Outline

1. Chapter 14: Graphs
   - Graph Data Structures
   - The Shortest Path Algorithms

Graphs

Graphs can represent airlines, electrical circuits, or computer networks.

A Graph G will consist of:

- A set V of **vertices** (nodes, points).
  (*Cities, circuit connections, computers*).
  We will use V to mean the **number** of vertices as well (mathematicians use *cardinality* notation, $|V|$).
- A set E of **edges** (lines connecting vertices).
  (*Air lanes, elements in a circuit, computer connections in a network*).
  We will use E to mean the **number** of edges as well.

Graphs

- A path is a series of edges connecting two vertices.
- In an undirected graph edges are "two-way streets"
- A connected graph is one in which every pair of vertices is connected by a path.
- A complete graph is one in which every pair of vertices is connected by an edge.
- Two vertices are adjacent if there is an edge connecting them.
- A cycle in a directed graph is a loop formed by adjacent vertices.

Graphs

In directed graphs:

- edges are "one-way streets" beginning at one vertex and ending at another.
- in-degree of a vertex = # of edges ending at that vertex.
- out-degree of a vertex = # of edges beginning at that vertex.
- A directed acyclic graph (DAG) is a directed graph containing no cycles.
- Vertex `B` is **adjacent** to vertex `A` if there is an edge from `A` to `B`.
- *Example*: A tree is a special type of DAG.

Graphs

A directed graph example

in-degree of A =

out-degree of C =

in-degree of D =

Is there a cycle?

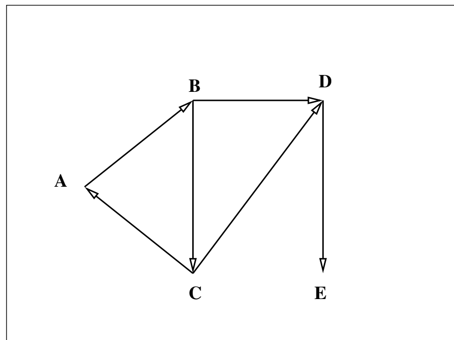

Graphs

- A graph is dense if it has many edges connecting vertices.
- A graph is sparse if it has much less than the maximum possible number of edges.
- The best implementation of a graph depends on how sparse it is.
- Two commonly used data structures to represent graphs are *adjacency matrix* and *adjacency list*.

Representing Graphs

Adjacency Matrices

- An adjacency matrix has rows and columns of zeros and ones. 1 in column i, row j means that an edge connects vertex i with vertex j (i.e. vertices i and j are adjacent).
- An adjacency matrix is used to implement a dense graph.
- It requires $\Theta(V^2)$ time to find all the edges (by checking every entry in the matrix).

Representing Graphs

Adjacency Matrices (Directed Graph)

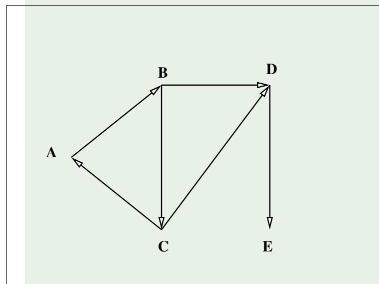

	A	B	C	D	E
A	0	1	0	0	0
B	0	0	1	1	0
C	1	0	0	1	0
D	0	0	0	0	1
E	0	0	0	0	0

Representing Graphs

Adjacency Matrices(Undirected Graph)

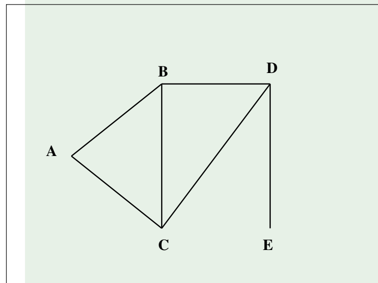

	A	B	C	D	E
A	0	1	1	0	0
B	1	0	1	1	0
C	1	1	0	1	0
D	0	1	1	0	1
E	0	0	0	1	0

The matrix is symmetric (entry at row i, column j is the same as at row j, column i), hence we need only half of the matrix to represent a graph (using diagonal to split it).

Representing Graphs

Adjacency Matrices(Undirected Graph)

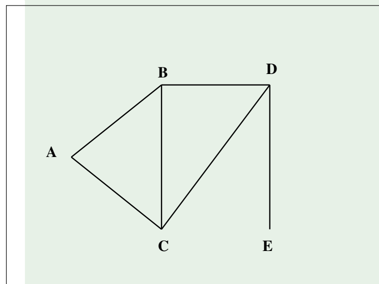

	B	C	D	E
A	1	1	0	0
B		1	1	0
C			1	0
D				1

Adjacency Lists

- An adjacency list gives each vertex an attribute which is a list of all the vertices adjacent to it.
- To represent a sparse graph, an adjacency list is more economical, since it only indicates where the edges are, not where they aren't.
- An adjacency list uses time $\Theta(V * E)$ to find all edges.

Adjacency Lists

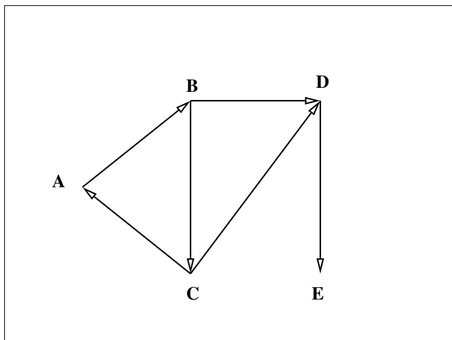

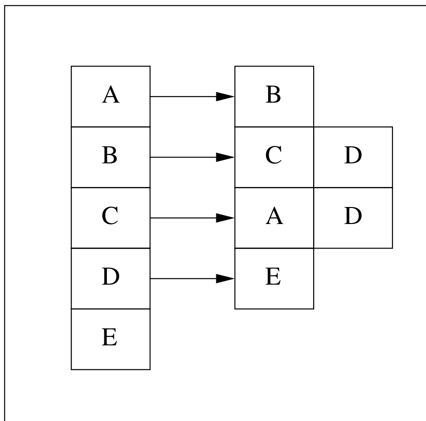

Adjacency Lists

Implementation of adjacency lists in Python:

- A list of lists.
- A dictionary.

Adjacency Lists: using Python list

Let's assume that the graph is weighted, and the weight of each edge is 1, then using Python list we can have the following:

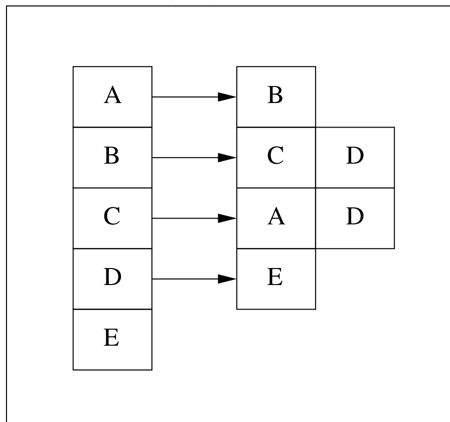


```
 g = [
['A',[('B',1)]],
['B',[('C',1),('D',1)]],
['C',[('A',1),('D',1)]],
['D',[('E',1)]],
['E',[]]]
```

Adjacency Lists: using Python dictionary

Let's assume that the graph is weighted, and the weight of each edge is 1, then using Python dictionary:

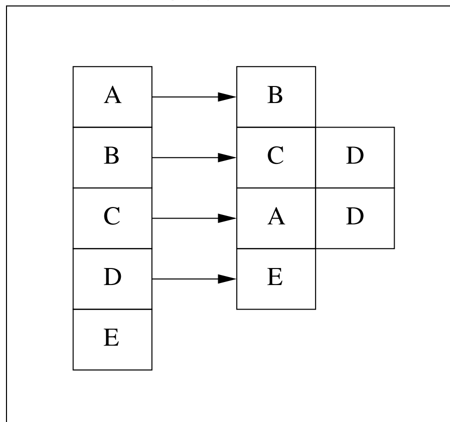


```
g = {
'A':{'B':1},
'B':{'C':1,'D':1},
'C':{'A':1,'D':1},
'D':{'E':1},
'E':{}}
```

Adjacency Lists: C++

Implementation of adjacency lists in C++:

- For static graphs (do not change): a two-dimensional array.
- For dynamic graphs: a list of lists (linked-list implementation).

Adjacency Matrix vs Adjacency List

Adjacency Matrix vs Adjacency List

- graph is dense → the adjacency matrix representation is preferred.
- graph is sparse → the adjacency list representation is preferred.
- Most graphs in real-world applications are sparse, hence the adjacency list representation is more commonly used.
- Using matrix representation to find all edges from a vertex we will need to examine V entries,
- Using list representation to find all edges from a vertex we will need to examine only the actual edges originating from the vertex.

Adjacency Matrix vs Adjacency List

Adjacency Matrix vs Adjacency List

- graph is dense → the adjacency matrix representation is preferred.
- graph is sparse → the adjacency list representation is preferred.
- Most graphs in real-world applications are sparse, hence the adjacency list representation is more commonly used.
- Using matrix representation to find all edges from a vertex we will need to examine V entries,
- Using list representation to find all edges from a vertex we will need to examine only the actual edges originating from the vertex.

Adjacency Matrix vs Adjacency List

Adjacency Matrix vs Adjacency List

- graph is dense → the adjacency matrix representation is preferred.
- graph is sparse → the adjacency list representation is preferred.
- Most graphs in real-world applications are sparse, hence the adjacency list representation is more commonly used.
- Using matrix representation to find all edges from a vertex we will need to examine V entries,
- Using list representation to find all edges from a vertex we will need to examine only the actual edges originating from the vertex.

Adjacency Matrix vs Adjacency List

Adjacency Matrix vs Adjacency List

- graph is dense → the adjacency matrix representation is preferred.
- graph is sparse → the adjacency list representation is preferred.
- Most graphs in real-world applications are sparse, hence the adjacency list representation is more commonly used.
- Using matrix representation to find all edges from a vertex we will need to examine `V` entries,
- Using list representation to find all edges from a vertex we will need to examine only the actual edges originating from the vertex.

Adjacency Matrix vs Adjacency List

Adjacency Matrix vs Adjacency List

- graph is dense $\rightarrow$ the adjacency matrix representation is preferred.
- graph is sparse $\rightarrow$ the adjacency list representation is preferred.
- Most graphs in real-world applications are sparse, hence the adjacency list representation is more commonly used.
- Using matrix representation to find all edges from a vertex we will need to examine `V` entries,
- Using list representation to find all edges from a vertex we will need to examine only the actual edges originating from the vertex.

Shortest Paths

Shortest path algorithms

Determining the shortest path between two vertices is a common problem for many applications.

Example: Maps of roads can be represented using graph (roads and intersections). Let's find a shortest route from one intersection to another, - it is a *weighted shortest path problem*.

Another case: a town/city where the length of each block is approximately the same. In this case we can ignore the length of the block and concentrate on minimization of the number of blocks to pass from one intersection to another (i.e. minimize the sum of edges with the same weight), - this is an *unweighted shortest path problem*.

Shortest or fastest route is a problem that must be solved every day by shipping and delivery companies.

Shortest Paths

Shortest path algorithms

Determining the shortest path between two vertices is a common problem for many applications.

Example: Maps of roads can be represented using graph (roads and intersections). Let's find a shortest route from one intersection to another, - it is a *weighted shortest path problem*.

Another case: a town/city where the length of each block is approximately the same. In this case we can ignore the length of the block and concentrate on minimization of the number of blocks to pass from one intersection to another (i.e. minimize the sum of edges with the same weight), - this is an *unweighted shortest path problem*.

Shortest or fastest route is a problem that must be solved every day by shipping and delivery companies.

Shortest Paths

Shortest path algorithms

Determining the shortest path between two vertices is a common problem for many applications.
Example: Maps of roads can be represented using graph (roads and intersections). Let's find a shortest route from one intersection to another, - it is a *weighted shortest path problem*.

Another case: a town/city where the length of each block is approximately the same. In this case we can ignore the length of the block and concentrate on minimization of the number of blocks to pass from one intersection to another (i.e. minimize the sum of edges with the same weight), - this is an *unweighted shortest path problem*.

Shortest or fastest route is a problem that must be solved every day by shipping and delivery companies.

Shortest Paths

Shortest path algorithms

Determining the shortest path between two vertices is a common problem for many applications.
Example: Maps of roads can be represented using graph (roads and intersections). Let's find a shortest route from one intersection to another, - it is a *weighted shortest path problem*.

Another case: a town/city where the length of each block is approximately the same. In this case we can ignore the length of the block and concentrate on minimization of the number of blocks to pass from one intersection to another (i.e. minimize the sum of edges with the same weight), - this is an *unweighted shortest path problem*.

Shortest or fastest route is a problem that must be solved every day by shipping and delivery companies.

The Unweighted Shortest Path (BFS)

This algorithm is usually referred to as a *Breadth First Search*.

It works on both directed and undirected graphs.

When using adjacency list representation with undirected graphs, each edge must appear in both lists.

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

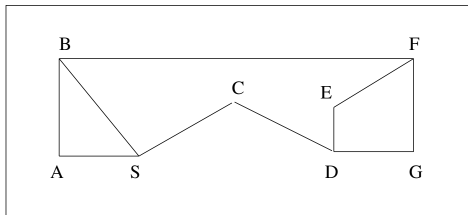


```
queue:
v:
w:
------------------------------------
|   | S | A | B | C | D | E | F | G |
------------------------------------
|par|   |   |   |   |   |   |   |   |
------------------------------------
|dis|   |   |   |   |   |   |   |   |
------------------------------------
```

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

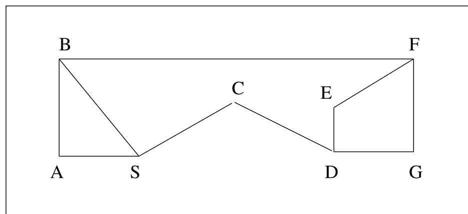

queue:
v:
w:

	S	A	B	C	D	E	F	G
par	-	N	N	N	N	N	N	N
dis								

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

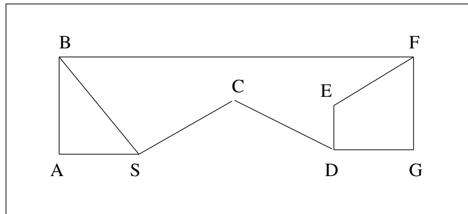

queue:
v:
w:

	S	A	B	C	D	E	F	G
par	-	N	N	N	N	N	N	N
dis	0							

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

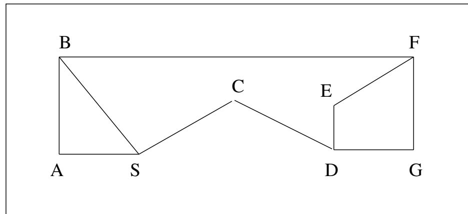

queue: S
v:
w:

	S	A	B	C	D	E	F	G
par	-	N	N	N	N	N	N	N
dis	0							

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

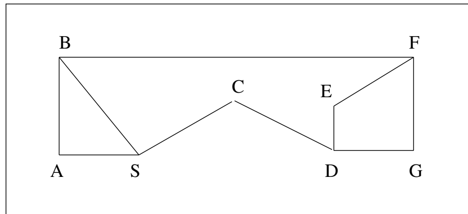


```
queue:S
v:
w:
-------------------------------------
|   | S | A | B | C | D | E | F | G |
-------------------------------------
|par| - | N | N | N | N | N | N | N |
-------------------------------------
|dis| 0 |   |   |   |   |   |   |   |
-------------------------------------
```

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

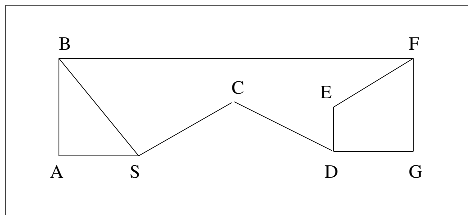

queue:
v: S
w:

	S	A	B	C	D	E	F	G
par	-	N	N	N	N	N	N	N
dis	0							

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

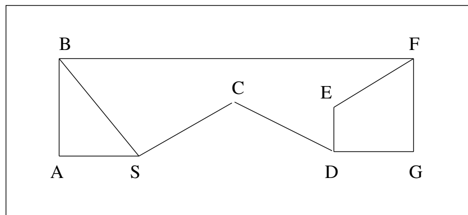

queue:
v: S
w: A

	S	A	B	C	D	E	F	G
par	-	N	N	N	N	N	N	N
dis	0							

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

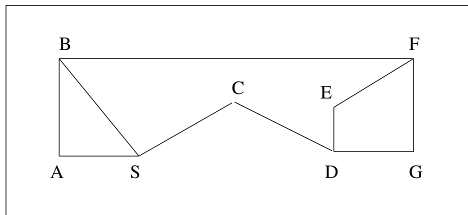


```
queue:
v:  S
w:  A
------------------------------------
|   | S | A | B | C | D | E | F | G |
------------------------------------
|par| - | N | N | N | N | N | N | N |
------------------------------------
|dis| 0 |   |   |   |   |   |   |   |
------------------------------------
```

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent ‘None’.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w’s parent is None:
         set w’s parent to v
         set w’s distance to
v’s distance + 1
         insert w into queue
```

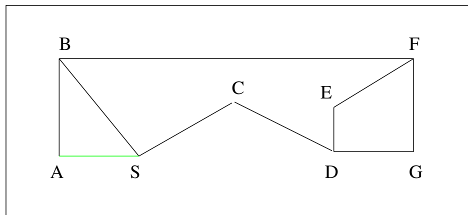

queue: A
v: S
w: A

	S	A	B	C	D	E	F	G
par	-	S	N	N	N	N	N	N
dis	0	1						

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

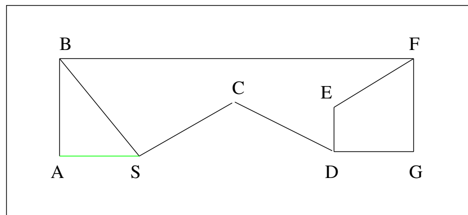

queue: A
v: S
w: B

	S	A	B	C	D	E	F	G
par	-	S	N	N	N	N	N	N
dis	0	1						

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

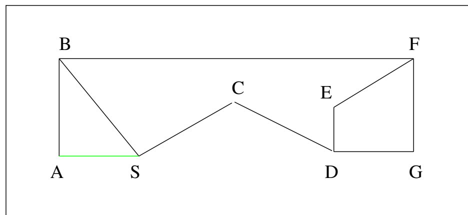

queue: A
v: S
w: B

	S	A	B	C	D	E	F	G
par	-	S	N	N	N	N	N	N
dis	0	1						

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

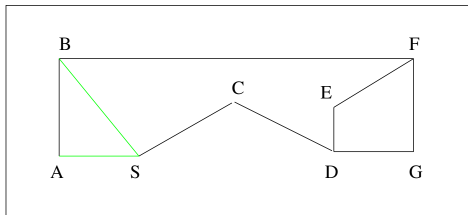

queue: BA
v: S
w: B

	S	A	B	C	D	E	F	G
par	-	S	S	N	N	N	N	N
dis	0	1	1					

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent ‘None’.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w’s parent is None:
         set w’s parent to v
         set w’s distance to
v’s distance + 1
         insert w into queue
```

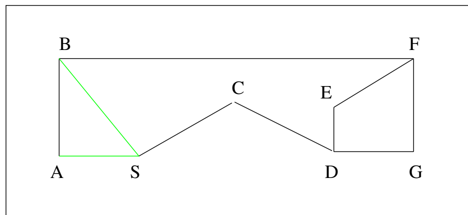

queue: BA
v: S
w: C

	S	A	B	C	D	E	F	G
par	-	S	S	N	N	N	N	N
dis	0	1	1					

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

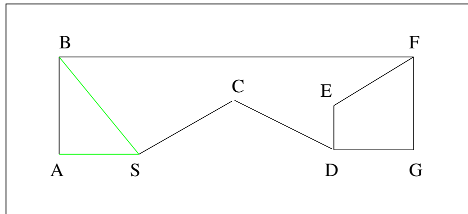

queue: BA
v: S
w: C

	S	A	B	C	D	E	F	G
par	-	S	S	N	N	N	N	N
dis	0	1	1					

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

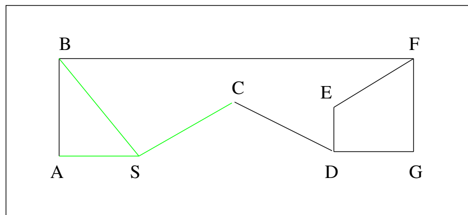

queue: CBA
v: S
w: C

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	N	N
dis	0	1	1	1				

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

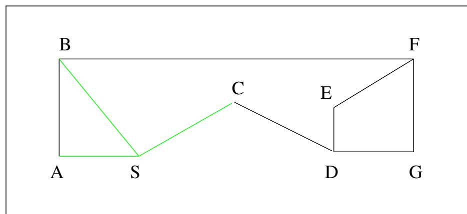

queue: CB
v: A
w:

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	N	N
dis	0	1	1	1				

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

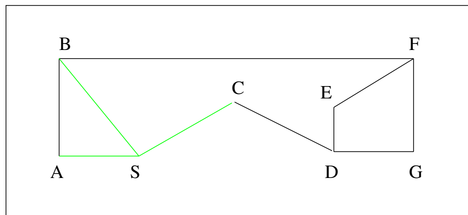

queue: CB
v: A
w: B

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	N	N
dis	0	1	1	1				

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

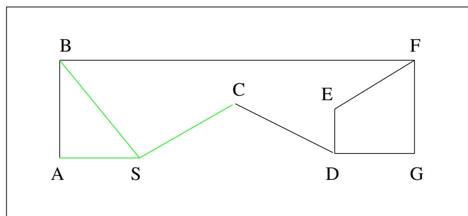

queue: CB
v: A
w: S

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	N	N
dis	0	1	1	1				

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

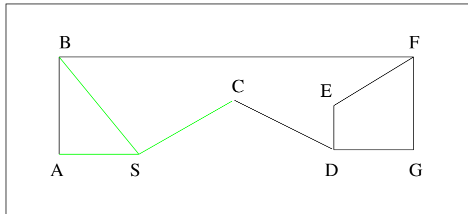

queue: C

v: B

w:

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	N	N
dis	0	1	1	1				

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

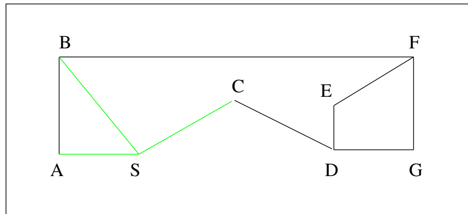

queue: C
v: B
w: F

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	N	N
dis	0	1	1	1				

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

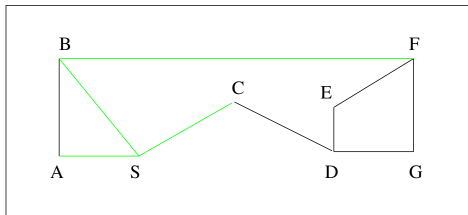

queue: FC
v: B
w: F

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	B	N
dis	0	1	1	1			2	

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

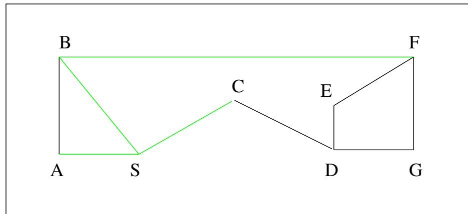

queue: F

v: C

w:

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	B	N
dis	0	1	1	1			2	

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

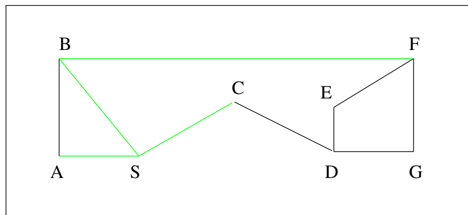

queue: F
v: C
w: D

	S	A	B	C	D	E	F	G
par	-	S	S	S	N	N	B	N
dis	0	1	1	1			2	

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

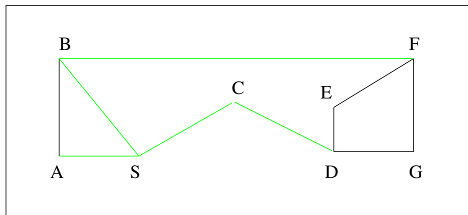

queue: DF
v: B
w: D

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	N	B	N
dis	0	1	1	1	2		2	

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

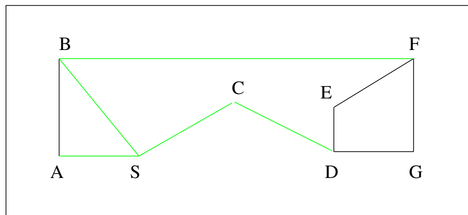

queue: D
v: F
w:

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	N	B	N
dis	0	1	1	1	2		2	

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

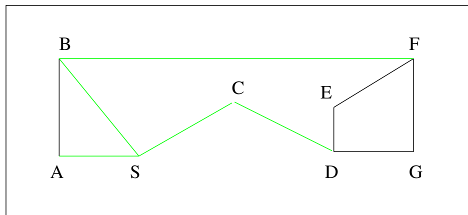

queue: D
v: F
w: E

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	N	B	N
dis	0	1	1	1	2		2	

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

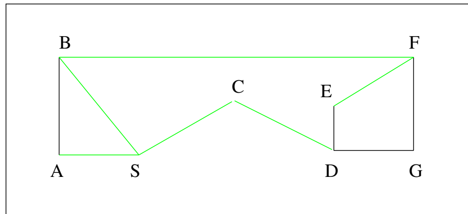

queue: ED
v: F
w: E

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	F	B	N
dis	0	1	1	1	2	3	2	

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

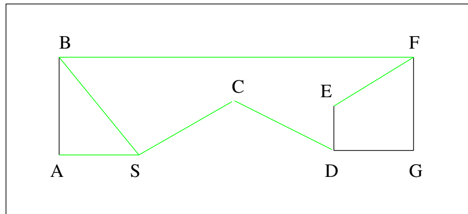

queue: ED
v: F
w: G

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	F	B	N
dis	0	1	1	1	2	3	2	

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

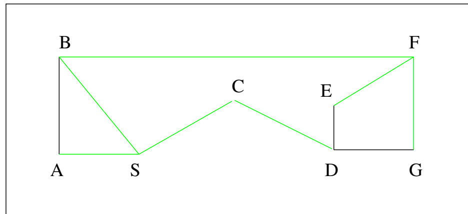

queue: GED
v: F
w: G

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	F	B	F
dis	0	1	1	1	2	3	2	3

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

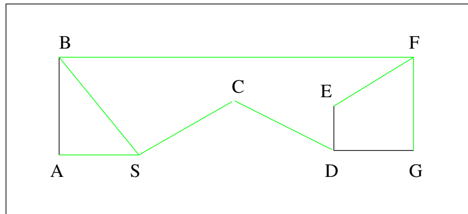

queue: GE
v: D
w:

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	F	B	F
dis	0	1	1	1	2	3	2	3

The Unweighted Shortest Path (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

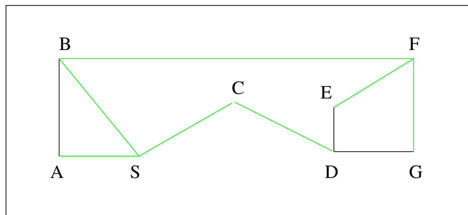

queue: G
v: E
w:

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	F	B	F
dis	0	1	1	1	2	3	2	3

THE UNWEIGHTED SHORTEST PATH (BFS)

```
set all vertices to have
parent 'None'.
set distance for source
vertex to 0
Insert the source vertex
into the queue.
while the queue is not
empty:
   dequeue a vertex v
   for each vertex w
adjacent to v:
      if w's parent is None:
         set w's parent to v
         set w's distance to
v's distance + 1
         insert w into queue
```

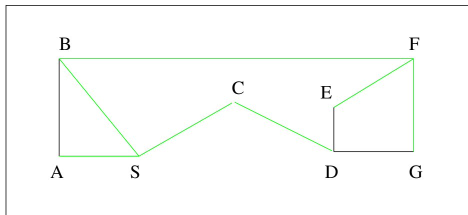


```
queue:
v:  G
w:
```

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	F	B	F
dis	0	1	1	1	2	3	2	3

The Unweighted Shortest Path (BFS)

To find the shortest path from source vertex S to any given vertex:
start with the specified vertex, move backward by following the parent vertices until we reach S

For example:
path from S to E is:
S → B → F → E

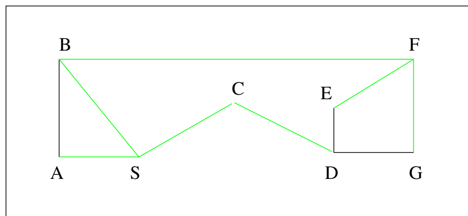

queue:
v: G
w:

	S	A	B	C	D	E	F	G
par	-	S	S	S	C	F	B	F
dis	0	1	1	1	2	3	2	3

Efficiency of BFS

We have two nested loops.

The outer `while` loop runs `V` times (each vertex is inserted into the queue exactly once, and removed through the while loop).

The inner `for` loop runs varies depending on how many adjacent vertices each vertex has (when we use adjacency list representation).

In an undirected graph, each edge is processed twice (once for each direction) and in directed graph, each edge is processed once, during the entire execution of the loop.

All the other steps require a constant time.

Hence, the running time of the algorithm is $\Theta(V + E)$.

This is a common pattern for graph algorithms. Any algorithm that processes each edge and each vertex in a constant number of times with all other operations being constant will have this run time.

Efficiency of BFS

We have two nested loops.
The outer `while` loop runs `V` times (each vertex is inserted into the queue exactly once, and removed through the while loop).

The inner `for` loop runs varies depending on how many adjacent vertices each vertex has (when we use adjacency list representation).

In an undirected graph, each edge is processed twice (once for each direction) and in directed graph, each edge is processed once, during the entire execution of the loop.

All the other steps require a constant time.

Hence, the running time of the algorithm is $\Theta(V + E)$.

This is a common pattern for graph algorithms. Any algorithm that processes each edge and each vertex in a constant number of times with all other operations being constant will have this run time.

Efficiency of BFS

We have two nested loops.
The outer `while` loop runs `V` times (each vertex is inserted into the queue exactly once, and removed through the while loop).

The inner `for` loop runs varies depending on how many adjacent vertices each vertex has (when we use adjacency list representation).

In an undirected graph, each edge is processed twice (once for each direction) and in directed graph, each edge is processed once, during the entire execution of the loop.

All the other steps require a constant time.

Hence, the running time of the algorithm is $\Theta(V + E)$.

This is a common pattern for graph algorithms. Any algorithm that processes each edge and each vertex in a constant number of times with all other operations being constant will have this run time.

Efficiency of BFS

We have two nested loops.
The outer `while` loop runs `V` times (each vertex is inserted into the queue exactly once, and removed through the while loop).

The inner `for` loop runs varies depending on how many adjacent vertices each vertex has (when we use adjacency list representation).

In an undirected graph, each edge is processed twice (once for each direction) and in directed graph, each edge is processed once, during the entire execution of the loop.

All the other steps require a constant time.

Hence, the running time of the algorithm is $\Theta(V + E)$.

This is a common pattern for graph algorithms. Any algorithm that processes each edge and each vertex in a constant number of times with all other operations being constant will have this run time.

Efficiency of BFS

We have two nested loops.
The outer `while` loop runs `V` times (each vertex is inserted into the queue exactly once, and removed through the while loop).

The inner `for` loop runs varies depending on how many adjacent vertices each vertex has (when we use adjacency list representation).

In an undirected graph, each edge is processed twice (once for each direction) and in directed graph, each edge is processed once, during the entire execution of the loop.

All the other steps require a constant time.

Hence, the running time of the algorithm is $\Theta(V + E)$.

This is a common pattern for graph algorithms. Any algorithm that processes each edge and each vertex in a constant number of times with all other operations being constant will have this run time.

Efficiency of BFS

We have two nested loops.
The outer `while` loop runs `V` times (each vertex is inserted into the queue exactly once, and removed through the while loop).

The inner `for` loop runs varies depending on how many adjacent vertices each vertex has (when we use adjacency list representation).

In an undirected graph, each edge is processed twice (once for each direction) and in directed graph, each edge is processed once, during the entire execution of the loop.

All the other steps require a constant time.

Hence, the running time of the algorithm is $\Theta(V + E)$.

This is a common pattern for graph algorithms. Any algorithm that processes each edge and each vertex in a constant number of times with all other operations being constant will have this run time.

Efficiency of BFS

We have two nested loops.
The outer `while` loop runs `V` times (each vertex is inserted into the queue exactly once, and removed through the while loop).

The inner `for` loop runs varies depending on how many adjacent vertices each vertex has (when we use adjacency list representation).

In an undirected graph, each edge is processed twice (once for each direction) and in directed graph, each edge is processed once, during the entire execution of the loop.

All the other steps require a constant time.

Hence, the running time of the algorithm is $\Theta(V + E)$.

This is a common pattern for graph algorithms. Any algorithm that processes each edge and each vertex in a constant number of times with all other operations being constant will have this run time.

The Weighted Shortest Path (Dijkstra)

Edgar Dijkstra's algorithm:

- set all vertices to have parent 'None'.
- set distance for all vertices to infinity
- set distance for source vertex to 0
- insert all vertices into a priority queue (by distance, smallest first).

while priority queue is not empty:

- dequeue a vertex v with the shortest distance

for each vertex w adjacent to v:

if w's distance > (v's distance + weight of edge v to w:

- set w's parent to v
- set w's distance to v's dist. + weight of edge v to w

The Weighted Shortest Path (Dijkstra)

```
Modified for unweighted graphs:
• set all vertices to have parent ‘None’.
• set distance for source vertex to 0
• Insert the source vertex into the queue.
while the queue is not empty:
   • dequeue a vertex v
   for each vertex w adjacent to v:
      if w’s parent is None:
         • set w’s parent to v
         • set w’s distance to v’s distance + 1
         • insert w into queue
```

The Weighted Shortest Path (Dijkstra)

Example

Use Dijkstra's algorithm for weighted graphs to find shortest paths from source vertex S to other vertices.

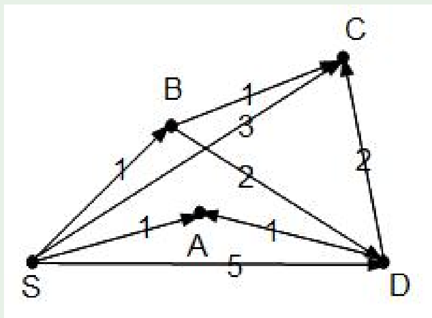

See DijkstrasExample.pdf.

The Weighted Shortest Path (Dijkstra)

Implementation comments

The algorithm asks for a priority queue.

Because we might need to update the vertex's parent and distance information, the priority queue implementation using a binary heap will not work (no efficient way to find a given vertex in the binary heap).

We can use a hash table to map the vertex to its position in the binary heap array/list allowing us to quickly find it, move the item up or down the tree, and then update the hash table to indicate the new position in the heap.

The Weighted Shortest Path (Dijkstra)

Efficiency of Dijkstra's algorithm

Analyzing the efficiency of Dijkstra's algorithm is a little bit more difficult.

Each vertex is removed exactly once from the priority queue. (similar to BFS)

What is different: we extract the vertices from the queue and priorities (along with the parents) might change after a vertex is inserted into the queue.

The Weighted Shortest Path (Dijkstra)

Efficiency of Dijkstra's algorithm

If we use a standard, array-based, list for the PQ: search for the smallest item and removal from PQ will require `V` steps.

adjustment: we can just mark an item as removed (then no need to shift elements). In this case the `while` loop requires `V*V` steps, plus `E` steps the `for` loop executes. Hence, the overall time is $\Theta(V^2 + E)$.

If we use a linked list for PQ: after we find the vertex, the removal is $\Theta(1)$.

The worst-case number of steps is $\frac{V(V-1)}{2}$, so the overall time is still $\Theta(V^2 + E)$

The Weighted Shortest Path (Dijkstra)

Efficiency of Dijkstra's algorithm

If we use the binary heap implementation along with a hash table to track where each item is located in the heap: to remove each item from the PQ and readjust the binary heap $\Theta(\lg V)$

As each edge is processed, the vertex it leads to may have its distance adjusted, requiring it be moved up/down the binary heap. Since the binary heap is a complete tree, $\Theta(\lg V)$ steps may be required to do so.

This gives us an overall running time of $\Theta((V + E) \lg V)$