Exercise 7: Lost in Complexity

Task 1: Why is everything so Hard?!? $(3 + 3 + 3 + 2)$

In this exercise, we always consider connected, simple, weighted graphs $G = (V, E, W)$, restrict message size to $\mathcal{O}(\log n)$ bits, and assess worst-case round complexity as a function of the (hop) diameter D and n. In the end, you have some useful tools for proving lower bounds in the context of communication complexity.

a) Show that finding any approximation to the (weighted) distance between a given pair of nodes $s, t \in V$ takes $\Omega(\sqrt{n}/\log^2 n + D)$ rounds.

   Hint: Use the same technique and graph as in the lecture, just change the weights.

b) Show that finding a Steiner tree requires $\Omega(\sqrt{n}/\log^2 n + D)$ rounds, regardless of the size of the subset T of nodes that needs to be connected to each other (unless it is 1).

   Hint: Attach some irrelevant nodes to the construction from a) for $|T| \leq n/2$ and to the MST construction for $|T| > n/2$.

c) For $s, t \in V$, an s-t cut is a subset $s \in S \subseteq V \setminus \{t\}$. The weight of the cut is the sum of weights of all edges $\{v, w\} \in E \cap (S \times (V \setminus S))$ crossing the cut. Show that finding any approximation to the weight of a minimum s-t cut takes $\Omega(\sqrt{n}/\log^2 n + D)$ rounds.

d) Conclude that finding an approximate maximum flow or even approximating the value of such a flow requires $\Omega(\sqrt{n}/\log^2 n + D)$ rounds.

Task 2: Harder, Better, Slower $(3 + 2 + 2)$

Consider graphs $G = (V, E)$ and message size $\mathcal{O}(\log n)$. In this exercise, we show that determining the diameter D of a graph more accurately than factor $3/2$ requires $\Omega(n/\log n)$ rounds or large messages.

a) For a set disjointness instance (x, y) and input size N, construct a graph with $\mathcal{O}(\sqrt{N})$ nodes that has diameter 2 if x and y encode disjoint sets, and diameter 3 otherwise. The graph must have a cut with $\mathcal{O}(\sqrt{N})$ edges between the parts encoding x and y, respectively.

   Hint: Start from $2k$ nodes $l_1, \ldots, l_k, r_1, \ldots, r_k$ where the edge $\{l_i, r_j\}$ is included if and only if $i = j$. Enumerate the bits of x in some convenient way as $x_{i,j}$ so that i and j stay in $\mathcal{O}(\sqrt{N})$; likewise for y. Then add edges so that for all $1 \leq i, j \leq k$ there is some path of length 2 from l_i to r_j (you may assume $i \neq j$) if $x_{i,j} = 0$ or $y_{i,j} = 0$, but not if $x_{i,j} = y_{i,j} = 1$.
   Finally, add two nodes l and r and some edges to make sure that the diameter is 2 if x and y encode disjoint sets and 3 otherwise.

b) Assume for a moment a distributed algorithm exists that uses B-bit messages to compute (or approximate) the diameter of such a graph in T rounds. Show that Alice and Bob can simulate each execution with a total communication complexity of $\mathcal{O}(\sqrt{N}BT)$.

   Hint: Don't think "It can't be that easy!".

c) Conclude that $T \in \Omega(\sqrt{N}/B)$ in the worst case, no matter what algorithm is used. Specifically, conclude that if $B \in \mathcal{O}(\log n)$, it requires $\Omega(n/\log n)$ rounds to determine the diameter of a graph more accurately than up to factor of $3/2$.

Task 3*: ...under a Heap of Presents $(1+1+1+1)$

weight	RGB
1	$(255, 255, 0)$
2	$(34, 139, 34)$
3	$(165, 42, 42)$
5	$(255, 0, 0)$
20	$(193, 255, 244)$

a) Determine an MST of the graph given in Figure **??**!

b) Color each MST edge. The edge colors are given in the table above, i.e., an edge of weight 1 has color $(255, 255, 0)$.

c) Look for other Christmas trees in the computer science literature!

Hint: xkcd.

d) Have a Merry Christmas and a Happy New Year!

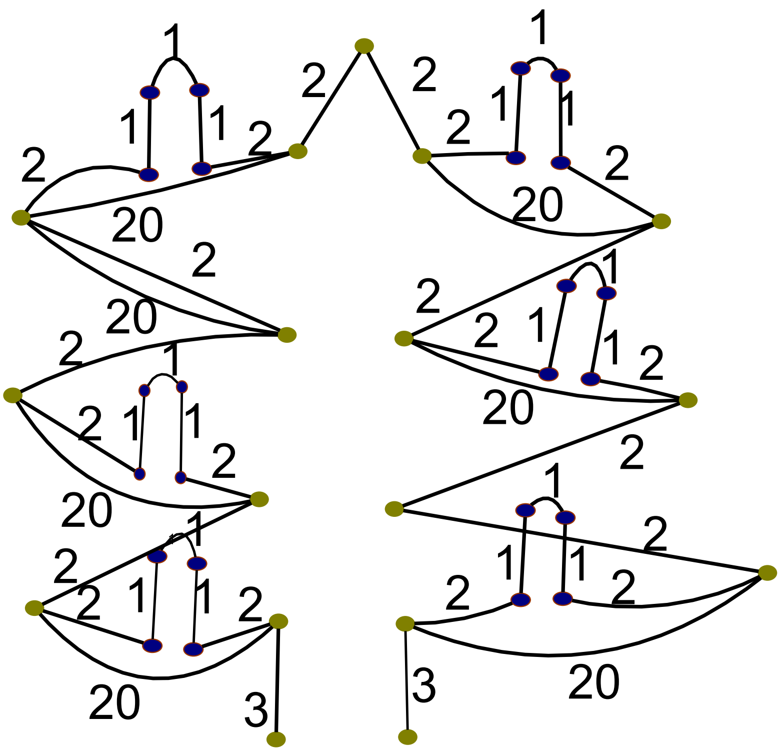

Figure 1: Poorly disguised Christmas tree.