

Lecture 12: Graph Algorithms

Tim LaRock

larock.t@northeastern.edu

bit.ly/cs3000syllabus

Business


Homework 2

Midterm

Homework 3

Homework 2

- Grades released earlier – overall good work!
 - Median: 34.75
 - Mean: 32.38
 - Histogram below
- Request regrades directly on gradescope!
 - You can also email me, but gradescope is much better and likely to be faster!
- Quite a few no submissions and partial submissions – if this was a mistake or miscommunication, you need to let me know ASAP!!
 - I will try to be accommodating, but it is your responsibility to make sure you turn things in correctly!



Midterm: Some high level stuff

Overall from initial grading it seems like people did well!

- We are aiming to have your grades by the end of the week. Thank you for your patience!

Pseudocode: high level, abstract description of an algorithm

- Focus on readability and helping understand the algorithm
- Someone reading it should be able to implement it in any language without knowing any other language
 - translation should be english \rightarrow implementation language, not implementation language1 \rightarrow implementation language2, since that requires knowing language1 which defeats the purpose!
- No strict syntax – when faced with options, choose the clearest and most concise that you can think of!
- NOT code

Recursive specification = Algorithm or recursive pseudocode

Recurrence relation = Runtime calculation like $T(n) = T(n/a) + O(f(n))$

Fibonacci Numbers: Recurrence Relation

$$f_n \begin{cases} 0 & \text{if } n = 0 \\ 1 & \text{if } n = 1 \\ f_{n-1} + f_{n-2} & \text{otherwise} \end{cases}$$

Recursive specification

```
Fib(n):  
  If n = 0:  
    return 0  
  ElseIf n = 1:  
    return 1  
  Else:  
    return Fib(n - 1) + Fib(n - 2)
```

What does the recurrence relation $T(n)$ look like?

$$T(0) = 1, T(1) = 1$$

$$T(n) = T(n - 1) + T(n - 2) + 1$$

```
class fibonacci
```

```
{  
  static int fib(int n)  
  {  
    if (n <= 1)  
      return n;  
    return fib(n-1) + fib(n-2);  
  }  
}
```

Not pseudocode!


Homework 3: Graphs 'n stuff

- Will be released after class
- Due next Monday June 1st at Midnight Boston time on Gradescope
 - No more canvas submissions!



- I really tried to make sure this one will be less time consuming!

Graphs: what are they?

Graphs: what are they?



Graphs: what are they?



Graphs

A graph $G = (V, E)$ consists of

- vertices $v \in V$ and
- edges $e \in E$, indicating two vertices u, v are connected

$n = |V|$, number of nodes


$m = |E|$, number of edges

A graph is:


- *directed* if the vertices in each edge are ordered
 - If we are being precise, we say “the edge *from* u to v ”
 - $(u, v) \in E \not\Rightarrow (v, u) \in E$
- *undirected* if its edges are not ordered.
 - “the edge *between* u and v ”
 - $(u, v) \in E \Rightarrow (v, u) \in E$

We will often refer to **vertices** as **nodes** and to **edges** as **links**. Consider these interchangeable!

Undirected



Directed



Question

Assume we have a directed graph $G = (V, E)$ that is *simple*, meaning there is at most one of each possible edge and no self-loops.


What is the maximum size of the set of edges E ?

$$n(n-1)$$


What about in an undirected graph?

$$\frac{n(n-1)}{2}$$

Undirected




Directed



$$E = n-1$$

Proof by contradiction


Given a simple directed graph $G = (V, E)$ with $n = |V|$ nodes, we want to prove that the maximum size of the edge set E is $|E| = n \cdot (n - 1)$.



Proof by contradiction

Given a simple directed graph $G = (V, E)$ with $n = |V|$ nodes, we want to prove that the maximum size of the edge set E is $|E| = n \cdot (n - 1)$.


Assume for contradiction that we have a graph with $|E| > n \cdot (n - 1)$.



Proof by contradiction

Given a simple directed graph $G = (V, E)$ with $n = |V|$ nodes, we want to prove that the maximum size of the edge set E is $|E| = n \cdot (n - 1)$.


Assume for contradiction that we have a graph with $|E| > n \cdot (n - 1)$. This implies that there is some node u with more than $(n - 1)$ neighbors.



Proof by contradiction

Given a simple directed graph $G = (V, E)$ with $n = |V|$ nodes, we want to prove that the maximum size of the edge set E is $|E| = n \cdot (n - 1)$.


Assume for contradiction that we have a graph with $|E| > n \cdot (n - 1)$. This implies that there is some node u with more than $(n - 1)$ neighbors. Since there are only n nodes in G , this implies that either u connects to some node twice, or it connects to itself.



Proof by contradiction


Given a simple directed graph $G = (V, E)$ with $n = |V|$ nodes, we want to prove that the maximum size of the edge set E is $|E| = n \cdot (n - 1)$.

Assume for contradiction that we have a graph with $|E| > n \cdot (n - 1)$. This implies that there is some node u with more than $(n - 1)$ neighbors. Since there are only n nodes in G , this implies that either u connects to some node twice, or it connects to itself. But we have assumed that G is simple. Contradiction!




Proof by contradiction steps

1. State the claim and all assumptions
2. Assume we have an example where it is not true
3. Show that this cannot be the case given the assumptions we made




Proof by contradiction steps

1. State the claim and all assumptions
 - Given a **simple directed graph** $G = (V, E)$ with $n = |V|$ nodes, we want to prove that **the maximum size of the edge set E is $|E| = n \cdot (n - 1)$.**
2. Assume we have an example where it is not true
3. Show that this cannot be the case given the assumptions we made




Proof by contradiction steps

1. State the claim and all assumptions
 - Given a **simple directed graph** $G = (V, E)$ with $n = |V|$ nodes, we want to prove that **the maximum size of the edge set E is $|E| = n \cdot (n - 1)$.**
2. Assume we have an example where it is not true
 - Assume for contradiction that we have a graph with $|E| > n \cdot (n - 1)$.
3. Show that this cannot be the case given the assumptions we made




Proof by contradiction steps

1. State the claim and all assumptions
 - Given a **simple directed graph** $G = (V, E)$ with $n = |V|$ nodes, we want to prove that **the maximum size of the edge set E is $|E| = n \cdot (n - 1)$.**
2. Assume we have an example where it is not true
 - Assume for contradiction that we have a graph with $|E| > n \cdot (n - 1)$.
3. Show that this cannot be the case given the assumptions we made
 - This implies that there is some node u with more than $(n - 1)$ neighbors...But we have assumed that G is simple. Contradiction!




Data Structures for Graphs

Undirected



Data Structures for Graphs



Data Structures for Graphs


Edgelist: A list of tuples (u, v) representing the edges in a graph G

- Advantage: Very simple to interpret
- Disadvantages:
 - Edge lookup/insertion/deletion is $O(m)$

Edgelist

(1,2)
(2,3)
(3,4)
(4,5)
(5,6)
(6,3)
(6,7)
(7,3)

Undirected



$$m = |E|$$

Data Structures for Graphs

Edgelist: A list of tuples (u, v) representing the edges in a graph G

- Advantage: Very simple to interpret
- Disadvantages:
 - Edge lookup/insertion/deletion is $O(m)$

Adjacency List: A list of lists where the first item is a node u and all items in the list are connected to u

- Advantages:
 - Stores same information as edgelist
 - Edge lookup/insertion/deletion can be as fast as $O(n)$
- Disadvantage: stores redundant info for undirected graphs


Edgelist

(1,2)
(2,3)
(3,4)
(4,5)
(5,6)
(6,3)
(6,7)
(7,3)

Adjacency List

[1, 2]
[2, 1, 3]
[3, 4, 6, 7]
[4, 3, 5]
[5, 4, 6]
[6, 3, 5, 7]
[7, 3, 6]

Undirected



Data Structures for Graphs

Edgelist: A list of tuples (u, v) representing the edges in a graph G

- Advantage: Very simple to interpret
- Disadvantages:
 - Edge lookup/insertion/deletion is $O(m)$

Adjacency List: A list of lists where the first item is a node u and all items in the list are connected to u

- Advantages:
 - Stores same information as edgelist
 - Edge lookup/insertion/deletion can be as fast as $O(n)$
- Disadvantage: stores redundant info for undirected graphs

Adjacency Matrix: A matrix $A[1..n, 1..n]$ where each entry $A[i, j]$ is 1 if an edge exists between nodes i and j and 0 otherwise


- Advantages:
 - Simple way to represent dense graphs (many entries 1)
 - Edge lookup/insertion/deletion is $O(1)$
 - Spectral graph analysis/linear algebraic operations
- Disadvantages:
 - Wastes space when many entries are 0
 - Stores redundant info for undirected graphs

Edgelist

(1,2)
(2,3)
(3,4)
(4,5)
(5,6)
(6,3)
(6,7)
(7,3)

$(i < j)$

Undirected



Adjacency List

$(u, i) \in E$

[1, 2]
[2, 1, 3]
[3, 4, 6, 7]
[4, 3, 5]
[5, 4, 6]
[6, 3, 5, 7]
[7, 3, 6]

$A[G, i] \neq \emptyset$

$(i, j) \in E$

$A[(i, j)] = 1$

Adjacency Matrix

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

Data Structures for Graphs

All of these data structures can be modified to make computations faster or more space efficient


- Example: Using a lookup table/dictionary to store an adjacency list would let us return the list of neighbors for a node in $O(1)$ time!

How we store a graph is a choice we make for every algorithm we design

- There is no one size fits all! Different problems will call for different data structures.
- Per Erickson: Usually we don't need arbitrary edge lookup, so it doesn't make sense to optimize for that all the time!

Edgelist

(1,2)
(2,3)
(3,4)
(4,5)
(5,6)
(6,3)
(6,7)
(7,3)



Adjacency List


[1, 2]
[2, 1, 3]
[3, 4, 6, 7]
[4, 3, 5]
[5, 4, 6]
[6, 3, 5, 7]
[7, 3, 6]

Adjacency Matrix


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

Paths through graphs

Undirected



Directed


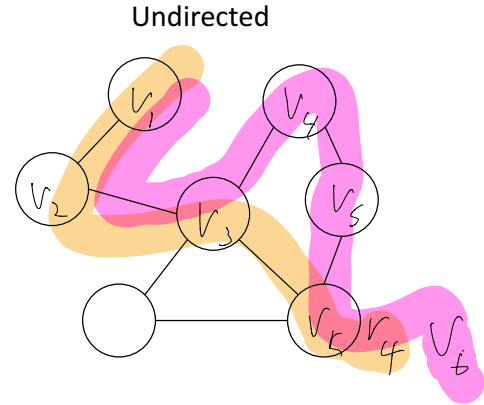


Paths through graphs

A *path* P from vertex v_1 to vertex v_k is an ordered sequence of consecutive edges from E where each node is visited at most once.

$$P = \{(v_1, v_2), (v_2, v_3), \dots, (v_{k-1}, v_k)\}$$

A path visiting k nodes has length $k - 1$, since the length is the number of edges traversed.



Paths through graphs

A *path* P from vertex v_1 to vertex v_k is an ordered sequence of consecutive edges from E where each node is visited at most once.


$$P = \{(v_1, v_2), (v_2, v_3), \dots, (v_{k-1}, v_k)\}$$

A path visiting k nodes has length $k - 1$, since the length is the number of edges traversed.


A *walk* through a graph is similar to a path, but nodes can be visited more than once. A walk is *closed* if it starts and ends with the same node; otherwise it is called *open*.

A *cycle* is a closed walk that visits any node except the first at most once.

Undirected




Directed




Reachability & Connectivity

A node v is *reachable* from a node u if there is a path from u to v .

Undirected



Directed




Reachability & Connectivity


A node v is *reachable* from a node u if there is a path from u to v .

A graph $G = (V, E)$ is *connected* if for every pair of nodes u, v , the node v is reachable from u .

Undirected



Directed



Reachability & Connectivity


A node v is *reachable* from a node u if there is a path from u to v .

A graph $G = (V, E)$ is *connected* if for every pair of nodes u, v , the node v is reachable from u .


In a directed graph, we have two types of connectivity:

- Strongly Connected: there is a path both from u to v and from v to u .
- Weakly Connected: there is a path either from u to v or from v to u

Undirected




Directed



Exploring a graph: Reachability

Assume we have an undirected graph $G = (V, E)$ and we want to determine whether the graph is connected.

We need an algorithm that will tell us whether every node is reachable from every other node.



Idea: traverse the graph edge by edge.

If we can reach every node without restarting, we know the graph is connected!


Exploring a graph: Breadth First Search

Queue is a FIFO data structure

```
BFS( $G = (V, E)$ ):  
   $Q \leftarrow$  empty queue  
  visited  $\leftarrow \emptyset$   
  
  Append node 1 to  $Q$   
  While  $Q$  is not empty:  
     $u \leftarrow$  next node in  $Q$   
    For  $v \in \text{Neighbors}(u)$ :  
      if  $v \notin$  visited  
        Append  $v$  to  $Q$   
    Add  $u$  to visited  
  
  If  $|\text{visited}| = |V|$ :  
    return True  
  Else:  
    return False
```

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

Dequeued (1)




$Q = \langle \langle \rangle, \langle \rangle, \langle \rangle, \langle \rangle, \langle \rangle, \langle \rangle, \langle \rangle \rangle$

Idea: traverse the graph edge by edge.

If we can reach every node without restarting, we know the graph is connected!

Breadth First Search Running time

```
BFS( $G = (V, E)$ ):  
   $Q \leftarrow$  empty queue  
  visited  $\leftarrow \emptyset$   
  
  Append node 1 to  $Q$   
  While  $Q$  is not empty:  
     $u \leftarrow$  next node in  $Q$   
    For  $v \in \text{Neighbors}(u)$ :  
      if  $v \notin$  visited  
        Append  $v$  to  $Q$   
    Add  $u$  to visited  
  
  If |visited| =  $|V|$ :  
    return True  
  Else:  
    return False
```



By definition, we visit every node once, so we immediately have $O(n)$ to start.

At each node, we check if each of its neighbors has been visited already.

Observation: this is the same as visiting every *edge*! Thus we also have $O(m)$.


Therefore, the running time of BFS is $O(n + m)$.

Exploring a graph: Depth First Search

LIFO

Last in first out

```
DFS( $G = (V, E)$ ):  
   $S \leftarrow$  empty stack  
  visited  $\leftarrow \emptyset$   
  
  Push node 1 onto  $S$   
  While  $S$  is not empty:  
     $u \leftarrow$  pop from  $S$   
    For  $v \in \text{Neighbors}(u)$ :  
      if  $v \notin$  visited:  
        Push  $v$  onto  $S$   
    Add  $u$  to visited  
  
  If  $|\text{visited}| = |V|$ :  
    return True  
  Else:  
    return False
```




Idea: traverse the graph edge by edge.

If we can reach every node without restarting, we know the graph is connected!


Depth First Search Running time

```
DFS( $G = (V, E)$ ):  
   $S \leftarrow$  empty stack  
  visited  $\leftarrow \emptyset$   
  
  Push node 1 onto  $S$   
  While  $S$  is not empty:  
     $u \leftarrow$  pop from  $S$   
    For  $v \in \text{Neighbors}(u)$ :  
      if  $v \notin$  visited:  
        Push  $v$  onto  $S$   
    Add  $u$  to visited  
  
  If  $|\text{visited}| = |V|$ :  
    return True  
  Else:  
    return False
```



Depth First Search Running time

```
DFS( $G = (V, E)$ ):  
   $S \leftarrow$  empty stack  
  visited  $\leftarrow \emptyset$   
  
  Push node 1 onto  $S$   
  While  $S$  is not empty:  
     $u \leftarrow$  pop from  $S$   
    For  $v \in \text{Neighbors}(u)$ :  
      if  $v \notin$  visited:  
        Push  $v$  onto  $S$   
    Add  $u$  to visited  
  
  If  $|\text{visited}| = |V|$ :  
    return True  
  Else:  
    return False
```



Same argument as BFS!

$$O(n + m)$$

Note: These algorithms have recursive equivalents!

```
DFS( $G = (V, E)$ ):  
   $S \leftarrow$  empty stack  
  visited  $\leftarrow \emptyset$   
  
  Push node 1 onto  $S$   
  While  $S$  is not empty:  
     $u \leftarrow$  pop from  $S$   
    For  $v \in \text{Neighbors}(u)$ :  
      if  $v \notin$  visited:  
        Push  $v$  onto  $S$   
    Add  $u$  to visited  
  
  If  $|\text{visited}| = |V|$ :  
    return True  
  Else:  
    return False
```

Our iterative DFS really just makes
the recursive stack explicit!


```
RECURSIVEDFS( $v$ ):  
  if  $v$  is unmarked  
    mark  $v$   
    for each edge  $vw$   
      RECURSIVEDFS( $w$ )
```

Exploring Connected Components

Subgraphs & Components

A graph $G' = (V', E')$ is a *subgraph* of another graph $G = (V, E)$ if $V' \subseteq V$ and $E' \subseteq E$.

A graph is trivially a subgraph of itself. We usually exclude this case and unless otherwise specified we mean *proper* subgraphs.




Subgraphs & Components

Every graph is made up of 1 or more *components*, which are maximal connected subgraphs.

Two nodes are in the same component if they are mutually reachable.




Nodes are in different components if they cannot be reached from one another.

We can use our exploration algorithms to find connected components!



Finding Undirected Components with DFS



```
ComponentsDFS( $G = (V, E)$ ):  
  component[v] = -1 For all  $v \in V$   
  comp = 1  
  S  $\leftarrow$  empty stack  
  visited  $\leftarrow \emptyset$   
  
  Push node 1 onto S  
  component[1] = comp  
  While S is not empty:  
    u  $\leftarrow$  pop from S  
    For  $v \in \text{Neighbors}(u)$ :  
      if  $v \notin \text{visited}$ :  
        Push v onto S  
        component[v] = comp  
    Add u to visited  
  
  If S is empty AND  $|\text{visited}| < |V|$ :  
    Choose a node  $v \in V - \text{visited}$   
    Push v onto S  
    comp = comp + 1  
    component[v] = comp
```



What about directed graphs?

Recall: Two types of connected components in directed graphs

1. Weakly connected: for every pair (u, v) , at least one node is reachable from the other.
2. Strongly connected: for every pair (u, v) , both nodes are reachable from the other



Finding Strongly Connected Components

$SCC(G = (V, E)):$

Let G^R be G with all edges "reversed"

Let $comp[u] \leftarrow -1$ for all u

Let $c \leftarrow 0$

For u from 1..n:

 if $comp[u] = -1$:

 Let S be the nodes found by $DFS(G, u)$



 Let T be the nodes found by $DFS(G^R, u)$

 // intersection of S and T is a SCC!

 label $comp[v] = c$ for all $v \in S \cap T$

 let $c \leftarrow c + 1$

Return $comp$



$S = \{8, 10\}$

$T = \{10, 11, 9\}$

Pause: What have we done so far?

We defined two graph traversal algorithms that can help us determine *reachability* between nodes and overall *connectivity* of a graph

- DFS: Stack based algorithm
- BFS: Queue based algorithm
- Both can be written either recursively or iteratively

We showed how we can use these algorithms to discover the *components* of a graph

- In undirected graphs, it is enough to just run our traversal algorithm until every node is visited once, assigning to a new component every time we “run out” of nodes
- In directed graphs, we need to check both directions to get *strongly connected* components.
 - We achieve this by cleverly running DFS from the same node twice, first on the input graph as usual, then on the graph with reversed edges. The intersection of the reachable sets for these two DFS calls is a strongly connected component!

Typology of Edges in DFS

For every node discovered during a DFS execution, we can keep track of its parent.

Typology of Edges in DFS

For every node discovered during a DFS execution, we can keep track of its parent.

```
DFS( $G = (V, E)$ ):  
   $S \leftarrow$  empty stack  
  visited  $\leftarrow \emptyset$   
  
  Push node 1 onto  $S$   
  While  $S$  is not empty:  
     $u \leftarrow$  pop from  $S$   
    For  $v \in \text{Neighbors}(u)$ :  
      if  $v \notin$  visited:  
        Push  $v$  onto  $S$   
        parent[ $v$ ]  $\leftarrow u$   
    Add  $u$  to visited  
  
  If |visited| =  $|V|$ :  
    return True  
  Else:  
    return False
```


Typology of Edges in DFS

For every node discovered during a DFS execution, we can keep track of its parent.

The graph of the parent-child relationships is a tree where each edge can be assigned to one of four types:

Tree edge:

- Explore new nodes

Forward edge:



- Ancestor to descendant

Backward edge:

- Descendant to ancestor

Cross edges:

- No ancestral relationship



Typology of Edges in DFS

For every node discovered during a DFS execution, we can keep track of its parent.

The graph of the parent-child relationships is a tree where each edge can be assigned to one of four types:

Tree edge: (u, a) , (u, b) , (b, c)

- Explore new nodes

Forward edge: (u, c)


- Ancestor to descendant

Backward edge: (a, u)


- Descendant to ancestor

Cross edges: (b, a)

- No ancestral relationship



Backwards edges identify *cycles* in the graph!



A cycle is a closed walk (starts and ends at the same vertex) that visits each vertex in the walk at most once.

Much more to come on graphs!

Tomorrow: Shortest paths and betweenness centrality

Suggested Reading assignment: Erickson up through Chapter 8.5

Homework 3 will be out shortly after class. Get started early!

Midterm grades coming later this week, please be patient!

