Question 1. Prove: for all $n \geq 1$, if G is a connected graph with n vertices and $n-1$ edges, then G has no cycles.

Question 2. Let $G = (V, E)$ be an undirected graph with n vertices and m edges, with no self-loops (that is, no edges of the form (v, v) from a vertex to itself). Show that if every vertex has degree at least $n/2$, the graph is connected. If it makes your proof easier, you may assume that n is even.

Question 3 (K&T 3.9). There's a natural intuition that two nodes that are far apart in a communication network—separated by many hops—have a more tenuous connection than two nodes that are close together. There are a number of algorithmic results that are based to some extent on different ways of making this notion precise. Here's one that involves the susceptibility of paths to the deletion of nodes.

Suppose that an undirected graph $G = (V, E)$ with n vertices contains two nodes s and t such that the distance between s and t is strictly greater than $n/2$. Show that there must exist some node v, not equal to either s or t, such that deleting v from G destroys all $s - t$ paths. (In other words, the graph obtained from G by deleting v contains no path from s to t.) Give an algorithm with running time $\Theta(m + n)$ to find such a node v. Describe your algorithm in prose and prove that it works correctly. Make sure your algorithmic description is clear and concise.

Question 4. Consider the family of undirected graphs $\mathcal{H}_k$ defined as follows. $\mathcal{H}_k$ has 2^k vertices labelled with the integers 0 through $2^k - 1$. Vertices u and v are connected by an edge if and only if the binary representations of u and v differ in exactly one bit position. For example, in $\mathcal{H}_4$, the vertices 5 and 13 are connected by an edge since $5 = 0101_2$ and $13 = 1101_2$ differ in the first bit position, but the rest of the bits are the same.

Consider doing a BFS in $\mathcal{H}_{10}$ starting at node 0. How many vertices are in L_6, that is, the sixth layer generated by the BFS? Give your answer together with either a proof, or the program you used to calculate the answer. Either approach will receive full credit. (*Hint* if you choose to write a program: to flip the jth bit of an integer `n`, you can use `n ^ (1 << j)`, that is, the bitwise XOR of `n` with 2^j. These operators are valid syntax in many languages such as Java, Python, and C.)

Question 5. As in the text, we say that the *distance* between two nodes u and v in a graph $G = (V, E)$ is the minimum number of edges in a path joining them; we'll denote this by $dist(u, v)$. We say that the *diameter* of G is the maximum distance between any pair of nodes, that is,

$$diam(G) = \max_{u,v \in V} dist(u, v).$$

Give an $O(n + m)$-time algorithm to find the diameter of an (undirected) **tree** $T = (V, E)$. Prove that your algorithm is correct.

Question 6. Does your algorithm from Question 5 work for **all** undirected graphs? If so, prove it. If not, provide a counterexample.

Question 7.

- On a scale of 1 to 10, how difficult was this assignment?

- How many hours would you estimate that you spent on it?

- Which was your favorite question? Your least favorite?

Extra credit: Approximating the diameter of a graph

Question 8. There's a good chance your linear-time algorithm in Question 5 doesn't extend to arbitrary undirected graphs. (If it does, you almost certainly have a STOC publication. Congratulations!) It turns out that, until very recently, we didn't have any method for computing the diameter of a graph that didn't first compute the shortest path between all pairs of nodes. When graphs are dense, all-pairs shortest paths is fairly expensive, so some people have explored quicker algorithms which *estimate* the diameter of the graph. Develop a linear-time algorithm that, given a graph G, returns a diameter estimate that is always within a factor of 1/2 of the true diameter. That is, if the true diameter is $diam(G)$ then you should return a value k where $diam(G)/2 \leq k \leq diam(G)$.