CS583 Lecture 09

Jana Kosecka

Graph Algorithms Topological Sort Strongly Connected Component Minimum Spanning Tree

Many slides here are based on E. Demaine , D. Luebke, Kleinberg-Tardos slides

Graph Algs. Continued

- Review BFS
- Application of BSF check bipartiteness
- Review DFS
- Check for cycles













Breadth-First Search

```
BFS(G, s) {
    initialize vertices;
    Q = \{s\};
                     // Q is a queue (duh); initialize
 to s
    while (Q not empty) {
        u = RemoveTop(Q);
        for each v \in u->adj {
             if (v->color == WHITE)
                 v->color = GREY;
                 v - d = u - d + 1;
                                  What does v->d represent?
                 v \rightarrow p = u;
                Enqueue (Q, v); What does v->p represent?
        }
        u->color = BLACK;
    }
}
```



BSF

- Check for bi-partite graphs
- Graphs representing relationships
- All nodes can belong to two subsets
- There are no edges between subsets

Depth-First Search

- *Depth-first search* is another strategy for exploring a graph
- Explore "deeper" in the graph whenever possible
- Edges are explored out of the most recently discovered vertex *v* that still has unexplored edges
- When all of *v*'s edges have been explored, backtrack to the vertex from which *v* was discovered

Depth-First Search

- Vertices initially colored white
- Then colored gray when discovered
- Then black when finished

Depth-First Search: The Code

ł

}

DFS_Visit(u)

```
DFS(G)
{
   for each vertex u ∈ G->V
   {
      u->color = WHITE;
   }
   time = 0;
   for each vertex u ∈ G->V
   {
      if (u->color ==
   WHITE)
      DFS_Visit(u);
   }
}
```

Depth-First Sort Analysis

- This running time argument is an informal example of *amortized analysis*
- "Charge" the exploration of edge to the edge:
- Each loop in DFS_Visit can be attributed to an edge in the graph
- Runs once/edge if directed graph, twice if undirected
- Thus loop will run in O(E) time, algorithm O(V+E)
- Considered linear for graph, b/c adj list requires O(V+E) storage
- Important to be comfortable with this kind of reasoning and analysis









 DFS introduces an important distinction among edges in the original graph: *Tree edge*: encounter new (white) vertex *Back edge*: from descendent to ancestor *Forward edge*: from ancestor to descendent

Not a tree edge, though From grey node to black node



DFS: Kinds of edges

 DFS introduces an important distinction among edges in the original graph: *Tree edge*: encounter new (white) vertex *Back edge*: from descendent to ancestor *Forward edge*: from ancestor to descendent *Cross edge*: between a tree or subtrees

From a grey node to a black node



DFS: Kinds of edges

• DFS introduces an important distinction among edges in the original graph: *Tree edge*: encounter new (white) vertex

Back edge: from descendent to ancestor *Forward edge*: from ancestor to descendent *Cross edge*: between a tree or subtrees

• Note: tree & back edges are important; most algorithms don't distinguish forward & cross

DFS: Kinds Of Edges

- Thm 23.9 (22.10 in 3rd edition): If G is undirected, a DFS produces only tree and back edges
- Suppose you have *u.d* < *v.d*
- Then search discovered u before v, so first time v is discovered it is white hence the edge
 (u,v) is a tree edge
- Otherwise the search already explored this edge in direction from *v to u*
- edge must actually be a back edge since
- *u* is still gray

DFS And Graph Cycles

- Thm: An undirected graph is *acyclic* iff a DFS yields no back edges
- If acyclic, no back edges (because a back edge implies a cycle
- If no back edges, acyclic No back edges implies only tree edges (*Why*?) Only tree edges implies we have a tree or a forest Which by definition is acyclic
- Thus, can run DFS to find whether a graph has a cycle











DFS And Cycles

- Running time: O(V+E)
- We can actually determine if cycles exist in O(V) time: In an undirected acyclic forest, |E| ≤ |V| - 1 So count the edges: if ever see |V| distinct edges, must have seen a back edge along the way Why not just test if |E| <|V| and answer the question in constant time?

We can have some isolated component nodes not connected by any edges to the rest of the graph



DFS and DAGs

- Argue that a directed graph G is acyclic iff a DFS of G yields no back edges:
- Forward: if G is acyclic, will be no back edges Trivial: a back edge implies a cycle
- Backward: if no back edges, G is acyclic
- Argue contrapositive: Suppose G has a cycle ⇒ we will show that DFS will produce a back edge
- Let *v* be the vertex on the cycle first discovered, and *u* be the predecessor of *v* on the cycle
- When *v* discovered, whole cycle is white
- Must visit everything reachable from *v* before returning from DFS-Visit()
- So path from $u \rightarrow v$ is descendant of v hence (gray \rightarrow gray), thus (u, v) is a back edge







Topological Sort Algorithm

```
Topological-Sort()
{
Run DFS
When a vertex is finished, output it
On the front of linked list
Vertices are output in reverse
topological order
}
• Time: O(V+E)
• Correctness: Want to prove that
(u,v) \in G \Rightarrow u \rightarrow f > v \rightarrow f
```



Correctness of Topological Sort

- Claim: $(u,v) \in G \Rightarrow u \rightarrow f > v \rightarrow f$
- Show that if there is an edge from u to v, finishing time of u is greater then v
- When (u,v) is explored, u is gray
 v = gray ⇒ (u,v) is back edge. Contradiction (Why?)
 hence v cannot be gray since there are no cycles
 v = white ⇒ v becomes descendent of u ⇒ v→f < u→f
 (since must finish v before backtracking and finishing u)
 v = black ⇒ v already finished ⇒ v→f < u→f



- Undirected graph is connected: there is a path from $u \mathchar`-> v$ For all u and v

- Directed graph is strongly connected if there is a path from u->v for all u and v
- How can we find strongly connected components of a directed graph





Strongly Connected Components

- Call DFS to compute finishing times *f*[*u*] of each vertex
- Create transpose graph (directions of edges reversed)
- Call DFS on the transpose, but in the main loop of DFS, consider vertices in the decreasing order of *f*[*u*]
- Output the vertices of each tree in the depth-first forest formed in line 3 as a separate strongly connected component



Strongly connected components

- Property: Suppose you have two SCC's C and C'. If there is an edge between C and C', then vertex of C visited first has the highest finishing number: f(C) > f(C') suppose there is an edge u -> v from C to C'
- If DFS is started at C it visits all vertices in C and C' before it gets "stuck".



Applications

• MST is fundamental problem with diverse applications.

Network design. telephone, electrical, hydraulic, TV cable, computer, road

Approximation algorithms for NP-hard problems. traveling salesperson problem, Steiner tree

Indirect applications.

- max bottleneck paths
- LDPC codes for error correction
- image registration with Renyi entropy
- learning salient features for real-time face verification
- reducing data storage in sequencing amino acids in a protein
- model locality of particle interactions in turbulent fluid flows
- autoconfig protocol for Ethernet bridging to avoid cycles in a network

Cluster analysis.











Minimum Spanning Tree

- MSTs satisfy the *optimal substructure* property: an optimal tree is composed of optimal subtrees
- Let T be an MST of G with an edge (*u*,*v*) in the middle Removing (*u*,*v*) partitions T into two trees T₁ and T₂
 - Claim: T_1 is an MST of $G_1 = (V_1, E_1)$, and T_2 is an MST of $G_2 = (V_2, E_2)$ (Do V_1 and V_2 share vertices? Why?)
 - Proof: $w(T) = w(u,v) + w(T_1) + w(T_2)$ (There can't be a better tree than T_1 or T_2 , or T would be suboptimal)



Minimum Spanning Tree

• Thm:

Let T be MST of G, and let $A \subseteq T$ be subtree of T Let (u,v) be min-weight edge connecting A to V-A Then $(u,v) \in T$

• Proof: in book (see Thm 23.1)

Prim's Algorithm

```
\begin{split} \text{MST-Prim}(G, w, r) \\ & Q = V[G]; \\ & \text{for each } u \in Q \\ & \text{key}[u] = \infty; \\ & \text{key}[r] = 0; \\ & p[r] = \text{NULL}; \\ & \text{while } (Q \text{ not empty}) \\ & u = \text{ExtractMin}(Q); \\ & \text{for each } v \in \text{Adj}[u] \\ & \text{if } (v \in Q \text{ and } w(u, v) < \text{key}[v]) \\ & p[v] = u; \\ & \text{key}[v] = w(u, v); \end{split}
```





































Review: Prim's Algorithm

```
MST-Prim(G, w, r)
Q = V[G];
for each u \in Q
key[u] = \infty;
key[r] = 0;
p[r] = NULL;
while (Q not empty)
u = ExtractMin(Q);
for each v \in Adj[u]
if (v \in Q \text{ and } w(u,v) < key[v])
p[v] = u;
DecreaseKey(v, w(u,v));
```



Review: Prim's Algorithm

MST-Prim(G, w, r) What will be the running time? Q = V[G];A: Depends on queue for each $u \in Q$ binary heap: O(E lg V) $key[u] = \infty;$ Fibonacci heap: O(V lg V + E) key[r] = 0;p[r] = NULL;while (Q not empty) u = ExtractMin(Q); for each $v \in \operatorname{Adj}[u]$ if $(v \in Q \text{ and } w(u, v) < \text{key}[v])$ p[v] = u;key[v] = w(u,v);

ExtractMin total number of calls O(V log V)

DecreaseKey total number of calls O(E log V)

Total number of calls $O(V \log V + E \log V) = O(E \log V)$ Think why we can combine things in the expression above

Minimum Weight Spanning Tree Kruskal's Algorithm

```
Kruskal()
{
    T = Ø;
    for each v ∈ V
        MakeSet(v);
    sort E by increasing edge weight w
    for each (u,v) ∈ E (in sorted order)
        if FindSet(u) ≠ FindSet(v)
        T = T U {{u,v}};
        Union(FindSet(u), FindSet(v));
}
```



- Want a data structure to support disjoint sets Collection of disjoint sets S = {S_i}, S_i ∩ S_i = Ø
- Need to support following operations: MakeSet(x): S = S U {{x}} Union(S_i, S_j): S = S - {S_i, S_j} U {S_i U S_j} FindSet(X): return S_i ∈ S such that x ∈ S_i
- Before discussing implementation details, we look at example application: MSTs















































Correctness Of Kruskal's Algorithm

- Sketch of a proof that this algorithm produces an MST for *T*:
 - Assume algorithm is wrong: result is not an MST
 - Then algorithm adds a wrong edge at some point
 - If it adds a wrong edge, there must be a lower weight edge (cut and paste argument)
 - But algorithm chooses lowest weight edge at each step -> Contradiction
- Again, important to be comfortable with cut and paste arguments

Kruskal's Algorithm

```
What will affect the running time?
Kruskal()
                                                   1 Sort
{
                                      O(V) MakeSet() calls
   T = \emptyset;
                                       O(E) FindSet() calls
   for each \mathbf{v} \in \mathbf{V}
                                        O(V) Union() calls
                             (Exactly how many Union()s?)
       MakeSet(v);
   sort E by increasing edge weight w
    for each (u,v) \in E (in sorted order)
       if FindSet(u) ≠ FindSet(v)
           T = T U \{\{u,v\}\};
           Union(FindSet(u), FindSet(v));
}
```

Kruskal's Algorithm: Running Time

```
To summarize:
Sort edges: O(E lg E)
O(V) MakeSet()'s
O(E) FindSet()'s and Union()'s
Upshot:
Best disjoint-set union algorithm makes above
3 operation stake O((V+E)·α(V)), α almost constant
(slowly growing function of V)
Since E >= V-1 then we have O(E·α(V))
Also since α(V) = O(lg V) = O(lg E)
```

Overall thus O(E lg E), almost linear w/o sorting

Disjoint Sets (ch 21)

- In Kruskal's alg., Connected Components
- Need to do set membership and set union efficiently
- Typical operations on disjoint sets

member(a,s)
insert(a,s)
delete(a,s)
union(s1, s2, s3)
find(a)
make-set(x)

- Analysis in terms on n number of make-set operations
- And m total number of make-set, find, union (more details Later)



Single-Source Shortest Path

- Problem: given a weighted directed graph G, find the minimum-weight path from a given source vertex s to another vertex v
- "Shortest-path" = minimum weight
- Weight of path is sum of edges
- E.g., a road map: what is the shortest path from Faixfax to Washington DC?















Dijkstra's Algorithm

```
Dijkstra(G)

for each v \in V

d[v] = \infty;

d[s] = 0; S = \emptyset; Q = V;

while (Q \neq \emptyset)

u = ExtractMin(Q) DecraseKey() called?

S = S \cup \{u\};

for each v \in u \rightarrow Adj[]

if (d[v] > d[u]+w(u,v))

d[v] = d[u]+w(u,v);

A: O(E \lg V) using binary heap for Q

Can acheive O(V \lg V + E) with Fibonacci heaps
```

