Homework 4

Topics in Graph Algorithms
CS598CCI, Spring 2020
Due: 4/20/2020 10am

Instructions and Policy: Each person should write up their own solutions independently. You need to indicate the names of the people you discussed a problem with. Solutions to most of these problems can be found from one source or the other. Try to solve on your own first, and cite your sources if you do use them.

Please write clearly and concisely. Refer to known facts. You should try to convince me that you know the solution, as quickly as possible.

Do as many problems as you can. I expect you to do at least 3 including Problems 1 and 5.

Problem 1. The notion of separators and well-linked sets plays an important role in several area including the graph theoretic notion called treewidth, routing and switching networks etc. The goal of this problem is to expose these notions in the context of (uniform) sparsest cut. We will use the notion of edge-well-linkedness which is related to sparsest cut in the sense we discussed in lectures via edge cuts. One needs the notion of vertex-well-linkedness for tree width.

Given a graph $G = (V, E)$ a set of vertices $X \subseteq V$ is well-linked if for any two disjoint sets $A, B \subset X$ of equal size, say k, there are $k = |A| = |B|$ edge-disjoint paths $P_1, P_2, \ldots, P_k$ between A and B with each vertex in $A \cup B$ as the end point of exactly one of these paths (think of the paths as defining a perfect matching between A and B).

- Given disjoint sets A, B of equal size k how can you check efficiently that there are k-edge disjoint A-B paths in G as above?
- Prove that X is well-linked in G iff for all $S \subset V$, $|\delta(S)| \geq \min\{|S \cap X|, |(V \setminus S) \cap X|\}$. In other words the sparsest cut with respect to X is at least 1.
- Suppose G is an 1-expander, that is $|\delta(S)| \geq |S|$ for all $S \subseteq V$ with $|S| \leq |V|/2$. Prove that V is well-linked as a corollary of the preceding part.
- Consider a $k \times k$ grid-graph. Let X be the set of vertices in the first row of the grid. Prove that X is well-linked in the graph. In fact it is easier to prove this by explicitly displaying the disjoint paths for any A, B rather than using the cut characterization.
- Suppose X is a well-linked set in G with $|X| = k$. Consider any disjoint sets $A, B \subset X$ with $|A| = |B| = k/(c \log k)$ for sufficiently large but fixed c. Consider any bijection $f : A \to B$ (in other words f corresponding to a perfect matching M between A and B). Prove that the multicommodity flow corresponding to f (sending 1 unit for each pair $(a, f(a))$, $a \in A$) can be routed in G with congestion at most c' for some fixed constant c'. Alternatively, the

maximum concurrent flow for the demand pairs $(a, f(a)), a \in A$ is at least $1/c'$. ***Hint:*** Use the flow-cut gap for sparsest cut that we showed in class.

In particular if G is an expander then for any two sets A, B with $A, B = \Omega(n/\log n)$, G acts as a switch between A and B (in the fractional sense). Recall that a switch acts as complete bipartite graph that can route any given (perfect) matching between A and B.

Problem 2. Let (V, d) be a finite metric space and let $w(uv), u \neq v$ be given non-negative weights on unordered pairs (think of w as edge weights of a complete graph on V). A line embedding of V is a mapping $f : V \to \mathbb{R}$ into the real line (a 1-dimensional embedding). A line-embedding f is a contraction if $|f(u)-f(v)| \leq d(u, v)$ for all $u, v \in V$. A contracting line embedding f has *average weighted* distortion α with respect to w if $\sum_{u,v} w(uv)d(u, v) \leq \alpha \sum_{u,v} w(uv)|f(u) - f(v)|$. Line embeddings are related to ℓ_1 embeddings but are weaker. These embeddings play an important role in flow-cut gaps for vertex capacities (see `https://epubs.siam.org/doi/abs/10.1137/05064299X?journalCode=smjcat`) and in further generalizations to polymatroidal networks (see `https://epubs.siam.org/doi/abs/10.1137/130906830`).

- Revisit Bourgain's proof for low distortion ℓ_1-embedding and show that given (V, d) and w there is randomized algorithm to produce a line embedding with average weighted distortion $O(\log n)$. In fact prove a stronger statement that if support of w is k (means that only k pairs have $w(uv) > 0$) then the average weighted distortion is $O(\log k)$.

- Use line embeddings in the previous part to derive an $O(\log k)$-approximation for non-uniform sparsest cut with k pairs by rounding the LP relaxation.

Problem 3. A finite metric (V, d) is a tree metric if there is an edge-weighted tree $T = (V, E)$ such that d is the shortest path metric on V induced by T.

- Prove that a tree metric isometrically emends in ℓ_1.

- Given (V, d), a tree metric d_T induced by a weighted tree $T = (V, E_T)$ is a dominating tree metric if $d_T(u, v) \geq d(u, v)$ for all $u, v \in V$. Recall that we claimed that there is an embedding of any finite n-point metric (V, d) into a probability distribution $\mathcal{D}$ over dominating tree metrics such that $E_{T\sim\mathcal{D}}[d_T(u, v)] \leq O(\log n)d(u, v)$ for all $u, v \in V$. Use this fact and the preceding part to argue that there exists an embedding of (V, d) in to ℓ_1 with worst-case distortion $O(\log n)$.

Problem 4. A metric (V, d) is a ring metric if it is the shortest path metric of a cycle on V with non-negative edge lengths.

- Prove that any ring metric isometrically embeds into ℓ_1.

- Suppose we have a non-uniform sparsest cut instance over a ring network. That is, the supply graph is a weighted cycle but the demand graph is arbitrary. Describe a simple combinatorial algorithm to find the sparsest cut in such an instance and justify its correctness.

- What is the connection between the preceding two parts?

Problem 5. Given a graph $G = (V, E)$ with edge-weights $c : E \to \mathbb{R}_+$, you wish to partition G into $G_1 = G[V_1], G_2 = G[V_2], G_3 = G[V_3]$ such that $\lfloor |V|/3 \rfloor \le |V_i| \le \lceil |V|/3 \rceil$ for $1 \le i \le 3$, and the cost of the edges between the partitions is minimized. Using an α-approximation for the sparsest cut problem, give a pseudo-approximation for this problem where you partition the graph into 3 pieces $G[V_1'], G[V_2'], G[V_3']$ such that $|V|/c_2 \le |V_i'| \le |V|/c_1$ for some constants $1 < c_1 < c_2$ and the cost of the edges between the partitions is $O(\alpha)\textsc{opt}$. What constants c_1, c_2 can you guarantee? Note that c_1 and c_2 should be *constants*, independent of the graph size. (Hint: this problem is similar to the one on partitioning into two pieces that is in Vazirani's book on applications of sparsest cut (Section 21.6.3).)