



Chapter 3

Graphs

CLRS 12-13



Slides by Kevin Wayne.
Copyright © 2005 Pearson-Addison Wesley.
All rights reserved.




3.1 Basic Definitions and Applications

Undirected Graphs

Undirected graph. $G = (V, E)$

- V = nodes.
- E = edges between pairs of nodes.
- Captures pairwise relationship between objects.
- Graph size parameters: $n = |V|$, $m = |E|$.



$V = \{ 1, 2, 3, 4, 5, 6, 7, 8 \}$

$E = \{ (1,2), (1,3), (2,3), (2,4), (2,5),$
 $(3,5), (3,7), (3,8), (4,5), (5,6), (7,8) \}$

$n = 8$

$m = 11$


Some Graph Applications

Graph	Nodes	Edges
transportation	street intersections	highways
communication	computers	fiber optic cables
World Wide Web	web pages	hyperlinks
social	people	relationships
food web	species	predator-prey
software systems	functions	function calls
scheduling	tasks	precedence constraints
circuits	gates	wires

World Wide Web

Web graph.

- Node: web page.
- Edge: hyperlink from one page to another.



9-11 Terrorist Network

Social network graph.


- Node: people.
- Edge: relationship between two people.



Ecological Food Web

Food web graph.

- Node = species.
- Edge = from prey to predator.




Reference: <http://www.twingroves.district96.k12.il.us/Wetlands/Salamander/SalGraphics/salfoodweb.gif>

Graph Representation: Adjacency Matrix

Adjacency matrix. n -by- n matrix with $A_{uv} = 1$ iff (u, v) is an edge.

- Two representations of each edge.
- Space proportional to n^2 .
- Checking if (u, v) is an edge takes $\Theta(1)$ time.
(Checking k pairs (u,v) will cost $\Theta(k)$ time.)
- Identifying all edges takes $\Theta(n^2)$ time.

n = number of vertices



	1	2	3	4	5	6	7	8
1	0	1	1	0	0	0	0	0
2	1	0	1	1	1	0	0	0
3	1	1	0	0	1	0	1	1
4	0	1	0	0	1	0	0	0
5	0	1	1	1	0	1	0	0
6	0	0	0	0	1	0	0	0
7	0	0	1	0	0	0	0	1
8	0	0	1	0	0	0	1	0



Graph Representation: Adjacency List

Adjacency list. Node indexed array of lists.

- Two representations of each edge.
- Space proportional to $m + n$.
- Checking if (u, v) is an edge takes $O(\deg(u))$ time.
(Checking k pairs (u, v) may cost up to $\Theta(kn)$ time.)
- Identifying all edges takes $\Theta(m + n)$ time.

degree = number of neighbours of u

n = number of vertices




Paths and Connectivity

Def. A **path** in an undirected graph $G = (V, E)$ is a sequence P of nodes $v_1, v_2, \dots, v_{k-1}, v_k$ with the property that each consecutive pair v_i, v_{i+1} is joined by an edge in E .


Def. A path is **simple** if all nodes are distinct.

Def. An undirected graph is **connected** if for every pair of nodes u and v , there is a path between u and v .



Cycles

Def. A **cycle** is a path $v_1, v_2, \dots, v_{k-1}, v_k$ in which $v_1 = v_k$, $k > 2$, and the first $k-1$ nodes are all distinct.




cycle $C = 1-2-4-5-3-1$

Trees

Def. An undirected graph is a **tree** if it is connected and does not contain a cycle.

Theorem. Let G be an undirected graph on n nodes. Any two of the following statements imply the third.


- G is connected.
- G does not contain a cycle.
- G has $n-1$ edges.



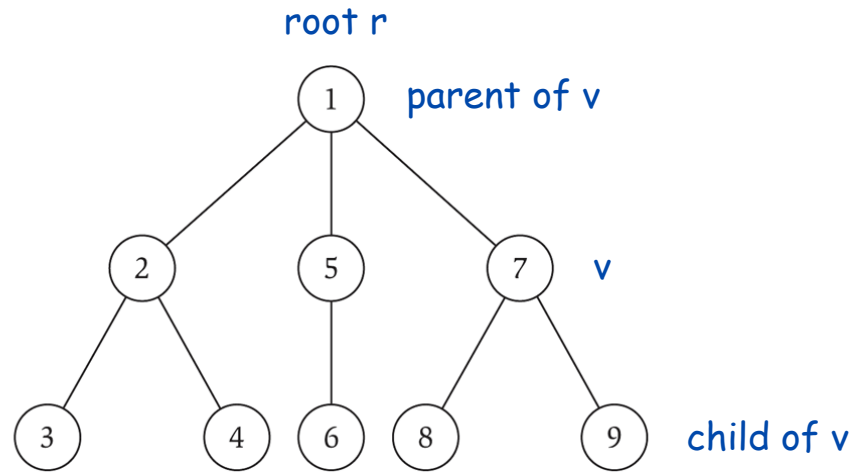
Rooted Trees

Rooted tree. Given a tree T , choose a root node r and orient each edge away from r .

Importance. Models hierarchical structure.




a tree



the same tree, rooted at 1

Phylogeny Trees

Phylogeny trees. Describe evolutionary history of species.



3.2 Graph Traversal


Connectivity

s-t connectivity problem. Given two nodes s and t , is there a path between s and t ?

s-t shortest path problem. Given two nodes s and t , what is the length (number of edges) of the shortest path between s and t ?


Applications.

- Facebook.
- Maze traversal.
- Erdos number.
- Fewest number of hops in a communication network.



Breadth First Search

BFS intuition. Explore outward from s in all possible directions, adding nodes one "layer" at a time.



BFS algorithm.

- $L_0 = \{ s \}$.
- L_1 = all neighbours of L_0 .
- L_2 = all nodes that do not belong to L_0 or L_1 , and that have an edge to a node in L_1 .
- L_{i+1} = all nodes that do not belong to an earlier layer, and that have an edge to a node in L_i .

Theorem. For each i , L_i consists of all nodes at distance exactly i from s . There is a path from s to t iff t appears in some layer.

Breadth First Search

BFS(s):

Set Discovered[s] = true and Discovered[v] = false for all other v

Initialize $L[0]$ to consist of the single element s

Set the layer counter $i=0$

Set the current BFS tree $T=\emptyset$

While $L[i]$ is not empty

 Initialize an empty list $L[i+1]$

 For each node $u \in L[i]$

 Consider each edge (u,v) incident to u

 If Discovered[v] = false then

 Set Discovered[v] = true

 Add edge (u,v) to the tree T

 Add v to the list $L[i+1]$

 Endif



 Endfor

 Increment the layer counter i by one


Endwhile

Breadth First Search


Property. Let T be a BFS tree of $G = (V, E)$, and let (x, y) be an edge of G . Then the level of x and y differ by at most 1.



(a)



(b)



(c)

Breadth First Search: Analysis

Theorem. The above implementation of BFS runs in $O(m + n)$ time if the graph is given by its adjacency list representation.


Pf.

- Easy to prove $O(n^2)$ running time:
 - at most n lists $L[i]$
 - each node occurs on at most one list
 - when we consider node u , there are $\leq n$ incident edges (u, v) , and we spend $O(1)$ processing each edge
- Actually runs in $O(m + n)$ time:
 - when we consider node u , there are $\deg(u)$ incident edges (u, v)
 - total time processing edges is $\sum_{u \in V} \deg(u) = 2m$ ▪

each edge (u, v) is counted exactly twice
in sum: once in $\deg(u)$ and once in $\deg(v)$

Connected Component

Connected component. Find all nodes reachable from s .





Connected component containing node 1 = $\{ 1, 2, 3, 4, 5, 6, 7, 8 \}$.

Flood Fill

Flood fill. Given lime green pixel in an image, change color of entire blob of neighbouring lime pixels to blue.

- Node: pixel.
- Edge: two neighbouring lime pixels.
- Blob: connected component of lime pixels.

recolour lime green blob to blue





Flood Fill

Flood fill. Given lime green pixel in an image, change color of entire blob of neighbouring lime pixels to blue.

- Node: pixel.
- Edge: two neighbouring lime pixels.
- Blob: connected component of lime pixels.

recolour lime green blob to blue



Connected Component

Connected component. Find all nodes reachable from s .


R will consist of nodes to which s has a path

Initially $R = \{s\}$

While there is an edge (u, v) where $u \in R$ and $v \notin R$

 Add v to R

Endwhile



it's safe to add v


Theorem. Upon termination, R is the connected component containing s .

- BFS = explore in order of distance from s .
- DFS = explore in a different way.

Chapter 3

Graphs

CLRS 12-13



Slides by Kevin Wayne.
Copyright © 2005 Pearson-Addison Wesley.
All rights reserved.

