Algorithms:

10:00-11:00 – Shan Muthukrishnan (Rutgers University):

Heavy Hitters on Data Streams and Recent Variants

Abstract:

The data stream model focused on processing data with sublinear storage, and one of the traditional tasks in this model is identifying the heavy hitters (items that appear with overwhelming frequency, HHs). In this talk, I will provide an overview of HH algorithms, and focus on some of the recent variants: HHs seen from modern software defined networking (SDNs), HHs with very high dimensional data motivated by web analytics, HHs with pan-private guarantees, and other notions of heavy hitters including H-Index variants and multigraph versions. This problem continues to represent what we can do efficiently under many computing, space, communication, and other constraints.

11:30-12:00 – Katrina Ligett (HUJI):

Accuracy First: Selecting a Differential Privacy Level for Accuracy-Constrained Empirical Risk Minimization

Abstract:

Traditional approaches to differential privacy assume a fixed privacy requirement ϵ for a computation, and attempt to maximize the accuracy of the computation subject to the privacy constraint. As differential privacy is increasingly deployed in practical settings, it may often be that there is instead a fixed accuracy requirement for a given computation and the data analyst would like to maximize the privacy of the computation subject to the accuracy constraint. This raises the question of how to find and run a maximally private empirical risk minimizer subject to a given accuracy requirement. We propose a general "noise reduction" framework that can apply to a variety of private empirical risk minimization (ERM) algorithms, using them to "search" the space of privacy levels to find the empirically strongest one that meets the accuracy constraint, incurring only logarithmic overhead in the number of privacy levels searched. The privacy analysis of our algorithm leads naturally to a version of differential privacy where the privacy parameters are dependent on the data, which we term ex-post privacy, and which is related to the recently introduced notion of privacy odometers. We also give an ex-post privacy analysis of the classical AboveThreshold privacy tool, modifying it to allow for

queries chosen depending on the database. Finally, we apply our approach to two common objectives, regularized linear and logistic regression, and empirically compare our noise reduction methods to (i) inverting the theoretical utility guarantees of standard private ERM algorithms and (ii) a stronger, empirical baseline based on binary search.

12:00-12:30 – Liad Blumrosen (HUJI):

Selling Complementary Goods.

Abstract:

We consider a price competition between two sellers of perfect-complement goods. Each seller posts a price for the good it sells, but the demand is determined according to the sum of prices. This is a classic model by Cournot (1838), who showed that in this setting a monopoly that sells both goods is better for the society than two competing sellers.

We show that non-trivial pure Nash equilibria always exist in this game. We also quantify Cournot's observation with respect to both the optimal welfare and the monopoly revenue. We then prove a series of mostly negative results regarding the convergence of best response dynamics to equilibria in such games.

Joint work with Moshe Babaioff and Noam Nisan

12:30-13:00 – Shahar Dobzinsky (Weizmann):

Computational Efficiency Requires Simple Taxation

Abstract:

We characterize the communication complexity of truthful mechanisms. Our departure point is the well known taxation principle. The taxation principle asserts that every truthful mechanism can be interpreted as follows: every player is presented with a menu that consists of a price for each bundle (the prices depend only on the valuations of the other players). Each player is allocated a bundle that maximizes his profit according to this menu. We define the taxation complexity of a truthful mechanism to be the logarithm of the maximum number of menus that may be presented to a player.

Our main finding is that in general the taxation complexity essentially equals the communication complexity. The proof consists of two main steps. First, we prove that for rich enough domains the taxation complexity is at most the communication

complexity. We then show that the taxation complexity is much smaller than the communication complexity only in ``pathological'' cases and provide a formal description of these extreme cases. We will also discuss some applications of our characterization.

14:30-15:00 – Yossi Azar (TAU):

Matching with delays

Abstract:

Traditional online algorithms deal typically with (1) inputs over a sequence with a goal of minimizing the total cost; (2) inputs over a time where the goal is to minimize some function of the time. Following the work of Emek, Kutten and Wattenhofer (STOC'16) we consider problems where the goal is to minimize the cost plus the waiting time (delay) until requests are served. In particular we deal with online min-cost matching with delays (matching players in game platforms) and online two sided min-cost matching with delays (matching costumers to drivers in transportation platforms such as Uber/Lyft). We provide logarithmic competitive algorithms for these problems.
Based on joint work with Chiplunkar, Kaplan (SODA' 17), Ashlagi, Charikar, Chiplunkar, Geri, Kaplan, Makhijani, Wang, Wattenhofer (APPROX' 17).

15:00-15:30 – Laszlo Kozma (TAU):

Selection from heaps, row-sorted matrices and X + Y using soft heaps

Abstract:

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).

Joint work with Haim Kaplan, Or Zamir, and Uri Zwick.

16:00-1630 – Eden Chlamtac (BGU):

The log-density method

Abstract:

Hardness of approximation for a given problem is proven by showing that it is hard to decide whether a given input belongs to some input family for which the optimum is at least C, or to some input family for which the optimum is at most S (for some S < C). In such a case, we know that no efficient algorithm can approximate the problem to within a factor greater than S/C.

For many problems, there is still a large gap between the best known approximations, and the strongest known hardness results. In recent years, we have seen that for some of these problems it appears to be hard to distinguish between a fully random input, and a random input in which there is a good hidden solution, when that solution is similar to the original input. By "similar", we mean that both the input and the hidden solution have the same "log-density", a problem-specific measure that determines statistical properties of both the input and the hidden solution.

This approach has allowed us to identify families of random inputs for various problems which we conjecture to be hard, and design algorithms inspired by these random input families which match the conjectured hardness even for general inputs. These problems range from finding dense subgraphs to expansion problems to optimisation variants of Label Cover. In some cases, our conjectured lower bounds are also matched by integrality gaps for powerful convex relaxations.

Based on various joint results with Aditya Bhaskara, Moses Charikar, Michael Dinitz, Uriel Feige, Robert Krauthgamer, Yury Makarychev, Pasin Manurangsi, Dana Moshkovitz, and Aravindan Vijayaraghavan.

16:30-17:30 – Bernhard Haeupler (Carnegie Mellon University):

Distributed Optimization Algorithms via Low-Congestion Shortcuts

Abstract:

How fast a non-local distributed optimization problem can be solved in a given network depends in a highly non-trivial manner on the topology of the network. This talk will introduce a simple graph structure, called low-congestion shortcuts, which often tightly characterize and capture this dependency. Low-congestion shortcuts furthermore make it easy to design near optimal distributed algorithms for a wide variety of problems. For example, this leads to MST and approximate min-cut and shortest-path algorithms which require only a near linear number of messages and

have the optimal $\tilde{O}(\sqrt{n} + D)$ running times on worst-case network topologies while also achieving a near instance-optimal $\tilde{O}(D)$ running times on planar, low-genus, low-treewidth and other non-pathological network topologies.