Advanced Graph Algorithms and Optimization Spring 2024

The Cut-Matching Game

R. Kyng & M. Probst Gutenberg *Problem Set 8— Monday, April 22th*

The exercises for this week will not count toward your grade, but you are highly encouraged to solve them all. This exercise sheet has exercises related to week 8. We encourage you to start early so you have time to go through everything.

To get feedback, you must hand in your solutions by 23:59 on May 2. Both hand-written and LaTeX solutions are acceptable, but we will only attempt to read legible text.

Exercise 1: Proof that $\mathbf{\Pi}^t$ is doubly-stochastic

In chapter 13, we stated that the matrix $\mathbf{\Pi}^t$ is doubly-stochastic. We now ask you to prove it.

Exercise 2: Run-time Analysis of the Cut-Matching Algorithm

Analyze the run-time of the Cut-Matching algorithm to establish a total run-time of $O(\log^2 n) \cdot T_{max-flow}(G) + \tilde{O}(m)$[1]. You may assume that computing a random vector $\boldsymbol{r}$ takes $O(n)$ time, and you may use without proof that using link-cut trees you can find a path decomposition in $\tilde{O}(m)$ time.

Exercise 3: Bounding the Diameter of an Expander

Let G be a connected, unweighted graph that is a ϕ-expander with regard to conductance. Prove that the diameter of G is $O(\log m/\phi)$. You may find it useful that $\phi \leq 1$ and that $e^x < 1 + x + x^2$ for $x < 1.79$.

Hint: Fix any pair of vertices s, t in G. Then think about running BFS explorations simultaneously from s and t.

Exercise 4: ASZ Vertex Sparsifiers

In this exercise, we consider a very simple way of obtaining a vertex sparsifier due to Andoni, Stein and Zhong. Recall that a vertex sparsifier of $G = (V, E)$ is a graph $H = (V', E', \boldsymbol{w})$ with less vertices than G that preserves some property of G with respect to a terminal set $T \subset V, V'$. In this exercise, the property we aim to preserve is the pairwise distance between vertices in T. Notice that the vertex sparsifier is a weighted graph, whereas we assume $G = (V, E)$ to be unweighted. Similarly, edge sparsifiers of unweighted graphs are often also weighted, for example in the case of cut/spectral sparsification. We assume that all shortest paths in G are unique, and that the maximum degree in G is Δ.

[1] Technically, we run the max flow procedure on graphs G with two additional vertices and n additional edges. You should ignore this subtlety as it will not impact the asymptotics for any max-flow algorithm known.

We now describe the construction of H given G and T. First, we add all the terminal vertices in T to V'. Then, we sample every vertex with probability $100 \log n/k$ and add all the sampled vertices to V'.To describe the edge set, we denote with $B_G(v, V')$ the largest ball around v that does not contain any vertex in V'.

We then describe a path collection $\mathcal{P}$. For every vertex u, $v \in B(u, V')$ and edge $(v, x) \in E$ add the path $\pi_G(u, v) \oplus (v, x)$ to $\mathcal{P}$ where $\pi_G(u, v)$ denotes the shortest path in between u and v in G. Finally, we say that the pivot $p(v) \in V'$ of each vertex v is the closest vertex to v in V'. We then let $\tilde{\mathcal{P}} = \{\pi_G(p(u) \oplus P \oplus \pi_G(v, p(v)) | P \in \mathcal{P} \text{ is a } u - v \text{ path}\}$ be the set of paths projected to the closest endpoints in V'. Finally, the edge set E' is constructed by adding an edge (u, v) of length $|P|$ for every $u - v$ path in $\tilde{\mathcal{P}}$. This concludes the description of the graph H.

1. Show that $|B(v, V')| \leq k$ with high probability for all v.
2. Argue that the expected number of vertices in V' is in $\tilde{O}(|T| + n/k)$.
3. Show that the size of the set $|\mathcal{P}| \leq \Delta \cdot n \cdot k$ with high probability.
4. Show that for every pair $u, v \in V'$, we have $\text{dist}_G(u, v) \leq \text{dist}_H(u, v) \leq 4 \, \text{dist}_G(u, v)$.

Exercise 5: Distance Oracles

Recall the notion of a cluster in the context of Thorup-Zwick distance oracles for a graph $G = (V, E)$. For every vertex w, we let $C(w) = \{v \in V | \, \text{dist}(v, w) < \text{dist}(v, p(v))\}$ i.e. the cluster of w contains all vertices that are closer to w than to their pivot. As previously, we assume that the distances (and thus shortest paths) are unique. Furthermore, assume that the distance $\text{dist}(v, p(v))$ of vertex v to its pivot is known for every $v \in V$.

1. Show that for a vertex w we can compute the cluster $C(w)$ by running a modified version of Dijkstra in time $\tilde{O}(|E(C(w))|)$.

Hint: First show that for a vertex $t \in C(w)$, all the vertices on its shortest path $\pi(t, w)$ to w are also in $C(w)$. Then, show that running the version of Dijkstra that does not relax edges from vertices u for which $\text{dist}(u, p(u)) \leq \text{dist}(w, u)$ finds the whole cluster $C(w)$.