

Scheduling

Claire Mathieu^{*1}, Nicole Megow^{*2}, Benjamin J. Moseley^{*3},
Frits C. R. Spijksma^{*4}, and Alexander Lindermayr^{†5}

1 CNRS, Paris, FR. clairemmathieu@gmail.com

2 University of Bremen, DE. nicole.megow@uni-bremen.de

3 Carnegie Mellon University, Pittsburgh, USA. moseleyb@andrew.cmu.edu

4 TU Eindhoven, NL. f.c.r.spijksma@tue.nl

5 Universität Bremen, DE. linderal@uni-bremen.de

Abstract

This report documents the program and outcomes of Dagstuhl Seminar 25121, “Scheduling”. The seminar focused on bridging traditional algorithmic scheduling with the emerging field of fairness in resource allocation. Scheduling is a longstanding research area that has been studied from both practical and theoretical perspectives in computer science, mathematical optimization, and operations research for over 70 years. Fairness has become a key concern in recent years, particularly in the context of resource allocation and scheduling, where it naturally arises in applications such as kidney exchange, school choice, and political districting. The seminar centered on three main themes: (1) fair allocation, (2) fairness versus quality of service, and (3) modeling aspects of fairness in scheduling.

Seminar March 16–21, 2025 – <https://www.dagstuhl.de/25121>

2012 ACM Subject Classification Theory of computation → Scheduling algorithms; Mathematics of computing → combinatorial optimization; Theory of computation → Approximation algorithms analysis

Keywords and phrases scheduling, fairness, mathematical optimization, algorithms and complexity, uncertainty

Digital Object Identifier 10.4230/DagRep.15.3.94

1 Executive Summary

Claire Mathieu (CNRS, Paris, FR)

Nicole Megow (University of Bremen, DE)

Benjamin J. Moseley (Carnegie Mellon University, Pittsburgh, USA)

Frits C. R. Spijksma (TU Eindhoven, NL)

License  Creative Commons BY 4.0 International license

© Claire Mathieu, Nicole Megow, Benjamin J. Moseley, and Frits C. R. Spijksma

This Dagstuhl Seminar was number 8 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. While there has been remarkable progress on algorithmic theory for fundamental scheduling problems, leading to insights for other fields as well, scheduling has proven to be an inspirational ground for new questions.

* Editor / Organizer

† Editorial Assistant / Collector



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

Scheduling, *Dagstuhl Reports*, Vol. 15, Issue 3, pp. 94–112

Editors: Claire Mathieu, Nicole Megow, Benjamin J. Moseley, Frits C. R. Spijksma, and Alexander Lindermayr



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

At this meeting, we have focussed on emerging models for **fairness in scheduling and resource allocation**. Traditionally, scheduling theory has focused on how to allocate resources to optimize quality of service guarantees, throughput, or efficiency. However, these objectives do not consider fairness to the underlying agents or entities.

An example of fairness considerations in government resource allocation can be observed in the distribution of healthcare resources, especially during times of crises like pandemic. A government may have a limited amount of resources available to distribute. One could distribute these to the areas affected most by the outbreak. It may also be important though to consider trade-offs between areas that traditionally have had disparities and are underfunded, ensuring vulnerable populations are not neglected. A government must balance immediate needs with long-term equability. Other examples of fairness in resource allocation and scheduling naturally arise in kidney exchange, school choice, tournament design, as well as political districting. These exciting and socially important problems demand to be better understood.

This seminar focused on three complementary themes.

- *Fair Allocation*. Fair allocation has taken center stage in multi-agent systems and economics over the past decade due to its significance both industrially and socially. Essentially, it addresses how to distribute items, whether they be goods or tasks, to agents in a way that leaves each content with their share. Notably, when dealing with indivisible items, perfect fairness metrics like envy-freeness and proportionality aren't always achievable. Recent research endeavors focus on creating algorithms that approximate these fairness standards. On the other hand, game theory delves into the challenge of fairly dividing resources among individuals with entitlements, a dilemma found in numerous real-world scenarios, from inheritance divisions to electronic frequency allocations. Fundamental to fair division is the belief that the involved parties, perhaps with the aid of a mediator, should carry out the allocation, as they best understand their value assessments. A classic example of a fair division method is the “divide and choose” algorithm, which ensures that two participants each feel they have received the most favorable portion. The vast landscape of fair division research extends this principle to more intricate contexts, adapting to varying goods, fairness criteria, player characteristics, and other evaluation standards. Fair allocation, resource allocation and scheduling are fields that build on one another as often algorithmic and analysis techniques in one find uses in the others.
- *Balancing Fairness and Quality of Service*. In the algorithms community, striking a balance between fairness and quality of service (QoS) is a pressing concern. While algorithms, particularly in sectors like finance, healthcare, and social networking, play a pivotal role in decision-making, ensuring equitable outcomes without compromising efficiency or performance is challenging. Fairness ensures that no group or individual is unfairly disadvantaged or discriminated against by algorithmic decisions, and it aims to create an even playing field across diverse sets of users or stakeholders. On the other hand, quality of service emphasizes responsiveness, reliability, and overall user satisfaction. Balancing these two elements is challenging analytically. The area requires a deep understanding of how to model the trade-offs and algorithmically balance quality of service and fairness.
- *Modeling Fairness*. Modeling fairness in scheduling and resource allocation presents a plethora of challenges. Scheduling and allocating resources inherently involves making decisions that prioritize certain tasks, individuals, or groups over others, which can inadvertently introduce biases or create disparities. One fundamental challenge lies in

defining what “fairness” actually means in varied contexts, as it can be subjective and differ across stakeholders. Even when fairness is well-defined, achieving it can sometimes conflict with optimizing for efficiency or maximum resource utilization. Additionally, when dealing with diverse sets of resources and stakeholders with distinct needs and preferences, ensuring equitable distribution becomes complex. There is also the issue of unseen biases in historical data, which, when used to train algorithms, can perpetuate past inequities. Furthermore, there is a constant need to balance immediate and long-term fairness, especially when resource availability fluctuates. Navigating these intricacies requires a deep understanding of real-world challenges to develop sound models for scheduling and resource allocation problems.

Organization of the Seminar. The seminar brought together 42 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 seminar, 29 talks of different lengths took place. Five keynote speakers gave an overview of the state-of-the art of the respective area or presented recent highlight results in 60 minutes:

- Adrian Vetta: Six Candidates Suffice to Win a Voter Majority
- Swati Gupta: Fair Resource Allocation from Theory to Practice
- Lars Rohwedder: The Santa Claus Problem: Three Perspectives
- Kavitha Telikepalli: Fair solutions to the house allocation problem
- Ulrike Schmidt-Kraepelin: Proportional Representation in Budget Allocation.

The remaining slots were filled with shorter talks of 30 minutes on various topics related to the intersection of fairness, social choice, and scheduling.

Outcome. Organizers and participants regard the seminar as a great success. The seminar 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 stimulating; participants actively met in working groups in the afternoon and evenings. It was remarked positively that a significant number of younger researchers (postdocs and PhD students) participated and integrated well.

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

2 Table of Contents

Executive Summary

Claire Mathieu, Nicole Megow, Benjamin J. Moseley, and Frits C. R. Spieksma . . . 94

Overview of Talks

Lossless Robustification of Packet Scheduling Algorithms <i>Yossi Azar</i>	99
Fair Strategic Facility Location with Predictions <i>Eric Balkanski</i>	99
Lift-and-Project Integrality Gaps for Santa Claus <i>Etienne Bamas</i>	100
Minimax Group Fairness in Strategic Classification <i>Emily Diana</i>	100
A Tight $(3/2 + \epsilon)$ -Approximation Algorithm for Demand Strip Packing <i>Franziska Eberle</i>	101
Students in highly competitive markets: the case of New York City specialized high schools <i>Yuri Faenza</i>	101
Fair Resource Allocation: From Theory to Practice <i>Swati Gupta</i>	102
Online Scheduling via Gradient Descent <i>Sungjin Im</i>	102
Fair solutions to the house allocation problem <i>Telikepalli Kavitha</i>	103
Supermodular Approximation of Norms and Applications <i>Thomas Kesselheim</i>	103
FPT Algorithms using Minimal Parameters for a Generalized Version of Maximin Shares <i>Alexandra Lassota</i>	104
A Little Clairvoyance Is All You Need <i>Alexander Lindermayr</i>	104
The Power of Proportional Fairness and Unifying Scheduling Algorithms for Group Completion Times <i>Nicole Megow</i>	105
Minimum Cost Adaptive Submodular Cover <i>Viswanath Nagarajan</i>	105
Near-Optimal PCM Wear-Leveling Under Adversarial Attacks <i>Seffi Naor</i>	106
Robust Gittins for Stochastic Scheduling <i>Heather Newman</i>	107
Fair Caching <i>Debmalya Panigrahi</i>	107

The Santa Claus Problem – Three Perspectives	
<i>Lars Rohwedder</i>	108
Optimal Online Discrepancy Minimization	
<i>Thomas Rothvoss</i>	108
Stochastic scheduling with Bernoulli-type jobs through policy stratification	
<i>Kevin Schewior</i>	108
Proportional Representation in Budget Allocation	
<i>Ulrike Schmidt-Kraepelin</i>	109
A new deterministic approximation for graph burning	
<i>Jiri Sgall</i>	109
A Simple Algorithm for Dynamic Carpooling with Recourse	
<i>Cliff Stein</i>	110
Six Candidates Suffice to Win a Voter Majority	
<i>Adrian Vetta</i>	110
The Power of Migrations in Dynamic Bin Packing	
<i>Rudy Zhou</i>	110
Participants	112

3 Overview of Talks

3.1 Lossless Robustification of Packet Scheduling Algorithms

Yossi Azar (Tel Aviv University, IL)

License © Creative Commons BY 4.0 International license
© Yossi Azar

Joint work of Yossi Azar, Or Vardi

Heuristics on what online algorithms should do at any given time can give large improvements to the performance of the algorithm. Today, such heuristics are mostly generated by some machine learning algorithm that was trained on what is hoped to be a similar input. We consider the online packets scheduling problem where unit size packets arrive over time, each is associated with a value and a deadline. The goal is to schedule the packets to maximize the value of the packets transmitted by their deadline. We consider an arbitrary algorithm (heuristic) and robustify it without loss. Specifically, we provide an algorithm that is at least as good as the heuristic for any input, while proving $O(1)$ competitiveness no matter how bad the heuristic is. For subclass of certain algorithms (called prediction upon arrival heuristic), we even provide a better robustness bound that provably cannot be achieved for general heuristics. Finally, we show that it is not possible to be as good as the prediction and remain $O(1)$ competitive if we consider the asynchronous model.

3.2 Fair Strategic Facility Location with Predictions

Eric Balkanski (Columbia University – New York, US)

License © Creative Commons BY 4.0 International license
© Eric Balkanski

Joint work of Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, Golnoosh Shahkarami, Xizhi Tan

Main reference Priyank Agrawal, Eric Balkanski, Vasilis Gkatzelis, Tingting Ou, Xizhi Tan: “Learning-Augmented Mechanism Design: Leveraging Predictions for Facility Location”, *Math. Oper. Res.*, Vol. 49(4), pp. 2626–2651, 2024.

URL <https://doi.org/10.1287/MOOR.2022.0225>

Main reference Eric Balkanski, Vasilis Gkatzelis, Golnoosh Shahkarami: “Randomized Strategic Facility Location with Predictions”, in *Proc. of the Advances in Neural Information Processing Systems*, Vol. 37, pp. 35639–35664, Curran Associates, Inc., 2024.

URL https://proceedings.neurips.cc/paper_files/paper/2024/file/3ec7806669b4048cdba4d1defc76ace3-Paper-Conference.pdf

In the strategic facility location problem, a set of agents report their locations in a metric space and the goal is to use these reports to open a new facility, minimizing an aggregate distance measure from the agents to the facility. However, agents are strategic and may misreport their locations to influence the facility’s placement in their favor. The aim is to design truthful mechanisms, ensuring agents cannot gain by misreporting. This problem was recently revisited through the learning-augmented framework, aiming to move beyond worst-case analysis and design truthful mechanisms that are augmented with (machine-learned) predictions. In this talk, I will focus on recent results for the egalitarian social cost objective, where the goal is to minimize the distance between the facility and the location of the agent who is the farthest from the facility.

3.3 Lift-and-Project Integrality Gaps for Santa Claus

Etienne Bamas (ETH Zürich, CH)

License © Creative Commons BY 4.0 International license
© Etienne Bamas

Main reference Étienne Bamas: “Lift-and-Project Integrality Gaps for Santa Claus”, in Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025, pp. 572–615, SIAM, 2025.

URL <https://doi.org/10.1137/1.9781611978322.18>

In this talk, I will focus on the MaxMinDegree Arborescence (MMDA) problem in layered directed graphs of depth $\ell \leq O(\log n / \log \log n)$, which is a key special case of the Santa Claus problem. The only way we have to solve the MMDA problem within a polylogarithmic factor is via an elegant recursive rounding of the $(\ell - 1)^{th}$ level of the Sherali-Adams hierarchy. However, it remains plausible that one could obtain a polylogarithmic approximation in polynomial time by using the same rounding with only 1 round of the Sherali-Adams hierarchy. As a main result, we rule out this possibility by constructing an MMDA instance of depth 3 for which a polynomial integrality gap survives 1 round of the Sherali-Adams hierarchy. This result is tight since it is known that after only 2 rounds the gap is at most polylogarithmic on depth-3 graphs. I will conclude the talk by related open problems.

3.4 Minimax Group Fairness in Strategic Classification

Emily Diana (TTIC – Chicago, US)

License © Creative Commons BY 4.0 International license
© Emily Diana

Joint work of Emily Diana, Saeed Sharifi-Malvajerdi, Ali Vakilian

Main reference Emily Diana, Saeed Sharifi-Malvajerdi, Ali Vakilian: “Minimax Group Fairness in Strategic Classification”, in Proc. of the IEEE Conference on Secure and Trustworthy Machine Learning, SaTML 2025, Copenhagen, Denmark, April 9-11, 2025, pp. 753–772, IEEE, 2025.

URL <https://doi.org/10.1109/SaTML64287.2025.00047>

In strategic classification, agents manipulate their features, at a cost, to receive a positive classification outcome from the learner’s classifier. The goal of the learner in such settings is to learn a classifier that is robust to strategic manipulations. While the majority of works in this domain consider accuracy as the primary objective of the learner, in this work, we consider learning objectives that have group fairness guarantees in addition to accuracy guarantees. We work with the minimax group fairness notion that asks for minimizing the maximal group error rate across population groups. Motivating examples will be focused on situations where agents are competing for resources and the classification decision influences allocation policies.

3.5 A Tight $(3/2 + \varepsilon)$ -Approximation Algorithm for Demand Strip Packing

Franziska Eberle (TU Berlin, DE)

License © Creative Commons BY 4.0 International license
 © Franziska Eberle
Joint work of Franziska Eberle, Felix Hommelsheim, Malin Rau, Stefan Walzer
Main reference Franziska Eberle, Felix Hommelsheim, Malin Rau, Stefan Walzer: “A Tight $(3/2 + \varepsilon)$ -Approximation Algorithm for Demand Strip Packing”, in Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025, pp. 641–699, SIAM, 2025.
URL <https://doi.org/10.1137/1.9781611978322.20>

We consider the Demand Strip Packing problem (DSP), in which we are given a set of jobs, each specified by a processing time and a demand. The task is to schedule all jobs such that they are finished before some deadline D while minimizing the peak demand, i.e., the maximum total demand of tasks executed at any point in time. DSP is closely related to the Strip Packing problem (SP), in which we are given a set of axis-aligned rectangles that must be packed into a strip of fixed width while minimizing the maximum height. DSP and SP are known to be NP-hard to approximate to within a factor below $\frac{3}{2}$.

To achieve the essentially best possible approximation guarantee, we prove a structural result. Any instance admits a solution with peak demand at most $(\frac{3}{2} + \varepsilon)\text{OPT}$ satisfying one of two properties. Either (i) the solution leaves a gap for a job with demand OPT and processing time $\mathcal{O}(\varepsilon D)$ or (ii) all jobs with demand greater than $\frac{\text{OPT}}{2}$ appear sorted by demand in immediate succession. We then provide two efficient algorithms that find a solution with maximum demand at most $(\frac{3}{2} + \varepsilon)\text{OPT}$ in the respective case. A central observation, which sets our approach apart from previous ones for DSP, is that the properties (i) and (ii) need not be efficiently decidable: We can simply run both algorithms and use whichever solution is the better one.

3.6 Students in highly competitive markets: the case of New York City specialized high schools

Yuri Faenza (Columbia University – New York, US)

License © Creative Commons BY 4.0 International license
 © Yuri Faenza
Joint work of Yuri Faenza, Swati Gupta, Xuan Zhang
Main reference Yuri Faenza, Swati Gupta, Xuan Zhang: “Discovering Opportunities in New York City’s Discovery Program: Disadvantaged Students in Highly Competitive Markets”, in Proc. of the 24th ACM Conference on Economics and Computation, EC 2023, London, United Kingdom, July 9-12, 2023, p. 585, ACM, 2023.
URL <https://doi.org/10.1145/3580507.3597762>

Eight among the most competitive high schools of the New York City Department of Education (NYC DOE) admit students only based on their score on a test, called SHSAT. 20% of these seats are reserved for students that the NYC DOE classifies, mostly following economic criteria, as disadvantaged. We show that the mechanism currently employed by the NYC DOE to assign these reserved seats creates a significant incentive for disadvantaged students to underperform, and we study alternatives. In particular, we highlight the superiority of one such alternative under the new ex-post hypothesis of High competitiveness (HC) of the market. We also give sufficient ex-ante conditions under which the HC hypothesis is satisfied with high probability. To prove such results, we rely on generalizations of Gale and Shapley’s

marriage model involving choice functions, and on the classical occupancy problem. Using 12 years of data, we show that the NYC DOE market that originated our work satisfies the HC hypothesis.

3.7 Fair Resource Allocation: From Theory to Practice

Swati Gupta (MIT – Cambridge, US)

License  Creative Commons BY 4.0 International license
© Swati Gupta

Joint work of Swati Gupta, Jai Moondra, Mohit Singh, Cheol Woo Kim, Shresth Verma, Madeleine Pollack, Lingkai Kong, Milind Tambe

Fairness in resource allocation is a fundamental problem that arises in a variety of domains, including healthcare, hiring, admissions, infrastructure development, recommendation systems, disaster management, and emergency response. Different ethical theories provide distinct lenses through which fairness can be understood and operationalized. In this talk, I will discuss (i) what it means to be fair in static and dynamic settings, depending on the application context, (ii) theoretical models for understanding noise and bias in data, and (iii) connections with law and policy. Through some of my recent work, I will discuss challenges related to differences in fairness objectives (e.g., how to find some “good” enough solutions across all objectives), navigating the space of human-AI collaboration (e.g., what should AI optimize?), and deviations from theoretical assumptions (e.g., of clean group memberships, discrimination models, etc).

3.8 Online Scheduling via Gradient Descent

Sungjin Im (University of California at Santa Cruz, US)

License  Creative Commons BY 4.0 International license
© Sungjin Im

Joint work of Qingyun Chen, Sungjin Im, Aditya Petety

Main reference Qingyun Chen, Sungjin Im, Aditya Petety: “Online Scheduling via Gradient Descent for Weighted Flow Time Minimization”, in Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025, pp. 3802–3841, SIAM, 2025.

URL <https://doi.org/10.1137/1.9781611978322.128>

In this talk, I will show how a generalization of the shortest remaining time first (SRPT) scheduling algorithm can be effectively used for various scheduling problems to minimize total weighted flow time. Essentially, SRPT can be interpreted as gradient descent on an estimate of the remaining jobs’ cost. In particular, we show that gradient descent is effective when the residual estimate possesses supermodularity, and that this supermodularity can be achieved when the scheduling constraints induce gross substitute valuations in the Walrasian Market.

3.9 Fair solutions to the house allocation problem

Telikepalli Kavitha (Tata Institute of Fundamental Research – Mumbai, IN)

License © Creative Commons BY 4.0 International license
© Telikepalli Kavitha
Joint work of Tamas Kiraly, Jannik Matuschke, Ildiko Schlotter, Ulrike Schmidt-Kraepelin
Main reference Telikepalli Kavitha, Tamás Király, Jannik Matuschke, Ildikó Schlotter, Ulrike Schmidt-Kraepelin: “The popular assignment problem: when cardinality is more important than popularity”, in Proc. of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 – 12, 2022, pp. 103–123, SIAM, 2022.
URL <https://doi.org/10.1137/1.9781611977073.6>

Matching problems with one-sided preferences are seen in many applications such as campus housing allocation in universities. Popularity is a well-studied notion of fairness that captures collective welfare. This talk will be on some simple algorithms to find popular solutions for matching problems in this model.

3.10 Supermodular Approximation of Norms and Applications

Thomas Kesselheim (Universität Bonn, DE)

License © Creative Commons BY 4.0 International license
© Thomas Kesselheim
Joint work of Thomas Kesselheim, Marco Molinaro, Sahil Singla
Main reference Thomas Kesselheim, Marco Molinaro, Sahil Singla: “Supermodular Approximation of Norms and Applications”, in Proc. of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24–28, 2024, pp. 1841–1852, ACM, 2024.
URL <https://doi.org/10.1145/3618260.3649734>

Many classic scheduling problems can be understood as minimizing a norm objective: Most prominently, the Makespan is nothing but the ℓ_∞ norm of the vector of machine loads. Every additive objective, like for example in Set Cover, can also be understood as an ℓ_1 norm. Over the years, a lot of results have been generalized to ℓ_p norms.

In this talk, we discuss techniques and results to go beyond ℓ_p norms. With a particular focus on online problems, we identify supermodularity—often reserved for combinatorial set functions and characterized by monotone gradients—as a defining feature. Every ℓ_p -norm is p -supermodular, meaning that its p^{th} power function exhibits supermodularity. The association of supermodularity with norms offers a new lens through which to view and construct algorithms.

For a large class of problems p -supermodularity is a sufficient criterion for developing good algorithms. Moreover, we show that every symmetric norm can be $O(\log m)$ -approximated by an $O(\log m)$ -supermodular norm, resulting in $O(\text{poly } \log m)$ -competitive algorithms for load balancing and covering with respect to an arbitrary monotone symmetric norm.

3.11 FPT Algorithms using Minimal Parameters for a Generalized Version of Maximin Shares

Alexandra Lassota (TU Eindhoven, NL)

License © Creative Commons BY 4.0 International license
© Alexandra Lassota

Joint work of Klaus Jansen, Alexandra Lassota, Malte Tutas, Adrian Vetta

Main reference Klaus Jansen, Alexandra Lassota, Malte Tutas, Adrian Vetta: “FPT Algorithms using Minimal Parameters for a Generalized Version of Maximin Shares”, CoRR, Vol. abs/2409.04225, 2024.

URL <https://doi.org/10.48550/ARXIV.2409.04225>

We study the computational complexity of fairly allocating indivisible, mixed-manna items. For basic measures of fairness, this problem is hard in general. The paradigm of fixed-parameter tractability (FPT) has led to new insights and improved algorithms for a variety of fair allocation problems. Our focus is designing FPT time algorithms for finding a best solution w.r.t. the fairness measure maximin shares (MMS). Furthermore, our techniques extend to finding allocations that optimize alternative objectives, such as minimizing the additive approximation, and maximizing some variants of global welfare. Our algorithms are actually designed for a more general MMS problem in machine scheduling. Here, each mixed-manna item (job) must be assigned to an agent (machine) and has a processing time and a deadline.

3.12 A Little Clairvoyance Is All You Need

Alexander Lindermayr (Universität Bremen, DE)

License © Creative Commons BY 4.0 International license
© Alexander Lindermayr

Joint work of Anupam Gupta, Haim Kaplan, Alexander Lindermayr, Jens Schlöter, Sorrachai Yingchareonthawornchai

We revisit the classical problem of minimizing the total *flow time* of jobs on a single machine in the online setting where jobs arrive over time. It has long been known that the Shortest Remaining Processing Time (SRPT) algorithm is optimal (i.e., 1-competitive) when the job sizes are known up-front [Schrage, 1968]. But in the non-clairvoyant setting where job sizes are revealed only when the job finishes, no algorithm can be constant-competitive [Motwani, Phillips, and Torng, 1994].

We consider the ε -clairvoyant setting, where $\varepsilon \in [0, 1]$, and each job’s processing time becomes known once its remaining processing time equals an ε fraction of its processing time. This captures settings where the system user uses the initial $(1 - \varepsilon)$ fraction of a job’s processing time to learn its true length, which it can then reveal to the algorithm. The model was proposed by Yingchareonthawornchai and Torng (2017), and it smoothly interpolates between the clairvoyant setting (when $\varepsilon = 1$) and the non-clairvoyant setting (when $\varepsilon = 0$). In a concrete sense, we are asking: *how much knowledge is required to circumvent the hardness of this problem?*

We show that *a little knowledge is enough*, and that a constant competitive algorithm exists for every constant $\varepsilon > 0$. More precisely, for all $\varepsilon \in (0, 1)$, we present an deterministic $\lceil 1/\varepsilon \rceil$ -competitive algorithm, which is optimal for deterministic algorithms. We also present a matching lower bound (up to a constant factor) for randomized algorithms.

Our algorithm to achieve this bound is remarkably simple and applies the “optimism in the face of uncertainty” principle. For each job, we form an optimistic estimate of its length, based on the information revealed thus far and run SRPT on these optimistic estimates. The

proof relies on maintaining a matching between the jobs in OPT’s queue and the algorithm’s queue, with small prefix expansion. We achieve this by carefully choosing a set of jobs *to arrive earlier than their release times* without changing the algorithm, and possibly helping the adversary. These early arrivals allow us to maintain structural properties inductively, giving us the tight guarantee.

3.13 The Power of Proportional Fairness and Unifying Scheduling Algorithms for Group Completion Times

Nicole Megow (*Universität Bremen, DE*)

License © Creative Commons BY 4.0 International license
© Nicole Megow

Joint work of Sven Jäger, Alexander Lindermayr, Zhenwei Liu, Nicole Megow

Main reference Emily Diana, Saeed Sharifi-Malvajerdi, Ali Vakilian: “Minimax Group Fairness in Strategic Classification”, in Proc. of the IEEE Conference on Secure and Trustworthy Machine Learning, SaTML 2025, Copenhagen, Denmark, April 9-11, 2025, pp. 753–772, IEEE, 2025.

URL <https://doi.org/10.1109/SaTML64287.2025.00047>

Main reference Sven Jäger, Alexander Lindermayr, Nicole Megow: “The Power of Proportional Fairness for Non-Clairvoyant Scheduling under Polyhedral Constraints”, in Proc. of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2025, New Orleans, LA, USA, January 12-15, 2025, pp. 3901–3930, SIAM, 2025.

URL <https://doi.org/10.1137/1.9781611978322.132>

We propose new abstract problems that unify a collection of scheduling and graph coloring problems with general min-sum objectives. Specifically, we consider the weighted sum of completion times over groups of entities (jobs, vertices, or edges), which generalizes two important objectives in scheduling: makespan and sum of weighted completion times.

We study these problems in both online and offline settings. In the non-clairvoyant online setting, we give a novel $O(\log g)$ -competitive algorithm, where g is the size of the largest group. This is the first non-trivial competitive bound for many problems with group completion time objective, and it is an exponential improvement over previous results for non-clairvoyant coflow scheduling. Notably, this bound is asymptotically best-possible. For offline scheduling, we provide powerful meta-frameworks that lead to new or stronger approximation algorithms for our new abstract problems and for previously well-studied special cases. In particular, we improve the approximation ratio from 13.5 to 10.874 for non-preemptive related machine scheduling and from $4 + \varepsilon$ to $2 + \varepsilon$ for preemptive unrelated machine scheduling (MOR 2012), and we improve the approximation ratio for sum coloring problems from 10.874 to 5.437 for perfect graphs and from 11.273 to 10.874 for interval graphs (TALG 2008).

3.14 Minimum Cost Adaptive Submodular Cover

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

License © Creative Commons BY 4.0 International license
© Viswanath Nagarajan

Joint work of Hessa Al-Thani, Yubing Cui, Blake Harris, Viswanath Nagarajan

Adaptive submodularity is a fundamental concept in stochastic optimization, with numerous applications such as sensor placement, hypothesis identification and viral marketing. We consider the problem of covering an adaptive-submodular function at minimum expected cost, which generalizes the classic set cover and submodular cover problems to the stochastic setting.

We show that the natural greedy policy has an approximation ratio of $4 \cdot (1 + \ln Q)$, where Q is the goal value. In fact, we consider a significantly more general objective of minimizing the p^{th} moment of the coverage cost, and show that the greedy policy *simultaneously* achieves a $(p + 1)^{p+1} \cdot (\ln Q + 1)^p$ approximation guarantee for all $p \geq 1$. All our approximation ratios are best possible up to constant factors (assuming $P \neq NP$). We also show that the greedy policy for minimizing expected cost has an approximation ratio at least $1.3 \cdot (1 + \ln Q)$ even when $Q = 1$, which invalidates a prior result on adaptive submodular cover. Moreover, our results extend to the setting where one wants to cover *multiple* adaptive-submodular functions, for which we obtain the same approximation guarantees.

3.15 Near-Optimal PCM Wear-Leveling Under Adversarial Attacks

Seffi Naor (Technion – Israel Institute of Technology – Haifa, IL)

License © Creative Commons BY 4.0 International license
© Seffi Naor

Joint work of Tomer Lange, Seffi Naor, Gala Yadgar

Main reference Tomer Lange, Joseph (Seffi) Naor, Gala Yadgar: “SSD Wear Leveling with Optimal Guarantees”, in Proc. of the 2024 Symposium on Simplicity in Algorithms, SOSA 2024, Alexandria, VA, USA, January 8-10, 2024, pp. 306–320, SIAM, 2024.

URL <https://doi.org/10.1137/1.9781611977936.28>

Phase change memory (PCM) is a promising memory technology known for its speed, high density, and durability. However, each PCM cell can endure only a limited number of erase and subsequent write operations before failing, and the failure of a single cell can limit the lifespan of the entire device. This vulnerability makes PCM particularly susceptible to adversarial attacks that induce excessive writes to accelerate device failure. To counter this, wear-leveling techniques aim to distribute write operations evenly across PCM cells.

In this paper we study the *online PCM utilization problem*, which seeks to maximize the number of write requests served before any cell reaches the erase limit. While extensively studied in the systems and architecture communities, this problem remains largely unexplored from a theoretical perspective. We bridge this gap by presenting a novel algorithm that leverages cell wear information to optimize PCM utilization. We prove that our algorithm achieves near-optimal worst-case guarantees and outperforms state-of-the-art practical solutions both theoretically and empirically, providing an efficient approach to prolonging PCM lifespan.

3.16 Robust Gittins for Stochastic Scheduling

Heather Newman (Carnegie Mellon University – Pittsburgh, US)

License © Creative Commons BY 4.0 International license
© Heather Newman

Joint work of Benjamin Moseley, Heather Newman, Kirk Pruhs, Rudy Zhou

Main reference Benjamin Moseley, Heather Newman, Kirk Pruhs, Rudy Zhou: “Robust Gittins for Stochastic Scheduling”, in Proc. of the Abstracts of the 2025 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, SIGMETRICS 2025, Stony Brook, NY, USA, June 9-13, 2025, pp. 166–168, ACM, 2025.

URL <https://doi.org/10.1145/3726854.3727315>

A common theme in stochastic optimization problems is that, theoretically, stochastic algorithms need to “know” relatively rich information about the underlying distributions. This is at odds with most applications, where distributions are rough predictions based on historical data. Thus, commonly, stochastic algorithms are making decisions using imperfect predicted distributions, while trying to optimize over some unknown true distributions.

We consider the fundamental problem of scheduling stochastic jobs preemptively on a single machine to minimize expected mean completion time in the setting where *the scheduler is only given imperfect predicted job size distributions*. If the predicted distributions are perfect, then it is known that this problem can be solved optimally by the Gittins index policy.

The goal of our work is to design a scheduling policy that is robust in the sense that it produces nearly optimal schedules even if there are modest discrepancies between the predicted distributions and the underlying real distributions. Our main contributions are:

- We show that the standard Gittins index policy is *not robust* in this sense. If the true distributions are perturbed by even an arbitrarily small amount, then running the Gittins index policy using the perturbed distributions can lead to an unbounded increase in mean completion time.
- We explain how to modify the Gittins index policy to make it robust, that is, to produce nearly optimal schedules, where the approximation depends on a new measure of error between the true and predicted distributions that we define.

Looking forward, the approach we develop here can be applied more broadly to many other stochastic optimization problems to better understand the impact of mispredictions, and lead to the development of new algorithms that are robust against such mispredictions.

3.17 Fair Caching

Debmalya Panigrahi (Duke University – Durham, US)

License © Creative Commons BY 4.0 International license
© Debmalya Panigrahi

Joint work of Anupam Gupta, Amit Kumar, Debmalya Panigrahi

Online convex paging models a broad class of cost functions for the classical paging problem. In particular, it naturally captures fairness constraints: e.g., that no specific page (or groups of pages) suffers an unfairly high number of evictions by considering ℓ_p norms of eviction vectors for $p > 1$. The case of the ℓ_∞ norm has also been of special interest, and is called min-max paging.

In this talk, I will discuss tight upper and lower bounds for the convex paging problem for a broad class of convex functions. Prior to our work, only fractional algorithms were known for this general setting. Moreover, our new results settle the competitive ratio for min-max paging and ℓ_p -norm paging for all values of $p > 1$.

3.18 The Santa Claus Problem – Three Perspectives

Lars Rohwedder (Maastricht University, NL)

License  Creative Commons BY 4.0 International license
© Lars Rohwedder

Santa Claus cannot accept that even a single child is unhappy on Christmas. Therefore, when he distributes his gifts, he maximizes the total value of gifts that the least happy child gets. This is a non-trivial task, especially when each gift j has a different value v_{ij} for each child i . This very natural problem, sometimes under the more serious name of max-min fair allocation, has seen significant attention in the last two decades. Yet, many questions about it remain widely open. We will survey developments on the problem using three different perspectives that demonstrate its versatile nature: First, we view it as a fair allocation problem, then as a scheduling problem, and finally as a network design problem.

3.19 Optimal Online Discrepancy Minimization

Thomas Rothvoss (University of Washington – Seattle, US)

License  Creative Commons BY 4.0 International license
© Thomas Rothvoss

Joint work of Thomas Rothvoss, Janardhan Kulkarni, Victor Reis
Main reference Janardhan Kulkarni, Victor Reis, Thomas Rothvoss: “Optimal Online Discrepancy Minimization”, CoRR, Vol. abs/2308.01406, 2023.
URL <https://doi.org/10.48550/ARXIV.2308.01406>

We prove that there exists an online algorithm that for any sequence of vectors $v_1, \dots, v_T \in \mathbb{R}^n$ with $\|v_i\|_2 \leq 1$, arriving one at a time, decides random signs $x_1, \dots, x_T \in \{-1, 1\}$ so that for every $t \leq T$, the prefix sum $\sum_{i=1}^t x_i v_i$ is $O(1)$ -subgaussian. This improves over the work of Alweiss, Liu and Sawhney who kept prefix sums $O(\sqrt{\log(nT)})$ -subgaussian. Our proof combines a generalization of Banaszczyk’s prefix balancing result to trees with a cloning argument to find distributions rather than single colorings.

3.20 Stochastic scheduling with Bernoulli-type jobs through policy stratification

Kevin Schewior (Universität zu Köln, DE)

License  Creative Commons BY 4.0 International license
© Kevin Schewior

Joint work of Antonios Antoniadis, Ruben Hoeksma, Kevin Schewior, Marc Uetz

This paper addresses the problem of computing a scheduling policy that minimizes the total expected completion time of a set of N jobs with stochastic processing times on m parallel identical machines. When all processing times follow Bernoulli-type distributions, Gupta et al. (SODA ’23) exhibited approximation algorithms with an approximation guarantee $\tilde{O}(\sqrt{m})$, where m is the number of machines and $\tilde{O}(\cdot)$ suppresses polylogarithmic factors in N , improving upon an earlier $O(m)$ approximation by Eberle et al. (OR Letters ’19) for a special case. The present paper shows that, quite unexpectedly, the problem with Bernoulli-type jobs admits a PTAS whenever the number of different job-size parameters is bounded by a constant. The result is based on a series of transformations of an optimal

scheduling policy to a “stratified” policy that makes scheduling decisions at specific points in time only, while losing only a negligible factor in expected cost. An optimal stratified policy is computed using dynamic programming. Two technical issues are solved, namely (i) to ensure that, with at most a slight delay, the stratified policy has an information advantage over the optimal policy, allowing it to simulate its decisions, and (ii) to ensure that the delays do not accumulate, thus solving the trade-off between the complexity of the scheduling policy and its expected cost. Our results also imply a quasi-polynomial $O(\log N)$ -approximation for the case with an arbitrary number of job sizes.

3.21 Proportional Representation in Budget Allocation

Ulrike Schmidt-Kraepelin (TU Eindhoven, NL)


License  Creative Commons BY 4.0 International license
© Ulrike Schmidt-Kraepelin

The ideal of proportional representation in social choice theory is easy to state yet challenging to formalize: any α -fraction of the population should have a say in determining an α -fraction of the outcome. This principle has gained significant attention in recent years and is arguably the most studied fairness notion in social choice theory today.

This talk explores proportional representation in the context of budget allocation—a broad framework capturing various models with wide-ranging applications, including apportionment, participatory budgeting, and committee elections. We will examine several formalizations of proportionality, introduce algorithms designed to achieve proportional outcomes, and highlight key open questions in the field. Beyond that, I hope to inspire discussion on how proportional representation might be relevant in settings beyond social choice theory.

3.22 A new deterministic approximation for graph burning

Jiri Sgall (Charles University – Prague, CZ)

License  Creative Commons BY 4.0 International license
© Jiri Sgall
Joint work of Matej Lieskovsky, Jiri Sgall

Graph Burning models information spreading in a given graph as a process such that in each step one node is infected (informed) and also the infection spreads to all neighbors of previously infected nodes. Formally, given a graph $G = (V, E)$, possibly with edge lengths, the burning number $b(G)$ is the minimum number g such that there exist nodes $v_0, \dots, v_{g-1} \in V$ satisfying the property that for each $u \in V$ there exists $i \in \{0, \dots, g-1\}$ so that the distance between u and v_i is at most i .

We present a simple deterministic 2.314-approximation algorithm for computing the burning number of a general graph, even with arbitrary edge lengths. This complements our previous more complicated randomized algorithm with the same approximation ratio.

3.23 A Simple Algorithm for Dynamic Carpooling with Recourse

Cliff Stein (Columbia University – New York, US)

License  Creative Commons BY 4.0 International license
© Cliff Stein

Joint work of Cliff Stein, Shyamal Patel, Yuval Efron

Main reference Yuval Efron, Shyamal Patel, Cliff Stein: “A Simple Algorithm for Dynamic Carpooling with Recourse”, in Proc. of the 2025 Symposium on Simplicity in Algorithms, SOSA 2025, New Orleans, LA, USA, January 13-15, 2025, pp. 196–201, SIAM, 2025.


URL <https://doi.org/10.1137/1.9781611978315.15>

We give an algorithm for the fully-dynamic carpooling problem with recourse: Edges arrive and depart online from a graph G with n nodes according to an adaptive adversary. Our goal is to maintain an orientation H of G that keeps the discrepancy, defined as $\max_{v \in V} |\deg_H^+(v) - \deg_H^-(v)|$, small at all times.

We present a simple algorithm and analysis for this problem with recourse based on cycles that simplifies and improves on a result of Gupta et al. [SODA ’22].

3.24 Six Candidates Suffice to Win a Voter Majority

Adrian Vetta (McGill University – Montreal, CA)

License  Creative Commons BY 4.0 International license
© Adrian Vetta

Joint work of Moses Charikar, Alexandra Lassota, Prasanna Ramakrishnan, Adrian Vetta, Kangning Wang

Main reference Moses Charikar, Alexandra Lassota, Prasanna Ramakrishnan, Adrian Vetta, Kangning Wang: “Six Candidates Suffice to Win a Voter Majority”, CoRR, Vol. abs/2411.03390, 2024.

URL <https://doi.org/10.48550/ARXIV.2411.03390>

Given an election of n voters with preference lists over m candidates, Elkind, Lang, and Saffidine (2011) defined a Condorcet winning set to be a collection of candidates that the majority of voters prefer over any individual candidate. Condorcet winning sets of cardinality one (a Condorcet winner) or cardinality two need not exist. We prove however that a Condorcet winning set of cardinality at most six exists in any election.

3.25 The Power of Migrations in Dynamic Bin Packing

Rudy Zhou (Carnegie Mellon University – Pittsburgh, US)

License  Creative Commons BY 4.0 International license
© Rudy Zhou

Joint work of Konstantina Mellou, Marco Molinaro, Rudy Zhou

Main reference Konstantina Mellou, Marco Molinaro, Rudy Zhou: “The Power of Migrations in Dynamic Bin Packing”, Proc. ACM Meas. Anal. Comput. Syst., Vol. 8(3), pp. 45:1–45:28, 2024.

URL <https://doi.org/10.1145/3700435>

In the *Dynamic Bin Packing* problem, n items arrive and depart the system in an online manner, and the goal is to maintain a good packing throughout. We consider the objective of minimizing the total active time, i.e., the sum of the number of open bins over all times. An important tool for maintaining an efficient packing in many applications is the use of *migrations*; e.g., transferring computing jobs across different machines. However, there are large gaps in our understanding of the approximability of dynamic bin packing with migrations. Prior work has covered the power of no migrations and $> n$ migrations, but we ask the question: What is the power of limited ($\leq n$) migrations?

Our first result is a dichotomy between no migrations and linear migrations: Using a sublinear number of migrations is asymptotically equivalent to doing *zero* migrations, where the competitive ratio grows with μ , the ratio of the largest to smallest item duration. On the other hand, we prove that for every $\alpha \in (0, 1]$, there is an algorithm that does $\approx \alpha n$ migrations and achieves competitive ratio $\approx 1/\alpha$ (in particular, independent of μ); we also show that this tradeoff is essentially best possible. This fills in the gap between zero migrations and $> n$ migrations in Dynamic Bin Packing.

Finally, in light of the above impossibility results, we introduce a new model that more directly captures the impact of migrations. Instead of limiting the number of migrations, each migration adds a delay of C time units to the item's duration; this commonly appears in settings where a blackout or set-up time is required before the item can restart its execution in the new bin. In this new model, we prove a $O(\min(\sqrt{C}, \mu))$ -approximation, and an almost matching lower bound. We also present preliminary experiments that indicate that our theoretical results are predictive of the practical performance of our algorithms.

Participants

- Antonios Antoniadis
University of Twente, NL
- Yossi Azar
Tel Aviv University, IL
- Eric Balkanski
Columbia University –
New York, US
- Etienne Bamas
ETH Zürich, CH
- Sanjoy Baruah
Washington University –
St. Louis, US
- Emily Diana
TTIC – Chicago, US
- Franziska Eberle
TU Berlin, DE
- Yuri Faenza
Columbia University –
New York, US
- Naveen Garg
Indian Institute of Technology –
New Delhi, IN
- Swati Gupta
MIT – Cambridge, US
- Sungjin Im
University of California –
Santa Cruz, US
- Thomas Kesselheim
Universität Bonn, DE
- Samir Khuller
Northwestern University –
Evanston, US
- Alexandra Lassota
TU Eindhoven, NL
- Alexander Lindermayr
Universität Bremen, DE
- Alberto Marchetti-Spaccamela
Sapienza University of Rome, IT
- Claire Mathieu
CNRS – Paris, FR
- Nicole Megow
Universität Bremen, DE
- Benjamin J. Moseley
Carnegie Mellon University –
Pittsburgh, US
- Viswanath Nagarajan
University of Michigan –
Ann Arbor, US
- Seffi Naor
Technion – Haifa, IL
- Heather Newman
Carnegie Mellon University –
Pittsburgh, US
- Debmalya Panigrahi
Duke University – Durham, US
- Kirk Pruhs
University of Pittsburgh, US
- Lars Rohwedder
Maastricht University, NL
- Thomas Rothvoss
University of Washington –
Seattle, US
- Kevin Schewior
Universität Köln, DE
- Ulrike Schmidt-Kraepelin
TU Eindhoven, NL
- Jiri Sgall
Charles University – Prague, CZ
- David Shmoys
Cornell University – Ithaca, US
- Martin Skutella
TU Berlin, DE
- Frits C. R. Spieksma
TU Eindhoven, NL
- Clifford Stein
Columbia University –
New York, US
- Leen Stougie
CWI – Amsterdam, NL
- Ola Svensson
EPFL – Lausanne, CH
- Kavitha Teliakapalli
TIFR Mumbai, IN
- Marc Uetz
University of Twente –
Enschede, NL
- Adrian Vetta
McGill University –
Montreal, CA
- Tjark Vredeveld
Maastricht Univ. School of
Business & Economics, NL
- Andreas Wiese
TU München, DE
- Hang Zhou
Ecole Polytechnique –
Palaiseau, FR
- Rudy Zhou
Carnegie Mellon University –
Pittsburgh, US

