Schedule

Dagstuhl Seminar

Scheduling & Fairness 2025

Monday

09:00 – 09:45	Welcome and Round of Introduction
09:45 – 10:15	Emily Diana: Minimax Group Fairness in Strategic Classification
10:15 – 10:45	Coffee break
10:45 – 11:45	Adrian Vetta: Six Candidates Suffice to Win a Voter Majority
12:15 – 13:30	Lunch
13:30 – 15:30	Collaboration Time
15:30 – 16:00	Coffee & Cake
16:00 – 16:30	Heather Newman: Robust Gittins for Stochastic Scheduling
16:30 – 17:00	Alexander Lindermayr: A Little Clairvoyance Is All You Need for Minimizing Total Flow Time
17:00 – 17:30	Cliff Stein: A Simple Algorithm for Dynamic Carpooling with Recourse
18:00 – 19:30	Dinner

Tuesday

09:00 – 10:00	Gupta, Swati: Fair Resource Allocation from Theory to Practice
10:00 – 10:30	Coffee break
10:30 – 11:00	Alexandra Lassota: FPT Algorithms using Minimal Parameters for a Generalized Version of Maximin Shares
11:00 – 11:30	Thomas Rothvoss: Optimal Online Discrepancy Minimization
11:30 – 12:00	Franziska Eberle: A Tight ($3/2 + \eps$)-Approximation Algorithm for Demand Strip Packing
12:15 – 13:30	Lunch
13:30 – 15:30	Collaboration Time
15:30 – 16:00	Coffee & Cake
16:00 – 16:30	Short Talks: Samir Khuller, Jiri Sgall, Seffi Naor
16:30 – 17:00	Kevin Schewior: A PTAS for stochastic parallel machine scheduling through policy stratification
17:00 – 18:00	**Open Problems** (Yuri, Leen, Swati)
18:00 – 19:30	Dinner

Wednesday

09:00 – 10:00	Lars Rohwedder: The Santa Claus Problem: Three Perspectives
10:00 – 10:30	Coffee break
10:30 – 11:00	Etienne Bamas: Lift-and-Project Integrality Gaps for Santa Claus
11:00 – 11:30	Yuri Faenza: Students in highly competitive markets: the case of New York City specialized high schools
11:30 – 12:00	Yossi Azar: Lossless Robustification of Packet Scheduling Algorithms
12:15 – 13:30	Lunch
meet 13:30	Hike tour or visit to Trier – free afternoon
15:30 – 16:00	Coffee & Cake
18:00 – 19:30	Dinner

Thursday

09:00 – 10:00	Telikepalli, Kavitha: Fair solutions to the house allocation problem
10:00 – 10:30	Coffee break
10:30 – 11:00	Sungjin Im: Online Scheduling via Gradient Descent
11:00 – 11:30	Debmalya Panigrahi: Learning-Augmented Algorithms via Convex Programming Duality.
11:30 – 12:00	Eric Balkanski: Fair Strategic Facility Location with Predictions
12:15 – 13:30	Lunch
13:30 – 15:30	Collaboration Time
15:30 – 16:00	Coffee & Cake
16:00 – 16:30	Thomas Kesselhein: Supermodular Approximation of Norms and Applications
16:30 – 17:00	Rudy Zhou: The Power of Migrations in Dynamic Bin Packing
17:00 – 17:30	Short talks: Naveen Garg, David Shmoys, Nicole Megow, Claire Mathieu
18:00 – 19:30	Dinner

Friday

09:00 – 10:00	Ulrike Schmidt-Kraepelin: Proportional Representation in Budget Allocation
10:00 – 10:30	Coffee break
10:30 – 11:00	Viswanath Nagarajan: Minimum Cost Adaptive Submodular Cover
11:00 – 12:00	Collaboration Time
12:15 – 13:30	Lunch

Emily Diana: Minimax Group Fairness in Strategic Classification

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. This paper will appear in the 2025 IEEE conference on Secure and Trustworthy Machine Learning, and a pre-print is available at https://arxiv.org/pdf/2410.02513 .

Adrian Vetta: Six Candidates Suffice to Win a Voter Majority

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

Heather Newman: Robust Gittins for Stochastic Scheduling

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 response 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 response 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.

Alexander Lindermayer: A Little Clairvoyance Is All You Need for Minimizing Total Flow Time

Abstract: We study the problem of preemptively minimizing the total flow time on a single machine where jobs arrive online over time. When processing times are known at arrival (called clairvoyant scheduling), scheduling at any time the job with the shortest remaining processing time is optimal. However, in non-clairvoyant scheduling, where the processing time of a job is only revealed at its completion, strong lower bounds rule out constant competitive algorithms. Therefore, this problem has been very successfully studied in models with resource augmentation or with additional information available upon the arrival of a job. We continue this line of research and consider a first model that captures practical scenarios such as profiling techniques that model a learning process while we process jobs. Specifically, for any \alpha between 0 and 1, an \alpha-clairvoyant algorithm learns about a job's processing time after it finished an \alpha-fraction of the job. We present a natural algorithm based on fair allocation and the optimism principle with a constant competitive ratio of ceiling(1/(1-\alpha). Moreover, we show that this ratio is optimal for deterministic \alpha-clairvoyant algorithms.

Joint work with Anupam Gupta, Haim Kaplan, Jens Schlöter, and Sorrachai Yingchareonthawornchai.

Swati Gupta: Fair Resource Allocation from Theory to Practice

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

Cliff Stein: A Simple Algorithm for Dynamic Carpooling with Recourse

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].

Joint work with Yuval Efron, Shyamal Patel.

Alexandra Lassota: FPT Algorithms using Minimal Parameters for a Generalized Version of Maximin Shares

Abstract: 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.

Thomas Rothvoss: Optimal Online Discrepancy Minimization

We prove that there exists an online algorithm that for any sequence of vectors $v_1,\ldots,v_T \in \setR^n$ with $\|v_i\|_2 \leq 1$, arriving one at a time, decides random signs $x_1,\ldots,x_T \in \{ -1,1\}$ so that for every $t \le T$, 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.

This is joint work with Janardhan Kulkarni and Victor Reis.

Franziska Eberle: A Tight ($3/2 + \eps$)-Approximation Algorithm for Demand Strip Packing

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 $\big(\frac32+\varepsilon\big)\textsc{Opt}$ satisfying one of two properties. Either (i) the solution leaves a gap for a job with demand $\textsc{Opt}$ and processing time $\mathcal O(\varepsilon D)$ or (ii) all jobs with demand greater than $\frac{\textsc{Opt}}2$ appear sorted by demand in immediate succession. We then provide two efficient algorithms that find a solution with maximum demand at most $\big(\frac32+\varepsilon\big)\textsc{Opt}$ 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.
This is joint work with Felix Hommelsheim, Malin Rau, and Stefan Walzer

Kevin Schewior: A PTAS for stochastic parallel machine scheduling through policy stratification.

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.

Lars Rohwedder: **The Santa Claus Problem: Three Perspectives**

Abstract: 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 vij 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.

Etienne Bamas: Lift-and-Project Integrality Gaps for Santa Claus

In this talk, I will focus on the MaxMinDegree Arborescence (MMDA) problem in layered directed graphs of depth \ell ≤ 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.

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

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. Joint work with Swati Gupta (MIT) and Xuan Zhang (Meta).

Yossi Azar: Lossless Robustification of Packet Scheduling Algorithms

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 a 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.

Kavitha Telikepalli: Fair solutions to the house allocation problem

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.

Sungjin Im: Online Scheduling via Gradient Descent

Abstract: 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.

Debmalya Panigrahi: Learning-Augmented Algorithms via Convex Programming Duality

Abstract: In the last Scheduling seminar at Dagstuhl, I presented a general framework for learning-augmented assignment problems such as online job scheduling, and mused that there might be some duality at play that was hiding underneath the surface. I will talk about some very recent work that in a sense follows up on that comment - we can now use an explicit convex programming duality to simplify, strengthen, and generalize to the online covering problem with applications such as fair routing on a network that were beyond the scope of the previous techniques. This is joint work with Ilan Cohen.

Eric Balkanski: Fair Strategic Facility Location with Predictions

Abstract: 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.

Thomas Kesselheim: Supermodular Approximation of Norms and Applications

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 $O(\poly\log m)$-competitive algorithms for load balancing and covering with respect to an arbitrary monotone symmetric norm.

Based on joint work with Marco Molinaro and Sahil Singla

Rudy Zhou: The Power of Migrations in Dynamic Bin Packing

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. However, prior work on Dynamic Bin Packing with migrations has focused on the regime where we can do many (> n) migrations. This is prohibitively many migrations in applications, so we are interested in algorithms that use migrations more judiciously. We give almost-tight bounds on algorithms that do < n migrations.

Further, we introduce a new model which more directly captures the cost of migrations: each time we migrate an item incurs an additive delay on the duration of that item, which occurs in systems applications, where migrating a job often requires data transfer and blackout times. We give tight bounds in this new model as well.

Finally, We also present preliminary experiments on real Microsoft Azure workloads which indicate that our theoretical results are predictive of practical performance of our algorithms. This is joint work with Konstantina Mellou and Marco Molinaro, and will appear in Sigmetrics 2025 (https://arxiv.org/abs/2408.13178v1).

Ulrike Schmidt-Kraepelin: Proportional Representation in Budget Allocation

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, committee elections, and donor coordination. 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.

Viswanath Nagarajan: Minimum Cost Adaptive Submodular Cover

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 minimum cost cover of adaptive-submodular functions, and provide a 4(1 + lnQ)-approximation algorithm, where Q is the goal value. In fact, we consider a significantly more general objective of minimizing the pth moment of the coverage cost, and show that our algorithm simultaneously achieves a (p + 1)^(p+1) (1+ln Q)^p approximation guarantee for all p ≥ 1. Moreover, our results also extend to the setting where one wants to cover multiple adaptive-submodular functions