

CSE 421 Winter 2025

Lecture 4: Asymptotics and Graph Search

Nathan Brunelle

<http://www.cs.uw.edu/421>

Complexity analysis

- Problem size n
 - **Worst-case complexity:**
maximum # steps algorithm takes on any input of size n
 - **Best-case complexity:**
minimum # steps algorithm takes on any input of size n
 - **Average-case complexity:**
Expected # steps algorithm takes on inputs of size n

Complexity


- The complexity of an algorithm associates a number $T(n)$, the worst/average-case/best time the algorithm takes, with each problem size n .
- Mathematically,
 - T is a function that maps positive integers giving problem size to positive real numbers giving number of steps.
- Sometimes we have more than one size parameter
 - e.g. n =# of vertices, m =# of edges in a graph.

Efficient = Polynomial Time

- Polynomial time
 - Running time $T(n) \leq cn^k + d$ for some $c, d, k \geq 0$
- Why polynomial time?
 - If problem size grows by at most a constant factor then so does the running time
 - e.g. $T(2n) \leq c(2n)^k + d = 2^k cn^k + d \leq 2^k(cn^k + d) = 2^k T(n)$
 - polynomial-time is exactly the set of running times that have this property
 - Typical running times are small degree polynomials, mostly less than n^3 , at worst n^6 , not n^{100}

Complexity

Time



O-notation etc

- Given two positive functions f and g
 - $f(n)$ is $O(g(n))$ iff there is a constant $c > 0$
so that $f(n)$ is eventually always $\leq c \cdot g(n)$
 - $f(n)$ is $o(g(n))$ iff for every constant $c > 0$
 $f(n)$ is eventually always $\leq c \cdot g(n)$
 - $f(n)$ is $\Omega(g(n))$ iff there is a constant $\varepsilon > 0$ so that $f(n) \geq \varepsilon \cdot g(n)$ for
infinitely many values of n
 - $f(n)$ is $\omega(g(n))$ iff for every constant $c > 0$
 $f(n)$ is eventually always $\geq c \cdot g(n)$
 - $f(n)$ is $\Theta(g(n))$ iff $f(n)$ is $O(g(n))$ and $f(n)$ is $\Omega(g(n))$

Note: The definition of “ $f(n)$ is $\Omega(g(n))$ ” is *almost* the same as “ $f(n)$ is **not** $o(g(n))$ ”

The definition of “ $f(n)$ is $\Omega(g(n))$ ” is *almost* the same as “ $f(n)$ is **not** $o(g(n))$ ”

Asymptotic Notation intuition

$f(n)$ is...


$O(g(n))$: ratio eventually below some line forever

$o(g(n))$: ratio eventually below every line forever

$\Omega(g(n))$: ratio eventually above some line forever

$\omega(g(n))$: ratio eventually above every line forever


$\Theta(g(n))$: both O and Ω




Introduction to Algorithms

- **Graph Search/Traversal**

Undirected Graph $G = (V, E)$



Directed Graph $G = (V, E)$



Graph Traversal

Learn the basic structure of a graph

Walk from a fixed starting vertex s to find all vertices reachable from s

Generic Graph Traversal Algorithm

Given: Graph $G = (V, E)$ vertex $s \in V$

Find: set R of vertices reachable from $s \in V$

Reachable(s):

 Add s to R

 while there is a $(u, v) \in E$ where $u \in R$ and $v \notin R$

 Add v to R

 return R

Generic Traversal Always Works

Claim: At termination, R is the set of nodes reachable from s

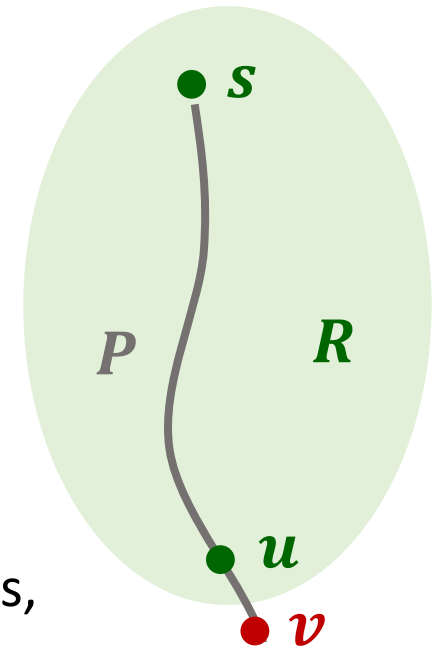
Proof

\subseteq : For every node $v \in R$ there is a path from s to v

- Induction based on edges found.
 - Base case: s is reachable from s
 - Inductive step: If there is a path to every member of R after i iterations, then there is a path to every member of R after $i + 1$ iterations

\supseteq : Suppose there is a node $w \notin R$ reachable from s via a path P

- Take **first** node v on P such that $v \notin R$
- Predecessor u of v in P satisfies
 - $u \in R$
 - $(u, v) \in E$
- But this contradicts the fact that the algorithm exited the while loop.



Graph Traversal

Learn the basic structure of a graph

Walk from a fixed starting vertex s to find all vertices reachable from s

Three states of vertices

- **unvisited**
- **visited/discovered** (in R)
- **fully-explored** (in R and all neighbors have been visited)

Breadth-First Search

Completely explore the vertices in order of their distance from **s**

Naturally implemented using a queue

BFS(s)

Global initialization: mark all vertices “unvisited”

BFS(s)

Mark s “visited”

Add s to Q

$i = 0$

Mark s as “layer i ”

while Q not empty

u = next item removed from Q

i = “layer of u ”

 for each edge (u, v)

 if (v is “unvisited”)

 mark v “visited”

 mark s as “layer $i + 1$ ”

 mark u “fully-explored”

Properties of BFS

$\text{BFS}(s)$ visits x iff there is a path in G from s to x .

Edges followed to undiscovered vertices define a
breadth first spanning tree of G

Layer i in this tree:

L_i = set of vertices u with shortest path in G from root s of length i .

Properties of BFS

Claim: For undirected graphs:

All edges join vertices on the same or adjacent layers of BFS tree

Proof: Suppose not...

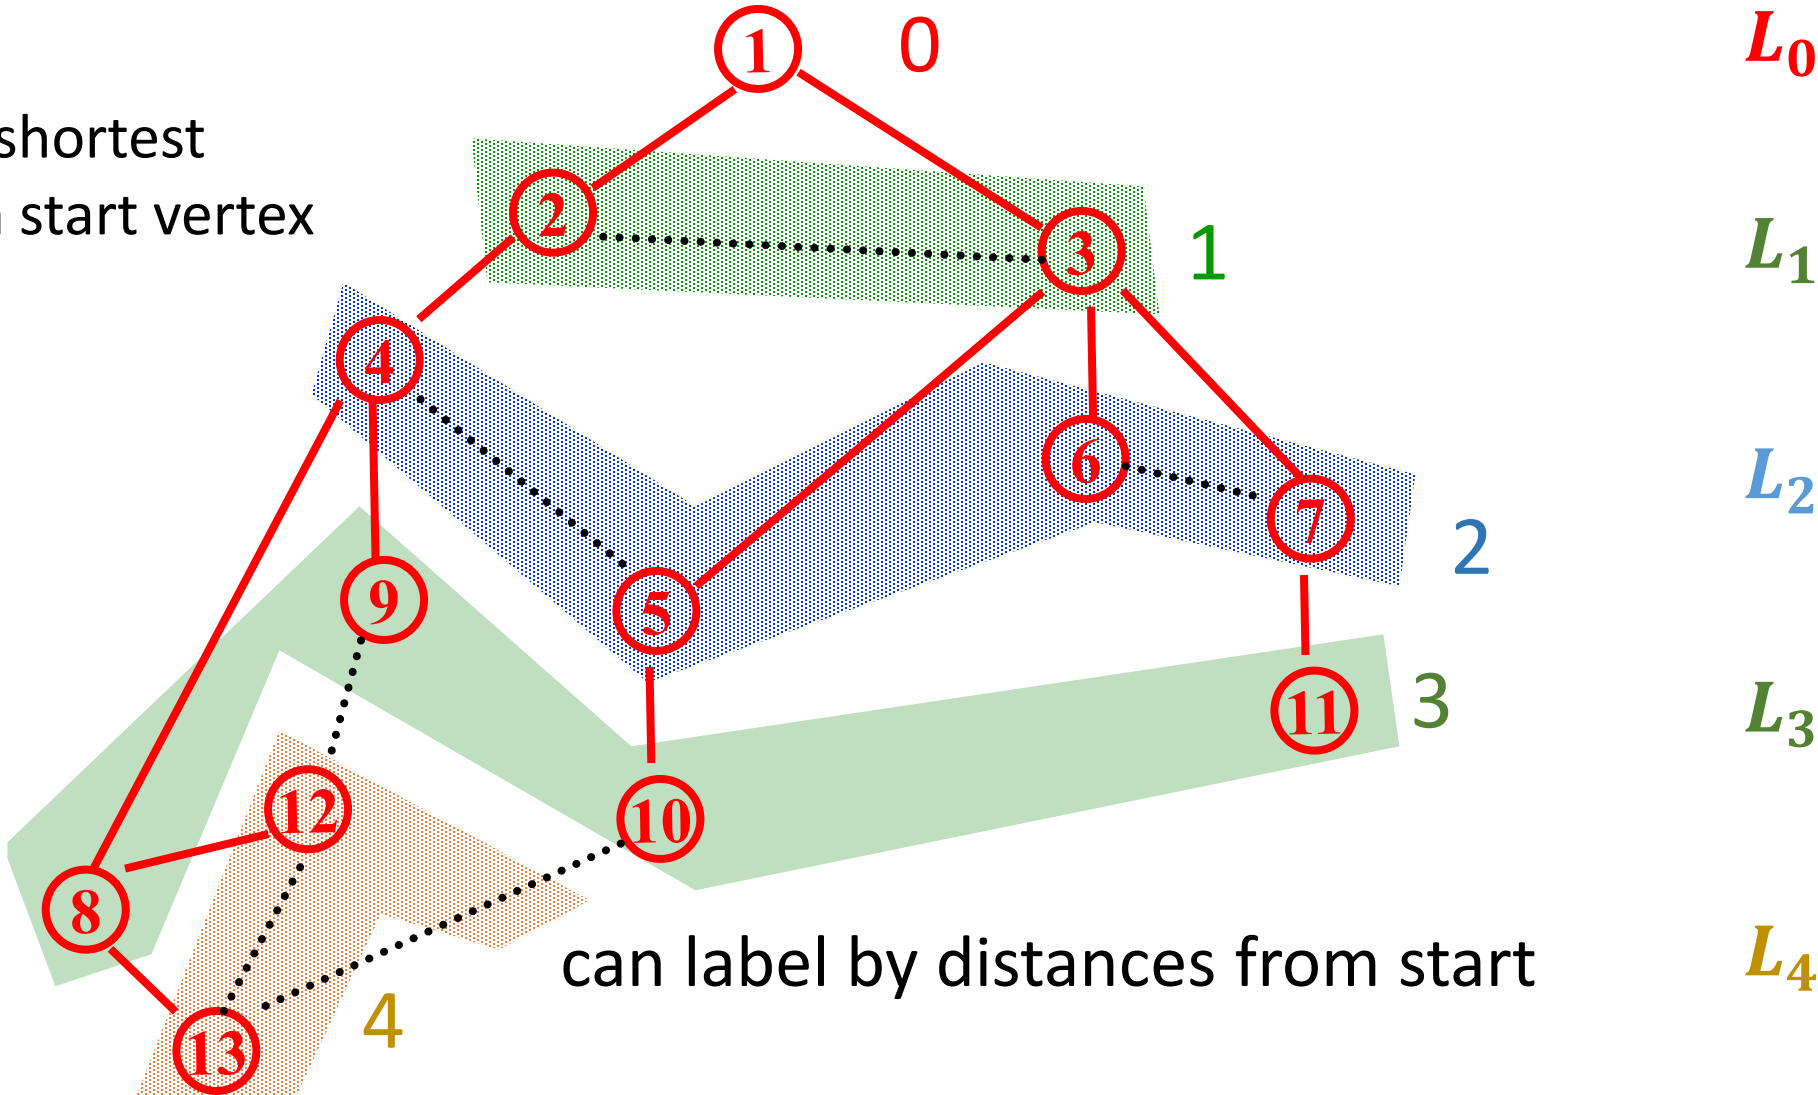
Then there would be vertices (x, y) s.t. $x \in L_i$ and $y \in L_j$ and $j > i + 1$.

Then, when vertices adjacent to x are considered in BFS, y would be added with layer $i + 1$ and not layer j .

Contradiction.

BFS Application: Shortest Paths

Tree gives shortest paths from start vertex



can label by distances from start

Undirected Graph Search Application: Connected Components

Want to answer questions of the form:

Given: vertices u and v in G

Is there a path from u to v ?

Idea: create array A s.t

$A[u]$ = smallest numbered vertex connected to u

Answer is yes iff $A[u] = A[v]$

Q: Why is this better than
an array $Path[u, v]$?

Undirected Graph Search Application: Connected Components

Initial state: all v unvisited

for s from 1 to n do:

if $\text{state}(s) \neq \text{fully-explored}$ then

BFS(s): setting $A[u] = s$ for each u found

(and marking u visited/fully-explored)

Total cost: $O(n + m)$

- Each vertex is touched once in outer procedure and edges examined in different BFS runs are disjoint
- Works also with Depth First Search ...

DFS(u) – Recursive Procedure

Global Initialization: mark all vertices "unvisited"

DFS(u)

mark u "visited" and add u to R

for each edge (u, v)

if (v is "unvisited")

DFS(v)

mark u "fully-explored"

Properties of DFS(s)

Like BFS(s):

- DFS(s) visits x iff there is a path in G from s to x
- Edges into undiscovered vertices define **depth-first spanning tree** of G

Unlike the BFS tree:


- the DFS spanning tree *isn't* minimum depth
- its levels *don't* reflect min distance from the root
- non-tree edges *never* join vertices on the same or adjacent levels

BUT...

Non-tree edges in DFS tree of undirected graphs

Claim: All non-tree edges join a vertex and one of its descendants/ancestors in the DFS tree

- In other words ... No “cross edges”.



No cross edges in DFS on undirected graphs

Claim: During $\text{DFS}(x)$ every vertex marked “visited” is a descendant of x in the DFS tree T

Claim: For every x, y in the DFS tree T , if (x, y) is an edge *not* in T then one of x or y is an ancestor of the other in T

Proof:

- One of $\text{DFS}(x)$ or $\text{DFS}(y)$ is called first, suppose WLOG that $\text{DFS}(x)$ was called before $\text{DFS}(y)$
- During $\text{DFS}(x)$, the edge (x, y) is examined
- Since (x, y) is a *not* an edge of T , y was already visited when edge (x, y) was examined during $\text{DFS}(x)$
- Therefore y was visited during the call to $\text{DFS}(x)$ so y is a descendant of x .

Applications of Graph Traversal: Bipartiteness Testing


Definition: An undirected graph G is **bipartite** iff we can color its vertices **red** and **green** so each edge has different color endpoints

Input: Undirected graph G

Goal: If G is bipartite, output a coloring;
otherwise, output “NOT Bipartite”.

Fact: Graph G contains an odd-length cycle \Rightarrow it is not bipartite

Just coloring the cycle part
of G is impossible



On a cycle the two colors must alternate, so

- **green** every 2nd vertex
- **red** every 2nd vertex

Can't have either if length is not divisible by 2.

Applications of Graph Traversal: Bipartiteness Testing

WLOG (“without loss of generality”): Can assume that G is connected

- Otherwise run on each component

Simple idea: start coloring nodes starting at a given node s

- Color s **red**
- Color all neighbors of s **green**
- Color all their neighbors **red**, etc.
- If you ever hit a node that was already colored
 - the **same** color as you want to color it, ignore it
 - the **opposite** color, output “NOT Bipartite” and halt

BFS gives Bipartiteness

Run BFS assigning all vertices from layer L_i the color $i \bmod 2$

- i.e., **red** if they are in an even layer, **green** if in an odd layer
- if no edge joining two vertices of the same color
 - then it is a good coloring
- otherwise
 - there is a bad edge; output “Not Bipartite”

Why is that “Not Bipartite” output correct?

Why does BFS work for Bipartiteness?


Recall: All edges join vertices on the same or adjacent BFS layers

⇒ Any “bad” edge must join two vertices u and v in the same layer

Say the layer with u and v is L_j

u and v have common ancestor at some level L_i for $i < j$

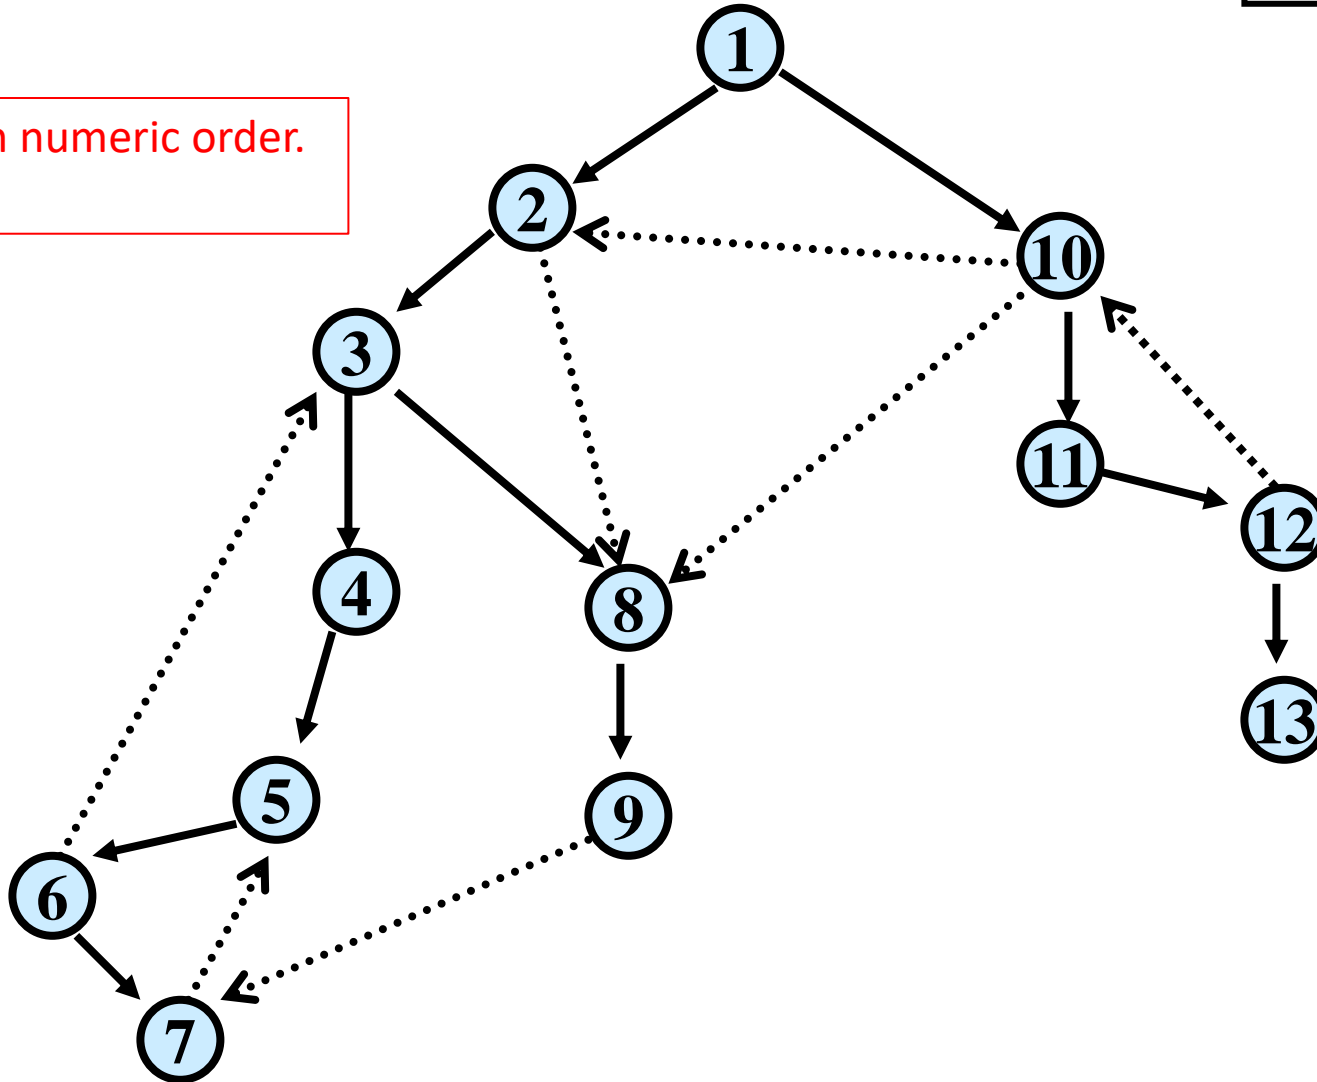
Odd cycle of length $2(j - i) + 1$
⇒ Not Bipartite



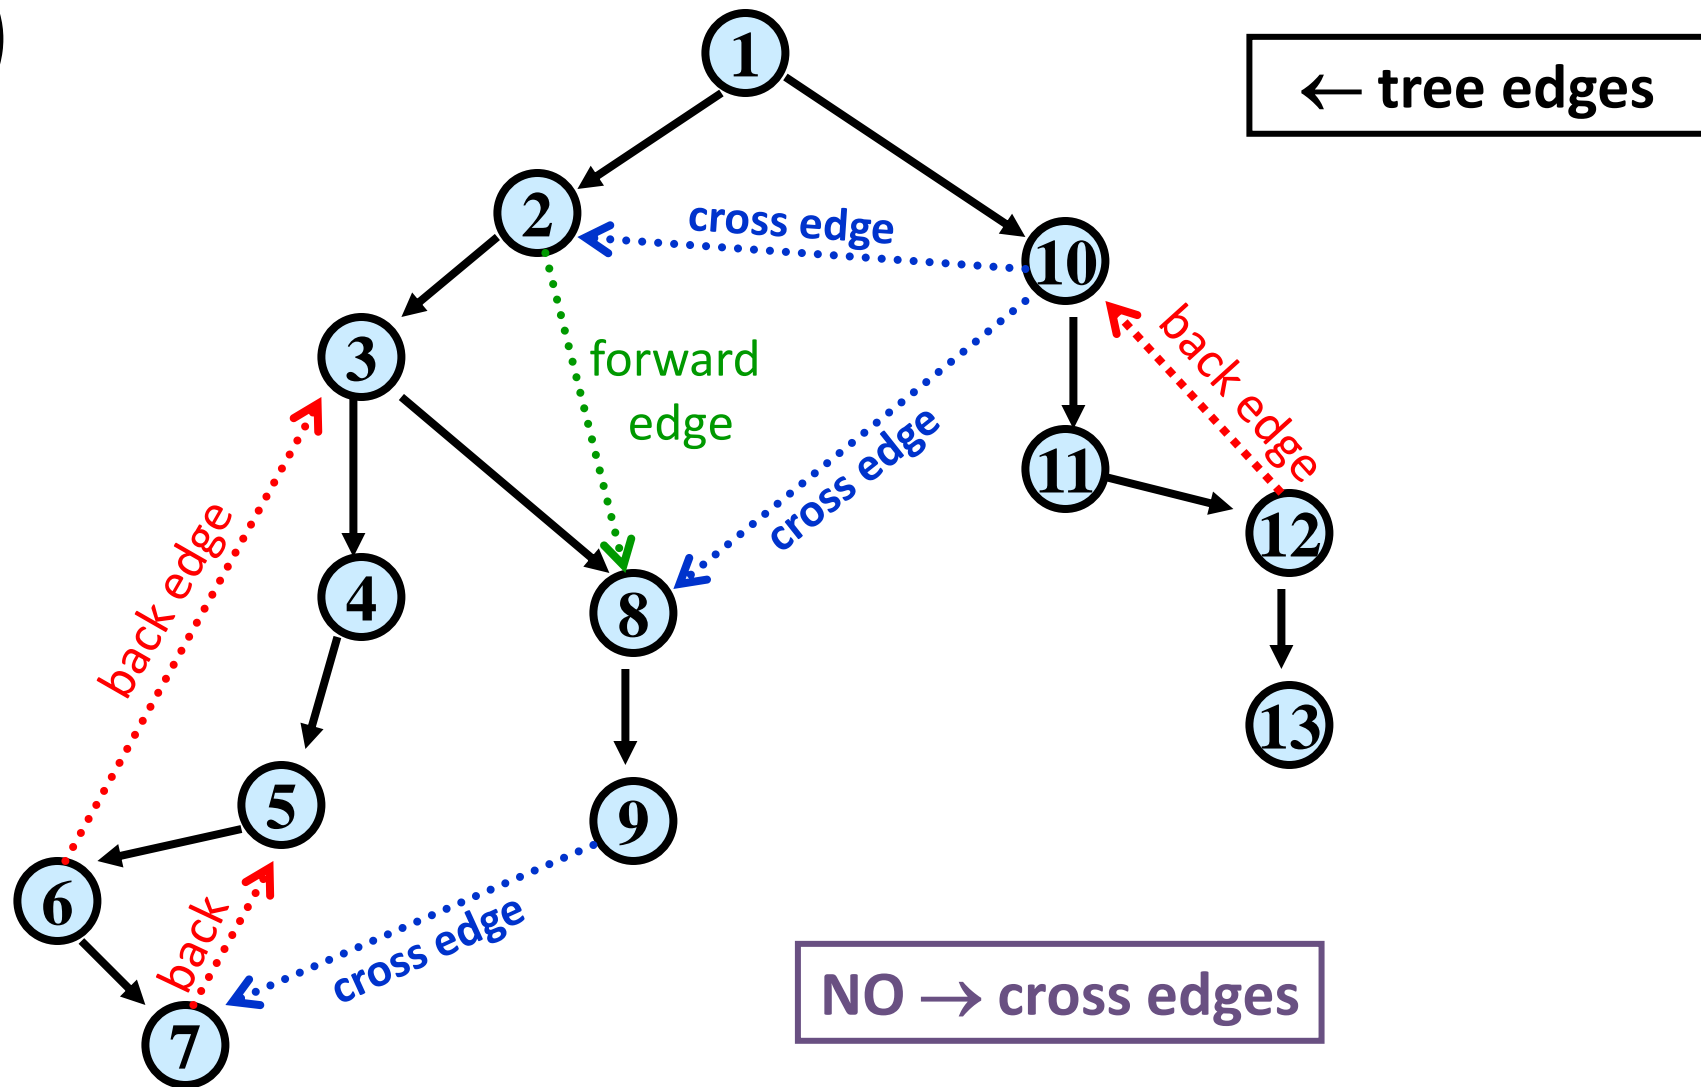
DFS(v) for a directed graph

← tree edges

Nodes were visited in numeric order.
How can we tell?



DFS(v)



Properties of Directed DFS

- Before $\text{DFS}(s)$ returns, it visits all previously unvisited vertices reachable via directed paths from s
- Every cycle contains a back edge in the DFS tree

Directed Acyclic Graphs

A directed graph $G = (V, E)$ is **acyclic** iff it has no directed cycles

Terminology: A **directed acyclic graph** is also called a **DAG**

After shrinking the strongly connected components of a directed graph to single vertices, the result is a DAG

Topological Sort

Given: a directed acyclic graph (DAG) $G = (V, E)$

Output: numbering of the vertices of G with distinct numbers from 1 to n so that edges only go from lower numbered to higher numbered vertices


Applications:

- nodes represent tasks
- edges represent precedence between tasks
- topological sort gives a sequential schedule for solving them

Nice algorithmic paradigm for general directed graphs:

- Process strongly connected components one-by-one in the order given by topological sort of the DAG you get from shrinking them.

Directed Acyclic Graph



In-degree 0 vertices

Claim: Every DAG has a vertex of in-degree 0

Proof: By contradiction

Suppose every vertex has some incoming edge

Consider following procedure:

while (true) do


$v = \text{some predecessor of } v$

- After $n + 1$ steps where $n = |V|$ there will be a repeated vertex
 - This yields a cycle, contradicting that it is a DAG.


Topological Sort

- Can do using DFS
- Alternative simpler idea:
 - Any vertex of in-degree 0 can be given number 1 to start
 - Remove it from the graph
 - Then give a vertex of in-degree 0 number 2
 - Etc.


Topological Sort




Topological Sort




Topological Sort




Topological Sort




Topological Sort




Topological Sort




Topological Sort




Topological Sort



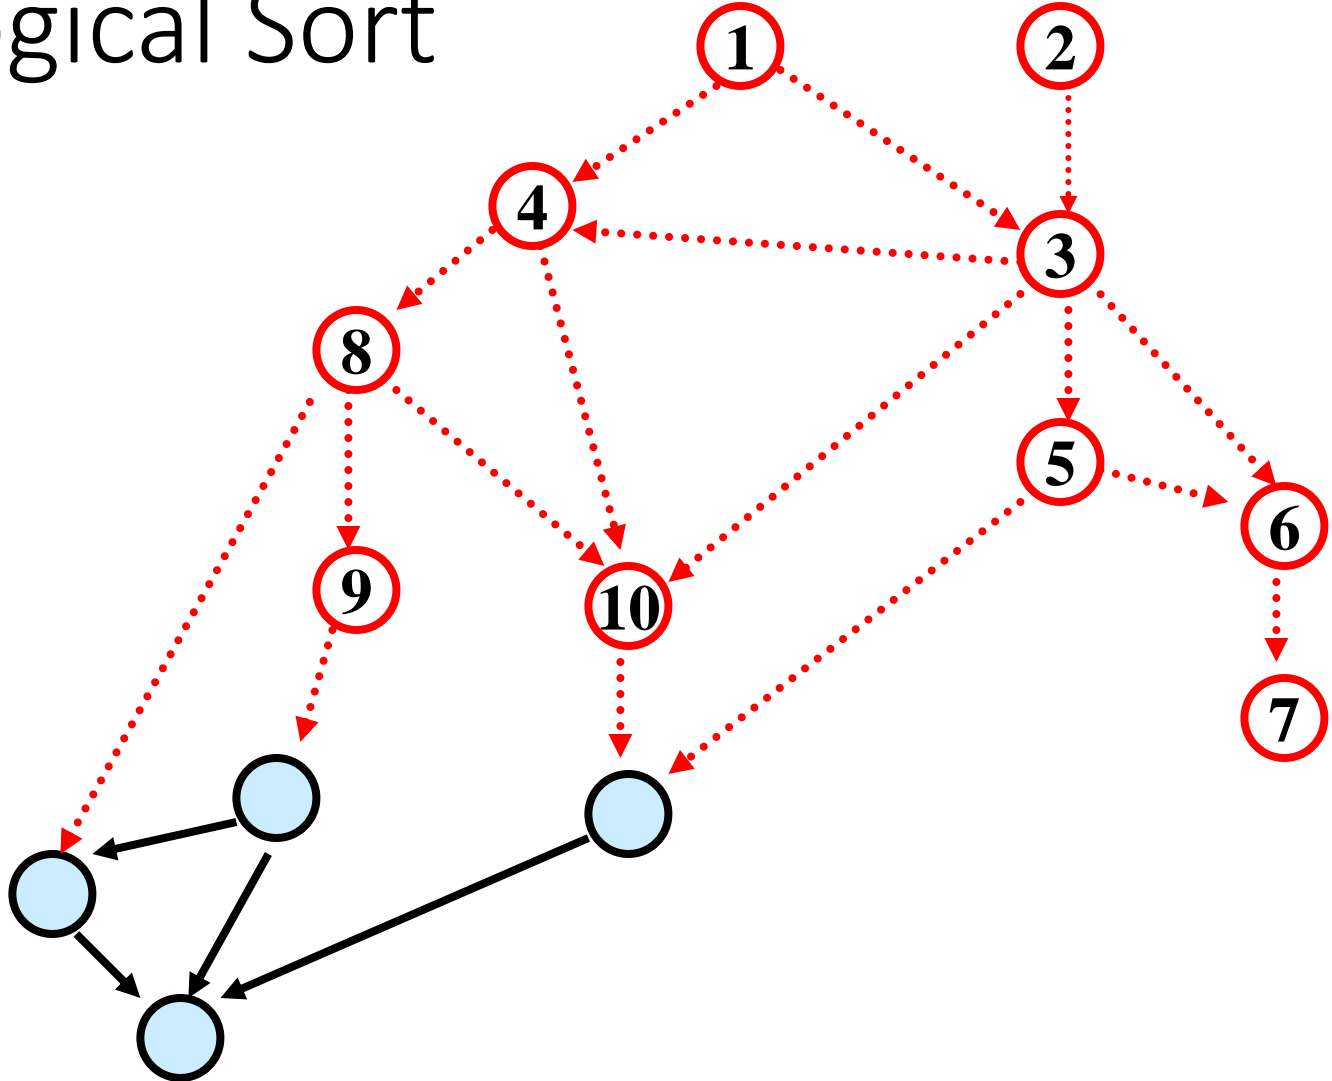
Topological Sort



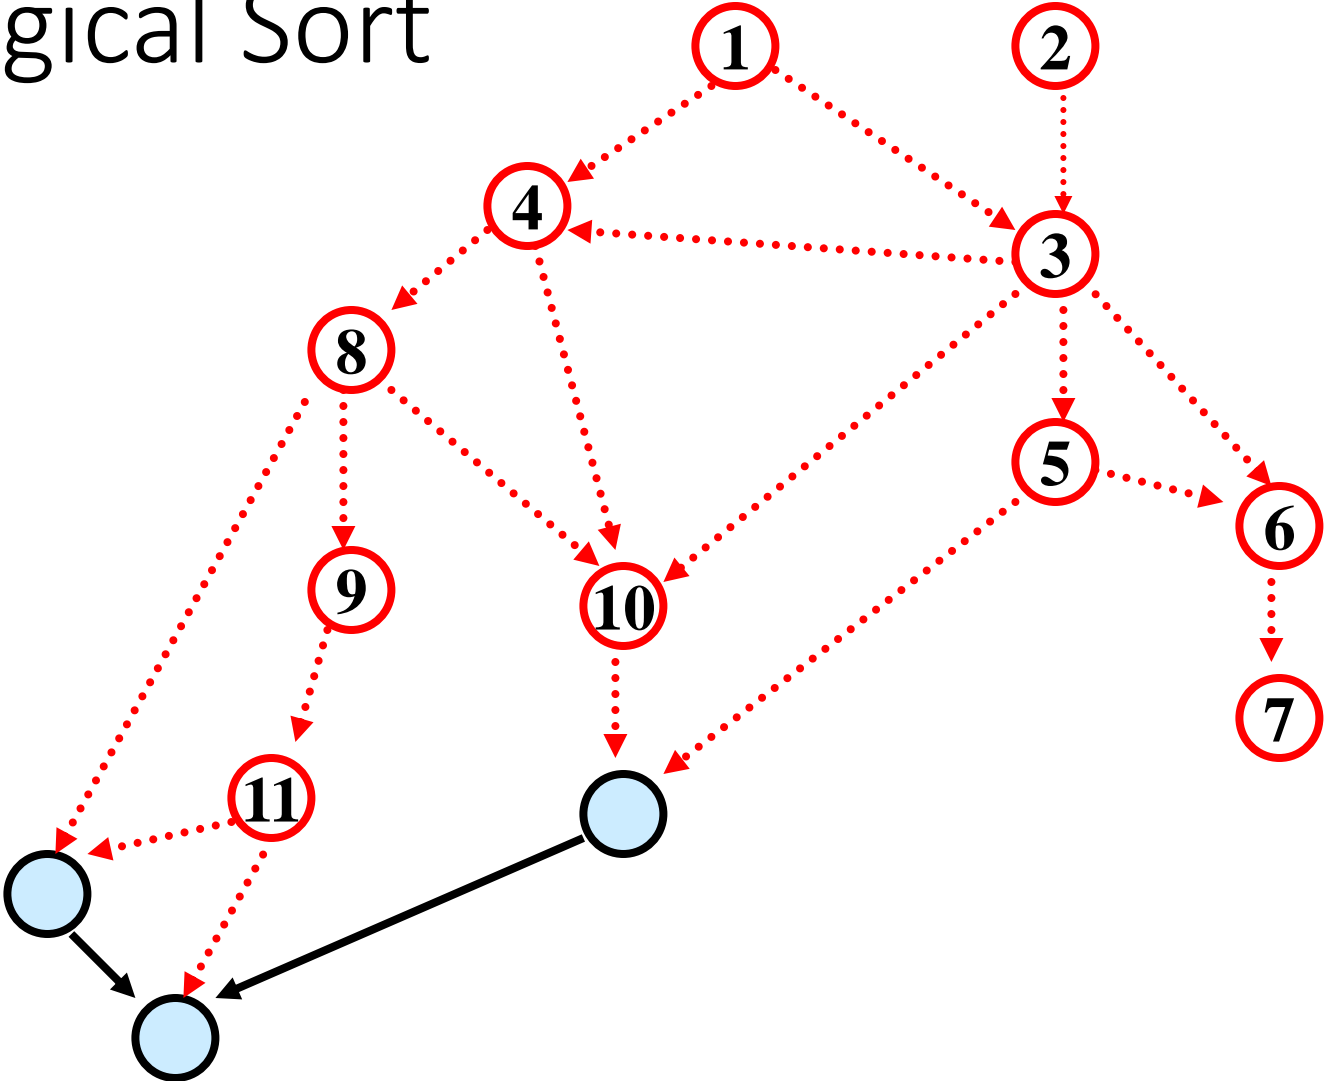
Topological Sort




Topological Sort




Topological Sort




Topological Sort



Topological Sort



Topological Sort



Implementing Topological Sort

- Go through all edges, computing array with in-degree for each vertex
 $O(m + n)$
- Maintain a list of vertices of in-degree **0**
- Remove any vertex in list and number it
- When a vertex is removed, decrease in-degree of each neighbor by **1** and add them to the list if their degree drops to **0**

Total cost: $O(m + n)$

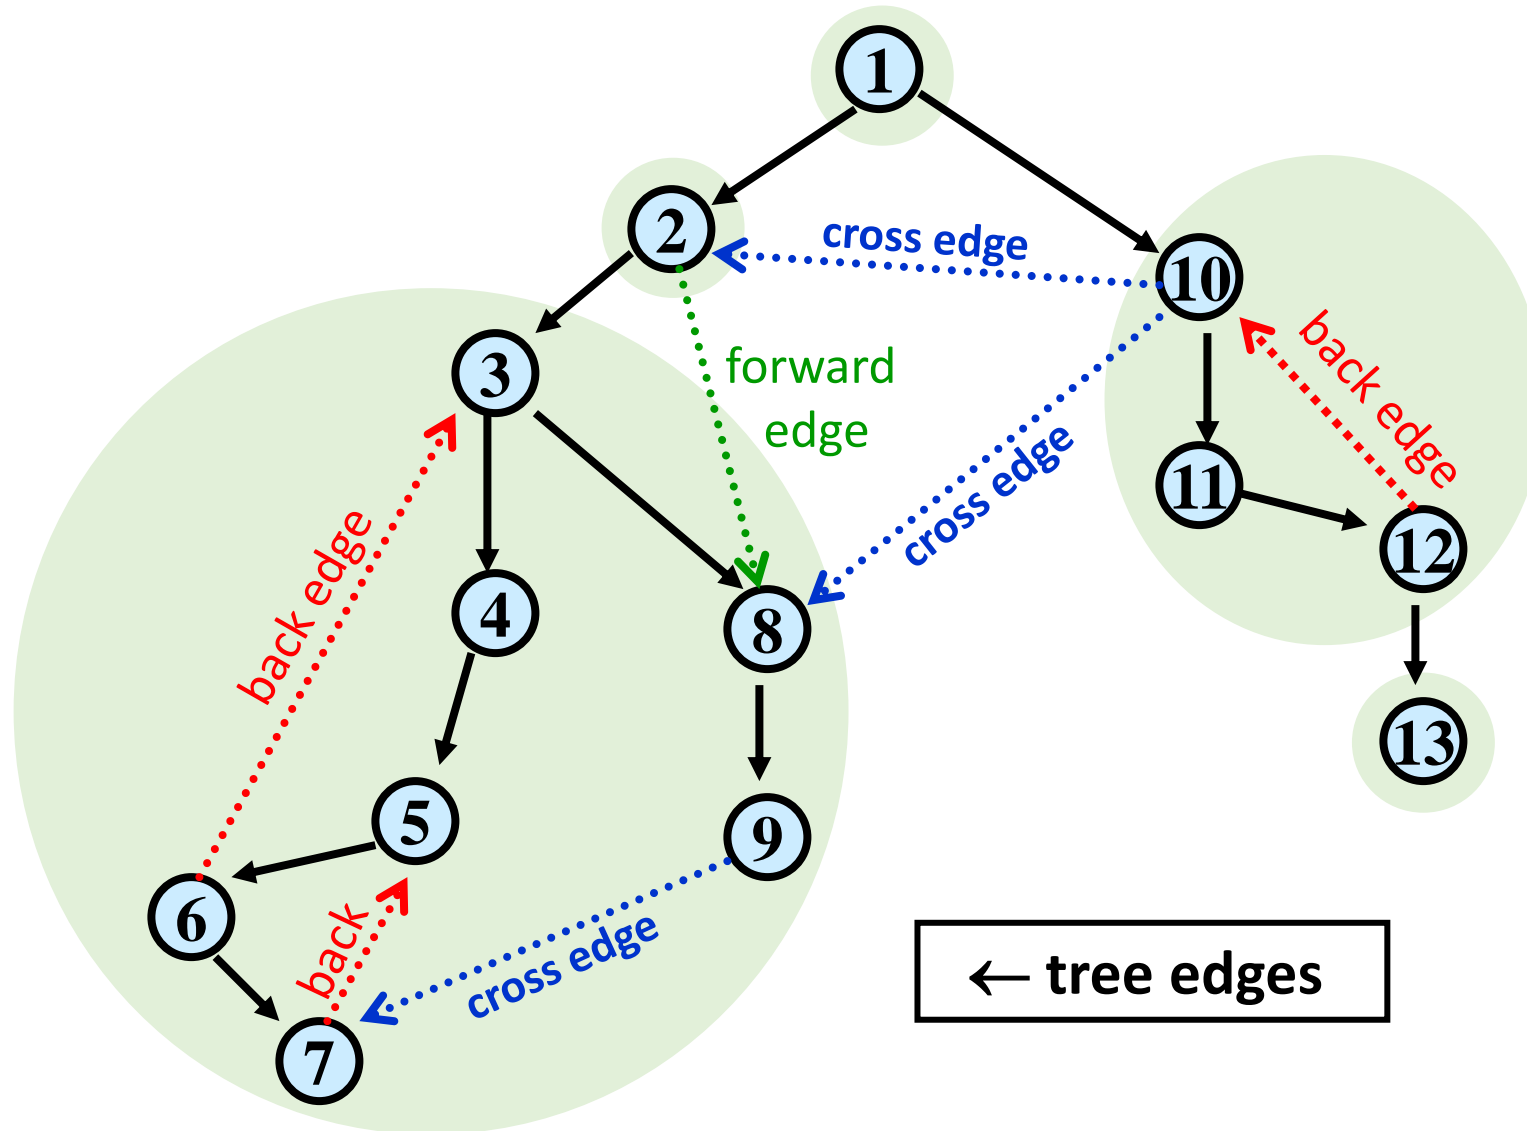
Strongly Connected Components of Directed Graphs

Defn: Vertices u and v are **strongly connected** iff they are on a directed cycle (there are paths from u to v and from v to u).


Defn: Can partition vertices of any directed graph into **strongly connected components**:

1. all pairs of vertices in the same component are strongly connected
 2. can't merge components and keep property 1
- Strongly connected components can be stored efficiently just like connected components
 - Can be found in $O(n + m)$ time using a DFS then a BFS
 - Do a depth-first sort, keeping track of the order nodes are marked “fully-explored”
 - Going in order from least recent to most recent, run connected components


Strongly Connected Components



Strongly Connected Components



Strongly Connected Components



Strongly-Connected Components Usage

Common algorithmic paradigm for general directed graphs:

- Process strongly connected components one-by-one in the order given by topological sort of the DAG you get from shrinking them.