Succinct Graph Structures and Applications **Spring 2022**

Exercise 1: May 03

Lecturer: Merav Parter

Girth and Short Cycles

Recall that the girth $g(G)$ of a graph G is the length of the shortest cycle in G. Erdős girth conjecture states that for every $k \geq 1$ and sufficiently large n, there exist n-vertex graphs with $\Omega(n^{1+1/k})$ edges and girth at least $2k+1$. A weaker lower bound can be shown via the probabilistic approach. Specifically, we will prove that there exists an n-vertex graph G^* with $\Omega(n^{1+1/(2k-1)})$ edges and girth $g(G)$ at least $2k+1$.

Exercise 1. The existence of G^* can be shown in two steps. (I) Consider a $G(n,p)$ graph[1] with $p = \Theta(1/n^{1-1/(2k-1)})$ and bound the expected (total) number of cycles of length $t \leq 2k$ in this graph. (II) Prove the existence of n-vertex graph G' with $\Theta(n^{1+1/(2k-1)})$ edges and a small number of cycles and turned it into the desired graph G^* while keeping the same order of the number of edges as in G'.

Multiplicative Spanners

We saw in class the Baswana-Sen algorithm for $(2k-1)$-spanners for any $k \geq 1$. This algorithm has k phases where in each phase $i \geq 1$ it computes a clustering $\mathcal{C}_i$ that consists of $O(n^{1-i/k})$ clusters. Each cluster has a center vertex, and all the vertices in a cluster are connected in the spanner via a depth-i tree rooted at the center. One of the benefits of the Baswana-Sen algorithm is that for vertices that stopped being clustered in level $i \geq 1$ it provides a stretch of at most $2i-1$ for each of their incident edges. See Fig. 8 of [1] for a short description of the algorithm.

Exercise 2. Our goal in this exercise is to compute a subgraph $H \subseteq G$ for an n-vertex graph $G = (V,E)$ with $\widetilde{O}(n^{1+1/k})$ edges such that for every vertex pair u,v such that $\text{dist}_G(u,v) \geq \sqrt{k}$, it would hold that $\text{dist}_H(u,v) = O(\sqrt{k}) \cdot \text{dist}_G(u,v)$. Note that the spanners we saw in class with $\widetilde{O}(k \cdot n^{1+1/k})$ edges provide a stretch of $2k-1$ for *every* pair of nodes. Thus our goal is to improve this stretch bound for sufficiently far vertex pairs in G. Consider the following algorithm for this purpose. Throughout assume that $k' = \sqrt{k}$ is an integer.

1. Run the first $k' = \sqrt{k}$ phases of the Baswana-Sen algorithm. Let $\mathcal{C}_{k'}$ be the level-k' clustering, $Z_{k'}$ be the cluster centers of the clusters of $\mathcal{C}_{k'}$. Also, let H_1 be the output subgraph containing all edges added by the algorithm throughout the first k' phases (this includes the internal trees of each cluster in $\mathcal{C}_i$ for $i \in [1,k']$, as well as the edges added due to unclustered vertices).
2. Consider the graph $G^* = (Z_{k'}, E^*)$ where $E^* = \{(u,v) \in Z_{k'} \mid \text{dist}_G(u,v) \leq 5\sqrt{k}\}$. That is, each vertex in G^* corresponds to a *center* of a cluster in $\mathcal{C}_{k'}$, and every two centers are connected by an edge in G^* if their distance in G is at most $5\sqrt{k}$.
3. Let $H^* \subseteq G^*$ be a $(2\sqrt{k}-1)$-spanner of G^*.
4. Let H_2 be the subgraph of G obtained by adding the u-v shortest path $\pi(u,v)$ in G for any $(u,v) \in H^*$. That is, $H_2 = \bigcup_{(u,v) \in H^*} \pi(u,v)$.

[1] In $G(n,p)$ graph, each of the $\binom{n}{2}$ edges exists with probability p.

5. Output $H = H_1 \cup H_2$.

(2a) Show that $|E(H)| = \widetilde{O}(\sqrt{k} \cdot n^{1+1/k})$ w.h.p.

(2b) Call a vertex u *clustered* if it belongs to one of the clusters of $\mathcal{C}_{k'}$, otherwise u is *unclustered*. Show that every edge incident to an unclustered vertex u has a stretch of at most $2\sqrt{k} - 1$ in H_1. That is, $\mathrm{dist}_H(u, v) \leq 2\sqrt{k} - 1$ for every $v \in N(u)$ where $N(u)$ are the neighbors of u in G.

(2c) Consider a pair of clustered vertices u and v at distance $\sqrt{k}$ in G. Show that z_u and z_v are neighbors in G^* where z_u, z_v are the centers of u, v (respectively) in $\mathcal{C}_{k'}$. Use it to deduce that $\mathrm{dist}_H(u, v) = O(k)$.

(2d) Show that $\mathrm{dist}_H(u, v) = O(k)$ for *every* pair of vertices u and v at distance $\sqrt{k}$ in G.

(2e) Use (5d) to deduce that $\mathrm{dist}_H(u, v) = O(\sqrt{k}) \cdot \mathrm{dist}_G(u, v)$ for every u, v satisfying that $\mathrm{dist}_G(u, v) \geq \sqrt{k}$.

Additive Spanners

Exercise 3. Given an unweighted n-vertex graph $G = (V, E)$, a 6-additive spanner $H \subseteq G$ is a subgraph satisfying that

$$\mathrm{dist}_H(u, v, H) \leq \mathrm{dist}_G(u, v) + 6, \forall u, v \in V \ .$$

We saw in class, the construction of +2 and +4 additive spanners. In this exercise, we will construct, in a step by step manner, a 6-additive spanner H with $\widetilde{O}(n^{4/3})$ edges.

1. The algorithm defines a degree threshold Δ_1, and adds all edges incident to vertices with degree at most Δ_1 to H. How large can Δ_1 be (i.e., so that the edge bound of $\widetilde{O}(n^{4/3})$ is kept)?

2. Next, the algorithm takes care of all shortest paths $\pi(u, v)$ that have at least *one* high-degree vertex with degree at least Δ_2. To do that, it samples each vertex $v \in V$ into a set Q independently with probability $C \cdot \log n / \Delta_2$ for some large constant C. A BFS tree rooted at each vertex $q \in Q$ is added to H. How small can Δ_2 be?

3. Finally, it remains to take care of paths that have no vertex with degree at least Δ_2. On each path $\pi(u, v)$, observe that all edges incident to low-degree vertices (vertices with degree at most Δ_1) are in H. Hence, when adding a shortest-path, we only "pay" for the number of missing edges (those that are incident to vertices with degree in $[\Delta_1, \Delta_2]$). We will take care of these paths in $O(\log n)$ phases. For every $i \in \{0, 1, \ldots, 2 \log n\}$, define the set Q_i by randomly including each vertex v into Q_i with probability $p_i = C \log n / (\Delta_1 \cdot 2^i)$ for sufficiently large constant C. For every vertex u that has a neighbor, say w, in Q_0, add one edge between u and w to the spanner.

   In each phase $i \geq 1$, we take care of all paths $\pi(u, v)$ that have $x \in [2^{i-1}, 2^i]$ edges that are missing in H. This is done as follows. For each $t_1 \in Q_0$ and each $t_2 \in Q_i$, add to H the shortest t_1-t_2 path in G that has at most 2^i missing edges in H. That is, among all paths between t_1 and t_2 in G that have at most 2^i edges that are not in H, pick the shortest one and add its edges to H.

(3a) Prove that H has $\widetilde{O}(n^{4/3})$ edges and **(3b)** show that H is a +6-spanner.

References

[1] Baswana, Surender and Kavitha, Telikepalli and Mehlhorn, Kurt and Pettie, Seth Additive spanners and (α, β)-spanners In *ACM Transactions on Algorithms (TALG)*, 1–26, 2010.