


CSE 332



Data Abstractions

Graphs 2:

Representing Graphs

Topological Sort





$$V = \{a\}, E = \emptyset$$

$$V = \{b, c\}, \\ E = \{(b, c)\}$$


Let's extend our terminology for **directed graphs**!




$$V = \{a\}, E = \emptyset$$



$$V = \{b, c\}, \\ E = \{(b, c)\}$$




$$V = \{d, e, f\}, \\ E = \{(f, e), (f, d)\}$$




$$V = \{g, h, i, j\}, \\ E = \{(g, h), (h, i), (g, j), \\ (i, h), (j, h), (i, j)\}$$

Let's extend our terminology for **directed graphs**!

A Lonely Graph




Complete Directed Graph




Some Questions

- How many edges can a **directed** graph with $|V| = n$ have?
- How many edges can a **directed** graph with $|V| = n$ and possible loops have?

A Lonely Graph



Complete Directed Graph




Some Questions

- How many edges can a **directed** graph with $|V| = n$ have?


$$|E| = n(n - 1)$$

- How many edges can a **directed** graph with $|V| = n$ and possible loops have?

A Lonely Graph



Complete Directed Graph




Some Questions

- How many edges can a **directed** graph with $|V| = n$ have?

$$|E| = n(n-1)$$

- How many edges can a **directed** graph with $|V| = n$ and possible loops have?


$$|E| = n^2$$



Definition (Degree)

The **degree** of a vertex in a graph is the number of vertices adjacent to it. In the above graph, we have:


a	b	c	d	e	f	g	h



Definition (Degree)

The **degree** of a vertex in a graph is the number of vertices adjacent to it. In the above graph, we have:

a	b	c	d	e	f	g	h
3	2	3	2	3	1	1	1




Definition (In & Out Degree)

The **in-degree** of a vertex, v , in a graph is $|\{(x, v) \mid (x, v) \in E, x \in V\}|$.

The **out-degree** of a vertex, v , in a graph is $|\{(v, x) \mid (x, v) \in E, x \in V\}|$.

	a	b	c	d	e	f	g	h
In-Degree								
Out-Degree								




Definition (In & Out Degree)

The **in-degree** of a vertex, v , in a graph is $|\{(x, v) \mid (x, v) \in E, x \in V\}|$.

The **out-degree** of a vertex, v , in a graph is $|\{(v, x) \mid (x, v) \in E, x \in V\}|$.

	a	b	c	d	e	f	g	h
In-Degree	1	2	2	1	0	1	1	0
Out-Degree								





Definition (In & Out Degree)

The **in-degree** of a vertex, v , in a graph is $|\{(x, v) \mid (x, v) \in E, x \in V\}|$.



The **out-degree** of a vertex, v , in a graph is $|\{(v, x) \mid (x, v) \in E, x \in V\}|$.

	a	b	c	d	e	f	g	h
In-Degree	1	2	2	1	0	1	1	0
Out-Degree	2	0	1	1	3	0	0	1

Paths?




Cycle




Definition (**Strongly** Connected Directed Graph)

We say a directed graph is **strongly connected** iff for every pair of vertices, $u, v \in V$, there is a path from u to v .



Strongly Connected!




Not Strongly Connected!

Definition (**Weakly** Connected Directed Graph)

We say a directed graph is **weakly connected** iff the underlying undirected graph is connected.

That is, if we “undirected the edges”, if the graph is connected, then the digraph is weakly connected.



Adjacency Matrix

	a	b	c	d
a	0	1	1	0
b	1	0	1	0
c	1	1	0	1
d	0	0	1	0


Adjacency List

a: [b] → [c] →
b: [a] → [c] →
c: [a] → [b] → [d] →
d: [c] →

Adjacency Matrix

	a	b	c	d
a	0	1	1	0
b	1	0	1	0
c	1	1	0	1
d	0	0	1	0

Adjacency List



Adjacency Matrix Properties

How long to...

- Get a vertex's out-edges? $\mathcal{O}(|V|)$
- Get a vertex's in-edges? $\mathcal{O}(|V|)$
- Check if an edge exists? $\mathcal{O}(1)$
- Insert an edge? $\mathcal{O}(1)$
- Delete an edge? $\mathcal{O}(1)$


Space Requirements: $\mathcal{O}(|V|^2)$

Adjacency Matrices are reasonable for dense graphs, but not otherwise.

Adjacency Matrix

	a	b	c	d
a	0	1	1	0
b	1	0	1	0
c	1	1	0	1
d	0	0	1	0

Adjacency List



Adjacency List Properties

How long to...


- Get a vertex's out-edges? $\mathcal{O}(d)$
- Get a vertex's in-edges? $\mathcal{O}(|E|)$
 - To fix this, keep a **second** adjacency list going the other way
- Check if an edge exists? $\mathcal{O}(d)$
- Insert an edge? $\mathcal{O}(1)$
- Delete an edge? $\mathcal{O}(d)$

Space Requirements: $\mathcal{O}(|V| + |E|)$

Adjacency Lists should be your goto choice.

Definition (DAG)

A **DAG** is a **directed**, **acyclic** graph.



By “acyclic”, we mean in the **directed** sense.

DAGs vs. Trees?

Is there a tree that isn't a DAG?

Is there a DAG that isn't a tree?

DAGs vs. Trees?

All trees are DAGs (remember, trees must be acyclic and connected!).

Not all DAGs are trees. See previous slide. Also, DAGs don't have to be connected!

Why DAGs?

They come up a lot in practice. Cycles can be icky. Examples:


- Any sort of scheduling problem (scheduling your courses, scheduling fork-join threads, ...)
- Causal Structures (Baysian Networks)
- Genealogy
- ...

Topological Sort


Given a DAG ($G = (V, E)$), output all the vertices in an order such that no vertex appears before any vertex that has an edge to it.

“Output an order to process the graph that meets all dependencies”

This is how we can allocate work in the ForkJoin model!



How Many Valid Topological Sorts?



- $T_1, T_2, T_3, T_4, T_5, T_6, T_7, T_8, T_9, T_{10}$
- $T_1, T_2, T_4, T_3, T_5, T_6, T_7, T_8, T_9, T_{10}$
- $T_1, T_2, T_5, T_4, T_3, T_6, T_7, T_8, T_9, T_{10}$
- $T_1, T_3, T_6, T_7, T_9, T_2, T_5, T_4, T_8, T_{10}$
- ...

Implementing Topological Sort

Throw all the **in-degrees** in a priority queue. `removeMin()` repeatedly.

- This works, but it's **too slow**.
- Insight: PriorityQueues must deal with negative numbers; indegree will never be negative!
- Instead: Split ready vs. not ready (0 vs. non-zero) sets
- The “ready set” is a worklist!


Setup

```
1 output = []
2 deps = {}
3 worklist = []
4 for (v : vertices) {
5     deps[v] = in-degree(v);
6     if (deps[v] == 0) {
7         worklist.add(v);
8     }
9 }
```

Do Work

```
1 while (worklist.hasWork()) {
2     v = worklist.next();
3     output.add(v);
4     for (w : neighbors(v)) {
5         deps[w] -= 1
6         if (deps[w] == 0) {
7             worklist.add(w);
8         }
9     }
10 }
```


worklist \leftarrow



output

o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]


worklist \leftarrow



output

o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]

worklist $\leftarrow \boxed{T_1 \quad T_8 \quad T_{10}} \leftarrow$




output

o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]

worklist \leftarrow

T_8	T_{10}
-------	----------


 \leftarrow



output

T_1									
$o[0]$	$o[1]$	$o[2]$	$o[3]$	$o[4]$	$o[5]$	$o[6]$	$o[7]$	$o[8]$	$o[9]$


worklist $\leftarrow \boxed{T_8 \quad T_{10} \quad T_3} \leftarrow$



output

T_1									
$o[0]$	$o[1]$	$o[2]$	$o[3]$	$o[4]$	$o[5]$	$o[6]$	$o[7]$	$o[8]$	$o[9]$


worklist $\leftarrow \boxed{T_{10} \quad T_3} \leftarrow$



output

T_1	T_8								
$o[0]$	$o[1]$	$o[2]$	$o[3]$	$o[4]$	$o[5]$	$o[6]$	$o[7]$	$o[8]$	$o[9]$


worklist $\leftarrow \boxed{T_{10} \quad T_3} \leftarrow$



output

T_1	T_8								
$o[0]$	$o[1]$	$o[2]$	$o[3]$	$o[4]$	$o[5]$	$o[6]$	$o[7]$	$o[8]$	$o[9]$


worklist $\leftarrow \boxed{T_3} \leftarrow$



output

T_1	T_8	T_{10}							
o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]


worklist \leftarrow



output

T_1	T_8	T_{10}	T_3						
$\circ[0]$	$\circ[1]$	$\circ[2]$	$\circ[3]$	$\circ[4]$	$\circ[5]$	$\circ[6]$	$\circ[7]$	$\circ[8]$	$\circ[9]$


worklist $\leftarrow \boxed{T_4} \leftarrow$



output

T_1	T_8	T_{10}	T_3						
$o[0]$	$o[1]$	$o[2]$	$o[3]$	$o[4]$	$o[5]$	$o[6]$	$o[7]$	$o[8]$	$o[9]$


worklist \leftarrow



output

T_1	T_8	T_{10}	T_3	T_4					
$\circ[0]$	$\circ[1]$	$\circ[2]$	$\circ[3]$	$\circ[4]$	$\circ[5]$	$\circ[6]$	$\circ[7]$	$\circ[8]$	$\circ[9]$


worklist $\leftarrow \boxed{T_2 \quad T_5} \leftarrow$



output

T_1	T_8	T_{10}	T_3	T_4					
$o[0]$	$o[1]$	$o[2]$	$o[3]$	$o[4]$	$o[5]$	$o[6]$	$o[7]$	$o[8]$	$o[9]$


worklist $\leftarrow \boxed{T_5} \leftarrow$



output

T_1	T_8	T_{10}	T_3	T_4	T_2				
$o[0]$	$o[1]$	$o[2]$	$o[3]$	$o[4]$	$o[5]$	$o[6]$	$o[7]$	$o[8]$	$o[9]$


worklist $\leftarrow \boxed{T_5 \quad T_7} \leftarrow$



output

T_1	T_8	T_{10}	T_3	T_4	T_2				
$\circ[0]$	$\circ[1]$	$\circ[2]$	$\circ[3]$	$\circ[4]$	$\circ[5]$	$\circ[6]$	$\circ[7]$	$\circ[8]$	$\circ[9]$


worklist $\leftarrow \boxed{T_7} \leftarrow$



output

T_1	T_8	T_{10}	T_3	T_4	T_2	T_5			
$o[0]$	$o[1]$	$o[2]$	$o[3]$	$o[4]$	$o[5]$	$o[6]$	$o[7]$	$o[8]$	$o[9]$


worklist $\leftarrow \boxed{T_7} \leftarrow$



output

T_1	T_8	T_{10}	T_3	T_4	T_2	T_5			
o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]


worklist \leftarrow



output

T_1	T_8	T_{10}	T_3	T_4	T_2	T_5	T_7		
o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]


worklist $\leftarrow \boxed{T_6 \quad T_9} \leftarrow$



output

T_1	T_8	T_{10}	T_3	T_4	T_2	T_5	T_7		
o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]


worklist $\leftarrow \boxed{T_9} \leftarrow$



output


T_1	T_8	T_{10}	T_3	T_4	T_2	T_5	T_7	T_6	
o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]

worklist \leftarrow



output

T_1	T_8	T_{10}	T_3	T_4	T_2	T_5	T_7	T_6	T_9
o[0]	o[1]	o[2]	o[3]	o[4]	o[5]	o[6]	o[7]	o[8]	o[9]



What happens if there is a cycle?

Our worklist will be empty before we've processed all of the vertices. (e.g., "there are no nodes ready to print next, but we haven't gone through all of them")

In this case: our algorithm should throw a "not a DAG exception".

Runtime?

- Setup: We follow every edge for every vertex: $\mathcal{O}(|V| + |E|)$
- We add/remove each vertex from the work list once: $\mathcal{O}(|V|)$
- We decrement each indegree until zero (once for each edge): $\mathcal{O}(|E|)$
- So, overall, it's graph linear!