Wednesday, June 19

	Wed-4A: Smallest Enclosing	**Wed-4B: Persistent Homology II**
9:00-9:20	**Probabilistic Smallest Enclosing Ball in High Dimensions via Subgradient Sampling** A. Krivošija and A. Munteanu We study a variant of the median problem for a collection of point sets in high dimensions. This generalizes the geometric median as well as the (probabilistic) smallest enclosing ball (pSEB) problems. Our main objective and motivation is to improve the previously best algorithm for the pSEB problem by reducing its exponential dependence on the dimension to linear. This is achieved via a novel combination of sampling techniques for clustering problems in metric spaces with the framework of stochastic subgradient descent. As a result, the algorithm becomes applicable to shape fitting problems in Hilbert spaces of unbounded dimension via kernel functions. We present an exemplary application by extending the support vector data description (SVDD) shape fitting method to the probabilistic case. This is done by simulating the pSEB algorithm implicitly in the feature space induced by the kernel function.	**Exact computation of the matching distance on 2-parameter persistence modules** Michael Kerber, Michael Lesnick and Steve Oudot The matching distance is a pseudometric on multi-parameter persistence modules, defined in terms of the weighted bottleneck distance on the restriction of the modules to affine lines. It is known that this distance is stable in a reasonable sense, and can be efficiently approximated, which makes it a promising tool for practical applications. In this work, we show that in the 2-parameter setting, the matching distance can be computed exactly in polynomial time. Our approach subdivides the space of affine lines into regions, via a line arrangement in the dual space. In each region, the matching distance restricts to a simple analytic function, whose maximum is easily computed. As a byproduct, our analysis establishes that the matching distance is a rational number, if the bigrades of the input modules are rational.
9:20-9:40	**Smallest k-Enclosing Rectangle Revisited** T. M. Chan and S. Har-Peled Given a set of n points in the plane, and a parameter k, we consider the problem of computing the minimum (perimeter or area) axis-aligned rectangle enclosing k points. We present the first near quadratic time algorithm for this problem, improving over the previous near-$O(n^{5/2})$-time algorithm by Kaplan, Roy, and Sharir [ESA 2017]. We provide an almost matching conditional lower bound, under the assumption that $(\min, +)$-convolution cannot be solved in truly subquadratic time. Furthermore, we present a new reduction (for either perimeter or area) that can make the time bound sensitive to k, giving near $O(nk)$ time. We also present a near linear time $(1+\varepsilon)$-approximation algorithm to the minimum area of the optimal rectangle containing k points. In addition, we study related problems including the 3-sided, arbitrarily oriented, weighted, and subset sum versions of the problem.	**Chunk Reduction for Multi-Parameter Persistent Homology** U. Fugacci and M. Kerber The extension of persistent homology to multi-parameter setups is an algorithmic challenge. Since most computation tasks scale badly with the size of the input complex, an important pre-processing step consists of simplifying the input while maintaining the homological information. We present an algorithm that drastically reduces the size of an input. Our approach is an extension of the chunk algorithm for persistent homology (Bauer et al., Topological Methods in Data Analysis and Visualization III, 2014). We show that our construction produces the smallest multi-filtered chain complex among all the complexes quasi-isomorphic to the input, improving on the guarantees of previous work in the context of discrete Morse theory. Our algorithm also offers an immediate parallelization scheme in shared memory. Already its sequential version compares favorably with existing simplification schemes, as we show by experimental evaluation.
9:40-10:00	**Computing Shapley Values in the Plane** S. Cabello and T. M. Chan We consider the problem of computing Shapley values for points in the plane, where each point is interpreted as a player, and the value of a coalition is defined by the area of usual geometric objects, such as the convex hull or the minimum axis-parallel bounding box. For sets of n points in the plane, we show how to compute in roughly $O(n^{3/2})$ time the Shapley values for the area of the minimum axis-parallel bounding box and the area of the union of the rectangles spanned by the origin and the input points. When the points form an increasing or decreasing chain, the running time can be improved to near-linear. In all these cases, we use linearity of the Shapley values and algebraic methods. We also show that Shapley values for the area of the convex hull or the minimum enclosing disk can be computed in $O(n^2)$ and $O(n^3)$ time, respectively. These problems are closely related to the model of stochastic point sets considered in computational geometry, but here we have to consider random insertion orders of the points instead of a probabilistic existence of points.	**Computing Persistent Homology of Flag Complexes via Strong Collapses** J-D. Boissonnat and S. Pritam In this article, we focus on the problem of computing Persistent Homology of a flag tower, i.e. a sequence of flag complexes connected by simplicial maps. We show that if we restrict the class of simplicial complexes to flag complexes, we can achieve decisive improvement in terms of time and space complexities with respect to previous work. We show that strong collapses of flag complexes can be computed in time $O(k^2 v^2)$ where v is the number of vertices of the complex and k is the maximal degree of its graph. Moreover we can strong collapse a flag complex knowing only its 1-skeleton and the resulting complex is also a flag complex. When we strong collapse the complexes in a flag tower, we obtain a reduced sequence that is also a flag tower we call the core flag tower. We then convert the core flag tower to an equivalent filtration to compute its PH. Here again, we only use the 1-skeletons of the complexes. The resulting method is simple and extremely efficient.

10:00-10:30	Coffee Break	
	Wed-5A: Combinatorial Geometry II	**Wed-5B: Optimization and Approximation**
10:30-10:50	**Ham-Sandwich cuts and center transversals in subspaces** Patrick Schnider The Ham-Sandwich theorem is a well-known result in geometry. It states that any d mass distributions in $\mathbb{R}^d$ can be simultaneously bisected by a hyperplane. The result is tight, that is, there are examples of $d+1$ mass distributions that cannot be simultaneously bisected by a single hyperplane. In this abstract we will study the following question: given a continuous assignment of mass distributions to certain subsets of $\mathbb{R}^d$, is there a subset on which we can bisect more masses than what is guaranteed by the Ham-Sandwich theorem? which we answer in the affirmative. We investigate two types of subsets. The first type are linear subspaces of $\mathbb{R}^d$, i.e., k-dimensional flats containing the origin. We show that for any continuous assignment of d mass distributions to the k-dimensional linear subspaces of $\mathbb{R}^d$, there is always a subspace on which we can simultaneously bisect the images of all d assignments. We extend this result to center transversals, a generalization of Ham-Sandwich cuts. As for Ham-Sandwich cuts, we further show that for $d-k+2$ masses, we can choose $k-1$ of the vectors defining the k-dimensional subspace in which the solution lies. The second type of subsets we consider are subsets that are determined by families of n hyperplanes in $\mathbb{R}^d$. Also in this case, we find a Ham-Sandwich-type result. In an attempt to solve a conjecture by Langerman about bisections with several cuts, we show that our underlying topological result can be used to prove this conjecture in a relaxed setting.	**Packing Disks into Disks with Optimal Worst-Case Density** S. P. Fekete and P. Keldenich and C. Scheffer We provide a tight result for a fundamental problem arising from packing disks into a circular container: The critical density of packing disks in a disk is 0.5. This implies that any set of (not necessarily equal) disks of total area $\delta \leq 1/2$ can always be packed into a disk of area 1; on the other hand, for any $\varepsilon > 0$ there are sets of disks of area $1/2+\varepsilon$ that cannot be packed. The proof uses a careful manual analysis, complemented by a minor automatic part that is based on interval arithmetic. Beyond the basic mathematical importance, our result is also useful as a blackbox lemma for the analysis of recursive packing algorithms.
10:50-11:10	**On the chromatic number of disjointness graphs of curves** János Pach and István Tomon Let $\omega(G)$ and $\chi(G)$ denote the clique number and chromatic number of a graph G, respectively. The disjointness graph of a family of curves (continuous arcs in the plane) is the graph whose vertices correspond to the curves and in which two vertices are joined by an edge if and only if the corresponding curves are disjoint. A curve is called x-monotone if every vertical line intersects it in at most one point. An x-monotone curve is grounded if its left endpoint lies on the y-axis. We prove that if G is the disjointness graph of a family of grounded x-monotone curves such that $\omega(G)=k$, then $\chi(G) \leq \binom{k+1}{2}$. If we only require that every curve is x-monotone and intersects the y-axis, then we have $\chi(G) \leq \frac{k+1}{2}\binom{k+2}{3}$. Both of these bounds are best possible. The construction showing the tightness of the last result settles a 25 years old problem: it yields that there exist K_k-free disjointness graphs of x-monotone curves such that any proper coloring of them uses at least $\Omega(k^4)$ colors. This matches the upper bound up to a constant factor.	**Preconditioning for the Geometric Transportation Problem** A. B. Khesin, A. Nikolov, and D. Paramonov In the geometric transportation problem, we are given a collection of points P in d-dimensional Euclidean space, and each point is given a supply of $\mu(p)$ units of mass, where $\mu(p)$ could be a positive or a negative integer, and the total sum of the supplies is 0. The goal is to find a flow (called a transportation map) that transports $\mu(p)$ units from any point p with $\mu(p) > 0$, and transports $-\mu(p)$ units into any point p with $\mu(p) < 0$. Moreover, the flow should minimize the total distance traveled by the transported mass. The optimal value is known as the transportation cost, or the Earth Mover's Distance (from the points with positive supply to those with negative supply). This problem has been widely studied in many fields of computer science: from theoretical work in computational geometry, to applications in computer vision, graphics, and machine learning. In this work we study approximation algorithms for the geometric transportation problem. We give an algorithm which, for any fixed dimension d, finds a $(1+\varepsilon)$-approximate transportation map in time nearly-linear in n, and polynomial in ε^{-1} and in the logarithm of the total supply. This is the first approximation scheme for the problem whose running time depends on n as $n \cdot \text{polylog}(n)$. Our techniques combine the generalized preconditioning framework of Sherman [SODA 2017], which is grounded in continuous optimization, with simple geometric arguments to first reduce the problem to a minimum cost flow problem on a sparse graph, and then to design a good preconditioner for this latter problem.

<table>
<tr><td>11:10-11:30</td><td>Semi-algebraic colorings of complete graphs
J. Fox, J. Pach, and A. Suk

We consider m-colorings of the edges of a complete graph, where each color class is defined semi-algebraically with bounded complexity. The case $m = 2$ was first studied by Alon et al., who applied this framework to obtain surprisingly strong Ramsey-type results for intersection graphs of geometric objects and for other graphs arising in computational geometry. Considering larger values of m is relevant, e.g., to problems concerning the number of distinct distances determined by a point set.

For $p \geq 3$ and $m \geq 2$, the classical Ramsey number $R(p; m)$ is the smallest positive integer n such that any m-coloring of the edges of K_n, the complete graph on n vertices, contains a monochromatic K_p. It is a longstanding open problem that goes back to Schur (1916) to decide whether $R(p; m) = 2^{O(m)}$, for a fixed p. We prove that this is true if each color class is defined semi-algebraically with bounded complexity, and that the order of magnitude of this bound is tight. Our proof is based on the Cutting Lemma of Chazelle et al., and on a Szemerédi-type regularity lemma for multicolored semi-algebraic graphs, which is of independent interest. The same technique is used to address the semi-algebraic variant of a more general Ramsey-type problem of Erdős and Shelah.</td><td>Algorithms for Metric Learning via Contrastive Embeddings
D. Ihara, N. Mohammadi and A. Sidiropoulos

We study the problem of supervised learning a metric space under discriminative constraints. Given a universe X and sets $\mathcal{S}, \mathcal{D} \subset \binom{X}{2}$ of similar and dissimilar pairs, we seek to find a mapping $f : X \to Y$, into some target metric space $M = (Y, \rho)$, such that similar objects are mapped to points at distance at most u, and dissimilar objects are mapped to points at distance at least ℓ. More generally, the goal is to find a mapping of maximum accuracy (that is, fraction of correctly classified pairs). We propose approximation algorithms for various versions of this problem, for the cases of Euclidean and tree metric spaces. For both of these target spaces, we obtain fully polynomial-time approximation schemes (FPTAS) for the case of perfect information. In the presence of imperfect information we present approximation algorithms that run in quasi-polynomial time (QPTAS). We also present an exact algorithm for learning line metric spaces with perfect information in polynomial time. Our algorithms use a combination of tools from metric embeddings and graph partitioning, that could be of independent interest.</td></tr>
<tr><td>11:30-11:40</td><td colspan="2">Break</td></tr>
<tr><td>11:40-12:40</td><td colspan="2">Invited Talk: A Geometric Data Structure from Neuroscience
Sanjoy Dasgupta

Abstract: An intriguing geometric primitive, "expand-and-sparsify", has been found in the olfactory system of the fly and several other organisms. It maps an input vector to a much higher-dimensional sparse representation, using a random linear transformation followed by winner-take-all thresholding.

I'll show that this representation has a variety of formal properties, such as locality preservation, that make it an attractive data structure for algorithms and machine learning. In particular, mimicking the fly's circuitry yields algorithms for similarity search and for novelty detection that have provable guarantees as well as having practical performance that is competitive with state-of-the-art methods.

This talk is based on work with Saket Navlakha (Salk Institute), Chuck Stevens (Salk Institute), and Chris Tosh (Columbia).

Bio: Sanjoy Dasgupta is a Professor of Computer Science and Engineering at UC San Diego, where he has been since 2002. He works on algorithmic statistics, with a particular focus on unsupervised and minimally supervised learning. He is author of a textbook, "Algorithms" (with Christos Papadimitriou and Umesh Vazirani).</td></tr>
</table>

See next page ...

12:40-2:30	Lunch on Your Own	
	Wed-6A: Graph Drawing I	**Wed-6B: Matching and Partitioning**
2:30-2:50	**Efficient Algorithms for Ortho-Radial Graph Drawing** B. Niedermann, I. Rutter, and M. Wolf Orthogonal drawings, i.e., embeddings of graphs into grids, are a classic topic in Graph Drawing. Often the goal is to find a drawing that minimizes the number of bends on the edges. A key ingredient for bend minimization algorithms is the existence of an *orthogonal representation* that allows to describe such drawings purely combinatorially by only listing the angles between the edges around each vertex and the directions of bends on the edges, but neglecting any kind of geometric information such as vertex coordinates or edge lengths. Barth et al. [2017] have established the existence of an analogous *ortho-radial representation* for *ortho-radial drawings*, which are embeddings into an ortho-radial grid, whose gridlines are concentric circles around the origin and straight-line spokes emanating from the origin but excluding the origin itself. While any orthogonal representation admits an orthogonal drawing, it is the circularity of the ortho-radial grid that makes the problem of characterizing valid ortho-radial representations all the more complex and interesting. Barth et al. prove such a characterization. However, the proof is existential and does not provide an efficient algorithm for testing whether a given ortho-radial representation is valid, let alone actually obtaining a drawing from an ortho-radial representation. In this paper we give quadratic-time algorithms for both of these tasks. They are based on a suitably constrained left-first DFS in planar graphs and several new insights on ortho-radial representations. Our validity check requires quadratic time, and a naive application of it would yield a quartic algorithm for constructing a drawing from a valid ortho-radial representation. Using further structural insights we speed up the drawing algorithm to quadratic running time.	**A Weighted Approach to the Maximum Cardinality Bipartite Matching Problem with Applications in Geometric Settings** N. Lahn and S. Raghvendra We present a weighted approach to compute a maximum cardinality matching in an arbitrary bipartite graph. Our main result is a new algorithm that takes as input a weighted bipartite graph $G(A \cup B, E)$ with edge weights of 0 or 1. Let $w \leq n$ be an upper bound on the weight of any matching in G. Consider the subgraph induced by all the edges of G with a weight 0. Suppose every connected component in this subgraph has $O(r)$ vertices and $O(mr/n)$ edges. We present an algorithm to compute a maximum cardinality matching in G in $\tilde{O}(m(\sqrt{w} + \sqrt{r} + \frac{wr}{n}))$ time. When all the edge weights are 1 (symmetrically when all weights are 0), our algorithm will be identical to the well-known Hopcroft-Karp (HK) algorithm, which runs in $O(m\sqrt{n})$ time. However, if we can carefully assign weights of 0 and 1 on its edges such that both w and r are sub-linear in n and $wr = O(n^{\gamma})$ for $\gamma < 3/2$, then we can compute maximum cardinality matching in G in $o(m\sqrt{n})$ time. Using our algorithm, we obtain a new $\tilde{O}(n^{4/3}/\varepsilon^4)$ time algorithm to compute an ε-approximate bottleneck matching of $A, B \subset \mathbb{R}^2$ and an $\frac{1}{\varepsilon^{O(d)}} n^{1+\frac{d-1}{2d-1}}$ poly log n time algorithm for computing ε-approximate bottleneck matching in d-dimensions. All previous algorithms take $\Omega(n^{3/2})$ time. Given any graph $G(A \cup B, E)$ that has an easily computable balanced vertex separator for every subgraph $G'(V', E')$ of size $\|V'\|^{\delta}$, for $\delta \in [1/2, 1)$, we can apply our algorithm to compute a maximum matching in $\tilde{O}(mn^{\frac{\delta}{1+\delta}})$ time improving upon the $O(m\sqrt{n})$ time taken by the HK-Algorithm.
2:50-3:10	**Bounded degree conjecture holds precisely for c-crossing-critical graphs with $c \leq 12$** D. Bokal, Z. Dvořák, P. Hliněný, J. Leaños, B. Mohar, T. Wiedera We study c-crossing-critical graphs, which are the minimal graphs that require at least c edge-crossings when drawn in the plane. For every fixed pair of integers with $c \geq 13$ and $d \geq 1$, we give first explicit constructions of c-crossing-critical graphs containing a vertex of degree greater than d. We also show that such unbounded degree constructions do not exist for $c \leq 12$, precisely, that there exists a constant D such that every c-crossing-critical graph with $c \leq 12$ has maximum degree at most D. Hence, the bounded maximum degree conjecture of c-crossing-critical graphs, which was generally disproved in 2010 by Dvořák and Mohar (without an explicit construction), holds true, surprisingly, exactly for the values $c \leq 12$.	**An Efficient Algorithm for Generalized Polynomial Partitioning and Its Applications** P. K. Agarwal, B. Aronov, E. Ezra, and J. Zahl In 2015, Guth proved that if S is a collection of n g-dimensional semi-algebraic sets in $\mathbb{R}^d$ and if $D \geq 1$ is an integer, then there is a d-variate polynomial P of degree at most D so that each connected component of $\mathbb{R}^d \setminus Z(P)$ intersects $O(n/D^{d-g})$ sets from S. Such a polynomial is called a *generalized partitioning polynomial*. We present a randomized algorithm that computes such polynomials efficiently—the expected running time of our algorithm is linear in $\|S\|$. Our approach exploits the technique of *quantifier elimination* combined with that of *ε-samples*. We present four applications of our result. The first is a data structure for answering point-enclosure queries among a family of semi-algebraic sets in $\mathbb{R}^d$ in $O(\log n)$ time, with storage complexity and expected preprocessing time of $O(n^{d+\varepsilon})$. The second is a data structure for answering range search queries with semi-algebraic ranges in $O(\log n)$ time, with $O(n^{t+\varepsilon})$ storage and expected preprocessing time, where $t > 0$ is an integer that depends on d and the description complexity of the ranges. The third is a data structure for answering vertical ray-shooting queries among semi-algebraic sets in $\mathbb{R}^d$ in $O(\log^2 n)$ time, with $O(n^{d+\varepsilon})$ storage and expected preprocessing time. The fourth is an efficient algorithm for cutting algebraic planar curves into pseudo-segments.

3:10-3:30	**$\mathbb{Z}_2$-Genus of Graphs and Minimum Rank of Partial Symmetric Matrices** R. Fulek and J. Kynčl The *genus* $g(G)$ of a graph G is the minimum g such that G has an embedding on the orientable surface M_g of genus g. A drawing of a graph on a surface is *independently even* if every pair of nonadjacent edges in the drawing crosses an even number of times. The $\mathbb{Z}_2$*-genus* of a graph G, denoted by $g_0(G)$, is the minimum g such that G has an independently even drawing on M_g. By a result of Battle, Harary, Kodama and Youngs from 1962, the graph genus is additive over 2-connected blocks. In 2013, Schaefer and Štefankovič proved that the $\mathbb{Z}_2$-genus of a graph is additive over 2-connected blocks as well, and asked whether this result can be extended to so-called 2-amalgamations, as an analogue of results by Decker, Glover, Huneke, and Stahl for the genus. We give the following partial answer. If $G = G_1 \cup G_2$, G_1 and G_2 intersect in two vertices u and v, and $G - u - v$ has k connected components (among which we count the edge uv if present), then $\|g_0(G) - (g_0(G_1) + g_0(G_2))\| \leq k + 1$. For complete bipartite graphs $K_{m,n}$, with $n \geq m \geq 3$, we prove that $\frac{g_0(K_{m,n})}{g(K_{m,n})} = 1 - O(\frac{1}{n})$. Similar results are proved also for the Euler $\mathbb{Z}_2$-genus. We express the $\mathbb{Z}_2$-genus of a graph using the minimum rank of partial symmetric matrices over $\mathbb{Z}_2$; a problem that might be of independent interest.	**Efficient Algorithms for Geometric Partial Matching** Pankaj K. Agarwal, Hsien-Chih Chang, Allen Xiao Let A and B be two point sets in the plane of sizes r and n respectively (assume $r \leq n$), and let k be a parameter.A matching between A and B is a family of pairs in $A \times B$ so that any point of $A \cup B$ appears in at most one pair. Given two positive integers p and q, we define the cost of matching M to be $c(M) = \sum_{(a,b)\in M} \|a - b\|_p^q$ where $\|\cdot\|_p$ is the L_p-norm. The geometric partial matching problem asks to find the minimum-cost size-k matching between A and B. We present efficient algorithms for geometric partial matching problem that work for any powers of L_p-norm matching objective: An exact algorithm that runs in $O((n + k^2) \operatorname{polylog} n)$ time, and a $(1 + \varepsilon)$-approximation algorithm that runs in $O((n + k\sqrt{k}) \operatorname{polylog} n \cdot \log \varepsilon^{-1})$ time. Both algorithms are based on the primal-dual flow augmentation scheme; the main improvements involve using dynamic data structures to achieve efficient flow augmentations. With similar techniques, we give an exact algorithm for the planar transportation problem running in $O(\min\{n^2, rn^{3/2}\} \operatorname{polylog} n)$ time.
3:30-4:00	Coffee/Snack Break	
	Wed-7A: Topology	**Wed-7B: Algorithm Complexity**
4:00-4:20	**Topologically Trivial Closed Walks in Directed Surface Graphs** Jeff Erickson and Yipu Wang Let G be a directed graph with n vertices and m edges, embedded on a surface S, possibly with boundary, with first Betti number β. We consider the complexity of finding closed directed walks in G that are either contractible (trivial in homotopy) or bounding (trivial in integer homology) in S. Specifically, we describe algorithms to determine whether G contains a simple contractible cycle in $O(n + m)$ time, or a contractible closed walk in $O(n + m)$ time, or a bounding closed walk in $O(\beta(n + m))$ time. Our algorithms rely on subtle relationships between strong connectivity in G and in the dual graph $G^\star$; our contractible-closed-walk algorithm also relies on a seminal topological result of Hass and Scott. We also prove that detecting simple bounding cycles is NP-hard. We also describe three polynomial-time algorithms to compute shortest contractible closed walks, depending on whether the fundamental group of the surface is free, abelian, or hyperbolic. A key step in our algorithm for hyperbolic surfaces is the construction of a context-free grammar with $O(g^2L^2)$ non-terminals that generates all contractible closed walks of length at most L, and only contractible closed walks, in a system of quads of genus $g \geq 2$. Finally, we show that computing shortest simple contractible cycles, shortest simple bounding cycles, and shortest bounding closed walks are all NP-hard.	**The One-Way Communication Complexity of Dynamic Time Warping Distance** V. Braverman, M. Charikar, W. Kuszmaul, D. P. Woodruff, and L. F. Yang We resolve the randomized one-way communication complexity of Dynamic Time Warping (DTW) distance. We show that there is an efficient one-way communication protocol using $\widetilde{O}(n/\alpha)$ bits for the problem of computing an α-approximation for DTW between strings x and y of length n, and we prove a lower bound of $\Omega(n/\alpha)$ bits for the same problem. Our communication protocol works for strings over an arbitrary metric of polynomial size and aspect ratio, and we optimize the logarithmic factors depending on properties of the underlying metric, such as when the points are low-dimensional integer vectors equipped with various metrics or have bounded doubling dimension. We also consider linear sketches of DTW, showing that such sketches must have size $\Omega(n)$.

4:20-4:40	**3-Manifold Triangulations with Small Treewidth** K. Huszár and J. Spreer Motivated by fixed-parameter tractable (FPT) problems in computational topology, we consider the treewidth $\mathrm{tw}(\mathcal{M})$ of a compact, connected 3-manifold $\mathcal{M}$, defined to be the minimum treewidth of the face pairing graph of any triangulation $\mathcal{T}$ of $\mathcal{M}$. In this setting the relationship between the topology of a 3-manifold and its treewidth is of particular interest. First, as a corollary of work of Jaco and Rubinstein, we prove that for any closed, orientable 3-manifold $\mathcal{M}$ the treewidth $\mathrm{tw}(\mathcal{M})$ is at most $4\mathfrak{g}(\mathcal{M}) - 2$, where $\mathfrak{g}(\mathcal{M})$ denotes Heegaard genus of $\mathcal{M}$. In combination with our earlier work with Wagner, this yields that for non-Haken manifolds the Heegaard genus and the treewidth are within a constant factor. Second, we characterize all 3-manifolds of treewidth one: These are precisely the lens spaces and a single other Seifert fibered space. Furthermore, we show that all remaining orientable Seifert fibered spaces over the 2-sphere or a non-orientable surface have treewidth two. In particular, for every spherical 3-manifold we exhibit a triangulation of treewidth at most two. Our results further validate the parameter of treewidth (and other related parameters such as cutwidth or congestion) to be useful for topological computing, and also shed more light on the scope of existing FPT-algorithms in the field.	**Upward Book Embeddings of st-Graphs** C. Binucci, G. Da Lozzo, E. Di Giacomo, W. Didimo, T. Mchedlidze, M. Patrignani We study *k-page upward book embeddings* (kUBEs) of *st*-graphs, that is, book embeddings of single-source single-sink directed acyclic graphs on k pages with the additional requirement that the vertices of the graph appear in a topological ordering along the spine of the book. We show that testing whether a graph admits a kUBE is NP-complete for $k \geq 3$. A hardness result for this problem was previously known only for $k = 6$ [Heath and Pemmaraju, 1999]. Motivated by this negative result, we focus our attention on $k = 2$. On the algorithmic side, we present polynomial-time algorithms for testing the existence of 2UBEs of planar *st*-graphs with branchwidth β and of plane *st*-graphs whose faces have a special structure. These algorithms run in $O(f(\beta) \cdot n + n^3)$ time and $O(n)$ time, respectively, where f is a singly-exponential function on β. Moreover, on the combinatorial side, we present two notable families of plane *st*-graphs that always admit an embedding-preserving 2UBE.
4:40-5:00	**When Convexity Helps Collapsing Complexes** D. Attali, A. Lieutier, and D. Salinas This paper illustrates how convexity hypotheses help collapsing simplicial complexes. We first consider a collection of compact convex sets and show that the nerve of the collection is collapsible whenever the union of sets in the collection is convex. We apply this result to prove that the Delaunay complex of a finite point set is collapsible. We then consider a convex domain defined as the convex hull of a finite point set. We show that if the point set samples sufficiently densely the domain, then both the Čech complex and the Rips complex of the point set are collapsible for a well-chosen scale parameter. A key ingredient in our proofs consists in building a filtration by sweeping space with a growing sphere whose center has been fixed and studying events occurring through the filtration. Since the filtration mimics the sublevel sets of a Morse function with a single critical point, we anticipate this work to lay the foundations for a non-smooth, discrete Morse Theory.	