Thursday, June 20

	Thu-8A: Contact and Surface Graphs	Thu-8B: Frechet Distance
9:00-9:20	**Near-optimal Algorithms for Shortest Paths in Weighted Unit-Disk Graphs** H. Wang, J. Xue We revisit a classical graph-theoretic problem, the *single-source shortest-path* (SSSP) problem, in weighted unit-disk graphs. We first propose an exact (and deterministic) algorithm which solves the problem in $O(n \log^2 n)$ time using linear space, where n is the number of the vertices of the graph. This significantly improves the previous deterministic algorithm by Cabello and Jejčič [CGTA'15] which uses $O(n^{1+\delta})$ time and $O(n^{1+\delta})$ space (for any small constant $\delta > 0$) and the previous randomized algorithm by Kaplan et al. [SODA'17] which uses $O(n \log^{12+o(1)} n)$ expected time and $O(n \log^3 n)$ space. More specifically, we show that if the 2D offline insertion-only (additively-)weighted nearest-neighbor problem with k operations (i.e., insertions and queries) can be solved in $f(k)$ time, then the SSSP problem in weighted unit-disk graphs can be solved in $O(n \log n + f(n))$ time. Using the same framework with some new ideas, we also obtain a $(1+\varepsilon)$-approximate algorithm for the problem, using $O(n \log n + n \log^2(1/\varepsilon))$ time and linear space. This improves the previous $(1+\varepsilon)$-approximate algorithm by Chan and Skrepetos [SoCG'18] which uses $O((1/\varepsilon)^2 n \log n)$ time and $O((1/\varepsilon)^2 n)$ space. Because of the $\Omega(n \log n)$-time lower bound of the problem (even when approximation is allowed), both of our algorithms are almost optimal.	**The VC Dimension of Metric Balls under Fréchet and Hausdorff Distances** A. Driemel, J. M. Phillips, I. Psarros The Vapnik-Chervonenkis dimension provides a notion of complexity for systems of sets. If the VC dimension is small, then knowing this can drastically simplify fundamental computational tasks such as classification, range counting, and density estimation through the use of sampling bounds. We analyze set systems where the ground set X is a set of polygonal curves in $\mathbb{R}^d$ and the sets $\mathcal{R}$ are metric balls defined by curve similarity metrics, such as the Fréchet distance and the Hausdorff distance, as well as their discrete counterparts. We derive upper and lower bounds on the VC dimension that imply useful sampling bounds in the setting that the number of curves is large, but the complexity of the individual curves is small. Our upper bounds are either near-quadratic or near-linear in the complexity of the curves that define the ranges and they are logarithmic in the complexity of the curves that define the ground set.
9:20-9:40	**Morphing Contact Representations of Graphs** Patrizio Angelini, Steven Chaplick, Sabine Cornelsen, Giordano Da Lozzo, Vincenzo Roselli We consider the problem of morphing between contact representations of a plane graph. In a contact representation of a plane graph, vertices are realized by internally disjoint elements from a family of connected geometric objects. Two such elements touch if and only if their corresponding vertices are adjacent. These touchings also induce the same embedding as in the graph. In a morph between two contact representations we insist that at each time step (continuously throughout the morph) we have a contact representation of the same type. We focus on the case when the geometric objects are triangles that are the lower-right half of axis-parallel rectangles. Such RT-representations exist for every plane graph and right triangles are one of the simplest families of shapes supporting this property. Thus, they provide a natural case to study regarding morphs of contact representations of plane graphs. We study piecewise linear morphs, where each step is a linear morph moving the endpoints of each triangle at constant speed along straight-line trajectories. We provide a polynomial-time algorithm that decides whether there is a piecewise linear morph between two RT-representations of a plane triangulation, and, if so, computes a morph with a quadratic number of linear morphs. As a direct consequence, we obtain that for 4-connected plane triangulations there is a morph between every pair of RT-representations where the "top-most" triangle in both representations corresponds to the same vertex. This shows that the realization space of such RT-representations of any 4-connected plane triangulation forms a connected set.	**Walking the Dog Fast in Practice: Algorithm Engineering of the Fréchet Distance** K. Bringmann, M. Künnemann and A. Nusser The Fréchet distance provides a natural and intuitive measure for the popular task of computing the similarity of two (polygonal) curves. While a simple algorithm computes it in near-quadratic time, a strongly subquadratic algorithm cannot exist unless the Strong Exponential Time Hypothesis fails. Still, fast practical implementations of the Fréchet distance, in particular for realistic input curves, are highly desirable. This has even lead to a designated competition, the ACM SIGSPATIAL GIS Cup 2017: Here, the challenge was to implement a near-neighbor data structure under the Fréchet distance. The bottleneck of the top three implementations turned out to be precisely the decision procedure for the Fréchet distance. In this work, we present a fast, certifying implementation for deciding the Fréchet distance, in order to (1) complement its pessimistic worst-case hardness by an empirical analysis on realistic input data and to (2) improve the state of the art for the GIS Cup challenge. We experimentally evaluate our implementation on a large benchmark consisting of several data sets (including handwritten characters and GPS trajectories). Compared to the winning implementation of the GIS Cup, we obtain running time improvements of up to more than two orders of magnitude for the decision procedure and of up to a factor of 30 for queries to the near-neighbor data structure.

9:40-10:00

Lower Bounds for Electrical Reduction on Surfaces
Hsien-Chih Chang, Marcos Cossarini, Jeff Erickson

We strengthen the connections between *electrical transformations* and *homotopy* from the planar setting—observed and studied since Steinitz—to arbitrary surfaces with punctures. As a result, we improve our earlier lower bound on the number of electrical transformations required to reduce an n-vertex graph on surface in the worst case [SOCG 2016] in two different directions. Our previous $\Omega(n^{3/2})$ lower bound applies only to *facial* electrical transformations on plane graphs with *no terminals*. First we provide a stronger $\Omega(n^2)$ lower bound when the planar graph has two or more terminals, which follows from a quadratic lower bound on the number of homotopy moves in the annulus. Our second result extends our earlier $\Omega(n^{3/2})$ lower bound to the wider class of *planar* electrical transformations, which preserve the planarity of the graph but may delete cycles that are not faces of the given embedding. This new lower bound follows from the observation that the *defect* of the medial graph of a planar graph is the same for all its planar embeddings.

Polyline Simplification has Cubic Complexity
K. Bringmann and B. R. Chaudhury

In the classic polyline simplification problem we want to replace a given polygonal curve P, consisting of n vertices, by a subsequence P' of k vertices from P such that the polygonal curves P and P' are 'close'. Closeness is usually measured using the Hausdorff or Fréchet distance. These distance measures can be applied *globally*, i.e., to the whole curves P and P', or *locally*, i.e., to each simplified subcurve and the line segment that it was replaced with separately (and then taking the maximum). We provide an $O(n^3)$ time algorithm for simplification under Global-Fréchet distance, improving the previous best algorithm by a factor of $\Omega(kn^2)$. We also provide evidence that in high dimensions cubic time is essentially optimal for all three problems (Local-Hausdorff, Local-Fréchet, and Global-Fréchet). Specifically, improving the cubic time to $O(n^{3-\epsilon}\mathrm{poly}(d))$ for polyline simplification over $(\mathbb{R}^d, L_p)$ for $p = 1$ would violate plausible conjectures. We obtain similar results for all $p \in [1, \infty), p \neq 2$. In total, in high dimensions and over general L_p-norms we resolve the complexity of polyline simplification with respect to Local-Hausdorff, Local-Fréchet, and Global-Fréchet, by providing new algorithms and conditional lower bounds.

10:00-10:30

Coffee Break

Thu-9A: Geometric Data Structures

Thu-9B: Robotics and Geometric Structures

10:30-10:50

A Spanner for the Day After
K. Buchin, S. Har-Peled and D. Oláh

We show how to construct $(1+\varepsilon)$-spanner over a set P of n points in $\mathbb{R}^d$ that is resilient to a catastrophic failure of nodes. Specifically, for prescribed parameters $\vartheta, \varepsilon \in (0,1)$, the computed spanner G has $O\left(\varepsilon^{-c}\vartheta^{-6} n \log n (\log\log n)^6\right)$ edges, where $c = O(d)$. Furthermore, for *any* k, and *any* deleted set $B \subseteq P$ of k points, the residual graph $G \setminus B$ is $(1+\varepsilon)$-spanner for all the points of P except for $(1+\vartheta)k$ of them. No previous constructions, beyond the trivial clique with $O(n^2)$ edges, were known such that only a tiny additional fraction (i.e., ϑ) lose their distance preserving connectivity.

Our construction works by first solving the exact problem in one dimension, and then showing a surprisingly simple and elegant construction in higher dimensions, that uses the one-dimensional construction in a black box fashion.

General techniques for approximate incidences and their application to the camera posing problem
D. Aiger, H. Kaplan, E. Kokiopoulou, M. Sharir, B. Zeisl

We consider the classical camera pose estimation problem that arises in many computer vision applications, in which we are given n 2D-3D correspondences between points in the scene and points in the camera image (some of which are incorrect associations), and where we aim to determine the camera pose (the position and orientation of the camera in the scene) from this data. We demonstrate that this posing problem can be reduced to the problem of computing ε-approximate incidences between two-dimensional surfaces (derived from the input correspondences) and points (on a grid) in a four-dimensional pose space. Similar reductions can be applied to other camera pose problems, as well as to similar problems in related application areas.

We describe and analyze three techniques for solving the resulting ε-approximate incidences problem in the context of our camera posing application. The first is a straightforward assignment of surfaces to the cells of a grid (of side-length ε) that they intersect. The second is a variant of a primal-dual technique, recently introduced by a subset of the authors [ESA17] for different (and simpler) applications. The third is a non-trivial generalization of a data structure Fonseca and Mount [CGTA2010], originally designed for the case of hyperplanes. We present and analyze this technique in full generality, and then apply it to the camera posing problem at hand.

We compare our methods experimentally on real and synthetic data. Our experiments show that for the typical values of n and ε, the primal-dual method is the fastest, also in practice.

10:50-11:10	**Searching for the Closest-pair in a Query Translate** J. Xue, Y. Li, S. Rahul, R. Janardan We consider a range-search variant of the closest-pair problem. Let Γ be a fixed shape in the plane. We are interested in storing a given set of n points in the plane in some data structure such that for any specified translate of Γ, the closest pair of points contained in the translate can be reported efficiently. We present results on this problem for two important settings: when Γ is a polygon (possibly with holes) and when Γ is a general convex body whose boundary is smooth. When Γ is a polygon, we present a data structure using $O(n)$ space and $O(\log n)$ query time, which is asymptotically optimal. When Γ is a general convex body with a smooth boundary, we give a near-optimal data structure using $O(n \log n)$ space and $O(\log^2 n)$ query time. Our results settle some open questions posed by Xue et al. at SoCG 2018.	**Rods and Rings: Soft Subdivision Planner for $\mathbb{R}^3 \times S^2$** C.-H. Hsu, Y.-J. Chiang and C. Yap We consider path planning for a rigid spatial robot moving amidst polyhedral obstacles. Our robot is either a rod or a ring. Being axially-symmetric, their configuration space is $\mathbb{R}^3 \times S^2$ with 5 degrees of freedom (DOF). Correct, complete and practical path planning for such robots is a long standing challenge in robotics. While the rod is one of the most widely studied spatial robots in path planning, the ring seems to be new, and a rare example of a non-simply-connected robot. This work provides rigorous and complete algorithms for these robots with theoretical guarantees. We implemented the algorithms in our open-source Core Library. Experiments show that they are practical, achieving near real-time performance. We compared our planner to state-of-the-art sampling planners in the Open Motion Planning Library (OMPL). Our subdivision path planner is based on the twin foundations of ε-exactness and soft predicates. Correct implementation is relatively easy. The technical innovations include subdivision atlases for S^2, introduction of Σ_2 representations for footprints, and extensions of our feature-based technique for “opening up the blackbox of collision detection”.
11:10-11:30	**Preprocessing Ambiguous Imprecise Points** I. van der Hoog, I. Kostitsyna, M. Löffler, B. Speckmann Let $\mathcal{R} = \{R_1, R_2, \ldots, R_n\}$ be a set of regions and let $X = \{x_1, x_2, \ldots, x_n\}$ be an (unknown) point set with $x_i \in R_i$. Region R_i represents the uncertainty region of x_i. We consider the following question: how fast can we establish order if we are allowed to preprocess the regions in $\mathcal{R}$? The *preprocessing model* of uncertainty uses two consecutive phases: a preprocessing phase which has access only to $\mathcal{R}$ followed by a reconstruction phase during which a desired structure on X is computed. Recent results in this model parametrize the reconstruction time by the *ply* of $\mathcal{R}$, which is the maximum overlap between the regions in $\mathcal{R}$. We introduce the *ambiguity* $\mathcal{A}(\mathcal{R})$ as a more fine-grained measure of the degree of overlap in $\mathcal{R}$. We show how to preprocess a set of d-dimensional disks in $O(n \log n)$ time such that we can sort X (if $d = 1$) and reconstruct a quadtree on X (if $d \geq 1$ but constant) in $O(\mathcal{A}(\mathcal{R}))$ time. If $\mathcal{A}(\mathcal{R})$ is sub-linear, then reporting the result dominates the running time of the reconstruction phase. However, we can still return a suitable data structure representing the result in $O(\mathcal{A}(\mathcal{R}))$ time. In one dimension, $\mathcal{R}$ is a set of intervals and the ambiguity is linked to interval entropy, which in turn relates to the well-studied problem of sorting under partial information. The number of comparisons necessary to find the linear order underlying a poset P is lower-bounded by the graph entropy of P. We show that if P is an interval order, then the ambiguity provides a constant-factor approximation of the graph entropy. This gives a lower bound of $\Omega(\mathcal{A}(\mathcal{R}))$ in all dimensions for the reconstruction phase (sorting or any proximity structure), independent of any preprocessing; hence our result is tight. Finally, our results imply that one can approximate the entropy of interval graphs in $O(n \log n)$ time, improving the $O(n^{2.5})$ bound by Cardinal *et al.*	**Optimal algorithm for geodesic farthest-point Voronoi diagrams** Luis Barba Let P be a simple polygon with n vertices. For any two points in P, the geodesic distance between them is the length of the shortest path that connects them among all paths contained in P. Given a set S of m sites being a subset of the vertices of P, we present the first randomized algorithm to compute the geodesic farthest-point Voronoi diagram of S in P running in expected $O(n+m)$ time. That is, a partition of P into cells, at most one cell per site, such that every point in a cell has the same farthest site with respect to the geodesic distance. This algorithm can be extended to run in expected $O(n + m \log m)$ time when S is an arbitrary set of m sites contained in P.
11:30-11:55	Break + Fast Forward	

Multimedia Sessions(The abstracts are not in order of the presentations)

11:55-12:45

Fréchet View – A Tool for Exploring Fréchet Distance Algorithms

Peter Schäfer

The **Fréchet-distance** is a similarity measure for geometric shapes. Alt and Godau presented the first algorithm for computing the Fréchet-distance and introduced a key concept, the Since then, numerous variants of the Fréchet-distance have been studied.

We present here an interactive, graphical tool for exploring some Fréchet-distance algorithms. Given two curves, users can experiment with the free-space diagram and compute the Fréchet-distance. The Fréchet-distance can be computed for two important classes of shapes: for polygonal curves in the plane, and for simple polygonal surfaces.

Finally, we demonstrate an implementation of a very recent concept, the ***k*-Fréchet-distance**.

A manual comparison of convex hull algorithms

Maarten Löffler

We have verified experimentally that there is at least one point set on which Andrew's algorithm (based on Graham's scan) to compute the convex hull of a set of points in the plane is significantly faster than a brute-force approach, thus supporting existing theoretical analysis with practical evidence. Specifically, we determined that executing Andrew's algorithm on the point set

$$P = \{(1,4),(2,8),(3,10),(4,1),(5,7),(6,3),(7,9),(8,5),(9,2),(10,6)\}$$

takes 41 minutes and 18 seconds; the brute-force approach takes 3 hours, 49 minutes, and 5 seconds.

Packing Geometric Objects with Optimal Worst-Case Density

A. T. Becker, S. P. Fekete, P. Keldenich, S. Morr, C. Scheffer

We motivate and visualize problems and methods for packing a set of objects into a given container, in particular a set of different-size circles or squares into a square or circular container. Questions of this type have attracted a considerable amount of attention and are known to be notoriously hard. We focus on a particularly simple criterion for deciding whether a set can be packed: comparing the total area A of all objects to the area C of the container. The *critical packing density* δ^* is the largest value A/C for which any set of area A can be packed into a container of area C. We describe algorithms that establish the critical density of squares in a square ($\delta^* = 0.5$), of circles in a square ($\delta^* = 0.5390\ldots$), regular octagons in a square ($\delta^* = 0.5685\ldots$), and circles in a circle ($\delta^* = 0.5$).

Properties of Minimal-Perimeter Polyominoes

G. Barequet and G. Ben-Shachar

In this video, we survey some results concerning polyominoes, which are sets of connected cells on the square lattice, and specifically, minimal-perimeter polyominoes, that are polyominoes with the minimal-perimeter from all polyominoes of the same size.