis of on-line algorithms is known to produce in unrealistic measures of performance for certain problems. In this talk I concentrate on the paging problem, providing the first theoretical proof of optimality of LRU using a new measure derived from first principles which better corresponds to observed practice. This can be extended to other settings. I also provide other alternative measures for analysis of online algorithms and show that they lead to promising results.

Better Bounds on Online Unit Clustering

Martin R. Ehmsen (Univ. of Southern Denmark - Odense, DK)

Unit Clustering is the problem of dividing a set of points from a metric space into a minimal number of subsets such that the points in each subset are enclosable by a unit ball. We continue the work, recently initiated by Chan and Zarrabi-Zadeh, on determining the competitive ratio of the online version of this problem. For the one-dimensional case, we develop a deterministic algorithm, improving the best known upper bound of $7/4$ by Epstein and van Stee to $5/3$. This narrows the gap to the best known lower bound of $8/5$ to only $1/15$. Our algorithm automatically leads to improvements in all higher dimensions as well. Finally, we strengthen the deterministic lower bound in two dimensions and higher from 2 to $13/6$.

Preemptive online scheduling with reordering

Leah Epstein (University of Haifa, IL)

We consider online preemptive scheduling of jobs, arriving one by one, on m parallel machines. A buffer of fixed size $K > 0$, which assists in partial reordering of the input, is available, to be used for the storage of at most K unscheduled jobs. We consider several variants of the problem. For general inputs and identical machines, we show that a buffer of size $\Theta(m)$ reduces the overall competitive ratio from $e/(e-1)$ to $4/3$. Surprisingly, the competitive ratio as a function of m is not monotone, unlike the case where $K = 0$.

Joint work of: Dósa, György; Epstein, Leah

Comparing First-Fit and Next-Fit for Online Edge Coloring

Lene Monrad Favrholdt (Univ. of Southern Denmark - Odense, DK)

We study the performance of the algorithms First-Fit and Next-Fit for two online edge coloring problems. In the min-coloring problem, all edges must be colored using as few edges as possible. In the max-coloring problem, a fixed number of colors is given, and as many edges as possible should be colored. Previous analysis using the competitive ratio has not separated the performance of First-Fit and Next-Fit, but intuition suggests that First-Fit should be better than Next-Fit. We compare First-Fit and Next-Fit using the relative worst order ratio, and show that First-Fit is better than Next-Fit for the min-coloring problem. For the max-coloring problem, we show that First-Fit and Next-Fit are not strictly comparable, i.e., there are graphs for which First-Fit is better than Next-Fit and graphs where Next-Fit is slightly better than First-Fit.

Keywords: Online edge coloring, competitive ratio, relative worst-order ratio, First-Fit, Next-Fit

Experimental Study of FPT Algorithms for the

Rudolf Fleischer (Fudan University - Shanghai, CN)

Finding minimum feedback vertex sets in directed graphs (the DFVS problem) is crucial in both theoretical and system research areas, for example, to find resource deadlocks. Unfortunately, the directed FVS problem is quite different from its undirected counterpart (the UFVS problem) and it is not very well understood. Since FVS problems are NP-hard, their theoretical study has recently focused on finding efficient fixed parameter tractable (FPT) algorithms. While UFVS has a quadratic kernel [5], DFVS was only recently shown to be in FPT [1], and it still remains an open question whether it has a polynomial size kernel.

Inspired by the recent kernelization techniques for UFVS by Thomasse [5], we propose five data reduction rules for DVFS, Edge Canonicalization, Self-Loop, Dummy Nodes, Chaining Nodes, and Flower, which can significantly reduce the size of the input graph and the parameter k . This leads to an algorithm for DFVS that runs in time $O(4k \ln O(1))$, where n and k are the number of vertices and the size of minimum feedback vertex set, respectively. In a comprehensive experimental study we compared our new algorithm with the previous FPT algorithm by Chen et al. [1].

We implemented a dense graph generator to test the running time, kernel size, memory usage, and recursion depth of this algorithm on dense random graphs. In particular, by controlling the optimum solution size k , we evaluated how the running time relates to the parameter k . Furthermore, we implemented a resource allocation graph (RAG) generator simulating the deadlock characteristics of processes in real complex systems [4]. We used Markov chains to control the distribution of processes and resources, and tested whether our algorithm could quickly resolve the deadlock cases. Finally, we implemented the approximation algorithm by Even et al. [2] to further accelerate the search of small feedback vertex sets in graphs generated by the RAG generator, and compared the results to Chen's algorithm. All implementations were done using LEDA [3].

Keywords: FPT, DFVS, data reduction rules

Joint work of: Fleischer, Rudolf; Wu, Xi

Two Topics in Adaptive Algorithms: Hulls and Strings

Robert Fraser (University of Waterloo, CA)

In this talk we cover two problems that we have been studying with adaptive analysis components.

First, we examine the planar convex hull problem, which has natural extensions from well known adaptive sorting algorithms. Next, we cover string matching, where we are performing online string matching upon a pattern with known distributions of characters.

Keywords: Adaptive analysis, convex hull, string matching, adaptive sorting

Hyperbolic Dovetailing

David G. Kirkpatrick (University of British Columbia - Vancouver, CA)

A familiar quandary, in many settings (computational and otherwise), arises when there are several *possible* alternatives for the solution of a problem, but no way of knowing which, if any, are viable for a particular problem instance. Faced with this uncertainty, we are forced to simulate the parallel exploration of alternatives through some kind of co-ordinated interleaving (*dovetailing*) process. The goal, as usual, is to find a solution with low total cost. Much of the existing work on such problems has assumed, implicitly or explicitly, that at most one of the alternatives is viable. This assumption provides support for a competitive analysis of algorithms (using the cost of the unique viable alternative as a benchmark).

However, just as it is unrealistic to analyse algorithms in terms of the worst case cost of the alternative solutions or their worst-case ordering (giving rise to competitive analysis), it is also unrealistic in many scenarios to make the worst-case assumption that at most one of the alternatives is viable. In this paper, we relax this assumption in revisiting several familiar dovetailing problems.

Our main contribution is the introduction of a novel process interleaving technique, called *hyperbolic dovetailing* that achieves a competitive ratio that is within a logarithmic factor of optimal on *all* inputs in the worst, average and expected cases, over all possible deterministic (and randomized) dovetailing schemes. We also show that no other dovetailing strategy can guarantee an asymptotically smaller competitive ratio for all inputs.

An interesting application of hyperbolic dovetailing arises in the design of what we call *input-thrifty* algorithms, algorithms that are designed to minimize the total precision of the input requested in order to evaluate some given predicate. We show that for some very basic predicates involving real numbers (such as certifying that the numbers are not all identical) we can use hyperbolic dovetailing to provide input-thrifty algorithms that are competitive, in this novel cost measure, with the best algorithms that solve these problems.

Parameterized complexity of geometric problems

Christian Knauer (FU Berlin, DE)

Parameterized complexity aims to design exact algorithms whose running times depend on certain parameters of the input data that are naturally related to the problem at hand and in a way capture its complexity.

A problem is called fixed-parameter tractable (FPT) with respect to a parameter k if there is an efficient algorithm to solve the problem for the cases where the parameter k is small. Another objective of this theory is to show that such algorithms are unlikely to exist for certain problems (and parameters).

Not many geometric problems have been studied from the parameterized complexity point of view. Most research has focused on special (combinatorial) parameters for geometric problems, like, e.g., the number of inner points (i.e., points in the interior of the convex hull) for the TSP problem or for the problem of computing minimum convex decompositions. Also, on the negative side, only few connections between geometric problems and known hard parameterized problems are known to date.

We provide a brief tour of results from parameterized complexity theory for various geometric problems with respect to various parameters (hyperplane depth, clustering, polygon guarding).

Keywords: Parameterized complexity, Computational geometry

Derandomizing Non-uniform Color-Coding

Alexander Langer (RWTH Aachen, DE)

Color-coding, as introduced by Alon, Yuster, and Zwick, is a well-known tool for algorithm design and can often be efficiently derandomized using universal hash functions. In the special case of only two colors, one can use (n, k) -universal sets for the derandomization. Here, we introduce (n, k, l) -universal sets that are typically smaller and can be constructed faster. Nevertheless, for some problems they are still sufficient for derandomization and faster deterministic algorithms can be obtained. This particularly works well when the color-coding does not use a uniform probability distribution. To exemplify the concept, we present an algorithm for the UNIQUE COVERAGE problem introduced by Demaine, Feige, Hajiaghayi, and Salavatipour. The example also shows how to extend the concept to multiple colors.

Joint work with Joachim Kneis and Peter Rossmanith.

A Comparison of Performance Measures for Online Algorithms

Kim Skak Larsen (Univ. of Southern Denmark - Odense, DK)

A systematic study of several recently suggested measures for online algorithms is provided in the context of a specific problem, namely, the two server problem on three colinear points. Even though the problem is simple, it encapsulates a core challenge in online algorithms which is to balance greediness and adaptability. We examine how these measures evaluate the Greedy Algorithm and

Lazy Double Coverage, commonly studied algorithms in the context of server problems. We examine Competitive Analysis, the Max/Max Ratio, the Random Order Ratio, Bijective Analysis and Relative Worst Order Analysis and determine how they compare the two algorithms. We find that by the Max/Max Ratio and Bijective Analysis, Greedy is the better algorithm. Under the other measures Lazy Double Coverage is better, though Relative Worst Order Analysis indicates that Greedy is sometimes better. Our results also provide the first proof of optimality of an algorithm under Relative Worst Order Analysis.

Joint work of: Larsen, Kim Skak; Boyar, Joan; Iranim, Sandy

Meta-Theorems for Kernelization

Daniel Lokshтанov (University of Bergen, NO)

Polynomial time preprocessing to reduce instance size is one of the most commonly deployed heuristics to tackle computationally hard problems. In a parameterized problem, every instance I comes with a positive integer k . The problem is said to admit a polynomial kernel if, in polynomial time, we can reduce the size of the instance I to a polynomial in k , while preserving the answer. In this paper, we show that all problems expressible in Counting Monadic Second Order Logic and satisfying a compactness property admit a polynomial kernel on graphs of bounded genus. Our second result is that all problems that have finite integer index and satisfy a weaker compactness condition admit a linear kernel on graphs of bounded genus. The study of kernels on planar graphs was initiated by a seminal paper of Alber, Fellows, and Niedermeier [J. ACM, 2004] who showed that PLANAR DOMINATING SET admits a linear kernel. Following this result, a multitude of problems have been shown to admit linear kernels on planar graphs by combining the ideas of Alber et al. with problem specific reduction rules. Our theorems unify and extend all previously known kernelization results on planar graph problems. Combining our theorems with the Erdos-Posa property we obtain various new results on linear kernels for a number of packing and covering problems.

Keywords: Kernelization

Joint work of: Bodlaender, Hans; Fomin Fedor; Lokshтанov, Daniel; Eelko, Penninkx; Saurabh, Saket; Thilikos, Dimitrios

Improving local search using parameterized complexity

Daniel Marx (Budapest Univ. of Technology & Economics, HU)

Local search algorithms iteratively improve solutions by checking whether it is possible to find a better solution in the local neighborhood of the current solution.

The local neighborhood is usually defined as the set of solutions that can be obtained by one (or more generally, at most k for some fixed k) elementary changes. Large values of k can give better results; however, the brute force search of the local neighborhood is not feasible for larger k .

Parameterized complexity gives a convenient framework for studying the question whether there is an efficient way of searching the local neighborhood. In the talk, I will briefly overview parameterized complexity, summarize recent results in this direction, and explain in more detail the analysis of the problem of finding minimum weight solutions for Boolean CSP.

Joint work with Andrei Krokhin.

Keywords: Parameterized complexity, local search, CSP

Selforganizing Linear Search Revisited

J. Ian Munro (University of Waterloo, CA)

In this meeting we focus on a variety of aspects of online methods and their relation to adaptiveness and parametrized complexity. One important issue is appropriate methods of comparing algorithms, including the notion of competitiveness. An issue often lost is that of the fundamental model and charging mechanism that online and offline methods are to follow. We revisit the work of Martinez et al and of the author that demonstrate that no algorithm for self-organizing linear search is better than n/lgn competitive if our model permits us to scan down the list (at least to the desired element) at a charge of one per element inspected, and then reorder the inspected values at no charge.

Keywords: Competitive, selforganizing, online

Adaptive algorithms for planar convex hull problems

Yoshio Okamoto (Tokyo Institute of Technology, JP)

We study problems in computational geometry from the viewpoint of adaptive algorithms. Adaptive algorithms have been extensively studied for the sorting problem, and in this paper we generalize the framework to geometric problems. To this end, we think of geometric problems as permutation (or rearranging) problems of arrays, and define the “presortedness” as a distance from the input array to the desired output array. We call an algorithm *adaptive* if it runs faster when a given input array is closer to the desired output, and furthermore it does not make use of any information of the presortedness. As a case study, we look into the planar convex hull problem for which we discover two natural formulations as permutation problems.

An interesting phenomenon that we prove is that for one formulation the problem can be solved adaptively, but for the other formulation no adaptive algorithm can be better than an optimal output-sensitive algorithm for the planar convex hull problem.

Joint work of: Hee-Kap Ahn; Okamoto, Yoshio

Parameterized Approximation Schemes for Batch Coloring

Hadas Shachnai (Technion - Haifa, IL)

Batch scheduling of conflicting jobs is modeled by *batch coloring* of a graph. Given an undirected graph and the number of colors required by each vertex, we need to find a proper batch coloring of the graph, namely, to partition the vertices to batches which are independent sets and assign to each batch a contiguous set of colors, whose size is equal to the maximum color requirement of any vertex in this batch.

When the objective is to minimize the sum of job completion times, we get the *batch sum coloring* problem; when we want to minimize the maximum completion time of any job (or, the *makespan*) we get the *max coloring* problem.

Given the hardness of batch coloring on general graphs, already for the special case of unit color requirements (known as *sum coloring* and the classic *graph coloring* problem, respectively), it is natural to seek out classes of graphs where effective solutions can be obtained efficiently.

We give the first (*parameterized*) *polynomial time approximation schemes* for batch sum coloring on several classes of “non-thick” graphs that arise in applications. This includes partial k -trees, planar graphs and other “flat” graphs.

Joint work with Magnus M. Halldorsson.

Keywords: Batch sum coloring, max coloring, parameterized approximation schemes