Program

Wednesday June 1, 2016

9:00 - 9:15 **Registration**

9:15 - 9:30 **Opening Remarks**

9:30 - 10:15 **Moshe Lewenstein -** *Orthogonal Range Searching and Text Indexing*

Text indexing, the problem in which one desires to preprocess a (usually large) text for future (shorter) queries, has been researched ever since the suffix tree was invented in the early 70's. With textual data continuing to increase and with changes in the way it is accessed, new data structures and new algorithmic methods are continuously required. Therefore, text indexing is of utmost importance and is a very active research domain.

Orthogonal range searching, classically associated with the computational geometry community, is one of the tools that has increasingly become important for various text indexingapplications. Initially, in the mid 90's there were a couple of results recognizing this connection. In the last few years we have seen an increase in use of this method and are reaching a deeper understanding of the range searching uses for text indexing.

10:15 - 10:45 **Coffee break**

10:45 - 11:30 **Gregory Kucherov -** *Recent Results on Computing Repeats in Strings*

Identifying repeats in strings is one of central topics in string algorithms. In this talk, we highlight some new developments in this area. After giving a quick overview of the history of the problem and the state-of-the-art, we present recent results on computing so-called maximal α-gapped repeats.

Joint work with Maxime Crochemore and Roman Kolpakov.

11:30 - 12:15 **Carola Doerr -** *Spotlight on the Analysis of Evolutionary Algorithms*

Evolutionary algorithms (EAs) form a powerful class of optimization techniques for both industrial and academic applications. A large research community exist that studies EAs from a scientific angle. In this talk, we show that EAs and other randomized search heuristics are also interesting from a mathematical point of view. We provide an example showing how running time analysis and complexity theory have inspired the development of a novel algorithm which provably outperforms existing techniques in evolutionary computation. Our example also affirmatively answers one of the main open questions in EA theory, namely the usefulness of crossover (i.e., the recombination of two or more search points) on simple fitness landscapes.

The talk is based on joint work with Benjamin Doerr (Ecole Polytechnique, France) and Franziska Ebel (Saarland University, Germany).

12:15 - 14:00 **Lunch (To be found in restaurants around. Please see our list of restaurants.)**

14:00 - 14:45 **Alan Roytman -** *Packing Small Vectors*

Online d-dimensional vector packing models many settings such as minimizing resources in data centers where jobs have multiple resource requirements (CPU, memory, etc.). However, no online d-dimensional vector packing algorithm can achieve a competitive ratio better than d. Fortunately, in many natural applications, vectors are relatively small, and thus the lower bound does not hold. For sufficiently small vectors, an $O(\log d)$-competitive algorithm was known. We improve this to a constant competitive ratio, arbitrarily close to e (where e is the base of the natural logarithm), given that vectors are sufficiently small.

We give improved results for the two dimensional case. For arbitrarily small vectors, the First Fit algorithm for two dimensional vector packing is no better than 2-competitive. We present a natural family of First Fit variants, and for optimized parameters get a competitive ratio of approximately 1.48 for sufficiently small vectors.

We improve upon the 1.48 competitive ratio – not via a First Fit variant – and give a competitive ratio arbitrarily close to 4/3 for packing small, two dimensional vectors. We show that no algorithm can achieve better than a 4/3 competitive ratio for two dimensional vectors, even if one allows the algorithm to split vectors among arbitrarily many bins.

Joint work with Yossi Azar, Ilan Reuven Cohen, and Amos Fiat.

14:45 - 15:30 **Adi Vardi -** *Make-to-Order Integrated Scheduling and Distribution*

Production and distribution are fundamental operational functions in supply chains. The main challenge is to design algorithms that optimize operational performance by jointly scheduling production and delivery of customer orders. In this paper we study a model of scheduling customer orders on multiple identical machines and their distribution to customers afterwards. The goal is to minimize the total time from release to distribution plus total distribution cost to the customers. We design the first poly-logarithmic competitive algorithm for the problem, improving upon previous algorithms with linear competitive ratios. Our model generalizes two fundamental problems: scheduling of jobs on multiple identical machines (where the goal function is to minimize the total flow time) as well as the TCP Acknowledgment problem.

15:30 - 16:00 **Coffee break**

16:00 - 16:45 **Florent Urrutia -** *Multi-Party Protocols, Information Complexity and Privacy*

We study computation protocols, especially in the multi-party setting, their communication complexity, and information theoretical tools that prove useful for their analysis. We first shortly overview the main definitions and the main known information theoretical tools that have proven useful in recent years for the analysis of two-party computation protocols. Since these tools do not readily extend to the multi-party case, we then introduce a new measure, called *Public Information Complexity* (PIC), as a tool for the study of multi-party computation protocols, and of quantities such as their communication complexity, or the amount of randomness they require in the context of information-theoretic private computations.

We are able to use this measure directly in the natural asynchronous message-passing *peer-to-peer* model and show a number of interesting properties and applications of our new notion: the Public Information Complexity is a lower bound on the Communication Complexity and an upper bound on the Information Complexity; the difference between the Public Information Complexity and the Information Complexity (IC) provides a lower bound on the amount of randomness used in a protocol; any communication protocol can be compressed to its Public Information Cost; an explicit calculation of the zero-error *Public Information Complexity* of the k-party, n-bit Parity function, where a player outputs the bit-wise parity of the inputs. The latter result establishes that the amount of randomness needed for a private protocol that computes this function is $\Omega(n)$.

Based on joint work with Iordanis Kerenidis and Adi Rosén.

16:45 - 17:30 **Benny Chor -** *Finding Relatives Almost Privately in the Moderator Model*

We study the problem of identification of relatives from genomic data, a service already being offered by a number of personal genomics companies. In this application, individuals who wish to discover distant (or lost) relatives, submit their DNA to a company, which analyzes it to detect familial relationships among pairs of individuals. This entire DNA information of all customers is therefore at the company's hands, posing a potential threat to their genomic privacy.

We introduce the moderator model, and present a two-step protocol for privately inferring relationships among a given set of individuals. In this setting, individuals submit only a relatively small number of bits, determined by their DNA. Based on this partial data, the moderator determines, with high probability, pairs of users that are potential relatives. These pairs of users will then engage in a two-party secure computation in order to determine whether they are indeed related.

Joint work with Orit Moskovich and Benny Pinkas.

Thursday June 2, 2016

9:30 - 10:15 **Pascal Koiran -** *Lower Bounds and Reconstruction Algorithms for Sums of Affine Powers*

A sum of affine powers is an expression of the form $f(x) = \sum_{i=1}^{s} \alpha_i (x - a_i)^{e_i}$. Although quite simple, this model is a generalization of two well-studied models: Waring decomposition and sparsest shift. For these three models there are natural extensions to several variables, but this talk will be focused on univariate polynomials.

We will explain how to prove lower bounds on the number of terms s needed to represent a given polynomial f. With some additional work, the same methods yield in several cases efficient algorithms for reconstructing the smallest expression for an input polynomial f. If time allows, we will also present some structural results connecting the three models.

The lower bounds are joint work with Neeraj Kayal, Timothée Pecatte and Chandan Saha (ICALP 2015). The other results are work in progress with Ignacio Garcia-Marco and Timothée Pecatte.

10:15 - 10:45 **Coffee break**

10:45 - 11:30 **Amnon Ta-Shma -** *Explicit Two-Source Extractors for Near-Logarithmic Min-Entropy*

We explicitly construct extractors for two independent n-bit sources of $(\log n)^{1+o(1)}$ min-entropy. Previous constructions required either $\text{polylog}(n)$ min-entropy or five sources.

Our result extends the breakthrough result of Chattopadhyay and Zuckerman and uses the non-malleable extractor of Cohen. The main new ingredient in our construction is a somewhere-random condenser with a small entropy gap, used as a sampler. We construct such somewhere-random condensers using the error reduction mechanism of Raz et al. together with the high-error, constant degree dispersers of Zuckerman.

Joint work with Dean Doron and Avraham Ben-Aroya.

11:30 - 12:15 **Frédéric Magniez -** *Stable Matching with Evolving Preferences*

We consider the problem of stable matching with dynamic preference lists. At each time-step, the preference list of some player may change by swapping random adjacent members. The goal of a central agency (algorithm) is to maintain an approximately stable matching, in terms of number of blocking pairs, at all time-steps. The changes in the preference lists are not reported to the algorithm, but must instead be probed explicitly. We design an algorithm that in expectation and with high probability maintains a matching that has at most $O((\log n)^2)$ blocking pairs.

Joint work with V. Kanade and N. Leonardos.

12:15 - 14:00 **Lunch (To be found in restaurants around. Please see our list of restaurants.)**

14:00 - 14:45 **Zvi Lotker -** *Social Networks and Plays*

In this talk I will describe the connection between social networks and theater plays. I will show how algorithms can analyze plays using social networks, and how plays can reveal an interesting algorithmic problem.

14:45 - 15:30 **Yinon Nahum -** *Social Network Cores*

Consider a preferential attachment model for network evolution that allows both node and edge arrival events: starting with an arbitrary graph G_0, at time t, with probability p_t a new node arrives and a new edge is added between the new node and an existing node, and with probability $1-p_t$ a new edge is added between two existing nodes. In both cases existing nodes are chosen at random according to preferential attachment, i.e., with probability proportional to their degree. Ordering the nodes of the network in order of their arrival, a *founders* set is a prefix of V, namely, the first k nodes for some k. For $\delta \in (0,1)$, the δ-*founders* set of the network at time t is the minimal founders set guaranteeing that the sum of degrees of its nodes is at least a δ fraction of the sum of degrees in the graph at time t. We show that for the common model where p_t is constant, i.e., when $p_t = p$ for every t and the network is sparse (with a linear number of edges) the size of the δ-founders set is concentrated around $\delta^{2/p} n_t$, and thus is linear in n_t, the number of nodes at time t. In contrast, we show that for $p_t = \min\{1, \frac{2c}{\ln t}\}$ and when the network is dense (with a super-linear number of edges), the size of the δ-founders set is sub-linear in n_t and concentrated around $\tilde{\Theta}((n_t)^{\eta})$, where $\eta = \delta^{1/c}$.

15:30 - 16:00 **Coffee break**

16:00 - 16:45 **Iordanis Kerenidis -** *Quantum Recommendation Systems*

We will review the power and limitations of quantum algorithms for linear algebraic problems. We will then give a concrete example of the power of quantum information by providing a quantum algorithm for Recommendation Systems that runs in time only poly-logarithmic in the dimension of the data. No prior knowledge of quantum information will be assumed.

19:00 - 21:30 **Reception (cocktail dînatoire)** - Ground floor of the Buffon building

Friday June 3, 2016

9:30 - 10:15 **Robert Krauthgamer -** *On Sketching Quadratic Forms*

I will introduce the problem of sketching a quadratic form: Given a square matrix A, create a succinct (randomized) sketch sk(A) that suffices to estimate, for any desired query vector x, the quadratic form $x^T Ax$ (without further access to A). Since a general matrix does not admit non-trivial sketches, the talk will focus on important special cases like positive semi-definite (PSD) matrices, or even when A is a graph Laplacian matrix and x is a 0-1 vector, in which case the quadratic form is just the value of a cut in a graph.

I will present near-tight bounds on the tradeoff between sketch size and accuracy. In particular, we will compare the “for each” guarantee (for each query x, with a constant probability the sketch succeeds) and the stronger “for all” guarantee (the sketch succeeds for all x's simultaneously).

Joint work with Alexandr Andoni, Jiecao Chen, Bo Qin, David P. Woodruff, and Qin Zhang.

10:15 - 10:45 **Coffee break**

10:45 - 11:30 **Amos Fiat -** *Dynamic Pricing*

In this talk we'll cover several recent results regarding dynamic pricing in a variety of settings. We show efficient dynamic pricing for online decision making covering metrical matchings, task systems, the k-server problem, unit demand markets, and load balancing.

Covering joint work with Ilan Cohen, Vincent Cohen-Addad, Alon Eden, Michal Feldman, and Alan Roytman.

11:30 - 12:15 **Vincent Cohen-Addad -** *One Size Fits All: Effectiveness of Local Search on Structured Data*

We analyze the performance of the Local Search algorithm for clustering on well behaved data. Since the seminal paper by Ostrovsky, Rabani, Schulman and Swamy [FOCS 2006], much progress has been done to characterize real-world instances. We distinguish three main definitions: (1) Distribution Stability (Awasthi, Blum, Sheffet, FOCS 2010); (2) Spectral Separability (Kumar, Kannan, FOCS 2010); (3) Perturbation Resilience (Bilu, Linial, ITCS 2010). We show that Local search preforms well on the instances with the aforementioned“stabilit” properties. Specifically, for the k-means and k-median objective, we show that Local Search exactly recovers the optimal clustering if the data set is $3 + \epsilon$-perturbation resilient, and is a PTAS for distribution stability and spectral separability. This matches the performance of existing approaches for the first two conditions and it is the first PTAS for instances satisfying the spectral separability condition. This sheds light on the success of Local Search in clustering applications.

12:15 - 14:00 **Lunch (To be found in restaurants around. Please see our list of restaurants.)**

14:00 - 14:45 **Allan Borodin -** *Rethinking Online and Greedy Algorithms*

Throughout history there has been an appreciation of the importance of simplicity in the arts and sciences. In the context of algorithm design, and in particular in approximation algorithms, social choice and algorithmic game theory, the importance of simplicity is currently more in vogue. In addition to applications where online or simple algorithms may be highly desirable or even necessary, recent results by Buchbinder and Feldman, and by Poloczek et al suggest that online and greedy algorithms offer the possibility of obtaining good” approximation ratios by “conceptually simple combinatorial algorithms” based on generalizations of the online model. I will present some examples of a renewed interest in the design of “conceptually simple algorithms”. In particular, I will revisit the pBT model we proposed for simple dynamic programming and how it applies to a new derandomization result of Buchbinder and Feldman.

14:45 - 15:30 **Shiri Chechik -** *Deterministic Decremental Single Source Shortest Paths: Beyond the $O(mn)$ Bound*

Computing shortest paths is one of the fundamental problems of graph algorithms. The goal of *dynamic* single source shortest paths (SSSP) is to maintain a shortest path tree from a fixed source s as the edges of the graph change over time. The most general case is the fully dynamic one, where each adversarial update inserts or deletes an arbitrary edge. The trivial algorithm is to recompute SSSP after every update in $O(m)$ time. For the fully dynamic case, no non-trivial algorithm is known. We can, however, improve upon the trivial algorithm by restricting the update sequence to be partially dynamic: only insertions (referred to as incremental), or only deletions (referred to as decremental).

In a seminal work, Even and Shiloach [JACM 1981] presented an exact solution to partially dynamic single source shortest path with only $O(mn)$ total update time over all edge deletions. Their classic algorithm was the best known result for the decremental SSSP problem for three decades, even when approximate shortest paths are allowed.

The first improvement over the Even-Shiloach algorithm was given by Bernstein and Roditty [SODA 2011], who for the case of an unweighted and undirected graph presented an approximate $(1+\epsilon)$ algorithm with constant query time and a total update time of $O(n^{2+o(1)})$. This was successively improved, with the state of the art achieving $O(m^{1+o(1)})$ total update time.

All of these algorithms, however, are randomized. In fact, all known improvements over the Even-Shiloach algorithm are randomized. All these algorithms maintain some truncated shortest path trees from a small subset of nodes. While in the randomized setting it is possible to "hide" these nodes from the adversary, in the deterministic setting this is impossible: the adversary can delete all edges touching these nodes, thus forcing the algorithm to choose a new set of nodes and incur a new computation of shortest paths.

In this talk I will present the first *deterministic* decremental SSSP algorithm that breaks the Even-Shiloach bound of $O(mn)$ total update time, for unweighted and undirected graphs.

Joint work with Aaron Bernstein.

15:30 - 15:45 **Coffee break**

15:45 - 16:30 **Guy Even -** *Sublinear Random Access Generators for Preferential Attachment Graphs*

We consider the problem of generating random graphs in evolving random graph models. In the standard approach, the whole graph is chosen randomly according to the distribution of the model before answering queries to the adjacency lists of the graph. Instead, we propose to answer queries by generating the graphs on-the-fly while respecting the probability space of the random graph model.

We focus on two random graph models: the Barabási-Albert Preferential Attachment model (BA-graphs) and the random recursive tree model. We present sublinear randomized generating algorithms for both models. Per query, the running time, the increase in space, and the number of random bits consumed are ploy $\log(n)$ with probability $1 - 1/\text{ploy}(n)$, where n denotes the number of vertices.

This result shows that, although the BA random graph model is defined sequentially, random access is possible without chronological evolution. In addition to a conceptual contribution, on-the-fly generation of random graphs can serve as a tool for simulating sublinear algorithms over large BA-graphs.

Joint work with Reut Levi, Moti Medina, and Adi Rosén.

16:30 - 17:15 **Pierre Fraigniaud -** *Distributed Testing of Excluded Subgraphs*

We study *property testing* in the context of *distributed computing*, under the classical `Congest` model. It is known that testing whether a graph is triangle-free can be done in a constant number of rounds, where the constant depends on how far the input graph is from being triangle-free. We show that, for every connected 4-node graph H, testing whether a graph is H-free can be done in a constant number of rounds too. The constant also depends on how far the input graph is from being H-free, and the dependence is identical to the one in the case of testing triangles. Hence, in particular, testing whether a graph is K_4-free, and testing whether a graph is C_4-free can be done in a constant number of rounds (where K_k denotes the k-node clique, and C_k denotes the k-node cycle). On the other hand, we show that testing K_k-freeness and C_k-freeness for $k \geq 5$ appear to be much harder. Specifically, we investigate two natural types of generic algorithms for testing H-freeness, called DFS tester and BFS tester. The latter captures the previously known algorithm to test the presence of triangles, while the former captures our generic algorithm to test the presence of a 4-node graph pattern H. We prove that both DFS and BFS testers fail to test K_k-freeness and C_k-freeness in a constant number of rounds for $k \geq 5$.

Joint work with Ivan Rapaport, Ville Salo, and Ioan Todinca.

17:15 - 18:30 **Farewell cakes**