13th Day on Computational Game Theory

Vrije Universiteit Amsterdam, 15 & 16 June 2023

Talks will take place in HG-12A00 on Thursday 15 June, and in HG-02A00 on Friday 16 June.

Thursday, 15 June, 13.40 – 14.25

Keynote talk by Neil Olver: *Understanding equilibria in flow-over-time traffic models*

Network congestion games are by now a standard, extensively studied topic in algorithmic game theory. Dynamic models of transport and network traffic, where traffic flows and congestion vary over time, bring with them many new challenges. This talk will give an overview of the deterministic queueing model, a fundamental model of this type with roots that go back to Ford and Fulkerson (from the perspective of network flows) and Vickrey (from the perspective of transportation economics). I will survey both older and very recent results about dynamic equilibria, which exhibit surprisingly rich behaviour and are still quite poorly understood.

Thursday, 15 June, 14.40 – 15.40

Lukas Graf: *Side-Constrained Dynamic Traffic Equilibria*

We study dynamic traffic assignment with side-constraints. We first give a counter-example to a key result from the literature regarding the existence of dynamic equilibria for volume-constrained traffic models in the classical edge-delay model. Our counter-example shows that the feasible flow space need not be convex and it further reveals that classical variational inequalities are not suited for the definition of side-constrained dynamic equilibria

We then propose a new framework for side-constrained dynamic equilibria based on the concept of feasible ε-deviations of flow particles in space and time. Under natural assumptions, we characterize the resulting equilibria by means of quasi-variational and variational inequalities, respectively. Finally, we establish first existence results for side-constrained dynamic equilibria for the non-convex setting of volume-constraints.

This is joint work with Tobias Harks.

Wouter Fokkema: *The Price of Anarchy for Matroid Congestion Games*

We consider atomic, symmetric congestion games with linear cost functions, for which it is known that the price of anarchy is at most (5n-2)/(2n+1). It was open if restricting the players strategy spaces to the bases of a matroid allows for improved upper bounds on the price of anarchy. We essentially answer this question negatively. Specifically, we consider the case in which players select edges of a graph to obtain a spanning tree. We show that the price of anarchy of this game equals the known upper bound (5n-2)/(2n+1) when the number of players n is equal to 2, 3, 4, or when n tends to infinity. That means that, for these values of n, restricting strategy spaces to bases of a matroid does not impact the price of anarchy of atomic, linear congestion games. It is conceivable that the same is true also for other values of n.

Svenja M. Griesbach: *Information Design for Congestion Games with Unknown Demand*

This paper considers non-atomic congestion games with affine costs and a single commodity whose demand is unknown to the players of the game, but can be observed by a principal. Prior to observing the demand, the principal commits to a public signaling scheme that governs to which extend (if at all) the information about the realized demand is communicated to the players. Upon receiving the signal, the players update their beliefs about the demand and respond by playing a Wardrop equilibrium for the resulting cost functions. We study the question of how to compute the signaling scheme that minimizes the total cost of the induced Wardrop equilibrium.

First, we devise a fully polynomial-time approximation scheme (FPTAS) for the case that the demand can only take two values. The FPTAS relies on several structural properties of the cost of the induced Wardrop equilibrium as a function of the updated belief of the distribution of demands. We show that this function is piecewise linear for any number of demands, and monotonic for two demands.

Second, we give a complete characterization of the graph structures where it is optimal for the principal to fully reveal the information about the realized demand to the players. Specifically, we show that this signaling scheme is optimal for all cost functions and probability distributions over demands if and only if the graph is series-parallel. Third, we propose an algorithm that computes the optimal signaling scheme for any number of demands whose time complexity is polynomial in the number of supports that occur in a Wardrop equilibrium for some demand. Finally, we conduct a computational study that tests this algorithm on real-world instances.

This is joint work with Martin Hoefer, Max Klimm, and Tim Koglin.

Thursday, 15 June, 16.15 – 17.15

Pieter Kleer: *Optimal Stopping Theory for a Distributionally Robust Seller*

We consider the classical problem of selling an item in an online fashion. Buyers arrive one-by-one and report a value, drawn from a known value distribution, that they are willing to pay for the item. Upon arrival of a buyer, we have to irrevocably decide whether or not to sell the item to them. The goal is to design a stopping rule, specifying which buyer to sell to, that maximizes the value at which the item is sold.

The assumption that the seller knows the full value distributions of the buyers is often unrealistic. Therefore, we consider a more robust version of the problem by taking a maximin perspective. We assume that the value distributions of the buyers come from a so-called ambiguity set. The seller first chooses a stopping rule, after which nature chooses value distributions from the ambiguity set that minimize the seller's expected payoff under the chosen stopping rule. The goal of the seller is now to choose a stopping rule that maximizes their expected maximin payoff, i.e., the value of the buyer to which the item is sold.

For various ambiguity sets described by mean and dispersion (variance or mean absolute deviation) information of the unknown distributions, we give closed-form expressions for the optimal robust threshold strategies, also generalizing work of Boshuizen and Hill (1992).

Joint work with Johan van Leeuwaarden (Tilburg University).

Elias Pitschmann: *Prophet Inequalities over Time*

We introduce an over-time variant of the well-known prophet inequality with i.i.d. random variables. Instead of stopping with one realized value at some point in the process, we decide for each step how long we select the value. Then we cannot select another value until this period is over. The goal is to maximize the expectation of the sum of selected values. We describe the structure of the optimal stopping rule and give upper and lower bounds on the prophet inequality. In online algorithms terminology, this corresponds to bounds on the competitive ratio of an online algorithm.

We give a surprisingly simple algorithm with a single threshold that results in a prophet inequality of approx. 0.396 for all input lengths n. Additionally, as our main result, we present a more advanced algorithm resulting in a prophet inequality of approx. 0.598 when the number of steps tends to infinity. We complement our results by an upper bound that shows that the best possible prophet inequality is at most the inverse of the golden ratio, which is approx. 0.618.

Anh Trieu: *Matching maximizing mechanism in a two-sided auction setting*

We consider a dominant strategy incentive compatible mechanism that maximizes the number of matchings in a two-sided auction setting. In this setting, m consumers have unit demand over identical indivisible items which are produced by n producers. Each consumer has their own valuation for the item and each producer has their own cost to produce the item. All valuations and costs are private information. We develop a matching rule that matches consumers to producers such that the number of matchings is maximized and for each matched pair, the consumer's valuation is higher or equal to the producer's cost. We also provide a payment rule that induces truthfulness. Then, we analyze the mechanism and compare its total number of matched pairs as well as its deficit to those of the VCG mechanism.

Thursday, 15 June, 17.30 – 18.30

Katharina Eickhoff: *Walrasian Prices – Computation and Properties*

Determining prices on a market where a seller aims to sell a finite set of indivisible distinguishable goods to a finite set of buyers is a well-studied problem with countless applications. Prices on the items are called Walrasian (equilibrium) prices if they admit an envy-free allocation of items to the buyers such that all items with positive prices are sold. Such Walrasian prices are guaranteed to exist if the buyers' valuations are gross substitute, a concept introduced by Kelso and Crawford.

A very natural process to determine prices of the goods is an ascending auction, where prices are raised step by step on overdemanded sets. Provided that all buyers have gross substitute valuations, Gul and Stacchetti have shown that an ascending auction which iteratively raises prices on inclusion-wise minimal maximal overdemanded sets terminates with the (unique) componentwise minimal Walrasian price vector. It is by now known that an inclusion-wise minimal maximal overdemanded set can be computed in polynomial time with tools from discrete convexity. We present a simple and purely combinatorial polynomial time algorithm for their computation: we show that the desired overdemanded sets can be computed using the exchange graph of the classical matroid partitioning algorithm.

Moreover, we show that for gross substitute valuations the component-wise minimal prices that admit an envy-free allocation are the same as the minimal Walrasian prices. This enables us to show a very natural monotonicity of minimal Walrasian prices with respect to changes in supply and demand.

Niklas Rieken: *Selling a Base of a Matroid*

We give a new presentation and simpler analysis of the clock auction of Bikhchandani, de Vries, Schummer, and Vohra [SODA'08] to sell a welfare-maximizing base of of a matroid at Vickrey prices. Since bidders can react to certain events in the auction, the VCG Theorem (and hence, truthfulness) does not apply immediately in contrast to a sealed-bid mechanism. The new analysis only requires a few matroid folklore theorems to establish all the good properties of the auction. Moreover, we can show some natural properties about market power of a bidder in simple matroid terms.

Joint work with Britta Peis.

Artem Tsikiridis: *Fractional Budget-feasible Mechanism Design*

We study procurement auction scenarios where bidders have a strict budget constraint. The task for the auctioneer is to purchase resources owned by rational bidders, while keeping the cost within a fixed, predetermined budget and maximizing a given valuation function. In this context, the cost of each resource is private, hence our objective is to design truthful mechanisms that provide a good approximation to the optimal value for the auctioneer and are budget-feasible, meaning that the sum of payments to bidders does not exceed the budget. This framework applies to many well-known optimization problems and is relevant in various application scenarios such as crowdsourcing platforms, where bidders are workers offering tasks, and influence maximization in social networks, where agents are influential users. We specifically concentrate on the environment in which agents are completely divisible or can be viewed as offering a fully divisible service to the auctioneer. We present a novel mechanism which achieves a better guarantee than the known lower bound for the indivisible setting. This separation result implies that fractional budget-feasible mechanisms can indeed achieve better approximations than their non-divisible counterparts.

Friday, 16 June, 09.00 – 09.45

Keynote talk by Edith Elkind: *Mind the gap: fair division with separation constraints*

Motivated by the social distancing rules, we consider the task of fairly sharing a divisible good among several agents in a setting where every two agents' shares must be separated. We first look at the case where the good is the [0, 1] segment (usually referred to as `cake'). In this model, the separation constraint is captured by specifying a parameter s such that for every pair of agents i, j their shares are separated by a segment of length at least s; intuitively, the cake is cut by a blunt knife of width s. We focus on the recently introduced fairness concept of maximin fair share, and show that each agent can be guaranteed her maximin fair share; however, computing the agents' fair shares is computationally hard. We then extend our analysis to richer models, such as a 2-dimensional cake (where we have additional restrictions on the shapes of agents' pieces) and graphical cake (where agents need to share edges of a graph), and obtain positive results for an ordinal relaxation of the maximin fair share solution concept.

Friday, 16 June, 10.00 – 11.00

Hannaneh Akrami: *Breaking the 3/4 Barrier for Approximate Maximin Share*

We study the fundamental problem of fairly allocating a set of indivisible goods among n agents with additive valuations using the desirable fairness notion of maximin share (MMS). MMS is the most popular share-based notion, in which an agent finds an allocation fair to her if she receives goods worth at least her MMS value. An allocation is called MMS if all agents find it fair.

However, since MMS allocations need not exist when $n>2$, a series of works have proved the existence of approximate MMS allocations with the current best factor of $3/4 + O(1/n)$. This factor has not improved much for more than half a decade due to the limitation of existing approaches. In this work, we prove the existence of $(3/4 + 3/4220)$-MMS allocations by developing new tools and techniques.

Javier Cembrano: *Improved Bounds for Single-Nomination Impartial Selection*

We give new bounds for the single-nomination model of impartial selection, a problem proposed by Holzman and Moulin (Econometrica, 2013). A selection mechanism, which may be randomized, selects one individual from a group of n based on nominations among members of the group; a mechanism is impartial if the selection of an individual is independent of nominations cast by that individual, and α-optimal if under any circumstance the expected number of nominations received by the selected individual is at least α times that received by any individual. In a many-nominations model, where individuals may cast an arbitrary number of nominations, the so-called permutation mechanism is 1/2-optimal, and this is best possible. In the single-nomination model, where each individual casts exactly one nomination, the permutation mechanism does better and prior to this work was known to be 67/108-optimal but no better than 2/3-optimal. We show that it is in fact 2/3-optimal for all n. This result is obtained via tight bounds on the performance of the mechanism for graphs with maximum degree Δ, for any Δ, which we prove using an adversarial argument. We then show that the permutation mechanism is not best possible; indeed, by combining the permutation mechanism, another mechanism called plurality with runner-up, and some new ideas, 2105/3147-optimality can be achieved for all n. We finally give new upper bounds on α for any α-optimal impartial mechanism. They improve on the existing upper bounds for all $n\geq 7$ and imply that no impartial mechanism can be better than 76/105-optimal for all n; they do not preclude the existence of a $(3/4-\varepsilon)$-optimal impartial mechanism for arbitrary $\varepsilon>0$ if n is large.

This is joint work with Felix Fischer and Max Klimm.

Konstantinos Varsos: *Forward Looking Best-Response Multiplicative Weights Update Methods for Bilinear Zero-sum Games*

In this work, we focus on extra gradient learning algorithms for finding Nash equilibria in bilinear zero-sum games. The proposed method, which can be formally considered as a variant of Optimistic Mirror Descent, uses a large learning rate for the intermediate gradient step which essentially leads to computing (approximate) best response strategies against the profile of the previous iteration. Although counter-intuitive at first sight due to the irrationally large, for an iterative algorithm, intermediate learning step, we prove that the method guarantees last-iterate convergence to an equilibrium. Particularly, we show that the algorithm reaches first an ϵ-approximate Nash equilibrium, by decreasing the Kullback-Leibler divergence of each iterate by a guaranteed size, until the method becomes a contracting map, and converges to the exact equilibrium. Furthermore, we provide experimental comparisons with the optimistic variant of the multiplicative weights update method and show that our algorithm has significant practical potential since it offers substantial gains in terms of accelerated convergence.

Friday, 16 June, 11.15 – 12.15

Gabriele Dragotto: *Integer Programming Games: Do We Really Need Them?*

This talk deals with Integer Programming Games (IPGs), i.e., simultaneous non-cooperative games where each player solves a parametrized mixed-integer optimization problem. The broad class of IPGs encompasses, among others, any finite game and provides a flexible framework to model games that cannot be efficiently represented compactly, for instance, in normal form. After providing some background motivation and context of applications, we introduce zero regrets, a general and efficient cutting-plane algorithm to compute, enumerate, and select Nash equilibria in a large class of IPGs.

Maximilian Stahlberg: *Complexity of equilibria in binary public goods games on undirected graphs*

We study the complexity of computing equilibria in binary public goods games on undirected graphs. In such a game, players correspond to vertices in a graph and face a binary choice of performing an action, or not. Each player's decision depends only on the number of neighbors in the graph who perform the action and is encoded by a per-player binary pattern. We show that games with decreasing patterns (where players only want to act up to a threshold number of adjacent players doing so) always have a pure Nash equilibrium and that one is reached from any starting profile by following a polynomially bounded sequence of best responses. For non-monotonic patterns where players want to act alone or alongside a fixed number of neighbors, we show that it is NP-hard to decide whether a pure Nash equilibrium exists.

Koosha Samieefar: *The Computational Complexity of Mixed Constrained Equilibria and Applications*

Extensive research has been conducted on the computational aspects of solution concepts such as Nash equilibrium. These studies have encompassed various scenarios, including those where the goal is to investigate the computational aspects of equilibria with additional properties. Furthermore, in order to address issues of tractability, attention has been given to approximate versions of these problems. Our work extends this direction by considering games with constraints in which players are subject to some form of restrictions on their strategic choices. We also delve into the relationship between Nash equilibria and constrained or social equilibria in this context, paying particular attention to their relationship in terms of totality and complexity. Our findings reveal that the computational complexity of finding an equilibrium varies significantly among games that feature slightly different strategic constraints. Apart from analyzing the computational aspects of these strategic constraints, we also demonstrate their application in modeling problems related to strategic resource allocation, while simultaneously addressing their relevance in behavioral game theory.

Friday, 16 June, 13.45 – 14.45

Niklas Hahn: *Online TSP with Predictions on Locations*

In the Online Traveling Salesperson Problem (OLTSP), a server has to serve requests which appear over time in a metric space, minimizing the total time spent. The server can move with unit speed through the space or wait at any point of it. We consider this problem augmented by (e.g. machine-learned) predictions on the locations of requests.

We propose an oracle-based algorithmic framework for this problem. Using it, we design online algorithms for general as well as specific metric spaces. Under perfect predictions, they improve the competitive ratio of the classical scenario. If the predictions are erroneous, their performance degrades gracefully based on the error. Nevertheless, our algorithms always achieve a ratio within a small constant factor compared to the best known competitive ratio in the classical scenario, regardless of the prediction error.

Golnoosh Shahkarami: *A Novel Prediction Setup for Online Speed-Scaling*

Given the rapid rise in energy demand by data centers and computing systems in general, it is fundamental to incorporate energy considerations when designing (scheduling) algorithms. Machine learning can be a useful approach in practice by predicting the future load of the system based on, for example, historical data. However, the effectiveness of such an approach highly depends on the quality of the predictions and can be quite far from optimal when predictions are sub-par. On the other hand, while providing a worst-case guarantee, classical online algorithms can be pessimistic for large classes of inputs arising in practice.

This paper, in the spirit of the new area of machine learning augmented algorithms, attempts to obtain the best of both worlds for the classical, deadline based, online speed-scaling problem: Based on the introduction of a novel prediction setup, we develop algorithms that (i) obtain provably low energy-consumption in the presence of adequate predictions, and (ii) are robust against inadequate predictions, and (iii) are smooth, i.e., their performance gradually degrades as the prediction error increases.

Sophie Rain: *Game Theory for Automated Verification of Protocols*

Game Theory is used in various different areas. However, only recently it found its way into the Blockchain Security community. In our work, we combine Game Theory, Automated Reasoning and Security aspects to formally prove a Blockchain protocol (game-theoretically) secure. This is necessary as even in the presence of ingenious cryptographic schemes certain unwanted behavior is possible. Such behavior is then punished, respectively the alternatives are incentivized. In our game-theoretic security framework a Blockchain protocol is translated to an Extensive Form Game (with symbolic terms as utilities), the game together with the intended course of action, the "honest" behavior, is then analyzed with respect to three sufficient game-theoretic security properties. The properties are variants of well-known game-theoretic properties such as Nash equilibria and subgame perfect strategies. Finally, everything is encoded as a first-order logic formula and handed to a so-called SMT solver.