

21 Disjoint Set Data Structure

Used, for example, to represent multiple strongly-connected components.

Since a strongly-connected component of a graph is the maximal set of vertices satisfying a property, a vertex/edge may only belong to one SCC.

Each object in a set needs to be able to identify if other elements belong to the same set. this is easy to do if each set has a representative object that all objects in the set point to.

```
1: procedure MAKESET(Object o) ▷ O(1)
2:   o.size ← 1
3:   o.parent ← 1
4: end procedure
```

```
1: procedure FINDSET(Object o) ▷ O(lg* n)
2:   if o.parent = 0 then return 0
3:   end if
4:   o.parent ← FINDSET(o.parent) ▷ path compression
5:   return o.parent
6: end procedure
```

```
1: procedure UNION(Object o, Object p) ▷ O(lg* n)
2:   x ← FINDSET(o)
3:   y ← FINDSET(p)
4:   if x.size ≥ y.size then
5:     y.parent ← x
6:     x.size ← x.size + y.size
7:   else
8:     x.parent ← y
9:     y.size ← y.size + x.size
10:  end if
11: end procedure
```

22 Graphs

BFS: $O(V+E)$. Use a queue. Add every vertex to the queue as it is discovered for the first time.

DFS: Use recursive calls to DFS. (note: many applications of DFS require keeping track of the discovery/finish times from calling DFS)

```

1: procedure DFS(Vertex v) ▷  $\Theta(V+E)$ 
2:   if v has been discovered then
3:     return
4:   end if
5:   discover v
6:   for all vertices u such that v has an edge to u do
7:     DFS(u)
8:   end for
9:   finish v
10: end procedure

```

In DFS(u):

- (u, v) edge is a back edge \iff v is discovered but not finished.
- (u, v) edge is a forward edge \iff v is discovered and finished and u is not finished.
- (u, v) edge is a cross edge \iff v is finished before u was discovered.

Example: With these discovery/finish times $(\overset{1}{A}(\overset{2}{B}(\overset{3}{C}A)\overset{4}{C}(\overset{5}{D})\overset{6}{D}(\overset{7}{E}D)\overset{8}{E})\overset{9}{B}D)\overset{10}{A}$ (where '(' indicates discovery and ')' indicates finishing), we have:

- (C, A) is a back edge because it occurs after (A and (C
- (E, D) is a cross edge because it occurs after (D,)D, and (E
- (A, D) is a forward edge because it occurs after (A, (D, and)D

Topological Sort: $\Theta(V+E)$. Call DFS(G). As each vertex finishes, add it to the beginning of a linked list. Return this list.

A strongly-connected component of G is a maximal set of vertices $C \subseteq V$ such that for every $u, v \in C$, $u \rightsquigarrow v$ and $v \rightsquigarrow u$.

To find strongly-connected components of graph G

```

1: procedure STRONGLYCONNECTEDCOMPONENTS ▷  $O(\Theta(V+E))$ 
2:   DFS(any vertex of G)
3:   for all vertices v in decreasing order of finish time do
4:      $G^T$ .DFS(v) ▷ All nodes discovered in this call to DFS belong to the
       same strongly-connected component
5:   end for
6: end procedure

```

An alternate form of topological sort exists, which is also $\Theta(V+E)$. In the case it is called on a graph with cycles, it produces a flawed ordering and has poorer failure modes.

```

1: procedure ALTOPOLOGICALSORT(Graph G) ▷  $\Theta(V+E)$ 
2:   Create a queue Topo
3:   Create a queue Q
4:   for all vertices  $v \in G$  do
5:     if  $v.indegree = 0$  then
6:       Q.ENQUEUE( $v$ )
7:     end if
8:   end for
9:   while Q nonempty do
10:    vertex  $u \leftarrow$  Q.DEQUEUE
11:    TOPO.ENQUEUE( $u$ )
12:    for all vertices  $v$  such that  $(u, v) \in G$  do
13:       $v.indegree \leftarrow v.indegree - 1$ 
14:      if  $v.indegree = 0$  then
15:        Q.ENQUEUE( $v$ )
16:      end if
17:    end for
18:  end while
19:  for all vertices  $v \in G$  do
20:    if  $v.indegree \neq 0$  then ▷ Found a cycle. Uh-oh!
21:      end if
22:  end for
23:  return Topo
24: end procedure

```

23 Minimum Spanning Trees

Kruskal's Algorithm makes use of the disjoint set data structure.

Prim's Algorithm makes use of a min-Priority Queue to retrieve vertices in non-descending order of weight; and also so that their weights can be decreased.

24 Single-source Shortest Paths

It's often useful to figure out the shortest paths within a graph G , from vertex v to every other vertex in G . This general problem has several variants.

```

1: procedure KRUSKALSMST(Graph G) ▷  $O(E \lg V)$ 
2:   for all vertices  $v \in G$  do
3:     MAKESET( $v$ )
4:   end for
5:   sort the edges of  $G$  in non-descending order by edge-weight
6:   for all edges  $(u, v)$  in non-descending order by weight do
7:     if FINDSET( $u$ )  $\neq$  FINDSET( $v$ ) then
8:       add edge  $(u, v)$  to MST
9:       UNION( $u, v$ )
10:    end if
11:  end for
12: end procedure

```

```

1: procedure PRIMSMST(Graph G) ▷  $O(E \lg V)$ 
2:   for all vertices  $v \in G$  do ▷ initialize all vertices
3:     key[ $v$ ]  $\leftarrow \infty$ 
4:      $\pi[v] \leftarrow null$ 
5:   end for
6:   key[some initial vertex]  $\leftarrow 0$ 
7:   PriorityQueue  $Q \leftarrow$  all vertices
8:   while  $Q$  not empty do
9:     vertex  $u \leftarrow$  EXTRACTMIN( $Q$ )
10:    for all vertices  $v$  such that  $(u, v)$  is an edge in  $G$  do
11:      if  $v \in Q$  and  $\text{weight}(u, v) < \text{key}[v]$  then
12:         $\pi[v] \leftarrow u$ 
13:        key[ $v$ ]  $\leftarrow \text{weight}(u, v)$ 
14:      end if
15:    end for
16:    add edge  $(\pi[u], u)$  to MST
17:  end while
18: end procedure

```

- Single-destination shortest paths: This variant reduces to a SSSP problem if you reverse the direction of every edge in G .
- Single-pair shortest paths: If you need to find the shortest path between vertices u and v in G , solving the SSSP general problem will give you an answer. While it may seem like overkill, no algorithm is known that will solve this problem better than a SSSP algorithm.

Shortest paths have an interesting property: If $a \rightsquigarrow c$ is the shortest path from a to c , and passes through vertex b , then $a \rightsquigarrow b$ is the shortest path from a to b and $b \rightsquigarrow c$ is the shortest path from b to c .

Some gotchas:

- Some SSSP algorithms don't produce correct results (or even halt) if the graph contains negative edge-weights. In particular, cycles of negative weight can create a shortest path with weight $-\infty$.
- Cycles of weight 0 contribute no savings to overall path weight, yet add edge traversals. So no SSSP should contain a 0-weight cycle.
- Positive weight cycles have higher weight than the same path without the cycles, so no SSSP will contain one.

For Graph G and starting vertex s , SSSP algorithms use these common functions:

```

1: procedure INITIALIZESINGLESOURCE(Graph  $G$ , Vertex  $v$ )           ▷  $\Theta(V)$ 
2:   for all vertices  $v \in G$  do
3:     distance[ $v$ ] ←  $\infty$ 
4:     parent[ $v$ ] ← null
5:   end for
6:   distance[ $s$ ] ← 0
7: end procedure

```

```

1: procedure RELAX(Source Vertex  $u$ , Destination Vertex  $v$ , Weight Function
    $w$ )                               ▷  $O(1)$ 
2:   if distance[ $v$ ] > distance[ $u$ ] +  $w(u, v)$  then
3:     distance[ $v$ ] ← distance[ $u$ ] +  $w(u, v)$ 
4:     parent[ $v$ ] ←  $u$ 
5:   end if
6: end procedure

```

Here's the first SSSP algorithm:

```

1: procedure BELLMANFORD(Graph G, Weight Function w, Starting Vertex
   s) ▷  $O(VE)$ 
2:   INITIALIZESINGLESOURCE(G, s)
3:   for i ← 1 to  $|V(G)| - 1$  do
4:     for all edge (u, v) in G do
5:       RELAX(u, v, w)
6:     end for
7:   end for
8:   for all edge (u, v) ∈ G do
9:     if distance[v] > distance[u] + w(u, v) then
10:      return false ▷ found a negative-weight cycle
11:     end if
12:   end for
13:   return true ▷ found no negative-weight cycles
14: end procedure

```

Though Bellman-Ford is slow, one benefit is that it can determine if there are negative-weight cycles in G , which are detectable if the shortest path keeps getting smaller past a fixed number of relaxation attempts.

DagShortestPaths only works on a DAG (directed acyclic graph), but is faster than Bellman-Ford.

```

1: procedure DAGSHORTESTPATHS(Graph G, Weight Function w, Starting
   Vertex s) ▷  $\Theta(V+E)$ 
2:   T ← a list of the vertices of G, sorted topologically
3:   INITIALIZESINGLESOURCE(G, s)
4:   for all vertices u ∈ T do
5:     for all vertices v such that (u, v) is an edge in G do
6:       RELAX(u, v, w)
7:     end for
8:   end for
9: end procedure

```

Dijkstra's/Dantzig's algorithm finds SSSPs in much the same way as Prim's algorithm finds MSTs. Whereas Prim's uses a Min-Priority Queue to keep track of minimum edge weights to a vertex, Dijkstra's associates with each vertex the minimum path weight to reach that vertex.

Dijkstra's algorithm will give the wrong answer if the graph it's used on has negative weights.

```

1: procedure DIJKSTRA(Graph G, Weight Function w, Starting Vertex s) ▷
   O(E lg V)
2:   INITIALIZESINGLESOURCE(G, s)
3:   MinPriorityQueue Q ← all vertices ∈ G
4:   while Q is not empty do
5:     Vertex u ← EXTRACTMIN(Q)
6:     for all vertices v such that (u, v) ∈ G do
7:       RELAX(u, v, w)
8:     end for
9:   end while
10: end procedure

```

25 All-Pairs Shortest Paths

The Floyd-Warshall algorithm is similar to the Bellman-Ford SSSP algorithm, but with more work and storage. Whereas B-F stored, for each destination vertex, what the parent and path-cost is on a path originating at some source vertex, F-W must store those two facts for each source vertex \times destination vertex pair.

Basically, you start with an adjacency matrix. Then, for each vertex, $\{v_1, v_2, \dots, v_n\}$ you derive a matrix from it that indicates the shortest paths from each vertex to each other vertex by storing only information about vertices that precede it on the path.

```

1: procedure FLOYDWARSHALL( $n \times n$  Matrix W)    ▷  $\Theta(V^3)$  time,  $O(V^2)$ 
   memory
2:    $n \leftarrow \text{rows}(W)$ 
3:    $D^0 \leftarrow W$ 
4:   for  $k \leftarrow 1$  to  $n$  do
5:     for  $i \leftarrow 1$  to  $n$  do
6:       for  $j \leftarrow 1$  to  $n$  do
7:          $d_{i,j}^k \leftarrow \min(d_{i,j}^{k-1}, d_{i,k}^{k-1} + d_{k,j}^{k-1})$ 
8:       end for
9:     end for
10:    Destroy  $D^{k-1}$ 
11:  end for
12:  return  $D^n$ 
13: end procedure

```

$i \rightsquigarrow k \rightsquigarrow j$ implies that $i \rightsquigarrow k$ and $k \rightsquigarrow j$. So if you're running Floyd-Warshall, you can save a lot of arithmetic by applying these rules:

1. if a row has ∞ in column k , then nothing in this row will change during

this iteration of the outer loop. i.e., $i \rightsquigarrow j$ does not pass through k if i has no path to k .

2. if a column has ∞ in row k , then nothing in that column will change during this k . This is because $i \rightsquigarrow j \rightsquigarrow k$ only if $j \rightsquigarrow k$.