Program

Thursday, April 4, 2024

9:00 - 9:25 *Gathering, Registration, and Coffee*

9:25 - 9:30 **Opening Remarks**

9:30 - 10:00 **Moni Naor (Weizmann) -** *From Donkeys to Kings in Tournaments*

A tournament is an orientation of a complete graph. A vertex that can reach every other vertex within two steps is called a king. We study the complexity of finding k kings in a tournament graph.

We show that the randomized query complexity of finding $k \leq 3$ kings is linear in n, the number of vertices. For both the deterministic and randomized cases it takes the same amount of queries (up to a constant) as finding a single king (the best-known deterministic algorithm makes $O(n^{3/2})$ queries). On the other hand, we show that finding $k > 3$ kings requires $\Omega(n^2)$ queries, even in the randomized case.

We consider the RAM model for $k > 3$: We show an algorithm that finds k kings in time $O(kn^2)$, which is optimal for constant values of k. Alternatively, one can also find k kings in time n^ω (the time for matrix multiplication). We provide evidence that this is the best possible for large k by suggesting a fine-grained reduction from a variant of the triangle detection problem.

Joint work with Amir Abboud, Tomer Grossman and Tomer Solomon

10:00 - 10:30 **Moshe Tennenholtz (Technion) -** *The Search for Stability: Learning Dynamics and the Relative Ranking Principle*

We study a game-theoretic information retrieval model in which strategic publishers aim to maximize their chances of being ranked first by the search engine while maintaining the integrity of their original documents. We show that the commonly used Probability Ranking Principle (PRP) ranking scheme results in an unstable environment where games often fail to reach pure Nash equilibrium. We propose the Relative Ranking Principle (RRP) as an alternative ranking principle and introduce two families of ranking functions that are instances of the RRP. We provide both theoretical and empirical evidence that these methods lead to a stable search ecosystem, by providing positive results on the learning dynamics convergence. We also define the publishers' and users' welfare, demonstrate a possible publisher-user trade-off, and provide means for a search system designer to control it.

Joint work with Omer Madmon, Idan Pipano and Itamar Reinman

10:30 - 11:00 **Keren Censor-Hillel (Technion) -** *Distributed Subgraph Finding: Known and New*

Finding small subgraphs is a fundamental task which is extensively studied in various computational settings. In this talk, I will chart what is known about variants of subgraph finding problems in distributed computing, present some new results, and discuss open problems and the challenges they bring.

11:0 - 11:30 *Coffee break*

11:30 - 12:00 **Dan Halperin (TAU) -** *Multi-Robot Motion Planning: The Easy, the Hard and the Uncharted*

Early results in robot motion planning had forecast a bleak future for the field by showing that problems with many degrees of freedom, and in particular those involving fleets of robots, are intractable. Then came sampling-based planners, which have been successfully, and often easily, solving a large variety of problems with many degrees of freedom. We strive to formally determine what makes a motion-planning problem with many degrees of freedom easy or hard. I'll describe our quest to resolve this (still wide open) problem, and some progress we have made in the context of multi-robot motion planning.

12:00 - 12:30 **Nadav Merlis (CREST, ENSAE Paris) -** *The Value of Reward Lookahead in Reinforcement Learning*

In reinforcement learning (RL), agents sequentially interact with changing environments while aiming to maximize the obtained rewards. Usually, rewards are observed only after acting, and so the goal is to maximize the expected cumulative reward. Yet, in many practical settings, reward information is observed in advance – prices are observed before performing transactions; nearby traffic information is partially known; and goals are oftentimes given to agents prior to the interaction. In this work, we aim to quantifiably analyze the value of such future reward information through the lens of competitive analysis. In particular, we measure the ratio between the value of standard RL agents and that of agents with partial future-reward lookahead. We characterize the worst-case reward distribution and derive exact ratios for the worst-case reward expectations. Surprisingly, the resulting ratios relate to known quantities in offline RL and reward-free exploration. We further provide tight bounds for the ratio given the worst-case dynamics. Our results cover the full spectrum between observing the immediate rewards before acting to observing all the rewards before the interaction starts.

12:30 - 14:00 *Lunch* (On your own)

14:00 - 14:30 **Noam Nissan (HUJI) -** *Serial Monopoly on Blockchains*

We study the following problem that is motivated by Blockchains where "miners" are serially given the monopoly for assembling transactions into the next block. Our model has a single good that is sold repeatedly every day where new demand for the good arrives every day. The novel element in our model is that all unsatisfied demand from one day remains in the system and is added to the new demand of the next day. Every day there is a new monopolist that gets to sell a fixed supply of the good and naturally chooses to do so at the monopolist's price for the combined demand. What will the dynamics of the prices chosen by the sequence of monopolists be? What level of efficiency will be obtained in the long term?

We start with a non-strategic analysis of users' behavior and our main result shows that prices keep fluctuating wildly and this is an endogenous property of the model and happens even when demand is stable with nothing stochastic in the model. These price fluctuations underscore the necessity of an analysis under strategic behavior of the users, which we show results in the prices being stable at the market equilibrium price.

14:30 - 15:00 **Niv Buchbinder (TAU) -** *Chasing Positive Bodies*

We study the problem of chasing positive bodies in ℓ_1: given a sequence of bodies $K_t \subset R^n$ revealed online, where each K_t is defined by a mixed packing-covering linear program, the goal is to (approximately) maintain a point $x_t \in K_t$ such that $\sum_t \|x_t - x_{t-1}\|_1$ is minimized. This captures the fully-dynamic low-recourse variant of any problem that can be expressed as a mixed packing-covering linear program and thus also the fractional version of many central

problems in dynamic algorithms such as set cover, load balancing, hyperedge orientation, minimum spanning tree, and matching.

We give an $O(\log n)$-competitive algorithm for this problem. This bypasses and improves exponentially over the lower bound of $\sqrt{n}$ known for general convex bodies. Our algorithm is based on iterated information projections, and, in contrast to general convex body chasing algorithms, is entirely memoryless.

We also show how to round our solution dynamically to obtain the first fully dynamic algorithms with competitive recourse for all the stated problems above; i.e. their recourse is less than the recourse of every other algorithm on every update sequence, up to polylogarithmic factors. This is a significantly stronger notion than the notion of absolute recourse in the dynamic algorithms literature

15:00 - 15:30 **Liat Peterfreund (HUJI) -** *From Standardization to Theory and Back*

Graph databases are becoming increasingly popular due to their natural data modeling, making them useful in expressing connections that are harder to express in the relational model. Indeed, graph databases are used in a plethora of domains ranging from social to biological networks, and for various use-cases including fraud detection and investigating journalism. Since 2019, GQL (Graph Query Language) is being developed under the auspices of ISO as the new standard for querying graph databases, akin to SQL for relational databases. In this talk, I will present a researcher's digest of GQL by describing its underlying theoretical model. I will demonstrate how we can use tools from formal language and automata theory to show the limitations of this new standard, which can hint at extensions for its next versions.

This talk is based on joint works with Nadime Francis, Amélie Gheerbrant, Paolo Guagliardo, Leonid Libkin, Victor Marsault , Wim Martens, Filip Murlak, Alexandra Rogova, and Domagoj Vrgoc.

15:30 - 16:00 *Coffee break*

16:00 - 16:30 **Joe Halpern (Cornell) -** *A Language-Based Decision Theory with Causality*

The talk discusses a sequence of four papers that give a language-based approach to decision theory.

The classical approach in decision theory (going back to Savage) is to place a preference order on acts, where an act is a function from states to outcomes. If the preference order satisfies appropriate postulates, then the decision maker can be viewed as acting as if he has a probability on states and a utility function on outcomes, and is maximizing expected utility. This framework implicitly assumes that the decision maker knows what the states and outcomes are. That isn't reasonable in a complex situation. For example, in trying to decide whether or not to attack Gaza, what are the states and what are the outcomes? Larry Blume, David Easley, and I consider a language-based framework in which actions are identified with (conditional) descriptions in a simple underlying language, while states and outcomes (along with probabilities and utilities) are constructed as part of a representation theorem. role of language in decision making, using it not only for the conditions that determine which actions are taken, but also the effects. We show that this approach has many benefits, both conceptual and pragmatic, of this approach. Among other things, it provides an elegant solution to framing problems.

Adam Bjorndahl and I then consider a special case of the BEH approach, where the set of actions is built from ones of the form do(p), for formulas p in the underlying language. The action do(p) is meant to represent an intervention that result in p being true. But this action is underspecified: there are many ways of making p true. We deal with this using the semantics

of counterfactuals defined in the philosophy community: roughly speaking, do(p) maps each state to the "closest" state where p is true. This notion of "closest" is is constructed as part of the representation theorem. Different agents can have different preferences regarding an act like do(raise the minimum wage to $15) because they have different notions regarding the closest world where the minimum wage is $15.

Finally, in work joint with Evan Piermont, we consider essentially the same language with do(ϕ), but now give a representation theorem that uses causal models, instead of "closest-world" semantics. This result actually goes through the Bjorndahl-Halpern result, showing a close relationship between the two approaches to giving semantics to counterfactuals.

The talk is self-contained, and does not presume any background in decision theory or causal modeling.

16:30 - 17:00 **Robert Krauthgamer (Weizmann) -** *Fully Scalable MPC Algorithms for Clustering in High Dimension*

We design new algorithms for k-clustering in high-dimensional Euclidean spaces. These algorithms run in the Massively Parallel Computation (MPC) model, and are fully scalable, meaning that the local memory in each machine is an arbitrarily small poly(n) for input size n, which importantly may be substantially smaller than k. Our algorithms are fast, i.e., take $O(1)$ rounds, and achieve $O(1)$-bicriteria approximation for k-Median and for k-Means, while previous work achieves only polylog(n)-bicriteria approximation or handles a special case.

Our results rely on a fast MPC algorithm for $O(1)$-approximation of facility location. A primary technical tool that we develop, and may be of independent interest, is a new MPC primitive for geometric aggregation, namely, computing certain statistics on an approximate neighborhood of every data point, which includes range counting and nearest-neighbor search.

Joint work with Artur Czumaj, Guichen Gao, Shaofeng H.-C. Jiang, and Pavel Vesely.

17:00- 17:30 **Kobbi Nissim (Georgetown) -** *A methodology for reconciling computer science and legal approaches to privacy*

As law and computer science interact in critical ways within sociotechnical systems, recognition is growing of significant gaps between these disciplines that create potential risks for privacy and data protection. These gaps need to be bridged to ensure that computer systems are designed and implemented to correctly address applicable legal requirements and that interpretations of legal concepts accurately reflect the capabilities and limitations of technical systems.

We will explore some of the gaps between the legal and computer science views of privacy and suggest directions towards mitigating them while respecting the values and principles of both disciplines.

Partly based on work in progress with Micah Altman and Aloni Cohen