

Terminologies

When we talk about upper bound on running time:

- Logarithmic time: $O(\log n)$
- Linear time: $O(n)$
- Quadratic time $O(n^2)$
- Cubic time $O(n^3)$
- Polynomial time: $O(n^k)$ for some constant k
 - $O(n \log n) \subseteq O(n^{1.1})$. So, an $O(n \log n)$ -time algorithm is also a polynomial time algorithm.
- Exponential time: $O(c^n)$ for some $c > 1$
- Sub-linear time: $o(n)$
- Sub-quadratic time: $o(n^2)$

Goal of Algorithm Design

- Design algorithms to minimize the order of the running time.

Goal of Algorithm Design

- Design algorithms to minimize the order of the running time.
- Using asymptotic analysis allows us to ignore the leading constants and lower order terms

Goal of Algorithm Design

- Design algorithms to minimize the order of the running time.
- Using asymptotic analysis allows us to ignore the leading constants and lower order terms
- Makes our life much easier! (E.g., the leading constant depends on the implementation, compiler and computer architecture of computer.)

Q: Does ignoring the leading constant cause any issues?

- e.g, how can we compare an algorithm with running time $0.1n^2$ with an algorithm with running time $1000n$?

Q: Does ignoring the leading constant cause any issues?

- e.g, how can we compare an algorithm with running time $0.1n^2$ with an algorithm with running time $1000n$?

A:

Q: Does ignoring the leading constant cause any issues?

- e.g, how can we compare an algorithm with running time $0.1n^2$ with an algorithm with running time $1000n$?

A:

- Sometimes yes

Q: Does ignoring the leading constant cause any issues?

- e.g, how can we compare an algorithm with running time $0.1n^2$ with an algorithm with running time $1000n$?

A:

- Sometimes yes
- However, when n is big enough, $1000n < 0.1n^2$

Q: Does ignoring the leading constant cause any issues?

- e.g, how can we compare an algorithm with running time $0.1n^2$ with an algorithm with running time $1000n$?

A:

- Sometimes yes
- However, when n is big enough, $1000n < 0.1n^2$
- For “natural” algorithms, constants are not so big!

Q: Does ignoring the leading constant cause any issues?

- e.g, how can we compare an algorithm with running time $0.1n^2$ with an algorithm with running time $1000n$?

A:

- Sometimes yes
- However, when n is big enough, $1000n < 0.1n^2$
- For “natural” algorithms, constants are not so big!
- So, for reasonably large n , algorithm with lower order running time beats algorithm with higher order running time.

CSE 431/531B: Algorithm Analysis and Design (Fall 2023)

Graph Basics

Lecturer: Kelin Luo

*Department of Computer Science and Engineering
University at Buffalo*

Outline

- 1 Graphs
- 2 Connectivity and Graph Traversal
 - Types of Graphs
- 3 Bipartite Graphs
 - Testing Bipartiteness
- 4 Topological Ordering

Examples of Graphs



Figure: Road Networks



Figure: Internet



Figure: Social Networks






Figure: Transition Graphs

(Undirected) Graph $G = (V, E)$




- V : set of vertices (nodes);
- E : pairwise relationships among V ;
 - (undirected) graphs: relationship is symmetric, E contains subsets of size 2

(Undirected) Graph $G = (V, E)$




- V : set of vertices (nodes);
 - $V = \{1, 2, 3, 4, 5, 6, 7, 8\}$
- E : pairwise relationships among V ;
 - (undirected) graphs: relationship is symmetric, E contains subsets of size 2
 - $E = \{\{1, 2\}, \{1, 3\}, \{2, 3\}, \{2, 4\}, \{2, 5\}, \{3, 5\}, \{3, 7\}, \{3, 8\}, \{4, 5\}, \{5, 6\}, \{7, 8\}\}$

Directed Graph $G = (V, E)$



- V : set of vertices (nodes);
 - $V = \{1, 2, 3, 4, 5, 6, 7, 8\}$
- E : pairwise relationships among V ;
 - **directed** graphs: relationship is asymmetric, E contains ordered pairs


Directed Graph $G = (V, E)$



- V : set of vertices (nodes);
 - $V = \{1, 2, 3, 4, 5, 6, 7, 8\}$
- E : pairwise relationships among V ;
 - **directed** graphs: relationship is asymmetric, E contains ordered pairs
 - $E = \{(1, 2), (1, 3), (3, 2), (4, 2), (2, 5), (5, 3), (3, 7), (3, 8), (4, 5), (5, 6), (6, 5), (8, 7)\}$

Abuse of Notations


- For (undirected) graphs, we often use (i, j) to denote the set $\{i, j\}$.
- We call (i, j) an unordered pair; in this case $(i, j) = (j, i)$.



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

- Social Network : Undirected
- Transition Graph : Directed
- Road Network : Directed or Undirected
- Internet : Directed or Undirected


Representation of Graphs



	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

- Adjacency matrix
 - $n \times n$ matrix, $A[u, v] = 1$ if $(u, v) \in E$ and $A[u, v] = 0$ otherwise
 - A is symmetric if graph is undirected

Representation of Graphs



1: [2] → [3]

6: [5]

2: [1] → [3] → [4] → [5]

7: [3] → [8]

3: [1] → [2] → [5] → [7] → [8]

4: [2] → [5]

8: [3] → [7]

5: [2] → [3] → [4] → [6]


- Adjacency matrix

- $n \times n$ matrix, $A[u, v] = 1$ if $(u, v) \in E$ and $A[u, v] = 0$ otherwise
- A is symmetric if graph is undirected

- Linked lists

- For every vertex v , there is a linked list containing all **neighbors** of v .

Representation of Graphs



1: [2 3]

6: [5]

2: [1 3 4 5]

7: [3 8]

3: [1 2 5 7 8]

8: [3 7]

4: [2 5]

5: [2 3 4 6]

$d : (2, 4, 5, 2, 4, 1, 2, 2)$

- Adjacency matrix

- $n \times n$ matrix, $A[u, v] = 1$ if $(u, v) \in E$ and $A[u, v] = 0$ otherwise
- A is symmetric if graph is undirected

- Linked lists

- For every vertex v , there is a linked list containing all **neighbors** of v .
- When graph is static, can use **array of variant-length arrays**.

Comparison of Two Representations

- Assuming we are dealing with undirected graphs
- n : number of vertices
- m : number of edges, assuming $n - 1 \leq m \leq n(n - 1)/2$
- d_v : number of neighbors of v

	Matrix	Linked Lists
memory usage		
time to check $(u, v) \in E$		
time to list all neighbors of v		

Comparison of Two Representations

- Assuming we are dealing with undirected graphs
- n : number of vertices
- m : number of edges, assuming $n - 1 \leq m \leq n(n - 1)/2$
- d_v : number of neighbors of v

	Matrix	Linked Lists
memory usage	$O(n^2)$	
time to check $(u, v) \in E$		
time to list all neighbors of v		

Comparison of Two Representations

- Assuming we are dealing with undirected graphs
- n : number of vertices
- m : number of edges, assuming $n - 1 \leq m \leq n(n - 1)/2$
- d_v : number of neighbors of v

	Matrix	Linked Lists
memory usage	$O(n^2)$	$O(m)$
time to check $(u, v) \in E$		
time to list all neighbors of v		

Comparison of Two Representations

- Assuming we are dealing with undirected graphs
- n : number of vertices
- m : number of edges, assuming $n - 1 \leq m \leq n(n - 1)/2$
- d_v : number of neighbors of v

	Matrix	Linked Lists
memory usage	$O(n^2)$	$O(m)$
time to check $(u, v) \in E$	$O(1)$	
time to list all neighbors of v		

Comparison of Two Representations

- Assuming we are dealing with undirected graphs
- n : number of vertices
- m : number of edges, assuming $n - 1 \leq m \leq n(n - 1)/2$
- d_v : number of neighbors of v

	Matrix	Linked Lists
memory usage	$O(n^2)$	$O(m)$
time to check $(u, v) \in E$	$O(1)$	$O(d_u)$
time to list all neighbors of v		

Comparison of Two Representations

- Assuming we are dealing with undirected graphs
- n : number of vertices
- m : number of edges, assuming $n - 1 \leq m \leq n(n - 1)/2$
- d_v : number of neighbors of v

	Matrix	Linked Lists
memory usage	$O(n^2)$	$O(m)$
time to check $(u, v) \in E$	$O(1)$	$O(d_u)$
time to list all neighbors of v	$O(n)$	

Comparison of Two Representations

- Assuming we are dealing with undirected graphs
- n : number of vertices
- m : number of edges, assuming $n - 1 \leq m \leq n(n - 1)/2$
- d_v : number of neighbors of v

	Matrix	Linked Lists
memory usage	$O(n^2)$	$O(m)$
time to check $(u, v) \in E$	$O(1)$	$O(d_u)$
time to list all neighbors of v	$O(n)$	$O(d_v)$

Outline

- 1 Graphs
- 2 Connectivity and Graph Traversal
 - Types of Graphs
- 3 Bipartite Graphs
 - Testing Bipartiteness
- 4 Topological Ordering

Connectivity Problem

Input: graph $G = (V, E)$, (using linked lists)
two vertices $s, t \in V$

Output: whether there is a path connecting s to t in G

Connectivity Problem

Input: graph $G = (V, E)$, (using linked lists)
two vertices $s, t \in V$

Output: whether there is a path connecting s to t in G

- Algorithm: starting from s , search for all vertices that are reachable from s and check if the set contains t

Connectivity Problem

Input: graph $G = (V, E)$, (using linked lists)
two vertices $s, t \in V$

Output: whether there is a path connecting s to t in G

- Algorithm: starting from s , search for all vertices that are reachable from s and check if the set contains t
- Breadth-First Search (BFS)

Connectivity Problem

Input: graph $G = (V, E)$, (using linked lists)
two vertices $s, t \in V$

Output: whether there is a path connecting s to t in G


- Algorithm: starting from s , search for all vertices that are reachable from s and check if the set contains t
 - Breadth-First Search (BFS)
 - Depth-First Search (DFS)

Breadth-First Search (BFS)

- Build layers $L_0, L_1, L_2, L_3, \dots$
- $L_0 = \{s\}$
- L_{j+1} contains all nodes that are not in $L_0 \cup L_1 \cup \dots \cup L_j$ and have an edge to a vertex in L_j


Breadth-First Search (BFS)

- Build layers $L_0, L_1, L_2, L_3, \dots$
- $L_0 = \{s\}$
- L_{j+1} contains all nodes that are not in $L_0 \cup L_1 \cup \dots \cup L_j$ and have an edge to a vertex in L_j




Breadth-First Search (BFS)

- Build layers $L_0, L_1, L_2, L_3, \dots$
- $L_0 = \{s\}$
- L_{j+1} contains all nodes that are not in $L_0 \cup L_1 \cup \dots \cup L_j$ and have an edge to a vertex in L_j




Breadth-First Search (BFS)

- Build layers $L_0, L_1, L_2, L_3, \dots$
- $L_0 = \{s\}$
- L_{j+1} contains all nodes that are not in $L_0 \cup L_1 \cup \dots \cup L_j$ and have an edge to a vertex in L_j



Breadth-First Search (BFS)

- Build layers $L_0, L_1, L_2, L_3, \dots$
- $L_0 = \{s\}$
- L_{j+1} contains all nodes that are not in $L_0 \cup L_1 \cup \dots \cup L_j$ and have an edge to a vertex in L_j





Implementing BFS using a Queue

BFS(s)



- 1: $head \leftarrow 1, tail \leftarrow 1, queue[1] \leftarrow s$
- 2: mark s as “visited” and all other vertices as “unvisited”
- 3: **while** $head \leq tail$ **do**
- 4: $v \leftarrow queue[head], head \leftarrow head + 1$
- 5: **for** all neighbors u of v **do**
- 6: **if** u is “unvisited” **then**
- 7: $tail \leftarrow tail + 1, queue[tail] = u$
- 8: mark u as “visited”

- Running time: $O(n + m)$.



Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue





Example of BFS via Queue



Example of BFS via Queue



Example of BFS via Queue



Example of BFS via Queue

