Advanced Algorithms 2024	11/04, 2024
Exercise 07	
Lecturer: Maximilian Probst	*Teaching Assistant: Jakob Nogler*

1 Streaming a Minimum Spanning Tree

In this problem, we will see algorithms for computing a minimum spanning tree in the streaming model. Throughout the question, you can assume that the weights of the edges in the input graph are polynomially-bounded non-negative integers.

1. Design a streaming algorithm for computing a 2-approximation for the MST of the graph, using $\tilde{O}(n)$ total memory.
2. Design a streaming algorithm for computing the MST of the graph, using $\tilde{O}(n)$ total memory and $O(\log^2 n)$ passes. In a k-pass algorithm, the algorithm is allowed to do k passes over the input graph.

2 Streaming 3-Connectivity

We have seen in class a $\tilde{O}(n)$-memory algorithm that solves the graph connectivity problem in the streaming setting, where we have a stream of edge arrivals and departures on a set V of n vertices. Devise and algorithm with $\tilde{O}(n)$ memory that solves the 3-connectivity problem in the same setting. Recall that a graph is 3-connected if it remains connected after the removal of any two edges.

3 Communication with a Coordinator

Recall the streaming algorithm for graph connectivity introduced in the lecture. In this algorithm, each node maintains a "sketch"—a compact, limited-space representation of its neighborhood of size $O(\log^4 n)$. These sketches are then linearly combined to solve the connectivity problem.

This approach follows what is known as the *coordinator model*. In this model, a set of nodes operates independently, each holding a piece of information needed to solve a certain problem. To find the solution, each node sends a compact representation of its information to a central coordinator. The coordinator then aggregates these inputs, performs the necessary computations, and outputs the result. What is usually optimized in the coordinator model is the amount of data that each node is allowed to send to the coordinator.

Consider the following alternative problem in the coordinator model. Let there be n nodes numbered $1, 2, \dots, n$. A set $S \subseteq \{1, 2, \dots, n\}$ of nodes is selected and revealed, where each node gets to know whether they are in S or not, but does not receive any information of the other members. Each node can send a B-bit message to the coordinator, but unfortunately, the coordinator receives the bit-wise AND of the sent messages, instead of receiving each of them. Devise a scheme with $B = O(\log^2 n)$ so that the coordinator can still approximate $s = |S|$ up to a 2-factor, with probability $1 - 1/n^2$.