Program

Wednesday June 15, 2022

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

9:00 - 9:15 **Opening Remarks**

9:15 - 9:30 **In Memoriam Benny Chor**

9:30 - 10:00 **Adrian Vladu -** *Decomposable Submodular Function Minimization via Maximum Flow*

We bridge discrete and continuous optimization approaches for minimizing sums of submodular functions. We provide improved running times for this problem by reducing it to a number of calls to a maximum flow oracle. When each function in the decomposition acts on $O(1)$ elements of the ground set, our running time is nearly linear.

We achieve this by providing a simple iterative method which can optimize to high precision any convex function defined on the submodular base polytope, provided we can efficiently minimize it on a certain cut polytope. We solve this minimization problem by lifting the solutions of a graph cut problem, which we obtain via a new efficient combinatorial reduction to maximum flow. This reduction is of independent interest and implies previously unknown bounds for the parametric minimum s,t-cut problem in multiple settings.

Joint work with Kyriakos Axiotis, Adam Karczmarz, Anish Mukherjee, and Piotr Sankowski.

10:00 - 10:30 **Simon Apers -** *Cut query algorithms with star contraction*

We study the complexity of determining the edge connectivity of a simple graph with cut queries. We show that (i) there is a bounded-error randomized algorithm that computes edge connectivity with $O(n)$ cut queries, and (ii) there is a bounded-error quantum algorithm that computes edge connectivity with $\tilde{O}(\sqrt{n})$ cut queries. To prove these results we introduce a new technique, called *star contraction*, to randomly contract edges of a graph while preserving non-trivial minimum cuts.

The $O(n)$ bound from item (i) was not known even for the simpler problem of connectivity, and is tight under the reasonable conjecture that the randomized communication complexity of connectivity is $\Omega(n \log n)$. The quantum algorithm from item (ii) gives a nearly-quadratic separation with the randomized complexity, and can alternatively be viewed as computing the edge connectivity of a simple graph with $\tilde{O}(\sqrt{n})$ matrix-vector multiplication queries to its adjacency matrix.

10:30 - 11:00 *Coffee break*

11:00 - 11:30 **Yaniv Sadeh -** *Weighted Load Balancing in TCAMs: Rulesets Minimizing Algorithms and Bounds*

Traffic splitting is a required functionality in networks, for example for load balancing over paths or servers, or by the source's access restrictions. The capacities of the servers (or the number of users with particular access restrictions) determine the sizes of the parts into which traffic should be split. A recent approach implements traffic splitting within the ternary content addressable memory (TCAM), which is often available in switches. It is important to reduce the amount of memory allocated for this task since TCAMs are power consuming.

In this talk I will present the model of the problem arising from the hardware's behavior, and focus on partitioning traffic using prefix-matching TCAMs (LPM). I will present two known optimal algorithms for computing the most compact LPM TCAMs, and review our recent results for finding similar partitions that are simpler to represent (compared to the desired partition) in the case when the memory is restricted to a predefined budget.

11:30 - 12:00 **Gregory Kucherov -** *Count-Min sketch: some new results*

Count-Min sketch is a hash-based data structure to represent a dynamically changing associative array of counters. We analyse the counting version of Count-Min under a stronger update rule known as conservative update, assuming the uniform distribution of input keys. We show that the accuracy of conservative update strategy undergoes a phase transition, depending on the number of distinct keys in the input as a fraction of the size of the Count-Min array. We prove that below the threshold, the relative error is asymptotically $o(1)$ (as opposed to the regular Count-Min strategy), whereas above the threshold, the relative error is $\Theta(1)$. The threshold corresponds to the peelability threshold of random k-uniform hypergraphs. We demonstrate that even for small number of keys, peelability of the underlying hypergraph is a crucial property to ensure the $o(1)$ error.

12:00 - 13:30 **Guy Even -** *An Extendable Data Structure for Incremental Stable Perfect Hashing*

We consider the problem of dynamically assigning n_t elements unique *hashcodes* in the range $[(1 + o(1))n_t]$. This problem is known as *perfect hashing* and is considered a fundamental building block in the design of more involved data structures.

The challenge we address is that of designing a data structure that meets several, seemingly opposing, requirements: (1) the range and the space of the data structure must be, at all times, proportional to the current cardinality n_t of the input set, and

(2) the hashcodes must be stable while an element is continuously in the set.

A simple argument shows that these two desiderata are impossible to achieve when arbitrary deletions and insertions are allowed.

In this paper, we show that one can achieve these requirements when only insertions occur and, more generally, when the hashcode range and the space are allowed to grow as a function of the maximum cardinality of the set until time t. The data structure executes all operations in worst case constant time with high probability and requires space that is within a constant factor of the lower bound.

Applications include: (1) A hash table design that does not need to move elements as its size grows.

(2) A compact dynamic dictionary with constant time operations requiring only $O(\log \log n)$ bits per element above the lower bound. This construction applies both to the extendable and non-extendable settings.

Joint work with Ioana Bercea. Accepted to STOC 2022.

12:30 - 14:00 *Lunch*

14:00 - 14:30 **Yishay Mansour -** *Cooperation and delay in reinforcement learning*

Much of the research in reinforcement learning focuses on single strategic agent optimizing its long term payoff. In this talk we will explore the possibilities of multiple agents to collaborate and improve their individual performance. We will also explore a related issue of delay when agent observe the outcomes with some delay, which is related to the communication between the agents. In this talk I will outline this research direction and focus on a few recent results both on cooperation and delay.

14:30 - 15:00 **Rotem Oshman -** *Massively Parallel Computing in a Heterogenous Memory Regime*

Massively-parallel graph algorithms have received extensive attention over the past decade, with research focusing on three memory regimes for the machines: the superlinear regime, where each machine has $n^{1+\gamma}$ memory for some $\gamma < 1$; the near-linear regime, with memory $O(npolylog(n))$; and the sublinear regime, with memory $O(n^{1-\gamma})$. The sublinear regime is the most desirable in practice, but some conditional hardness results point towards its limitations.

In this work we study a *heterogeneous* model, in which the memory size of the machines varies. We focus our attention on the heterogeneous setting created by adding a single near-linear machine to the sublinear MPC model, and show that even a single large machine suffices to circumvent most of the conditional hardness results for the sublinear regime. Specifically, in this setting, for graphs with n vertices and m edges, we give:

- an MST algorithm that runs in $O(\log\log(m/n))$ rounds;
- an algorithm that constructs an $O(k)$-spanner of size $O(n^{1+1/k})$ in $O(1)$ rounds; and
- a maximal-matching algorithm that runs in $O(\sqrt{\log(m/n)}\log\log(m/n))$ rounds.

We also point out that the best known near-linear MPC algorithms for several other graph problems which are conjectured to be hard in the sublinear regime (minimum cut, maximal independent set, and vertex coloring) can easily be transformed to work in the heterogeneous MPC model with a single near-linear machine, while retaining their original round complexity in the near-linear regime. If the large machine is allowed to have *superlinear* memory, all of the problems above can be solved in $O(1)$ rounds.

15:00 - 15:30 **Robi Krauthgamer -** *Streaming Algorithms for Geometric Steiner Forest*

I will discuss the Steiner forest problem in the Euclidean plane, where the input is a multiset of points, partitioned into k color classes, and the goal is to find a minimum-cost Euclidean graph G such that every color class is connected. We study this problem in dynamic streams, where the input is provided by a stream of insertions and deletions of colored points from the discrete grid $[\Delta]^2$.

Our main result is a single-pass streaming algorithm that uses $poly(k \log \Delta)$ space and time, and estimates the cost of an optimal Steiner forest solution within ratio arbitrarily close to the famous Euclidean Steiner ratio α_2 (currently $1.1547 \le \alpha_2 \le 1.214$). Our approach relies on a novel combination of streaming techniques, like sampling and linear sketching, with the classical dynamic-programming framework for geometric optimization problems, which usually requires large memory and has so far not been applied in the streaming setting.

Time permitting, I will also discuss possible directions for future work.

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

15:30 - 16:00 **Ami Paz -** *Beyond Worst-Case Analysis of Dynamic Networks*

Smoothed analysis is a framework suggested for mediating gaps between worst-case and average-case complexities. In the case of dynamic networks, this technique aims to explain the gaps between real-world networks that function well despite being dynamic, and the strong theoretical lower bounds for arbitrary networks. In this talk, I will present our recent applications of smoothed analysis to dynamic networks. Our work studies different models of adversarial behaviour, and concerns problems such as information propagation and load balancing.

Based on joint works with Seth Gilbert, Uri Meir, and Gregory Schwartzman.

16:00 - 16:30 *Coffee break*

16:30 - 17:00 **Sophie Laplante -** *Certificate games*

Consider the following two-player game for a Boolean function. The players are given inputs with different function values and are asked to output some index i such that $x_i \neq y_i$. (This is the well-known Karchmer Wigderson relation.) We study the setting where the players cannot communicate. We define Certificate Game complexity $CG(f)$ to be the multiplicative inverse of the probability of winning this game, for the best strategy, in the worst case.

Certificate game complexity, in its public-coin, shared entanglement, and non-signaling variants, all sit between certificate complexity $C(f)$ and randomized query complexity $R(f)$, above, and (fractional) block sensitivity, below. For total functions, these measures are all polynomially related, but we show an exponential separation between R and CG in the public coin model, for a partial function.

The Tribes (AND composed with OR) function illustrates the power of shared randomness, giving a quadratic separation between the private coin and public coin variants of the certificate games. This also separates R and the public coin variant.

The quantum measure highlights an interesting and surprising difference between classical and quantum query models. Whereas the public coin certificate game complexity is bounded from above by randomized query complexity, the quantum certificate game complexity can be quadratically larger than quantum query complexity (for the OR function). This says that it can be quadratically easier to solve the OR function than it is for two players with inputs with different function values to identify a specific query where their computations diverge. The proof uses non-signaling, a property from quantum information that holds for all strategies with shared entanglement.

This is joint work with Sourav Chakraborty, Anna Gal, Rajat Mittal and Anupa Sunny.

17:00 - 17:30 **Pierre Meyer -** *Sublinear Secure Computation from New Assumptions*

Secure computation enables mutually distrusting parties to jointly compute a function on their secret inputs, while revealing nothing beyond the function output. A long-running challenge is understanding the required communication complexity of such protocols—in particular, when communication can be *sublinear* in the circuit representation size of the desired function. For certain functions, such as Private Information Retrieval (PIR), this question extends to even sublinearity in the input size.

We develop new techniques expanding the set of computational assumptions for sublinear communication in both settings:

- **Circuit size.** We present sublinear-communication protocols for secure evaluation of general layered circuits, given any 2-round rate-1 batch oblivious transfer (OT) protocol with a particular "decomposability" property. In particular, this condition can be shown to hold for the recent batch OT protocols of (Brakerski et al. Eurocrypt 2022), in turn yielding a new sublinear secure computation feasibility: from Quadratic Residuosity (QR) together with polynomial-noise-rate Learning Parity with Noise (LPN).

  Our approach constitutes a departure from existing paths toward sublinear secure computation, all based on fully homomorphic encryption or homomorphic secret sharing.
- **Input size.** We construct single-server PIR based on the Computational Diffie-Hellman (CDH) assumption, with *polylogarithmic* communication in the database input size n. Previous constructions from CDH required communication $\Omega(n)$. In hindsight, our construction comprises of a relatively simple combination of existing tools from the literature.

In this talk, we focus on the first setting.

17:30 - 18:00 **Benny Applebaum -** *Secret Sharing, Slice Formulas, and Monotone Real Circuits*

Suppose that you wish to compute an arbitrary monotone boolean function f over n-bit inputs by a formula F over "slice gates" of arbitrary fan-in. Here a k-slice gate takes arbitrary values on inputs of Hamming weight k, rejects lighter inputs, and accepts heavier inputs. How small can the formula F be? Is it possible to beat the naive 2^n size? Can we achieve a sub-exponential size of $2^{n^{1-\epsilon}}$? We relate this question to recent advances in general secret-sharing schemes and to problems in the domain of Real Communication Complexity. By using these connections, we prove that every monotone function can be computed by a formula over slices of size 1.5^n, and that some functions require a formula of size $2^{\Omega(n/log^2 n)}$.

Based on joint work with Amos Beimel, Oded Nir, Naty Peter, and Toniann Pitassi.

19:00 - *Social Event*

Thursday June 16, 2022

9:00 - 9:30 **Amos Korman -** *On the Search Efficiency of Common Animal Movement Patterns*

Lévy walks are random walk processes whose step length follows a long-tailed power-law distribution. Due to their abundance as movement patterns of biological organisms, significant theoretical efforts have been devoted to identify the foraging circumstances that would make such patterns advantageous [Viswanathan et al. Nature, 1999]. Despite numerous attempts, however, no analytical argument was found indicating that Lévy walks are, in any manner, preferable strategies in higher dimensions than one. In this work, we show that the optimality of inverse-square Lévy walks in two dimensions becomes striking when targets are sparse and unpredictable in size, and when detection is weak. Specifically, we prove that under the *intermittent model*, in which detection is possible between steps, but not while moving ballistically, this strategy optimally finds sparse targets of any size and shape. That is, in a square torus of area n, and assuming that the detection radius is normalized to 1, the strategy finds any connected set of diameter D in $\tilde{O}(n/D)$ expected time, whereas $\Omega(n/D)$ is an unconditional lower bound on the expected time, that holds even when assuming that the shape and size of the target are known. Furthermore, this particular Lévy process stands in stark contrast to any other intermittent Lévy walk, which fails to efficiently find either large targets or small ones. Our results shed new light on the Lévy foraging hypothesis and are thus expected to affect future experiments on animals performingLévy walks.

- This talk is based on a joint work with Brieuc Guinard, which appeared at *Science Advances* in 2021.

9:30 - 10:00 **Robin Vacus -** *Early Adapting to Trends: Self-Stabilizing Information Spread using Passive Communication*

How to efficiently and reliably spread information in a system is one of the most fundamental problems in distributed computing as well as in the animal world. Inspired by biological scenarios, we consider the self-stabilizing *bit-dissemination* problem, in the extremely constrained model of *passive communication*. We prove that it can be solved in a number of rounds poly-logarithmic in the number of agents, with high probability. Our proposed protocol has a natural appeal as it is based on letting agents estimate the current tendency direction of the dynamics, and then adapt to the emerging trend.

10:00 - 10:30 **Janna Burman -** *Time-Optimal Self-Stabilizing Leader Election in Population Protocols*

We consider the standard population protocol model, where (*a priori*) indistinguishable and anonymous agents interact in pairs according to uniformly random scheduling. The *self-stabilizing leader election* problem requires the protocol to converge on a single leader agent from *any* possible initial configuration. We initiate the study of time complexity of population protocols solving this problem in its original setting: with probability 1, in a complete communication graph. The only previously known protocol by Cai, Izumi, and Wada [Theor. Comput. Syst. 50] runs in expected parallel time $\Theta(n^2)$ and has the optimal number of n states in a population of n agents. The existing protocol has the additional property that it becomes silent, i.e., the agents' states eventually stop changing.

Observing that any silent protocol solving self-stabilizing leader election requires $\Omega(n)$ expected parallel time, we introduce a silent protocol that uses optimal $O(n)$ parallel time and states. Without any silence constraints, we show that it is possible to solve self-stabilizing leader election in asymptotically optimal expected parallel time of $O(\log n)$, but using at least exponential states (a quasipolynomial number of bits). All of our protocols (and also that of Cai et al.) work by solving the more difficult *ranking* problem: assigning agents the ranks $1, \ldots, n$.

10:30 - 11:00 *Coffee break*

11:00 - 11:30 **Nati Linial -** *Graphs as geometric objects*

It may seem quite obvious that graphs carry a lot of geometric structure. Don't we learn in algorithm classes how to solve all-pairs-shortest-paths, minimum spanning trees etc.? However, in this talk I will try to impress on you the idea that there is much more to be discovered here. For example: Let $G = (V, E)$ be a graph and let w be a mapping from E to the positive reals. Let us restrict ourselves to the case where no ties occur and so for every two distinct vertices u, v there is a unique w-shortest path that connects them (uniqueness holds with probability 1). Let w' be another set of positive weights on E. We consider w and w' equivalent if they define the same system of shortest paths. Question: Given a graph G how many such different equivalence classes can it have at least/at most? Also, we say that a path system on G is consistent if it is closed under taking subpaths. Question: Does every consistent path system on G come from a function w as above? The answer is an emphatic NO.

The new results are from joint work with my student Daniel Cizma.

11:30 - 12:00 **Pierre Fraigniaud -** *Distributed Speedup Theorems*

The talk will present the round-reduction technique used for establishing lower bounds in the context of distributed computing in networks, as introduced in the speedup theorem by Brandt [PODC 2019], and used by Balliu et al. [FOCS 2019] for establishing lower bounds for maximal matchings and maximal independent sets in the LOCAL model. Next, the talk will show how to extend the round-reduction technique to other models of distributed computing, and in particular how to establish a speedup theorem in the context of asynchronous distributed shared-memory computing with crash-prone processes.

Joint work with Ami Paz (CNRS) and Sergio Rajsbaum (UNAM).

12:00 - 12:30 **Vianney Perchet -** *An algorithmic solution to the Blotto game using multi-marginal couplings*

We describe an efficient algorithm to compute solutions for the general two-player Blotto game on n battlefields with heterogeneous values. While explicit constructions for such solutions have been limited to specific, largely symmetric or homogeneous, setups, this algorithmic resolution covers the most general situation to date: value-asymmetric game with asymmetric budget. The proposed algorithm rests on recent theoretical advances regarding Sinkhorn iterations for matrix and tensor scaling. An important case which had been out of reach of previous attempts is that of heterogeneous but symmetric battlefield values with asymmetric budget. In this case, the Blotto game is constant-sum so optimal solutions exist, and our algorithm samples from an ϵ-optimal solution in time $O(n^2 + \epsilon^{-4})$, independently of budgets and battlefield values. In the case of asymmetric values where optimal solutions need not exist but Nash equilibria do, our algorithm samples from an ϵ-Nash equilibrium with similar complexity but where implicit constants depend on various parameters of the game such as battlefield values.

12:30 - 14:00 *Lunch*

14:00 - 14:30 **Uri Feige -** *On maximin fair allocation of indivisible items to three agents*

We consider the allocation of indivisible goods to n agents with additive valuation functions, in a setting with no money. A natural fairness requirement is that every agent will get a bundle that she values at least as high as her maximin share (MMS), namely, the value she could guarantee to herself if she were to partition the items into n bundles of her choice, and receive the least valuable of these bundles. Unfortunately, for three or more agents, there are allocation instances in which no MMS allocation exists.

For the case of three agents, we present sufficient conditions for an MMS allocation to exist. We also consider ρ_3, the largest approximation ratio for which there always is an allocation giving each of the three agents at least a ρ_3 fraction of her MMS. We show that $11/12 \leq \rho_3 \leq 39/40$, and discuss the prospects of pinning down the exact value of ρ_3.

Based on joint works with Moshe Babaioff, Alexey Norkin, Ariel Sapir, Laliv Tauber.

14:30 - 15:00 **Michal Feldman -** *Prophet and Secretary Online Algorithms for Graph Matching*

A common tension in market scenarios is choosing the right timing to commit to a decision. This tension is captured by the mathematical literature of optimal stopping theory, within two models that have become to be known as the secretary problem and the prophet inequality. In these models, a sequence of originally-unknown values arrive one by one. Upon arrival, the online algorithm observes the value and should decide whether or not to accept it. In secretary settings, the value sequence is arbitrary, but the values arrive in a uniformly random order. In prophet settings, every value is drawn from a known probability distribution, but the arrival order is arbitrary.

In this talk I will shortly review the basic settings of secretary and prophet, and will present extensions to graph matching. These include matching in bipartite graphs with 1-sided vertex arrival and matching in general graphs (with general arrival).

Based on joint work with Tomer Ezra, Nick Gravin and Zhihao Tang.

15:00 - 15:30 **Liat Peterfreund -** *Querying incomplete numerical data*

Queries with aggregation and arithmetic operations, as well as incomplete data, are common in real-world databases, but we lack a good understanding of how they should interact. On one hand, systems based on SQL provide ad-hoc rules for numerical nulls, on the other, theoretical research largely concentrates on the standard notions of certain and possible answers which in the presence of numerical attributes and aggregates are often meaningless. In this work, we define a principled compositional framework for databases with numerical nulls and answering queries with arithmetic and aggregations over them. We assume that missing values are given by probability distributions associated with marked nulls which yields a model of probabilistic bag databases. We concentrate on queries that resemble standard SQL with arithmetic and aggregation and show that they are measurable, and that their outputs have a finite representation. Moreover, since the classical forms of answers provide little information in the numerical setting, we look at the probability that numerical values in output tuples belong to specific intervals. Even though their exact computation is intractable, we show efficient approximation algorithms to compute such probabilities.

The talk is based on joint work with Marco Console and Leonid Libkin.

15:30 - 16:00 *Coffee break*

16:00 - 16:30 **Allan Borodin -** *Prophet Matching in the Probe-Commit Model*

We consider the online bipartite stochastic matching problem with known i.d. (independently distributed) online vertex arrivals. In this problem, when an online vertex arrives, its weighted edges must be probed (queried) to determine if they exist, based on known edge probabilities. Our algorithms operate in the probe-commit model, in that if a probed edge exists, it must be used in the matching. Additionally, each online node has a downward-closed probing constraint on its adjacent edges which indicates which sequences of edge probes are allowable.

Our setting generalizes the commonly studied patience (or time-out) constraint which limits the number of probes that can be made to an online node's adjacent edges. Most notably, this includes non-uniform edge probing costs and a knapsack/budget constraint. We extend a recently introduced configuration LP to the known i.d. setting, and also provide the first proof that it is a relaxation of an optimal offline probing algorithm (the offline adaptive benchmark). Using this LP, we establish the following competitive ratio results against the offline adaptive benchmark:

1. A tight $\frac{1}{2}$ ratio when the arrival ordering π is chosen adversarially.
2. A $1-1/e$ ratio when the arrival ordering π is chosen u.a.r. (uniformly at random).

If π is generated adversarially, we generalize the prophet inequality matching problem. If π is u.a.r., we generalize the prophet secretary matching problem. Both results improve upon the previous best competitive ratio of 0.46 in the more restricted known i.i.d. (independent and identically distributed) arrival model against the standard offline adaptive benchmark due to Brubach et al. We are the first to study the prophet secretary matching problem in the context of probing, and our $1-1/e$ ratio matches the best known result without probing due to Ehsani et al. This result also applies to the unconstrained bipartite matching probe-commit problem, where we match the best known result due to Gamlath et al.

16:30 - 17:00 **David Naori -** *Online Weighted Matching with a Sample*

We study the greedy-based online algorithm for edge-weighted matching with (one-sided) vertex arrivals in bipartite graphs, and edge arrivals in general graphs. This algorithm was first studied more than a decade ago by Korula and Pál for the bipartite case in the random-order model. While the weighted bipartite matching problem is solved in the random-order model, this is not the case in recent and exciting online models in which the online player is provided with a sample, and the arrival order is adversarial. The greedy-based algorithm is arguably the most natural and practical algorithm to be applied in these models. Despite its simplicity and appeal, and despite being studied in multiple works, the greedy-based algorithm was not fully understood in any of the studied online models, and its actual performance remained an open question for more than a decade.

We provide a thorough analysis of the greedy-based algorithm in several online models. For vertex arrivals in bipartite graphs, we characterize the exact competitive-ratio of this algorithm in the random-order model, for any arrival order of the vertices subsequent to the sampling phase (adversarial and random orders in particular). We use it to derive tight analysis in the recent adversarial-order model with a sample (AOS model) for any sample size, providing the first result in this model beyond the simple secretary problem. Then, we generalize and strengthen the black box method of converting results in the random-order model to single-sample prophet inequalities, and use it to derive the state-of-the-art single-sample prophet inequality for the problem. Finally, we use our new techniques to analyze the greedy-based algorithm for edge arrivals in general graphs and derive results in all the mentioned online models. In this case as well, we improve upon the state-of-the-art single-sample prophet inequality.

17:00 - 17:30 **Yossi Azar -** *Flow Time Scheduling with Uncertain Processing Time*

We consider the problem of online scheduling on a single machine to minimize unweighted and weighted flow time. The existing algorithms for these problems require exact knowledge of the processing time of each job. This assumption is crucial, as even a slight perturbation of the processing time would lead to polynomial competitive ratio. However, this assumption very rarely holds in real-life scenarios. We present a competitive algorithm (the competitive ratio is a function of the distortion) for unweighted flow time that do not require knowledge of the distortion in advance. For the weighted flow time we present competitive algorithms but, in this case, we need to know (an upper bound on) the distortion in advance.

This is joint work with Stefano Leonardi and Noam Touitou based on papers appear in STOC 21 and SODA 2022.

17:30 - 18:00 **Boaz Patt-Shamir -** *Competitive Disengagement*

Motivated by placement of jobs in physical machines, we introduce and analyze the problem of online recoloring, or online disengagement. In this problem, we are given a set of n weighted vertices and an initial k-coloring of the vertices (vertices represent jobs, and colors represent physical machines). Edges, representing conflicts between jobs, arrive in an online fashion. Following each edge arrival, the algorithm must output a proper k-coloring of the vertices. Each recoloring of a vertex costs the weight of that vertex, and the goal is to minimize the overall cost.

We consider a couple of polynomially-solvable coloring variants. First, for 2-coloring bipartite graphs we present an $O(\log n)$-competitive deterministic algorithm and an $\Omega(\log n)$ lower bound on the competitive ratio of randomized algorithms. Second, for $(\Delta + 1)$-coloring we present tight bounds of $\Theta(\Delta)$ and $\Theta(\log \Delta)$ on the competitive ratios of deterministic and randomized algorithms, respectively (where Δ denotes the maximum degree). All algorithms are applicable in the weighted case, and all lower bounds hold in the restricted unweighted case.

Joint work with Yossi Azar, Chay Machluf and Noam Touitou.