CIS 1210—Data Structures and Algorithms—Spring 2024

BFS & DFS—Tuesday, March 4 / Wednesday, March 5

Readings

- Lecture Notes Chapter 16: Graph Traversals: BFS and DFS

Review: Graph Representations

Let $G = (V, E)$, where $|V| = n$ and $|E| = m$. In other words, G contains n vertices and m edges.

Adjacency Matrix

We can represent G with an $n \times n$ matrix A, where $A[i][j] = 1$ if there is an edge (i, j) and 0 otherwise. The primary advantage of an adjacency matrix is that we can check whether or not there is an edge between two vertices in $O(1)$ time since all we have to do is index into the matrix. On the other hand, the matrix uses $\Theta(n^2)$ space, and it takes $\Theta(n)$ time to enumerate the neighbors of a vertex v since we must iterate through an entire row in the matrix.

Adjacency List

We can also represent G as an adjacency list, where we have an array A of $|V|$ lists such that $A[u]$ contains a LinkedList of vertices v, such that for every vertex v_i in v, $(u, v_i) \in E$. (Note that if we have a directed graph, we can store two LinkedLists at index u, one for u's in-neighbors and one for u's out-neighbors.) The primary advantage of an adjacency list is that we only use $\Theta(n + m)$ space, which is better than the $\Theta(n^2)$ space usage of an adjacency matrix especially if $m \ll n^2$. Furthermore, it only takes $O(deg(v))$ time to enumerate the neighbors of a vertex v, but on the other hand, checking whether $(u, v) \in E$ also takes $O(deg(u))$ time since we have to traverse the LinkedList. In CIS 1210, we usually use adjacency lists.

Review: BFS (Breadth First Search)

At a high-level, in BFS, we explore the graph in "layers," so we explore "wide" before exploring "deep." We begin at a node v (layer 0); then explore the children of v (layer 1); then the children of the nodes in layer 1 (which make up layer 2); and so on. In other words, we explore **all** nodes at layer L_i before exploring **any** nodes at layer L_{i+1}. Because we explore the graph in layers, as seen in the pseudocode, we use a queue to implement this algorithm since it is FIFO. Note that the output of BFS is a BFS tree, and that as written, BFS is not guaranteed to visit every node in the graph.

The running time of BFS is $O(n + m)$. This is because we add and remove each node from the queue at most once, and for each node, we only examine its neighbors exactly once.

Review: DFS (Depth First Search)

In DFS, we explore "deep" before exploring "wide." We begin at some node v and examine its neighbors. When we encounter a neighbor that has not been visited yet, we visit it. Once we arrive at a node whose neighbors have all been visited, we "backtrack," via returning from recursive calls, until we reach a node that still has unvisited neighbors. Because of this, DFS can be implemented recursively as seen in this pseudocode, but we can also implement it iteratively by using a stack since it is LIFO. The iterative pseudocode for DFS with a stack is the BFS pseudocode, except instead of a queue we use a stack, instead of dequeue we use pop, and instead of enqueue we use push. We know that DFS will visit every node in the graph, so

the output of DFS is a DFS forest. Note that DFS also introduces the concept of colors and finishing times, which makes it useful for other graph applications later on.

The running time analysis for DFS is similar to that of BFS; for each vertex v, we iterate through its neighbors to yield a runtime of $O(n + m)$.

Given a DFS forest of an input graph, we can classify its edges into four types, each of which can provide important information about the graph. The four types of edges are:

1. Tree Edges: Edges in the DFS forest

2. Back Edges: Edges (u, v) where v is a ancestor of u

3. Forward Edges: Non-forest edges (u, v) where v is a descendant of u

4. Cross Edges: All remaining edges

Problems

Problem 1

Design an $O(n + m)$ algorithm to find the shortest path between nodes u and v in a connected, unweighted graph.

Problem 2

The CIS Department wants to assign prerequisites to their n classes. It has a list of m prerequisite pairings, such as CIS 110 being a prerequisite for CIS 120. Given the list of prerequisite pairings, design an algorithm that determines if this list of pairings is a valid list of prerequisites. For example, (CIS 110, CIS 120), (CIS 120, CIS 121), (CIS 121, CIS 110) is not a valid list of prerequisites. What is the runtime of your algorithm?