

Shortest Paths and Minimum Spanning Trees

David Kauchak
cs302
Spring 2012



Admin



Dijkstra's algorithm

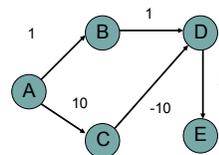
```

DIJKSTRA( $G, s$ )
1  for all  $v \in V$ 
2      $dist[v] \leftarrow \infty$ 
3      $prev[v] \leftarrow null$ 
4   $dist[s] \leftarrow 0$ 
5   $Q \leftarrow MAKEHEAP(V)$ 
6  while !EMPTY( $Q$ )
7      $u \leftarrow EXTRACTMIN(Q)$ 
8     for all edges  $(u, v) \in E$ 
9         if  $dist[v] > dist[u] + w(u, v)$ 
10             $dist[v] \leftarrow dist[u] + w(u, v)$ 
11            DECREASEKEY( $Q, v, dist[v]$ )
12             $prev[v] \leftarrow u$ 

```

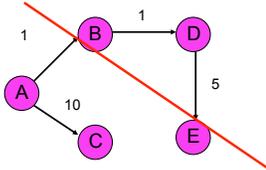


What about Dijkstra's on...?



What about Dijkstra's on...?

Dijkstra's algorithm only works for positive edge weights



Bounding the distance

- Another invariant: For each vertex v , $dist[v]$ is an upper bound on the actual shortest distance
 - start off at ∞
 - only update the value if we find a shorter distance
- An update procedure

$$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$$

$$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$$

Can we ever go wrong applying this update rule?

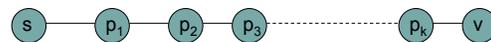
- We can apply this rule as many times as we want and will never underestimate $dist[v]$

When will $dist[v]$ be right?

- If u is along the shortest path to v and $dist[u]$ is correct

$$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- Consider the shortest path from s to v



$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- What happens if we update all of the vertices with the above update?

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- What happens if we update all of the vertices with the above update?

correct

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- What happens if we update all of the vertices with the above update?

correct correct

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- Does the order that we update the vertices matter?

correct correct

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- How many times do we have to do this for vertex p_i to have the correct shortest path from s ?
 - i times

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- How many times do we have to do this for vertex p_i to have the correct shortest path from s ?
 - i times

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- How many times do we have to do this for vertex p_i to have the correct shortest path from s ?
 - i times

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- How many times do we have to do this for vertex p_i to have the correct shortest path from s ?
 - i times

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- How many times do we have to do this for vertex p_i to have the correct shortest path from s ?
 - i times

$dist[v] = \min \{dist[v], dist[u] + w(u, v)\}$

- $dist[v]$ will be right if u is along the shortest path to v and $dist[u]$ is correct
- What is the longest (vertex-wise) the path from s to any node v can be?
 - $|V| - 1$ edges/vertices

Bellman-Ford algorithm

```

BELLMAN-FORD( $G, s$ )
1 for all  $v \in V$ 
2    $dist[v] \leftarrow \infty$ 
3    $prev[v] \leftarrow null$ 
4  $dist[s] \leftarrow 0$ 
5 for  $i \leftarrow 1$  to  $|V| - 1$ 
6   for all edges  $(u, v) \in E$ 
7     if  $dist[v] > dist[u] + w(u, v)$ 
8        $dist[v] \leftarrow dist[u] + w(u, v)$ 
9        $prev[v] \leftarrow u$ 
10 for all edges  $(u, v) \in E$ 
11   if  $dist[v] > dist[u] + w(u, v)$ 
12     return false
    
```

Bellman-Ford algorithm

```

BELLMAN-FORD( $G, s$ )
1 for all  $v \in V$ 
2    $dist[v] \leftarrow \infty$ 
3    $prev[v] \leftarrow null$ 
4  $dist[s] \leftarrow 0$ 
5 for  $i \leftarrow 1$  to  $|V| - 1$ 
6   for all edges  $(u, v) \in E$ 
7     if  $dist[v] > dist[u] + w(u, v)$ 
8        $dist[v] \leftarrow dist[u] + w(u, v)$ 
9        $prev[v] \leftarrow u$ 
10 for all edges  $(u, v) \in E$ 
11   if  $dist[v] > dist[u] + w(u, v)$ 
12     return false
    
```

Initialize all the distances

do it $|V| - 1$ times

iterate over all edges/vertices and apply update rule

Bellman-Ford algorithm

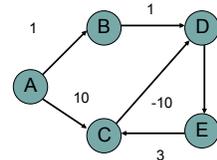
```

BELLMAN-FORD( $G, s$ )
1  for all  $v \in V$ 
2       $dist[v] \leftarrow \infty$ 
3       $prev[v] \leftarrow null$ 
4   $dist[s] \leftarrow 0$ 
5  for  $i \leftarrow 1$  to  $|V| - 1$ 
6      for all edges  $(u, v) \in E$ 
7          if  $dist[v] > dist[u] + w(u, v)$ 
8               $dist[v] \leftarrow dist[u] + w(u, v)$ 
9               $prev[v] \leftarrow u$ 
10 for all edges  $(u, v) \in E$ 
11     if  $dist[v] > dist[u] + w(u, v)$ 
12         return false
    
```

check for negative cycles

Negative cycles

What is the shortest path from a to e?

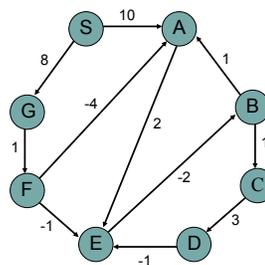


Bellman-Ford algorithm

```

BELLMAN-FORD( $G, s$ )
1  for all  $v \in V$ 
2       $dist[v] \leftarrow \infty$ 
3       $prev[v] \leftarrow null$ 
4   $dist[s] \leftarrow 0$ 
5  for  $i \leftarrow 1$  to  $|V| - 1$ 
6      for all edges  $(u, v) \in E$ 
7          if  $dist[v] > dist[u] + w(u, v)$ 
8               $dist[v] \leftarrow dist[u] + w(u, v)$ 
9               $prev[v] \leftarrow u$ 
10 for all edges  $(u, v) \in E$ 
11     if  $dist[v] > dist[u] + w(u, v)$ 
12         return false
    
```

Bellman-Ford algorithm



How many edges is the shortest path from s to:

A:

Bellman-Ford algorithm

How many edges is the shortest path from s to:

A: 3

Bellman-Ford algorithm

How many edges is the shortest path from s to:

A: 3

B:

Bellman-Ford algorithm

How many edges is the shortest path from s to:

A: 3

B: 5

Bellman-Ford algorithm

How many edges is the shortest path from s to:

A: 3

B: 5

D:

Bellman-Ford algorithm

How many edges is the shortest path from s to:

- A: 3
- B: 5
- D: 7

Bellman-Ford algorithm

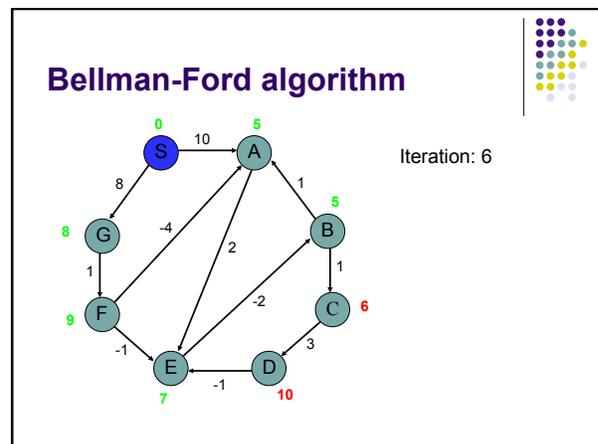
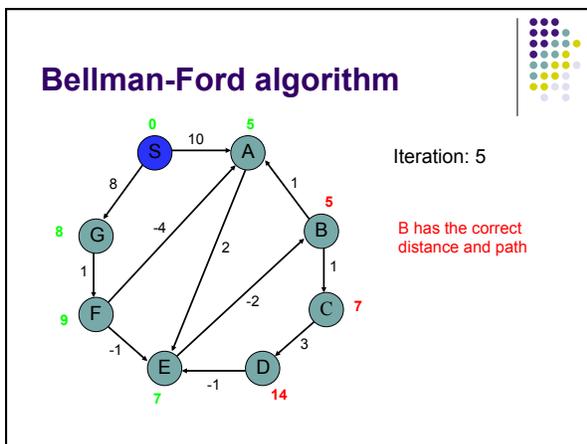
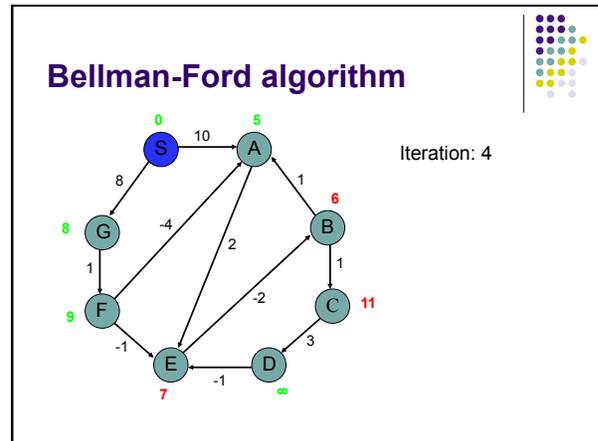
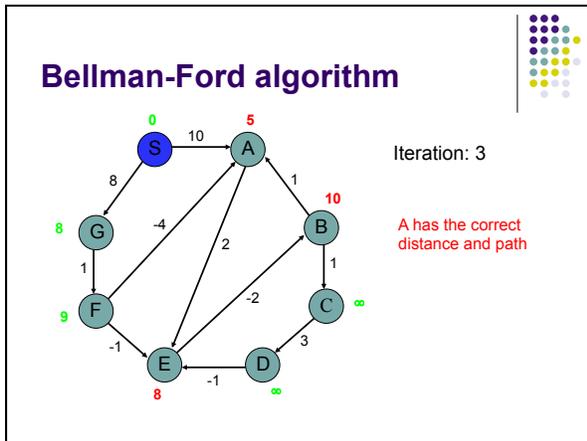
Iteration: 0

Bellman-Ford algorithm

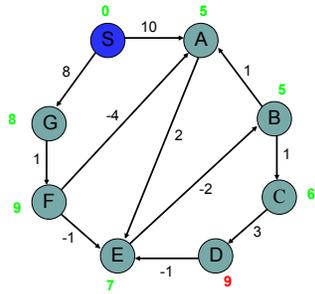
Iteration: 1

Bellman-Ford algorithm

Iteration: 2



Bellman-Ford algorithm



Iteration: 7

D (and all other nodes) have the correct distance and path

Correctness of Bellman-Ford

Loop invariant:

```

BELLMAN-FORD( $G, s$ )
1 for all  $v \in V$ 
2    $dist[v] \leftarrow \infty$ 
3    $prev[v] \leftarrow null$ 
4    $dist[s] \leftarrow 0$ 
5 for  $i \leftarrow 1$  to  $|V| - 1$ 
6   for all edges  $(u, v) \in E$ 
7     if  $dist[v] > dist[u] + w(u, v)$ 
8        $dist[v] \leftarrow dist[u] + w(u, v)$ 
9        $prev[v] \leftarrow u$ 
10 for all edges  $(u, v) \in E$ 
11   if  $dist[v] > dist[u] + w(u, v)$ 
12     return false
    
```

Correctness of Bellman-Ford

Loop invariant: After iteration i , all vertices with shortest paths from s of length i edges or less have correct distances

```

BELLMAN-FORD( $G, s$ )
1 for all  $v \in V$ 
2    $dist[v] \leftarrow \infty$ 
3    $prev[v] \leftarrow null$ 
4    $dist[s] \leftarrow 0$ 
5 for  $i \leftarrow 1$  to  $|V| - 1$ 
6   for all edges  $(u, v) \in E$ 
7     if  $dist[v] > dist[u] + w(u, v)$ 
8        $dist[v] \leftarrow dist[u] + w(u, v)$ 
9        $prev[v] \leftarrow u$ 
10 for all edges  $(u, v) \in E$ 
11   if  $dist[v] > dist[u] + w(u, v)$ 
12     return false
    
```

Runtime of Bellman-Ford

```

BELLMAN-FORD( $G, s$ )
1 for all  $v \in V$ 
2    $dist[v] \leftarrow \infty$ 
3    $prev[v] \leftarrow null$ 
4    $dist[s] \leftarrow 0$ 
5 for  $i \leftarrow 1$  to  $|V| - 1$ 
6   for all edges  $(u, v) \in E$ 
7     if  $dist[v] > dist[u] + w(u, v)$ 
8        $dist[v] \leftarrow dist[u] + w(u, v)$ 
9        $prev[v] \leftarrow u$ 
10 for all edges  $(u, v) \in E$ 
11   if  $dist[v] > dist[u] + w(u, v)$ 
12     return false
    
```

$$O(|V| |E|)$$

Runtime of Bellman-Ford

```

BELLMAN-FORD( $G, s$ )
1  for all  $v \in V$ 
2      $dist[v] \leftarrow \infty$ 
3      $prev[v] \leftarrow null$ 
4   $dist[s] \leftarrow 0$ 
5  for  $i \leftarrow 1$  to  $|V| - 1$ 
6     for all edges  $(u, v) \in E$ 
7         if  $dist[v] > dist[u] + w(u, v)$ 
8              $dist[v] \leftarrow dist[u] + w(u, v)$ 
9              $prev[v] \leftarrow u$ 
10 for all edges  $(u, v) \in E$ 
11     if  $dist[v] > dist[u] + w(u, v)$ 
12         return false

```

Can you modify the algorithm to run faster (in some circumstances)?

All pairs shortest paths

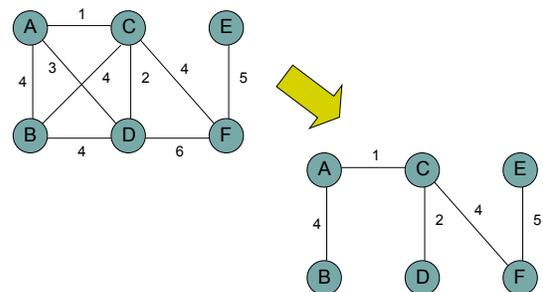
- Simple approach
 - Call Bellman-Ford $|V|$ times
 - $O(|V|^2 |E|)$
- Floyd-Warshall – $\Theta(|V|^3)$
- Johnson's algorithm – $O(|V|^2 \log |V| + |V| |E|)$

Minimum spanning trees

- What is the lowest weight set of edges that connects all vertices of an undirected graph with positive weights
- Input: An undirected, positive weight graph, $G=(V, E)$
- Output: A tree $T=(V, E')$ where $E' \subseteq E$ that minimizes

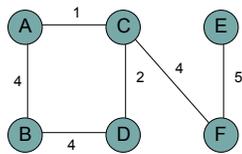
$$weight(T) = \sum_{e \in E'} w_e$$

MST example



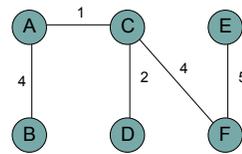
MSTs

Can an MST have a cycle?



MSTs

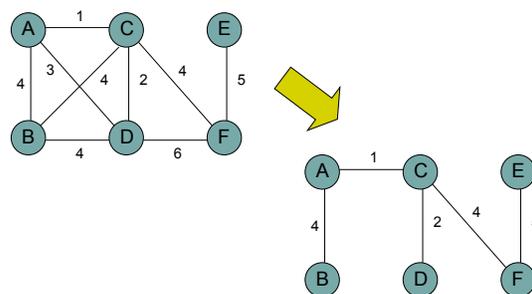
Can an MST have a cycle?



Applications?

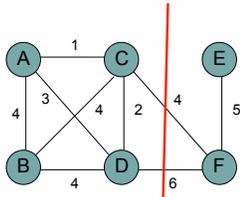
- Connectivity
 - Networks (e.g. communications)
 - Circuit design/wiring
- hub/spoke models (e.g. flights, transportation)
- Traveling salesman problem?

Algorithm ideas?



Cuts

- A cut is a partitioning of the vertices into two sets S and V-S
- An edge “crosses” the cut if it connects a vertex $u \in V$ and $v \in V-S$



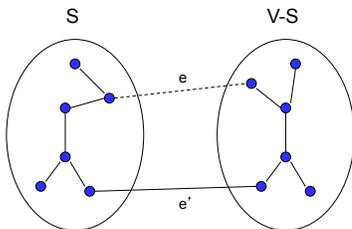
Minimum cut property

Given a partition S, let edge e be the minimum cost edge that **crosses** the partition. *Every* minimum spanning tree contains edge e.

Prove this!

Minimum cut property

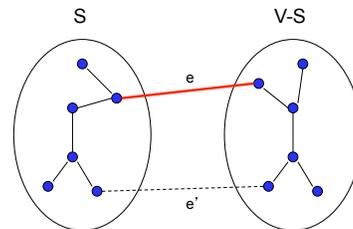
Given a partition S, let edge e be the minimum cost edge that **crosses** the partition. *Every* minimum spanning tree contains edge e.



Consider an MST with edge e' that is not the minimum edge

Minimum cut property

Given a partition S, let edge e be the minimum cost edge that **crosses** the partition. *Every* minimum spanning tree contains edge e.



Using e instead of e', still connects the graph, but produces a tree with smaller weights

Kruskal's algorithm

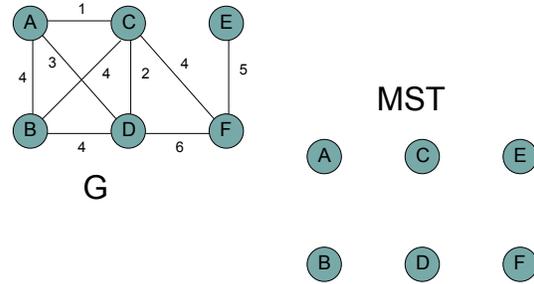
Given a partition S , let edge e be the minimum cost edge that **crosses** the partition. Every minimum spanning tree contains edge e .

```

KRUSKAL( $G$ )
1 for all  $v \in V$ 
2   MAKESET( $v$ )
3  $T \leftarrow \{\}$ 
4 sort the edges of  $E$  by weight
5 for all edges  $(u, v) \in E$  in increasing order of weight
6   if FIND-SET( $u$ )  $\neq$  FIND-SET( $v$ )
7     add edge to  $T$ 
8     UNION(FIND-SET( $u$ ), FIND-SET( $v$ ))
    
```

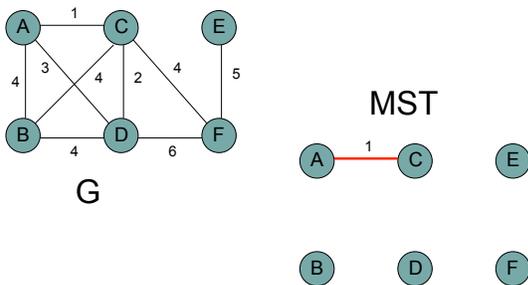
Kruskal's algorithm

Add smallest edge that connects two sets not already connected



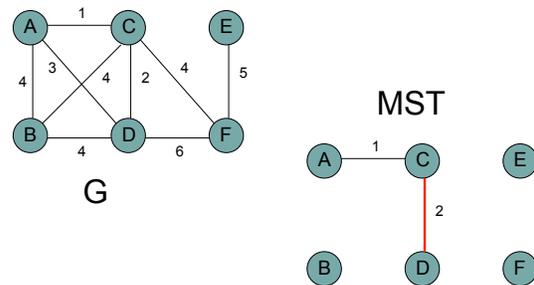
Kruskal's algorithm

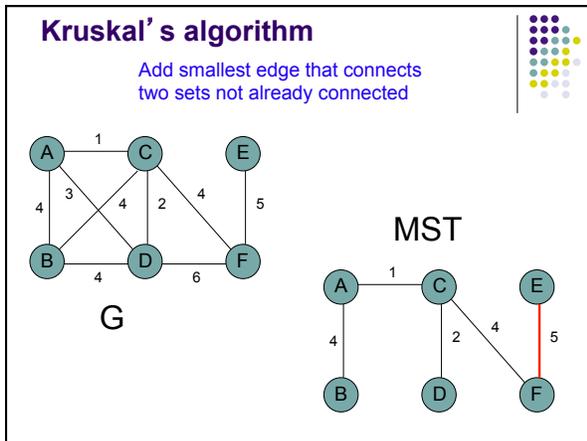
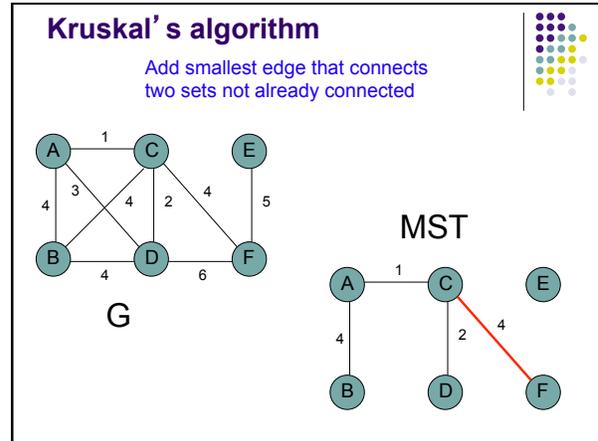
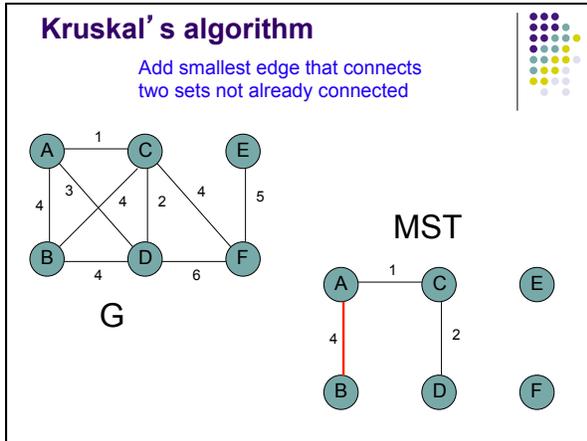
Add smallest edge that connects two sets not already connected



Kruskal's algorithm

Add smallest edge that connects two sets not already connected





Correctness of Kruskal's

- Never adds an edge that connects already connected vertices
- Always adds lowest cost edge to connect two sets. By min cut property, that edge must be part of the MST

```

KRUSKAL(G)
1  for all v in V
2      MAKESET(v)
3  T ← {}
4  sort the edges of E by weight
5  for all edges (u,v) in E in increasing order of weight
6      if FIND-SET(u) ≠ FIND-SET(v)
7          add edge to T
8          UNION(FIND-SET(u),FIND-SET(v))
    
```

Running time of Kruskal's

```

KRUSKAL(G)
1 for all v in V
2   MAKESET(v)
3 T ← {}
4 sort the edges of E by weight
5 for all edges (u, v) in E in increasing order of weight
6   if FIND-SET(u) ≠ FIND-SET(v)
7     add edge to T
8     UNION(FIND-SET(u), FIND-SET(v))
    
```

Running time of Kruskal's

```

KRUSKAL(G)
1 for all v in V           |V| calls to MakeSet
2   MAKESET(v)
3 T ← {}
4 sort the edges of E by weight   O(|E| log |E|)
5 for all edges (u, v) in E in increasing order of weight
6   if FIND-SET(u) ≠ FIND-SET(v)  2 |E| calls to FindSet
7     add edge to T
8     UNION(FIND-SET(u), FIND-SET(v))  |V| calls to Union
    
```

Running time of Kruskal's

Disjoint set data structure

$O(|E| \log |E|) +$

	MakeSet	FindSet E calls	Union V calls	Total
Linked lists	V	$O(V E)$	V	$O(V E + E \log E)$ $O(V E)$
Linked lists + heuristics	V	$O(E \log V)$	V	$O(E \log V + E \log E)$ $O(E \log E)$

Prim's algorithm

```

PRIM(G, r)
1 for all v in V
2   key[v] ← ∞
3   prev[v] ← null
4 key[r] ← 0
5 H ← MAKEHEAP(key)
6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) in E
10    if !visited[v] and w(u, v) < key[v]
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
    
```

Prim's algorithm

```

PRIM(G, r)
1 for all v in V
2   key[v] ← ∞
3   prev[v] ← null
4 key[r] ← 0
5 H ← MAKEHEAP(key)
6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) in E
10    if !visited[v] and w(u, v) < key[v]
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
            
```

```

Dijkstra(G, s)
1 for all v in V
2   dist[v] ← ∞
3   prev[v] ← null
4 dist[s] ← 0
5 Q ← MAKEHEAP(V)
6 while !EMPTY(Q)
7   u ← EXTRACTMIN(Q)
8   for all edges (u, v) in E
9     if dist[u] > dist[u] + w(u, v)
10      dist[v] ← dist[u] + w(u, v)
11      DECREASEKEY(Q, v, dist[v])
12     prev[v] ← u
            
```

Prim's algorithm

```

PRIM(G, r)
1 for all v in V
2   key[v] ← ∞
3   prev[v] ← null
4 key[r] ← 0
5 H ← MAKEHEAP(key)
6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) in E
10    if !visited[v] and w(u, v) < key[v]
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
            
```

Prim's algorithm

Start at some root node and build out the MST by adding the lowest weighted edge at the frontier

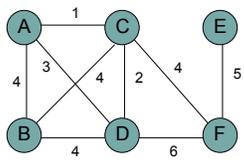
```

PRIM(G, r)
1 for all v in V
2   key[v] ← ∞
3   prev[v] ← null
4 key[r] ← 0
5 H ← MAKEHEAP(key)
6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) in E
10    if !visited[v] and w(u, v) < key[v]
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
            
```

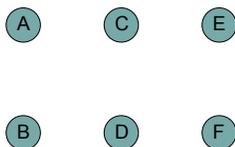
Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) in E
10    if !visited[v] and w(u, v) < key[v]
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
            
```



MST



Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
  
```

MST

A	C	E
∞	∞	∞
B	D	F
∞	∞	0

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
  
```

MST

A	C	E
∞	4	5
B	D	F
∞	6	0

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
  
```

MST

A	C	E
∞	4	5
B	D	F
∞	6	0

Note: A red line connects node C to node F.

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11      DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
  
```

MST

A	C	E
1	4	5
B	D	F
4	2	0

Note: A red line connects node C to node F.

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11      DECREASE-KEY(v, w(u, v))
12      prev[v] ← u
  
```

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11      DECREASE-KEY(v, w(u, v))
12      prev[v] ← u
  
```

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11      DECREASE-KEY(v, w(u, v))
12      prev[v] ← u
  
```

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11      DECREASE-KEY(v, w(u, v))
12      prev[v] ← u
  
```

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11     DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
    
```

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11     DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
    
```

Prim's

```

6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) ∈ E
10    if !visited[v] and w(u, v) < key(v)
11     DECREASE-KEY(v, w(u, v))
12    prev[v] ← u
    
```

Correctness of Prim's?

- Can we use the min-cut property?
 - Given a partition S, let edge e be the minimum cost edge that **crosses** the partition. Every minimum spanning tree contains edge e.
- Let S be the set of vertices visited so far
- The only time we add a new edge is if it's the lowest weight edge from S to V-S

Running time of Prim's

```

PRIM(G, r)
1 for all v in V
2   key[v] ← ∞
3   prev[v] ← null
4 key[r] ← 0
5 H ← MAKEHEAP(key)
6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) in E
10    if !visited[v] and w(u, v) < key[v]
11      DECREASE-KEY(v, w(u, v))
12      prev[v] ← u
    
```

Running time of Prim's

```

PRIM(G, r)
1 for all v in V
2   key[v] ← ∞
3   prev[v] ← null
4 key[r] ← 0
5 H ← MAKEHEAP(key)
6 while !Empty(H)
7   u ← EXTRACT-MIN(H)
8   visited[u] ← true
9   for each edge (u, v) in E
10    if !visited[v] and w(u, v) < key[v]
11      DECREASE-KEY(v, w(u, v))
12      prev[v] ← u
    
```

$\Theta(|V|)$
 $\Theta(|V|)$
 $|V|$ calls to Extract-Min
 $|E|$ calls to Decrease-Key

Running time of Prim's

Same as Dijkstra's algorithm

	1 MakeHeap	$ V $ ExtractMin	$ E $ DecreaseKey	Total
Array	$O(V)$	$O(V ^2)$	$O(E)$	$O(V ^2)$
Bin heap	$O(V)$	$O(V \log V)$	$O(E \log V)$	$O((V + E) \log V)$ $O(E \log V)$
Fib heap	$O(V)$	$O(V \log V)$	$O(E)$	$O(V \log V + E)$

Kruskal's: $O(|E| \log |E|)$