Minimum Spanning Tree

By Jonathan Davis

STUDENTS-HUB.com

A spanning tree of a graph is just a subgraph that contains all the vertices and is a tree.

A graph may have many spanning trees.







Minimum Spanning Trees

A Minimum Spanning Tree (MST) is a subgraph of an undirected graph such that the subgraph spans (includes) all nodes, is connected, is acyclic, and has minimum total edge weight

Complete Graph



Minimum Spanning Tree

Definition

• A Minimum Spanning Tree (MST) is a subgraph of an undirected graph such that the subgraph spans (includes) all nodes, is connected, is acyclic, and has minimum total edge weight

Algorithms for Obtaining the Minimum Spanning Tree

- Prim's Algorithm
- Kruskal's Algorithm

This algorithm starts with one node. It then, one by one, adds a node that is unconnected to the new graph to the new graph, each time selecting the node whose connecting edge has the smallest weight out of the available nodes' connecting edges. The steps are:

The new graph is constructed - with one node from the old graph.
While new graph has fewer than n nodes,

1. Find the node from the old graph with the smallest connecting edge to the new graph,

2. Add it to the new graph

Every step will have joined one node, so that at the end we will have one graph with all the nodes and it will be a minimum spanning tree of the original graph.



























Minimum Spanning Tree





Analysis of Prim's Algorithm

Running Time = $O(m + n \log n)$ (m = edges, n = nodes)

If a heap is not used, the run time will be $O(n^2)$ instead of $O(m + n \log n)$. However, using a heap complicates the code since you're complicating the data structure. A Fibonacci heap is the best kind of heap to use, but again, it complicates the code.

Unlike Kruskal's, it doesn't need to see all of the graph at once. It can deal with it one piece at a time. It also doesn't need to worry if adding an edge will create a cycle since this algorithm deals primarily with the nodes, and not the edges.

For this algorithm the number of nodes needs to be kept to a minimum in addition to the number of edges. For small graphs, the edges matter more, while for large graphs the number of nodes matters more.

This algorithm creates a forest of trees. Initially the forest consists of n single node trees (and no edges). At each step, we add one edge (the cheapest one) so that it joins two trees together. If it were to form a cycle, it would simply link two nodes that were already part of a single connected tree, so that this edge would not be needed. The steps are:

- 1. The forest is constructed with each node in a separate tree.
- 2. The edges are placed in a priority queue.
- 3. Until we've added n-1 edges,
 - 1. Extract the cheapest edge from the queue,
 - 2. If it forms a cycle, reject it,
- 3. Else add it to the forest. Adding it to the forest will join two trees together.
- Every step will have joined two trees in the forest together, so that at the end, there will only be one tree in T.









Sort Edges

(in reality they are placed in a priority queue - not sorted - but sorting them makes the algorithm easier to visualize)



STUDENTS-HUB.com

















































Uploaded By: anonymous

3

3

С

F

5

Ι

1

Walk-Through



Consider an undirected, weight graph

STUDENTS-HUB.com



Sort the edges by increasing edge weight

	edge	d_v	
((D,E)	1	
((D,G)	2	
((E,G)	3	
((C,D)	3	
((G,H)	3	
((C,F)	3	
((B,C)	4	

edge	d_{v}	
(B ,E)	4	
(B,F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	

STUDENTS-HUB.com





edge	d_v			
(D,E)	1	\checkmark		
(D,G)	2			
(E,G)	3			
(C,D)	3			
(G,H)	3			
(C,F)	3			
(B,C)	4			

edge	d_{v}	
(B , E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





0				
edge	d_v			
(D,E)	1	\checkmark		
(D,G)	2	\checkmark		
(E,G)	3			
(C,D)	3			
(G,H)	3			
(C,F)	3			
(B,C)	4			

edge	d_{v}	
(B , E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	

Select first |V|–1 edges which do not generate a cycle



edge	d_v		edge	d_{v}	
(D,E)	1	\checkmark	(B,E)	4	
(D,G)	2	\checkmark	(B,F)	4	
(E,G)	3	χ	(B,H)	4	
(C,D)	3		(A,H)	5	
(G,H)	3		(D,F)	6	
(C,F)	3		(A,B)	8	
(B,C)	4		(A,F)	10	

Accepting edge (E,G) would create a cycle





•				
edge	d_v			
(D,E)	1	\checkmark		
(D,G)	2	\checkmark		
(E,G)	3	χ		
(C,D)	3	V		
(G,H)	3			
(C,F)	3			
(B,C)	4			

edge	d_{v}	
(B ,E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v			
(D,E)	1	\checkmark		
(D,G)	2	V		
(E,G)	3	χ		
(C,D)	3	V		
(G,H)	3	V		
(C,F)	3			
(B,C)	4			

edge	d_{v}	
(B ,E)	4	
(B,F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





-			
edge	d_v		
(D,E)	1	\checkmark	
(D,G)	2	\checkmark	
(E,G)	3	χ	
(C,D)	3	\checkmark	
(G,H)	3	\checkmark	
(C,F)	3	\checkmark	
(B,C)	4		

edge	d_{v}	
(B , E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	χ
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B , E)	4	
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	χ
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B,E)	4	X
(B , F)	4	
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	X
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B,E)	4	χ
(B , F)	4	χ
(B,H)	4	
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	χ
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B,E)	4	χ
(B,F)	4	χ
(B,H)	4	χ
(A,H)	5	
(D,F)	6	
(A,B)	8	
(A,F)	10	





edge	d_v	
(D,E)	1	\checkmark
(D,G)	2	\checkmark
(E,G)	3	X
(C,D)	3	\checkmark
(G,H)	3	\checkmark
(C,F)	3	\checkmark
(B,C)	4	\checkmark

edge	d_{v}	
(B , E)	4	χ
(B , F)	4	χ
(B,H)	4	χ
(A,H)	5	\checkmark
(D,F)	6	
(A,B)	8	
(A,F)	10	



Select first |V|–1 edges which do not generate a cycle

edge	d_v	
(D,E)	1	1
(D,G)	2	
(E,G)	3	χ
(C,D)	3	
(G,H)	3	\checkmark
(C,F)	3	
(B,C)	4	

edge	d_v		
(B,E)	4	X	
(B,F)	4	X	
(B,H)	4	χ	
(A,H)	5	\checkmark	
(D,F)	6		
(A,B)	8		considered
(A,F)	10		J

Done

Total Cost = $\Sigma d_v = 21$

Analysis of Kruskal's Algorithm

Running Time = $O(m \log n)$ (m = edges, n = nodes)

Testing if an edge creates a cycle can be slow unless a complicated data structure called a "union-find" structure is used.

It usually only has to check a small fraction of the edges, but in some cases (like if there was a vertex connected to the graph by only one edge and it was the longest edge) it would have to check all the edges.

This algorithm works best, of course, if the number of edges is kept to a minimum.

Conclusion

Kruskal's and Boruvka's have better running times if the number of edges is low, while Prim's has a better running time if both the number of edges and the number of nodes are low.

Boruvka's avoids the complicated data structures needed for the other two algorithms.

So, of course, the best algorithm depends on the graph and if you want to bear the cost of complex data structures.

The best algorithm that I know of is a hybrid of Boruvka's and Prim's, which I did not examine here. It does O(log log n) passes of Boruvka's and then switches to Prim's, resulting in a running time of O(m log log n). So, it's the fastest algorithm, but would, of course, require the Fibonacci heap for Prim's which Boruvka's avoids when used by itself. However, uncorder to keep things simple, I did not explore iteherenonymous