

Graph Representations and Depth-First Search

1 Graphs

A graph is a set of vertices and edges connecting those vertices. Formally, we define a graph G as $G = (V, E)$ where $E \subseteq V \times V$. For ease of analysis, the variables n and m typically stand for the number of vertices and edges, respectively. Graphs can come in two flavors, directed or undirected. If a graph is undirected, it must satisfy the property that $(i, j) \in E \iff (j, i) \in E$ (i.e., all edges are bidirectional). In undirected graphs, $m \leq \frac{n(n-1)}{2}$. In directed graphs, $m \leq n(n-1)$. Thus, $m = O(n^2)$ and $\log m = O(\log n)$. A connected graph is a graph in which for any two nodes u and v there exists a path from u to v . For an undirected connected graph $m \geq n - 1$. A sparse graph is a graph with few edges (for example, $\Theta(n)$ edges) while a dense graph is a graph with many edges (for example, $m = \Theta(n^2)$).

1.1 Representation

A common issue is the topic of how to represent a graph's edges in memory. There are two standard methods for this task.

An **adjacency matrix** uses an arbitrary ordering of the vertices from 1 to $|V|$. The matrix consists of an $n \times n$ binary matrix such that the (i, j) -th element is 1 if (i, j) is an edge in the graph, 0 otherwise.

An **adjacency list** consists of an array A of $|V|$ lists, such that

$$A[u]$$

contains a linked list of vertices v such that $(u, v) \in E$ (the neighbors of u). In the case of a directed graph, it's also helpful to distinguish between outgoing and ingoing edges by storing two different lists at $A[u]$: a list of v such that $(u, v) \in E$ (the out-neighbors of u) as well as a list of v such that $(v, u) \in E$ (the in-neighbors of u).

What are the tradeoffs between these two methods? To help our analysis, let $\text{deg}(v)$ denote the **degree** of v , or the number of vertices connected to v . In a directed graph, we can distinguish between out-degree and in-degree, which respectively count the number of outgoing and incoming edges.

- The adjacency matrix can check if (i, j) is an edge in G in constant time, whereas the adjacency list representation must iterate through up to $\text{deg}(i)$ list entries.

- The adjacency matrix takes $\Theta(n^2)$ space, whereas the adjacency list takes $\Theta(m + n)$ space.
- The adjacency matrix takes $\Theta(n)$ operations to enumerate the neighbors of a vertex v since it must iterate across an entire row of the matrix. The adjacency list takes $\text{deg}(v)$ time.

What's a good rule of thumb for picking the implementation? One useful property is the sparsity of the graph's edges. If the graph is **sparse**, and the number of edges is considerably less than the max ($m \ll n^2$), then the adjacency list is a good idea. If the graph is dense and the number of edges is nearly n^2 , then the matrix representation makes sense because it speeds up lookups without too much space overhead. Of course, some applications will have lots of space to spare, making the matrix feasible no matter the structure of the graphs. Other applications may prefer adjacency lists even for dense graphs. Choosing the appropriate structure is a balancing act of requirements and priorities.

2 Depth First Search (DFS)

Given a starting vertex, it's desirable to find all vertices reachable from the start. There are many algorithms to do this, the simplest of which is depth-first search. As the name implies, DFS enumerates the deepest paths, only backtracking when it hits a dead end or an already-explored section of the graph. DFS by itself is fairly simple, so we introduce some augmentations to the basic algorithm.

- To prevent loops, DFS keeps track of a "color" attribute for each vertex. Unvisited vertices are white by default. Vertices that have been visited but still may be backtracked to are colored gray. Vertices which are completely processed are colored black. The algorithm can then prevent loops by skipping non-white vertices
- Instead of just marking visited vertices, the algorithm also keeps track of the tree generated by the depth-first traversal. It does so by marking the "parent" of each visited vertex, aka the vertex that DFS visited immediately prior to visiting the child.
- The augmented DFS also marks two auto-incrementing timestamps d and f to indicate when a node was first discovered and finished.

The algorithm takes as input a start vertex s and a starting timestamp t , and returns the timestamp at which the algorithm finishes. Let $N(s)$ denote the neighbors of s ; for a directed graph, let $N_{\text{out}}(s)$ denote the out-neighbors of s .

There are multiple ways we can search using DFS. One way is to search from some source node s , which will give us a set of black nodes reachable from s and white nodes unreachable from s .

Another way to use DFS is to search over the entire graph, choosing some white node and finding everything we can reach from that node, and repeating until we have no white nodes remaining. In an undirected graph this will give us all of the connected components.

Algorithm 1: $\text{init}(G)$

```
for  $v \in G$  do  
   $\text{color}(v) \leftarrow \text{white}$   
   $d(v), f(v) \leftarrow \infty$   
   $p(v) \leftarrow \text{nil}$   
end for
```

Algorithm 2: $\text{DFS}(s, t)$: $s \in V$ is white, $t = \text{time}$

```
 $\text{color}(s) \leftarrow \text{gray}$   
//  $d(s)$  is the discovery time of  $s$   
 $d(s) \leftarrow t$   
 $t \leftarrow t + 1$   
for  $v \in N(s)$  do  
  if  $\text{color}(v) = \text{white}$  then  
     $p(v) \leftarrow s$   
    // Update  $t$  to be the finish time of DFS starting at  $v$   
     $t \leftarrow \text{DFS}(v, t)$   
     $t \leftarrow t + 1$   
  end if  
end for  
Finish time:  
 $f(s) \leftarrow t$   
//  $s$  is finished  
 $\text{color}(s) \leftarrow \text{black}$ 
```

Algorithm 3: $\text{DFS}(s)$: DFS from a source node s

```
 $\text{init}(G)$   
 $\text{DFS}(s, 1)$ 
```

Algorithm 4: $\text{DFS}(G)$: DFS on an entire graph G

```
 $\text{init}(G)$   
 $t \leftarrow 1$   
for  $v \in G$  do  
  if  $\text{color}(v) = \text{white}$  then  
     $t \leftarrow \text{DFS}(v, t)$   
     $t \leftarrow t + 1$   
  end if  
end for
```

2.1 Runtime of DFS


We will now look at the runtime for the standard DFS algorithm (2).

Everything above the loop runs in $O(1)$ time per node visit. Excluding the recursive call, everything inside of the for loop takes $O(1)$ time every time an edge is scanned. Everything after the for loop also runs in $O(1)$ time per node visit.


We can express the runtime of DFS as $O(\# \text{ of node visits} + \# \text{ of edge scans})$. Assume we have a graph with n nodes and m edges. We know that the $\#$ of node visits is $\leq n$, since we only visit white nodes and whenever we visit a node we change its color from white to gray and never change it back to white again. We also know that an edge (u, v) is scanned only when u or v is visited. Since every node is visited at most once, we know that an edge (u, v) is scanned at most twice (or only once for directed graphs). Thus, $\#$ of edges scanned is $O(m)$, and the overall runtime of DFS is $O(m + n)$.

2.2 DFS Example

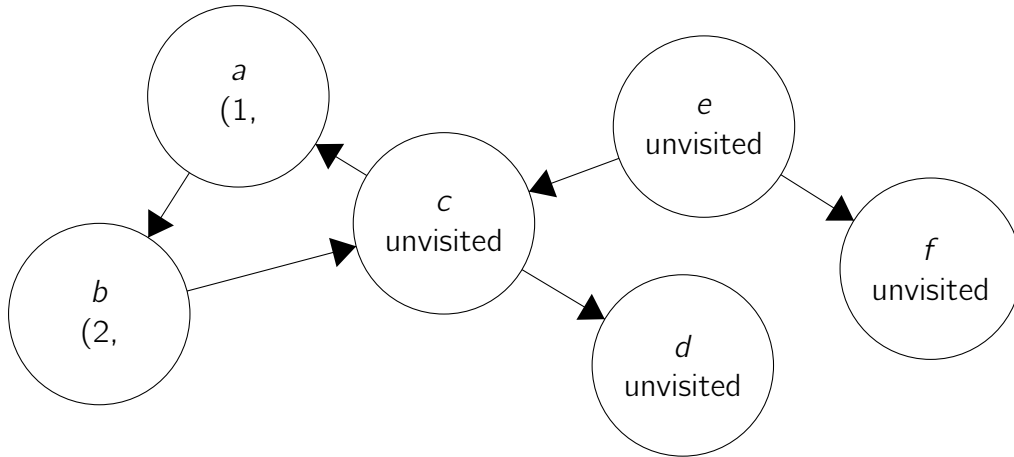
We will now try running DFS on the example graph below.



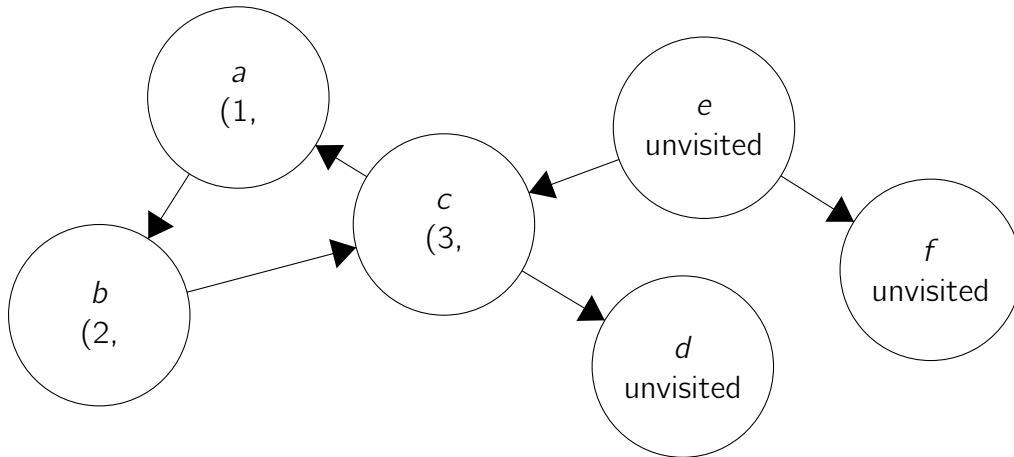
We mark all of the nodes as unvisited and start at a white node, in our case node a.




From node *a* we will visit all of *a*'s children, namely node *b*.




We now visit *b*'s child, node *c*.



Node *c* has two children that we must visit. When we try to visit node *a* we find that node *a* has already been visited (and would be colored gray, as we are in the process of searching *a*'s children), so we do not continue searching down that path. We will next search *c*'s second child, node *d*.



Since node d has no children, we return back to its parent node, c, and continue to go back up the path we took, marking nodes with a finish time when we have searched all of their children.



Once we reach our first source node a we find that we have searched all of its children, so we look in the graph to see if there are any unvisited nodes remaining. For our example, we start with a new source node e and run DFS to completion.

