Guest Editors' Foreword

Bettina Speckmann¹ · Csaba D. Tóth²

Received: 22 April 2020 / Published online: 26 May 2020

© Springer Science+Business Media, LLC, part of Springer Nature 2020

The 34th International Symposium on Computational Geometry (SoCG) was held in Budapest, Hungary, June 11–14, 2018, as part of Computational Geometry Week. A record number of 206 papers were submitted to SoCG 2018, of which the Program Committee accepted 73 for presentation. This special issue of *Discrete & Computational Geometry* contains the full versions of seven especially strong papers from the symposium, recommended by the Program Committee. These papers were submitted, reviewed, and revised according to the usual high standards of *D&CG*.

We thank the authors of all submitted papers for revising and polishing their work. We are grateful to the anonymous referees for their dedication and expertise, which ensured the high quality of the papers in this special issue.

The papers appear here in alphabetical order of the names of the first authors. In the remainder of this foreword we briefly introduce all papers, arranging them into broad topics.

Geometric Hitting Sets and Set Covers Capacitated covering problems in geometric range spaces are variants of the set cover problem, with constraints on the maximum number of points assigned to each range in the cover. *Sayan Bandyapadhyay, Santanu Bhowmick, Tanmay Inamdar, and Kasturi Varadarajan* consider covering a set of points in $\mathbb{R}^d$ by the minimum number of balls, subject to capacity constraints for each ball. The problem is known to be hard to approximate in Euclidean spaces. The authors present bicriteria approximation results, approximating the cost while expanding the balls by a constant factor (but not changing the capacities), using a novel rounding scheme for the natural LP-relaxation of these problems that can handle balls of widely different radii.

The Hadwiger–Debrunner number $HD_d(p, q)$, $q \geq d + 1$, is the minimum number of points that can pierce any family of convex bodies in $\mathbb{R}^d$ with the (p, q) -property

Bettina Speckmann
b.speckmann@tue.nl

Csaba D. Tóth
csaba.toth@csun.edu

¹ Department of Mathematics and Computer Science, TU Eindhoven, Eindhoven, The Netherlands

² Department of Mathematics, California State University Northridge, Los Angeles, CA, USA

(i.e., among any p sets some q have nonempty intersection). However, the $(p, 2)$ -property does not imply bounded piercing number for convex bodies in general, except for families with Helly number 2, such as axis-aligned boxes. *Chaya Keller* and *Shakhar Smorodinsky* bootstrap from a $(p, 2)$ -theorem for axis-aligned boxes in $\mathbb{R}^d$, and prove that $\text{HD}_{d\text{-box}}(p, q) = p - q + 1$ when $q = \Omega(\log^{d-1} p)$. The exact value of $\text{HD}_{d\text{-box}}(p, q)$ was previously known only for $d = 2$ and $q \geq \sqrt{2p}$.

The colorful Helly theorem states that if a family $\mathcal{F}$ of convex sets in $\mathbb{R}^d$ is partitioned into $d + 1$ color classes $\mathcal{F} = \bigcup_{i=1}^{d+1} \mathcal{F}_i$, and any rainbow selection of $d + 1$ sets (one from each color class) has a nonempty intersection, then one of the families also has a nonempty intersection. However, the theorem does not say anything about the remaining d color classes. *Leonardo Martínez-Sandoval*, *Edgardo Roldán-Pensado*, and *Natan Rubin* derive stronger conclusions under the hypothesis of the colorful Helly theorem, and show that one additional color class can be pierced with a constant number of points, or all sets in $\mathcal{F}$ can be stabbed by a constant number of lines; both constants depend only on the dimension d . This result may be the first to explore a sense of “transversal dimension” of convex sets.

Geometric and Topological Graph Theory Every planar graph can be embedded in $\mathbb{R}^2$ with straight-line edges such that the vertices have integer coordinates. The minimum area of an axis-aligned bounding box of such an embedding measures the complexity of the graph. The area required for trees already lends itself to interesting problems on hierarchical subdivisions and recurrence relations. *Timothy M. Chan* improves several long-standing bounds for trees, using bootstrapping and basis reduction. Among other results, he improves the area bound for trees from $O(n \log n)$ to $n \exp(O(\sqrt{\log \log n \log \log \log n}))$, and for binary trees from $O(n \log \log n)$ to $n \exp(O(\log^* n))$.

The celebrated *crossing lemma* gives a lower bound of $\Omega(e^3/n^2)$ on the number of edge crossings in any proper drawing of a simple graph with n vertices and $e \geq 4n$ edges. For multigraphs, with maximum multiplicity m , a tight lower bound $\Omega(e^3/(mn^2))$ was established by Székely in 1997. *János Pach* and *Géza Tóth* settle a recent conjecture by Bekos, Kaufmann, and Raftopoulou, and show that the crossing number of multigraphs is at least $\Omega(e^3/n^2)$, for $e \geq 4n$, independent of the multiplicity, under three natural conditions: any two edges cross at most once (*single-crossing*), adjacent edges do not cross (*local starlikeness*), and there is some vertex in the interior and some vertex in the exterior of every Jordan curve formed by two parallel edges (*separatedness*). This result has already generated follow-up work on variants where some of these conditions are dropped.

It is easy to see that the intersection graph of convex sets in $\mathbb{R}^2$ is always a string graph, that is, the intersection graph of continuous arcs in $\mathbb{R}^2$; but not all string graphs are intersection graphs of convex set. *János Pach*, *Bruce Reed*, and *Yelena Yuditsky* show that almost all string graphs on n vertices are intersection graphs of convex sets in $\mathbb{R}^2$, as n goes to infinity. This is a consequence of a beautiful and important structural result, conjectured by Janson and Uzzell (2017), which states that the vertices of almost all string graphs can be partitioned into five cliques so that there is no edge between at least one pair of cliques.

Machine Learning and Data Analysis Kernel density estimates are the cornerstone of several data-analysis algorithms. A *kernel* is a bivariate function K on a finite set $P \subset \mathbb{R}^d$, and the *kernel density estimate* KDE_P at a point $x \in P$ is the average of $K(p, x)$ over all $p \in P$. One can approximate KDE_P by KDE_Q for some much smaller sample $Q \subset \mathbb{R}^d$; the set Q is an ε -coreset for KDE_P if $\|\text{KDE}_P - \text{KDE}_Q\|_\infty \leq \varepsilon$. Wai Ming Tai and Jeff Phillips construct a coreset of size $O(d^{1/2}\varepsilon^{-1}(\log \varepsilon^{-1})^{1/2})$ and derive a nearly-matching lower bound of $\Omega(\min\{d^{1/2}\varepsilon^{-1}, \varepsilon^{-2}\})$ for positive definite kernels under mild assumptions. The proof for the upper bound yields a polynomial-time algorithm to construct a suitable coreset.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.