Program

Monday November 20, 2017

Morning - Beit Hatefutzot, Zeevi Auditorium

Afternoon - Shenkar building, Melamed 006 / Holcblat 007

8:30 - 9:00 **Gathering and Registration**

9:00 - 9:10 **Opening**

9:10 - 9:40 **Moni Naor -** *Fiat-Shamir and Proofs of Non-Memberships*

In this talk I will discuss old and new work related to the famed Fiat-Shamir Crypto 1986 paper. In particular I will describe work related to proofs of non-membership in a set (with Asaf Ziv) and collecting information of passwords while maintaining privacy (with Eyal Ronen and Benny Pinkas).

9:40 - 10:10 **Benny Pinkas -** *From Broadcast Encryption to Multi-Party Set Intersection* In the talk I'll revisit the famed results of Fiat-Naor and Chor-Fiat-Naor on broadcast encryption and traitor tracing, and describe how they influenced my work on private set intersection, and in particular recent work on secure computation of the intersection of multiple sets.

10:10 - 10:40 **Ran Canetti -** *Using Diffie and Hellman to Wield the Magic of Fiat and Shamir*

In 1986 Fiat and Shamir threw a major challenge to the cryptographic arena: They showed how to turn an interactive identification scheme into a non-interactive signature scheme by using a publicly computable function that is "sufficiently random-looking" instead of the interaction. The idea proved very powerful and was used over and over again. However, despite many attempts in the 31 years that followed, the cryptographic community was unable to come up with constructions that would be provably sufficient for the Fiat-Shamir transformation, under some clear computational hardness assumptions. In fact, over the years some strong impossibility results regarding the existence of these elusive "magic functions" were proven in some cases, as well as barriers to their provability in other cases. Instead, the cryptographic community resorted to the not-so-realistic modeling of these magic, publicly computable functions as truly random functions.

First signs that these magic functions might indeed be constructed in an analyzable way were given in a work by Kalai, Rothblum and Rothblum at Crypto 17: They constructed a function family that suffices for the Fiat-Shamir transformation, when applied to statistically sound proofs. However that construction uses very strong variants of program obfuscation that are not known to exist under more standard assumptions.

Building on the techniques of KRR17, we construct simple and efficient families of functions that suffice for the Fiat-Shamir transform, under the assumption that a variant of the Diffie-Hellman problem (or, alternatively, a variant of the Learning With Errors problem) is solvable by polytime algorithms only with exponentially small probability. As corollaries we get non-interactive zero-knowledge protocols and succinct delegation schemes based on different assumptions than previously known.

Joint work with Yilei Chen, Leonid Reyzin, and Ron Rothblum

10:40 - 11:00 *Coffee break*

11:00 - 11:40 **Avrim Blum -** *Learning about Agents and Mechanisms from Opaque Transactions*

In this talk I will discuss the problem of trying to learn the requirements and preferences of economic agents by observing the outcomes of an allocation mechanism whose rules you also donÕt initially know. As an example, consider observing web pages where the agents are advertisers and the winners are those whose ads show up on the given page. We know these ads are placed based on bids and other constraints given to some auction mechanism, but we do not get to see these bids and constraints. What we would like to do is from repeated observations of this type to learn what the requirements and preferences of the agents are. Or consider observing the input-output behavior of a cloud computing service, where the input consists of a set of agents requesting service, and the output tells us which actually received service and which did not. In this case, we assume the agents who did not receive service were not served due to overlap of their resource needs with higher-priority requests. From such input-output behavior, we would like to learn the underlying structure. Our goal will be from observing a series of such interactions to try to learn both the needs and preferences of the agents and perhaps also the rules of the allocation mechanism. This talk is based on work joint with Yishay Mansour and Jamie Morgenstern, as well as work joint with undergraduate student Michael Liang.

11:40 - 12:20 **Dan Feldman -** *Provable Real-Time Learning with applications to Robotics*

A coreset (or, core-set) for a given problem is a "compressed" representation of its input, in the sense that a solution for the problem with the (small) coreset as input would yield a provable $(1+\epsilon)$ factor approximation to the problem with the original (large) input.

Using traditional techniques, a coreset usually implies provable linear time algorithms for the corresponding optimization problem, which can be computed in parallel, via one pass over Big data on the cloud, and using only logarithmic space (i.e, in the streaming model).

In this talk I will forge links between coresets for machine learning, computational geometry, and robotics. In particular, we use geometric techniques such as John Ellipsoid to compute small coresets for optimal space exploration and logistic regression.

Finally, experimental results and videos from the first autonomous toy-quadcopters in a supermarket (Rami Levy, Nesher) will be presented.

Joint work with Murad Tukan and Elad Tolichensky.

12:30 - 14:00 *Lunch* (Please see our list for a selection of restaurants on campus.)

14:00 - 14:30 **Michal Feldman -** *Working with Amos on Pricing is Priceless* -
Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs

We present a general framework for stochastic online maximization problems with combinatorial feasibility constraints. The framework establishes prophet inequalities by constructing price-based online approximation algorithms, a natural extension of threshold algorithms for settings beyond binary selection. Our analysis takes the form of an extension theorem: we derive sufficient conditions on prices when all weights are known in advance, then prove that the resulting approximation guarantees extend directly to stochastic settings. Our framework unifies and simplifies much of the existing literature on prophet inequalities and posted price mechanisms, and is used to derive new and improved results for combinatorial markets (with and without complements), multi-dimensional matroids, and sparse packing problems. Finally, we highlight a surprising connection between the smoothness framework for bounding

the price of anarchy of mechanisms and our framework, and show that many smooth mechanisms can be recast as posted price mechanisms with comparable performance guarantees.

Joint work with Paul Dutting, Thomas Kesselheim and Brendan Lucier.

14:30 - 15:00 **Uri Nadav -** *Bid Limited Targeting in Ad Auctions*

This paper analyzes a mechanism for selling items in auctions in which the auctioneer specifies a cap on the ratio between the maximum and minimum bids that bidders may use. Such a mechanism is widely used in the online advertising business through the caps that companies impose on the minimum and maximum bid multipliers that advertisers may use in targeting. When biddersÕ values are independent and identically distributed, using this mechanism results in higher revenue than allowing bidders to condition their bids on the targeting information in an arbitrary way and also almost always results in higher revenue than not allowing bidders to target. Choosing the optimal cap on the ratio between the maximum bid and the minimum bid can also be more important than introducing additional competition in the auction. However, if biddersÕ values are not identically distributed, pure-strategy equilibria may fail to exist.

15:00 - 15:30 **Stefano Leonardi -** *Approximately Optimal Mechanisms in Two-Sided Markets*

I can clearly remember the first time I heard about algorithmic auction design. It was about 15 years ago from Amos while sitting together at a cafè on Tel-Aviv shoreline.

I'll talk at AmosFest indeed about algorithmic auction design in which both sides - buyers and sellers - can act strategically. Unfortunately, Myerson and Satterthwaite proved in 1983 that social welfare maximising two-sided auctions which are incentive compatible (IC), individually rational (IR) and budget balanced (BB) do not exist even in the bayesian setting.

In this talk, I'll discuss sequential posted price two-sided mechanisms which provide a provable good approximation to the optimal social welfare while obeying the IC, IR and BB constraints. The mechanisms work for any number of buyers and sellers with independent distributions and matroid constraints for unit supply buyers and sellers. These mechanisms can also be extended to two-sided combinatorial auctions with additive and XOS valuations. Finally, I'll discuss the harder problem of optimising the gain from trade that can be obtained from the exchange.

Based on joint works with Riccardo Colini-Baldeschi, Paul Goldberg, Bart de Keijzer, Tim Roughgarden and Stefano Turchetta.

15:30 - 16:10 **Nati Linial -** *Hypertrees*

An elementary collapse in a graph G is a step in which we delete a vertex of degree one and the single edge that contains it. We say that G is collapsible if by repeatedly applying such steps it is possible to eliminate all of G's edges. Clearly G is collapsible iff it is a forest, i.e., iff it is acyclic. All these concepts have higher-dimensional counterparts, but this equivalence is lost. Namely, for $d > 1$ a d-dimensional collapsible complex is acyclic, but not vice versa. In this talk I will discuss these structures and the ways in which higher-dimensional hypertrees differ from (classical, one-dimensional) trees.

This talk is based on joint works with R. Meshulam, M. Rosenthal, Y. Rabinovich, I. Newman and Y. Peled.

16:10 - 16:40 *Coffee break*

16:40 - 17:10 **A Session of Short Announcements**

17:10 - 17:45 **Anna Kralin -** *The top* 10 *reasons it's awesome to work with Amos Fiat*

I will survey the top 10 reasons that it is awesome to work with Amos.

20:00 - **PARTY -** *at Amos Fiat's residence, please see local information page.*

Tuesday November 21, 2017

All day - Beit Hatefutzot, Zeevi Auditorium

9:00 - 9:40 **Howard Karloff -** *On the Optimal Bmi*

This papers studies both the optimal body-mass index (BMI) for individuals as well as alternative formulas for BMI (which is classically defined as mass/height^2). We use the NIH-AARP study of approximately $550,000$ people who entered the study at age 50-70, of whom approximately $116,000$ died within 12.9 years, the (minimum) duration of the study.

By treating BMI as a continuous variable and estimating its interaction effects with other demographic and health variables, we are able to compute for each study participant a "personalized optimal BMI," defined as the BMI which, based on the individual's covariates, is associated with the lowest risk of death. The averages of these personalized optimal BMIs across the subjects in our study are 25.7 for men and 26.3 for women. By contrast, traditional medical advice classifies BMI as "normal" if it lies in the interval $[18.5, 25)$, with anyone having a BMI 25 or higher classified as "overweight" or "obese." Our higher-than-traditional recommended BMI is more in line with recent papers suggesting that the optimal BMI lies in the interval $[25, 30)$.

Second, we study the question, "For which constant α is mass/height$^\alpha$ the best formula for BMI?", in the sense of giving the most accurate estimates of the risk of death. Many people have questioned the use of 2 as the exponent on height. Here we show that the "best" exponent alpha is less than 2 for both men and women, and that the difference between using the optimal exponent and using 2 is statistically significant. Interestingly, we cannot exclude the possibility that the optimal exponent is 1, which would yield the simple formula of mass/height.

9:40 - 10:20 **Alon Eden -** *Harnessing prices for efficiency*

Classic mechanism design settings deal with agents who bid simultaneously for some abstract service, and according to these bids, get served and charged. The mechanism designer desires to design a mechanism which guarantees to optimize some (maximization or minimization) goal function. This is typically done by offering the Òright" pricing and allocation algorithms which incentivize the agents to act in some coordinated manner. The designer assumes receiving all the bids at once and being able to force the agents to agree to the outcome. However, in many real life settings, these assumptions are simply not true; Not always can the designer get all the input in advance, and moreover, it is not common to have a bid-like structure to markets which sell a specific service.

We consider a more realistic market structure, where agents may arrive dynamically, in an unknown order, and the seller can only set prices on the services to be purchased. For instance, a city municipality might price parking slots in order to reduce parking congestion in city blocks and minimize walking distance to driversÕ destinations. This problem, among others, was considered in a line of work of papers by Amos (some co-authored with me). Surprisingly, many "strategic" variants of online problems were shown to have near-optimal pricing algorithms. In this talk, I'll give an overview of these results.

10:20 -11:00 **Yishay Mansour -** *Adaptive pricing, online learning and metric movement cost*

We start by considering a seller with an unlimited supply of a single good, who is faced with a stream of T buyers. Each buyer has a window of time in which she would like to purchase, and would buy at the lowest price in that window, provided that this price is lower than

her private value (and otherwise, would not buy at all). We relate this pricing problem to computing a strategy for a Multi-Arm Bandit problem, where there is a metric of the actions, and the algorithm has a cost which is a metric when it switches between actions. We show an optimal regret minimization for MAB for the metric setting, and relate it back to the pricing problem.

Based on joint works with Michal Feldman, Roi Livni, Tomer Koren and Aviv Zohar.

11:00- 11:30 *Coffee break*

11:30 - 12:10 **Yuval Rabani -** *LP Relaxations for Reordering Buffer Management*

We present LP relaxations for uniform and non-uniform reordering buffer management, and their application to the design and analysis of competitive online algorithms and polynomial time approximation algorithms for these problems.

12:10 - 12:40 **Elias Koutsoupias -** *The infinite server problem*

I will discuss some recent results on variants of the k-server problem, focusing on the infinite server problem. In this online problem, infinitely many servers reside initially at a particular point of a metric space and service a sequence of requests. Surprisingly, the major open question is whether there exists an online algorithm with bounded competitive ratio. The best known lower bound is only 3.146. I will discuss a tight connection between this problem and the (h, k)-server problem, in which an online algorithm with k servers competes against an offline algorithm with h servers. Of particular interest is the infinite server problem on the line, which turns out to be equivalent to the seemingly easier case in which all requests are in a fixed interval away from the original position of the servers. Unfortunately, classical approaches (work function algorithm, balancing algorithms, double coverage) fail even for this special case.

12:40 - 14:00 *Lunch* (Please see our list for a selection of restaurants on campus.)

14:00 - 14:40 **Noam Nisan -** *The communication complexity of cake-cutting*

We study classic cake-cutting problems, but in discrete models rather than using infinite-precision real values, specifically, focusing on their communication complexity. Using general discrete simulations of classical infinite-precision protocols (Robertson-Webb and moving-knife), we roughly partition the various fair-allocation problems into 3 classes: "easy" (constant number of rounds of logarithmic many bits), "medium" (poly-log total communication), and "hard".

Our main technical result concerns two of the "medium" problems (perfect allocation for 2 players and equitable allocation for any number of players) which we prove are not in the "easy" class. Our main open problem is to separate the "hard" from the "medium" classes.

Joint work with Simina Brânzei

14:40 - 15:20 **Valerie King -** *The communication cost of broadcasting*

It was long thought that to broadcast a message in a network where each node knows only its neighbors would require either a flooding of the network in which the message is sent along every edge or the construction of a spanning tree through the use of very long messages which communicate which nodes have already been seen. The number of bits of communication in either of these techniques is at least m where m is the number of edges, possibly much higher than the number of nodes n in the graph, which is clearly a lower bound.

A method developed in the study of streaming and also dynamic graph data structures for connectivity gives rise to a fairly simple Monte Carlo protocol to create a spanning tree (and a minimum spanning tree) in a distributed network with $\tilde{O}(n)$ communication, if the network is synchronous. My student Ali Mashreghi and I have recently found the first method for asynchronous networks which is sublinear in m (for large m).

Once the broadcast tree is build, the MST is not hard to build within these communication constraints.

I will talk about the ideas behind these protocols and related work investigating the complexity of this problem under varying assumptions.

15:20 - 15:50 *Coffee break*

15:50 - 16:30 **Kira Goldner -** *The Space Between Single- and Multi-Dimensional Mechanism Design*

Consider the problem of selling items to a unit-demand buyer. Most work on maximizing seller revenue considers either a setting that is single-dimensional, such as where the items are identical, or multi-dimensional, where the items are heterogenous. With respect to revenue-optimal mechanisms, these settings sit at extreme ends of a spectrum: from simple and fully characterized (single-dimensional) to complex and nebulous (multi-dimensional).

Starting in ÒThe FedEx Problem,Ó we identify a fascinating middle ground. Each buyer has a value v and a deadline d, and requires service by their deadline. We characterize the optimal mechanism for selling to a single buyer, finding it to be significantly more complex than the take-it-or-leave-it price of single-dimensional settings, but far more tractable than multi-dimensional settings.

In subsequent works, we have come to better understand a number of problems that further depict this space between single-dimensional and multi-dimensional, with respect to both (1) our ability to pin down the revenue-optimal mechanisms and (2) the degree of randomization (or menu-complexity) required.

16:30 - 17:10 **Uri Zwick -** *Selection from heaps, row-sorted matrices and $X + Y$ using soft heaps*

We use soft heaps to obtain simpler optimal algorithms for selecting the k-th smallest item, and the set of k smallest items, from a heap-ordered tree, from a collection of sorted lists, and from $X+Y$, where X and Y are two unsorted sets. Our results match, and in some ways extend and improve, classical results of Frederickson (1993) and Frederickson and Johnson (1982). In particular, for selecting the k-th smallest item, or the set of k smallest items, from a collection of m sorted lists we obtain a new optimal *"output-sensitive"* algorithm that performs only $O(m + \sum_{i=1}^{m} \log(k_i + 1))$ comparisons, where k_i is the number of items of the i-th list that belong to the overall set of k smallest items.

Joint work with Haim Kaplan, László Kozma and Or Zamir.