OPEN PROBLEMS FROM WADS 2023

WADS 2023

August 3, 2023

1 Are k-Guarding Sets Splittable?

Pat Morin (Carleton University)

For a pair of points p and q in a simple polygon P, we say that p *sees* q (and vice-versa) if the line segment pq is contained in P. A set S of points in P *k-guards* P if, for every point x in P, there are at least k points in S that see x.

Open Problem 1. Is there an integer k such that, for every polygon P, every set S that k-guards P can be split into two sets that each 1-guard P?

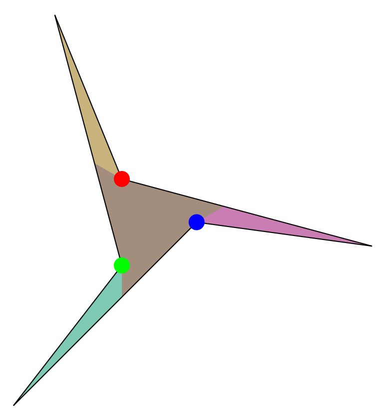

Godfried's favourite polygon (GFP) shows that, if k is defined, it is greater than 2. The reflex vertices of GFP are a set of 3 vertices that 2-guard GFP, but none of the three 1-guards GFP.

It is also possible to show that $k > 3$ by constructing an example in which S contains 5 guards $p_0, \ldots, p_4$ in convex position inside a polygon P that has spikes so that, for each i, there is a spike that can only be seen by p_i, p_{i+1}, p_{i+2} and another spike that can only be seen by p_i, p_{i+2}, p_{i+3}.

2 Mixing Graphs

Marek Chrobak (University of California, Riverside)

A *micro mixer* is a gadget with two inputs and two outputs. Each of the inputs accepts a liquid, the micro mixer mixes the two input liquids perfectly and the resulting mixture comes out of both outputs. A mixing network is a directed acyclic graph in which each vertex either a source of out-degree 1, a sink of in-degree 1, or a micro mixer (with in-degree and out-degree both equal to 2). The edge incident to each source is labelled with a 0 or 1 and when a micro mixer has two incoming edges with labels a and b, its two outgoing edges are both labelled $(a+b)/2$.

The *mixing problem* takes as input a vector $(x_1,\dots,x_n)$ of rational numbers in the interval $[0,1]$ and asks if there exists a mixing network with n sinks whose incoming edges are labelled $x_1,\dots,x_n$.

Open Problem 2. Is the mixing problem decideable?

3 Min-Load Cluster

Zac Friggstad (University of Alberta)

Let (X,d) be a metric space. The *min-load k-clustering problem* takes as input two multisets (C,S) of X (clients and servers) and asks to partitition C into k multisets $\{C_s : s\in S\}$ indexed by the elements of S so that $\max_{s\in S}\sum_{c\in C_s} d(c,s)$ is minimized.

An easy reduction from Partition shows that this problem is NP-hard when the metric space is $X := \mathbb{R}^1$ with the usual distance function $d(x,y) := |x-y|$ even when $S := \{\{0,0\}\}$. The only positive result known is a PTAS for the continuous $\mathbb{R}^1$ case in which $X := \mathbb{R}^1$ and $S := \mathbb{R}^1$ (so the clients are numbers and the servers can be located anywhere).

Open Problem 3. Prove something about the complexity of min-load k-clustering in some interesting metric space (for example, $\mathbb{R}^2$ with Euclidean distances).

4 Set-Cover by Rotations

Anthony Wirth (University of Melbourne)

Consider the following special case of SetCover. Let $X := \{0,\dots,n-1\}$ and let S be a subset of X. For each $i\in X$, let $S_i := \{(s+i) \bmod n : s\in S\}$. The problem is to find the minimum size of a set $I\subseteq X$ such that $\bigcup_{i\in I} S_i = X$.

Open Problem 4. What is the computational complexity of this problem?

5 Differentially Private All-Pairs Shortest Distance

Chengyuan Deng (Rutgers University)

For this problem the input is an edge-weighted graph $G := (V,E,w)$ where $w : E\to\mathbb{R}^+$. This gives rise to an all-pairs distance matrix A where $A(v,w)$ gives the length of the shortest path from v to w. The output of the algorithm should be a differentially private distance matrix $\hat{A}$ that minimizes $c(\hat{A}) := \max_{v,w\in V} |\hat{A}_{v,w} - A_{v,w}|$.

A paper published at NeurIPS 2022 proves the existence of an ϵ-differentially-private matrix $\hat{A}$ that achieves $c(\hat{(A)}) \leq \tilde{O}(n^{1/3})$. A paper published at SoDA 2023 proves the existence of an (ϵ, δ)-DP matrix $\hat{A}$ that achieves $c(\hat{A}) \leq \tilde{O}(n^{1/2})$. The author has a new unpublished result that shows a lower bound of $\Omega(n^{1/4})$ for any (ϵ, δ)-DP matrix $\hat{A}$.

Open Problem 5. Close the gap between the $\Omega(n^{1/4})$ lower bound and the $\tilde{O}(n^{1/2})$ upper bound for (ϵ, δ)-differential privacy.

6 Depth-Reduction for Explainable k-Center

Chengyuan Deng (Rutgers University)

This question is about depth-reduction for clustering using axis-parallel cuts. If you insist on axis-parallel cuts and upper-bound the maximum depth of the decision/explation tree, then this results in polynomially unbounded error for k-median and k-mean clustering. For k-center clustering, the only known lower bound is a factor of 2.

Open Problem 6. Is this tight?

7 Optimal Covering of Coloured Rectangular Tiles

Bogdan Armaselu (Spectral MD)

The input to this problem is a rectangle that has been tiled with n smaller rectangles, each of which is assigned a colour in $\{1, \ldots, k\}$ and a shape S. The problem is to find a rigid transformation of S that intersects the maximum number of different colours.

It was observed by several members of the audience that this problem can be solved in polynomial time by creating, for each tile, a 3-dimensional shape whose boundary is piecewise helical and then computing the arrangment of all n such shapes.