CSE 373 Data Structures Lecture 21

Readings and References

- Reading
 - > Section 9.3, Section 10.3.4

Some slides based on: CSE 326 by S. Wolfman, 2000

Path

• A path is a list of vertices $\{v_1, v_2, ..., v_n\}$ such that (v_i, v_{i+1}) is in **E** for all $0 \le i < n$.



Path cost and Path length

- *Path cost*: the sum of the costs of each edge
- Path length: the number of edges in the path
 - Path length is the unweighted path cost (each edge = 1)



Shortest Path Problems

- Given a graph G = (V, E) and a "source" vertex s in V, find the minimum cost paths from s to every vertex in V
- Many variations:
 - > unweighted vs. weighted
 - > cyclic vs. acyclic
 - > pos. weights only vs. pos. and neg. weights
 - > etc

Why study shortest path problems?

- Traveling on a budget: What is the cheapest airline schedule from Seattle to city X?
- Optimizing routing of packets on the internet:
 - Vertices are routers and edges are network links with different delays. What is the routing path with smallest total delay?
- Shipping: Find which highways and roads to take to minimize total delay due to traffic

Unweighted Shortest Path Problem

Problem: Given a "source" vertex s in an unweighted graph
G = (V,E), find the shortest path from s to all vertices in G



Breadth-First Search Solution

 <u>Basic Idea</u>: Starting at node s, find vertices that can be reached using 0, 1, 2, 3, ..., N-1 edges (works even for cyclic graphs!)



Breadth-First Search Algorithm

- Uses a queue to track vertices that are "nearby"
- source vertex is s

```
Distance[s] := 0
Enqueue(Q,s); Mark(s)
while queue is not empty do
   X := Dequeue(Q);
   for each vertex Y adjacent to X do
        if Y is unmarked then
        Distance[Y] := Distance[X] + 1;
        Previous[Y] := X;
        Enqueue(Q,Y); Mark(Y);
```

• Running time = O(|V| + |E|)

12/4/02



Q = C





Q = D E B



Q = B G



Q = F



Q = H

What if edges have weights?

- Breadth First Search does not work anymore
 - minimum cost path may have more edges than minimum length path



Shortest Paths - Lecture 21

12/4/02

Dijkstra's Algorithm for Weighted Shortest Path

- Classic algorithm for solving shortest path in weighted graphs (without negative weights)
- A greedy algorithm (irrevocably makes decisions without considering future consequences)
- Each vertex has a cost for path from initial vertex

Dijkstra's Shortest Path Algorithm

- Initialize the cost of s to 0, and all the rest of the nodes to ∞
- Initialize set S to be \varnothing
 - > S is the set of nodes to which we have a shortest path
- While S is not all vertices
 - Select the node A with the lowest cost that is not in S and identify the node as now being in S
 - > for each node B adjacent to A
 - if cost(A)+cost(A,B) < B's currently known cost
 - set cost(B) = cost(A)+cost(A,B)
 - set previous(B) = A so that we can remember the path



Pick vertex not in S with lowest cost.





Pick vertex not in S with lowest cost





Pick vertex not in S with lowest cost and update neighbors



Pick vertex not in S with lowest cost and update neighbors



Pick vertex not in S with lowest cost and update neighbors



Pick vertex not in S with lowest cost and update neighbors



Pick vertex not in S with lowest cost and update neighbors

Data Structures



Priority queue for finding finding and deleting lowest cost vertex and for decreasing costs (Binary Heap works)

Priority Queue



Before the update, but after find min.

Priority Queue



Priority Queue



Time Complexity

- n vertices and m edges
- Initialize data structures O(n+m)
- Find min cost vertices O(n log n)
 - n delete mins
- Update costs O(m log n)
 - > Potentially m updates
- Update previous pointers O(m)
 - > Potentially m updates
- Total time O((n + m) log n) very fast.

Does It Always Work?

- Dijkstra's algorithm is an example of a greedy algorithm
- Greedy algorithms always make choices that currently seem the best
 - Short-sighted no consideration of long-term or global issues
 - > Locally optimal does not always mean globally optimal
- In Dijkstra's case choose the least cost node, but what if there is another path through other vertices that is cheaper?



• If the path to G is the next shortest path, the path to P must be at least as long. Therefore, any path through P to G cannot be shorter!

Inside the Cloud (Proof)

- Everything inside the cloud has the correct shortest path
- Proof is by induction on the number of nodes in the cloud:
 - Base case: Initial cloud is just the source with shortest path 0
 - Inductive hypothesis: cloud of k-1 nodes all have shortest paths
 - Inductive step: choose the least cost node G à has to be the shortest path to G (previous slide).
 Add k-th node G to the cloud

All Pairs Shortest Path

 Given a edge weighted directed graph G = (V,E) find for all u,v in V the length of the shortest path from u to v. Use matrix representation.



: = infinity Shortest Paths - Lecture 21

12/4/02

Matrix Representation

- C[i,j] = the cost of the edge (i,j)
 - > C[i,i] = 0 because no cost to stay where you are
 - > C[i,j] = infinity (:) if no edge from i to j.

Floyd – Warshall Algorithm

```
All_Pairs_Shortest_Path {
for k = 1 to n do
  for i = 1 to n do
    for j = 1 to n do
        C[i,j] := min(C[i,j], C[i,k] + C[k,j]);
}
Note x + : = : by definition
```

On termination C[i,j] is the length of the shortest path from i to j.

The Computation



Proof of Correctness

- After the k-th time through the loop C[i,j] is the length of the shortest path that only passes through vertices numbered 1,2,...,k.
 - Let C_k[i,j] be C[i,j] after k time through the loop.
- Base case: k = 0. C₀[i,j] is the cost of an edge that does not pass through any vertices.

Inductive Step

- Assume true for k-1.
 - A shortest path from i to j that only goes through vertices 1,2, ..., k does not go through vertex k at all.

• $C_{k}[i,j] = C_{k-1}[i,j]$

- All shortest paths from i to j that only goes through vertices 1,2, ..., k must go through vertex k.
 - $C_{k}[i,j] = C_{k-1}[i,k] + C_{k-1}[k,j]$

Cloud Argument



Time Complexity of All Pairs Shortest Path

- n is the number of vertices
- Three nested loops. O(n³)
 - > Shortest paths can be found too (see the book).
- Repeated Dijkstra's algorithm
 - > $O(n(n + m)\log n)$ (= $O(n^3 \log n)$ for dense graphs).
 - > Run Dijkstra starting at each vertex.
 - Dijkstra also gives the shortest paths not just their lengths.