



CHAPTER 13

GRAPH ALGORITHMS

ACKNOWLEDGEMENT: THESE SLIDES ARE ADAPTED FROM SLIDES PROVIDED WITH DATA STRUCTURES AND ALGORITHMS IN C++, GOODRICH, TAMASSIA AND MOUNT (WILEY 2004) AND SLIDES FROM NANCY M. AMATO



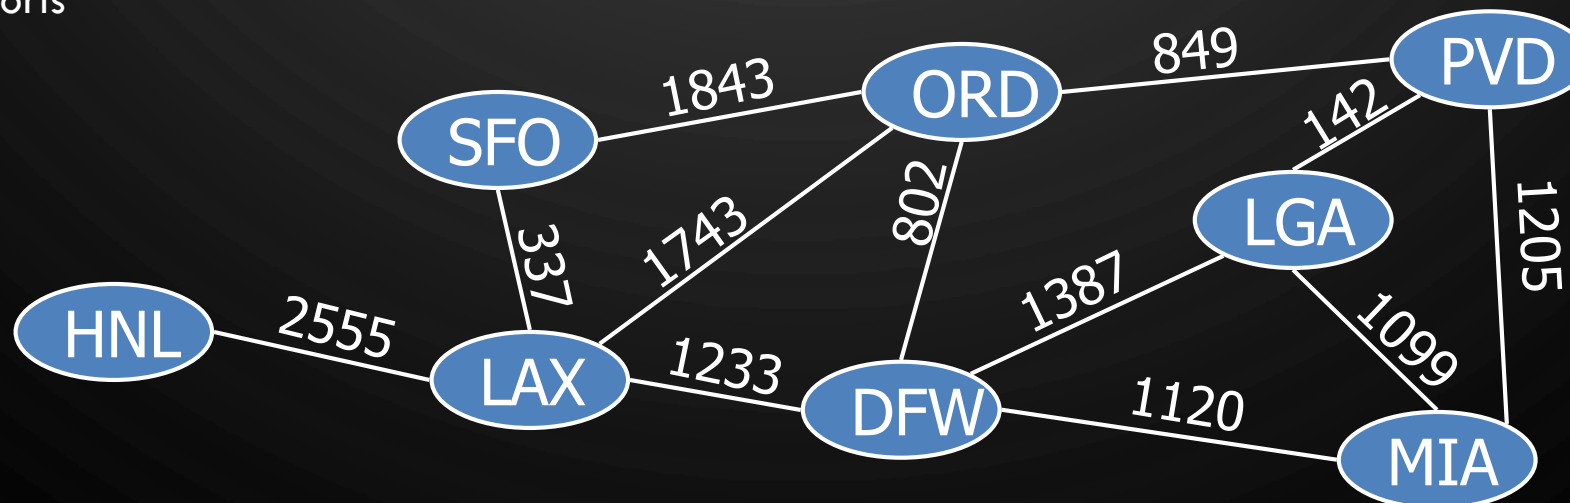
MINIMUM SPANNING TREES



REMINDER WEIGHTED GRAPHS

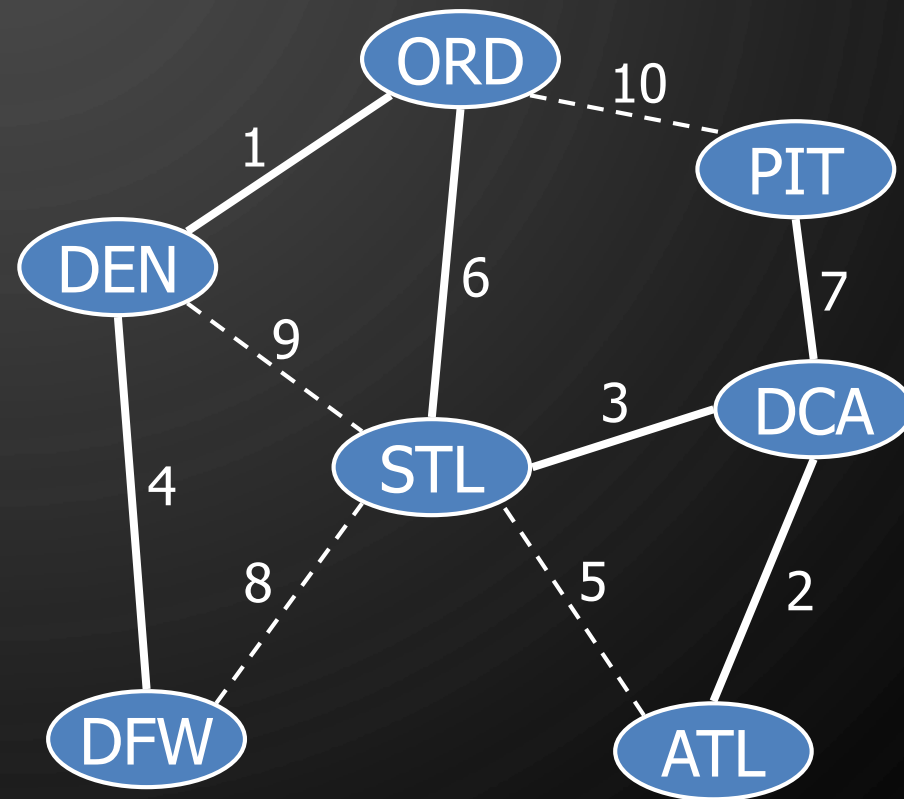


- In a weighted graph, each edge has an associated numerical value, called the weight of the edge
- Edge weights may represent, distances, costs, etc.
- Example:
 - In a flight route graph, the weight of an edge represents the distance in miles between the endpoint airports



MINIMUM SPANNING TREE

- Spanning subgraph
 - Subgraph of a graph G containing all the vertices of G
- Spanning tree
 - Spanning subgraph that is itself a (free) tree
- Minimum spanning tree (MST)
 - Spanning tree of a weighted graph with minimum total edge weight
- Applications
 - Communications networks
 - Transportation networks




EXERCISE

MST



- Show an MST of the following graph.




CYCLE PROPERTY

- Cycle Property:


- Let T be a minimum spanning tree of a weighted graph G
- Let e be an edge of G that is not in T and C let be the cycle formed by e with T
- For every edge f of C , $weight(f) \leq weight(e)$

- Proof:

- By contradiction
- If $weight(f) > weight(e)$ we can get a spanning tree of smaller weight by replacing e with f

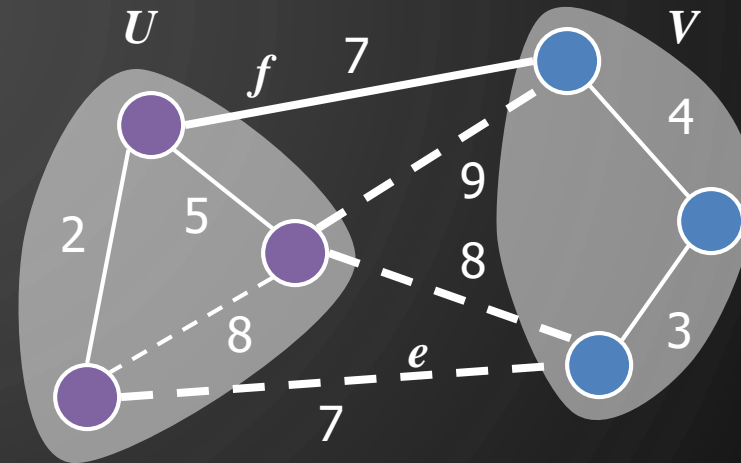


Replacing f with e yields a better spanning tree




PARTITION PROPERTY

- Partition Property:
 - Consider a partition of the vertices of G into subsets U and V
 - Let e be an edge of minimum weight across the partition
 - There is a minimum spanning tree of G containing edge e
- Proof:
 - Let T be an MST of G
 - If T does not contain e , consider the cycle C formed by e with T and let f be an edge of C across the partition
 - By the cycle property,
$$\text{weight}(f) \leq \text{weight}(e)$$
 - Thus, $\text{weight}(f) = \text{weight}(e)$
 - We obtain another MST by replacing f with e

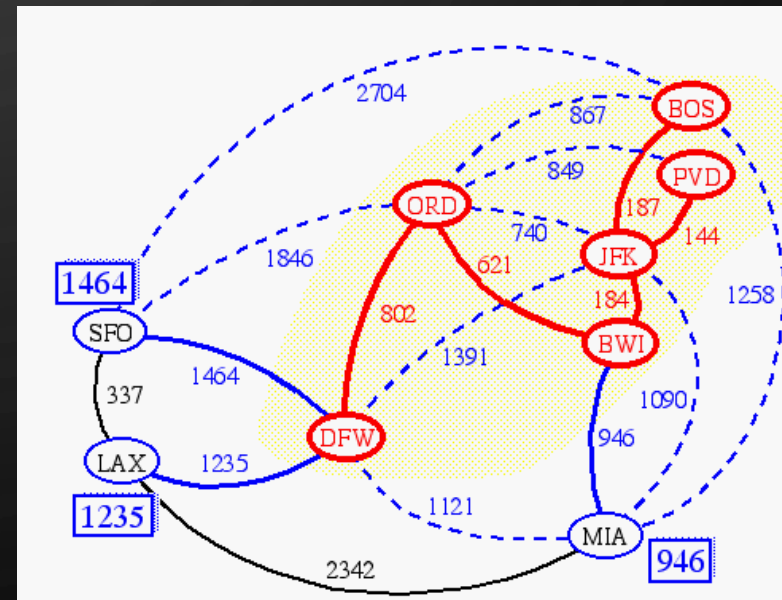


Replacing f with e yields another MST



PRIM-JARNIK'S ALGORITHM

- We pick an arbitrary vertex s and we grow the MST as a cloud of vertices, starting from s
- We store with each vertex v a label $d(v) =$ the smallest weight of an edge connecting v to a vertex in the cloud
- At each step:
 - We add to the cloud the vertex u outside the cloud with the smallest distance label
 - We update the labels of the vertices adjacent to u



PRIM-JARNIK'S ALGORITHM

- An adaptable priority queue stores the vertices outside the cloud
 - Key: distance, $D[v]$
 - Element: vertex v
 - $Q.replaceKey(i, k)$ changes the key of an item
- We store three labels with each vertex v :
 - Distance $D[v]$
 - Parent edge in MST $P[v]$
 - Locator in priority queue

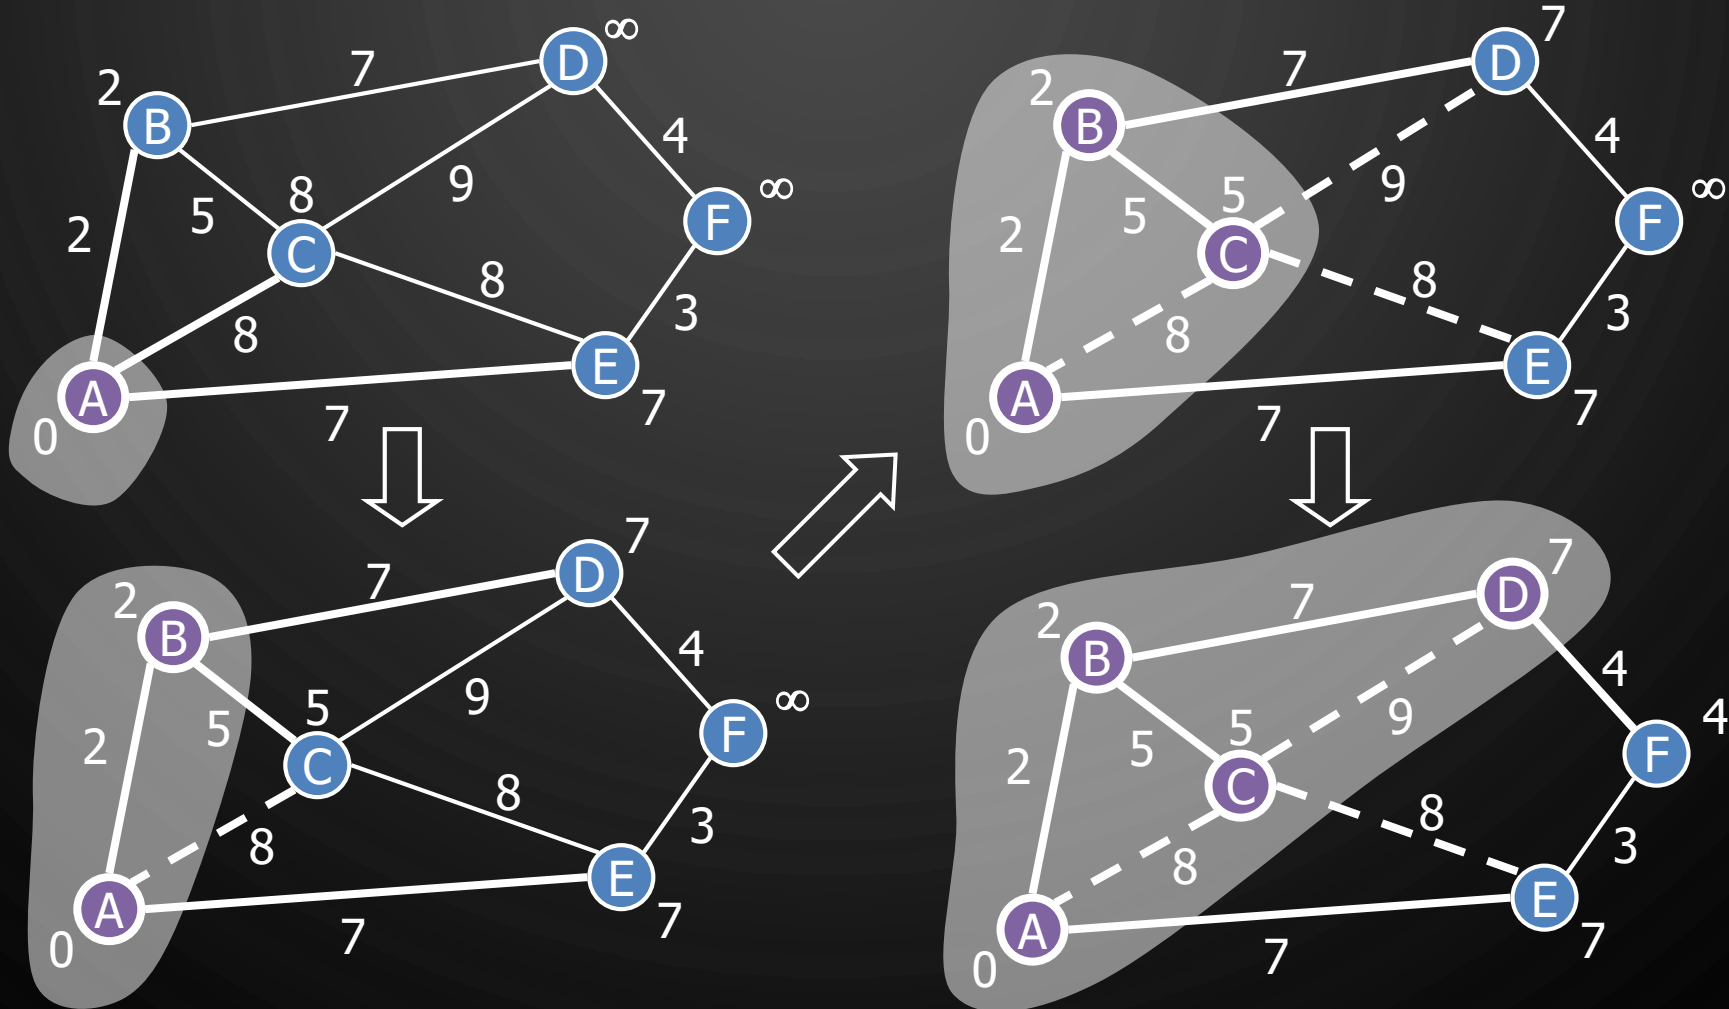
Algorithm PrimJarnikMST(G)

Input: A weighted connected graph G

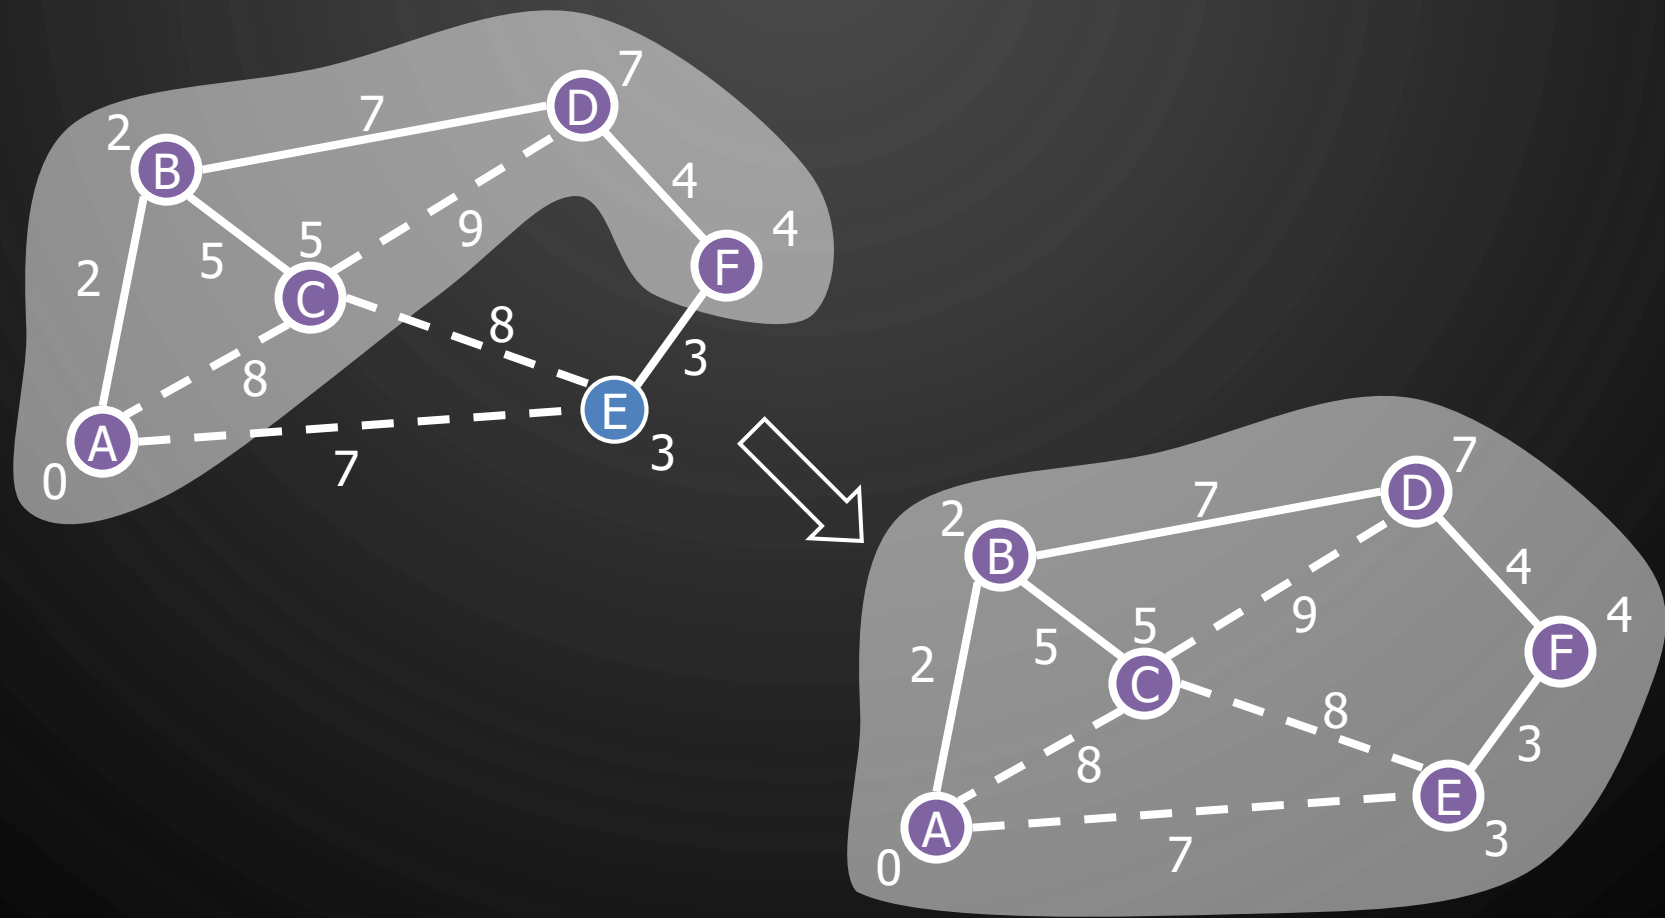
Output: A minimum spanning tree T of G

1. Pick any vertex v of G
2. $D[v] \leftarrow 0; P[v] \leftarrow \emptyset$
3. **for each** vertex $u \neq v$ **do**
4. $D[u] \leftarrow \infty; P[u] \leftarrow \emptyset$
5. $T \leftarrow \emptyset$
6. Priority queue Q of vertices with $D[u]$ as the key
7. **while** $\neg Q.empty()$ **do**
8. $u \leftarrow Q.removeMin()$
9. Add vertex u and edge $P[u]$ to T
10. **for each** $e \in u.incidentEdges()$ **do**
11. $z \leftarrow e.opposite(u)$
12. **if** $e.weight() < D[z]$
13. $D[z] \leftarrow e.weight(); P[z] \leftarrow e$
14. $Q.replaceKey(z, D[z])$
15. **return** T

EXAMPLE



EXAMPLE

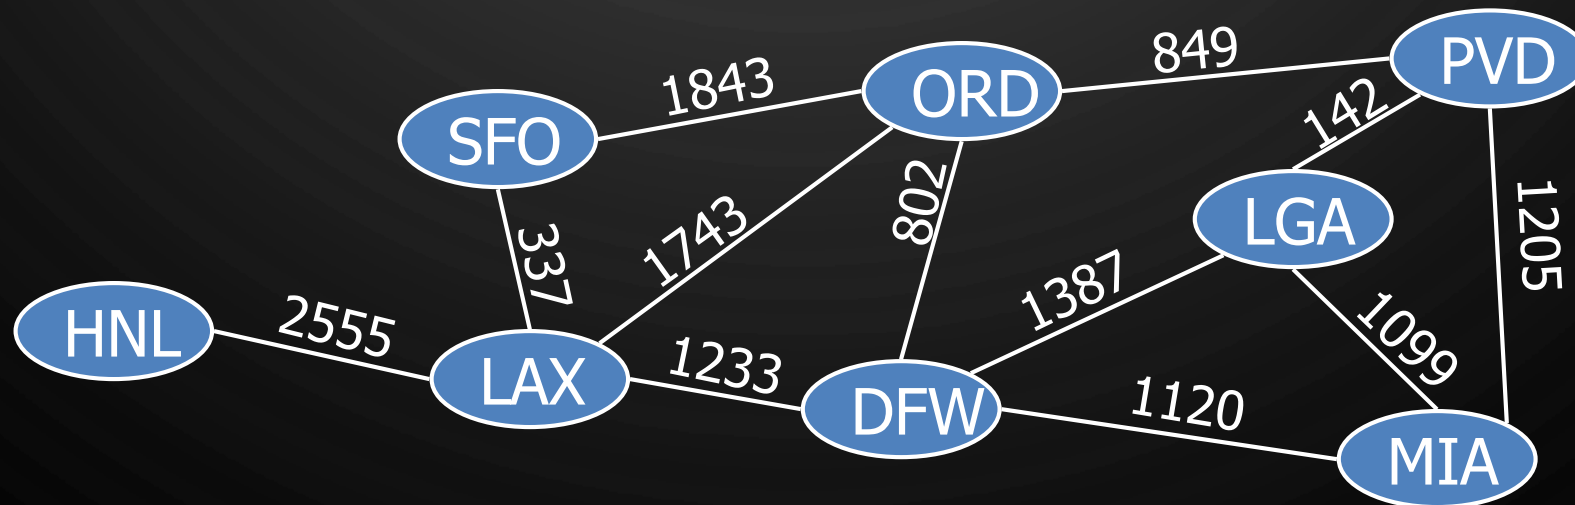


EXERCISE

PRIM'S MST ALGORITHM



- Show how Prim's MST algorithm works on the following graph, assuming you start with SFO
 - Show how the MST evolves in each iteration (a separate figure for each iteration).



ANALYSIS

- Graph operations
 - Method `incidentEdges` is called once for each vertex
- Label operations
 - We set/get the distance, parent and locator labels of vertex z $O(\deg(z))$ times
 - Setting/getting a label takes $O(1)$ time
- Priority queue operations
 - Each vertex is inserted once into and removed once from the priority queue, where each insertion or removal takes $O(\log n)$ time
 - The key of a vertex w in the priority queue is modified at most $\deg(w)$ times, where each key change takes $O(\log n)$ time
- Prim-Jarnik's algorithm runs in $O((n + m) \log n)$ time provided the graph is represented by the adjacency list structure
 - Recall that $\sum_v \deg(v) = 2m$
- The running time is $O(m \log n)$ since the graph is connected


KRUSKAL'S ALGORITHMS

- A priority queue stores the edges outside the cloud
 - Key: weight
 - Element: edge
- At the end of the algorithm
 - We are left with one cloud that encompasses the MST
 - A tree T which is our MST

Algorithm KruskalMST(G)

1. **for each** vertex $v \in G.vertices()$ **do**
2. Define a cluster $C(v) \leftarrow \{v\}$
3. Initialize a priority queue Q of edges using the weights as keys
4. $T \leftarrow \emptyset$
5. **while** T has fewer than $n - 1$ edges **do**
6. $(u, v) \leftarrow Q.removeMin()$
7. **if** $C(v) \neq C(u)$ **then**
8. Add (u, v) to T
9. Merge $C(v)$ and $C(u)$
10. **return** T


DATA STRUCTURE FOR KRUSKAL'S ALGORITHM



- The algorithm maintains a forest of trees
- An edge is accepted if it connects distinct trees
- We need a data structure that maintains a partition, i.e., a collection of disjoint sets, with the operations:
 - $\text{find}(u)$: return the set storing u
 - $\text{union}(u, v)$: replace the sets storing u and v with their union

REPRESENTATION OF A PARTITION


- Each set is stored in a sequence
- Each element has a reference back to the set
 - Operation $find(u)$ takes $O(1)$ time, and returns the set of which u is a member.
 - In operation $union(u, v)$, we move the elements of the smaller set to the sequence of the larger set and update their references
 - The time for operation $union(u, v)$ is $O(\min(n_u, n_v))$, where n_u and n_v are the sizes of the sets storing u and v
- Whenever an element is processed, it goes into a set of size at least double, hence each element is processed at most $\log n$ times



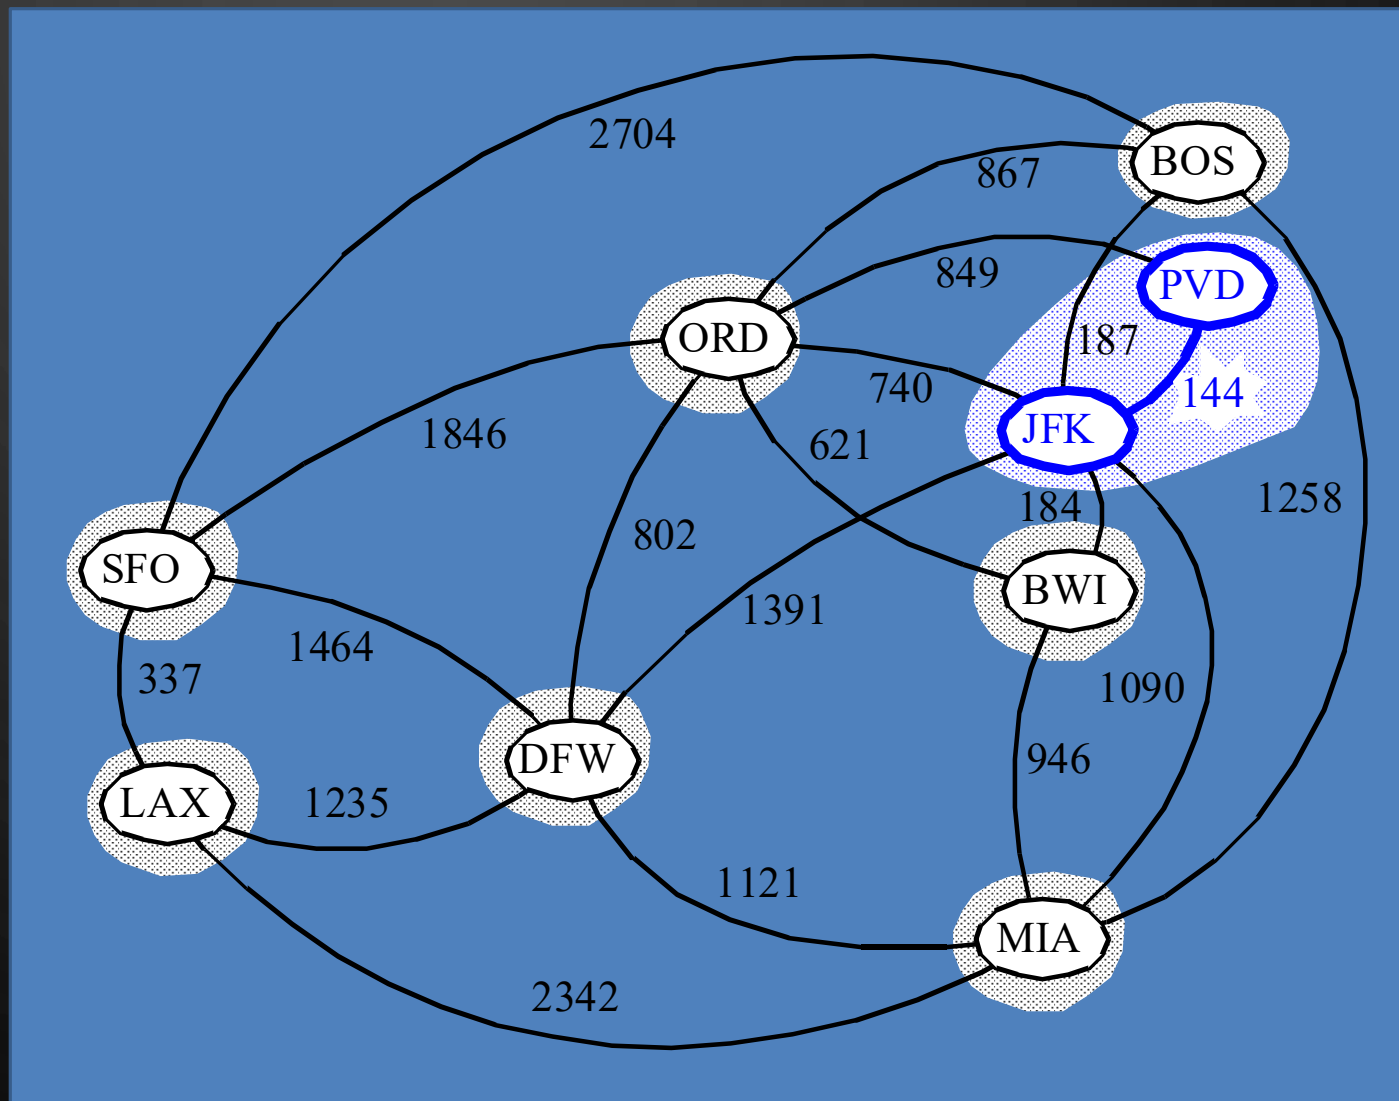
ANALYSIS

- A partition-based version of Kruskal's Algorithm performs cluster merges as unions and tests as finds.
- Running time
 - There will be at most m removals from the priority queue - $O(m \log n)$
 - Each vertex can be merged at most $\log n$ times, as the clouds tend to "double" in size - $O(n \log n)$
 - Total: $O((n + m) \log n)$

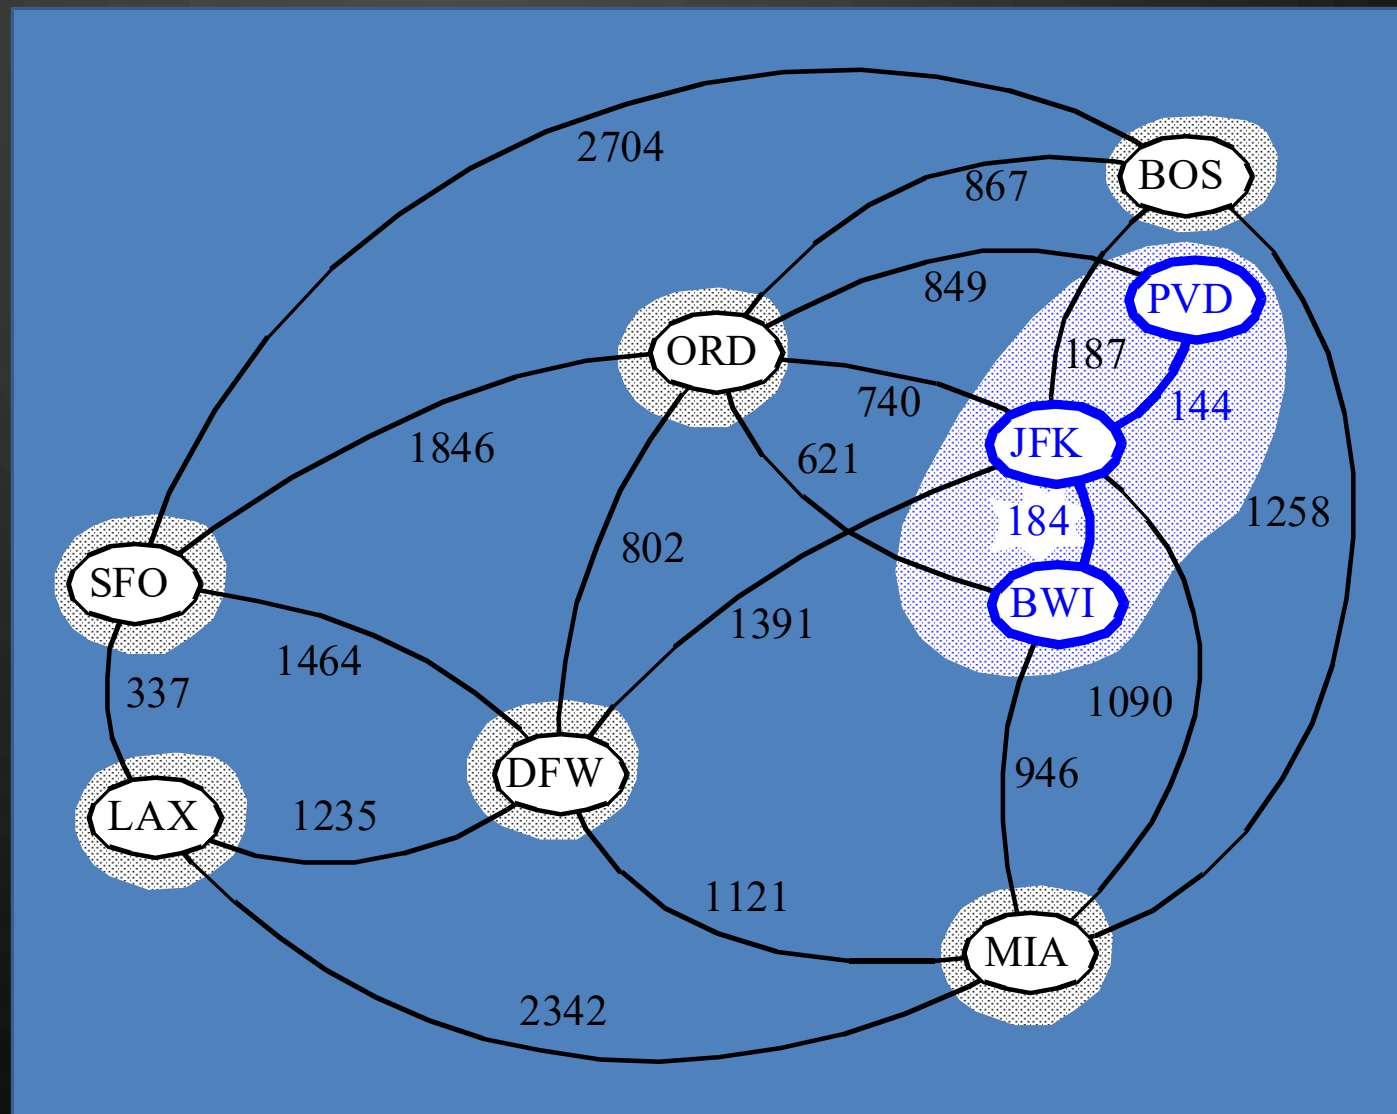
EXAMPLE



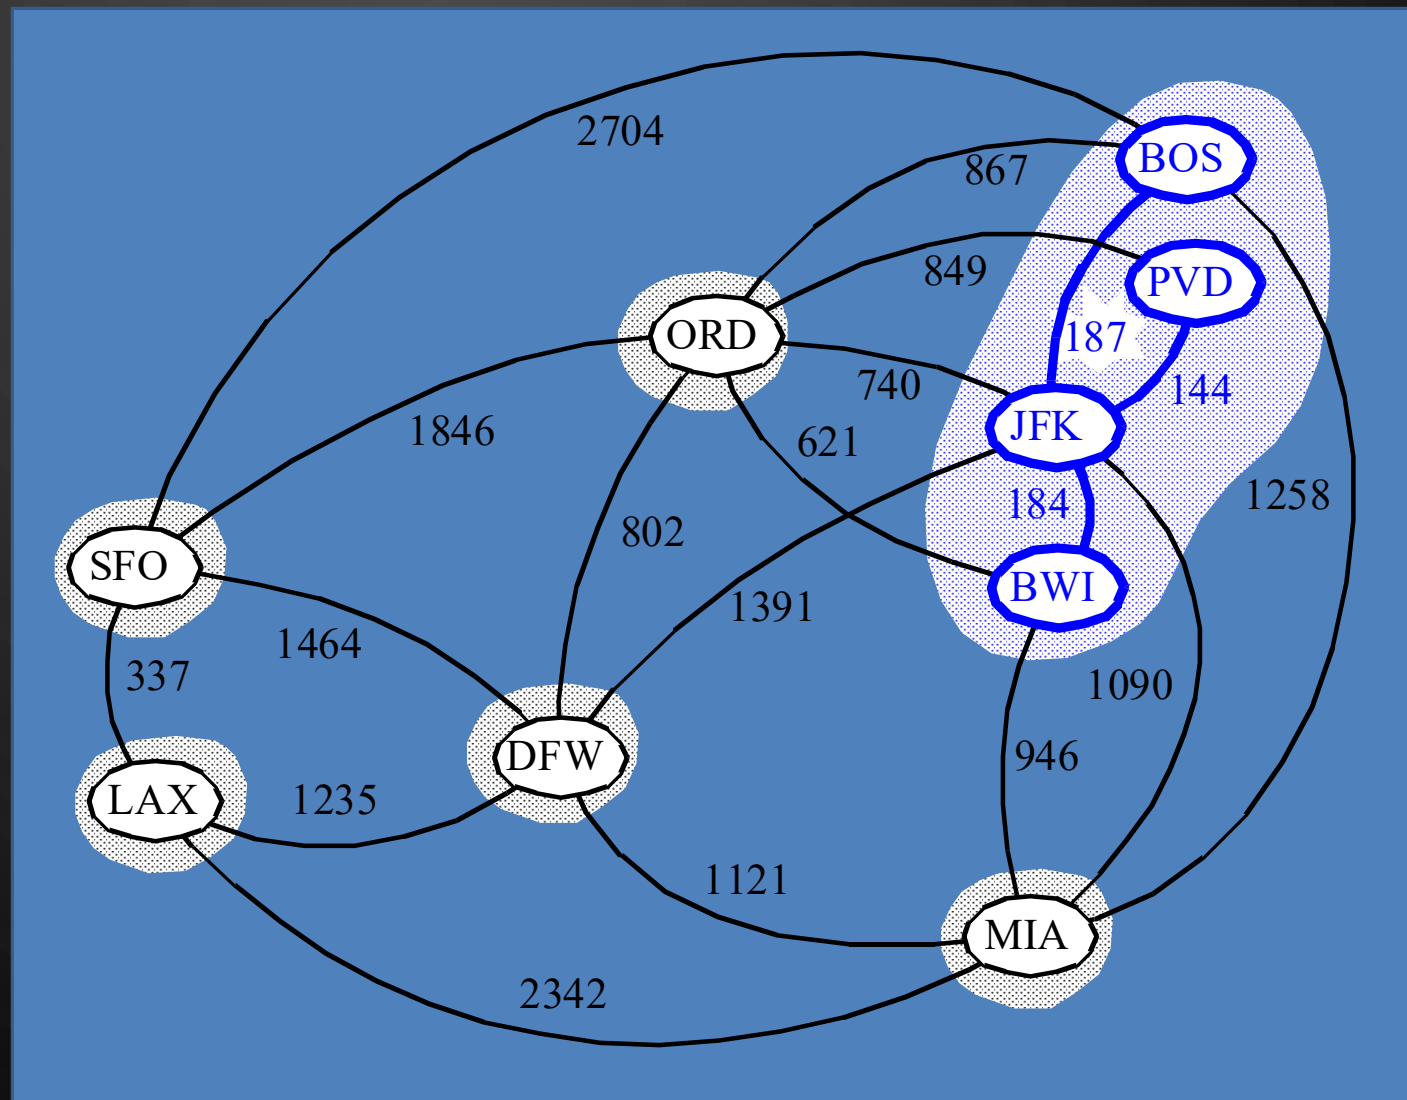
EXAMPLE



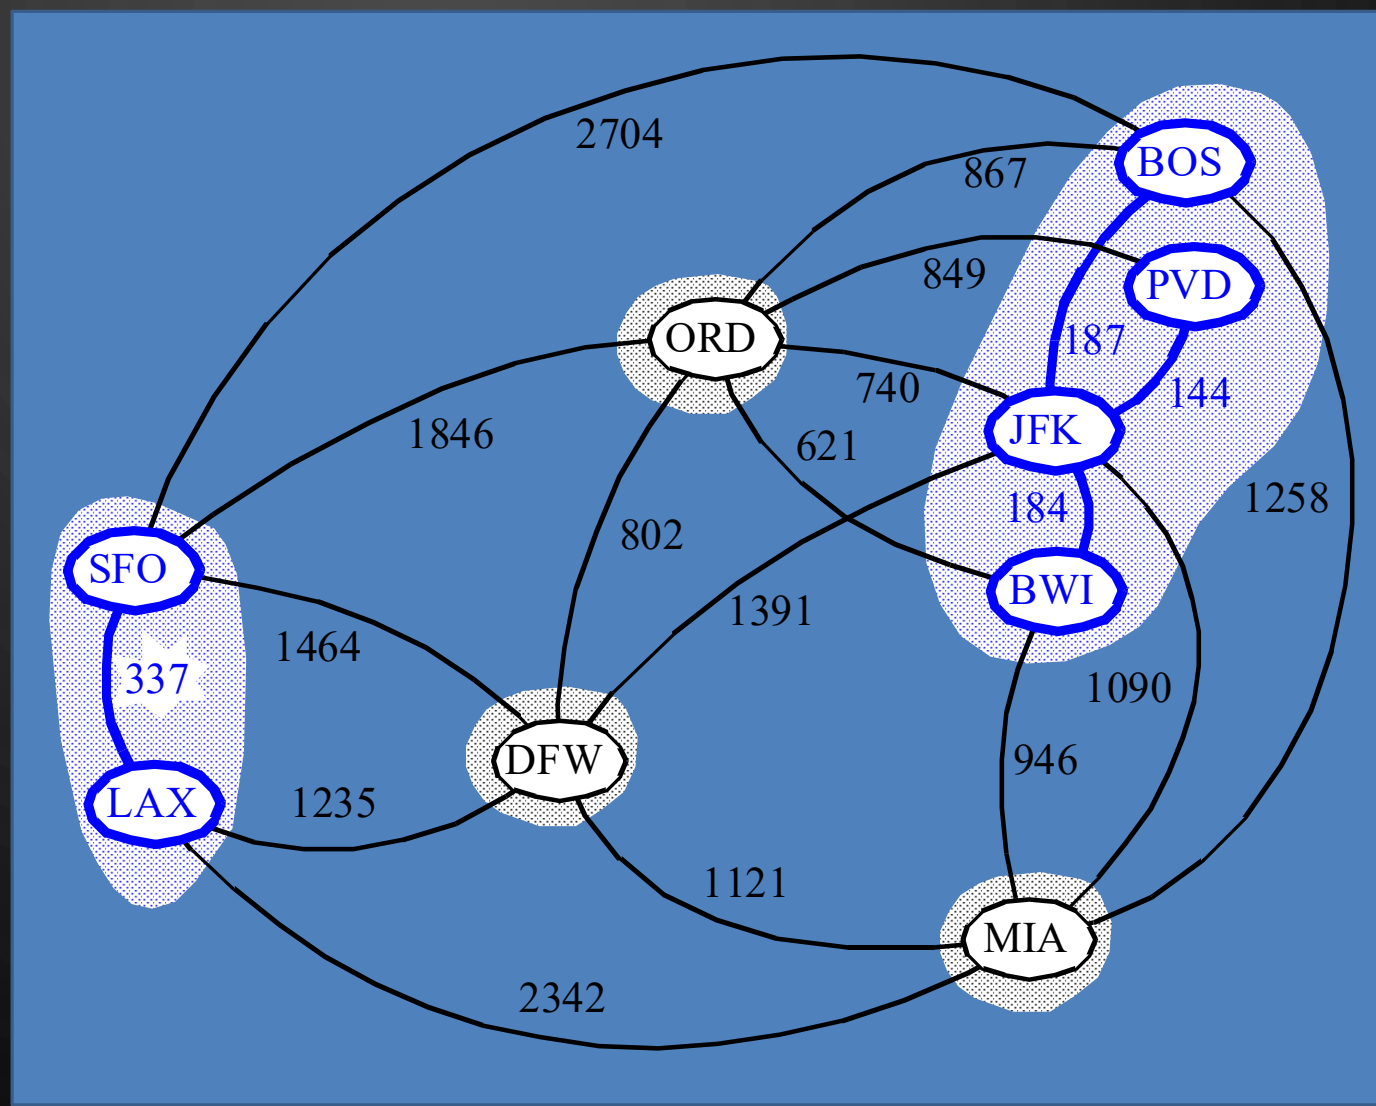
EXAMPLE



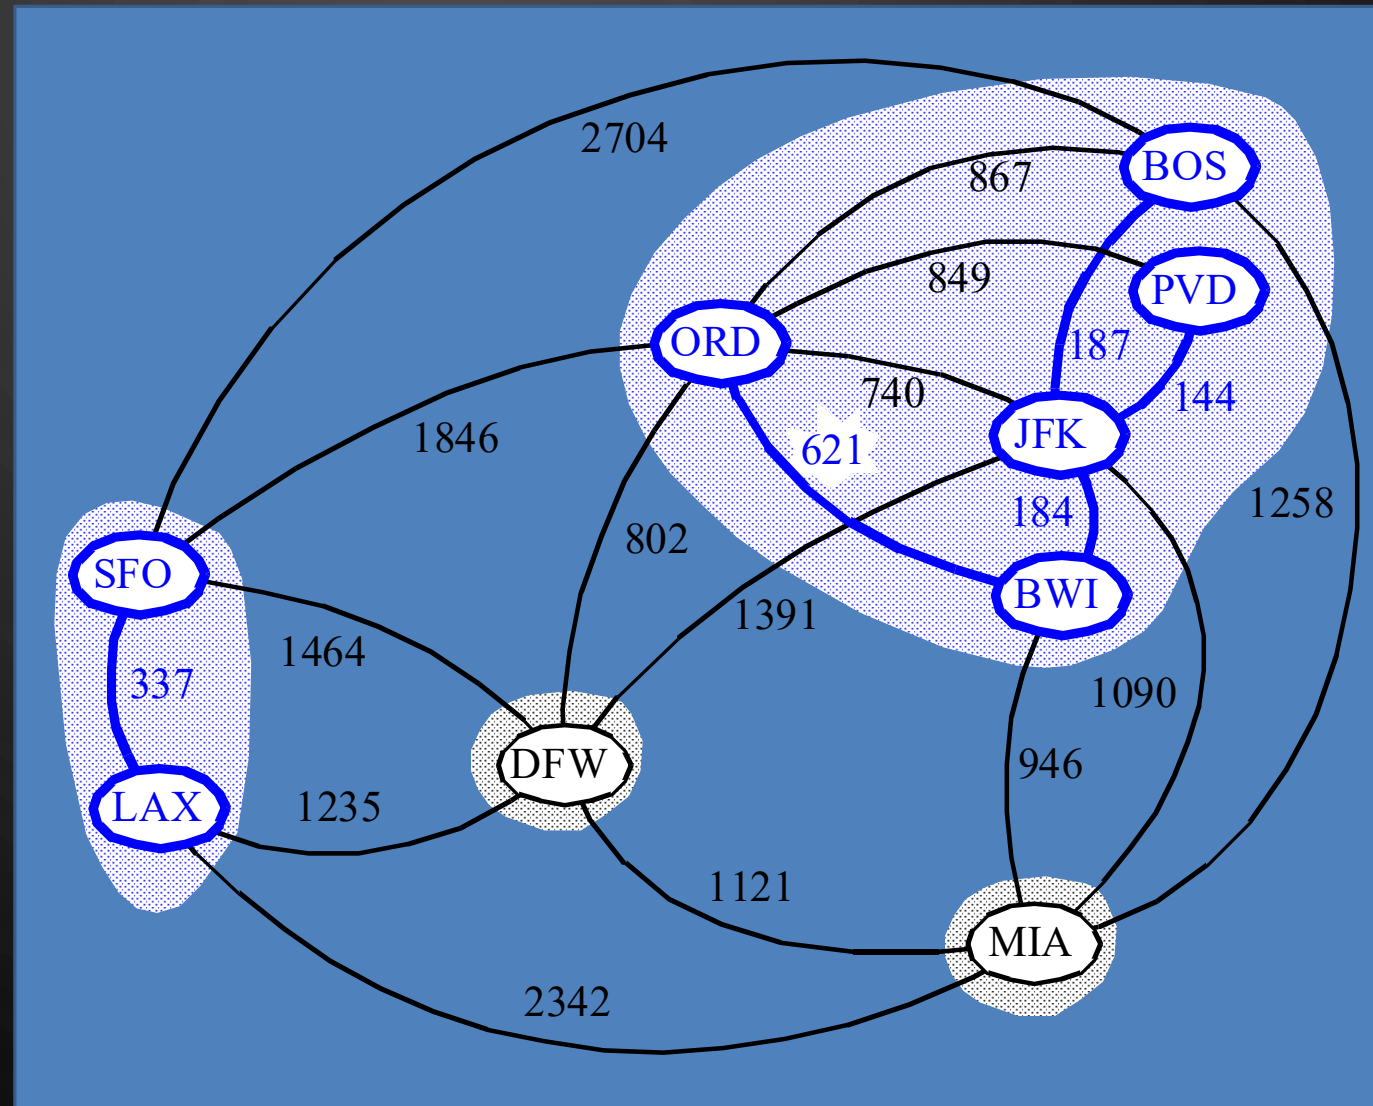
EXAMPLE



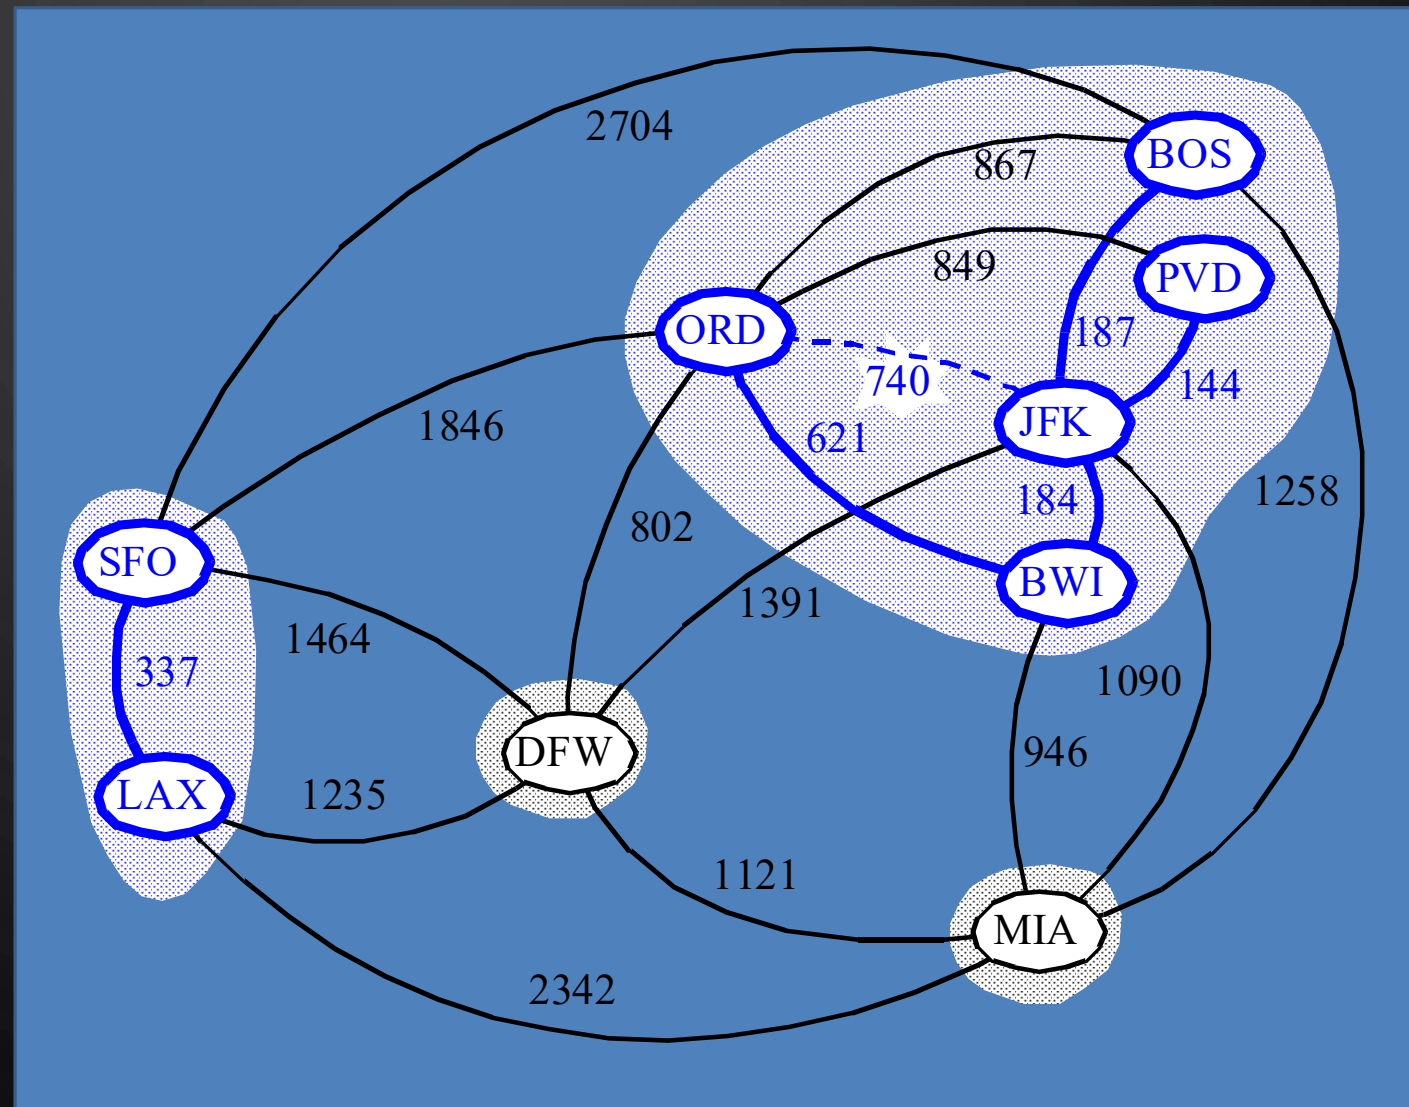
EXAMPLE



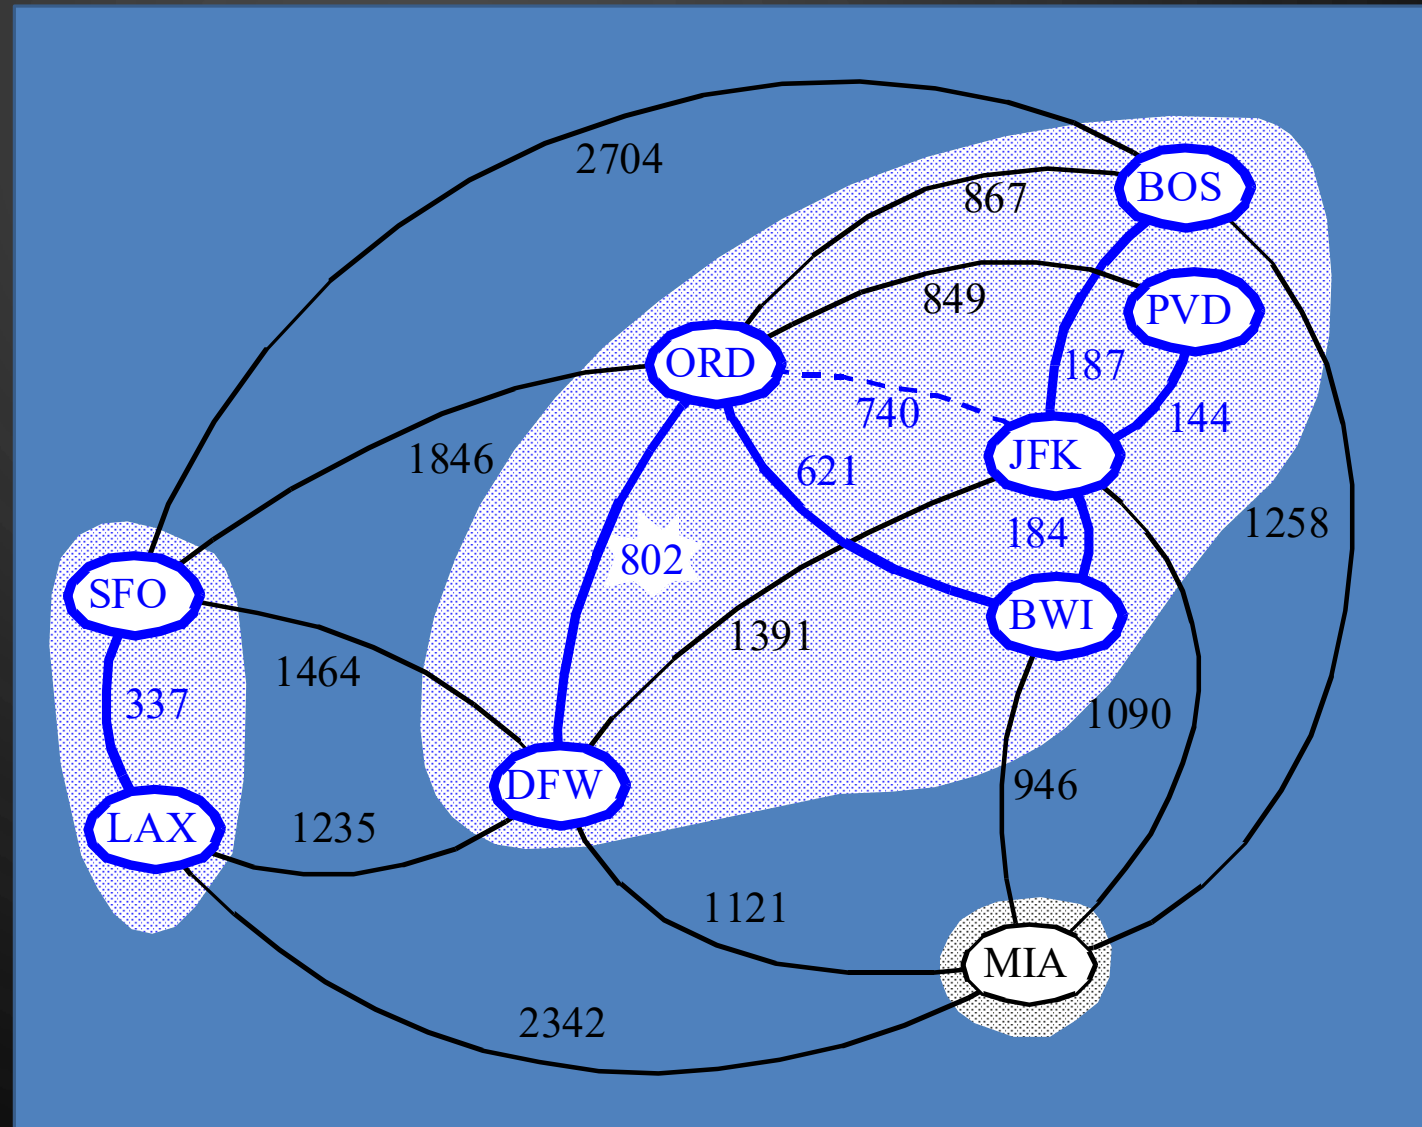
EXAMPLE



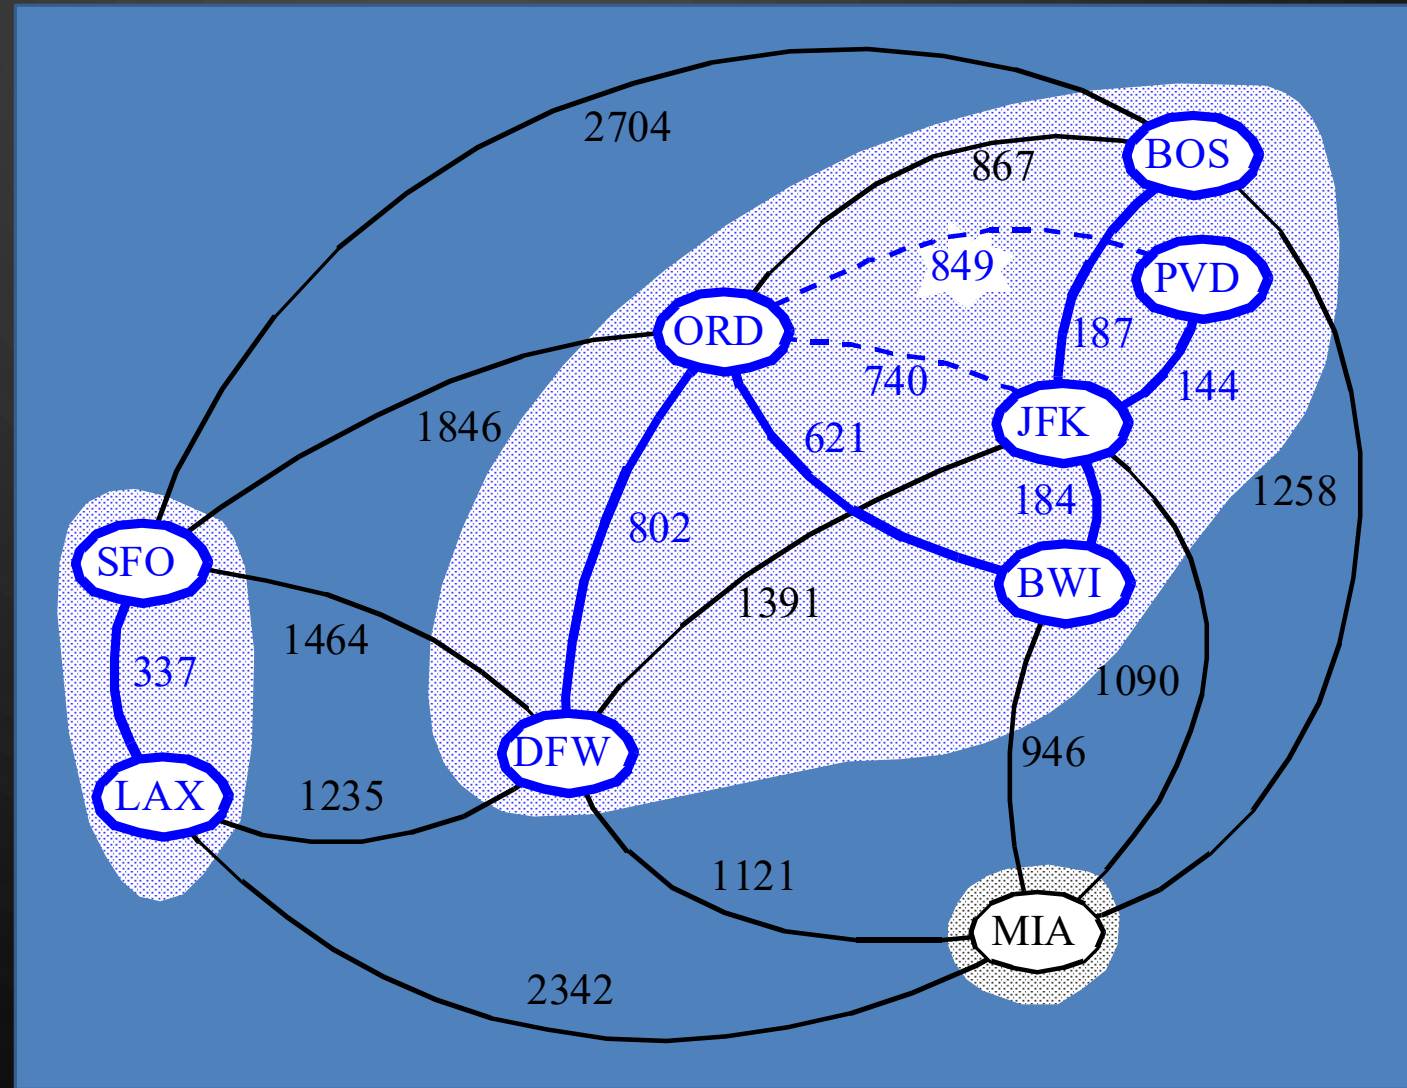
EXAMPLE



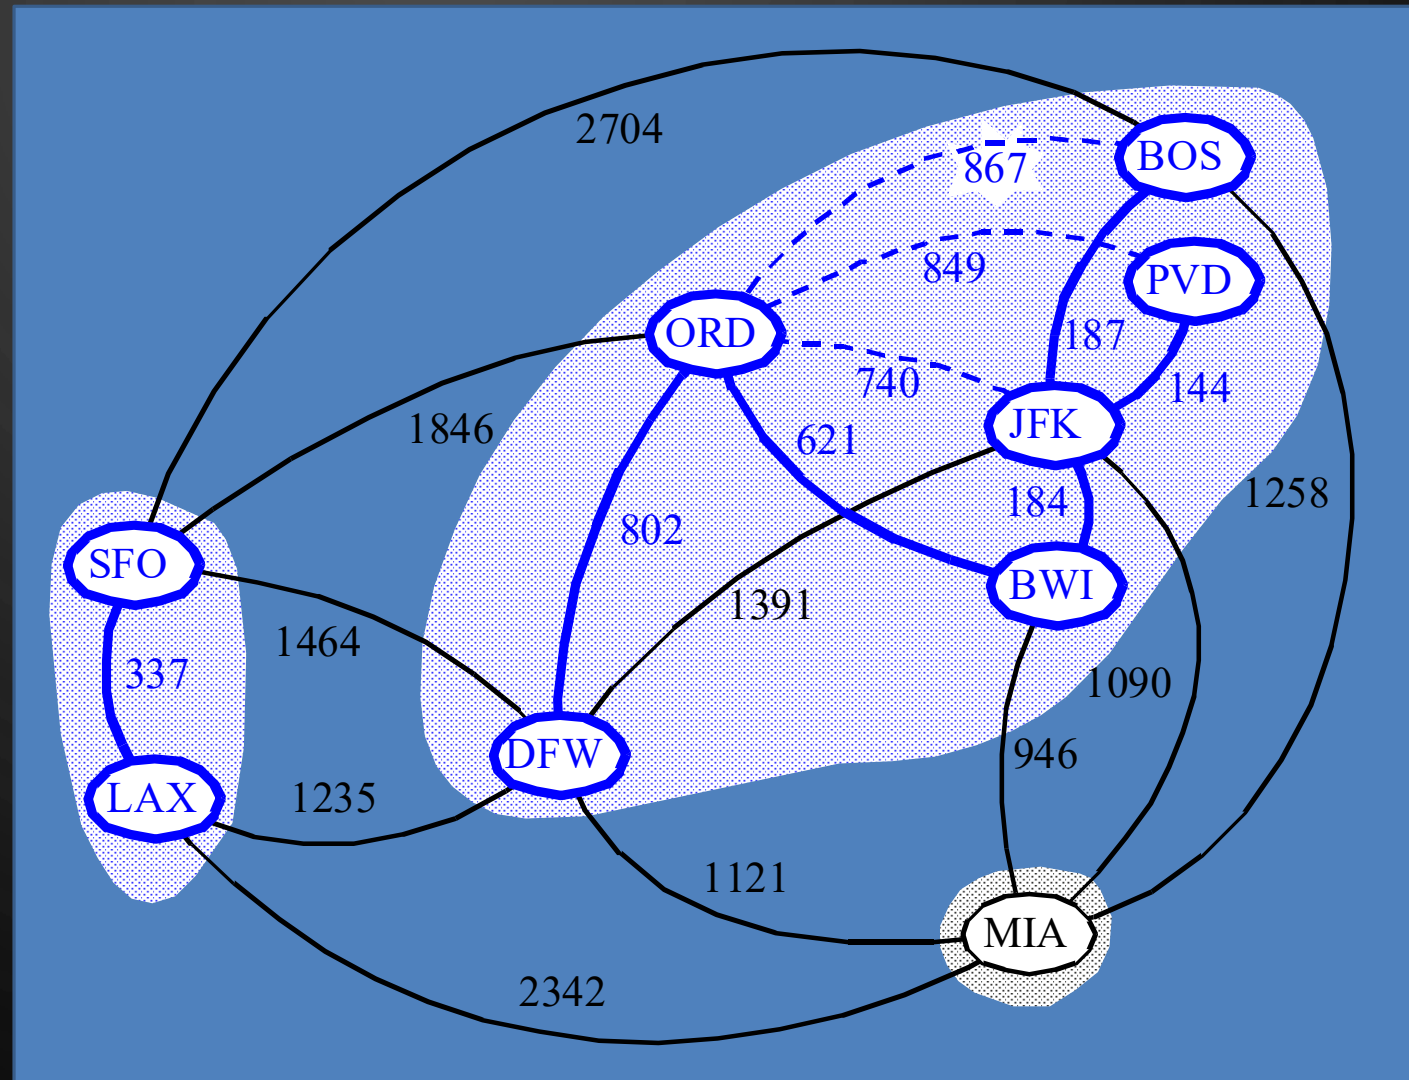
EXAMPLE



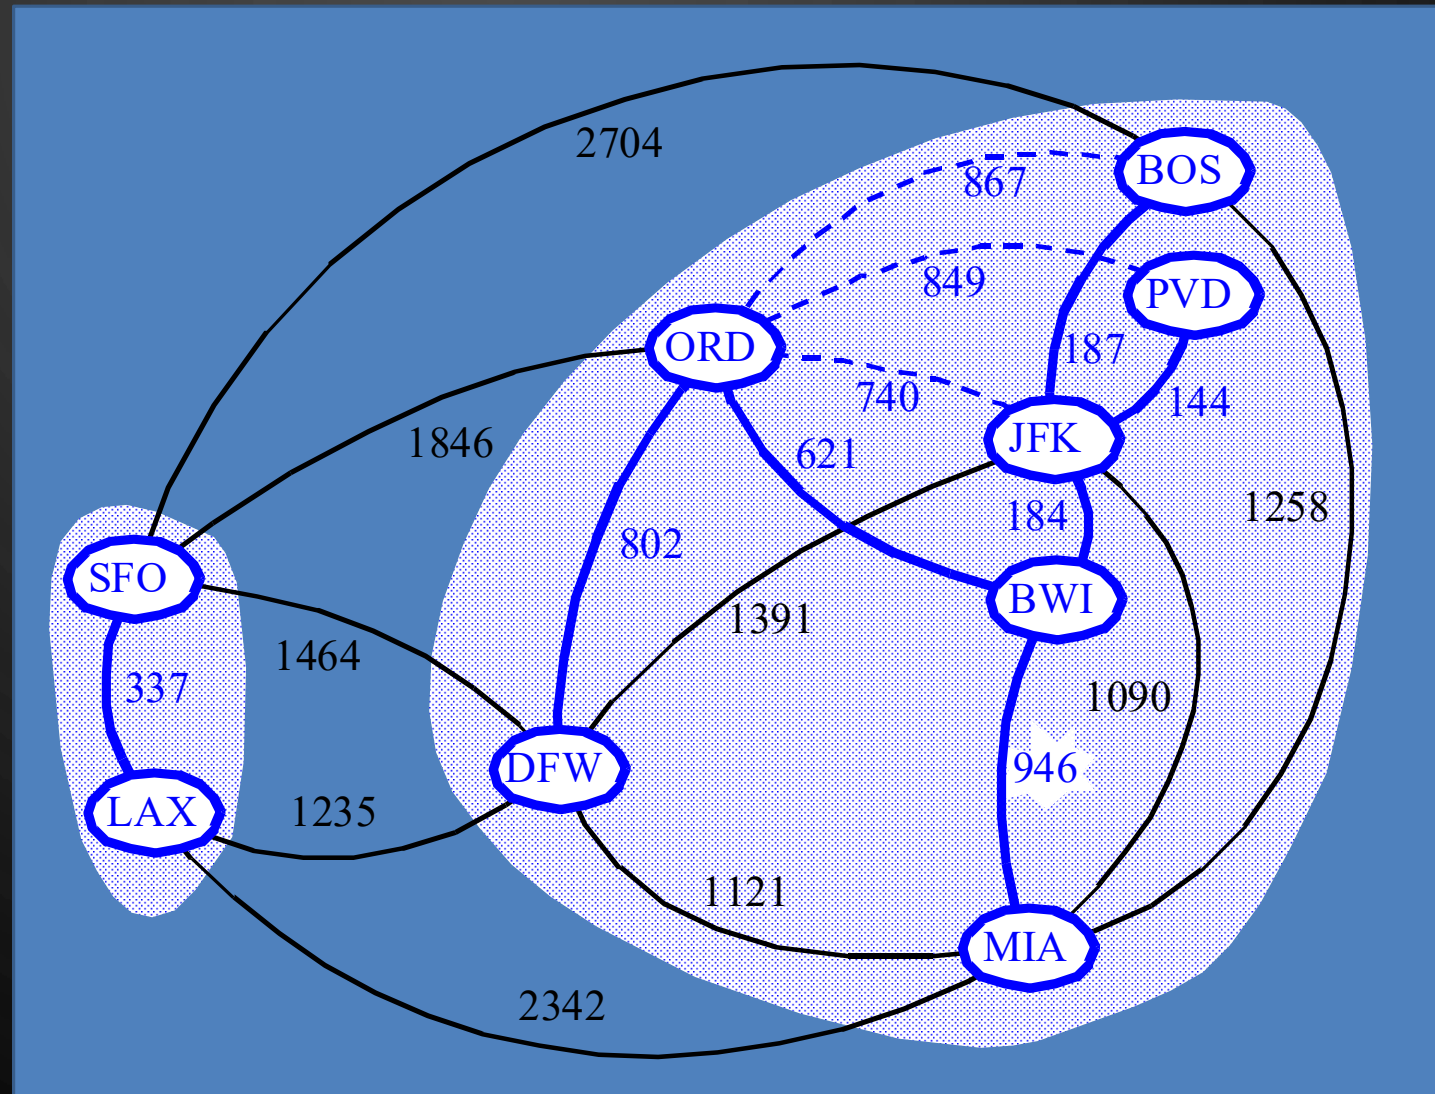
EXAMPLE



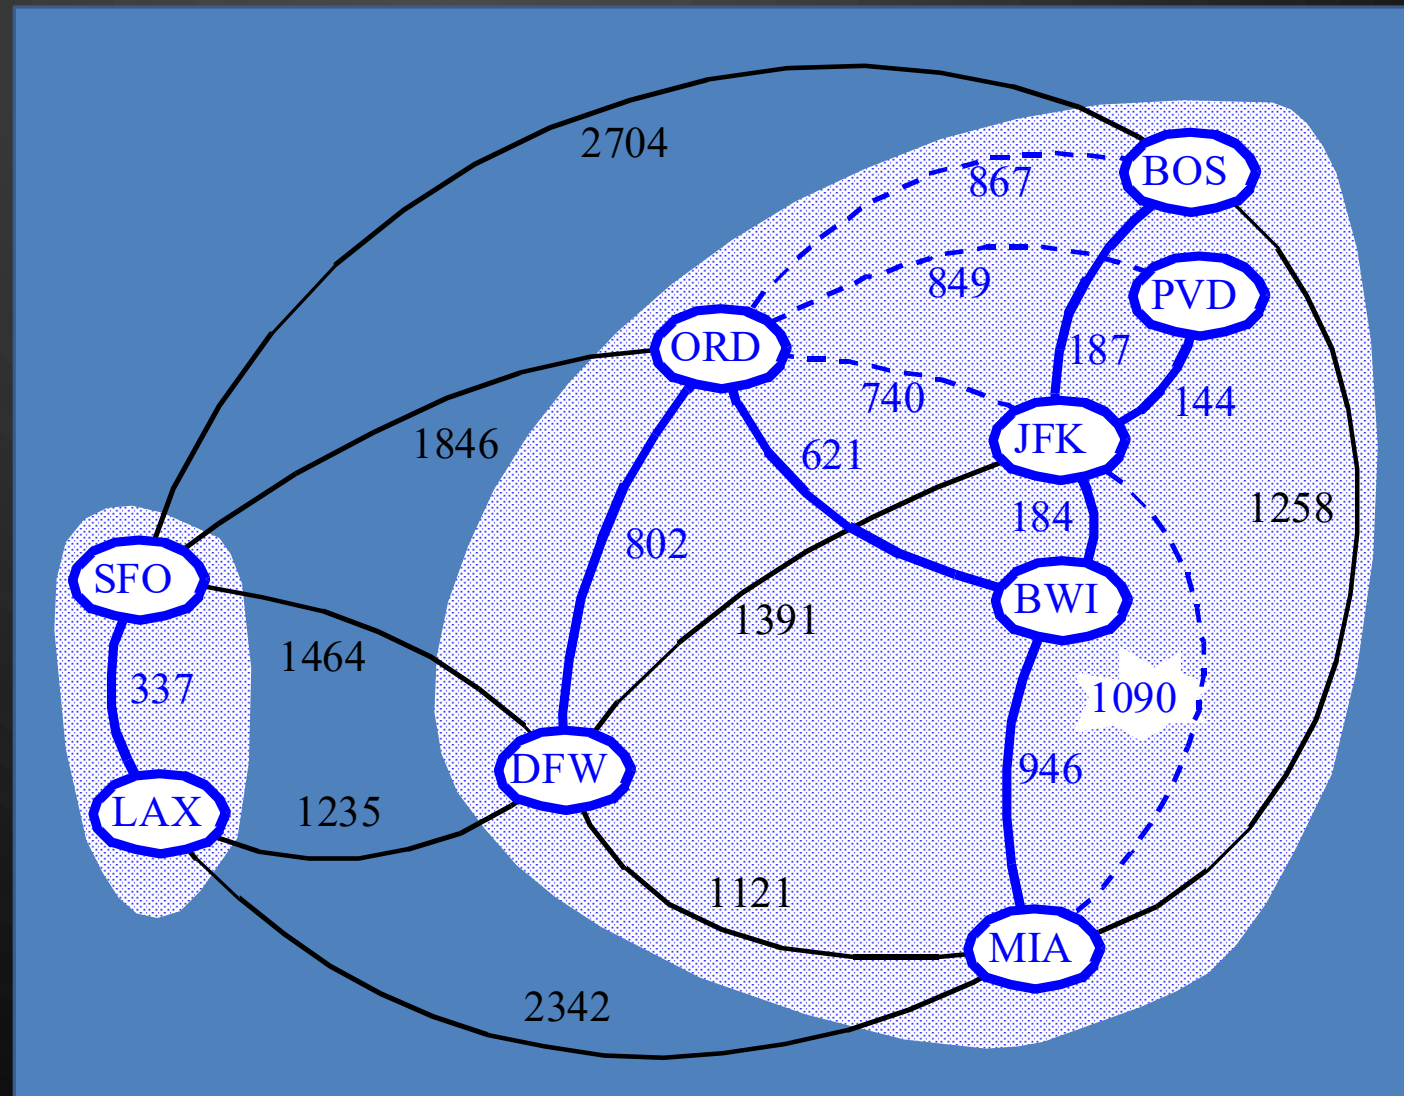
EXAMPLE



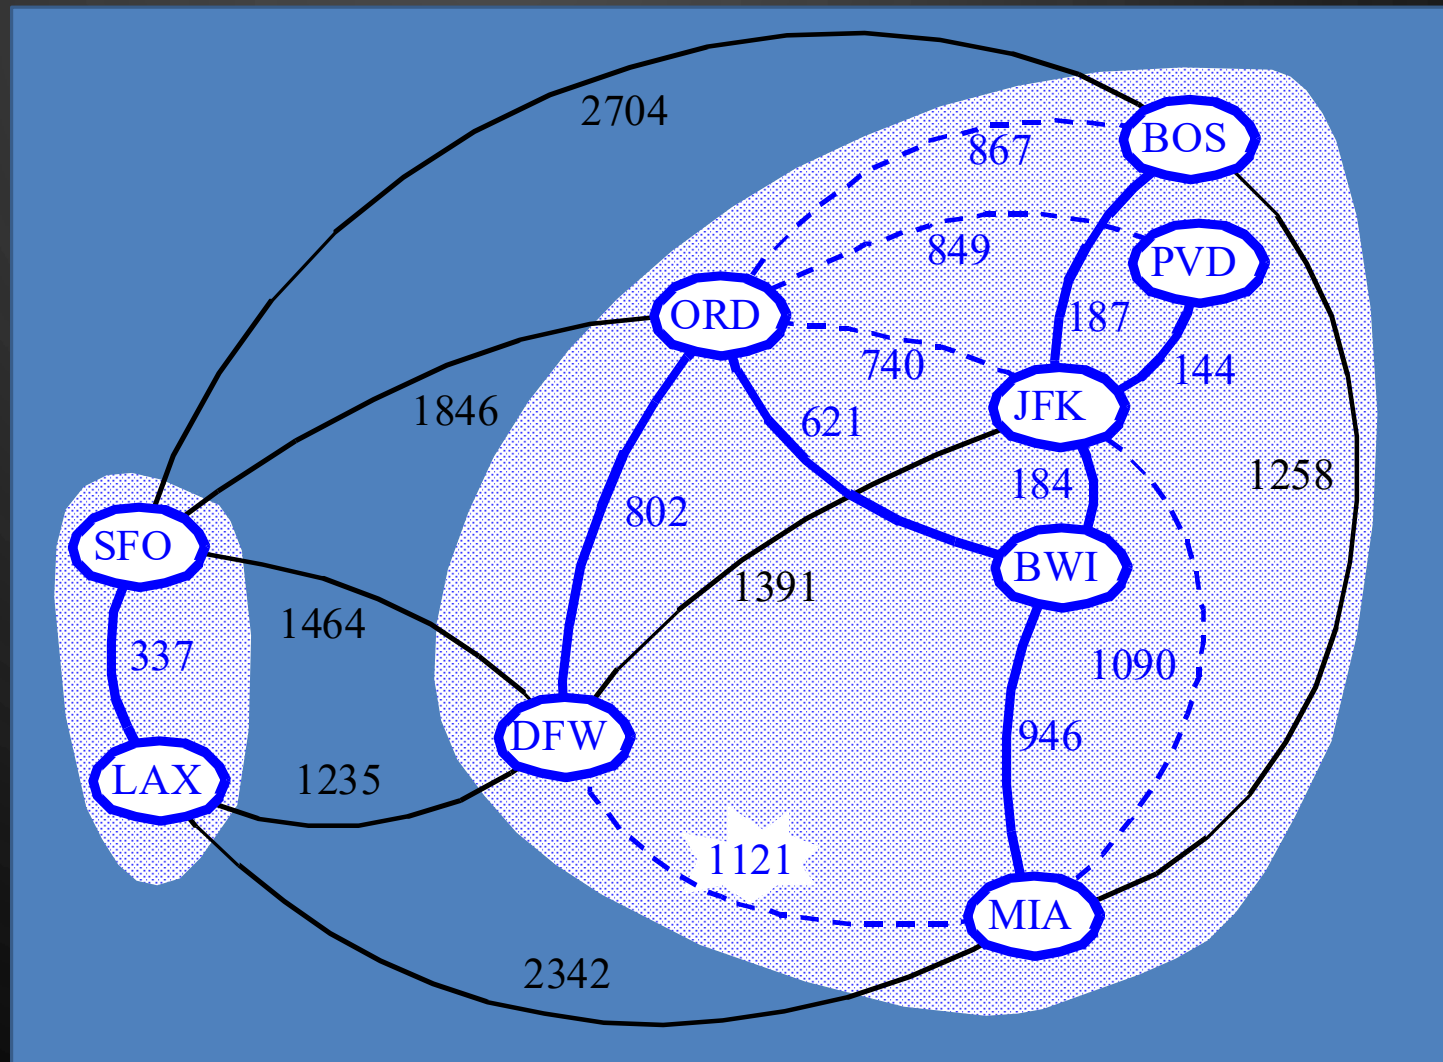
EXAMPLE




EXAMPLE



EXAMPLE



EXAMPLE

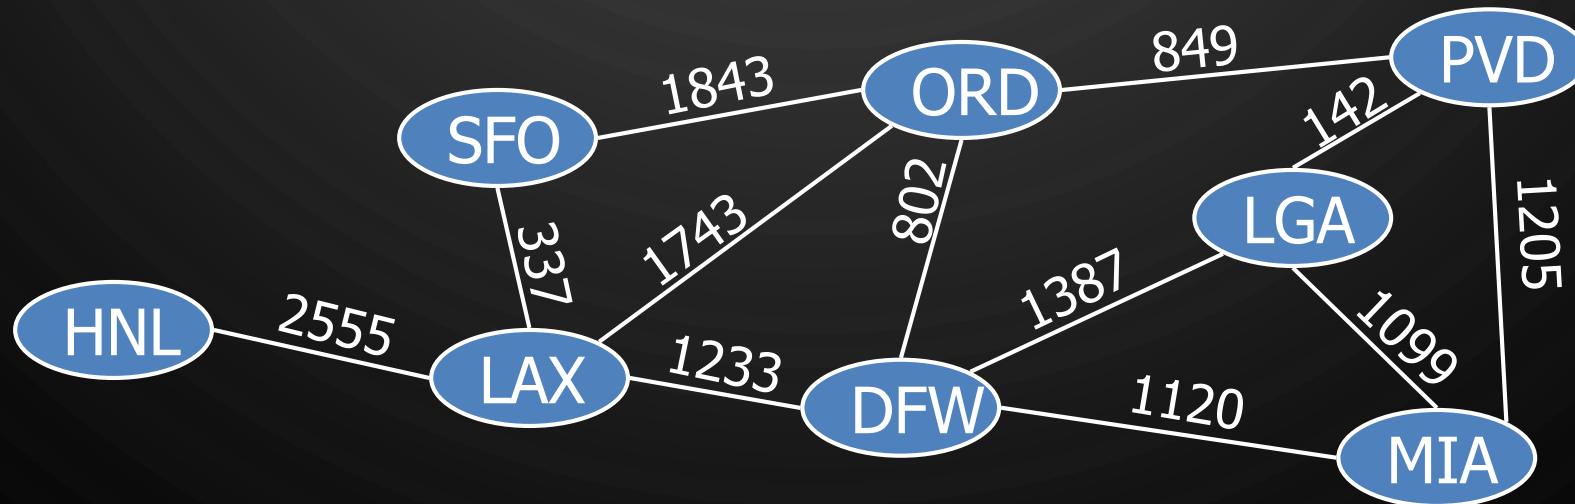


EXERCISE


KRUSKAL'S MST ALGORITHM



- Show how Kruskal's MST algorithm works on the following graph.
 - Show how the MST evolves in each iteration (a separate figure for each iteration).



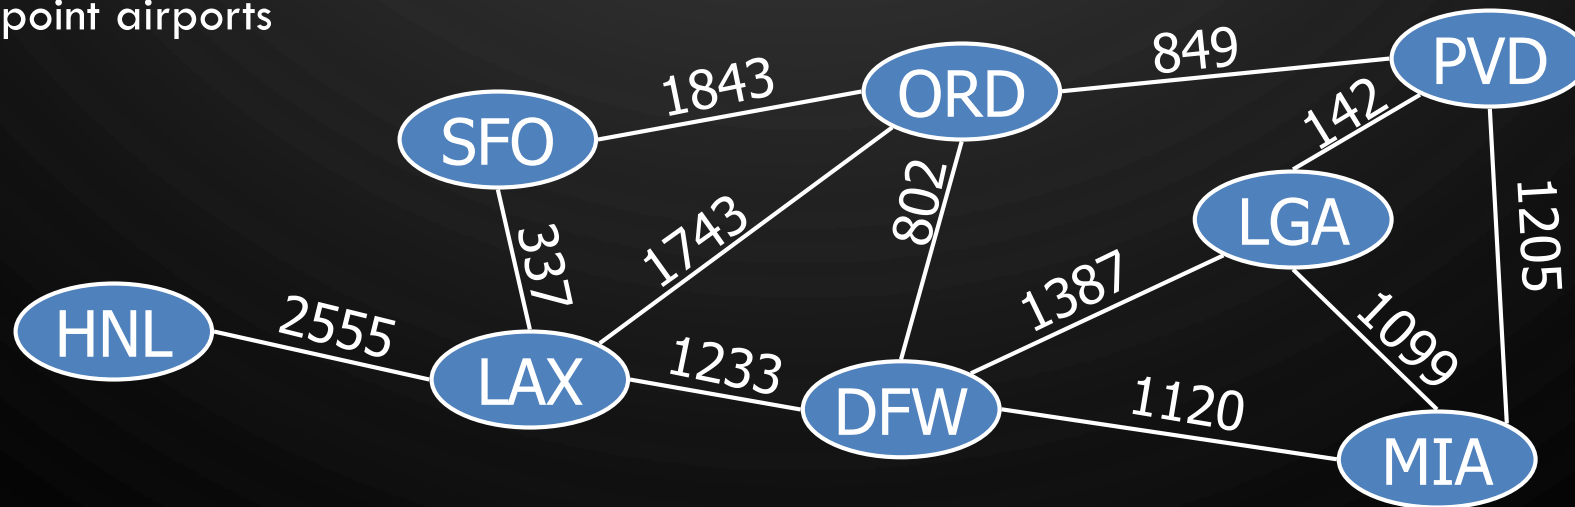
SHORTEST PATHS



WEIGHTED GRAPHS



- In a weighted graph, each edge has an associated numerical value, called the weight of the edge
- Edge weights may represent, distances, costs, etc.
- Example:
 - In a flight route graph, the weight of an edge represents the distance in miles between the endpoint airports



SHORTEST PATH PROBLEM



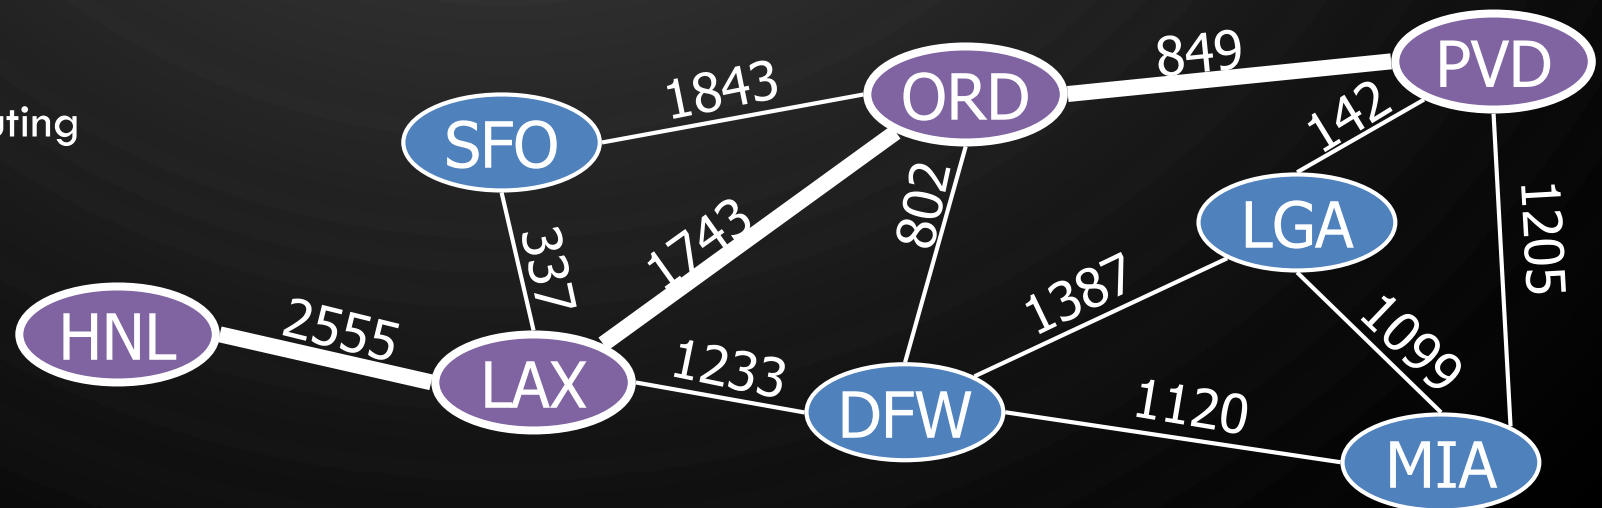
- Given a weighted graph and two vertices u and v , we want to find a path of minimum total weight between u and v .
 - Length of a path is the sum of the weights of its edges.

- Example:

- Shortest path between Providence and Honolulu

- Applications

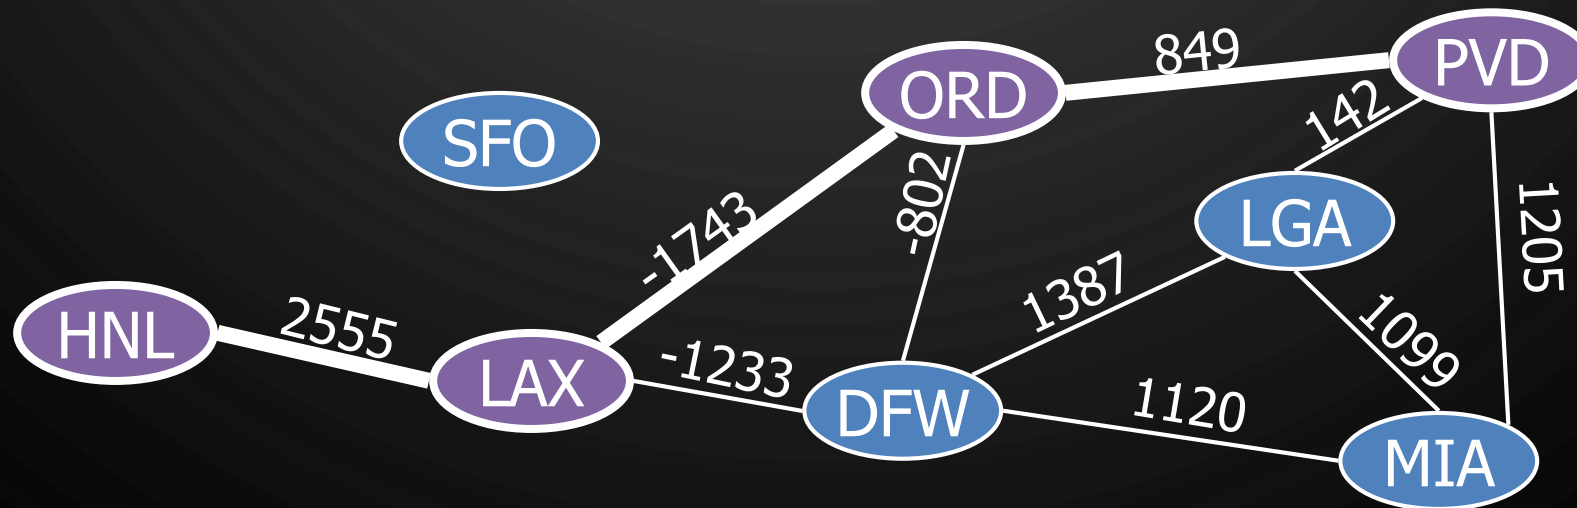
- Internet packet routing
 - Flight reservations
 - Driving directions



SHORTEST PATH PROBLEM



- If there is no path from v to u , we denote the distance between them by $d(v, u) = \infty$
- What if there is a negative-weight cycle in the graph?



SHORTEST PATH PROPERTIES



- Property 1:


- A subpath of a shortest path is itself a shortest path

- Property 2:

- There is a tree of shortest paths from a start vertex to all the other vertices

- Example:

- Tree of shortest paths from Providence



DIJKSTRA'S ALGORITHM





- The distance of a vertex v from a vertex s is the length of a shortest path between s and v
- Dijkstra's algorithm computes the distances of all the vertices from a given start vertex s (single-source shortest paths)
- Assumptions:
 - The graph is connected
 - The edges are undirected
 - The edge weights are nonnegative
- Extremely similar to Prim-Jarnik's MST Algorithm
- We grow a "cloud" of vertices, beginning with s and eventually covering all the vertices
- We store with each vertex v a label $D[v]$ representing the distance of v from s in the subgraph consisting of the cloud and its adjacent vertices
- The label $D[v]$ is initialized to positive infinity
- At each step
 - We add to the cloud the vertex u outside the cloud with the smallest distance label, $D[v]$
 - We update the labels of the vertices adjacent to u , in a process called edge relaxation

EDGE RELAXATION

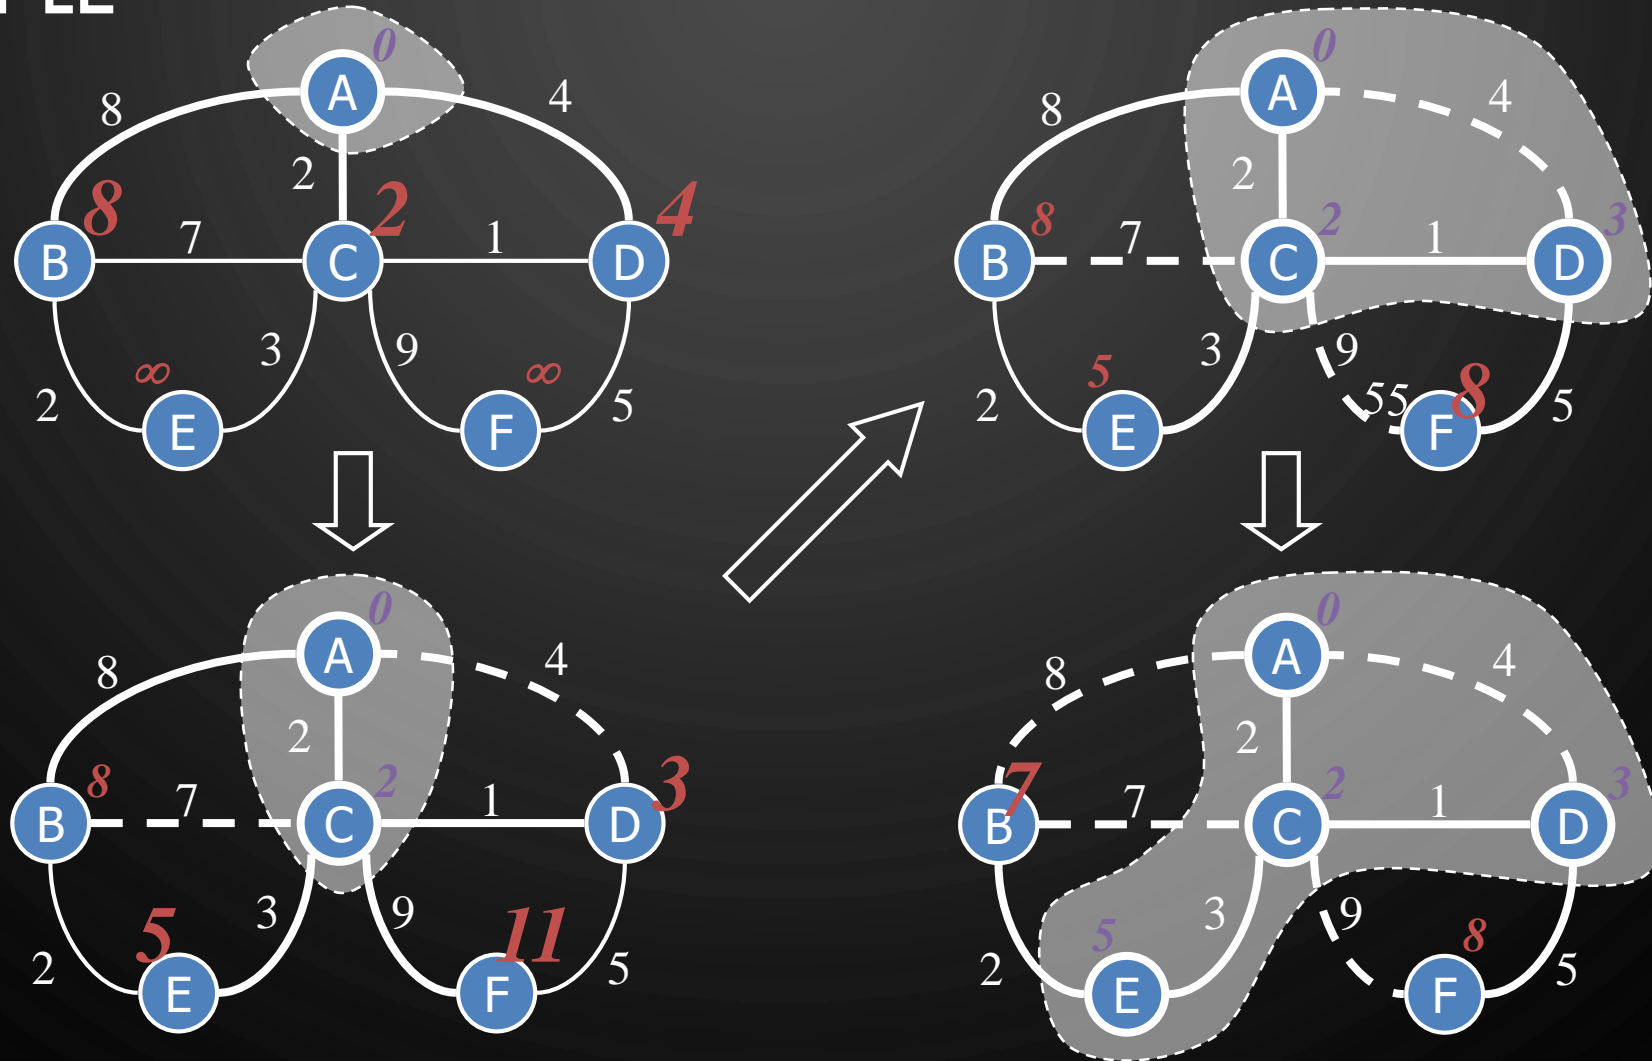


- Consider an edge $e = (u, z)$ such that
 - u is the vertex most recently added to the cloud
 - z is not in the cloud
- The relaxation of edge e updates distance $D[z]$ as follows:
- $D[z] \leftarrow \min(D[z], D[u] + e.weight())$





1. Pull in one of the vertices with **red labels**
2. The relaxation of edges updates the labels of **LARGER** font size

EXAMPLE



EXAMPLE




EXERCISE


DIJKSTRA'S ALGORITHM



- Show how Dijkstra's algorithm works on the following graph, assuming you start with SFO, i.e., $s = \text{SFO}$.
 - Show how the labels are updated in each iteration (a separate figure for each iteration).



DIJKSTRA'S ALGORITHM



- A locator-based priority queue stores the vertices outside the cloud
 - Key: distance
 - Element: vertex
- We store with each vertex:
 - distance $D[v]$ label
 - locator in priority queue


Algorithm Dijkstras sssp(G, s)

Input: A simple undirected weighted graph G with nonnegative edge weights and a source vertex s

Output: A label $D[v]$ for each vertex v of G , such that $D[u]$ is the length of the shortest path from s to v

1. $D[s] \leftarrow 0; D[v] \leftarrow \infty$ for each vertex $v \neq s$
2. Let priority queue Q contain all the vertices of G using $D[v]$ as the key
3. **while** $\neg Q.empty()$ **do** *// $O(n)$ iterations*
4. *//pull a new vertex u in the cloud*
5. $u \leftarrow Q.removeMin()$ *// $O(\log n)$*
6. **for each** edge $e \in u.incidentEdges()$ **do** *// $O(deg(u))$ iterations*
7. *//relax edge e*
8. $z \leftarrow e.opposite(u)$
9. **if** $D[u] + e.weight() < D[z]$ **then**
10. $D[z] \leftarrow D[u] + e.weight()$
11. $Q.updateKey(z)$ *// $O(\log n)$*

ANALYSIS



- Graph operations
 - Method `incidentEdges` is called once for each vertex
- Label operations
 - We set/get the distance and locator labels of vertex z $O(\deg(z))$ times
 - Setting/getting a label takes $O(1)$ time
- Priority queue operations
 - Each vertex is inserted once into and removed once from the priority queue, where each insertion or removal takes $O(\log n)$ time
 - The key of a vertex in the priority queue is modified at most $\deg(w)$ times, where each key change takes $O(\log n)$ time
- Dijkstra's algorithm runs in $O((n + m) \log n)$ time provided the graph is represented by the adjacency list structure
 - Recall that $\sum_v \deg v = 2m$
- The running time can also be expressed as $O(m \log n)$ since the graph is connected
- The running time can be expressed as a function of n , $O(n^2 \log n)$

EXTENSION




- Using the template method pattern, we can extend Dijkstra's algorithm to return a tree of shortest paths from the start vertex to all other vertices
- We store with each vertex a third label:
 - parent edge in the shortest path tree
- Parents are all initialized to null
- In the edge relaxation step, we update the parent label as well

WHY DIJKSTRA'S ALGORITHM WORKS




- Dijkstra's algorithm is based on the greedy method. It adds vertices by increasing distance.
- Proof by contradiction
 - Suppose it didn't find all shortest distances. Let F be the first wrong vertex the algorithm processed.
 - When the previous node, D , on the true shortest path was considered, its distance was correct.
 - But the edge (D, F) was relaxed at that time!
 - Thus, so long as $D[F] > D[D]$, F 's distance cannot be wrong. That is, there is no wrong vertex.




WHY IT DOESN'T WORK FOR NEGATIVE-WEIGHT EDGES

- Dijkstra's algorithm is based on the greedy method. It adds vertices by increasing distance.
 - If a node with a negative incident edge were to be added late to the cloud, it could mess up distances for vertices already in the cloud.



C's true distance is 1, but it is already in the cloud with $D[C] = 2$!

BELLMAN-FORD ALGORITHM




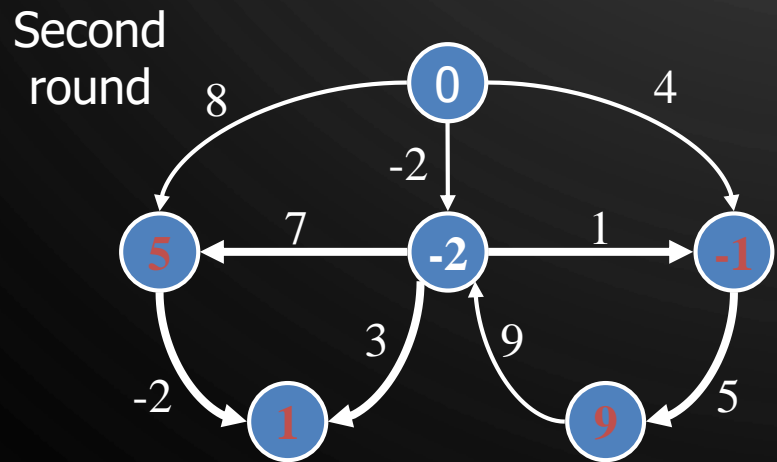


- Works even with negative-weight edges
- Must assume directed edges (for otherwise we would have negative-weight cycles)
- Iteration i finds all shortest paths that use i edges.
- Running time: $O(nm)$
- Can be extended to detect a negative-weight cycle if it exists
 - How?

Algorithm BellmanFord(G, s)

1. Initialize $D[s] \leftarrow 0$ and $D[v] \leftarrow \infty$ for all vertices $v \neq s$
2. **for** $i \leftarrow 1 \dots n - 1$ **do**
3. **for each** $e \in G.\text{edges}()$ **do**
4. //relax edge e
5. $u \leftarrow e.\text{source}()$; $z \leftarrow e.\text{target}()$
6. **if** $D[u] + e.\text{weight}() < D[z]$ **then**
7. $D[z] \leftarrow D[u] + e.\text{weight}()$

- Nodes are labeled with their $D[v]$ values

BELLMAN-FORD EXAMPLE




EXERCISE


BELLMAN-FORD'S ALGORITHM



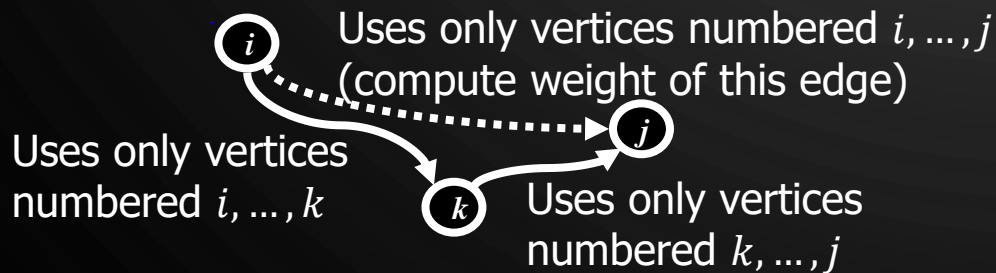
- Show how Bellman-Ford's algorithm works on the following graph, assuming you start with the top node
 - Show how the labels are updated in each iteration (a separate figure for each iteration).



ALL-PAIRS SHORTEST PATHS



- Find the distance between every pair of vertices in a weighted directed graph G
- We can make n calls to Dijkstra's algorithm (if no negative edges), which takes $O(nm \log n)$ time.
- Likewise, n calls to Bellman-Ford would take $O(n^2m)$ time.
- We can achieve $O(n^3)$ time using dynamic programming (similar to the Floyd-Warshall algorithm).



Algorithm AllPairsShortestPath(G)

Input: Graph G with vertices labeled $1, \dots, n$

Output: Distances $D[i, j]$ of shortest path lengths between each pair of vertices

1. **for each** vertex pair (i, j) **do**
2. **if** $i = j$ **then**
3. $D_0[i, i] \leftarrow 0$
4. **else if** $e = (i, j)$ is an edge in G **then**
5. $D_0[i, j] \leftarrow e.\text{weight}()$
6. **else**
7. $D_0[i, j] \leftarrow \infty$
8. **for** $k \leftarrow 1 \dots n$ **do**
9. **for** $i \leftarrow 1 \dots n$ **do**
10. **for** $j \leftarrow 1 \dots n$ **do**
11. $D_k[i, j] \leftarrow \min(D_{k-1}[i, j], D_{k-1}[i, k] + D_{k-1}[k, j])$
12. **return** D_n