Greedy Graph Algorithms

T. M. Murali

September 10, 12, 19, 2018

Algorithm Design

- Start discussion of different ways of designing algorithms.
- Greedy algorithms, divide and conquer, dynamic programming, flow-based approaches.
- Discuss principles that can solve a variety of problem types.
- Design an algorithm, prove its correctness, analyse its complexity.

Algorithm Design

- Start discussion of different ways of designing algorithms.
- Greedy algorithms, divide and conquer, dynamic programming, flow-based approaches.
- Discuss principles that can solve a variety of problem types.
- Design an algorithm, prove its correctness, analyse its complexity.
- Greedy algorithms: make the current best choice.
    - First discuss greedy graph algorithms.
    - Will discuss greedy algorithms for scheduling (Chapters 4.1 to 4.3) later in the semester.

Shortest Paths Problem

- $G(V, E)$ is a connected directed graph. Each edge e has a length $l(e) \geq 0$.
- V has n nodes and E has m edges.
- *Length of a path* P is the sum of the lengths of the edges in P.
- Goal is to determine the shortest path from a specified start node s to each node in V.
- Aside: If G is undirected, convert to a directed graph by replacing each edge in G by two directed edges.

Shortest Paths Problem

- $G(V, E)$ is a connected directed graph. Each edge e has a length $l(e) \geq 0$.
- V has n nodes and E has m edges.
- *Length of a path* P is the sum of the lengths of the edges in P.
- Goal is to determine the shortest path from a specified start node s to each node in V.
- Aside: If G is undirected, convert to a directed graph by replacing each edge in G by two directed edges.

SHORTEST PATHS

INSTANCE: A directed graph $G(V, E)$, a function $l : E \rightarrow \mathbb{R}^+$, and a node $s \in V$

SOLUTION: A set $\{P_u, u \in V\}$ of paths, where P_u is the shortest path in G from s to u.

Shortest Paths Problem Instance

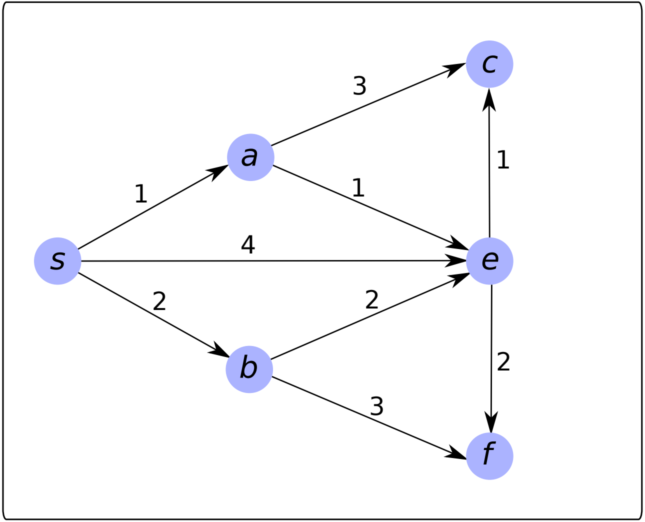

Generalizing BFS

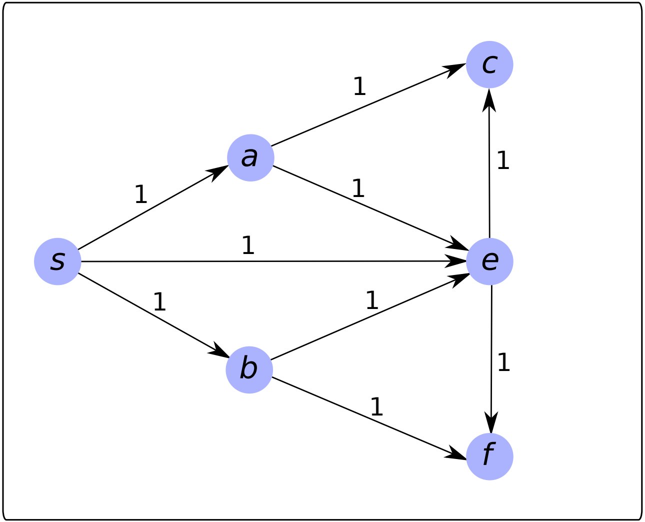

Generalizing BFS

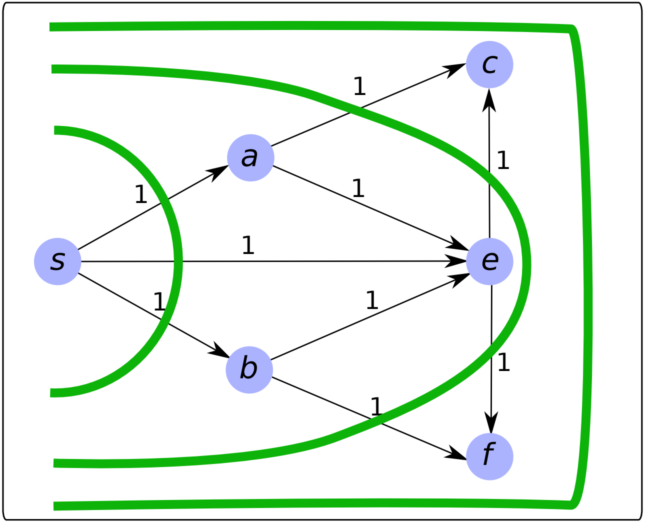

Unweighted graph: Use BFS. Process nodes in non-decreasing order of distance.

Generalizing BFS

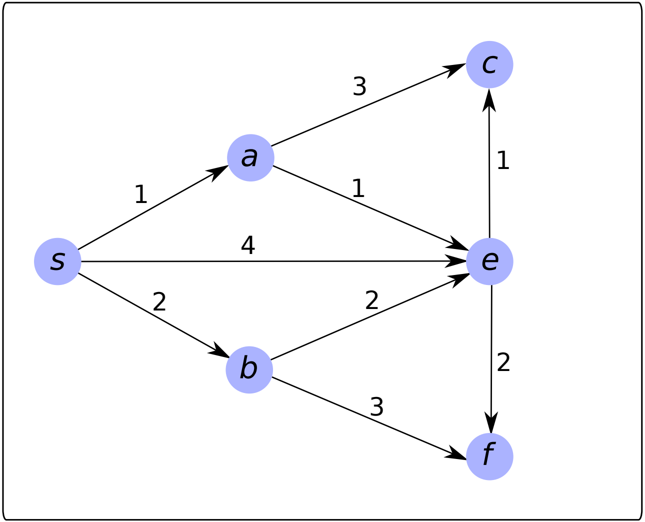

Weighted graph: Edge weights are integers. Can we make the graph unweighted?

Generalizing BFS

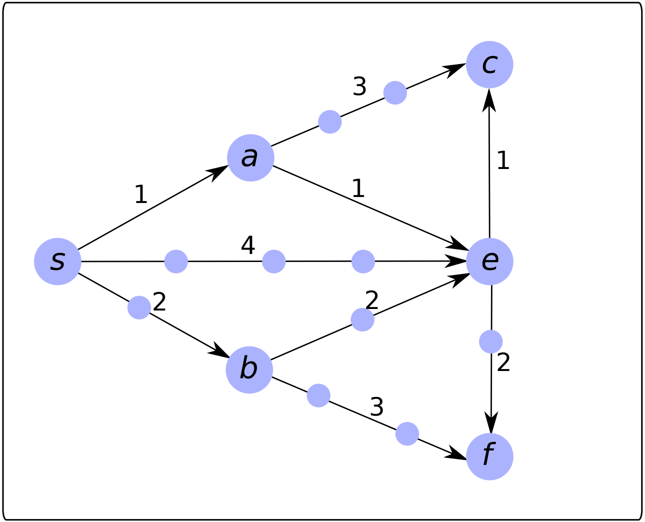

Add dummy nodes: Edge of weight w gets $w - 1$ nodes.

Generalizing BFS

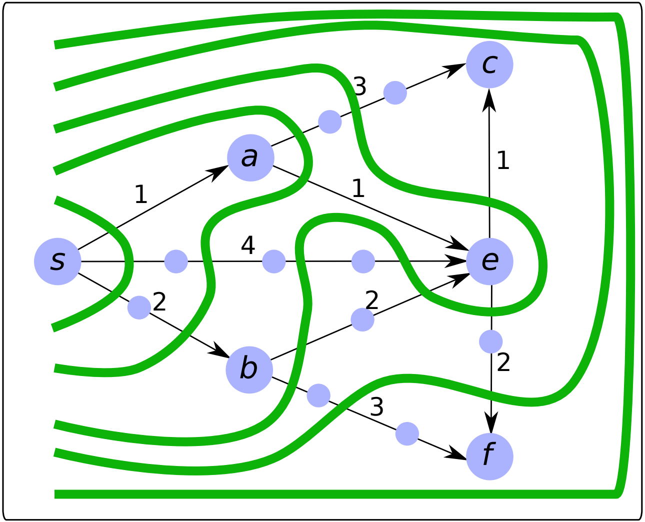

Dummy nodes: BFS computes shortest paths correctly. Running time is *Pseudo-polynomial time*: depends on input *values*.

Generalizing BFS

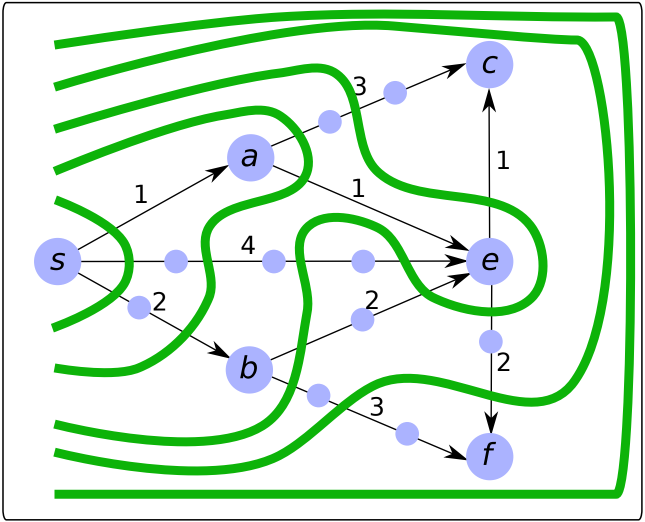

Dummy nodes: BFS computes shortest paths correctly. Running time is $O(m + n + \sum_{e \in E} l(e))$. *Pseudo-polynomial time*: depends on input *values*.

Generalizing BFS to Dijkstra's Algorithm

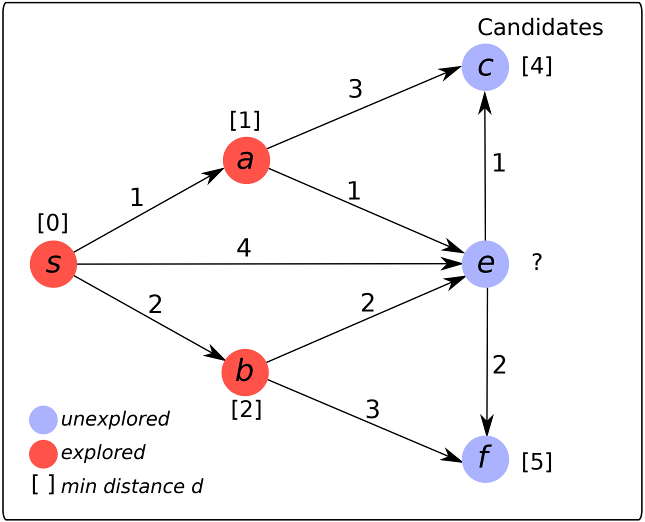

Like BFS: explore nodes in non-increasing order of distance from s. Once a node is explored, its distance is fixed.

Generalizing BFS to Dijkstra's Algorithm

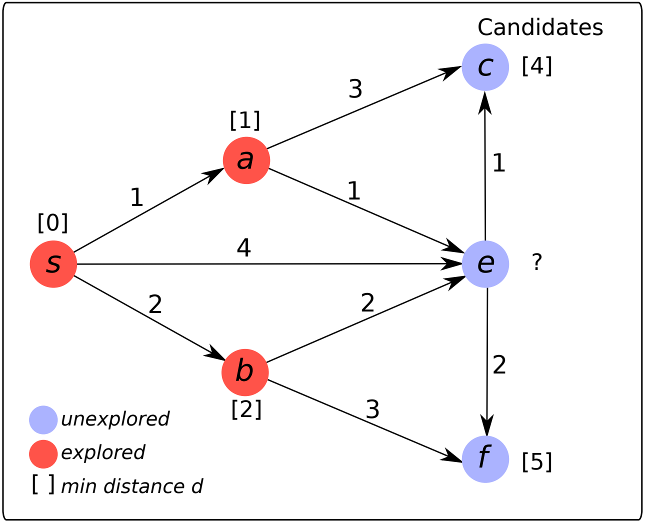

Unlike BFS: Layers are not uniform. Which node to process next? Candidates are nodes with an edge from a explored node.

Generalizing BFS to Dijkstra's Algorithm

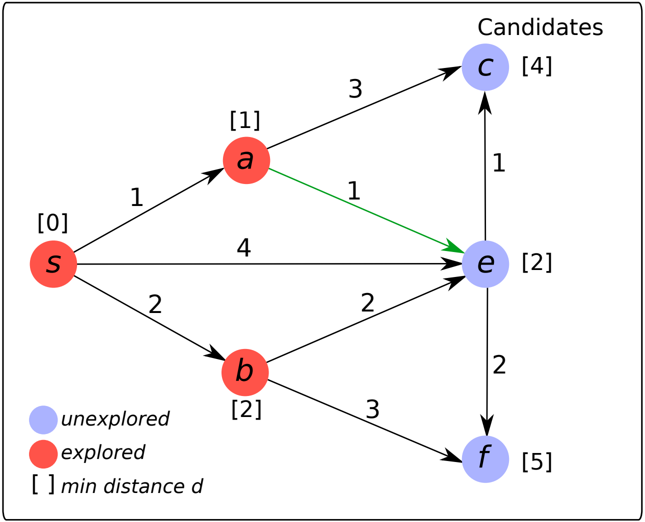

For each unexplored node, determine "best" preceding explored node.

Generalizing BFS to Dijkstra's Algorithm

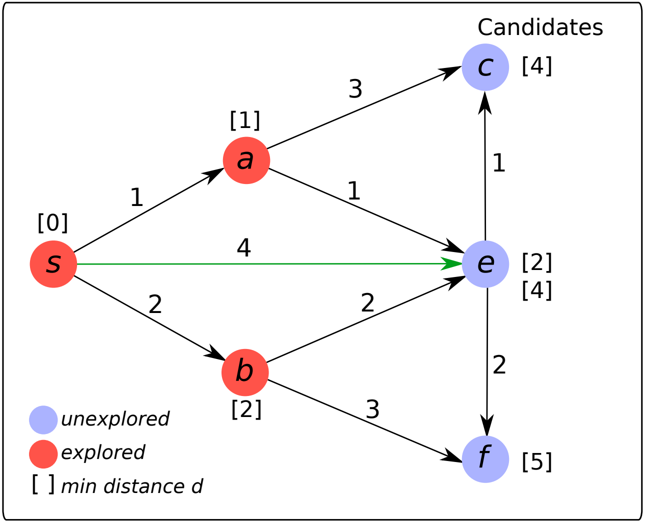

For each unexplored node, determine "best" preceding explored node.

Generalizing BFS to Dijkstra's Algorithm

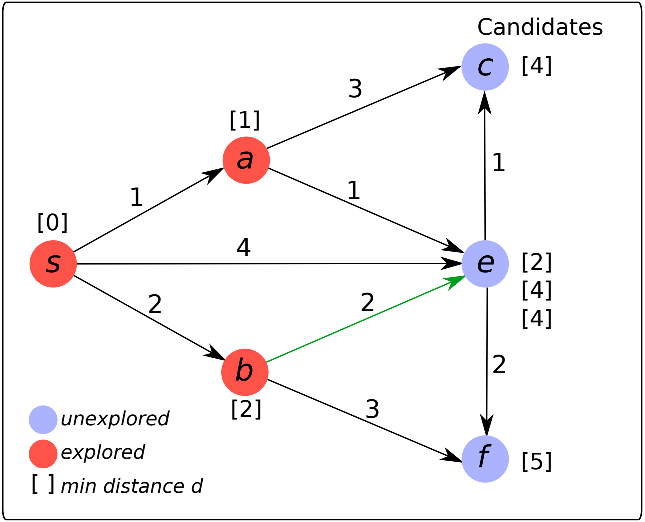

For each unexplored node, determine "best" preceding explored node.

Generalizing BFS to Dijkstra's Algorithm

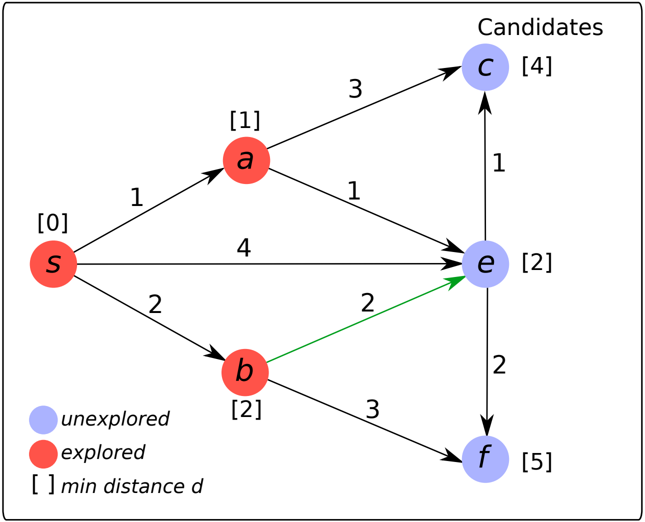

For each unexplored node, determine "best" preceding explored node. Record shortest path length only through explored nodes.

Generalizing BFS to Dijkstra's Algorithm

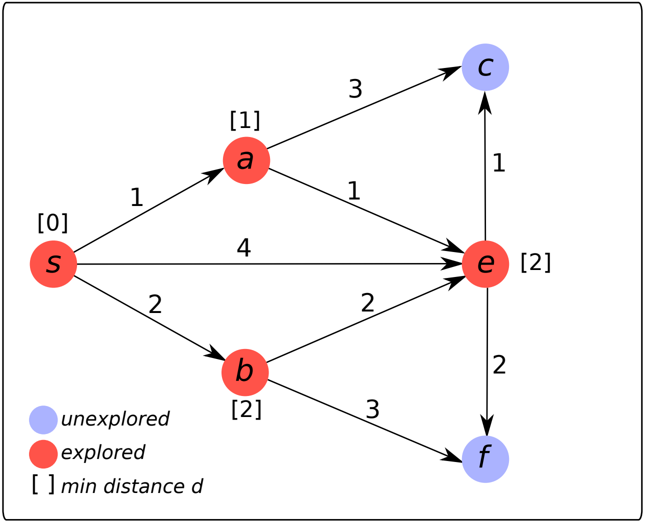

Explore node with smallest path length only through explored nodes.

Generalizing BFS to Dijkstra's Algorithm

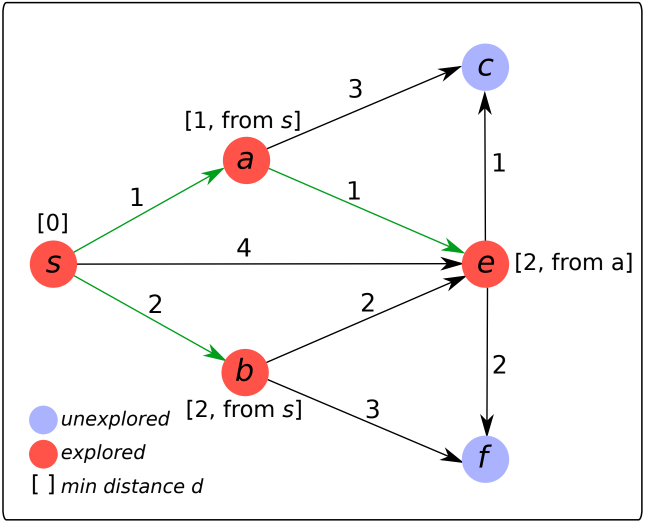

Like BFS: Record previous node in the computed path.

Generalizing BFS to Dijkstra's Algorithm

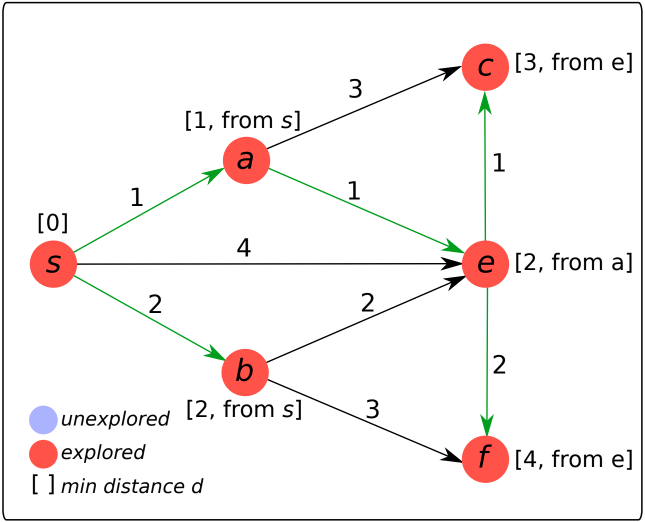

Follow previous nodes to compute shortest path. Like BFS: these edges form a tree.

Idea Underlying Dijkstra's Algorithm

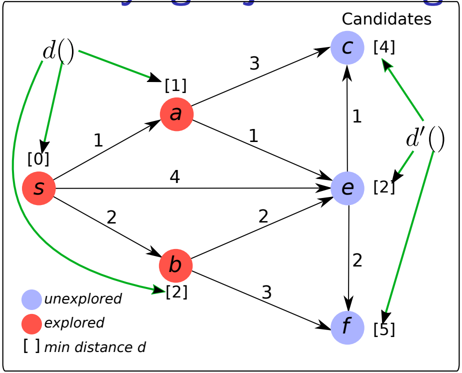

- Maintain a set S of explored nodes.
  - For each node $u \in S$, compute a value $d(u)$, which (we will prove) is the length of the shortest path from s to u.
  - For each node $x \notin S$, maintain a value $d'(x)$, which is the length of the shortest path from s to x using only the nodes in S (and x, of course).

Idea Underlying Dijkstra's Algorithm

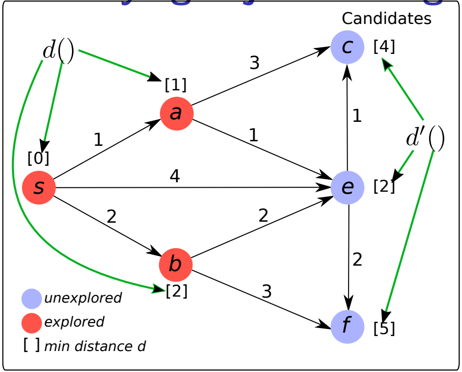

- Maintain a set S of explored nodes.
  - For each node $u \in S$, compute a value $d(u)$, which (we will prove) is the length of the shortest path from s to u.
  - For each node $x \notin S$, maintain a value $d'(x)$, which is the length of the shortest path from s to x using only the nodes in S (and x, of course).
- "Greedily" add a node v to S that has the smallest value of $d'(v)$ (is closest to s using only nodes in S).

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x): u \in S} (d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V - S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

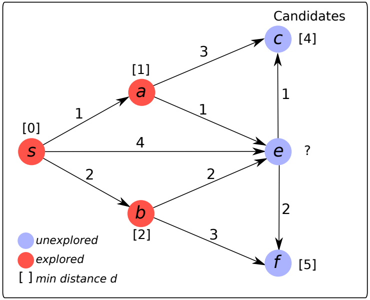

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
1: $S = \{s\}$ and $d(s) = 0$
2: while $S \neq V$ do
3:     for every node $x \in V - S$ do
4:         Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
5:     Set $v = \arg\min_{x \in V-S} d'(x)$
6:     Add $v$ to $S$ and set $d(v) = d'(v)$
```

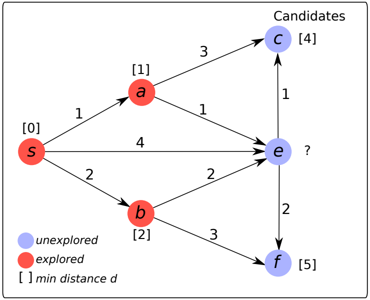

- How do we parse $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$?

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

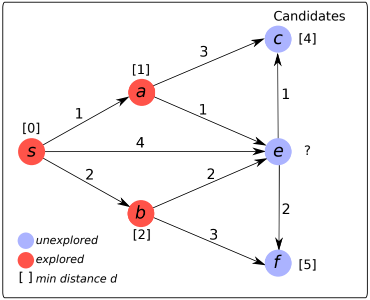

- How do we parse $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$?
  - The algorithm is examining a particular (unexplored) node x in $V - S$.

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM$(G, l, s)$
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V - S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

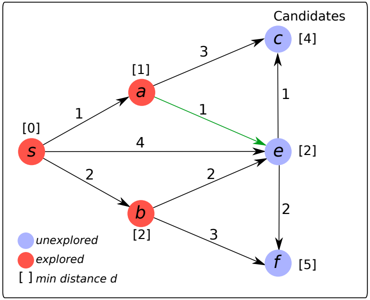

- How do we parse $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$?
    - The algorithm is examining a particular (unexplored) node x in $V - S$.
    - Argument of min runs over all edges of the type (u, x), where u is in S (i.e., u is explored).

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

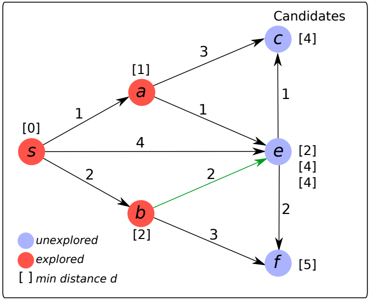

- How do we parse $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$?
  - The algorithm is examining a particular (unexplored) node x in $V - S$.
  - Argument of min runs over all edges of the type (u, x), where u is in S (i.e., u is explored).
  - For each such edge, we compute the length of the shortest path from s to x via u, which is $d(u) + l(u, x)$.

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V - S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

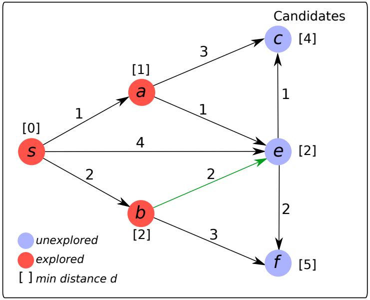

- How do we parse $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$?
    - The algorithm is examining a particular (unexplored) node x in $V - S$.
    - Argument of min runs over all edges of the type (u, x), where u is in S (i.e., u is explored).
    - For each such edge, we compute the length of the shortest path from s to x via u, which is $d(u) + l(u, x)$.
    - We store the smallest of these values in $d'(x)$.

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM$(G, l, s)$
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x): u \in S}(d(u) + l(u, x))$
    Set $v = \arg\min_{x \in V - S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

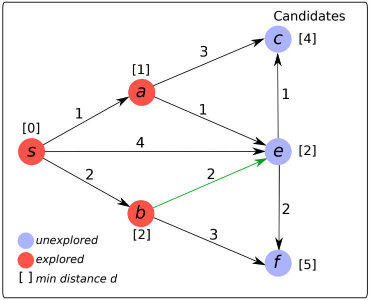

- How do we parse $v = \arg\min_{x \in V - S} d'(x)$?

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM$(G, l, s)$
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

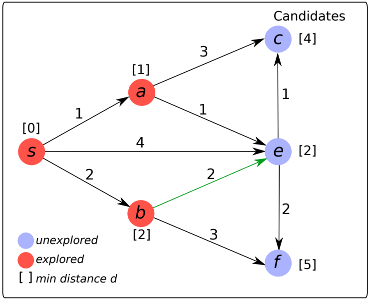

- How do we parse $v = \arg\min_{x \in V-S} d'(x)$?
    - Run over all (unexplored) nodes x in $V - S$.

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM$(G, l, s)$
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x): u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

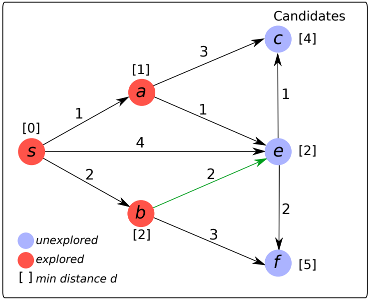

- How do we parse $v = \arg\min_{x \in V-S} d'(x)$?
  - Run over all (unexplored) nodes x in $V - S$.
  - Examine the d' values for these nodes.

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM$(G, l, s)$
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

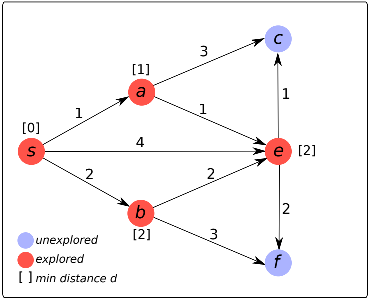

- How do we parse $v = \arg\min_{x \in V-S} d'(x)$?
  - Run over all (unexplored) nodes x in $V - S$.
  - Examine the d' values for these nodes.
  - Return the *argument* (i.e., the node) that has the smallest value of $d'(x)$.

Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

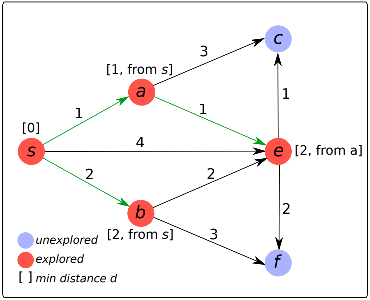

- How do we parse $v = \arg\min_{x \in V-S} d'(x)$?
  - Run over all (unexplored) nodes x in $V - S$.
  - Examine the d' values for these nodes.
  - Return the *argument* (i.e., the node) that has the smallest value of $d'(x)$.
- To compute the shortest paths: when adding a node v to S, store the predecessor u that minimises $d'(v)$.

Proof of Correctness

- Let P_u be the path computed by the algorithm for a node u.
- Claim: P_u is the shortest path from s to u.
- Prove by induction on the size of S.

Proof of Correctness

- Let P_u be the path computed by the algorithm for a node u.
- Claim: P_u is the shortest path from s to u.
- Prove by induction on the size of S.
  - Base case: $|S| = 1$. The only node in S is s.
  - Inductive hypothesis:

Proof of Correctness

- Let P_u be the path computed by the algorithm for a node u.
- Claim: P_u is the shortest path from s to u.
- Prove by induction on the size of S.
    - Base case: $|S| = 1$. The only node in S is s.
    - Inductive hypothesis: The algorithm has correctly computed P_u for all nodes $u \in S$.

Proof of Correctness

- Let P_u be the path computed by the algorithm for a node u.
- Claim: P_u is the shortest path from s to u.
- Prove by induction on the size of S.
  - Base case: $|S| = 1$. The only node in S is s.
  - Inductive hypothesis: The algorithm has correctly computed P_u for all nodes $u \in S$.
  - Inductive step: we add the node v to S. Let u be the v's predecessor on the path P_v. Could there be a shorter path P from s to v? We must prove this cannot be the case.

Proof of Correctness

- Let P_u be the path computed by the algorithm for a node u.
- Claim: P_u is the shortest path from s to u.
- Prove by induction on the size of S.
  - Base case: $|S| = 1$. The only node in S is s.
  - Inductive hypothesis: The algorithm has correctly computed P_u for all nodes $u \in S$.
  - Inductive step: we add the node v to S. Let u be the v's predecessor on the path P_v. Could there be a shorter path P from s to v? We must prove this cannot be the case.

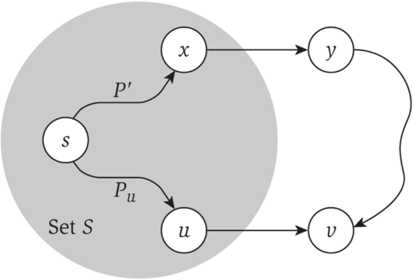

The alternate s–v path P through x and y is already too long by the time it has left the set S.

Figure 4.8 The shortest path P and an alternate s–v path P through the node y.

Comments about Dijkstra's Algorithm

- Algorithm cannot handle negative edge lengths. We will discuss the Bellman-Ford algorithm in a few weeks.
- Union of shortest paths output by Dijkstra's algorithm forms a tree. Why?

Comments about Dijkstra's Algorithm

- Algorithm cannot handle negative edge lengths. We will discuss the Bellman-Ford algorithm in a few weeks.
- Union of shortest paths output by Dijkstra's algorithm forms a tree. Why?
- Union of shortest paths from a fixed source s forms a tree; paths not necessarily computed by Dijkstra's algorithm.

Running time of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

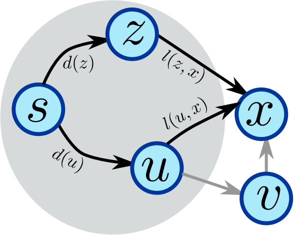

- How many iterations are there of the while loop?

Running time of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
S = {s} and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

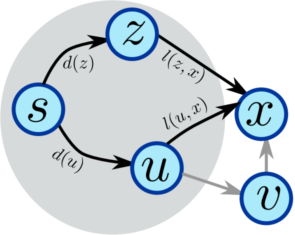

- How many iterations are there of the while loop? $n - 1$.

Running time of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

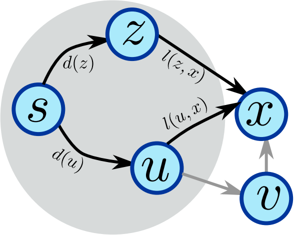

- How many iterations are there of the while loop? $n - 1$.
- In each iteration, for each node $x \in V - S$, compute

$$d'(x) = \min_{(u,x), u \in S}(d(u) + l(u,x))$$

Running time of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V - S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

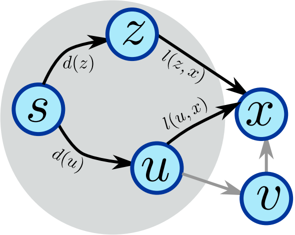

- How many iterations are there of the while loop? $n - 1$.
- In each iteration, for each node $x \in V - S$, compute

$$d'(x) = \min_{(u,x), u \in S}(d(u) + l(u,x))$$

- Running time per iteration is

Running time of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

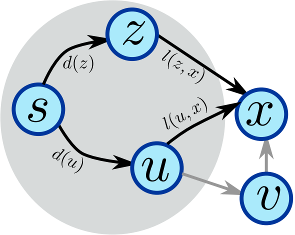

- How many iterations are there of the while loop? $n - 1$.
- In each iteration, for each node $x \in V - S$, compute

$$d'(x) = \min_{(u,x), u \in S}(d(u) + l(u,x))$$

- Running time per iteration is $O(m)$, since the algorithm processes each edge (u,x) in the graph exactly once (when computing $d'(x)$).
- The overall running time is $O(nm)$.

A Faster implementation of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

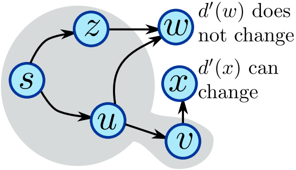

A Faster implementation of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u, x))$
    Set $v = \arg\min_{x \in V - S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

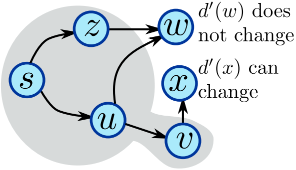

- Observation: If we add v to S, $d'(x)$ changes only if (v, x) is an edge in G.

A Faster implementation of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

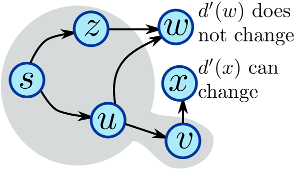

- Observation: If we add v to S, $d'(x)$ changes only if (v, x) is an edge in G.
- Idea: For each node $x \in V - S$, store the current value of $d'(x)$. Upon adding a node v to S, update $d'()$ only for neighbours of v.

A Faster implementation of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V - S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

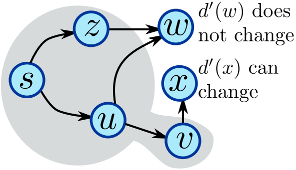

- Observation: If we add v to S, $d'(x)$ changes only if (v, x) is an edge in G.
- Idea: For each node $x \in V - S$, store the current value of $d'(x)$. Upon adding a node v to S, update $d'()$ only for neighbours of v.
- How do we efficiently compute $v = \arg\min_{x \in V - S} d'(x)$?

A Faster implementation of Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
$S = \{s\}$ and $d(s) = 0$
while $S \neq V$ do
    for every node $x \in V - S$ do
        Set $d'(x) = \min_{(u,x):u \in S}(d(u) + l(u,x))$
    Set $v = \arg\min_{x \in V-S} d'(x)$
    Add $v$ to $S$ and set $d(v) = d'(v)$
```

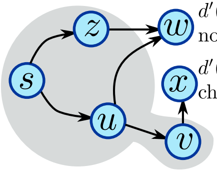

$d'(w)$ does not change

$d'(x)$ can change

- Observation: If we add v to S, $d'(x)$ changes only if (v, x) is an edge in G.
- Idea: For each node $x \in V - S$, store the current value of $d'(x)$. Upon adding a node v to S, update $d'()$ only for neighbours of v.
- How do we efficiently compute $v = \arg\min_{x \in V-S} d'(x)$?
- Use a priority queue!

Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
1: INSERT($Q, s, 0$).
2: while $S \neq V$ do
3:     $(v, d'(v))$ = EXTRACTMIN($Q$)
4:     Add $v$ to $S$ and set $d(v) = d'(v)$
5:     for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
6:         if $d(v) + l(v, x) < d'(x)$ then
7:             $d'(x) = d(v) + l(v, x)$
8:             CHANGEKEY($Q, x, d'(x)$)
```

- For each node $x \in V - S$, store the pair $(x, d'(x))$ in a priority queue Q with $d'(x)$ as the key.
- Determine the next node v to add to S using EXTRACTMIN (line 3).
- After adding v to S, for each node $x \in V - S$ such that there is an edge from v to x, check if $d'(x)$ should be updated, i.e., if there is a shortest path from s to x via v (lines 5–8).
- In line 8, if x is not in Q, simply insert it.

Running Time of Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
INSERT($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v))$ = EXTRACTMIN($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            CHANGEKEY($Q, x, d'(x)$)
```

- How many times does the algorithm invoke EXTRACTMIN?

Running Time of Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
INSERT($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v)) =$ EXTRACTMIN($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            CHANGEKEY($Q, x, d'(x)$)
```

- How many times does the algorithm invoke EXTRACTMIN? $n - 1$ times.

Running Time of Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
Insert($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v)) = $ ExtractMin($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            ChangeKey($Q, x, d'(x)$)
```

- How many times does the algorithm invoke ExtractMin? $n - 1$ times.
- For every node v, what is the running time of step 5?

Running Time of Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
INSERT($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v)) = $ EXTRACTMIN($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            CHANGEKEY($Q, x, d'(x)$)
```

- How many times does the algorithm invoke ExtractMin? $n - 1$ times.
- For every node v, what is the running time of step 5? $O(d_v)$, the number of *outgoing* neighbours of v.

Running Time of Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
INSERT($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v))$ = EXTRACTMIN($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            CHANGEKEY($Q, x, d'(x)$)
```

- How many times does the algorithm invoke EXTRACTMIN? $n - 1$ times.
- For every node v, what is the running time of step 5? $O(d_v)$, the number of *outgoing* neighbours of v.
- What is the total running time of step 5?

Running Time of Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
INSERT($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v))$ = EXTRACTMIN($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            CHANGEKEY($Q, x, d'(x)$)
```

- How many times does the algorithm invoke EXTRACTMIN? $n - 1$ times.
- For every node v, what is the running time of step 5? $O(d_v)$, the number of *outgoing* neighbours of v.
- What is the total running time of step 5? $\sum_{v \in V} O(d_v) = O(m)$.

Running Time of Faster Dijkstra's Algorithm

```
Dijkstra's Algorithm($G, l, s$)
Insert($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v)) = $ ExtractMin($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            ChangeKey($Q, x, d'(x)$)
```

- How many times does the algorithm invoke ExtractMin? $n - 1$ times.
- For every node v, what is the running time of step 5? $O(d_v)$, the number of *outgoing* neighbours of v.
- What is the total running time of step 5? $\sum_{v \in V} O(d_v) = O(m)$.
- How many times does the algorithm invoke ChangeKey?

Running Time of Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
1: INSERT($Q, s, 0$).
2: while $S \neq V$ do
3:     $(v, d'(v)) =$ EXTRACTMIN($Q$)
4:     Add $v$ to $S$ and set $d(v) = d'(v)$
5:     for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
6:         if $d(v) + l(v, x) < d'(x)$ then
7:             $d'(x) = d(v) + l(v, x)$
8:             CHANGEKEY($Q, x, d'(x)$)
```

- How many times does the algorithm invoke ExtractMin? $n - 1$ times.
- For every node v, what is the running time of step 5? $O(d_v)$, the number of *outgoing* neighbours of v.
- What is the total running time of step 5? $\sum_{v \in V} O(d_v) = O(m)$.
- How many times does the algorithm invoke ChangeKey? At most m times.

Running Time of Faster Dijkstra's Algorithm

```
Dijkstra's Algorithm($G, l, s$)
Insert($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v)) =$ ExtractMin($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            ChangeKey($Q, x, d'(x)$)
```

- How many times does the algorithm invoke ExtractMin? $n - 1$ times.
- For every node v, what is the running time of step 5? $O(d_v)$, the number of *outgoing* neighbours of v.
- What is the total running time of step 5? $\sum_{v \in V} O(d_v) = O(m)$.
- How many times does the algorithm invoke ChangeKey? At most m times.
- What is total running time of the algorithm?

Running Time of Faster Dijkstra's Algorithm

```
Dijkstra's Algorithm($G, l, s$)
Insert($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v))$ = ExtractMin($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            ChangeKey($Q, x, d'(x)$)
```

- How many times does the algorithm invoke ExtractMin? $n - 1$ times.
- For every node v, what is the running time of step 5? $O(d_v)$, the number of *outgoing* neighbours of v.
- What is the total running time of step 5? $\sum_{v \in V} O(d_v) = O(m)$.
- How many times does the algorithm invoke ChangeKey? At most m times.
- What is total running time of the algorithm? $O(m \log n)$.

Running Time of Faster Dijkstra's Algorithm

```
DIJKSTRA'S ALGORITHM($G, l, s$)
INSERT($Q, s, 0$).
while $S \neq V$ do
    $(v, d'(v))$ = EXTRACTMIN($Q$)
    Add $v$ to $S$ and set $d(v) = d'(v)$
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $d(v) + l(v, x) < d'(x)$ then
            $d'(x) = d(v) + l(v, x)$
            CHANGEKEY($Q, x, d'(x)$)
```

- How many times does the algorithm invoke EXTRACTMIN? $n - 1$ times.
- For every node v, what is the running time of step 5? $O(d_v)$, the number of *outgoing* neighbours of v.
- What is the total running time of step 5? $\sum_{v \in V} O(d_v) = O(m)$.
- How many times does the algorithm invoke CHANGEKEY? At most m times.
- What is total running time of the algorithm? $O(m \log n)$.
- State of the art: Fibonacci heaps achieve a running time of $O(m)$ for all CHANGEKEY operations, for a running time of $O(n \log n + m)$.

Network Design

- Connect a set of nodes using a set of edges with certain properties.
- Input is usually a graph and the desired network (the output) should use subset of edges in the graph.
- Example: connect all nodes using a cycle of shortest total length.

Network Design

- Connect a set of nodes using a set of edges with certain properties.
- Input is usually a graph and the desired network (the output) should use subset of edges in the graph.
- Example: connect all nodes using a cycle of shortest total length. This problem is the NP-complete traveling salesman problem.

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

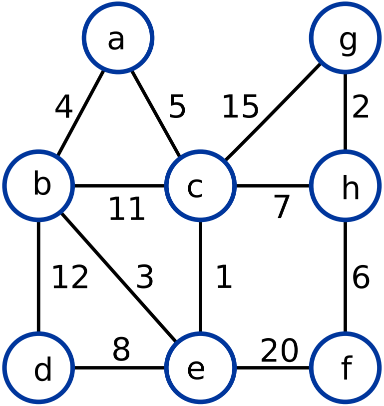

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

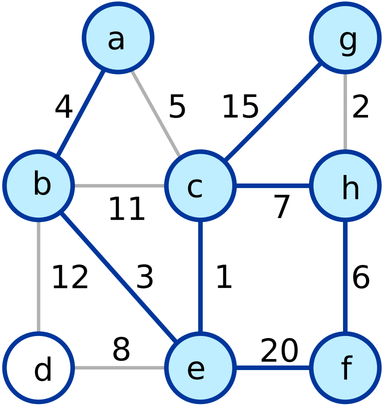

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

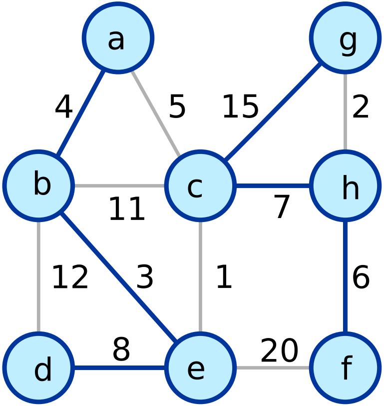

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

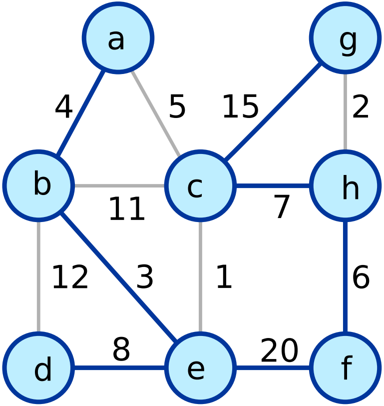

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

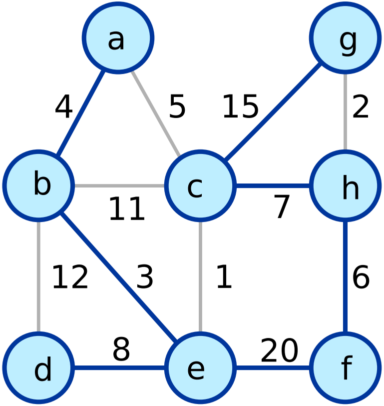

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

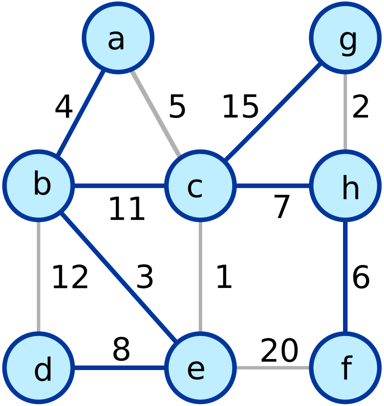

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

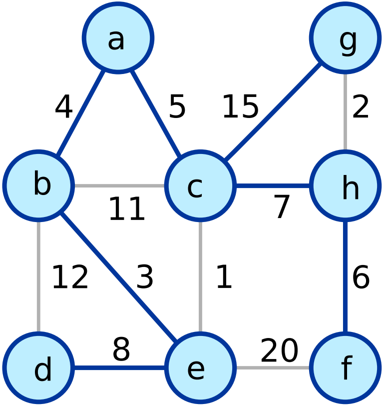

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

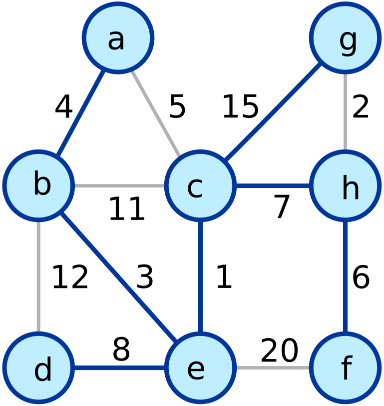

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

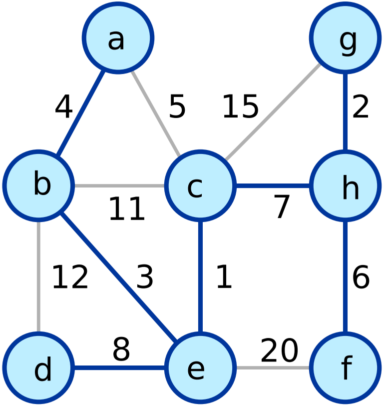

Minimum Spanning Tree (MST)

- Given an undirected graph $G(V, E)$ with a cost $c(e) > 0$ associated with each edge $e \in E$.
- Find a subset T of edges such that the graph (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

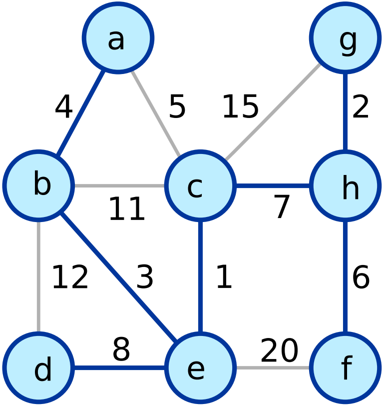

Minimum Spanning Tree

INSTANCE: An undirected graph $G(V, E)$ and a function $c : E \rightarrow \mathbb{R}^+$

SOLUTION: A set $T \subseteq E$ of edges such that (V, T) is connected and the cost $\sum_{e \in T} c(e)$ is as small as possible.

- Claim: If T is a minimum-cost solution to this problem then (V, T) is a tree.
- A subset T of E is a *spanning tree* of G if (V, T) is a tree.

Greedy Algorithm for the MST Problem

- Template: process edges in some order. Add an edge to T if tree property is not violated.

Greedy Algorithm for the MST Problem

- Template: process edges in some order. Add an edge to T if tree property is not violated.

  Increasing cost order Process edges in increasing order of cost. Discard an edge if it creates a cycle.

  Dijkstra-like Start from a node s and grow T outward from s: add the node that can be attached most cheaply to current tree.

  Decreasing cost order Delete edges in order of decreasing cost as long as graph remains connected.

Greedy Algorithm for the MST Problem

- Template: process edges in some order. Add an edge to T if tree property is not violated.

  Increasing cost order Process edges in increasing order of cost. Discard an edge if it creates a cycle.

  Dijkstra-like Start from a node s and grow T outward from s: add the node that can be attached most cheaply to current tree.

  Decreasing cost order Delete edges in order of decreasing cost as long as graph remains connected.

- Which of these algorithms works?

Greedy Algorithm for the MST Problem

- Template: process edges in some order. Add an edge to T if tree property is not violated.

  Increasing cost order: Process edges in increasing order of cost. Discard an edge if it creates a cycle. Kruskal's algorithm

  Dijkstra-like: Start from a node s and grow T outward from s: add the node that can be attached most cheaply to current tree. Prim's algorithm

  Decreasing cost order: Delete edges in order of decreasing cost as long as graph remains connected. Reverse-Delete algorithm

- Which of these algorithms works? All of them!

Greedy Algorithm for the MST Problem

- Template: process edges in some order. Add an edge to T if tree property is not violated.
  - Increasing cost order: Process edges in increasing order of cost. Discard an edge if it creates a cycle. Kruskal's algorithm
  - Dijkstra-like: Start from a node s and grow T outward from s: add the node that can be attached most cheaply to current tree. Prim's algorithm
  - Decreasing cost order: Delete edges in order of decreasing cost as long as graph remains connected. Reverse-Delete algorithm
- Which of these algorithms works? All of them!
- Simplifying assumption: all edge costs are distinct.

Characterising MSTs

- Does the edge of smallest cost belong to an MST?

Characterising MSTs

- Does the edge of smallest cost belong to an MST? Yes. Why?

Characterising MSTs

- Does the edge of smallest cost belong to an MST? Yes. Why?
    - Wrong proof: because Kruskal's algorithm adds it. We have not yet proved correctness of Kruskal's algorithm!
    - Correct proof: will work it out soon.
- Which edges must belong to an MST?

Characterising MSTs

- Does the edge of smallest cost belong to an MST? Yes. Why?
  - Wrong proof: because Kruskal's algorithm adds it. We have not yet proved correctness of Kruskal's algorithm!
  - Correct proof: will work it out soon.
- Which edges must belong to an MST?
  - What happens when we delete an edge from an MST?
  - MST breaks up into sub-trees.
  - Which edge should we add to join them?

Characterising MSTs

- Does the edge of smallest cost belong to an MST? Yes. Why?
    - Wrong proof: because Kruskal's algorithm adds it. We have not yet proved correctness of Kruskal's algorithm!
    - Correct proof: will work it out soon.
- Which edges must belong to an MST?
    - What happens when we delete an edge from an MST?
    - MST breaks up into sub-trees.
    - Which edge should we add to join them?
- Which edges cannot belong to an MST?

Characterising MSTs

- Does the edge of smallest cost belong to an MST? Yes. Why?
  - Wrong proof: because Kruskal's algorithm adds it. We have not yet proved correctness of Kruskal's algorithm!
  - Correct proof: will work it out soon.
- Which edges must belong to an MST?
  - What happens when we delete an edge from an MST?
  - MST breaks up into sub-trees.
  - Which edge should we add to join them?
- Which edges cannot belong to an MST?
  - What happens when we add an edge to an MST?
  - We obtain a cycle.
  - Which edge in the cycle can we be sure does not belong to an MST?

Graph Cuts

- A *cut* in a graph $G(V, E)$ is a set of edges whose removal disconnects the graph (into two or more connected components).
- Every set $S \subset V$ (S cannot be empty or the entire set V) has a corresponding cut: $cut(S)$ is the set of edges (v, w) such that $v \in S$ and $w \in V - S$.

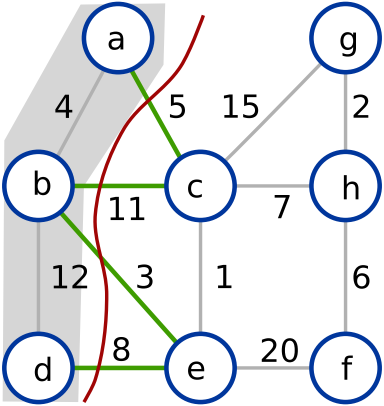

Graph Cuts

- A *cut* in a graph $G(V, E)$ is a set of edges whose removal disconnects the graph (into two or more connected components).
- Every set $S \subset V$ (S cannot be empty or the entire set V) has a corresponding cut: $cut(S)$ is the set of edges (v, w) such that $v \in S$ and $w \in V - S$.
- cut(S) is a "cut" because deleting the edges in cut(S) disconnects S from $V - S$.

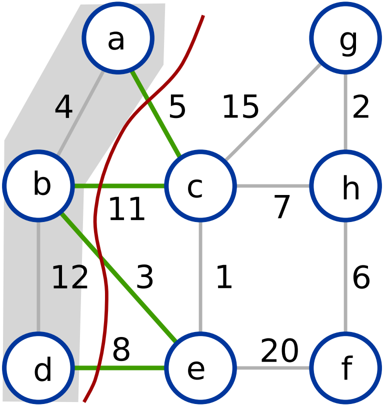

Graph Cuts

- A *cut* in a graph $G(V, E)$ is a set of edges whose removal disconnects the graph (into two or more connected components).
- Every set $S \subset V$ (S cannot be empty or the entire set V) has a corresponding cut: $cut(S)$ is the set of edges (v, w) such that $v \in S$ and $w \in V - S$.
- cut(S) is a "cut" because deleting the edges in cut(S) disconnects S from $V - S$.

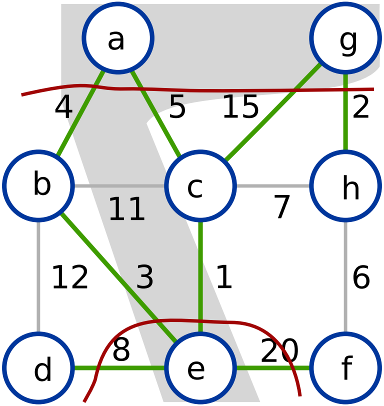

Graph Cuts

- A *cut* in a graph $G(V, E)$ is a set of edges whose removal disconnects the graph (into two or more connected components).
- Every set $S \subset V$ (S cannot be empty or the entire set V) has a corresponding cut: $cut(S)$ is the set of edges (v, w) such that $v \in S$ and $w \in V - S$.
- cut(S) is a "cut" because deleting the edges in cut(S) disconnects S from $V - S$.

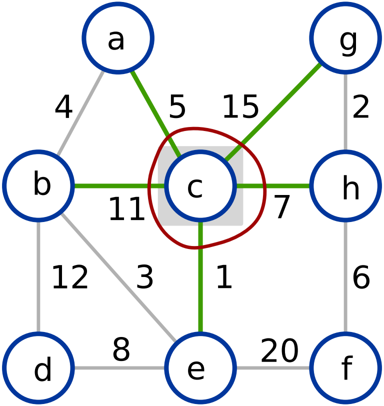

Cut Property

- When is it safe to include an edge in an MST?

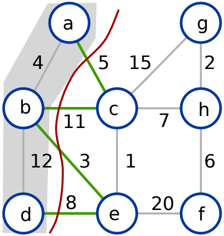

Cut Property

- When is it safe to include an edge in an MST?
- Claim: For every $S \subset V, S \neq \emptyset$, every MST contains the cheapest edge in cut(S).

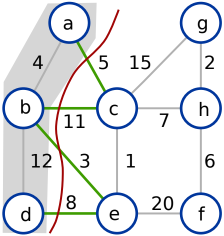

Cut Property

- When is it safe to include an edge in an MST?
- Claim: For every $S \subset V, S \neq \emptyset$, every MST contains the cheapest edge in cut(S).
- Proof by contradiction using exchange argument.

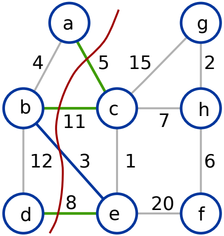

Cut Property

- When is it safe to include an edge in an MST?
- Claim: For every $S \subset V, S \neq \emptyset$, every MST contains the cheapest edge in cut(S).
- Proof by contradiction using exchange argument.
- How do you state the contradiction to the claim?

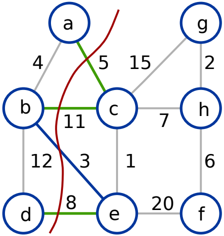

Cut Property

- When is it safe to include an edge in an MST?
- Claim: For every $S \subset V, S \neq \emptyset$, every MST contains the cheapest edge in cut(S).
- Proof by contradiction using exchange argument.
- How do you state the contradiction to the claim? There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).

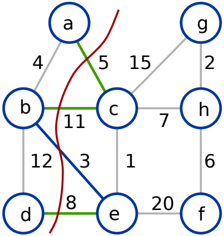

Cut Property

- When is it safe to include an edge in an MST?
- Claim: For every $S \subset V, S \neq \emptyset$, every MST contains the cheapest edge in cut(S).
- Proof by contradiction using exchange argument.
- How do you state the contradiction to the claim? There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).
    - Let $e = (u, v)$ be the cheapest edge in cut(S).

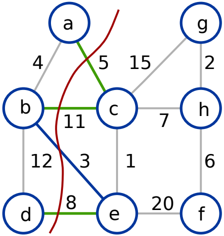

Cut Property

- When is it safe to include an edge in an MST?
- Claim: For every $S \subset V, S \neq \emptyset$, every MST contains the cheapest edge in cut(S).
- Proof by contradiction using exchange argument.
- How do you state the contradiction to the claim? There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).
    - Let $e = (u, v)$ be the cheapest edge in cut(S).
- Proof strategy: If T does not contain e, show that there is a tree with smaller cost than T that contains e.

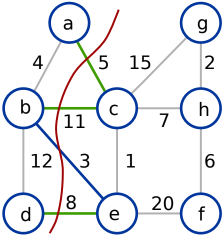

Proof of Cut Property

- There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).
- Proof strategy: If T does not contain e, show that there is a tree with smaller cost than T that contains e.

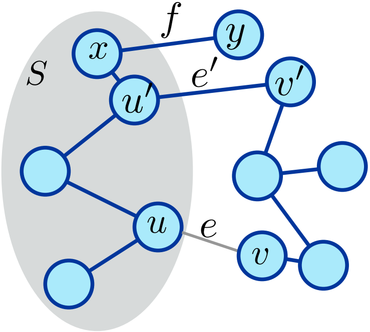

Proof of Cut Property

- There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).
- Proof strategy: If T does not contain e, show that there is a tree with smaller cost than T that contains e.

- Wrong proof:
    - Since T is spanning, it must contain *some* edge, e.g., f, in cut(S).
    - $T - \{f\} \cup \{e\}$ has smaller cost than T but

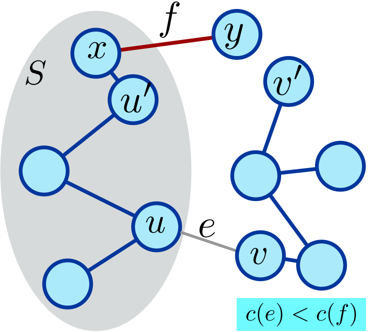

Proof of Cut Property

- There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).
- Proof strategy: If T does not contain e, show that there is a tree with smaller cost than T that contains e.

- Wrong proof:
  - Since T is spanning, it must contain *some* edge, e.g., f, in cut(S).
  - $T - \{f\} \cup \{e\}$ has smaller cost than T but may not be a spanning tree.

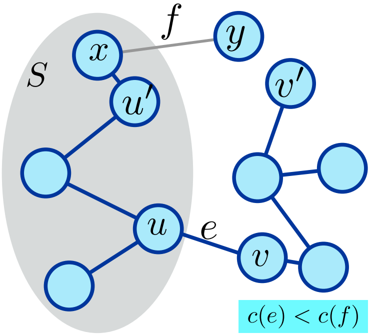

Proof of Cut Property

- There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).
- Proof strategy: If T does not contain e, show that there is a tree with smaller cost than T that contains e.

- Wrong proof:
  - Since T is spanning, it must contain *some* edge, e.g., f, in cut(S).
  - $T - \{f\} \cup \{e\}$ has smaller cost than T but may not be a spanning tree.
- Correct proof:
  - Add e to T forming a cycle.

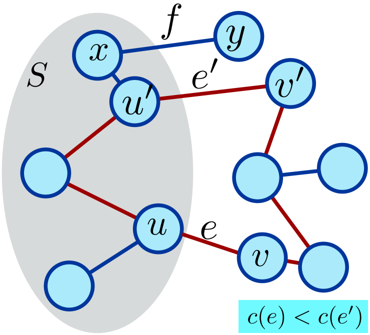

Proof of Cut Property

- There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).
- Proof strategy: If T does not contain e, show that there is a tree with smaller cost than T that contains e.

- Wrong proof:
  - Since T is spanning, it must contain *some* edge, e.g., f, in cut(S).
  - $T - \{f\} \cup \{e\}$ has smaller cost than T but may not be a spanning tree.
- Correct proof:
  - Add e to T forming a cycle.
  - This *cycle* must contain an edge e' in cut(S).

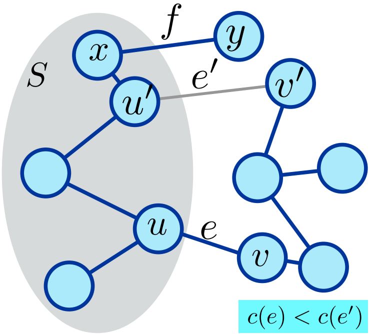

Proof of Cut Property

- There is a set $S \subset V$ and an MST T such that T does not contain the cheapest edge in cut(S).
- Proof strategy: If T does not contain e, show that there is a tree with smaller cost than T that contains e.

- Wrong proof:
  - Since T is spanning, it must contain *some* edge, e.g., f, in cut(S).
  - $T - \{f\} \cup \{e\}$ has smaller cost than T but may not be a spanning tree.
- Correct proof:
  - Add e to T forming a cycle.
  - This *cycle* must contain an edge e' in cut(S).
  - $T - \{e'\} \cup \{e\}$ has smaller cost than T and is a spanning tree.

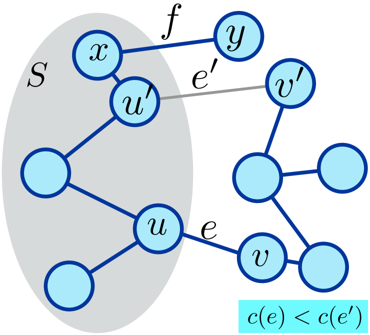

Prim's Algorithm

- Maintain a tree (S, T), i.e. a set of nodes and a set of edges, which we will show will always be a tree.
- Start with an arbitrary node $s \in S$.

Prim's Algorithm

- Maintain a tree (S, T), i.e. a set of nodes and a set of edges, which we will show will always be a tree.
- Start with an arbitrary node $s \in S$.

```
PRIM'S ALGORITHM($G, c, s$)
$S = \{s\}$ and $T = \emptyset$
while $S \neq V$ do
    Compute $(u, v) = \arg\min_{(u,v): u \in S, v \in V - S} c(u, v)$
    Add the node $v$ to $S$ and add the edge $(u, v)$ to $T$.
```

Prim's Algorithm

- Maintain a tree (S, T), i.e. a set of nodes and a set of edges, which we will show will always be a tree.
- Start with an arbitrary node $s \in S$.

```
PRIM'S ALGORITHM($G, c, s$)
S = {s} and T = ∅
while $S \neq V$ do
    Compute $(u, v) = \arg\min_{(u,v): u \in S, v \in V - S} c(u, v)$
    Add the node $v$ to $S$ and add the edge $(u, v)$ to $T$.
```

- Note that

$$\arg \min_{(u,v), u \in S, v \in V - S} c(u, v) \equiv \arg \min_{(u,v) \in \text{cut}(S)} c(u, v).$$

- In other words, in each step Prim's algorithm computes and adds the cheapest edge in the current value of $\text{cut}(S)$.

Example of Prim's Algorithm

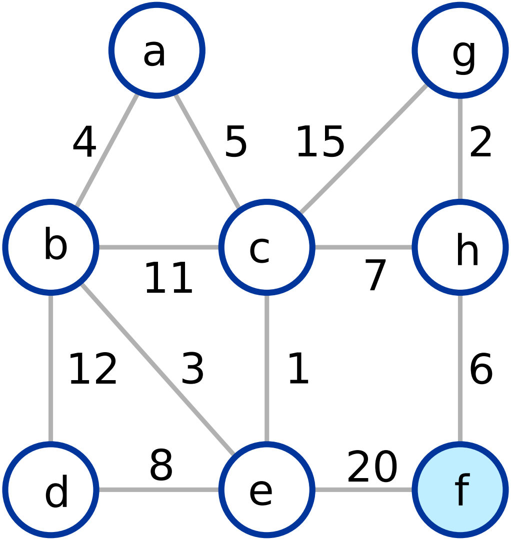

Example of Prim's Algorithm

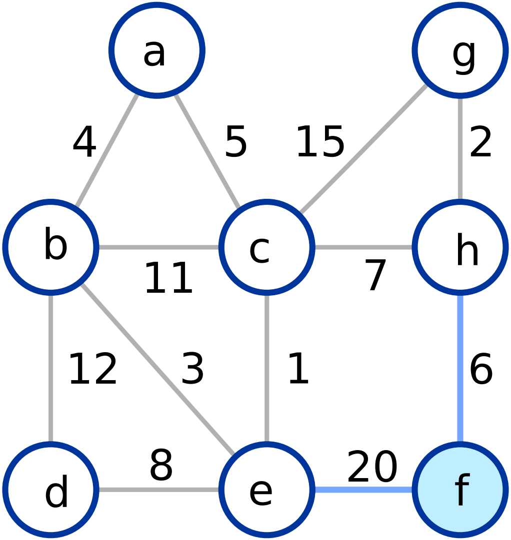

Example of Prim's Algorithm

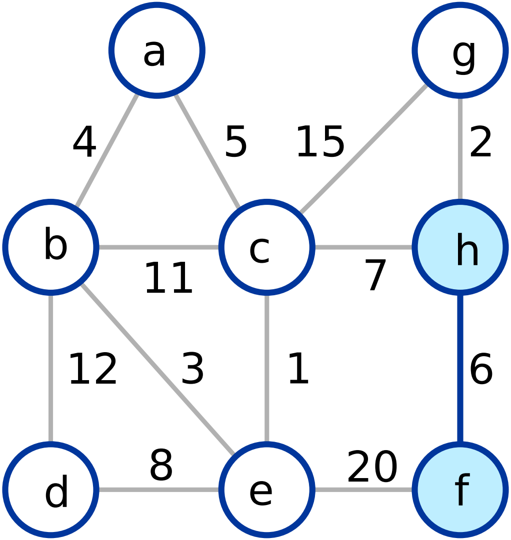

Example of Prim's Algorithm

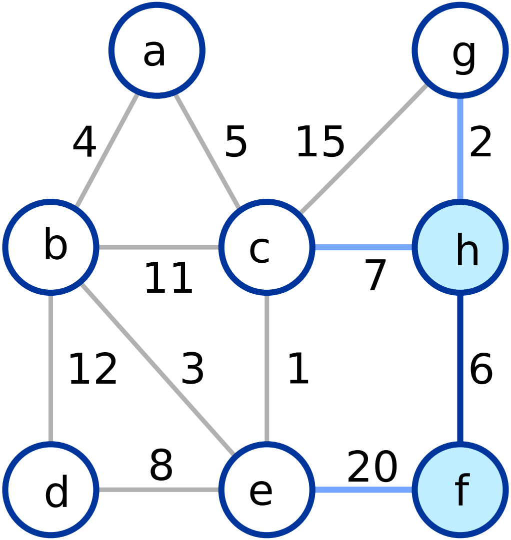

Example of Prim's Algorithm

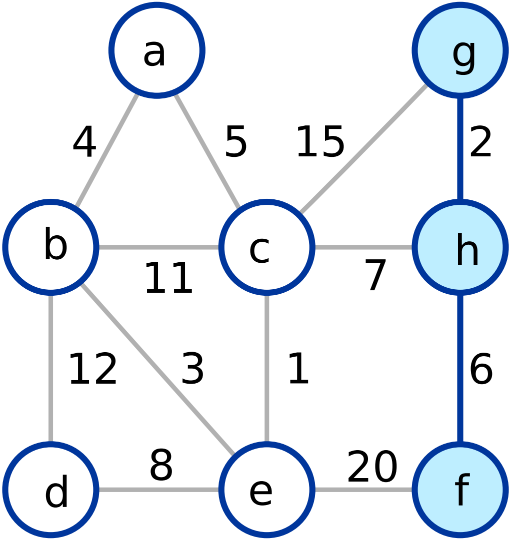

Example of Prim's Algorithm

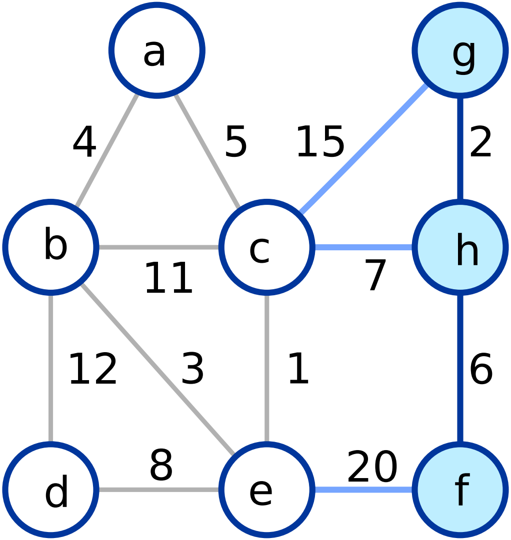

Example of Prim's Algorithm

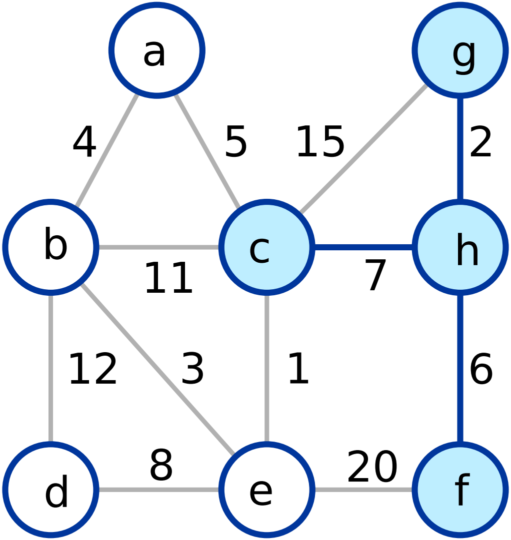

Example of Prim's Algorithm

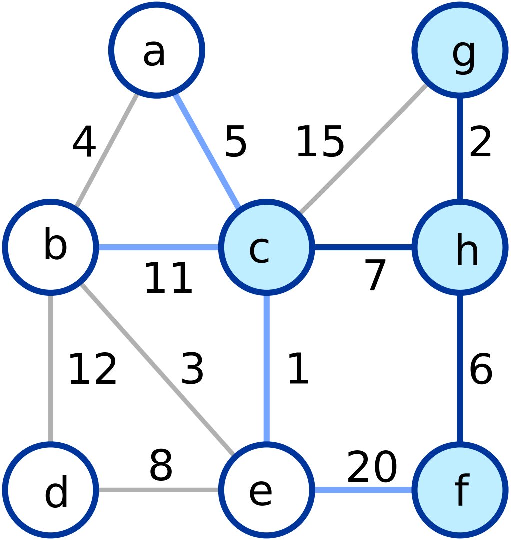

Example of Prim's Algorithm

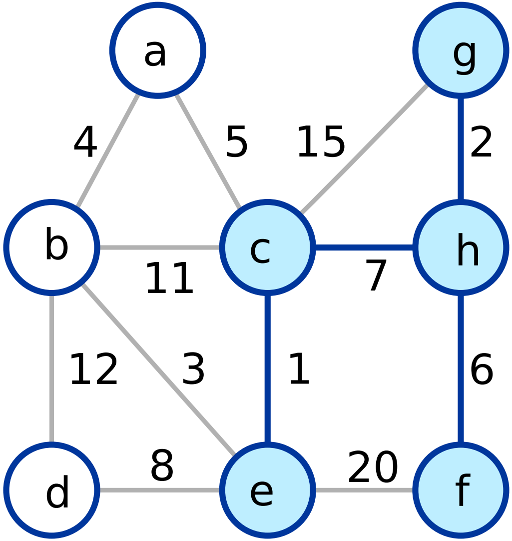

Example of Prim's Algorithm

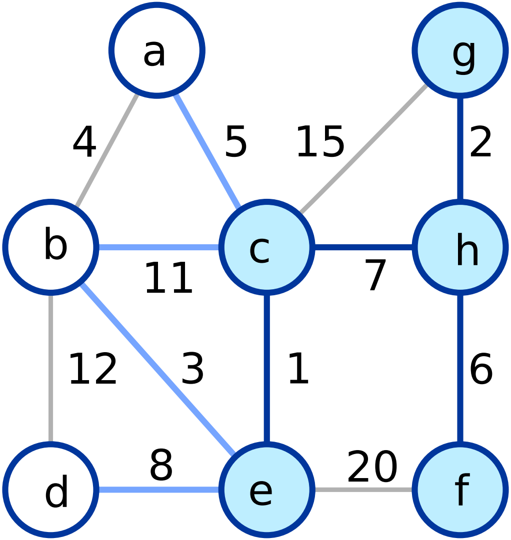

Example of Prim's Algorithm

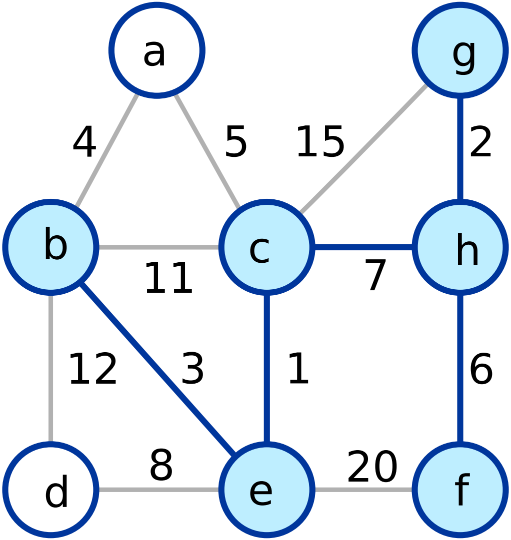

Example of Prim's Algorithm

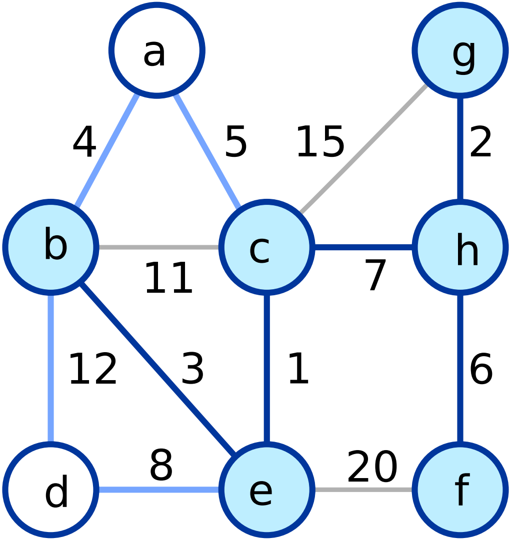

Example of Prim's Algorithm

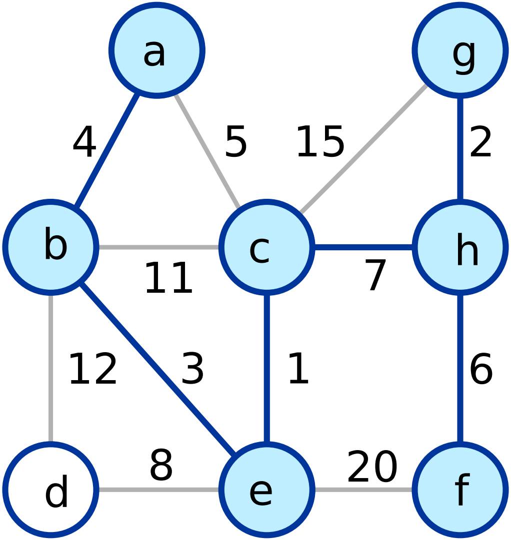

Example of Prim's Algorithm

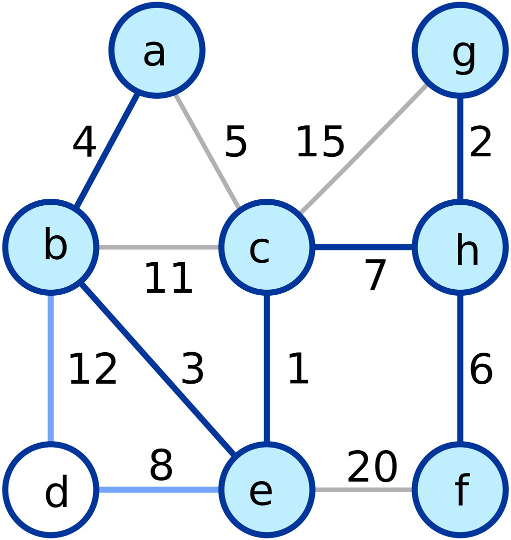

Example of Prim's Algorithm

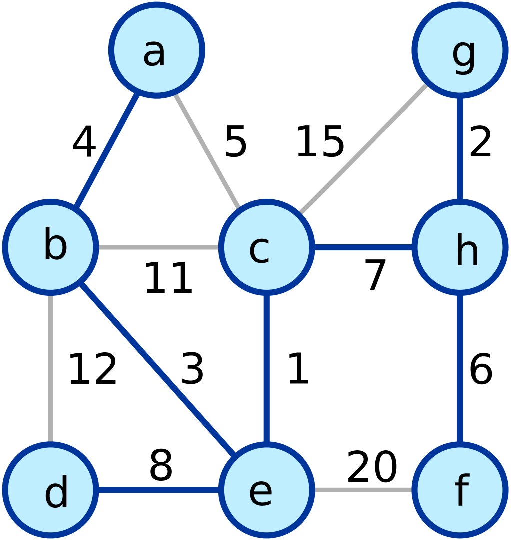

Optimality of Prim's Algorithm

```
PRIM'S ALGORITHM($G, c, s$)
$S = \{s\}$ and and $T = \emptyset$
while $S \neq V$ do
    Compute $(u, v) = \arg\min_{(u,v) \in \text{cut}(S)} c(u, v)$
    Add the node $v$ to $S$ and add the edge $(u, v)$ to $T$.
```

- Claim: Prim's algorithm outputs an MST.

Optimality of Prim's Algorithm

```
PRIM'S ALGORITHM($G, c, s$)
$S = \{s\}$ and and $T = \emptyset$
while $S \neq V$ do
    Compute $(u, v) = \arg\min_{(u,v) \in \mathrm{cut}(S)} c(u, v)$
    Add the node $v$ to $S$ and add the edge $(u, v)$ to $T$.
```

- Claim: Prim's algorithm outputs an MST.
  1. Prove that every edge inserted satisfies the cut property.
  2. Prove that the graph constructed is a spanning tree.

Optimality of Prim's Algorithm

```
PRIM'S ALGORITHM($G, c, s$)
$S = \{s\}$ and and $T = \emptyset$
while $S \neq V$ do
    Compute $(u, v) = \arg\min_{(u,v) \in \text{cut}(S)} c(u, v)$
    Add the node $v$ to $S$ and add the edge $(u, v)$ to $T$.
```

- Claim: Prim's algorithm outputs an MST.
  1. Prove that every edge inserted satisfies the cut property.
     - ⋆ By construction, in each iteration (u, v) is the cheapest edge in cut(S) for the current value of S.
  2. Prove that the graph constructed is a spanning tree.

Optimality of Prim's Algorithm

```
PRIM'S ALGORITHM($G, c, s$)
$S = \{s\}$ and and $T = \emptyset$
while $S \neq V$ do
    Compute $(u, v) = \arg\min_{(u,v) \in \mathrm{cut}(S)} c(u, v)$
    Add the node $v$ to $S$ and add the edge $(u, v)$ to $T$.
```

- Claim: Prim's algorithm outputs an MST.
  1. Prove that every edge inserted satisfies the cut property.
     - By construction, in each iteration (u, v) is the cheapest edge in $\mathrm{cut}(S)$ for the current value of S.
  2. Prove that the graph constructed is a spanning tree.
     - Why are there no cycles in (V, T)?

Optimality of Prim's Algorithm

```
PRIM'S ALGORITHM($G, c, s$)
$S = \{s\}$ and and $T = \emptyset$
while $S \neq V$ do
    Compute $(u, v) = \arg\min_{(u,v) \in \text{cut}(S)} c(u, v)$
    Add the node $v$ to $S$ and add the edge $(u, v)$ to $T$.
```

- Claim: Prim's algorithm outputs an MST.
  1. Prove that every edge inserted satisfies the cut property.
     - ★ By construction, in each iteration (u, v) is the cheapest edge in cut(S) for the current value of S.
  2. Prove that the graph constructed is a spanning tree.
     - ★ Why are there no cycles in (V, T)?
     - ★ Why is (V, T) connected?

Kruskal's Algorithm

- Start with an empty set T of edges.
- Process edges in E in increasing order of cost.
- Add the next edge e to T only if adding e does not create a cycle. Discard e if it creates a cycle.

Example of Kruskal's Algorithm

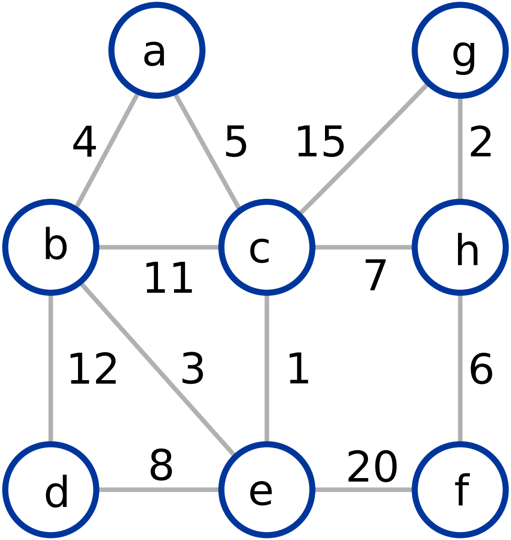

Example of Kruskal's Algorithm

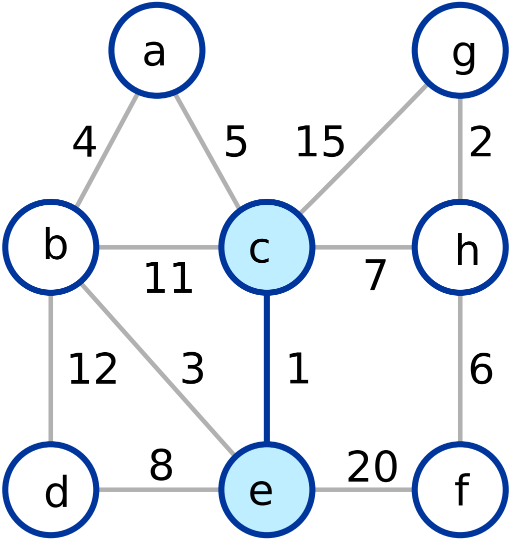

Example of Kruskal's Algorithm

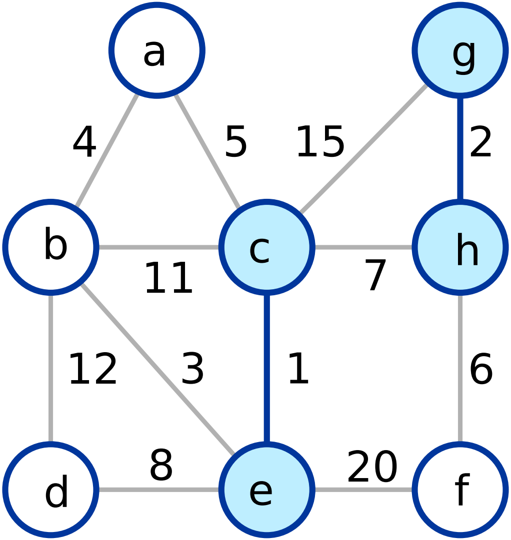

Example of Kruskal's Algorithm

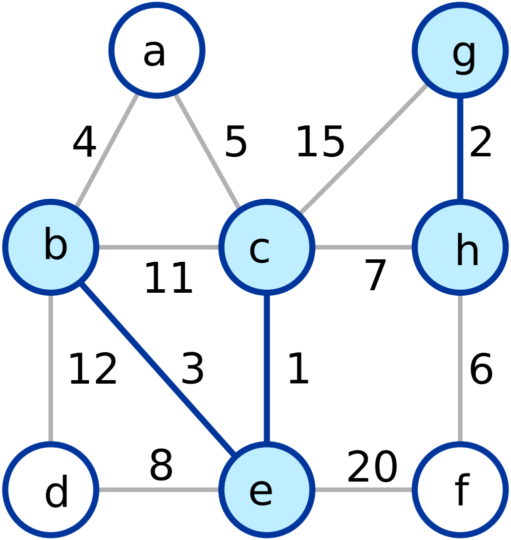

Example of Kruskal's Algorithm

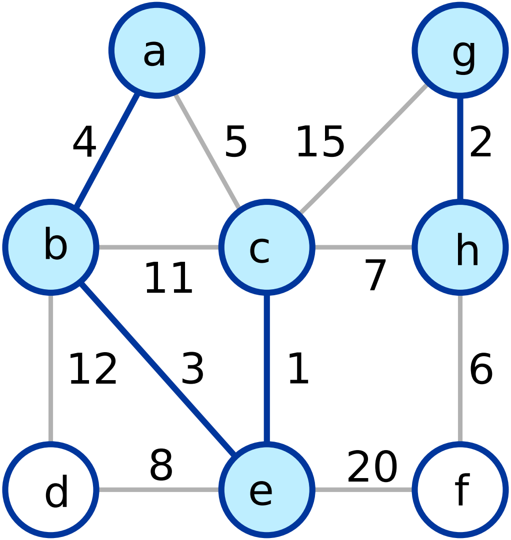

Example of Kruskal's Algorithm

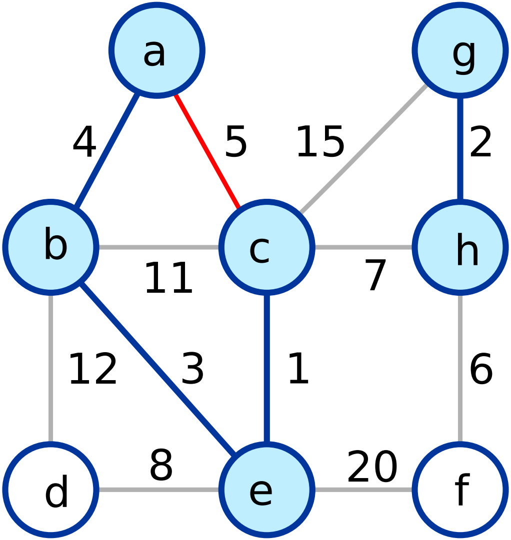

Example of Kruskal's Algorithm

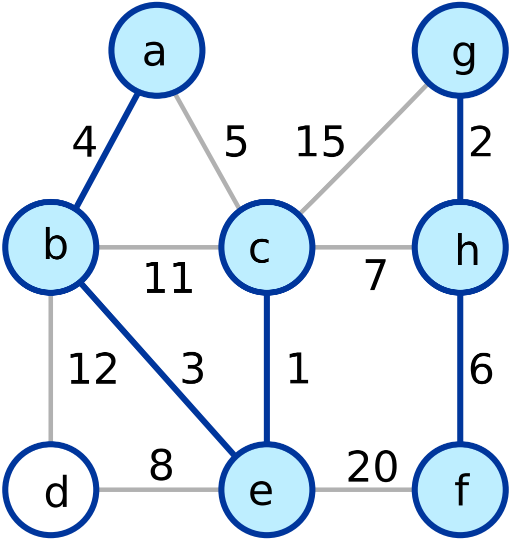

Example of Kruskal's Algorithm

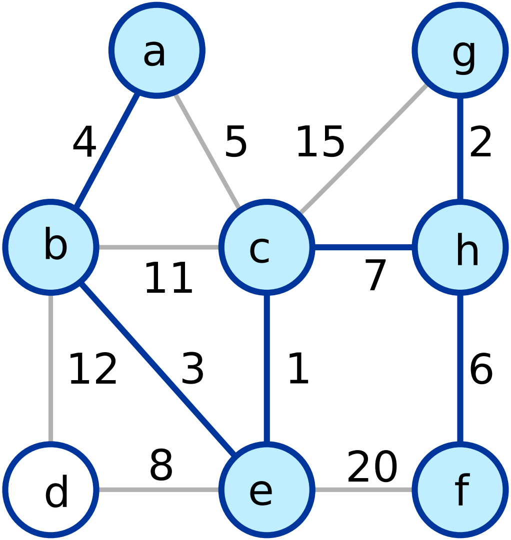

Example of Kruskal's Algorithm

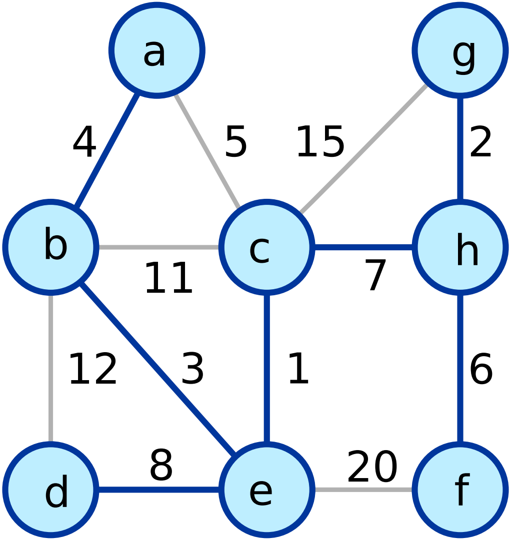

Example of Kruskal's Algorithm

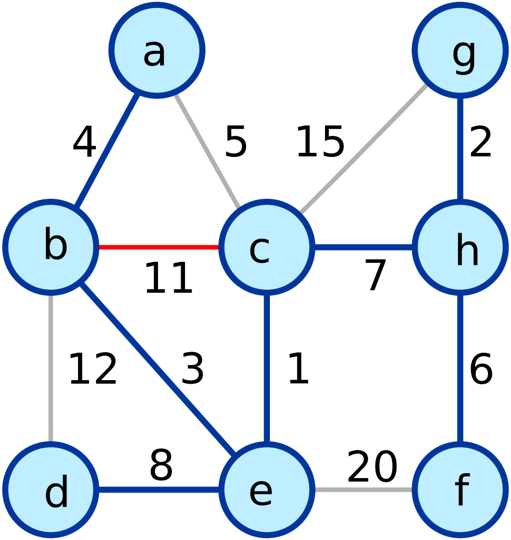

Example of Kruskal's Algorithm

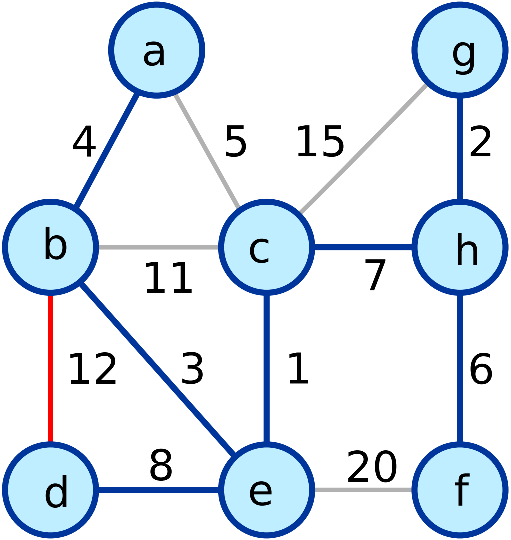

Example of Kruskal's Algorithm

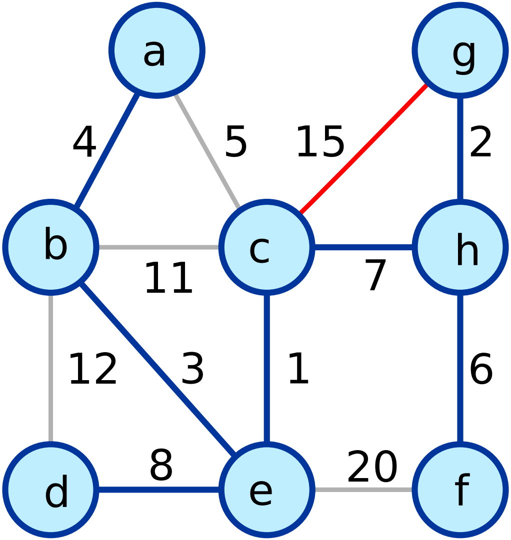

Example of Kruskal's Algorithm

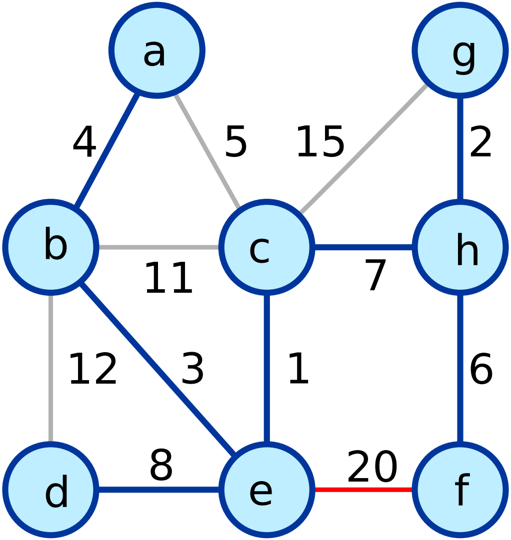

Example of Kruskal's Algorithm

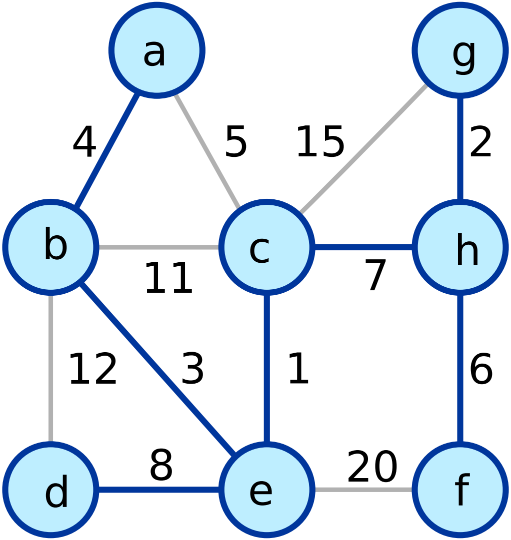

Optimality of Kruskal's Algorithm

- Kruskal's algorithm:
  - Start with an empty set T of edges.
  - Process edges in E in increasing order of cost.
  - Add the next edge e to T only if adding e does not create a cycle. Discard e if it creates a cycle.
- Note: at any iteration, T may contain several connected components and each node in V is in some component.
- Claim: Kruskal's algorithm outputs an MST.

Optimality of Kruskal's Algorithm

- Kruskal's algorithm:
  - Start with an empty set T of edges.
  - Process edges in E in increasing order of cost.
  - Add the next edge e to T only if adding e does not create a cycle. Discard e if it creates a cycle.
- Note: at any iteration, T may contain several connected components and each node in V is in some component.
- Claim: Kruskal's algorithm outputs an MST.
  1. For every edge e added, demonstrate the existence of a set $S \subset V$ (and $V - S$) such that e and S satisfy the cut property, i.e., e is the cheapest edge in cut(S).
  2. Prove that the algorithm computes a spanning tree.

Optimality of Kruskal's Algorithm

- Kruskal's algorithm:
  - Start with an empty set T of edges.
  - Process edges in E in increasing order of cost.
  - Add the next edge e to T only if adding e does not create a cycle. Discard e if it creates a cycle.
- Note: at any iteration, T may contain several connected components and each node in V is in some component.
- Claim: Kruskal's algorithm outputs an MST.
  1. For every edge e added, demonstrate the existence of a set $S \subset V$ (and $V - S$) such that e and S satisfy the cut property, i.e., e is the cheapest edge in cut(S).
     - If $e = (u, v)$, let S be the set of nodes connected to u in the current graph T.
  2. Prove that the algorithm computes a spanning tree.

Optimality of Kruskal's Algorithm

- Kruskal's algorithm:
  - Start with an empty set T of edges.
  - Process edges in E in increasing order of cost.
  - Add the next edge e to T only if adding e does not create a cycle. Discard e if it creates a cycle.
- Note: at any iteration, T may contain several connected components and each node in V is in some component.
- Claim: Kruskal's algorithm outputs an MST.
  1. For every edge e added, demonstrate the existence of a set $S \subset V$ (and $V - S$) such that e and S satisfy the cut property, i.e., e is the cheapest edge in cut(S).
     - If $e = (u, v)$, let S be the set of nodes connected to u in the current graph T.
     - Why is e the cheapest edge in cut(S)?
  2. Prove that the algorithm computes a spanning tree.

Optimality of Kruskal's Algorithm

- Kruskal's algorithm:
  - Start with an empty set T of edges.
  - Process edges in E in increasing order of cost.
  - Add the next edge e to T only if adding e does not create a cycle. Discard e if it creates a cycle.
- Note: at any iteration, T may contain several connected components and each node in V is in some component.
- Claim: Kruskal's algorithm outputs an MST.
  1. For every edge e added, demonstrate the existence of a set $S \subset V$ (and $V - S$) such that e and S satisfy the cut property, i.e., e is the cheapest edge in cut(S).
     - ★ If $e = (u, v)$, let S be the set of nodes connected to u in the current graph T.
     - ★ Why is e the cheapest edge in cut(S)?
  2. Prove that the algorithm computes a spanning tree.
     - ★ (V, T) contains no cycles by construction.

Optimality of Kruskal's Algorithm

- Kruskal's algorithm:
  - Start with an empty set T of edges.
  - Process edges in E in increasing order of cost.
  - Add the next edge e to T only if adding e does not create a cycle. Discard e if it creates a cycle.
- Note: at any iteration, T may contain several connected components and each node in V is in some component.
- Claim: Kruskal's algorithm outputs an MST.
  1. For every edge e added, demonstrate the existence of a set $S \subset V$ (and $V - S$) such that e and S satisfy the cut property, i.e., e is the cheapest edge in cut(S).
     - ⋆ If $e = (u, v)$, let S be the set of nodes connected to u in the current graph T.
     - ⋆ Why is e the cheapest edge in cut(S)?
  2. Prove that the algorithm computes a spanning tree.
     - ⋆ (V, T) contains no cycles by construction.
     - ⋆ If (V, T) is not connected, then exists a subset S of nodes not connected to $V - S$. What is the contradiction?

Cycle Property

- When can we be sure that an edge cannot be in *any* MST?

Cycle Property

- When can we be sure that an edge cannot be in *any* MST?
- Let C be any cycle in G and let $e = (v, w)$ be the most expensive edge in C.
- Claim: e does not belong to any MST of G.

Cycle Property

- When can we be sure that an edge cannot be in *any* MST?
- Let C be any cycle in G and let $e = (v, w)$ be the most expensive edge in C.
- Claim: e does not belong to any MST of G.
- Proof: exchange argument. If a supposed MST T contains e, show that there is a tree with smaller cost than T that does not contain e.

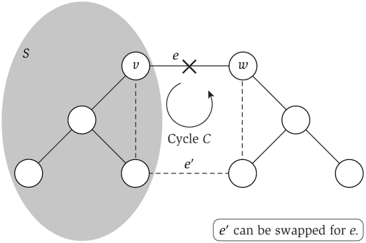

Figure 4.11 Swapping the edge e' for the edge e in the spanning tree T, as described in the proof of (4.20).

Optimality of the Reverse-Delete Algorithm

- Reverse-Delete algorithm: Maintain a set E' of edges.
  - Start with $E' = E$.
  - Process edges in decreasing order of cost.
  - Delete the next edge e from E' only if (V, E') is connected after deletion.
  - Stop after processing all the edges.
- Claim: the Reverse-Delete algorithm outputs an MST.

Optimality of the Reverse-Delete Algorithm

- Reverse-Delete algorithm: Maintain a set E' of edges.
  - Start with $E' = E$.
  - Process edges in decreasing order of cost.
  - Delete the next edge e from E' only if (V, E') is connected after deletion.
  - Stop after processing all the edges.
- Claim: the Reverse-Delete algorithm outputs an MST.
  1. Show that every edge deleted belongs to no MST.

  2. Prove that the graph remaining at the end is a spanning tree.

Optimality of the Reverse-Delete Algorithm

- Reverse-Delete algorithm: Maintain a set E' of edges.
    - Start with $E' = E$.
    - Process edges in decreasing order of cost.
    - Delete the next edge e from E' only if (V, E') is connected after deletion.
    - Stop after processing all the edges.
- Claim: the Reverse-Delete algorithm outputs an MST.
    1. Show that every edge deleted belongs to no MST.
        - A deleted edge must belong to some cycle C.
        - Since the edge is the first encountered by the algorithm, it is the most expensive edge in C.
    2. Prove that the graph remaining at the end is a spanning tree.

Optimality of the Reverse-Delete Algorithm

- Reverse-Delete algorithm: Maintain a set E' of edges.
  - Start with $E' = E$.
  - Process edges in decreasing order of cost.
  - Delete the next edge e from E' only if (V, E') is connected after deletion.
  - Stop after processing all the edges.
- Claim: the Reverse-Delete algorithm outputs an MST.
  1. Show that every edge deleted belongs to no MST.
     - A deleted edge must belong to some cycle C.
     - Since the edge is the first encountered by the algorithm, it is the most expensive edge in C.
  2. Prove that the graph remaining at the end is a spanning tree.
     - (V, E') is connected at the end, by construction.

Optimality of the Reverse-Delete Algorithm

- Reverse-Delete algorithm: Maintain a set E' of edges.
  - Start with $E' = E$.
  - Process edges in decreasing order of cost.
  - Delete the next edge e from E' only if (V, E') is connected after deletion.
  - Stop after processing all the edges.
- Claim: the Reverse-Delete algorithm outputs an MST.
  1. Show that every edge deleted belongs to no MST.
     - A deleted edge must belong to some cycle C.
     - Since the edge is the first encountered by the algorithm, it is the most expensive edge in C.
  2. Prove that the graph remaining at the end is a spanning tree.
     - (V, E') is connected at the end, by construction.
     - If (V, E') contains a cycle, consider the costliest edge in that cycle. The algorithm would have deleted that edge.

Comments on MST Algorithms

- To handle multiple edges with the same length, perturb each length by a random infinitesimal amount. Read the textbook.
- *Any* algorithm that constructs a spanning tree by including edges that satisfy the cut property and deleting edges that satisfy the cycle property will yield an MST!

Implementing Prim's Algorithm

```
PRIM'S ALGORITHM($G, c, s$)
$S = \{s\}$ and $T = \emptyset$
while $S \neq V$ do
    Compute $(u, v) = \arg\min_{(u,v): u \in S, v \in V-S} c(u, v)$
    Add the node $v$ to $S$ and add the edge $(u, v)$ to $T$.
```

- Implementation and analysis are very similar to Dijkstra's algorithm.
- Maintain S and store attachment costs $a(v) = \min_{e \in \text{cut}(S)} c(e)$ for every node $v \in V - S$ in a priority queue.
- At each step, extract the node v with the minimum attachment cost from the priority queue and update the attachment costs of the neighbours of v.

Final Version of Prim's Algorithm

```
PRIM'S ALGORITHM($G, c, s$)
Insert($Q, s, 0, \emptyset$)
while $S \neq V$ do
    $(v, a(v), u)$ = ExtractMin($Q$)
    Add node $v$ to $S$ and edge $(u, v)$ to $T$.
    for every node $x \in V - S$ such that $(v, x)$ is an edge in $G$ do
        if $c(v, x) < a(x)$ then
            $a(x) = c(v, x)$
            ChangeKey($Q, x, a(x), v$)
```

- Q is a priority queue.
- Each element in Q is a triple: the node, its attachment cost, and its predecessor in the MST.
- In Step 8, if x is not already in Q, simply insert $(x, a(x), v)$ into Q.
- Total of $n - 1$ ExtractMin and m ChangeKey operations, yielding a running time of $O(m \log n)$.

Implementing Kruskal's Algorithm

- Start with an empty set T of edges.
- Process edges in E in increasing order of cost.
- Add the next edge e to T only if adding e does not create a cycle.

Implementing Kruskal's Algorithm

- Start with an empty set T of edges.
- Process edges in E in increasing order of cost.
- Add the next edge e to T only if adding e does not create a cycle.

- Sorting edges takes $O(m \log n)$ time.
- Key question: "Does adding $e = (u, v)$ to T create a cycle?"
  - Maintain set of connected components of T.
  - $\textsc{Find}(u)$: return the name of the connected component of T that u belongs to.
  - $\textsc{Union}(A, B)$: merge connected components A and B.

Analysing Kruskal's Algorithm

- How many FIND invocations does Kruskal's algorithm need?

Analysing Kruskal's Algorithm

- How many FIND invocations does Kruskal's algorithm need? $2m$.
- How many UNION invocations does Kruskal's algorithm need?

Analysing Kruskal's Algorithm

- How many FIND invocations does Kruskal's algorithm need? $2m$.
- How many UNION invocations does Kruskal's algorithm need? $n - 1$.

Analysing Kruskal's Algorithm

- How many FIND invocations does Kruskal's algorithm need? $2m$.
- How many UNION invocations does Kruskal's algorithm need? $n - 1$.
- Textbook describes two implementations of UNION-FIND: (see appendix to this set of slides)
    - Each FIND takes $O(1)$ time, k invocations of UNION take $O(k \log k)$ time in total.
    - Each FIND takes $O(\log n)$ time and each invocation of UNION takes $O(1)$ time.

Analysing Kruskal's Algorithm

- How many FIND invocations does Kruskal's algorithm need? $2m$.
- How many UNION invocations does Kruskal's algorithm need? $n-1$.
- Textbook describes two implementations of UNION-FIND: (see appendix to this set of slides)
  - Each FIND takes $O(1)$ time, k invocations of UNION take $O(k \log k)$ time in total.
  - Each FIND takes $O(\log n)$ time and each invocation of UNION takes $O(1)$ time.
- Total running time of Kruskal's algorithm is $O(m \log n)$.

Comments on Union-Find and MST

- The UNION-FIND data structure is useful to maintain the connected components of a graph as edges are added to the graph.
- The data structure does not support edge deletion efficiently.
- Current best algorithm for MST runs in $O(m\alpha(m, n))$ time (Chazelle 2000) and $O(m)$ randomised time (Karger, Klein, and Tarjan, 1995).
- Holy grail: $O(m)$ deterministic algorithm for MST.

Union-Find Data Structure

- Abstraction of the data structure needed by Kruskal's algorithm.
- Maintain disjoint subsets of elements from a universe U of n elements.
- Each subset has an name. We will set a set's name to be the identity of some element in it.
- Support three operations:
  1. MakeUnionFind(U): initialise the data structure with elements in U.
  2. Find(u): return the identity of the subset that contains u.
  3. Union(A, B): merge the sets named A and B into one set.

Union-Find Data Structure: Implementation 1

- Store all the elements of U in an array COMPONENT.
  - Assume identities of elements are integers from 1 to n.
  - COMPONENT$[s]$ is the name of the set containing s.
- Implementing the operations:

Union-Find Data Structure: Implementation 1

- Store all the elements of U in an array COMPONENT.
  - Assume identities of elements are integers from 1 to n.
  - COMPONENT$[s]$ is the name of the set containing s.
- Implementing the operations:
  1. MAKEUNIONFIND(U): For each $s \in U$, set COMPONENT$[s] = s$ in $O(n)$ time.
  2. FIND(s): return COMPONENT$[s]$ in $O(1)$ time.
  3. UNION(A, B): merge B into A by scanning COMPONENT and updating each index whose value is B to the value A. Takes $O(n)$ time.

Union-Find Data Structure: Implementation 1

- Store all the elements of U in an array Component.
  - Assume identities of elements are integers from 1 to n.
  - Component$[s]$ is the name of the set containing s.
- Implementing the operations:
  1. MakeUnionFind(U): For each $s \in U$, set Component$[s] = s$ in $O(n)$ time.
  2. Find(s): return Component$[s]$ in $O(1)$ time.
  3. Union(A, B): merge B into A by scanning Component and updating each index whose value is B to the value A. Takes $O(n)$ time.
- Union is very slow because

Union-Find Data Structure: Implementation 1

- Store all the elements of U in an array COMPONENT.
  - Assume identities of elements are integers from 1 to n.
  - COMPONENT$[s]$ is the name of the set containing s.
- Implementing the operations:
  1. MAKEUNIONFIND(U): For each $s \in U$, set COMPONENT$[s] = s$ in $O(n)$ time.
  2. FIND(s): return COMPONENT$[s]$ in $O(1)$ time.
  3. UNION(A, B): merge B into A by scanning COMPONENT and updating each index whose value is B to the value A. Takes $O(n)$ time.
- UNION is very slow because we cannot efficiently find the elements that belong to a set.

Union-Find Data Structure: Implementation 2

- Optimisation 1: Use an array ELEMENTS
  - Indices of ELEMENTS range from 1 to n.
  - ELEMENTS$[s]$ stores the elements in the subset named s in a list.
- Execute UNION(A, B) by merging B into A in two steps:
  1. Updating COMPONENT for elements of B in $O(|B|)$ time.
  2. Append ELEMENTS$[B]$ to ELEMENTS$[A]$ in $O(1)$ time.
- UNION takes $\Omega(n)$ in the worst-case.

Union-Find Data Structure: Implementation 2

- Optimisation 1: Use an array ELEMENTS
  - Indices of ELEMENTS range from 1 to n.
  - ELEMENTS$[s]$ stores the elements in the subset named s in a list.
- Execute UNION(A, B) by merging B into A in two steps:
  1. Updating COMPONENT for elements of B in $O(|B|)$ time.
  2. Append ELEMENTS$[B]$ to ELEMENTS$[A]$ in $O(1)$ time.
- UNION takes $\Omega(n)$ in the worst-case.
- Optimisation 2: Store size of each set in an array (say, SIZE). If SIZE$[B] \leq$ SIZE$[A]$, merge B into A. Otherwise merge A into B. Update SIZE.

Jnion-Find Data Structure: Analysis of Implementation

- MAKEUNIONFIND(S) and FIND(u) are as before.

nion-Find Data Structure: Analysis of Implementation

- MakeUnionFind(S) and Find(u) are as before.
- Union(A, B): Running time is proportional to the size of the smaller set, which may be $\Omega(n)$.

nion-Find Data Structure: Analysis of Implementation

- MakeUnionFind(S) and Find(u) are as before.
- Union(A, B): Running time is proportional to the size of the smaller set, which may be $\Omega(n)$.
- Any sequence of k Union operations takes $O(k \log k)$ time.

nion-Find Data Structure: Analysis of Implementation

- MakeUnionFind(S) and Find(u) are as before.
- Union(A, B): Running time is proportional to the size of the smaller set, which may be $\Omega(n)$.
- Any sequence of k Union operations takes $O(k \log k)$ time.
    - k Union operations touch at most $2k$ elements.

nion-Find Data Structure: Analysis of Implementation

- MAKEUNIONFIND(S) and FIND(u) are as before.
- UNION(A, B): Running time is proportional to the size of the smaller set, which may be $\Omega(n)$.
- Any sequence of k UNION operations takes $O(k \log k)$ time.
  - k UNION operations touch at most $2k$ elements.
  - Intuition: running time of UNION is dominated by updates to COMPONENT. Charge each update to the element being updated and bound number of charges per element.

nion-Find Data Structure: Analysis of Implementation

- MAKEUNIONFIND(S) and FIND(u) are as before.
- UNION(A, B): Running time is proportional to the size of the smaller set, which may be $\Omega(n)$.
- Any sequence of k UNION operations takes $O(k \log k)$ time.
    - k UNION operations touch at most $2k$ elements.
    - Intuition: running time of UNION is dominated by updates to COMPONENT. Charge each update to the element being updated and bound number of charges per element.
    - Consider any element s. Every time s's set identity is updated, the size of the set containing s at least doubles $\Rightarrow$ s's set can change at most $\log(2k)$ times $\Rightarrow$ the total work done in k UNION operations is $O(k \log k)$.

nion-Find Data Structure: Analysis of Implementation

- MakeUnionFind(S) and Find(u) are as before.
- Union(A, B): Running time is proportional to the size of the smaller set, which may be $\Omega(n)$.
- Any sequence of k Union operations takes $O(k \log k)$ time.
    - k Union operations touch at most $2k$ elements.
    - Intuition: running time of Union is dominated by updates to Component. Charge each update to the element being updated and bound number of charges per element.
    - Consider any element s. Every time s's set identity is updated, the size of the set containing s at least doubles $\Rightarrow$ s's set can change at most $\log(2k)$ times $\Rightarrow$ the total work done in k Union operations is $O(k \log k)$.
- Find is fast in the worst case, Union is fast in an amortised sense. Can we make both operations worst-case efficient?

Union-Find Data Structure: Implementation 3

- Goal: Implement FIND in $O(\log n)$ and UNION in $O(1)$ worst-case time.

Union-Find Data Structure: Implementation 3

- Goal: Implement FIND in $O(\log n)$ and UNION in $O(1)$ worst-case time.
- Represent each subset in a tree using pointers:
    - Each tree node contains an element and a pointer to a parent.
    - The identity of the set is the identity of the element at the root.

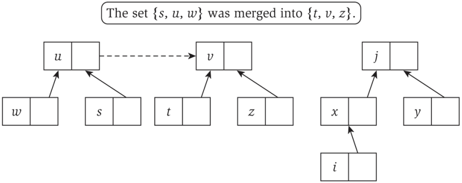

Figure 4.12 A `Union-Find` data structure using pointers. The data structure has only two sets at the moment, named after nodes v and j. The dashed arrow from u to v is the result of the last `Union` operation. To answer a `Find` query, we follow the arrows until we get to a node that has no outgoing arrow. For example, answering the query `Find`(i) would involve following the arrows i to x, and then x to j.

Union-Find Data Structure: Implementation 3

- Goal: Implement FIND in $O(\log n)$ and UNION in $O(1)$ worst-case time.
- Represent each subset in a tree using pointers:
    - Each tree node contains an element and a pointer to a parent.
    - The identity of the set is the identity of the element at the root.
- Implementing FIND(u): follow pointers from u to the root of u's tree.

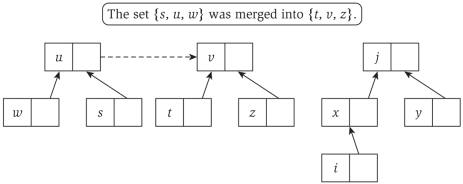

Figure 4.12 A `Union-Find` data structure using pointers. The data structure has only two sets at the moment, named after nodes v and j. The dashed arrow from u to v is the result of the last `Union` operation. To answer a `Find` query, we follow the arrows until we get to a node that has no outgoing arrow. For example, answering the query `Find`(i) would involve following the arrows i to x, and then x to j.

Union-Find Data Structure: Implementation 3

- Goal: Implement FIND in $O(\log n)$ and UNION in $O(1)$ worst-case time.
- Represent each subset in a tree using pointers:
  - Each tree node contains an element and a pointer to a parent.
  - The identity of the set is the identity of the element at the root.
- Implementing FIND(u): follow pointers from u to the root of u's tree.
- Implementing UNION(A, B): make smaller tree's root a child of the larger tree's root. Takes $O(1)$ time.

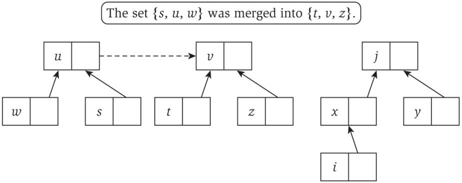

Figure 4.12 A `Union-Find` data structure using pointers. The data structure has only two sets at the moment, named after nodes v and j. The dashed arrow from u to v is the result of the last `Union` operation. To answer a `Find` query, we follow the arrows until we get to a node that has no outgoing arrow. For example, answering the query `Find`(i) would involve following the arrows i to x, and then x to j.

Union-Find Data Structure: Find in Implementation 3

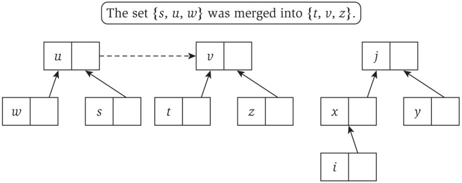

Figure 4.12 A `Union-Find` data structure using pointers. The data structure has only two sets at the moment, named after nodes v and j. The dashed arrow from u to v is the result of the last `Union` operation. To answer a `Find` query, we follow the arrows until we get to a node that has no outgoing arrow. For example, answering the query `Find`(i) would involve following the arrows i to x, and then x to j.

- Why does FIND(u) take $O(\log n)$ time?

Union-Find Data Structure: Find in Implementation 3

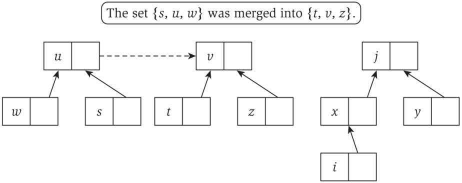

Figure 4.12 A `Union-Find` data structure using pointers. The data structure has only two sets at the moment, named after nodes v and j. The dashed arrow from u to v is the result of the last `Union` operation. To answer a `Find` query, we follow the arrows until we get to a node that has no outgoing arrow. For example, answering the query `Find`(i) would involve following the arrows i to x, and then x to j.

- Why does FIND(u) take $O(\log n)$ time?
- Number of pointers followed equals the number of times the identity of the set containing u changed.
- Every time u's set's identity changes, the set at least doubles in size ⇒ there are $O(\log n)$ pointers followed.

Union-Find Data Structure: Improving Implementation

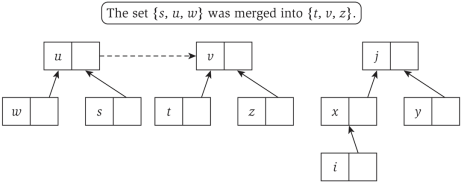

Figure 4.12 A `Union-Find` data structure using pointers. The data structure has only two sets at the moment, named after nodes v and j. The dashed arrow from u to v is the result of the last `Union` operation. To answer a `Find` query, we follow the arrows until we get to a node that has no outgoing arrow. For example, answering the query `Find`(i) would involve following the arrows i to x, and then x to j.

- Every time we invoke FIND(u), we follow the same set of pointers.

Union-Find Data Structure: Improving Implementation

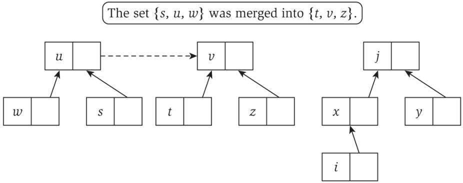
The set $\{s, u, w\}$ was merged into $\{t, v, z\}$.

Figure 4.12 A `Union-Find` data structure using pointers. The data structure has only two sets at the moment, named after nodes v and j. The dashed arrow from u to v is the result of the last `Union` operation. To answer a `Find` query, we follow the arrows until we get to a node that has no outgoing arrow. For example, answering the query `Find`(i) would involve following the arrows i to x, and then x to j.

- Every time we invoke $\text{Find}(u)$, we follow the same set of pointers.
- Path compression: make all nodes visited by $\text{Find}(u)$ children of the root.

Union-Find Data Structure: Improving Implementation

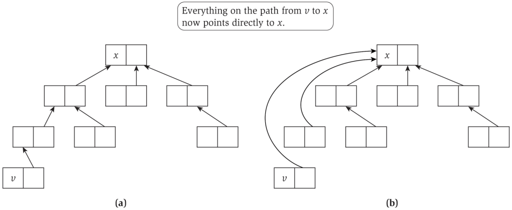

Figure 4.13 (a) An instance of a `Union-Find` data structure; and (b) the result of the operation `Find`(v) on this structure, using path compression.

- Every time we invoke FIND(u), we follow the same set of pointers.
- Path compression: make all nodes visited by FIND(u) children of the root.

Union-Find Data Structure: Improving Implementation

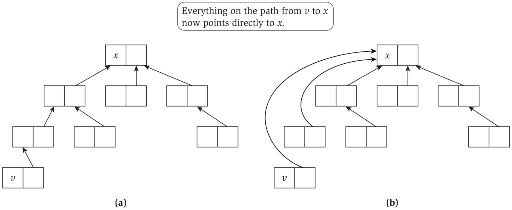

Figure 4.13 (a) An instance of a `Union-Find` data structure; and (b) the result of the operation `Find`(v) on this structure, using path compression.

- Every time we invoke $\textsc{Find}(u)$, we follow the same set of pointers.
- Path compression: make all nodes visited by $\textsc{Find}(u)$ children of the root.
- Can prove that total time taken by n $\textsc{Find}$ operations is $O(n\alpha(n))$, where $\alpha(n)$ is the inverse of the Ackermann function, and grows e-x-t-r-e-m-e-l-y s-l-o-w-l-y with n.