Tuesday, June 18

9:00-9:10 — Welcome

9:10-9:40

Best Paper: **Almost Tight Lower Bounds for Hard Cutting Problems in Embedded Graphs**

V. Cohen-Addad, É. Colin de Verdière, D. Marx and A. de Mesmay

We prove essentially tight lower bounds, conditionally to the Exponential Time Hypothesis, for two fundamental but seemingly very different cutting problems on surface-embedded graphs: the Shortest Cut Graph problem and the Multiway Cut problem.

A cut graph of a graph G embedded on a surface S is a subgraph of G whose removal from S leaves a disk. We consider the problem of deciding whether an unweighted graph embedded on a surface of genus g has a cut graph of length at most a given value. We prove a time lower bound for this problem of $n^{\Omega(g/\log g)}$ conditionally to ETH. In other words, the first $n^{O(g)}$-time algorithm by Erickson and Har-Peled [SoCG 2002, Discr. Comput. Geom. 2004] is essentially optimal. We also prove that the problem is W[1]-hard when parameterized by the genus, answering a 17-year old question of these authors.

A multiway cut of an undirected graph G with t distinguished vertices, called *terminals*, is a set of edges whose removal disconnects all pairs of terminals. We consider the problem of deciding whether an unweighted graph G has a multiway cut of weight at most a given value. We prove a time lower bound for this problem of $n^{\Omega(\sqrt{gt+g^2}/\log(gt))}$, conditionally to ETH, for any choice of the genus $g \geq 0$ of the graph and the number of terminals $t \geq 4$. In other words, the algorithm by the second author [Algorithmica 2017] (for the more general multicut problem) is essentially optimal; this extends the lower bound by the third author [ICALP 2012] (for the planar case).

Reductions to planar problems usually involve a grid-like structure. The main novel idea for our results is to understand what structures instead of grids are needed if we want to exploit optimally a certain value g of the genus.

9:40-10:50 — Fast Forward/Coffee Break

Tue-2A: Data Structures I

10:50-11:10

Dynamic Planar Point Location in External Memory

J. I. Munro and Y.Nekrich

In this paper we describe a fully-dynamic data structure for the planar point location problem in the external memory model. Our data structure supports queries in $O(\log_B n(\log\log_B n)^3))$ I/Os and updates in $O(\log_B n(\log\log_B n)^2))$ amortized I/Os, where n is the number of segments in the subdivision and B is the block size. This is the first dynamic data structure with almost-optimal query cost. For comparison all previously known results for this problem require $O(\log_B^2 n)$ I/Os to answer queries. Our result almost matches the best known upper bound in the internal-memory model.

11:10-11:30

A Divide-and-Conquer Algorithm for Two-Point L_1 Shortest Path Queries in Polygonal Domains

Haitao Wang

Let $\mathcal{P}$ be a polygonal domain of h holes and n vertices. We study the problem of constructing a data structure that can compute a shortest path between s and t in $\mathcal{P}$ under the L_1 metric for any two query points s and t. To do so, a standard approach is to first find a set of n_s "gateways" for s and a set of n_t "gateways" for t such that there exist a shortest s-t path containing a gateway of s and a gateway of t, and then compute a shortest s-t path using these gateways. Previous algorithms all take quadratic $O(n_s \cdot n_t)$ time to solve this problem. In this paper, we propose a divide-and-conquer technique that solves the problem in $O(n_s + n_t \log n_s)$ time. As a consequence, we construct a data structure of $O(n + (h^2 \log^3 h/\log\log h))$ size in $O(n + (h^2 \log^4 h/\log\log h))$ time such that each query can be answered in $O(\log n)$ time.

Tue-2B: Persistent Homology I

10:50-11:10

DTM-based Filtrations

H. Anai, F. Chazal, M. Glisse, Y. Ike, H. Inakoshi, R. Tinarrage and Y. Umeda

Despite strong stability properties, the persistent homology of filtrations classically used in Topological Data Analysis, such as, e.g. the Čech or Vietoris-Rips filtrations, are very sensitive to the presence of outliers in the data from which they are computed. In this paper, we introduce and study a new family of filtrations, the DTM-filtrations, built on top of point clouds in the Euclidean space which are more robust to noise and outliers. The approach adopted in this work relies on the notion of distance-to-measure functions and extends some previous work on the approximation of such functions.

11:10-11:30

Topological Data Analysis in Information Space

Herbert Edelsbrunner, Ziga Virk, Hubert Wagner

Various kinds of data are routinely represented as discrete probability distributions. Examples include text documents summarized by histograms of word occurrences and images represented as histograms of oriented gradients. Viewing a discrete probability distribution as a point in the standard simplex of the appropriate dimension, we can understand collections of such objects in geometric and topological terms. Importantly, instead of using the standard Euclidean distance, we look into dissimilarity measures with information-theoretic justification, and we develop the theory needed for applying topological data analysis in this setting. In doing so, we emphasize constructions that enable usage of existing computational topology software in this context.

11:30-11:50	**Maintaining the Union of Unit Discs under Insertions with Near-Optimal Overhead** Pankaj K. Agarwal, Ravid Cohen, Dan Halperin and Wolfgang Mulzer We present efficient data structures for problems on unit discs and arcs of their boundary in the plane. (i) We give an output-sensitive algorithm for the dynamic maintenance of the union of n unit discs under insertions in $O(k \log^2 n)$ update time and $O(n)$ space, where k is the combinatorial complexity of the structural change in the union due to the insertion of the new disc. (ii) As part of the solution of (i) we devise a fully dynamic data structure for the maintenance of lower envelopes of pseudo-lines, which we believe is of independent interest. The structure has $O(\log^2 n)$ update time and $O(\log n)$ vertical ray shooting query time. To achieve this performance, we devise a new algorithm for finding the intersection between two lower envelopes of pseudo-lines in $O(\log n)$ time, using *tentative* binary search; the lower envelopes are special in that at $x = -\infty$ any pseudo-line contributing to the first envelope lies below every pseudo-line contributing to the second envelope. (iii) We also present a dynamic range searching structure for a set of circular arcs of unit radius (not necessarily on the boundary of the union of the corresponding discs), where the ranges are unit discs, with $O(n \log n)$ preprocessing time, $O(n^{1/2+\varepsilon} + \ell)$ query time and $O(\log^2 n)$ amortized update time, where ℓ is the size of the output and for any $\varepsilon > 0$. The structure requires $O(n)$ storage space	**On the Metric Distortion of Embedding Persistence Diagrams into separable Hilbert spaces** M. Carrière and U. Bauer Persistence diagrams are important descriptors in Topological Data Analysis. Due to the nonlinearity of the space of persistence diagrams equipped with their *diagram distances*, most of the recent attempts at using persistence diagrams in machine learning have been done through kernel methods, i.e., embeddings of persistence diagrams into Reproducing Kernel Hilbert Spaces, in which all computations can be performed easily. Since persistence diagrams enjoy theoretical stability guarantees for the diagram distances, the *metric properties* of the feature map, i.e., the relationship between the Hilbert distance and the diagram distances, are of central interest for understanding if the persistence diagram guarantees carry over to the embedding. In this article, we study the possibility of embedding persistence diagrams into separable Hilbert spaces with bi-Lipschitz maps. In particular, we show that for several stable embeddings into infinite-dimensional Hilbert spaces defined in the literature, any lower bound must depend on the cardinalities of the persistence diagrams, and that when the Hilbert space is finite dimensional, finding a bi-Lipschitz embedding is impossible, even when restricting the persistence diagrams to have bounded cardinalities.
11:50-12:00	Break	
	Tue-3A: Combinatorial Geometry I	**Tue-3B: ε-Nets and VC Dimension**
12:00-12:20	**On the Complexity of the k-Level in Arrangements of Pseudoplanes** M. Sharir and C. Ziv A classical open problem in combinatorial geometry is to obtain tight asymptotic bounds on the maximum number of *k-level* vertices in an arrangement of n hyperplanes in $\mathbb{R}^d$ (vertices with exactly k of the hyperplanes passing below them). This is essentially a dual version of the *k-set problem*, which, in a primal setting, seeks bounds for the maximum number of *k-sets* determined by n points in $\mathbb{R}^d$, where a k-set is a subset of size k that can be separated from its complement by a hyperplane. The k-set problem is still wide open even in the plane. In three dimensions, the best known upper and lower bounds are, respectively, $O(nk^{3/2})$ and $nk \cdot 2^{\Omega(\sqrt{\log k})}$. In its dual version, the problem can be generalized by replacing hyperplanes by other families of surfaces (or curves in the planes). Reasonably sharp bounds have been obtained for curves in the plane, but the known upper bounds are rather weak for more general surfaces, already in three dimensions, except for the case of triangles. The best known general bound, due to Chan is $O(n^{2.997})$, for families of surfaces that satisfy certain (fairly weak) properties. In this paper we consider the case of *pseudoplanes* in $\mathbb{R}^3$ (defined in detail in the introduction), and establish the upper bound $O(nk^{5/3})$ for the number of k-level vertices in an arrangement of n pseudoplanes. The bound is obtained by establishing suitable (and nontrivial) extensions of dual versions of classical tools that have been used in studying the primal k-set problem, such as the Lovász Lemma and the Crossing Lemma.	**On weak ε-nets and the Radon number** S. Moran and A. Yehudayoff We show that the Radon number characterizes the existence of weak nets in separable convexity spaces (an abstraction of the Euclidean notion of convexity). The construction of weak nets when the Radon number is finite is based on Helly's property and on metric properties of VC classes. The lower bound on the size of weak nets when the Radon number is large relies on the chromatic number of the Kneser graph. As an application, we prove an amplification result for weak ε-nets.

12:20-12:40	**On grids in point-line arrangements in the plane** M. Mirzaei and A. Suk The famous Szemerédi-Trotter theorem states that any arrangement of n points and n lines in the plane determines $O(n^{4/3})$ incidences, and this bound is tight. In this paper, we prove the following Turán-type result for point-line incidence. Let $\mathcal{L}_a$ and $\mathcal{L}_b$ be two sets of t lines in the plane and let $P = \{\ell_a \cap \ell_b : \ell_a \in \mathcal{L}_a, \ell_b \in \mathcal{L}_b\}$ be the set of intersection points between $\mathcal{L}_a$ and $\mathcal{L}_b$. We say that $(P, \mathcal{L}_a \cup \mathcal{L}_b)$ forms a *natural* $t \times t$ *grid* if $\|P\| = t^2$, and $conv(P)$ does not contain the intersection point of some two lines in $\mathcal{L}_a$ and does not contain the intersection point of some two lines in $\mathcal{L}_b$. For fixed $t > 1$, we show that any arrangement of n points and n lines in the plane that does not contain a natural $t \times t$ grid determines $O(n^{\frac{4}{3}-\varepsilon})$ incidences, where $\varepsilon = \varepsilon(t) > 0$. We also provide a construction of n points and n lines in the plane that does not contain a natural 2×2 grid and determines at least $\Omega(n^{1+\frac{1}{14}})$ incidences.	**Distribution-Sensitive Bounds on Relative Approximations of Geometric Ranges** Y. Tao and Y. Wang A family $\mathcal{R}$ of ranges and a set X of points, all in $\mathbb{R}^d$, together define a range space $(X, \mathcal{R}\|_X)$, where $\mathcal{R}\|_X = \{X \cap h \mid h \in \mathcal{R}\}$. We want to find a structure to estimate the quantity $\|X \cap h\|/\|X\|$ for any range $h \in \mathcal{R}$ with the (ρ, ϵ)-*guarantee*: (i) if $\|X \cap h\|/\|X\| > \rho$, the estimate must have a relative error ϵ; (ii) otherwise, the estimate must have an absolute error $\rho\epsilon$. The objective is to minimize the size of the structure. Currently, the dominant solution is to compute a relative (ρ, ϵ)-approximation, which is a subset of X with $\tilde{O}(\lambda/(\rho\epsilon^2))$ points, where λ is the VC-dimension of $(X, \mathcal{R}\|_X)$, and $\tilde{O}$ hides polylog factors. This paper shows a more general bound sensitive to the content of X. We give a structure that stores $O(\log(1/\rho))$ integers plus $\tilde{O}(\theta \cdot (\lambda/\epsilon^2))$ points of X, where θ — called the *disagreement coefficient* — measures how much the ranges differ from each other in their intersections with X. The value of θ is between 1 and $1/\rho$, such that our space bound is never worse than that of relative (ρ, ϵ)-approximations, but we improve the latter's $1/\rho$ term whenever $\theta = o(\frac{1}{\rho \log(1/\rho)})$. We also prove that, in the worst case, summaries with the $(\rho, 1/2)$-guarantee must consume $\Omega(\theta)$ words even for $d = 2$ and $\lambda \leq 3$. We then constrain $\mathcal{R}$ to be the set of halfspaces in $\mathbb{R}^d$ for a constant d, and prove the existence of structures with $o(1/(\rho\epsilon^2))$ size offering (ρ, ϵ)-guarantees, when X is generated from various stochastic distributions. This is the first formal justification on why the term $1/\rho$ is not compulsory for "realistic" inputs.
12:40-1:00	**The Crossing Tverberg Theorem** R. Fulek and B. Gärtner and A. Kupavskii and P. Valtr and U. Wagner The Tverberg theorem is one of the cornerstones of discrete geometry. It states that, given a set X of at least $(d+1)(r-1)+1$ points in $\mathbb{R}^d$, one can find a partition $X = X_1 \cup \ldots \cup X_r$ of X, such that the convex hulls of the X_i, $i = 1, \ldots, r$, all share a common point. In this paper, we prove a strengthening of this theorem that guarantees a partition which, in addition to the above, has the property that the boundaries of full-dimensional convex hulls have pairwise nonempty intersections. Possible generalizations and algorithmic aspects are also discussed. As a concrete application, we show that any n points in the plane in general position span $\lfloor n/3 \rfloor$ vertex-disjoint triangles that are pairwise crossing, meaning that their boundaries have pairwise nonempty intersections; this number is clearly best possible. A previous result of Alvarez-Rebollar et al. guarantees $\lfloor n/6 \rfloor$ pairwise crossing triangles. Our result generalizes to a result about simplices in $\mathbb{R}^d, d \geq 2$.	**Journey to the Center of the Point Set** S. Har-Peled and M. Jones We revisit an algorithm of Clarkson et al. (Internat. J. Comput. Geom. Appl., 6.03 (1996) 357), that computes (roughly) a $1/(4d^2)$-centerpoint in $\widetilde{O}(d^9)$ time, for a point set in $\Re^d$, where $\widetilde{O}$ hides polylogarithmic terms. We present an improved algorithm that computes (roughly) a $1/d^2$-centerpoint with running time $\widetilde{O}(d^7)$. While the improvements are (arguably) mild, it is the first progress on this well known problem in over twenty years. The new algorithm is simpler, and the running time bound follows by a simple random walk argument, which we believe to be of independent interest. We also present several new applications of the improved centerpoint algorithm.