Report from Dagstuhl Seminar 20081

Scheduling

Edited by
Nicole Megow[1], David Shmoys[2], and Ola Svensson[3]

1 **Universität Bremen, DE,** nicole.megow@uni-bremen.de
2 **Cornell University – Ithaca, US,** david.shmoys@cornell.edu
3 **EPFL – Lausanne, CH,** ola.svensson@epfl.ch

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 20081 "Scheduling". The seminar focused on the interplay between scheduling problems and problems that arise in the management of transportation and traffic. Important aspects at the intersection of these two research directions include data-driven approaches in dynamic decision-making, scheduling in combination with routing, shared mobility, and coordination versus competition.

Seminar February 16–21, 2020 – http://www.dagstuhl.de/20081
2012 ACM Subject Classification Theory of computation → Scheduling algorithms,Applied computing → Transportation, Mathematics of computing → combinatorial optimization, Theory of computation → Approximation algorithms analysis
Keywords and phrases scheduling, optimization, approximation algorithms, routing, transportation, mechanism design
Digital Object Identifier 10.4230/DagRep.10.2.50
Edited in cooperation with Franziska Eberle

1 Executive Summary

Nicole Megow (Universität Bremen, DE)
David Shmoys (Cornell University – Ithaca, US)
Ola Svensson (EPFL – Lausanne, CH)

This seminar was the sixth in a series of Dagstuhl "Scheduling" seminars (since 2008). Scheduling is a major research field that is studied from a practical and theoretical perspective in computer science, mathematical optimization, and operations research. Applications range from traditional production scheduling and project planning to the newly arising resource management tasks in the advent of internet technology and shared resources.

This edition of the seminar focused on the interplay between scheduling problems and problems that arise in the management of traffic. There are several notable aspects of the scheduling problems that arise particularly in this context:

- the role of dynamic decision-making in which data-driven approaches emerge (especially those that have stochastic elements in modelling multi-stage decision-making);
- the interplay between scheduling aspects and what might be viewed as routing aspects, providing a spacial component to the nature of the scheduling problem;
- the tension between questions of coordination and competition that arise from the fact that, for many of the issues in this domain, there are significant questions that depend on the extent to which the traffic can be centrally coordinated.

Since the community working on the intersection of scheduling and traffic is itself rather broad, the seminar focused on researchers whose methodological focus relies on tools from the theoretical design of algorithms, on mathematical optimization methods, and on the combination of optimization and game-theoretic approaches.

Organization of the Seminar. The workshop brought together 59 researchers from theoretical computer science, mathematical optimization and operations research. The participants consisted of both senior and junior researchers, including a number of postdocs and advanced PhD students.

During the five days of the workshop, 31 talks of different lengths took place. Four keynote speakers gave an overview of the state-of-the art of the respective area in 60 minutes:

- Shuchi Chawla: Mechanisms for resource allocation
- Benjamin Moseley: Combinatorial Optimization Augmented with Machine Learning
- Evá Tardos: Learning in Games and in Queueing Systems
- Vera Traub: Approximation algorithms for traveling salesman problems.

The remaining slots were filled with shorter talks of 30 minutes on various topics related to scheduling, routing, transportation, mechanism design, learning, and applications in practice. Another highlight of the workshop was a historical note given by Jan Karel Lenstra with his view on the dynamic development of the area of scheduling in the past 60 years. Further, in the beginning of the week, open problem sessions were held. Throughout the week, a few sessions with spotlight talks of 8 minutes gave participants the chance to announce recent results and invite for discussions. The schedule left ample free time that was actively used for fruitful discussions and joint research.

Outcome. Organizers and participants regard the workshop as a great success. The workshop achieved the goal to bring together the related communities, share the state-of-the art research and discuss the current major challenges. The talks were excellent and very stimulating; participants actively met in working groups in the afternoon and evenings. It was remarked very positively that a significant number of younger researchers (postdocs and PhD students) participated and integrated very well.

The organizers wish to express their gratitude towards the Scientific Directorate and the administration of the Dagstuhl Center for their great support for this workshop.

2 Table of Contents

Executive Summary
Nicole Megow, David Shmoys, and Ola Svensson . . . 50

Overview of Talks

Minimizing Energy Consumption on Multiple Machines with Sleep States
Antonios Antoniadis . . . 54

General Framework for Metric Optimization Problems with Delay or with Deadlines
Yossi Azar . . . 54

Learning Augmented Energy Minimization via Speed Scaling
Etienne Bamas . . . 55

Mechanisms for resource allocation
Shuchi Chawla . . . 55

The online k-taxi problem
Christian Coester . . . 56

A Tale of Santa Claus, Hypergraphs and Matroids
Sami Davies . . . 56

Optimally Handling Commitment Issues in Online Throughput Maximization
Franziska Eberle . . . 57

Pricing in Resource Allocation Games Based on Duality Gaps
Tobias Harks . . . 57

Scheduling stochastic jobs with release dates on a single machine
Sven Jäger . . . 58

Online Learning with Vector Costs and Bandits with Knapsacks
Thomas Kesselheim . . . 58

Equilibria in Atomic Splittable Congestion Games
Max Klimm . . . 59

Non-Clairvoyant Precedence Constrained Scheduling
Amit Kumar . . . 59

Online Vehicle Routing, the edge of optimization in large scale applications
Sebastien Martin . . . 60

Malleable Scheduling Beyond Identical Machines
Jannik Matuschke . . . 60

Combinatorial Optimization Augmented with Machine Learning
Benjamin J. Moseley . . . 61

Group Fairness in Network Design and Combinatorial Optimization
Kamesh Munagala . . . 61

Scheduling Bidirectional Traffic
Rolf H. Möhring . . . 62

Stochastic Makespan Minimization
Viswanath Nagarajan . . . 62

Online Algorithms via Projections
Seffi Naor 63

Improved Approximation Algorithms for Inventory Problems
Neil Olver 63

Sample-Based Prophet Inequalities
Kevin Schewior 64

Online Metric Algorithms with Untrusted Predictions
Bertrand Simon 64

Fixed-Order Scheduling on Parallel Machines
René Sitters 65

Approximation Algorithms for Replenishment Problems with Fixed Turnover Times
Leen Stougie 65

Learning in Games and in Queueing Systems
Éva Tardos 66

Approximation algorithms for traveling salesman problems
Vera Traub 66

Greed...Is Good For Scheduling Under Uncertainty (2nd ed.)
Marc Uetz 67

Nash flows over Time with Spillback and Kinematic Waves
Laura Vargas Koch 67

A Water-Filling Primal-Dual Algorithm for Approximating Non-Linear Covering Problems
Jose Verschae 68

Dynamic Approximate Maximum Independent Set of Intervals, Hypercubes and Hyperrectangles
Andreas Wiese 68

Open Problems

Deterministic min-cost matching with delays
Yossi Azar 69

On-line Routing
Sanjoy Baruah 70

Integrality Gap of the natural LP for $P2|\text{prec}|C_{\max}$
Xinrui Jia 70

The Busy Time Problem
Samir Khuller 71

Location Routing with Depot Capacities
Jannik Matuschke 71

Scheduling on two types of machines
Bertrand Simon 72

Participants 74

3 Overview of Talks

3.1 Minimizing Energy Consumption on Multiple Machines with Sleep States

Antonios Antoniadis (MPI für Informatik – Saarbrücken, DE)

Joint work of Antonios Antoniadis, Naveen Garg, Gunjan Kumar, Nikhil Kumar
Main reference Antonios Antoniadis, Naveen Garg, Gunjan Kumar, Nikhil Kumar: "Parallel Machine Scheduling to Minimize Energy Consumption", in Proc. of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020, pp. 2758–2769, SIAM, 2020.
URL http://dx.doi.org/10.1137/1.9781611975994.168

The talk is about the problem of minimizing energy consumption on multiple machines with sleep states: We are given n jobs, each with a release date, a deadline and a processing time, and wish to schedule them on m parallel machines so as to minimize the total energy consumed. Machines can enter a sleep state and they consume no energy in this state. Each machine requires L units of energy to awaken from the sleep state and in its active state the machine can process jobs and consumes a unit of energy per unit time.

The core of the talk revolves around giving a 2-approximation algorithm for the single machine case. The algorithm is based on the solution of a linear programming relaxation, which can be decomposed into a convex combination of integer solutions. However none of them may be feasible since the linear programming relaxation has a strictly positive integrality gap. We discuss how such an integer solution can nevertheless be turned into a feasible solution without increasing the cost by too much. We conclude the talk by outlining how these ideas can be extended in order to obtain the first known constant approximation for the multiprocessor setting (with an approximation factor of 3).

3.2 General Framework for Metric Optimization Problems with Delay or with Deadlines

Yossi Azar (Tel Aviv University, IL)

Joint work of Yossi Azar, Noam Touitou
Main reference Yossi Azar, Noam Touitou: "General Framework for Metric Optimization Problems with Delay or with Deadlines", in Proc. of the 60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019, pp. 60–71, IEEE Computer Society, 2019.
URL https://doi.org/10.1109/FOCS.2019.00013

In this paper, we present a framework used to construct and analyze algorithms for online optimization problems with deadlines or with delay over a metric space. Using this framework, we present algorithms for several different problems. We present an $O(D^2)$ -competitive deterministic algorithm for online multilevel aggregation with delay on a tree of depth D, an exponential improvement over the $O(D^4 2^D)$ -competitive algorithm of Bienkowski et al. (ESA '16), where the only previously-known improvement was for the special case of deadlines by Buchbinder et al. (SODA '17). We also present an $O(\log^2 n)$ -competitive randomized algorithm for online service with delay over any general metric space of n points, improving upon the $O(\log^4 n)$ -competitive algorithm by Azar et al. (STOC '17). In addition, we present the problem of online facility location with deadlines. In this problem, requests arrive over time in a metric space, and need to be served until their deadlines by facilities that are opened momentarily for some cost. We also consider the problem of facility location with delay, in which the deadlines are replaced with arbitrary delay functions. For those problems, we present $O(\log^2 n)$ -competitive algorithms, with n the number of points in the metric space. The algorithmic framework we present includes techniques for the design of algorithms as well as techniques for their analysis.

3.3 Learning Augmented Energy Minimization via Speed Scaling

Etienne Bamas (EPFL – Lausanne, CH)

As power management has become a primary concern in modern data centers, computing resources are being scaled dynamically to minimize energy consumption. A classic speed scaling problem of Yao et al. (FOCS 1995) consists in scaling the speed of a processor dynamically to minimize energy under the constraint of finishing all the jobs in-between their release time and their deadline. This problem has been well studied both in offline and online settings. In this work, we initiate the study of the online speed scaling problem in a new framework in which machine learning predictions about the future can be integrated naturally. Inspired by recent work on learning-augmented online algorithms, we propose an algorithm which incorporates predictions in a black-box manner and outperforms any online algorithm if the accuracy is high, yet maintains provable guarantees if the prediction is very inaccurate.

3.4 Mechanisms for resource allocation

Shuchi Chawla (University of Wisconsin – Madison, US)

In this talk we will consider the problem of allocating multiple items to buyers who have values/preferences for subsets of items. We will consider a strategic setting where buyers can misreport their values so as to get as large an allocation as cheaply as possible. The algorithm knows the distributions from which buyers draw their values, but not the actual values themselves. Our goal is to design truthful mechanisms. I will talk about two objectives: social welfare maximization, where we want to maximize the sum of all buyers' values from their allocations; and revenue maximization, where we want to maximize the total payment the mechanism receives from all the buyers. In some contexts we will consider the online version of the problem.

This is a survey talk. I will describe what is known for the above problems under different assumptions on the instances, and what the main open problems are. The emphasis will be on algorithmic rather than economic issues.

3.5 The online k-taxi problem

Christian Coester (CWI – Amsterdam, NL)

Joint work of Christian Coester, Elias Koutsoupias
Main reference Christian Coester, Elias Koutsoupias: "The online k-taxi problem", in Proc. of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019, pp. 1136–1147, ACM, 2019.
URL http://dx.doi.org/10.1145/3313276.3316370

In the k-taxi problem, a generalization of the k-server problem, an algorithm controls k taxis to serve a sequence of requests in a metric space. A request consists of two points s and t, representing a passenger that wants to be carried by a taxi from s to t. For each request and without knowledge of future requests, the algorithm has to select a taxi to transport the passenger. The goal is to minimize the total distance traveled by all taxis. The problem comes in two flavors, called the easy and the hard k-taxi problem: In the easy k-taxi problem, the cost is defined as the total distance traveled by the taxis; in the hard k-taxi problem, the cost is only the distance of empty runs.

The easy k-taxi problem is exactly equivalent to the k-server problem. The talk will focus mostly on the hard version, which is substantially more difficult. For hierarchically separated trees, I will present a memoryless randomized algorithm with optimal competitive ratio $2^k - 1$. This implies an $O(2^k \log n)$-competitive algorithm for arbitrary n-point metrics, which is the first competitive algorithm for the hard k-taxi problem for general finite metrics and general k. I will also describe main ideas of an algorithm based on growing, shrinking and shifting regions which achieves a constant competitive ratio for three taxis on the line (abstracting the scheduling of three elevators).

3.6 A Tale of Santa Claus, Hypergraphs and Matroids

Sami Davies (University of Washington – Seattle, US)

Joint work of Sami Davies, Thomas Rothvoss, and Yihao Zhang
Main reference Sami Davies, Thomas Rothvoss, Yihao Zhang: "A Tale of Santa Claus, Hypergraphs and Matroids", in Proc. of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020, pp. 2748–2757, SIAM, 2020.
URL http://dx.doi.org/10.1137/1.9781611975994.167

A well-known problem in scheduling and approximation algorithms is the Santa Claus problem. Suppose that Santa Claus has a set of gifts, and he wants to distribute them among a set of children so that the least happy child is made as happy as possible. Here, the value that a child i has for a present j is of the form $p_{ij} \in \{0, p_j\}$. A polynomial time algorithm by Annamalai et al. gives a 12.33-approximation and is based on a modification of Haxell's hypergraph matching argument. In this paper, we introduce a matroid version of the Santa Claus problem. Our algorithm is also based on Haxell's augmenting tree, but with the introduction of the matroid structure we solve a more general problem with cleaner methods. Our result can then be used as a blackbox to obtain a $(4 + \varepsilon)$-approximation for Santa Claus. This factor also compares against a natural, compact LP for Santa Claus.

3.7 Optimally Handling Commitment Issues in Online Throughput Maximization

Franziska Eberle (Universität Bremen, DE)

Joint work of Franziska Eberle, Nicole Megow, and Kevin Schewior
Main reference Franziska Eberle, Nicole Megow, Kevin Schewior: "Optimally handling commitment issues in online throughput maximization", CoRR, Vol. abs/1912.10769, 2019.
URL https://arxiv.org/abs/1912.10769

We consider a fundamental online scheduling problem in which jobs with processing times and deadlines arrive online over time at their release dates. The task is to determine a feasible preemptive schedule on m machines that maximizes the number of jobs that complete before their deadline. Due to strong impossibility results for competitive analysis, it is commonly required that jobs contain some *slack* $\varepsilon > 0$, which means that the feasible time window for scheduling a job is at least $1 + \varepsilon$ times its processing time. In this paper, we answer the question on how to handle commitment requirements which enforce that a scheduler has to guarantee at a certain point in time the completion of admitted jobs. This is very relevant, e.g., in providing cloud-computing services and disallows last-minute rejections of critical tasks. We present the first online algorithm for handling commitment on parallel machines for small slack ε. When the scheduler must commit upon starting a job, the algorithm is $\Theta(\frac{1}{\varepsilon})$-competitive. Somewhat surprisingly, this is the same optimal performance bound (up to constants) as for scheduling without commitment on a single machine. If commitment decisions must be made before a job's slack becomes less than a δ-fraction of its size, we prove a competitive ratio of $O(\frac{\varepsilon}{\delta(\varepsilon--\delta)})$ for $0 < \delta < \varepsilon$. This result nicely interpolates between commitment upon starting a job and commitment upon arrival. For the latter commitment model, it is known that no (randomized) online algorithms does admit any bounded competitive ratio.

3.8 Pricing in Resource Allocation Games Based on Duality Gaps

Tobias Harks (Universität Augsburg, DE)

Main reference Tobias Harks: "Pricing in Resource Allocation Games Based on Duality Gaps", CoRR, Vol. abs/1907.01976, 2019.
URL http://arxiv.org/abs/1907.01976

We consider a basic resource allocation game, where the players' strategy spaces are subsets of R^m and cost/utility functions are parameterized by some common vector $u \in R^m$ and, otherwise, only depend on the own strategy choice. A strategy of a player can be interpreted as a vector of resource consumption and a joint strategy profile naturally leads to an aggregate consumption vector. Resources can be priced, that is, the game is augmented by a price vector $\lambda \in R^m_+$ and players have quasi-linear overall costs/utilities meaning that in addition to the original costs/utilities, a player needs to pay the corresponding price per consumed unit. We investigate the following question: for which aggregated consumption vectors u can we find prices λ that induce an equilibrium realizing the targeted consumption profile?

For answering this question, we revisit a well-known duality-based framework and derive several characterizations of the existence of such u and λ. We show that the characterization can help to unify parts of three largely independent streams in the literature – tolls in transportation systems, Walrasian market equilibria and congestion control in communication networks. Besides reproving existing results we establish novel existence results by using methods from polyhedral combinatorics, global optimization and discrete convexity.

3.9 Scheduling stochastic jobs with release dates on a single machine

Sven Jäger (TU Berlin, DE)

We consider the problem of minimizing the expected sum of weighted job completion times when jobs have stochastic processing times and may arrive over time. While in the offline model, all jobs are known upfront, jobs are only revealed at their release dates in the online model. The problem on identical parallel machines has been considered by Möhring, Schulz, and Uetz (1999), Megow, Uetz, and Vredeveld (2006), and Schulz (2008). Möhring, Schulz, and Uetz observed that their offline policy has a performance guarantee of 3 in the case of a single machine. A refined analysis of the policies developed by Schulz shows that in the single-machine case they are a 2.619-competitive deterministic online policy and a 2-competitive randomized online policy. These are also the best known performance guarantees of any offline policy. In the talk I will sketch how to obtain a (randomized) competitive ratio below 2 in the special case of NBUE processing times. This is based on a similar analysis as by Goemans, Queyranne, Schulz, Skutella, and Wang (2002) for the problem with deterministic processing times.

3.10 Online Learning with Vector Costs and Bandits with Knapsacks

Thomas Kesselheim (Universität Bonn, DE)

Joint work of Thomas Kesselheim , Sahil Singla

We introduce an online learning problem with vector costs (OLVC). Akin to online generalized load balancing, in each time step an algorithm chooses one of n actions and then incurs a vector cost $[0,1]^d$, which depends on the chosen action. The goal of the online algorithm is to minimize the ℓ_p norm of the sum of its cost vectors. The difference is that incurred costs are not known until after having chosen the action. This way, the setting generalizes the classical online learning setting, which is captured by $d = 1$.

We study OLVC in both stochastic and adversarial arrival settings, and give a general procedure to reduce the problem from d dimensions to a single dimension. This allows us to use classical online learning algorithms in both full and bandit feedback models to obtain (near) optimal results. In particular, we obtain a single algorithm (up to the choice of learning rate) that gives sublinear regret for stochastic arrivals and a tight $O(\min\{p, \log d\})$ competitive ratio for adversarial arrivals.

The OLVC problem also occurs as a natural subproblem when trying to solve the popular Bandits with Knapsacks (BWK) problem. This connection allows us to use our OLVCp techniques to obtain (near) optimal results for BWK in both stochastic and adversarial settings. In particular, we obtain a tight $O(\log d \cdot \log T)$ competitive ratio algorithm for adversarial BWK, which improves over the $O(d \cdot \log T)$ competitive ratio algorithm of Immorlica et al. (2019).

3.11 Equilibria in Atomic Splittable Congestion Games

Max Klimm (HU Berlin, DE)

Joint work of Max Klimm, Philipp Warode
Main reference Max Klimm, Philipp Warode: "Complexity and Parametric Computation of Equilibria in Atomic Splittable Congestion Games via Weighted Block Laplacians", in Proc. of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020, pp. 2728–2747, SIAM, 2020.
URL httpd://doi.org/10.1137/1.9781611975994.166

We settle the complexity of computing an equilibrium in atomic splittable congestion games with player-specific affine cost functions showing that it is PPAD-complete. To prove that the problem is contained in PPAD, we develop a homotopy method that traces an equilibrium for varying flow demands of the players. A key technique is to describe the evolution of the equilibrium locally by a novel block Laplacian matrix. This leads to a path following formulation where states correspond to supports that are feasible for some demands and neighboring supports are feasible for increased or decreased flow demands. A closer investigation of the block Laplacian system allows to orient the states giving rise to unique predecessor and successor states thus putting the problem into PPAD. For the PPAD-hardness, we reduce from computing an approximate equilibrium of a bimatrix win-lose game. As a byproduct of our reduction we further show that computing a multiclass Wardrop equilibrium with class-dependent affine cost functions is PPAD-complete as well. As a byproduct of our PPAD-completeness proof, we obtain an algorithm that computes all equilibria parametrized by the players' flow demands. For player-specific costs, this computation may require several increases and decreases of the demands leading to an algorithm that runs in polynomial space but exponential time. For player-independent costs only demand increases are necessary. If the coefficients be,i are in general position, this yields an algorithm computing all equilibria as a function of the flow demand running in time polynomial in the output.

3.12 Non-Clairvoyant Precedence Constrained Scheduling

Amit Kumar (Indian Institute of Technology – New Dehli, IN)

Joint work of Naveen Garg, Anupam Gupta, Amit Kumar, Sahil Singla
Main reference Naveen Garg, Anupam Gupta, Amit Kumar, Sahil Singla: "Non-Clairvoyant Precedence Constrained Scheduling", in Proc. of the 46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, July 9-12, 2019, Patras, Greece, LIPIcs, Vol. 132, pp. 63:1–63:14, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.
URL https://doi.org/10.4230/LIPIcs.ICALP.2019.63

We consider the online problem of scheduling jobs on identical machines, where jobs have precedence constraints. We are interested in the demanding setting where the jobs sizes are not known up-front, but are revealed only upon completion (the non-clairvoyant setting). Such precedence-constrained scheduling problems routinely arise in map-reduce and large-scale optimization. For minimizing the total weighted completion time, we give a constant-competitive algorithm. And for total weighted flow-time, we give an $O(1/\epsilon^2)$-competitive algorithm under $(1 + epsilon)$-speed augmentation and a natural ""no-surprises" assumption on release dates of jobs (which we show is necessary in this context). Our algorithm proceeds by assigning virtual rates to all waiting jobs, including the ones which are dependent on other uncompleted jobs. We then use these virtual rates to decide on the actual rates of minimal jobs (i.e., jobs which do not have dependencies and hence are eligible to run). Interestingly, the virtual rates are obtained by allocating time in a fair manner, using a Eisenberg-Gale-type convex program (which we can solve optimally using a primal-dual scheme). The optimality condition of this convex program allows us to show dual-fitting proofs more easily, without having to guess and hand-craft the duals. This idea of using fair virtual rates may have broader applicability in scheduling problems.

3.13 Online Vehicle Routing, the edge of optimization in large scale applications

Sebastien Martin (LYFT – New York, US)

Joint work of Dimitris Bertsimas, Patrick Jaillet, and Sébastien Martin
Main reference Dimitris Bertsimas, Patrick Jaillet, Sébastien Martin: "Online Vehicle Routing: The Edge of Optimization in Large-Scale Applications", Oper. Res., Vol. 67(1), pp. 143–162, 2019.
URL https://doi.org/10.1287/opre.2018.1763

With the emergence of ride-sharing companies that offer transportation on demand at a large scale and the increasing availability of corresponding demand data sets, new challenges arise to develop routing optimization algorithms that can solve massive problems in real time. In this paper, we develop an optimization framework, coupled with a novel and generalizable backbone algorithm, that allows us to dispatch in real time thousands of taxis serving more than 25,000 customers per hour. We provide evidence from historical simulations using New York City routing network and yellow cab data to show that our algorithms improve upon the performance of existing heuristics in such real-world settings.

3.14 Malleable Scheduling Beyond Identical Machines

Jannik Matuschke (KU Leuven, BE)

Joint work of Dimitris Fotakis, Jannik Matuschke, Orestis Papadigenopoulos
Main reference Dimitris Fotakis, Jannik Matuschke, Orestis Papadigenopoulos: "Malleable Scheduling Beyond Identical Machines", in Proc. of the Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2019, September 20-22, 2019, Massachusetts Institute of Technology, Cambridge, MA, USA, LIPIcs, Vol. 145, pp. 17:1–17:14, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.
URL http://dx.doi.org/10.4230/LIPIcs.APPROX-RANDOM.2019.17

In malleable scheduling, jobs can be executed simultaneously on multiple machines with the processing time depending on the number of allocated machines. Each job is required to be executed non-preemptively and in unison, i.e., it has to occupy the same time interval on all its allocated machines. In this talk, we discuss a generalization of malleable scheduling, in which a function $f(S, j)$ determines the processing time of job j on machine subset S. We derive a constant factor approximation for minimizing the makespan in the case that $f(S, j)$ can be expressed in terms of job-dependent machine speeds and fulfills a non-decreasing workload assumption for each fixed job j. We also discuss further generalizations and open problems.

3.15 Combinatorial Optimization Augmented with Machine Learning

Benjamin J. Moseley (Carnegie Mellon University – Pittsburgh, US)

Combinatorial optimization often focuses on optimizing for the worst-case. However, the best algorithm to use depends on the "relative inputs", which is application specific and often does not have a formal definition.

The talk gives a new theoretical model for designing algorithms that are tailored to inputs for the application at hand. In the model, learning is performed on past problem instances to make predictions on future instances. These predictions are incorporated into the design and analysis of the algorithm. The predictions can be used to achieve "instance-optimal" algorithm design when the predictions are accurate and the algorithm's performance gracefully degrades when there is error in the prediction.

The talk will apply this framework to applications in online algorithm design and give algorithms with theoretical performance that goes beyond worst-case analysis. The majority of the talk will focus on load balancing on unrelated machines.

3.16 Group Fairness in Network Design and Combinatorial Optimization

Kamesh Munagala (Duke University – Durham, US)

Joint work of Yu Cheng, Zhihao Jiang, Kamesh Munagala, Kangning Wang
Main reference Yu Cheng, Zhihao Jiang, Kamesh Munagala, Kangning Wang: "Group Fairness in Committee Selection", in Proc. of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019, pp. 263–279, ACM, 2019.
URL https://doi.org/10.1145/3328526.3329577

Consider the following classical network design model. There are n clients in a multi-graph with a single sink node. Each edge has a cost to buy, and a length if bought; typically, costlier edges have smaller lengths. There is a budget B on the total cost of edges bought. Given a set of bought edges, the distance of a client to the sink is the shortest path according to the edge lengths. Such a model captures buy-at-bulk network design and facility location as special cases.

Rather than pose this as a standard optimization problem, we ask a different question: Suppose a provider is allocating budget B to build this network, how should it do so in a manner that is fair to the clients? We consider a classical model of group fairness termed the core in cooperative game theory: If each client contributes its share B/n amount of budget as tax money, no subset of clients should be able to pool their tax money to build a different network that simultaneously improves all their distances to the sink. The question is: Does such a solution always exist, or approximately exist?

We consider an abstract "committee selection" model from social choice literature that captures not only the above problem, but other combinatorial optimization problems where we need to provision public resources subject to combinatorial constraints, in order to provide utility to clients. For this general model, we show that an approximately fair solution always exists, where the approximation scales down the tax money each client can use for deviation by only a constant factor. Our existence result relies on rounding an interesting fractional relaxation to this problem. In certain cases such as the facility location problem, it also implies a polynomial time algorithm. We conclude with several open questions.

3.17 Scheduling Bidirectional Traffic

Rolf H. Möhring (TU Berlin, DE)

Joint work of Elisabeth Lübbecke, Marco Lübbecke, Rolf H. Möhring
Main reference Elisabeth Lübbecke, Marco E. Lübbecke, Rolf H. Möhring: "Ship Traffic Optimization for the Kiel Canal", Oper. Res., Vol. 67(3), pp. 791–812, 2019.
URL https://doi.org/10.1287/opre.2018.1814

We introduce, discuss, and solve a hard practical optimization problem that deals with routing bidirectional traffic. This situation occurs in train traffic on a single track with sidings, ship traffic in a canal, or bidirectional data communication. We have developed a combinatorial algorithm that provides a unified view of routing and scheduling that combines joint (global) and sequential (local) solution approaches to allocate scarce network resources to a stream of online arriving vehicles in a collision-free manner. Computational experiments on real traffic data with results obtained by human expert planners show that our algorithm improves upon manual planning by 25%.

This combination of routing and scheduling leads to a new class of scheduling problems, and we will also address some complexity and approximation results for this class.

3.18 Stochastic Makespan Minimization

Viswanath Nagarajan (University of Michigan – Ann Arbor, US)

Joint work of Anupam Gupta, Amit Kumar, Viswanath Nagarajan, Xiangkun Shen
Main reference Anupam Gupta, Amit Kumar, Viswanath Nagarajan, Xiangkun Shen: "Stochastic Makespan Minimization in Structured Set Systems (Extended Abstract)", in Proc. of the Integer Programming and Combinatorial Optimization – 21st International Conference, IPCO 2020, London, UK, June 8-10, 2020, Proceedings, Lecture Notes in Computer Science, Vol. 12125, pp. 158–170, Springer, 2020.
URL http://dx.doi.org/10.1007/978-3-030-45771-6_13

We consider stochastic combinatorial optimization problems where the objective is to minimize the expected makespan. First, we provide a constant-factor approximation algorithm for stochastic makespan minimization on unrelated machines. Second, we provide an $O(\log \log m)$ approximation algorithm for stochastic resource allocation problems with some geometric structure, such as intervals in a line, paths in a tree and rectangles/disks in the plane. Both results utilize (i) an exponential-size LP based on the cumulant generating function and (ii) an iterative rounding algorithm.

3.19 Online Algorithms via Projections

Seffi Naor (Technion – Haifa, IL)

Joint work of Niv Buchbinder, Anupam Gupta, Marco Molinaro, Joseph (Seffi) Naor
Main reference Niv Buchbinder, Anupam Gupta, Marco Molinaro, Joseph (Seffi) Naor: "k-Servers with a Smile: Online Algorithms via Projections", in Proc. of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019, pp. 98–116, SIAM, 2019.
URL https://doi.org/10.1137/1.9781611975482.7

We present a new/old approach to the design of online algorithms via Bregman projections. This approach is applicable to a wide range of online problems and we discuss connections to previous work on online primal-dual algorithms. In particular, the k-server problem on trees and HSTs is considered. The projection-based algorithm for this problem turns out to have a competitive ratio that matches some of the recent results given by Bubeck et al. (STOC 2018), whose algorithm uses mirror-descent-based continuous dynamics prescribed via a differential inclusion.

3.20 Improved Approximation Algorithms for Inventory Problems

Neil Olver (London School of Economics and Political Science, GB)

Joint work of Thomas Bosman, Neil Olver
Main reference Thomas Bosman, Neil Olver: "Improved Approximation Algorithms for Inventory Problems", in Proc. of the Integer Programming and Combinatorial Optimization – 21st International Conference, IPCO 2020, London, UK, June 8-10, 2020, Proceedings, Lecture Notes in Computer Science, Vol. 12125, pp. 91–103, Springer, 2020.
URL https://doi.org/10.1007/978-3-030-45771-6_8

We give new approximation algorithms for the submodular joint replenishment problem and the inventory routing problem, using an iterative rounding approach. In both problems, we are given a set of N items and a discrete time horizon of T days in which given demands for the items must be satisfied. Ordering a set of items incurs a cost according to a set function, with properties depending on the problem under consideration. Demand for an item at time t can be satisfied by an order on any day prior to t, but a holding cost is charged for storing the items during the intermediate period; the goal is to minimize the sum of the ordering and holding cost. Our approximation factor for both problems is $O(\log \log \min(N, T))$; this improves exponentially on the previous best results.

3.21 Sample-Based Prophet Inequalities

Kevin Schewior (Universität Köln, DE)

Joint work of José R. Correa, Paul Dütting, Felix A. Fischer, and Kevin Schewior
Main reference José R. Correa, Paul Dütting, Felix A. Fischer, Kevin Schewior: "Prophet Inequalities for I.I.D. Random Variables from an Unknown Distribution", in Proc. of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019, pp. 3–17, ACM, 2019.
URL https://doi.org/10.1145/3328526.3329627

Consider a gambler who observes independent draws from a known sequence of positive-valued distributions. Upon observation of any value, the gambler has to decide whether to keep it as final reward or to discard it forever. The prophet inequality (Krengel, Sucheston and Garling 1978) states that there is a strategy whose expected accepted value is at least half of the expected maximum of all draws (the prophet's value). We first review the recent result (Wang 2018) that the following simple strategy also achieves the same ratio: Sample one value from each of the distributions, and set their maximum as a threshold for accepting any value.

We then turn to the case in which all distributions are identical. It is a simple corollary from results on the secretary problem that a ratio of $1/e$ is achievable without samples (again with respect to the expected maximum). We show using Ramsey's Theorem that this is best possible. We also show that knowing $O(n^2)$ samples is essentially as good as knowing the distribution, meaning that a ratio of 0.745 can be approached in that case (Correa et al. 2017). For the remainder of the talk, we work towards understanding the case with $O(n)$ samples, but, unlike for the distinct-distributions case, open questions remain.

3.22 Online Metric Algorithms with Untrusted Predictions

Bertrand Simon (Universität Bremen, DE)

Joint work of Antonios Antoniadis, Christian Coester, Marek Eliáš, Adam Polak, and Bertrand Simon
Main reference Antonios Antoniadis, Christian Coester, Marek Eliáš, Adam Polak, Bertrand Simon: "Online metric algorithms with untrusted predictions", CoRR, Vol. abs/2003.02144, 2020.
URL https://arxiv.org/abs/2003.02144

Machine-learned predictors, although achieving very good results for inputs resembling training data, cannot possibly provide perfect predictions in all situations. Still, decision-making systems that are based on such predictors need not only to benefit from good predictions but also to achieve a decent performance when the predictions are inadequate. In this paper, we propose a prediction setup for Metrical Task Systems (MTS), a broad class of online decision-making problems including, e.g., caching, k-server and convex body chasing. We utilize results from the theory of online algorithms to show how to make the setup robust. We extend our setup in two ways, (1) adapting it beyond MTS to the online matching on the line problem, and (2) specifically for caching, to achieve an improved dependence on the prediction error. Finally, we present an empirical evaluation of our methods on real world datasets, which suggests practicality.

3.23 Fixed-Order Scheduling on Parallel Machines

René Sitters (VU University of Amsterdam, NL)

Joint work of René Sitters, Thomas Bosman, Dario Frascaria, Neil Olver, René Sitters, and Leen Stougie
Main reference Thomas Bosman, Dario Frascaria, Neil Olver, René Sitters, Leen Stougie: "Fixed-Order Scheduling on Parallel Machines", in Proc. of the Integer Programming and Combinatorial Optimization – 20th International Conference, IPCO 2019, Ann Arbor, MI, USA, May 22-24, 2019, Proceedings, Lecture Notes in Computer Science, Vol. 11480, pp. 88–100, Springer, 2019.
URL https://doi.org/10.1007/978-3-030-17953-3_7

We consider the following natural scheduling problem: Given a sequence of jobs with weights and processing times, one needs to assign each job to one of m identical machines in order to minimize the sum of weighted completion times. The twist is that for machine the jobs assigned to it must obey the order of the input sequence, as is the case in multi-server queuing systems. We establish a constant factor approximation algorithm for this (strongly NP-hard) problem.

3.24 Approximation Algorithms for Replenishment Problems with Fixed Turnover Times

Leen Stougie (CWI – Amsterdam, NL)

Joint work of Thomas Bosman, Martijn van EE, Yang Jiao, Alberto Marchetti Spaccamela, R Ravi, Leen Stougie
Main reference Thomas Bosman, Martijn van Ee, Yang Jiao, Alberto Marchetti-Spaccamela, R. Ravi, Leen Stougie: "Approximation Algorithms for Replenishment Problems with Fixed Turnover Times", in Proc. of the LATIN 2018: Theoretical Informatics – 13th Latin American Symposium, Buenos Aires, Argentina, April 16-19, 2018, Proceedings, Lecture Notes in Computer Science, Vol. 10807, pp. 217–230, Springer, 2018.
URL https://doi.org/10.1007/978-3-319-77404-6_17

We introduce and study a class of optimization problems we call replenishment problems with fixed turnover times: a very natural model that has received little attention in the literature. Clients with capacity for storing a certain commodity are located at various places; at each client the commodity depletes within a certain time, the turnover time, which is constant but can vary between locations. Clients should never run empty. The natural feature that makes this problem interesting is that we may schedule a replenishment (well) before a client becomes empty, but then the next replenishment will be due earlier also. This added workload needs to be balanced against the cost of routing vehicles to do the replenishments. In this paper, we focus on the aspect of minimizing routing costs. However, the framework of recurring tasks, in which the next job of a task must be done within a fixed amount of time after the previous one is much more general and gives an adequate model for many practical situations. Note that our problem has an infinite time horizon. However, it can be fully characterized by a compact input, containing only the location of each client and a turnover time. This makes determining its computational complexity highly challenging and indeed it remains essentially unresolved. We study the problem for two objectives: min-avg minimizes the average tour cost and min-max minimizes the maximum tour cost over all days. For min-max we derive a logarithmic factor approximation for the problem on general metrics and a 6-approximation for the problem on trees, for which we have a proof of NP-hardness. For min-avg we present a logarithmic factor approximation on general metrics, a 2-approximation for trees, and a pseudopolynomial time algorithm for the line. Many intriguing problems remain open.

In this lecture I will explain the model and the complexity issues and give an intuitive idea of the approximation results on a tree.

3.25 Learning in Games and in Queueing Systems

Éva Tardos (Cornell University – Ithaca, US)

Over the last two decades we have developed good understanding how to quantify the impact of strategic user behavior on overall performance in many games (including traffic routing as well as online auctions), and showed that the resulting bounds extend to repeated games assuming players use a form of no-regret learning that helps them adapt to the environment. In this talk we will review these results, and study this phenomenon in the context of a game modeling queuing systems: routers compete for servers, where packets that do not get service will be resent at future rounds, resulting in a system where the number of packets at each round depends on the success of the routers in the previous rounds. In joint work with Jason Gaitonde, we analyze the resulting highly dependent random process and find that if the capacity of the servers is high enough to allow a centralized and knowledgeable scheduler to get all packets served even with double the packet arrival rate, then learning can help the queues in coordinating their behavior, the expected number of packets in the queues will remain bounded throughout time, assuming older packets have priority.

3.26 Approximation algorithms for traveling salesman problems

Vera Traub (Universität Bonn, DE)

In the traveling salesman problem we are given a finite set of cities with pairwise non-negative distances. The task is to find a shortest tour that visits all cities and returns to the starting point. The distances between cities can either be symmetric (TSP) or asymmetric (ATSP). We will also consider the path version, which is the generalization of the traveling salesman problem in which the endpoints of the tour are given and distinct.

For most of these traveling salesman problems improved approximation algorithms have been found during the past few years. The only exception is the symmetric TSP, where Christofides' classical 3/2-approximation from the 70's remains the best known approximation algorithm.

In this talk we survey the recent progress on approximation algorithms for both the symmetric and the asymmetric traveling salesman problem, as well as their path versions.

3.27 Greed...Is Good For Scheduling Under Uncertainty (2nd ed.)

Marc Uetz (University of Twente – Enschede, NL)

Joint work of Varun Gupta, Benjamin Moseley, Marc Uetz, Qiaomin Xie
Main reference Varun Gupta, Benjamin Moseley, Marc Uetz, Qiaomin Xie: "Greed Works—Online Algorithms for Unrelated Machine Stochastic Scheduling", Mathematics of Operations Research, Vol. 45(2), pp. 497–516, 2020.
URL https://doi.org/10.1287/moor.2019.0999

This spotlight talk reports about an update of a 2017 IPCO paper together with Varun Gupta, Ben Moseley, and Qiaomin Xie. In that updated paper, we show that a rather simple and intuitive greedy algorithm performs surprisingly well for a classical scheduling problem, namely minimizing the total weighted completion time on unrelated machines. In fact we give the first results for this problem when jobs are allowed to appear online over time, and have uncertain job sizes. Due to recent simplifications and improvements in both algorithm and analysis, our new performance bounds even improve upon previously best known results for the deterministic online problem, from the 1990s. The algorithm's basic idea is a greedy assignment of jobs to machines, just mimicking a "nominal" schedule where stochastic processing times are replaced by expectations. Moreover, the main idea for making this algorithm also competitive for uncertain job sizes, is is to adhere as much as possible to that nominal schedule. The analysis is based on dual fitting.

3.28 Nash flows over Time with Spillback and Kinematic Waves

Laura Vargas Koch (RWTH Aachen, DE)

Joint work of Leon Sering and Laura Vargas Koch
Main reference Leon Sering, Laura Vargas Koch: "Nash Flows Over Time with Spillback", in Proc. of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019, pp. 935–945, SIAM, 2019.
URL https://doi.org/10.1137/1.9781611975482.57

Modeling traffic in road networks is a widely studied but challenging problem, especially under the assumption that drivers act selfishly. A common approach is the deterministic queuing model, for which the structure of dynamic equilibria has been studied extensively in the last couple of years. The basic idea is to model traffic by a continuous flow that travels over time through a network, in which the arcs are endowed with transit times and capacities. Whenever the flow rate exceeds the capacity the flow particles build up a queue. So far it was not possible to represent the real-world phenomena spillback and kinematic waves in this model. By introducing a storage capacity arcs can become full, and thus, might block preceding arcs, i.e., spillback occurs. Furthermore, we model kinematic waves by upstream moving flows over time representing the gaps between vehicles. We carry over the main results of the original model to our generalization, i.e., we characterize Nash flows over time by sequences of particular static flows, so-called spillback thin flows. Furthermore, we give a constructive proof for the existence.

3.29 A Water-Filling Primal-Dual Algorithm for Approximating Non-Linear Covering Problems

Jose Verschae (O'Higgins University – Rancagua, CL)

Joint work of Andrés Fielbaum, Ignacio Morales, José Verschae
Main reference Andrés Fielbaum, Ignacio Morales, José Verschae: "A Water-Filling Primal-Dual Algorithm for Approximating Non-Linear Covering Problems", CoRR, Vol. abs/1912.12151, 2019.
URL https://arxiv.org/abs/1912.12151

Obtaining strong linear relaxations of capacitated covering problems constitute a major technical challenge even for simple settings. For one of the most basic cases, the Knapsack-Cover (Min-Knapsack) problem, the relaxation based on *knapsack-cover inequalities* achieves an integrality gap of 2. These inequalities are exploited in more general problems, many of which admit primal-dual approximation algorithms.

Inspired by problems from power and transport systems, we introduce a general setting in which items can be taken fractionally to cover a given demand. The cost incurred by an item is given by an arbitrary non-decreasing function of the chosen fraction. We generalize the knapsack-cover inequalities to this setting an use them to obtain a $(2+\varepsilon)$-approximate primal-dual algorithm. Our procedure has a natural interpretation as a bucket-filling algorithm, which effectively balances the difficulties given by having different slopes in the cost functions: when some superior portion of an item presents a low slope, it helps to increase the priority with which the inferior portions may be taken. We also present a rounding algorithm with an approximation guarantee of 2.

We generalize our algorithm to the Unsplittable Flow-Cover problem on a line, also for the setting where items can be taken fractionally. For this problem we obtain a $(4 + \varepsilon)$-approximation algorithm in polynomial time, almost matching the 4-approximation known for the classical setting.

3.30 Dynamic Approximate Maximum Independent Set of Intervals, Hypercubes and Hyperrectangles

Andreas Wiese (University of Chile, CL)

Joint work of Monika Henzinger, Stefan Neumann, Andreas Wiese
Main reference Monika Henzinger, Stefan Neumann, Andreas Wiese: "Dynamic Approximate Maximum Independent Set of Intervals, Hypercubes and Hyperrectangles", CoRR, Vol. abs/2003.02605, 2020.
URL https://arxiv.org/abs/2003.02605

Independent set is a fundamental problem in combinatorial optimization. While in general graphs the problem is essentially inapproximable, for many important graph classes there are approximation algorithms known in the offline setting. These graph classes include interval graphs and geometric intersection graphs, where vertices correspond to intervals/geometric objects and an edge indicates that the two corresponding objects intersect. We present the first dynamic approximation algorithms for independent set of intervals and geometric objects. They work in the fully dynamic model where in each update an interval/geometric object is inserted or deleted. Our algorithms are deterministic and have worst-case update times that are polylogarithmic for constant d and ϵ. We achieve the following approximation ratios:

- For independent set of intervals, we maintain $(1+\epsilon)$-approximate solutions for the unweighted and the weighted case.
- For independent set of d-dimensional hypercubes we maintain $(1+\epsilon)2^d$-approximate solutions in the unweighted case and $O(2^d)$-approximate solutions in the weighted case. Also, we show that for maintaining unweighted $(1+\epsilon)$-approximate solutions one needs polynomial update time for $d \geq 2$ if the ETH holds.
- For weighted d-dimensional hyperrectangles we present a dynamic algorithm with approximation ratio $(1+\epsilon)\log^{d-1} N$, assuming that the coordinates of all input hyperrectangles are in $[0,N]^d$ and each of their edges has length at least 1.

4 Open Problems

4.1 Deterministic min-cost matching with delays

Yossi Azar (Tel Aviv University, IL)

We are given a metric space. Requests arrive over time. Requests need to be matched in pairs. Pending requests can be matched at any time after their arrival. The goal is to minimize the cost of the matching (distance between matched requests) plus the total delay of all requests (the delay of a request is the time between its arrival until it is matched). Denote the size of the metric space by n and the number of requests by m.

Known Results

- $O(\log n)$ randomized competitive algorithm for arbitrary metric space [1, 2]
- A lower bound of $\Omega(\log n/\log\log n)$ for randomized algorithm (even for the metric space of n integer point $[0,n]$ on the line) [3]
- $O(m^{0.59})$ deterministic competitive algorithm and $O(n)$ deterministic competitive algorithms. No lower bounds better then the randomized lower bound are known. [4, 5, 6]

Open: Close the gap between $\log n$ and n (or $m^{0.59}$) for deterministic algorithms.

Variant (bi-chromatic): Requests are red or blue (suppliers vs costumers) and the matching is always between requests of different colors.

Known results: Similar results to the previous model except of a weaker lower bound of $\Omega(\sqrt{\log n/\log\log n})$ for the (randomized) competitive ratio [3,6].

Open: close the gap between $\log n$ and n (or $m^{0.59}$) for deterministic algorithms for the bi-chromatic case.

References

1 Yuval Emek, Shay Kutten, and Roger Wattenhofer *Online matching: haste makes waste!*. STOC 2016: 333-344.
2 Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. *Polylogarithmic bounds on the competitiveness of min-cost perfect matching with delays*. SODA 2017: 1038-1050.
3 Itai Ashlagi, Yossi Azar, Moses Charikar, Ashish Chiplunkar, Ofir Geri, Haim Kaplan, Rahul M. Makhijani, Yuyi Wang, and Roger Wattenhofer. *Min-cost bipartite perfect matching with delays*. APPROX-RANDOM 2017: 1:1-1:20

4 Marcin Bienkowski, Artur Kraska, and Pawel Schmidt. *A match in time saves nine: Deterministic online matching with delays.* WAOA 2017: 132-146.
5 Marcin Bienkowski, Artur Kraska, Hsiang-Hsuan Liu, and Pawel Schmidt. *A primal-dual online deterministic algorithm for matching with delays.* WAOA 2018: 51-68.
6 Yossi Azar, Amit Jacob Fanani: *Deterministic Min-Cost Matching with Delays.* Theory Comput. Syst. 64(4): 572-592 (2020) and WAOA 2018: 21-35

4.2 On-line Routing

Sanjoy Baruah (Washington Univerity in St. Louis, US)

Consider the following on-line routing problem on graphs. Each edge e of a directed graph is characterized by two upper bounds on the delay that will be encountered upon traversing it: an upper bound $c_W(e)$ on the maximum delay under all circumstances, and a (smaller) upper bound $c_T(e)$ on the maximum delay one would encounter under all "typical" (i.e., non-pathological) circumstances. The actual delay that will be encountered upon traversing an edge is unknown prior to actually traversing that edge.

Given a source vertex s, a destination vertex t,and a delay bound D, the objective is to travel from s to t such that one is guaranteed to arrive at t within D time units of leaving s, whilst simultaneously minimizing the duration taken in doing so under all typical (non-pathological) circumstances.

Algorithms have previously been proposed [1] for obtaining optimal such routes; however it is unknown whether or not the actual number of edges in all such optimal routes is polynomial in the problem specification.

References

1 Agrawal, K and Baruah, S.: *Adaptive Real-Time Routing in Polynomial Time*, Proceedings of the IEEE Real-Time Systems Symposium (RTSS 2019), Hong Kong, December 2019. IEEE Computer Society Press.

4.3 Integrality Gap of the natural LP for $P2|\text{prec}|C_{\max}$

Xinrui Jia (EPFL – Lausanne, CH)

This is an open problem stated in Kulkarni et al. 2020.

The problem is whether the natural LP relaxation for $P2|\text{prec}|C_{\max}$ has a large integrality gap when raised to $o(\log n)$ levels of the Sherali-Adams hierarchy. In the paper, a different scheduling problem, denoted $1|r_j, d_j| \sum_j p_j U'_j$, was demonstrated to have a large integrality gap when raised to $o(\log n)$ levels. This is the scheduling problem on one machine where jobs have release times r_j, deadlines d_j, processing times p_j, and the objective is to schedule the jobs non-preemptively for p'_j units of time $0 \le p'_j \le p_j$, to minimize $\sum_{j \in J}(p_j - -p'_j)$.

The authors present a way to transform instances of $1|r_j, d_j| \sum_j p_j U'_j$ into instances of $P2|\text{prec}|C_{\max}$, and show that an $o(\log n)$-level Sherali-Adams lift does not lead to a $(1+\epsilon)$ approximation for $1|r_j, d_j| \sum_j p_j U'_j$. The authors believe that the same instance used is the right one for showing an integrality gap for $P2|\text{prec}|C_{\max}$.

References

1 Kulkarni J., Li S., Tarnawski J., and Ye M. *Hierarchy-Based Algorithms for Minimizing Makespan under Precedence and Communication Constraints.* SODA, 2020.

4.4 The Busy Time Problem

Samir Khuller (Northwestern University – Evanston, US)

Given a set of jobs with release times, deadlines and processing times. The goal is to partition the jobs into bundles such that each bundle consists of a set of jobs that can be executed non-preemptively on a batch machine with batch capacity B. The total cost of this schedule is the sum of lengths of all the bundles. Each job has to be scheduled respecting its release time and deadline. The goal is to find a minimum cost schedule. Since the number of bundles is unbounded, a feasible solution can simply put each job alone in a bundle, or any B jobs in a bundle.

We are interested in approximation algorithms for this problem. The paper below presents a fairly simple "Greedy tracking" algorithm for this problem with a factor 3 guarantee. The worst example known for this algorithm is 2, so its true performance might be better.

References

1 Chang J., Khuller S., and Mukherjee K. *LP rounding and combinatorial algorithms for minimizing active and busy time.*. J. Sched. 20(6): 657-680 (2017)

4.5 Location Routing with Depot Capacities

Jannik Matuschke (KU Leuven, BE)

In the classic Facility Location Problem, we are given a set of possible locations for facilities with associated opening costs and a set of clients. We have to decide on which facilities to open. Then each client is served by its closest open facility and we have to pay the sum of all opening and connection costs. A drawback of this classic problem is that it assumes every client to receive its own dedicated connection, whereas in practice, several nearby clients can easily be served by the same vehicle. Ignoring these synergies and thus over-estimating connection costs can lead to inferior solutions whose cost significantly exceeds that of an optimal solution. This motivates the study of Location Routing, a combination of Facility Location and Vehicle Routing, in which clients are served from tours originating at open facilities.

Capacitated Location Routing Problem

Input: a set of clients C, a set of facilities F, a metric d on $C \cup F$, opening costs costs f_i for each $i \in F$, a facility capacity $B \in \mathbb{Z}_+$, a vehicle capacity $U \in \mathbb{Z}_+$

Task: Find a set facility S and a set of tours[1] $\mathcal{T}$ such that

1. $C \subseteq \bigcup_{T \in \mathcal{T}} V(T)$ (every client is served by a tour),
2. $|V(T) \cap S| = 1$ for all $T \in \mathcal{T}$ (every tour contains an open facility),
3. $|V(T) \cap C| \leq U$ for all $T \in \mathcal{T}$ (every tour serves at most U clients),
4. $\bigcup_{T \in \mathcal{T}: i \in V(T)} |V(T) \cap C| \leq B$ for all $i \in S$ (every facility serves at most B clients),

minimizing the total cost $\sum_{i \in S} f_i + \sum_{T \in \mathcal{T}} d(T)$.

Open question: Is there a constant factor approximation algorithm for the Capacitated Location Routing Problem?

Known results. For the case that facilities are uncapacitated ($B = \infty$), a constant factor approximation is known [1, 2], based on the combination of two lower bounds derived from instances of Facility Location and Minimum Spanning Tree, respectively. For the general problem with capacitated facilities, similar lower bounds can be combined with the classic LP rounding approach for unrelated machine scheduling [3] to obtain the following bifactor approximation:

Theorem 1. [4] There is an algorithm that, given an instance of Capacitated Location Routing and a number $\gamma \in (0, 1)$, computes in polynomial time a solution fulfilling conditions 1 to 3, with cost at most $(2 + 6/(1 - \gamma))\,\mathrm{OPT}$ and a maximum facility load of $(3/2 + 1/\gamma)B$.

It has further been observed that the known lower bounds based on Facility Location/Minimum Spanning Tree are not sufficient to derive constant factor approximations without exceeding the capacities [4].

References

1 R. Ravi R. and Sinha A. *Approximation algorithms for problems combining facility location and network design.* Operations Research, 54:73-81, 2006.
2 Harks T., König F.G., and Matuschke J. *Approximation algorithms for capacitated location routing.* Transportation Science, 47:3-22, 2013.
3 Lenstra J.K., Shmoys D.B., and Tardos E. *Approximation algorithms for scheduling unrelated parallel machines.* Mathematical Programming, 46:259-271, 1990.
4 Demleitner A. *Approximation Algorithms for Location Routing with Depot Capacities.* Master Thesis, Technische Universität München, 2019.

4.6 Scheduling on two types of machines

Bertrand Simon (Universität Bremen, DE)

We consider the problem of scheduling tasks with precedence constraints on several machines in order to minimize the makespan. The originality is that these machines are composed of two types of unrelated machines. Hence, this problem lies between the setting of identical and

[1] A tour T consists of a node set $V(T) \subseteq C \cup F$ and a permutation σ of $V(T)$. The length of T is $d(T) = \sum_{v \in V(T)} d(v, \sigma(v))$.

unrelated machines. Current results for the offline setting include a $3 + 2\sqrt{2}$-approximation (~ 5.8), and a conditional lower bound of 3 on the approximation ratio, assuming a variant of the unique games conjecture. Hence, there is still a gap between these bounds.

References

1 Fagnon V., Kacem I., Lucarelli G., and Simon B. *Scheduling on Hybrid Platforms: Improved Approximability Window.* LATIN, 2020.

2 Kedad-Sidhoum S., Monna F., and Trystram d. *Scheduling tasks with precedence constraints on hybrid multi-core machines.* IPDPS Workshop, 2015.

Participants

- Antonios Antoniadis
MPI für Informatik – Saarbrücken, DE
- Yossi Azar
Tel Aviv University, IL
- Etienne Bamas
EPFL – Lausanne, CH
- Sanjoy K. Baruah
Washington Univerity in St. Louis, US
- Shuchi Chawla
University of Wisconsin – Madison, US
- Christian Coester
CWI – Amsterdam, NL
- Sami Davies
University of Washington – Seattle, US
- Christoph Dürr
UPMC – Paris, FR
- Franziska Eberle
Universität Bremen, DE
- Thomas Erlebach
University of Leicester, GB
- Naveen Garg
Indian Institute of Technology – New Delhi, IN
- Tobias Harks
Univeristät Augsburg, DE
- Ruben Hoeksma
University of Twente – Enschede, NL
- Sungjin Im
University of California – Merced, US
- Sven Jäger
TU Berlin, DE
- Xinrui Jia
EPFL – Lausanne, CH
- Thomas Kesselheim
Universität Bonn, DE
- Samir Khuller
Northwestern University – Evanston, US
- Max Klimm
HU Berlin, DE
- Peter Kling
Universität Hamburg, DE
- Amit Kumar
Indian Inst. of Technology – New Dehli, IN
- Marilena Leichter
TU München, DE
- Jan Karel Lenstra
CWI – Amsterdam, NL
- Alberto Marchetti-Spaccamela
Sapienza University of Rome, IT
- Sebastien Martin
LYFT – New York, US
- Jannik Matuschke
KU Leuven, BE
- Nicole Megow
Universität Bremen, DE
- Rolf H. Möhring
TU Berlin, DE
- Sarah Morell
TU Berlin, DE
- Benjamin J. Moseley
Carnegie Mellon University – Pittsburgh, US
- Kamesh Munagala
Duke University – Durham, US
- Viswanath Nagarajan
University of Michigan – Ann Arbor, US
- Seffi Naor
Technion – Haifa, IL
- Neil Olver
London School of Economics and Political Science, GB
- Britta Peis
RWTH Aachen, DE
- Kirk Pruhs
University of Pittsburgh, US
- Jens Quedenfeld
TU München, DE
- Shijin Rajakrishnan
Cornell University – Ithaca, US
- Lars Rohwedder
EPFL – Lausanne, CH
- Thomas Rothvoss
University of Washington – Seattle, US
- Guido Schäfer
CWI – Amsterdam, NL
- Kevin Schewior
Universität Köln, DE
- Jiri Sgall
Charles University – Prague, CZ
- David Shmoys
Cornell University, US
- Bertrand Simon
Universität Bremen, DE
- René Sitters
VU University of Amsterdam, NL
- Martin Skutella
TU Berlin, DE
- Clifford Stein
Columbia University, US
- Leen Stougie
CWI – Amsterdam, NL
- Ola Svensson
EPFL – Lausanne, CH
- Éva Tardos
Cornell University – Ithaca, US
- Vera Traub
Universität Bonn, DE
- Marc Uetz
University of Twente – Enschede, NL
- Rob van Stee
Universität Siegen, DE
- Laura Vargas Koch
RWTH Aachen, DE
- Victor Verdugo
London School of Economics, GB
- Jose Verschae
Pontifical Catholic University of Chile – Santiago, CL
- Tjark Vredeveld
Maastricht University, NL
- Andreas Wiese
Universidad de Chile – Santiago, CL


KENT STATE