Program

Wednesday October 3, 2018

8:30 - 8:50 **Gathering and Registration**

8:50 - 9:00 **Opening Remarks**

9:00 - 9:30 **Amos Korman -** *Cooperative Search Games: Symmetric Equilibria, Robustness, and Price of Anarchy*

Assume that a treasure is placed in one of M boxes according to a known distribution and that k searchers are searching for it in parallel during T rounds. We study the question of how to incentivize selfish players so that the success probability, i.e., the probability that at least one player finds the treasure, would be maximized. We identify a reward policy that yields the best possible price of anarchy, which turns out to be precisely $(1-(1-1/k)^k)^{-1}$. This exclusive policy yields high levels of competition, as it rewards a player (with a reward of 1) only if it finds the treasure strictly before others.

Together with an appropriate reward policy, a central entity can suggest players to play particular strategies at equilibrium. For such purposes, we advocate the use of symmetric equilibria. Besides being fair, we argue that symmetric equilibria can also become highly robust to crashes of players. Indeed, in many cases, despite the fact that some small fraction of players crash (or refuse to participate), symmetric equilibria remain efficient in terms of their group performances and, at the same time, serve as approximate equilibria. We show that this principle holds for a class of games, which we denote by socially monotone games. This applies in particular for our search game, assuming the natural sharing policy, in which all players that simultaneously find the treasure for the first time equally share the reward of 1. For the exclusive policy, however, this does not apply, yet we show that the symmetric strategy with the highest success probability is both at equilibrium and, under mild assumptions, robust to crashes. Finally, for the sharing policy, we present an algorithm to construct an efficient and robust approximate equilibrium.

Joint work with Yoav Rodeh.

9:30 - 10:00 **Avi Cohen -** *Fault-Prone Hotelling Games*

The n-player Hotelling game calls for each player to choose a point on the line segment, so as to maximize the size of his Voronoi cell. This work studies fault-tolerant versions of the Hotelling game. Two fault models are studied. The first assumes that the environment is prone to failure: disconnections occur at a random points on the line, each splitting it into two separate segments and modifying each player's Voronoi cell accordingly. The second fault model assumes the players are prone to failure: each player is removed from the game with i.i.d. probability, changing the payoffs of the remaining players accordingly. For each variant of the game considered, an explicit construction of all existing Nash equilibria is provided. Somewhat surprisingly, the potential occurrence of failures can have a stabilizing effect on the game.

Joint work with Chen Avin, Zvi Lotker and David Peleg.

10:00 - 10:30 **Miklos Santha -** *Strategies for quantum races*

We initiate the study of quantum races, games where two or more quantum computers compete to solve a computational problem. While the problem of dueling algorithms has been studied for classical deterministic algorithms before, the quantum case presents additional sources of uncertainty for the players. The foremost among these is that players do not know if they have solved the problem until they measure their quantum state. This question of "when to measure?" presents a very interesting strategic problem. We develop a game-theoretic model of a multiplayer quantum race, and find an approximate Nash equilibrium where all players play the same strategy. In the two-party case, we further show that this strategy is nearly optimal in terms of payoff among all symmetric Nash equilibria. One application of our results is to the stability of the Bitcoin protocol when mining is done by quantum computers.

The technical part of the talk is purely classical. It is joint work with Troy Lee and Maharshi Ray.

10:30 - 11:00 *Coffee break*

11:00 - 11:30 **Amnon Ta-Shma -** *Memory augmented Markovian walks and Explicit Parity samplers giving almost optimal binary codes*

I will show an explicit construction of a binary error correcting code with relative distance $\frac{1-\epsilon}{2}$ and relative rate $\epsilon^{2+o(1)}$. This comes close to the Gilbert-Varshamov bound that shows such codes with rate ϵ^2 exist, and the LP lower bound that shows rate $\frac{\epsilon^2}{\log \frac{1}{\epsilon}}$ is necessary. Previous explicit constructions had rate about ϵ^3, and this is the first explicit construction to get that close to the Gilbert-Varshamov bound. The main tool we use is "Parity Samplers". A parity sampler is a collection of sets $\{S_i \subset \Lambda\}$ with the property that for every "test" $A \subset \Lambda$ of a given constant density ϵ_0, the probability a set S_i from the collection falls into the test set A an *even* number of times is about half. A *sparse* parity sampler immediately implies a good code with distance close to $\frac{1}{2}$. The complete t-complex of all sequences of cardinality t is a good parity sampler, but with too many sets in the collection. Rozenman and Wigderson, and independently Alon, used random walks over expanders to explicitly construct sparse parity samplers, and their construction implies explicit codes with relative rate ϵ^4. In the last part of the talk I will explain how one can get better explicit parity samplers (and therefore also better explicit codes) using a variant of the zig-zag product. In the random walk sampler, there exist many sets with substantial overlap. One way to look at the zig-zag product is that it takes a sub-collection of the random walk sampler, and this sub-collection has a smaller overlap between sets in the collection. The zig-zag product achieves that by keeping a small internal state. I will show that by enlarging the internal state one can further reduce the overlap, and as a result improve the quality of the parity sampler. One may view this process as a memory augmented Markovian process.

11:30 - 12:00 **Inbal Livni Navon -** *List decoding using double samplers*

Samplers are bipartite graphs in which large sets of vertices are expanding. Sampler graphs are used in the ABNNR code construction to amplify the distance of an error correcting code. This construction has a unique decoding algorithm, but it is not known to be efficiently list decodable.

In this talk, I will introduce "double samplers", which extend sampler graphs to three layers, and show a polynomial time algorithm which list decodes the ABNNR code construction on a double sampler. Double samplers can be constructed from high dimensional expanders, and currently it is the only known construction.

Joint work with Irit Dinur, Prahladh Harsha, Tali Kaufman and Amnon Ta-Shma.

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

14:00 - 14:30 **Martin Golumbic -** *Total Coloring of Rooted Path Graphs*

A *total coloring* of a graph is an assignment of colors to both its vertices and edges so that adjacent or incident elements acquire distinct colors. In this note, we give a simple greedy algorithm to totally color a rooted path graph G with at most $\Delta(G) + 2$ colors, where $\Delta(G)$ is the maximum vertex degree of G. Our algorithm is inspired by a method by Bojarshinov [2001] for interval graphs and provides a new proof that the Total Coloring Conjecture, posed independently by Behzad [1965] and Vizing [1968], holds for rooted path graphs. In the process, we also prove a useful property of greedy neighborhood coloring for chordal graphs.

14:30 - 15:00 **Christophe Paul -** *A polynomial Turing kernel to compute the cut-width of semi-complete digraphs*

The existence of a fixed-parameterized algorithm is equivalent to the existence of a kernelization algorithm, that is a polynomial time that given a parameterized instance compute an equivalent instance, so-called kernel, the size of which is bounded by the parameter. However, recent lower bound techniques allows to establish, under some standard complexity hypothesis, the non-existence of a polynomial size kernel. This raised the question of designing polynomial time algorithms that instead of reducing the input instance to one kernel, computes polynomially many kernel from which the original instance can be solved. This is known under the concept of Turing-kernlization. In this talk, we provide such an algorithm for the parameterized cut-width problem in semi-complete digraphs. To date, this provides one of the very few examples of non-trivial Turing kernelization.

This work was done with Florian Barbero (Montpellier University) and Michal Pilipczuk (Warsaw University)

15:00 - 15:30 **Keren Censor-Hillel -** *Distributed Spanner Approximation*

I will talk about recent results on distributed constructions of approximate minimum spanners of various types. Our work provides lower bounds for deterministic and randomized α-approximations for the minimum directed k-spanner problem for $k \geq 5$, which, together with known and new algorithmic results imply a strict separation between the LOCAL and CONGEST models. Notably, to the best of our knowledge, this is the first separation between these models for a *local approximation problem.* Similarly, a separation between the directed and undirected cases is implied. We also provide a lower bound for the minimum *weighted* k-spanner problem for $k \geq 4$, for directed or undirected graphs. In addition, we show lower bounds for the minimum weighted 2-spanner problem in the CONGEST and LOCAL models.

On the algorithmic side, we show a polylog-round $(1 + \epsilon)$-approximation algorithm for minimum k-spanners, and a new distributed construction of minimum 2-spanners that uses only polynomial local computations. The latter has a *guaranteed* approximation ratio of $O(\log(m/n))$ for a graph with n vertices and m edges, which matches the best known ratio for polynomial time sequential algorithms [Kortsarz and Peleg, 1994], and is tight if we restrict ourselves to polynomial local computations. The algorithm completes in $O(\log n \log \Delta)$ rounds w.h.p, where Δ is the maximum degree in the graph. Our approach allows us to extend our algorithm to work also for the *directed*, *weighted*, and *client-server* variants of the problem. It also provides a CONGEST algorithm for the minimum dominating set problem, with a *guaranteed* $O(\log \Delta)$ approximation ratio.

Joint work with Michal Dory.

15:30 - 16:00 **Michel Habib -** *Graph classes and forbidden patterns*

Characterization of graph classes via a vertex ordering that avoids some patterns are well known for chordal, interval graph, proper interval graphs and comparability graphs. Most intriguing is that the best known algorithms to recognize these classes in fact produce such an ordering. From the seminal work in Damaschke [1990] and Wood [2004] that pointed out this subject, we know that the recognition problem of all patterns of size 3 is polynomial Hell, Mohar and Rafiey [2014]. Furthermore Duffus, Ginn and Rödl [1995] showed that almost all 2-connected patterns are NP-complete to recognize. Hell et al conjectured a dichotomy property on this recognition problem, but it seems to be false Nešetřil [2017]. This paper is devoted to a systematic study of the classes of graphs that can be defined using forbidden patterns on three vertices. We give a complete characterization of all the graph classes than can be defined that way, not only this allows to give another proof of the Hell, Mohar and Rafiey's result, but also to greatly improve the complexity of the recognition. In fact we prove that except for the well-know classes of triangle-free graphs and comparability graphs and their complement all other classes can be linearly recognized. We will finish the talk by studying patterns on four vertices and describe the landscape including a series of conjectures.

Joint work with Laurent Feuilloley.

16:00 - 16:30 *Coffee break*

16:30 - 17:00 **Guy Even -** *Optimal Distributed Weighted Set Cover Approximation*

We present a time-optimal deterministic distributed algorithm for approximating a minimum weight vertex cover in hypergraphs of rank f. This problem is equivalent to the Minimum Weight Set Cover Problem in which the frequency of every element is bounded by f. The approximation factor of our algorithm is $(f+\epsilon)$. Let Δ denote the maximum degree in the hypergraph. Our algorithm runs in the CONGEST model and requires $O(\log\Delta/\log\log\Delta)$ rounds, for constants $\epsilon \in (0,1]$ and $f \in \mathbb{N}^+$. This is the first distributed algorithm for this problem whose running time does not depend on the vertex weights or the number of vertices. Thus adding another member to the exclusive family of *provably optimal* distributed algorithms.

For constant values of f and ϵ, our algorithm improves over the $(f+\epsilon)$-approximation algorithm of Kuhn et al. (2006) whose running time is $O(\log\Delta+\log W)$, where W is the ratio between the largest and smallest vertex weights in the graph.

Joint work with Ran Ben Basat, Ken-ichi Kawarabayashi, and Gregory Schwartzman.

17:00 - 17:30 **Yonatan Nakar -** *On the Testability of Graph Partition Properties*

In this work we study the testability of a family of graph partition properties that generalizes a family previously studied by Goldreich, Goldwasser, and Ron (Journal of the ACM, 1998). While the family studied by Goldreich, Goldwasser, and Ron includes a variety of natural properties, such as k-colorability and containing a large cut, it does not include other properties of interest, such as split graphs, and more generally (p,q)-colorable graphs. The generalization we consider allows us to impose constraints on the edge-densities within and between parts (relative to the sizes of the parts). We denote the family studied in this work by GPP. We first show that all properties in GPP have a testing algorithm whose query complexity is polynomial in 1/epsilon, where epsilon is the given proximity parameter (and there is no dependence on the size of the graph). As the testing algorithm has two-sided error, we next address the question of which properties in GPP can be tested with one-sided error and query complexity polynomial in 1/epsilon. We answer this question by establishing a characterization result. Namely, we define a subfamily $\mathrm{GPP}_{0,1}$ of GPP and show that a property P in GPP is testable by a one-sided error algorithm that has query complexity $\mathrm{poly}(1/\epsilon)$ if and only if P in $\mathrm{GPP}_{0,1}$.

17:30 - 18:00 **Pierre Ohlmann -** *Unifying circuit lower bounds using Hankel matrix and parse trees*

Arithmetic circuit are the algebraic analogue of boolean circuits. As a natural model for computing multivariate polynomials, they have been extensively studied. The most important open question in the field of algebraic complexity theory is that of separating the classes VP and VNP, the analogues of P and NP. More precisely, can one obtain super-polynomial lower bounds for circuits computing a given explicit polynomial ?

Despite decades of efforts, this question yet seems out of reach, the best general lower bound being only slightly super-linear. The most common approach is to prove lower bounds for restricted classes of circuits, such as monotone or constant-depth circuits. Another approach would be removing relations arising from the mathematical structure underlying the computations, making it harder for circuits to compute polynomials and thus conceivably easier to obtain lower bounds. In this line of thought, Nisan (1991) was able to obtain breakthrough results in the context of non-commutative computations, separating circuits and formulas and characterizing the minimal size of Algebraic Branching Programs (ABP).

Likewise, circuits for which the multiplication is assumed to be non-associative, meaning that (xy)x and x(yx) are different monomials, have been considered. General exponential lower bounds can be proved in this setting. We highlight a syntactical equivalence between non-associative circuits and acyclic Multiplicity Tree Automata (MTA), which leads to a general algebraic characterization of the size of the smallest non-associative circuit computing a given non-associative polynomial.

As a direct consequence of this characterization, we unify several recent circuit lower bounds in the non-commutative setting. Going deeper in the comprehension of this new algebraic tool, we are able to considerably extend a known lower bound to a class which is very close to general.

Thursday October 4, 2018

9:00 - 9:30 **Shiri Chechik -** *Near-Optimal Approximate Decremental All Pairs Shortest Paths*

Computing shortest paths is one of the fundamental problems of graph algorithms. The goal of *dynamic* all pairs shortest paths (APSP) is to maintain shortest path trees from all vertices as the edges of the graph change over time. The algorithm is said to be decremental if it handles only deletions, incremental if it handles only insertions and fully dynamic if it handles both deletions and insertions. In this talk I will present a near optimal decremental algorithm that maintains approximate all pairs shortest paths.

9:30 - 10:00 **Or Zamir -** *Dynamic ordered sets with approximate queries*

We consider word RAM data structures for maintaining ordered sets of integers whose select and rank operations are allowed to return approximate results, i.e., ranks or items whose rank differ by at most D from the exact answer, where $D = D(n)$ is an error parameter. While exact select and rank have identical operation times, both in comparison-based and word RAM implementations, an interesting separation emerges between approximate versions of these operations in the word RAM model. It turns out that approximate select is much easier than approximate rank. We also consider approximate nearest neighbors and approximate priority queues. We present optimal bounds for all these approximate problems with matching cell-probe lower bounds for select, rank, and nearest neighbors, and an equivalence to a well-studied static problem for approximate priority queues.

(Joint work with Mikkel Thorup and Uri Zwick)

10:00 - 10:30 **Benny Chor -** *Is inversion symmetry of chromosomes a law of nature*

Erwin Chargaff, the Austro-Hungarian biochemist, has made, in 1950, the important observation that the numbers of bases in DNA satisfy #A=#T and #G=#C. This played a crucial role in realizing that DNA has a two strand structure with base-pair bindings (of A to T and C to G), as proposed in 1953 by Crick and Watson.

Another paper, in 1968, has revealed the second Chargaff rule (SCR) stating that the same sets of identities hold for each long enough single DNA strand. But whereas the 1st rule can in hindsight be justified by base-pair binding, the SCR has remained a curious puzzle.

I will discuss generalizations of the second Chargaff rule to k-mers of DNA bases, and some possible conjectures on the origins of this "rule".

Joint work with Sagi Shopper, Saharon Rosset, David Horn.

10:30 - 11:00 *Coffee break*

11:00 - 11:30 **Laurent Feuilloley -** *Local verification of global proofs*

In this work we study the cost of local and global proofs on distributed verification. In this setting the nodes of a distributed system are provided with a proof for the correctness of the state of the system, and the nodes need to verify this proof by looking at only their local neighborhood in the system. Previous works have studied the model where each node is given its own, possibly unique, part of the proof as input. The cost of a proof is the maximum size of an individual label. We compare this model to a model where each node has access to the same global proof, and the cost is the size of this global proof.

Joint work with Juho Hirvonen.

11:30 - 12:00 **Ami Paz -** *Redundancy in Distributed Proofs*

Distributed proofs are mechanisms enabling the nodes of a network to collectively and efficiently check the correctness of Boolean predicates on the structure of the network (e.g. having a specific diameter), or on data-structures distributed over the nodes (e.g. a spanning tree). We show that many network predicates have distributed proofs offering a high level of redundancy, explicitly or implicitly. We use this remarkable property of distributed proofs for establishing perfect tradeoffs between the size of the certificate stored at every node, and the number of rounds of the verification protocol.

Joint work with Laurent Feuilloley, Pierre Fraigniaud, Juho Hirvone and Mor Perry.

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

14:00 - 14:30 **Zvi Lotker -** *Was Abraham Chinese?*

From Abraham to one-child policy. A journey through time and culture using Percolation theory. See more at https://www.youtube.com/watch?v=a1saDaRWiaU .

This is a joint work with David Peleg.

14:30 - 15:00 **Serge Abiteboul -** *Explanations and Transparency in Collaborative Workflows*

We pursue an investigation of data-driven collaborative work-flows. In the model, peers can access and update local data, causing side effects on other peers' data. In this paper, we study means of explaining to a peer her local view of a global run, both at runtime and statically. We consider the notion of "scenario for a given peer" that is a subrun observationally equivalent to the original run for that peer. Because such a scenario can sometimes differ significantly from what happens in the actual run, thus providing a misleading explanation, we introduce and study a faithfulness requirement that ensures closer adherence to the global run. We show that there is a unique minimal faithful scenario, that explains what is happening in the global run by extracting only the portion relevant to the peer. With regard to static explanations, we consider the problem of synthesizing, for each peer, a "view program" whose runs generate exactly the peer's observations of the global runs. Assuming some conditions desirable in their own right, namely transparency and boundedness, we show that such a view program exists and can be synthesized. As an added benefit, the view program rules provide provenance information for the updates observed by the peer.

Joint work with Pierre Bourhis and Victor Vianu.

15:00 - 15:30 **Tova Milo -** *Experts-in-the-Loop in Data Management*

One of the foremost challenges for information technology over the last years has been to explore, understand, and extract useful information from large amounts of data. To assist in this challenging task, a new research field called data-centered crowdsourcing has emerged, aiming to effectively delegate a wide range of analysis tasks to human workers (called crowd) and to improve the efficiency of this process. What may be achieved with the help of the crowd naturally depends on the properties of the human workers. This talk will focus on the high end of workers - knowledgeable crowds - and in particular domain knowledge experts, and examine how they can be used for assisting in solving data management problems. Specifically we examine two complementary dimensions: (1) How domain experts can help in improving the *data* itself, e.g. by gathering missing data and improving the quality of existing data, and (2) How they can assist in gathering *meta-data* that facilitate improved data processing. In each of the two dimensions we zoom on one typical problem-instance and investigate how experts can be effectively harnesses to improve data processing in this instance

15:30 - 16:00 **Pierre Senellart -** *Connecting Width and Structure in Knowledge Compilation*

Several query evaluation tasks can be done via knowledge compilation: the query result is compiled as a lineage circuit from which the answer can be determined. For such tasks, it is important to leverage some width parameters of the circuit, such as bounded treewidth or pathwidth, to convert the circuit to structured classes, e.g., deterministic structured NNFs (d-SDNNFs) or OBDDs. In this work, we show how to connect the width of circuits to the size of their structured representation, through upper and lower bounds.

(joint work with Antoine Amarilli and Mikaël Monet)

16:00 - 16:30 *Coffee break*

16:30 - 17:00 **Ioana Bercea -** *Improved bounds for the Traveling Salesman Problem with Neighborhoods*

Given a set of n disks of radius R in the Euclidean plane, the Traveling Salesman Problem With Neighborhoods (TSPN) on uniform disks asks for the shortest tour that visits all of the disks. The problem is a generalization of the classical Traveling Salesman Problem(TSP) on points and has been widely studied in the literature. For the case of disjoint uniform disks of radius R, Dumitrescu and Mitchell [2003] show that the optimal TSP tour on the centers of the disks is a 3.547-approximation to the TSPN version. The core of their analysis is based on bounding the detour that the optimal TSPN tour has to make in order to visit the centers of each disk and shows that it is at most $2Rn$ in the worst case. Häme, Hyytiä and Hakula [2011] asked whether this bound is tight when R is small and conjectured that it is at most $\sqrt{3}Rn$.

In this talk, we will further investigate this question and derive structural properties of the optimal TSPN tour to describe the cases in which the bound is smaller than $2Rn$. Specifically, we show that if the optimal TSPN tour is not a straight line, at least one of the following is guaranteed to be true: the bound is smaller than $2Rn$ or the TSP on the centers is a 2-approximation (best we can get with this heuristic). This leads to an improved overall approximation factor for the problem. Along the way, we will uncover one way in which the TSPN problem is structurally different from the classical TSP.

17:00 - 17:30 **Karthik C. S. -** *On Complexity of Closest Pair in Euclidean Metric*

Given a set of n points in R^d, the Closest Pair problem asks to find a pair of distinct points in the set that are closest in the L_p-metric. Closest Pair is a fundamental problem in Computational Geometry and understanding its fine-grained complexity in the Euclidean metric when $d = \text{polylog}(n)$ has recently received attention. In this talk, we will see that under the Strong Exponential Time Hypothesis, for every $\epsilon > 0$, no algorithm running in time $n^{2-\epsilon}$ can solve the Closest Pair problem in the Euclidean metric, even when $d = \text{polylog}(n)$. The talk will be based on a joint work with Pasin Manurangsi.

18:00 - 20:00 *Reception*

Friday October 5, 2018

9:00 - 9:30 **Yishay Mansour -** *Submultiplicative Glivenko-Cantelli and Uniform Convergence of Revenues*

In this work we derive a variant of the classic Glivenko-Cantelli Theorem, which asserts uniform convergence of the empirical Cumulative Distribution Function (CDF) to the CDF of the underlying distribution. Our variant allows for tighter convergence bounds for extreme values of the CDF.

We apply our bound in the context of revenue learning, which is a well-studied problem in economics and algorithmic game theory. We derive sample-complexity bounds on the uniform convergence rate of the empirical revenues to the true revenues, assuming a bound on the k-th moment of the valuations, for any (possibly fractional) $k > 1$.

For uniform convergence in the limit, we give a complete characterization and a zero-one law: if the first moment of the valuations is finite, then uniform convergence almost surely occurs; conversely, if the first moment is infinite, then uniform convergence almost never occurs.

Join work with Noga Alon, Moshe Babaioff, Yannai A. Gonczarowski, Shay Moran, and Amir Yehudayoff.

9:30 - 10:00 **Guy Rothblum -** *On Fairness in Automated Classification*

As algorithms increasingly inform and influence decisions made about individuals, it becomes increasingly important to address concerns that these algorithms might be discriminatory.

The increasing reach of algorithmic classification and decision making into our daily lives Ñ from advertising to candidate filtering for jobs to incarceration Ñ has given rise to an explosion of research into the ethics embodied by these algorithms; in a word, are they ÒfairÓ? But what is fairness? Can we test for it? Can we achieve it? Are there limits?

The burgeoning study of algorithmic fairness attempts to tackle these challenges from a foundational perspective. This talk will give a brief survey about some of this progress and the many challenges that remain.

Based on joint works with Cynthia Dwork, Úrsula Hébert-Johnson, Michael P. Kim, Omer Reingold and Gal Yona.

10:00 - 10:30 **Iordanis Kerenidis -** *Quantum algorithms for Optimisation and Machine Learning*

We will review some recent results on quantum algorithms for optimisation, including efficient algorithms for gradient descent and interior point methods, and for machine learning applications, including recommendation systems and classification.

10:30 - 11:00 *Coffee break*

11:00 - 11:30 **Claire Mathieu -** *Approximation schemes for rank aggregation*

Given a collection of rankings, how does one produce a satisfactory full ranking? I will present approximation schemes for several variants of the problem.

11:30 - 12:00 **Amos Fiat -** *c-single crossing Interdependent valuations*

We consider a setting where an auctioneer sells a single item to n potential agents with *interdependent values*. That is, each agent has her own private signal, and the valuation of each agent is a known function of all n private signals. This captures settings such as valuations for artwork, oil drilling rights, broadcast rights, and many more.

In the interdependent value setting, all previous work has assumed a so-called *single-crossing condition*. Single-crossing means that the impact of agent i's private signal, s_i, on her own valuation is greater than the impact of s_i on the valuation of any other agent. It is known that without the single-crossing condition an efficient outcome cannot be obtained. We study welfare maximization for interdependent valuations through the lens of approximation.

We show that, in general, without the single-crossing condition, one cannot hope to approximate the optimal social welfare any better than the approximation given by assigning the item to a random bidder. Consequently, we introduce a relaxed version of single-crossing, *c-single-crossing*, parameterized by $c \geq 1$, which means that the impact of s_i on the valuation of agent i is at least $1/c$ times the impact of s_i on the valuation of any other agent ($c = 1$ is single-crossing). Using this parameterized notion, we obtain a host of positive results. We also consider interdependent settings when valuations are concave and give improved results. Joint work with Alon Eden, Michal Feldman, and Kira Goldner.

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

14:00 - 14:30 **Sylvain Perifel -** *Lempel-Ziv: a "one-bit catastrophe" but not a tragedy*

The robustness of the famous compression algorithm of Lempel and Ziv is still not well understood: in particular, until now it was unknown whether the addition of one bit in front of a compressible word could make it incompressible. This talk will answer that question, advertised by Jack Lutz under the name "one-bit catastrophe" and which has been around since at least 1998. We will show that a "well" compressible word remains compressible when a bit is added in front of it, but some "few" compressible words indeed become incompressible.

This is a joint work with Guillaume Lagarde.

14:30 - 15:00 **Danny Vainstein -** *The Price of Clustering in Bin-Packing with Applications to Bin-Packing with Delays*

One of the most significant algorithmic challenges in the "big data era" is handling instances that are too big to be processed by a single machine. The common practice in this regard is to partition the massive problem instance into smaller ones and process each one of them separately. In some cases, the solutions for the smaller instances are later on assembled into a solution for the whole instance, but in many cases this last stage cannot be pursued (e.g., because it is too costly, because of locality issues, or due to privacy considerations). Motivated by this phenomenon, we consider the following natural combinatorial question: Given a bin-packing instance (namely, a set of items with sizes in $(0, 1]$ that should be packed into unit capacity bins) I and a partition $\{I_i\}_i$ of I into clusters, how large is the ratio $\sum_i OPT(I_i)/OPT(I)$, where $OPT(J)$ denotes the optimal number of bins into which the items in J can be packed? In this talk, we investigate the supremum of this ratio over all instances I and partitions $\{I_i\}_i$, referred to as the bin-packing price of clustering (PoC). It is trivial to observe that if each cluster contains only one tiny item (and hence, $OPT(I_i) = 1$), then the PoC is unbounded. On the other hand, a relatively straightforward argument shows that under the constraint that $OPT(I_i) \geq 2$, the PoC is 2. Our main challenge was to determine whether the PoC drops below 2 when $OPT(I_i) > 2$. In addition, one may hope that $\lim_{k\to\infty} PoC(k) = 1$, where $PoC(k)$ denotes the PoC under the restriction to clusters I_i with $OPT(I_i) \geq k$. We resolve the former question affirmatively and the latter one negatively: Our main results are that $PoC(k) \leq 1.951$ for $k > 2$ and that $\lim_{k\to\infty} PoC(k) \leq 1.751$; we then show that this cannot be significantly improved, proving that $PoC(k) \geq 1.691$ for any value of k. In addition to the immediate contribution of this combinatorial result to "big data" kind of applications, it turns out that it is useful also for an interesting online problem called bin packing with delays.

Joint work with Yossi Azar and Yuval Emek.

15:00 - 15:30 **Yossi Azar -** *The Price of Bounded Preemption*

We provide a tight bound for the "price of preemption" for scheduling jobs on a single or multiple machines. The input consists of a set of jobs to be scheduled and of an integer parameter k >0. Each job has a release time, deadline, length (also called processing time) and value associated with it. The goal is to feasibly schedule a subset of the jobs so that their total value is maximal; while preemption of a job is permitted, a job may be preempted no more than k times. The price of preemption is the worst possible (i.e., largest) ratio of the optimal non-bounded-preemptive scheduling to the optimal k-bounded-preemptive scheduling.

Our results show that allowing at most k preemptions suffices to guarantee a logarithmic (to the base k+1) factor giving us an upper bound for the price; a specific scenario serves to prove the tightness of this bound. We further determine the price when no preemptions are permitted at all (i.e., k=0).

As part of the proof, we introduce the notion of the "Bounded-Degree Ancestor-Independent Sub-Forest (BAS)". We investigate the problem of computing the maximal-value BAS of a given forest and give a tight bound for the loss factor, which is of the order log n / log (k+1), where n is the size of the original forest and k is the bound on the degree of the sub-forest.

Joint work with Noga Alon and Mark Berlin.

15:30 - 16:00 **Noam Touitou -** *Set Cover and Vertex Cover with Delay*

The set cover problem is one of the most fundamental problems in computer science. We present the problem of online set cover with delay (SCD). A family of sets with costs and a universe of elements are known in advance. Requests then arrive over time on the elements, and each request accumulates delay cost until served by the algorithm through buying a containing set. A request can only be served by sets that are bought after the request's arrival, and thus a set may be bought an unbounded number of times over the course of the algorithm. This property sets SCD apart from previous considerations of set cover in the online setting, in which there are no delays, elements are covered immediately, and sets stay bought permanently. This allows SCD to describe an unbounded process, with an unlimited number of requests for any given universe.

For the SCD problem, we show an O(log k log n)-competitive randomized algorithm, where n is the number of elements and k is the maximum number of sets containing any single element. We also show a lower bound of Omega(sqrtlog k) and Omega(sqrtlog n) on the competitiveness of any algorithm for SCD. For the special case of Vertex Cover with Delay (VCD), we show a simple 3-competitive deterministic algorithm. The O(log k log n)-competitive algorithm is based on exponential weights combined with the max operator (in contrast to most algorithms employing exponential weights, which use summation). The lower bound is described by a recursive construction.

Joint work with Yossi Azar, Ashish Chiplunkar and Shay Kutten.

16:00 - 16:30 *Coffee break*

16:30 - 17:00 **Haim Kaplan -** *Clustering in Hypergraphs to Minimize Average Edge Service Time*

We study the problem of clustering the vertices of a weighted hypergraph such that on average the vertices of each edge can be covered by a small number of clusters. This problem has many applications such as for designing medical tests, clustering files on disk servers, and placing network services on servers. The edges of the hypergraph model groups of items that are likely to be needed together, and the optimization criteria which we use can be interpreted as the average delay (or cost) to serve the items of a typical edge. We describe and analyze algorithms for this problem for the case in which the clusters have to be disjoint and for the case where clusters can overlap. The analysis is often subtle and reveals interesting structure and invariants that one can utilize.

17:00 - 17:30 **Uri Zwick -** *One-way communication complexity of connectivity with private randomness*

Each vertex of an undirected graph sees the edges incident to it and can send a message to a referee. Given all the messages received from the vertices, the referee should be able to decide whether the graph is connected or not. Using shared randomness, this can be done, with high probability, with each vertex sending only$O(\log^3 n)$ bits. We show that using private randomness this can be done with each vertex sending $\tilde{O}(n^{1/2})$ bits. In the best known deterministic solution each vertex sends n/2 bits and it is open whether this is best possible.

Joint work with Jacob Holm, Valerie King and Mikkel Thorup.