Program

Tuesday, May 9, 2023

9:00 - 9:15 *Gathering, Registration*

9:15 - 9:30 **Opening Remarks**

9:30 - 10:00 **Sergio Rajsbaum -** *A Combinatorial Topology Approach to Arrow's Impossibility Theorem*

Abstract: Social choice theory is a long-standing discipline in the fields of philosophy and politics. However, social choice theory took off within mathematical economics in the middle of the XX century with the momentous formulation of Keneth Arrow's impossibility theorem, showing that some apparently natural and undemanding conditions relating social decisions to individual preferences in a democratic way cannot be simultaneously satisfied. We present a combinatorial topology perspective of this theorem, providing new insights useful to compute domain restrictions, inspired by distributed computing techniques. Joint work with Armajac Raventos presented in PODC 2022.

10:00 - 10:30 **Uri Meir -** *Resilience of 3-Majority Dynamics to Non-Uniform Schedulers*

Abstract: In recent years there has been great interest in networks of passive, computationally-weak nodes, whose interactions are controlled by the outside environment; examples include population protocols, chemical reactions networks (CRNs), DNA computing, and more. Such networks are usually studied under one of two extreme regimes: the schedule of interactions is either assumed to be adversarial, or it is assumed to be chosen uniformly at random. In this paper we study an intermediate regime, where the interaction at each step is chosen from some not-necessarily-uniform distribution: we introduce the definition of a (p, ε)-scheduler, where the distribution that the scheduler chooses at every round can be arbitrary, but it must have ℓ_p-distance at most ε from the uniform distribution. We ask how far from uniform we can get before the dynamics of the model break down.

In this talk, we focus on the 3-majority dynamics, a type of chemical reaction network where the nodes of the network interact in triplets. Each node initially has an opinion of either X or Y, and when a triplet of nodes interact, all three nodes change their opinion to the majority of their three opinions. It is known that under a uniformly random scheduler, if we have an initial gap of $\Omega\left(\sqrt{n \log n}\right)$ in favor of one value, then w.h.p. all nodes converge to the majority value within $O\left(n \log n\right)$ steps.

For the 3-majority dynamics, we prove that among all non-uniform schedulers with a given ℓ_1-distance or ℓ_∞-distance to the uniform scheduler, the worst case is a scheduler that creates a partition in the network, disconnecting some nodes from the rest: under any (p, ε)-scheduler, if the scheduler's distance from uniform only suffices to disconnect a set of size at most S nodes and we start from a configuration with a gap of $\Omega\left(S + \sqrt{n \log n}\right)$ in favor of one value, then we are guaranteed that all but $O(S)$ nodes will convert to the majority value. We also show that creating a partition is not necessary to cause the system to converge to the wrong value, or to fail to converge at all.

Joint work with Rotem Oshman, Ofer Shayevitz and Yuval Volkov, presented in ITCS 2023.

10:30 - 10:50 *Coffee break*

10:50 - 11:20 **Pierre Fraigniaud -** *Energy efficient distributed algorithms* Abstract: This talk will provide an overview of different models of distributed computing whose objective it to measure the energy consumption of distributed algorithms. The objective of this approach is to develop frugal algorithms (i.e., algorithms consuming little energy) without limiting their ability to solve standard problems studied in the context of network computing. In particular, the presentation will highlight some trade-offs between energy consumption and running time, for distributed algorithms on networks.

11:20 - 11:50 **David Hay -** *Network-Level IoT Security*

Abstract: Computer networks have undergone and continue to experience a significant transformation, whereby billions of low-cost devices are being connected to the network to provide additional functionality and better user experience. Unlike traditional network devices, these devices, collectively known as the "Internet of Things" (IoT), typically have very limited computational, memory, and power resources. These IoT devices became a major security concern, both due to human factors and technical challenges in deploying security mechanisms on devices with low resources. The number and diversity of IoT devices create a huge attack surface that is often exploited by attackers to launch large-scale attacks, sometimes using well-known vulnerabilities.

This talk will highlight the security concerns of IoT devices from a networking perspective and explore how to secure IoT devices using allowlists, in which communication between a device and an endpoint is prohibited unless that endpoint appears in the corresponding allowlist. We will show how to obtain and maintain these allowlists in different settings. Finally, we will discuss deployment options for such a solution (namely, within the internet gateway, as a virtual network function within the ISP network, or a combination of the two).

11:50 - 12:20 **Robin Vacus -** *Stochastic Self-stabilizing Information Spread*

Abstract: Efficiently and reliably spreading information is a fundamental task in many distributed systems, including in the animal world. It is particularly hard when the communication is noisy and constrained, when informed and uninformed agents are indistinguishable, and in the absence of global organization. The "Self-stabilizing information spread" framework of problems was introduced to capture this challenge in a minimal way. In this talk, we will describe several elegant protocols that have been proposed to address it, and present a few impossibility results.

12:20 - 14:15 *Lunch* (On your own)

14:15 - 14:45 **Adrian Vladu -** *Interior Point Methods with a Gradient Oracle*

Abstract: We provide an interior point method based on quasi-Newton iterations, which only requires first-order access to a strongly self-concordant barrier function. To achieve this, we extend the techniques of Dunagan-Harvey [STOC '07] to maintain a preconditioner, while using only first-order information. We measure the quality of this preconditioner in terms of its relative excentricity to the unknown Hessian matrix, and we generalize these techniques to convex functions with a slowly-changing Hessian. We combine this with an interior point method to show that, given first-order access to an appropriate barrier function for a convex set K, we can solve well-conditioned linear optimization problems over K to ε precision in time $\widetilde{O}\left(\left(\mathcal{T}+n^2\right)\sqrt{n\nu}\log\left(1/\varepsilon\right)\right)$, where ν is the self-concordance parameter of the barrier function, and $\mathcal{T}$ is the time required to make a gradient query.

As a consequence we show that linear optimization over n-dimensional convex sets can be performed using $\widetilde{O}(n)$ gradient queries, with an additional cubic running time overhead; this parallels the running time achieved by state of the art algorithms for cutting plane methods when replacing separation oracles with first-order oracles for an appropriate barrier function. In addition we obtain improved running times for solving semidefinite programs in the case where the number of input matrices is much larger than their dimension.

14:45 - 15:15 **Talya Eden -** *Triangle Counting with Local Edge Differential Privacy*

ABSTRACT: Many deployments of differential privacy in industry are in the local model, where each party releases its private information via a differentially private randomizer. We will discuss triangle counting in the noninteractive and interactive local model with edge differential privacy (that, intuitively, requires that the outputs of the algorithm on graphs that differ in one edge be indistinguishable). In this model, each party's local view consists of the adjacency list of one vertex. We will cover new results both in the noninteractive and interactive setting, focusing on a tight lower bound on the error of noninteractive protocols. Our proof uses a reconstruction attack with a new class of linear queries and a novel mix-and-match strategy that runs the local randomizers with different completions of their adjacency lists.

Based on joint work with Quaquan Liu, Sofya Raskhodnikova and Adam Smith.

15:15 - 15:45 **Dani Dorfman -** *Expander Decomposition with Fewer Inter-Cluster Edges Using a Spectral Cut Player*

ABSTRACT: A (ϕ, ϵ)-expander-decomposition of a graph G (with n vertices and m edges) is a partition of V into clusters $V_1, \ldots, V_k$ with conductance $\Phi(G[V_i]) \geq \phi$, such that there are at most ϵm inter-cluster edges. Such a decomposition plays a crucial role in many graph algorithms. We give a randomized $\tilde{O}(m/\phi)$ time algorithm for computing a $(\phi, \phi \log^2 n)$-expander decomposition. This improves upon the $(\phi, \phi \log^3 n)$-expander decomposition also obtained in $\tilde{O}(m/\phi)$ time by [Saranurak and Wang, SODA 2019] (SW) and brings the number of inter-cluster edges within logarithmic factor of optimal.

One crucial component of SW's algorithm is non-stop version of the cut-matching game of [Khandekar, Rao, Vazirani, JACM 2009] (KRV): The cut player does not stop when it gets from the matching player an unbalanced sparse cut, but continues to play on a trimmed part of the large side. The crux of our improvement is the design of a non-stop version of the cleverer cut player of [Orecchia, Schulman, Vazirani, Vishnoi, STOC 2008] (OSVV). The cut player of OSSV uses a more sophisticated random walk, a subtle potential function, and spectral arguments. Designing and analysing a non-stop version of this game was an explicit open question asked by SW.

15:45 - 16:15 *Coffee break*

16:15 - 16:45 **Seffi Naor -** *Online Rounding of Bipartite Matchings*

ABSTRACT: Two complementary facets of the online bipartite matching problem are discussed. (1) For numerous online bipartite matching problems, such as edge-weighted matching and matching under two-sided vertex arrivals, state-of-the-art fractional algorithms outperform their randomized integral counterparts. Thus, a natural question is whether we can achieve lossless online rounding of fractional solutions in this setting. Even though lossless online rounding is impossible in general, randomized algorithms do induce fractional algorithms of the same competitive ratio, which by definition are losslessly roundable online. This motivates the addition of constraints that decrease the "online integrality gap", thus allowing for lossless online rounding. We characterize a set of non-convex constraints which allow for such lossless online rounding and allow for better competitive ratios than yielded by deterministic algorithms. (2) In a different vein, we study the problem of rounding fractional bipartite matchings in online settings. We assume that a fractional solution is already

generated for us online by a black box (via a fractional algorithm, or some machine-learned advice) and provided as part of the input, which we then wish to round. We provide improved bounds on the rounding ratio and discuss several applications.

16:45 - 17:15 **Vianney Perchet -** *Online Matching in Geometric Random Graphs*

ABSTRACT: In online advertisement, ad campaigns are sequentially displayed to users. Both users and campaigns have inherent features, and the former is eligible to the latter if they are "similar enough". We model these interactions as a bipartite geometric random graph: the features of the $2N$ vertices (N users and N campaigns) are drawn independently in a metric space and an edge is present between a campaign and a user node if the distance between their features is smaller than c/N, where $c > 0$ is the parameter of the model.

Our contributions are two-fold. In the one-dimensional case, with uniform distribution over the segment $[0, 1]$, we derive the size of the optimal offline matching in these bi-partite random geometric graphs, and we build an algorithm achieving it (as a benchmark), and analyze precisely its performance.

We then turn to the online setting where one side of the graph is known at the beginning while the other part is revealed sequentially. We study the number of matches of the online algorithm Closest, which matches any incoming point to its closest available neighbor. We show that its performances can be compared to its fluid limit, completely described as a solution of an explicit PDE. From the latter, we can compute the competitive ratio of Closest.

Wednesday May 10, 2023

9:30 - 10:00 **Alex B. Grilo -** *Are quantum-safe zero-knowledge protocols also NISQ-safe*

Abstract: The traditional definition of quantum zero-knowledge stipulates that the knowledge gained by any quantum polynomial-time verifier in an interactive protocol can be simulated by a quantum polynomial-time algorithm. One drawback of this definition is that it allows the simulator to consume significantly more computational resources than the verifier. In this talk, we will discuss new definitions that mitigate this issue, and show (im)possibilities of zero-knowledge protocols under this new assumption.

This is a joint work with Prabhanjan Ananth.

10:00 - 10:30 **Quoc-Huy Vu -** *Public-key Encryption in a Quantum World*

Abstract: It is an important question to find constructions of quantum cryptographic protocols which rely on weaker computational assumptions than classical protocols. In this talk, we focus on quantum public-key encryption (PKE), which is quantum counterparts of classical PKE, and discuss the (im)possibility of constructing quantum PKE from OWF (and weaker assumptions).

This is based on joint work with Khashayar Barooti, Samuel Bouaziz–Ermann, Alex B. Grilo, Loïs Huguenin-Dumittan, Giulio Malavolta, Or Sattath, Damien Vergnaud and Michael Walter.

10:30 - 10:50 *Coffee break*

10:50 - 11:20 **Simon Apers -** *Elfs, trees and quantum walks*

Abstract: We study an elementary Markov process on graphs based on electric flow sampling (elfs). The elfs process repeatedly samples from an electric flow on a graph. While the sinks of the flow are fixed, the source is updated using the electric flow sample, and the process ends when it hits a sink vertex. We argue that this process naturally connects to many key quantities of interest. E.g., we describe a random walk coupling which implies that the elfs process has the same arrival distribution as a random walk. We also analyze the electric hitting time, which is the expected time before the process hits a sink vertex. As our main technical contribution, we show that the electric hitting time on trees is logarithmic in the graph size and weights. The initial motivation behind the elfs process is that quantum walks can sample from electric flows, and they can hence implement this process very naturally. This yields a quantum walk algorithm for sampling from the random walk arrival distribution, which has widespread applications. It complements the existing line of quantum walk search algorithms which only return an element from the sink, but yield no insight in the distribution of the returned element. By our bound on the electric hitting time on trees, the quantum walk algorithm on trees requires quadratically fewer steps than the random walk hitting time, up to polylog factors.

Joint work with Stephen Piddock (arXiv:2211.16379).

11:20 - 11:50 **Geoffroy Couteau -** *Random Sources in Private Computation*

Abstract: We consider multi-party information-theoretic private computation. Such computation inherently requires the use of local randomness by the parties, and the question of minimizing the total number of random bits used for given private computations has received considerable attention in the literature, see, e.g., [KR94, KM96, KOR98, BDSPV99, JLR03, DPP16, RU19, KOP+19]. In this work we are interested in another question: given a private computation, we ask how many of the players need to have access to a random source, and how many of them can be deterministic parties. We are further interested in the possible interplay between the number of random sources in the system and the total number of random bits necessary for the computation.

We give a number of results. We first show that, perhaps surprisingly, t players (rather than $t+1$) with access to a random source are sufficient for the information-theoretic t-private computation of any deterministic functionality over n players for any $t < n/2$; by a result of [KM96], this is best possible. This means that, counterintuitively, while private computation is impossible without randomness, it is possible to have a private computation even when the adversary can control all parties who can toss coins (and therefore sees all random coins). For randomized functionalities, we show that $t+1$ random sources are necessary (and sufficient). We then turn to the question of the possible interplay between the number of random sources and the necessary number of random bits. Since for only very few settings in private computation meaningful bounds on the number of necessary random bits are known, we consider the AND function, for which some such bounds are known. We give a new protocol to 1-privately compute the n-player AND function, which uses a single random source and 6 random bits tossed by that source. This improves, upon the currently best-known results [KOP+21], at the same time the number of sources and the number of random bits ([KOP+21] gives a 2-source, 8-bit protocol). This result may give some evidence that for 1-privacy, using the minimum necessary number of sources one can also achieve the necessary minimum number of random bits. We believe however that our protocol is of independent interest to the study of randomness in private computation.

11:50 - 12:20 **Lianna Hambardzumyan -** *An improved protocol for Exactly-N with more than* 3 *players*

Abstract: The Exactly-N problem in the number-on-forehead (NOF) communication setting asks k players, each of whom can see every input but their own, if the k input numbers add up to N. Chandra, Furst, and Lipton (STOC 1983) introduced the Exactly-N problem and showed that its k-party NOF communication complexity exactly corresponds to the number of colors needed to color $[N]^{k-1}$ avoiding monochromatic corners. The latter is a generalization of van der Waerden's famous coloring problem concerning arithmetic progressions (APs).

Until recently, the best-known corner-free colorings came via a reduction to the constructions of AP-free sets. New results by Linial and Shraibman (CCC 2021) and Green (New Zealand Journal of Mathematics 2021) changed this picture by providing improved protocols for Exactly-N in the $k = 3$ case. The communication complexity point of view is key to these improvements.

In this work, we design an *explicit* protocol for Exactly-N that works for any number of players. Based on that, we give an *improved* protocol, which matches Green's construction for $k = 3$ players but works for any number of players. Consequently, our protocol yields the first improvement in the higher-order terms of higher-dimensional corner-free colorings since the original construction of Rankin from 1961.

The talk is based on joint work with Toniann Pitassi, Suhail Sherif, Morgan Shirley, and Adi Shraibman.

12:20 - 14:15 *Lunch* (On your own)

14:15 - 14:45 **Michal Feldman -** *Algorithmic Contract Design*

Abstract: Contract design captures situations where a principal delegates the execution of a costly task to an agent. To complete the task, the agent chooses an action from a set of costly actions. The principal can only observe the outcome, which is stochastically determined by the chosen action. The principal incentivizes the desired action through a

contract, that specifies payments based on the observed outcome. In this talk, I will survey two papers on *combinatorial* contracts, which highlight different sources of complexity that arise in contract design. The first (FOCSÕ21) is where the agent can choose any subset of a given set of actions; the second (STOCÕ23) is where the principal motivates a team of agents. We provide (approximation) algorithms and hardness results for the optimal contract problem in these scenarios.

Based on joint work with Tomer Ezra, Paul Duetting and Thomas Kesselheim.

14:45 - 15:15 **Alon Eden -** *Private Interdependent Valuations*

Abstract: The celebrated Interdependent values (IDV) model captures settings where buyers have partial information regarding their value for the item being sold. In the IDV model, each buyer i has a private piece of information, a signal s_i, and their value is determined by all signals. This is represented by a valuation function $v_i(s_1, \text{É}, s_n)$. While this model is more realistic than the private values model (that assumes buyers know their value), it is much less understood. Most works in the IDV model assume that the valuation function v_i is public information.

I will present recent works where both the signal and the valuation function are considered to be private information. In this setting, assuming the valuation functions satisfy the standard single-crossing condition can only guarantee trivial approximation. I will discuss how the recently introduced Submodularity-over-signals (SOS) condition can be leveraged to obtain a constant factor approximation to the optimal welfare.

Based on joint-works with Michal Feldman, Kira Goldner, Simon Mauras, Divyarthi Mohan and Shuran Zheng.

15:15 - 15:45 **Simon Mauras -** *Truthful Matching with Online Items and Offline Agents*

Abstract: We study truthful mechanisms for welfare maximization in online bipartite matching. In our (multi-parameter) setting, every buyer is associated with a (possibly private) desired set of items, and has a private value for being assigned an item in her desired set. Unlike most online matching settings, where agents arrive online, in our setting the items arrive online in an adversarial order while the buyers are present for the entire duration of the process. This poses a significant challenge to the design of truthful mechanisms, due to the ability of buyers to strategize over future rounds. We provide an almost full picture of the competitive ratios in different scenarios, including myopic vs. non-myopic agents, tardy vs. prompt payments, and private vs. public desired sets. Among other results, we identify the frontier for which the celebrated $e/(e-1)$ competitive ratio for the vertex-weighted online matching of Karp, Vazirani and Vazirani extends to truthful agents and online items.

15:45 - 16:15 *Coffee Break*

16:15 - 16:45 **Claire Mathieu -** *A Tight* $(1.5 + \epsilon)$*-Approximation for Unsplittable Capacitated Vehicle Routing on Trees*

Abstract: In the unsplittable capacitated vehicle routing problem (UCVRP) on trees, we are given a rooted tree with edge weights and a subset of vertices of the tree called terminals. Each terminal is associated with a positive demand between 0 and 1. The goal is to find a minimum length collection of tours starting and ending at the root of the tree such that the demand of each terminal is covered by a single tour (i.e., the demand cannot be split), and the total demand of the terminals in each tour does not exceed the capacity of 1. For the special case when all terminals have equal demands, a long line of research culminated in a quasi-polynomial time approximation scheme [Jayaprakash and Salavatipour, SODA 2022] and a polynomial time approximation scheme [Mathieu and Zhou, ICALP 2022]. We

study the general case when the terminals have arbitrary demands, giving a polynomial time $(1.5 + \epsilon)$-approximation algorithm for the UCVRP on trees. This is the first improvement upon the 2-approximation algorithm more than 30 years ago [Labbé, Laporte, and Mercure, Operations Research, 1991]. Our approximation ratio is essentially best possible, since it is NP-hard to approximate the UCVRP on trees to better than a 1.5 factor.

This is joint work with Hang Zhou.

16:45 - 17:15 **Hang Zhou -** *Euclidean Capacitated Vehicle Routing*

ABSTRACT: In the capacitated vehicle routing problem, we are given a metric space with a vertex called depot and a set of vertices called terminals. The goal is to find a minimum length collection of tours starting and ending at the depot such that each tour visits at most k terminals, and each terminal is visited by some tour.

I am going to talk about recent progress on the capacitated vehicle routing problem in the two-dimensional Euclidean plane. This is joint work with Claire Mathieu [RSA 2023], with Fabrizio Grandoni and Claire Mathieu [ITCS 2023], and with Zipei Nie [work in submission].

17:15 - 17:45 **Ami Paz -** *A Speedup Theorem for Asynchronous Computation*

ABSTRACT: Speedup theorems have recently gained increasing attention in studying distributed graph algorithms in synchronous systems. Using topological tools, we present a variant of this technique that applies to asynchronous shared memory systems. At the core of our technique is a round reduction theorem: given a distributed task Π, we define a new task called the closure of Π, such that if Π is solvable in t rounds, then its closure is solvable in $t-1$ rounds.

We illustrate the power of our speedup theorem by providing new proof of the wait-free impossibility of consensus using read/write registers. This is done merely by showing that the closure of consensus is consensus itself. The simplicity of our technique allows us to study additional communication objects, namely test&set and binary consensus. We analyze the approximate agreement task in systems augmented with the two objects and show that while these objects are more powerful than read/write registers from the computability perspective, they do not help to reduce the time complexity of solving approximate agreement.

Based on joint work with Pierre Fraigniaud and Sergio Rajsbaum.

19:00 - *Social Event*

Thursday May 11, 2023

9:30 - 10:00 **Spyros Angelopoulos -** *Online Bin Packing with Predictions*

Abstract: Bin packing is a classic optimization problem with a wide range of applications, from load balancing to supply chain management. In this work, we study the online variant of the problem, in which a sequence of items of various sizes must be placed into a minimum number of bins of uniform capacity. The online algorithm is enhanced with a (potentially erroneous) prediction concerning the frequency of item sizes in the sequence. We design and analyze online algorithms with efficient tradeoffs between the consistency (i.e., the competitive ratio assuming no prediction error) and the robustness (i.e., the competitive ratio under adversarial error), and whose performance degrades near-optimally as a function of the prediction error. This is the first theoretical and experimental study of online bin packing under competitive analysis, in the realistic setting of learnable predictions. Previous work addressed only extreme cases with respect to the prediction error, and relied on overly powerful and error-free oracles. Joint work with Shahin Kamali and Kimia Shadkami.

10:00 - 10:30 **Haim Kaplan -** *Almost Tight Bounds for Online Facility Location in the Random-Order Model*

Abstract: We study the online facility location problem with uniform facility costs in the random-order model. Meyerson's algorithm [FOCS'01] is arguably the most natural and simple online algorithm for the problem with several advantages and appealing properties. Its analysis in the random-order model is one of the cornerstones of random-order analysis beyond the secretary problem. Meyerson's algorithm was shown to be (asymptotically) optimal in the standard worst-case adversarial-order model and 8-competitive in the random order model. While this bound in the random-order model is the long-standing state-of-the-art, it is not known to be tight, and the true competitive-ratio of Meyerson's algorithm remained an open question for more than two decades.

We resolve this question and prove tight bounds on the competitive-ratio of Meyerson's algorithm in the random-order model, showing that it is exactly 4-competitive. Following our tight analysis, we introduce a generic parameterized version of Meyerson's algorithm that retains all the advantages of the original version. We show that the best algorithm in this family is exactly 3-competitive. On the other hand, we show that no online algorithm for this problem can achieve a competitive-ratio better than 2. Finally, we prove that the algorithms in this family are robust to partial adversarial arrival orders.

10:30 - 11:50 *Coffee break*

10:50 - 11:20 **Allan Borodin -** *Revisiting an old problem (online interval selection) and where it takes us*

Abstract: Online interval selection is a well studied problem that still has some research challenges. We will consider interval selection (and its relation to call admission on the line) with respect to a number of different assumptions, extensions, and online computational and input models. We revisit the interval selection algorithms of Adler and Azar, and Emek, Halldo'rsson and Rose'n to motivate more generally the question of memory vs competitiveness in online algorithms.

Work in progress with Chris Karavasilis.

11:20 - 11:50 **Shahar Lewkovicz -** *List Update with Delays or Time Windows*

Abstract: We consider the problem of List Update, one of the most fundamental problems in online algorithms and competitive analysis. Informally, we are given a list of elements and

requests for these elements that arrive over time. Our goal is to serve these requests, at a cost equivalent to their position in the list, with the option of moving them towards the head of the list. Sleator and Tarjan introduced the famous "Move to Front" algorithm (wherein any requested element is immediately moved to the head of the list) and showed that it is 2-competitive. While this bound is excellent, the absolute cost of the algorithm's solution may be very large (e.g., requesting the last half elements of the list would result in a solution cost that is quadratic in the length of the list). Thus, we consider the more general problem wherein every request arrives with a deadline and must be served, not immediately, but rather before the deadline. We further allow the algorithm to serve multiple requests simultaneously - thereby allowing a large decrease in the solutions' total costs. We denote this problem as List Update with Time Windows. While this generalization benefits from lower solution costs, it requires new types of algorithms. In particular, for the simple example of requesting the last half elements of the list with overlapping time windows, Move-to-Front fails. In this work we show an O(1) competitive algorithm. The algorithm is natural but the analysis is a bit complicated and a novel potential function is required. Thereafter we consider the more general problem of List Update with Delays in which the deadlines are replaced with arbitrary delay functions. This problem includes as a special case the prize collecting version in which a request might not be served (up to some deadline) and instead suffers an arbitrary given penalty. Here we also establish an O(1) competitive algorithm for general delays. The algorithm for the delay version is more complex and its analysis is significantly more involved. Joint work with Yossi Azar and Danny Vainstein.

11:50 - 12:20 **Chien-Chung Huang -** *FPT-Algorithms for the ℓ-Matchoid Problem with a Coverage Objective*

Abstract: We consider the problem of optimizing a coverage function under a $\ell-$matchoid of rank k. We design fixed-parameter algorithms as well as streaming algorithms to compute an exact solution. Unlike previous work that presumes linear representativity of matroids, we consider the general oracle model. For the special case where the coverage function is linear, we give a deterministic fixed-parameter algorithm parameterized by ℓ and k. This result, combined with the lower bounds of Lovasz, and Jensen and Korte demonstrates a separation between the $\ell-$matchoid and the matroid $\ell-$parity problems in the setting of fixed-parameter tractability. For a general coverage function, we give both deterministic and randomized fixed-parameter algorithms, parameterized by ℓ and z, where z is the number of points covered in an optimal solution. This result, combined with existing lower bounds, also provides a new separation between the space and time complexity of maximizing an arbitrary submodular function and a coverage function in the value oracle model.

This is joint-work with Justin Ward (University of Queen Mary university of London).

12:20 - 14:15 *Lunch* (On your own)

14:15 - 14:45 **Yossi Azar -** *Multi Layer Peeling for Linear Arrangement and Hierarchical Clustering*

Abstract: We present a new multi-layer peeling technique to cluster points in a metric space. A well-known non-parametric objective is to embed the metric space into a simpler structured metric space such as a line (i.e., Linear Arrangement) or a binary tree (i.e., Hierarchical Clustering). Points which are close in the metric space should be mapped to close points/leaves in the line/tree; similarly, points which are far in the metric space should be far in the line or on the tree. In particular we consider the Maximum Linear Arrangement problem and the Maximum Hierarchical Clustering problem applied to metrics.

We design approximation schemes ($1-\epsilon$ approximation for any constant $\epsilon > 0$) for these objectives. In particular this shows that by considering metrics one may significantly improve former approximations (0.5 for Max Linear Arrangement and 0.74 for Max Hierarchical

Clustering). Our main technique, which is called multi-layer peeling, consists of recursively peeling off points which are far from the "core" of the metric space. The recursion ends once the core becomes a sufficiently densely weighted metric space (i.e. the average distance is at least a constant times the diameter) or once it becomes negligible with respect to its inner contribution to the objective. Interestingly, the algorithm in the Linear Arrangement case is much more involved than that in the Hierarchical Clustering case and uses a significantly more delicate peeling.

Joint work with Danny Vainstein

14:45 - 15:15 **Yaniv Sadeh -** *Dynamic Binary Search Trees: Improved Lower Bounds for the Greedy-Future Algorithm*

ABSTRACT: Binary search trees (BSTs) are one of the most basic and widely used data structures. The best static tree for serving a sequence of queries (searches) can be computed by dynamic programming. In contrast, when the BSTs are allowed to be dynamic (i.e. change by rotations between searches), we still do not know how to compute the optimal algorithm (OPT) for a given sequence. One of the candidate algorithms whose serving cost is suspected to be optimal up-to a (multiplicative) constant factor is known by the name Greedy Future (GF). In an equivalent geometric way of representing queries on BSTs, GF is in fact equivalent to another algorithm called Geometric Greedy (GG). Most of the results on GF are obtained using the geometric model and the study of GG. Despite this intensive recent fruitful research, the best lower bound we have on the competitive ratio of GF is 4/3. Furthermore, it has been conjectured that the additive gap between the cost of GF and OPT is only linear in the number of queries. In this paper we prove a lower bound of 2 on the competitive ratio of GF, and we prove that the additive gap between the cost of GF and OPT can be $\Omega(m \cdot \log\log n)$ where n is the number of items in the tree and m is the number of queries.

15:15 - 15:45 **Daniel Vaz -** *Good and Efficient Approximation for the Sparsest Cut Problem in Bounded-Treewidth Graphs*

ABSTRACT: Sparsest cut is a fundamental graph problem, which models a general notion of balanced separator of a graph, and has uses in graph theory and divide-and-conquer approaches. In it, we are given an edge-capacitated graph, together with demands over pair of vertices, and we want to find a cut that minimizes the ratio between the capacities and demands across the cut. In other words, we aim to separate as much demand as possible using little cut capacity. For general graphs, the best known approximation factor is $\widetilde{O}(\sqrt{\log n})$, and the problem is known to have no constant-approximation under the unique games conjecture.

In this talk, we focus on the simpler setting of bounded-treewidth graphs, and present new constant-factor approximation algorithms. Known algorithms in this setting are either inefficient or have large approximation factors. We will see that these algorithms can be framed as specific cases of a general algorithm with uses beyond sparsest cut, and then show how to combine the best of both approaches to obtain efficient and good approximations to the problem. As a result, we give the first constant-factor approximation algorithm running in FPT time.

15:45 - 16:15 *Coffee Break*

16:15 - 16:45 **Tomasz Ponitka -** *On Optimal Tradeoffs between EFX and Nash Welfare*

ABSTRACT: A major problem in fair division is how to allocate a set of indivisible resources among agents fairly and efficiently. We give optimal tradeoffs between fairness and efficiency, with respect to well-studied measures of fairness and efficiency – envy freeness up to any item (EFX) for fairness, and Nash welfare for efficiency. Our results improve upon the current

state of the art, for both additive and subadditive valuations. For additive valuations, we show the existence of allocations that are simultaneously α-EFX and guarantee a $(1/(\alpha+1))$-fraction of the maximum Nash welfare, for any $\alpha \in [0, 1]$. For $\alpha \in [0, \varphi - 1 \approx 0.618]$ these are complete allocations (all items are assigned), whereas for larger α these are partial allocations (some items may be unassigned). We partially extend this to subadditive valuations where we show the existence of complete allocations that give α-EFX and a $(1/(\alpha + 1))$-fraction of the maximum Nash welfare (as above), for any z$\alpha \in [0, 1/2]$. We also give impossibility results that show that our tradeoffs are tight, even with respect to partial allocations.

16:45 - 17:15 **Amos Fiat -** *EFX Allocations on Graphs*

Abstract: We study *envy freeness up to any good (EFX)* in settings where valuations can be represented via a graph of arbitrary size where vertices correspond to agents and edges to items. An item (edge) has zero marginal value to all agents (vertices) not incident to the edge. Each vertex may have an arbitrary monotone valuation on the set of incident edges. We first consider allocations that correspond to orientations of the edges, where we show that EFX does not always exist, and furthermore that it is NP-complete to decide whether an EFX orientation exists. Our main result is that (EFX) allocations exist for this setting. This is one of the few cases where EFX allocations are known to exist for more than 3 agents.

Joint work with George Christodoulou, Elias Koutsoupias, Alkmini Sgouritsa.