Graphs : DFS + Properties[1]

1 Depth First Search (DFS)

We start graph algorithms with the pretty intuitive, but surprisingly powerful, *depth first search* (DFS). This algorithm solves the reachabilty problem, but then in one swoop solves much more. It also runs in $O(n+m)$ time. Let's get to it.

1.1 DFS from a vertex

Our journey starts from a given vertex v in the graph. As explorers in the past have done, we begin by first planting a "flag" at v. In algorithmic terms, we have a **global** Boolean variable visited$[x]$ for every vertex x initialized to 0 (false). When starting at v, our first action is to set visited$[v] = 1$. Subsequently, we investigate v's out-neighbors, and if we encounter any "unflagged" neighbor x, we proceed to x remembering (via a thread, say, to give a physical mnemonic) that we came from v, and then just repeat the procedure from x. If at some point we see that all of x's neighbors are flagged, then we ravel the metaphorical thread back to the place where x was called from (in this case v), and then continue the investigation of v's other neighbors. All this thread business is sweetly handled by recursion. Here is the pseudocode for DFS from a vertex.

```
1:  global visited[1 : n] initialized to 0.
2:  global F initialized to ⊥.
3:  procedure DFS(G, v): ▷ We assume V = {1, 2, ..., n}
4:        ▷ F will be an out-tree rooted at v
5:        visited[v] ← 1. ▷ Mark v visited.
6:        for u neighbor of v do: ▷ In an arbitrary, but fixed order
7:            if visited[u] = 0 then:
8:                Add edge (v, u) to F.
9:                DFS(G, u)
10:       return F
```

As you can see from the pseudocode, there is some extra stuff that the algorithm is building. Namely, the **global** object F. It begins with being $\bot$, and then in Line 8 edges are added to F. These edges are indeed the "threads" that were alluded to above. At the end of the call of $\text{DFS}(G, v)$, the algorithm returns this collection of edges. The following theorem shows that this collection indeed has a lot of structure, and contains information about the *connectivity* properties of G.

Theorem 1. A vertex $x \in F$ if and only if x is reachable from v in G. F is a *rooted out-tree* rooted at v.

[1] *Lecture notes by Deeparnab Chakrabarty. Last modified : 19th Mar, 2022*
These have not gone through scrutiny and may contain errors. If you find any, or have any other comments, please email me at deeparnab@dartmouth.edu. Highly appreciated!

Proof. First we prove if visited$[x] = 1$, then x is reachable from v. In fact we will show that x is reachable from v in F. The proof is by induction on the *time* at which visited$[x]$ was set to 1. Imagine every time the algorithm runs Line 5, we increment time by 1. At time 0, this was set visited$[v] = 1$ and v is reachable from v in F. Now pick a vertex x whose visited$[x] = 1$ is set at time t. This happens because of some $y \in V$ such that (a) $(y, x) \in E$, and (b) the run of DFS(G, y) calls DFS(G, x). In that case, (a) visited$[y] = 1$ has been set strictly before time t implying, by induction, y is reachable from v in F, and (b) we add edge (y, x) to F, which then implies x is reachable from v in F.

Now for the other direction. Suppose there exists a vertex x which is reachable from v in G but visited$[x] = 0$. Since x is reachable from v, there is a path $(v = v_0, v_1, \ldots, v_k = x)$ in G. Let us pick the *last* vertex v_i in this path which has visited$[v_i] = 1$; clearly $0 \le i < k$. Since visited$[v_i] = 1$, we have run DFS(G, v_i). But the for-loop in the algorithm would then call DFS(G, v_{i+1}) since visited$[v_{i+1}] = 0$. But that would set visited$[v_{i+1}] = 1$, and once visited a vertex is never "un-visited". This is a contradiction, and thus all vertices x reachable from v have visited$[x] = 1$.

The first part above shows that for every $x \in F$, there is a path from v to x in F. To show F is a rooted out-tree, we need to show that if we ignore the directions of the edges of F, the underlying undirected graph is a tree. Suppose not, and suppose there is a cycle C in the underlying undirected graph of F. Let the edge (x, y) be the *last* directed edge being added to this cycle. That is, DFS(G, x) called DFS(G, y). In particular, at this time, visited$[y] = 0$ (otherwise, the algorithm wouldn't call DFS(G, y)). However, there is an edge (y, z) in the (undirected) cycle C. That is, either $(y, z) \in F$ or $(z, y) \in F$. And this has occurred before DFS(G, x) is going to call DFS(G, y). However, the presence of this edge (either (y, z) or (z, y)) *implies* visited$[y] = 1$. This is a contradiction since visited$[y] = 0$. This proves that F is a rooted out-tree rooted at v. $\square$

Theorem 2. The REACHABLE? problem can be solved in $O(n + m)$ time.

Proof. Theorem 1 implies this as a corollary. Given vertex u and v, we can check if v is reachable from u by just running DFS(G, u), and checking if visited$[v] = 1$ or not. To get the path, we can use the tree F.

What is the running time of DFS? It is a recursive algorithm, so it is not completely trivial to see this. Let us fix an arbitrary vertex x, and let us figure out the time taken in the running of its pseudocode *other* than the DFS calls it makes. That is, the time for running Line 5 to Line 8. We see that the maximum time is taken in the for-loop, and this costs $O(1 + \deg^+(x))$ time (the $+1$ is to take care of the corner case when $\deg^+(x) = 0$). The second, and the key, observation is that in *all* the calls of DFS(G, v), a call DFS(G, x) for any vertex x is made at most once (exactly once for vertices reachable from v). This is because, once DFS(G, x) is called, visited$[x]$ is set to 1 which prevents any further calls. Thus, the total time taken by DFS(G, v) is at most $\sum_{x \in V}$ the total time taken by "non-recursive" calls of DFS(G, x), which is $\sum_{x \in V} O(1 + \deg^+(x))$. This evaluates to $O(n + m)$. $\square$

1.2 DFS on the whole graph

The next algorithm is a *traversal* over all vertices of the graphs using the subroutine DFS(G, v) repeatedly. This is called the *depth first traversal* algorithm of the graph G, but is also called the depth first search (or simply DFS) of G. The input to this algorithm is the graph G and a permutation/ordering σ of the vertices. This permutation tells the algorithm the order in which to "explore" vertices, that is, to run DFS(G, v).

The output to this algorithm has a lot of things; these objects contain surprising amounts of information about G, as we will see below.

- One output is a couple of vectors $\mathsf{first}[1 : n]$ and $\mathsf{last}[1 : n]$ where for any vertex v, $\mathsf{first}[v]$ notes the "time" at which the algorithm starts exploring from v, that is, $\mathrm{DFS}(G, v)$ is called, and $\mathsf{last}[v]$ denotes the "time" the exploring ends, that is, the subroutine $\mathrm{DFS}(G, v)$ terminates.

- The other output is a *"directed forest"* F spanning all the vertices of G. Each weakly connected component in F is a rooted out-tree, directed (even when G is undirected) away from the root. Together with this we store the scalar fcomp which counts the number of these weakly connected components of F, the array $\mathsf{root}[1 : \mathsf{fcomp}]$ where $\mathsf{root}(i)$ will store the root of the ith tree in F, and the array $\mathsf{Fcomp}[1 : n]$ where $\mathsf{Fcomp}[v]$ contains a number between 1 and fcomp indicating the tree in which v exists.

The algorithm is simple: it has a for-loop going over all vertices in the order σ; if the vertex is unvisited, then we run $\mathrm{DFS}(G, v)$ on it starting a new tree rooted from v. We end when there are no more vertices left. Here is the full pseudocode, where we have enhanced $\mathrm{DFS}(G, v)$ to take care of what we need.

```
 1: procedure DFS(G, σ[1 : n]): ▷ σ is an ordering of the vertices
 2:     global array visited[1 : n] initialized to all 0.
 3:     global array first[1 : n], last[1 : n], root[1 : n], Fcomp[1 : n] initialized to all 0.
 4:     global scalar fcomp, time initialized to 0.
 5:     global F initialized to ∅.

 6:     for v in σ do:
 7:         if visited[v] = 0 then:▷ v hasn't been visited yet:
 8:             fcomp ← fcomp + 1 ▷ Increase the number of trees in the forest
 9:             root[fcomp] ← v ▷ Set v to be the root of the new tree
10:             DFS(G, v)

11: procedure DFS(G, v):
12:     visited[v] ← 1
13:     Fcomp[v] ← fcomp. ▷ Set v's tree in the forest
14:     time ← time + 1.
15:     Set first[v] ← time. ▷ Start exploring.
16:     for u neighbor of v do: ▷ In an arbitrary, but fixed order
17:         if visited[u] = 0 then:
18:             Add edge (v, u) to the forest F.
19:             ▷ It will be added to the fcompth component.
20:             DFS(G, u)
21:     time ← time + 1.
22:     Set last[v] ← time.
```

As we argued in the case of $\mathrm{DFS}(G, v)$, one can see that $\mathrm{DFS}(G, \sigma)$ takes $O(n + m)$ time as well.

Claim 1. The running time of $\mathrm{DFS}(G, \sigma)$ is $O(m + n)$ for any σ.

Figure 1: The edges that appear in the forest are marked in solid, while the remaining edges are dotted. The first and last are noted near the vertices. In the third figure on the right, the interval is the $[\text{first}[v], \text{last}[v]]$ interval.
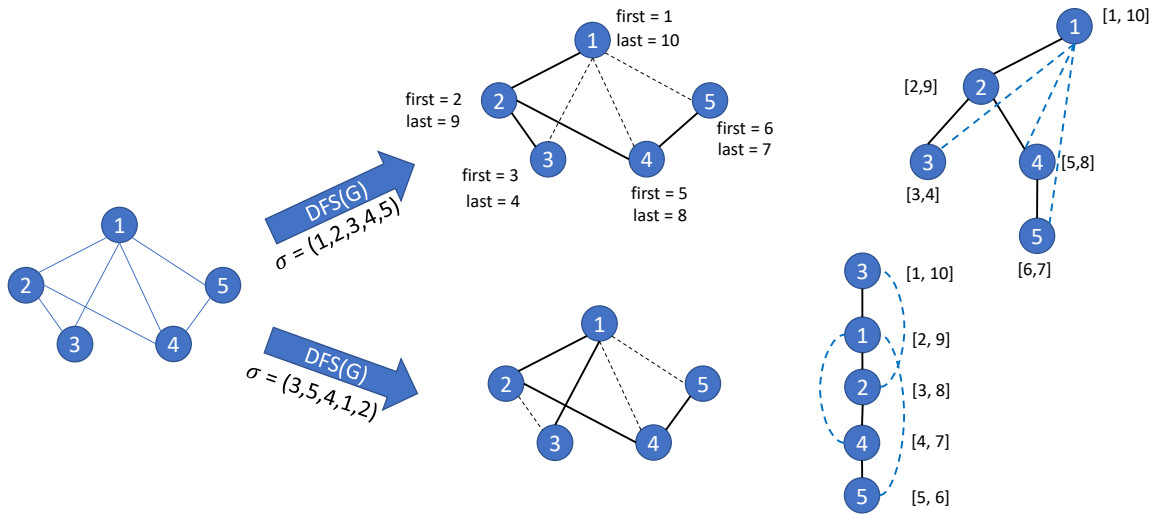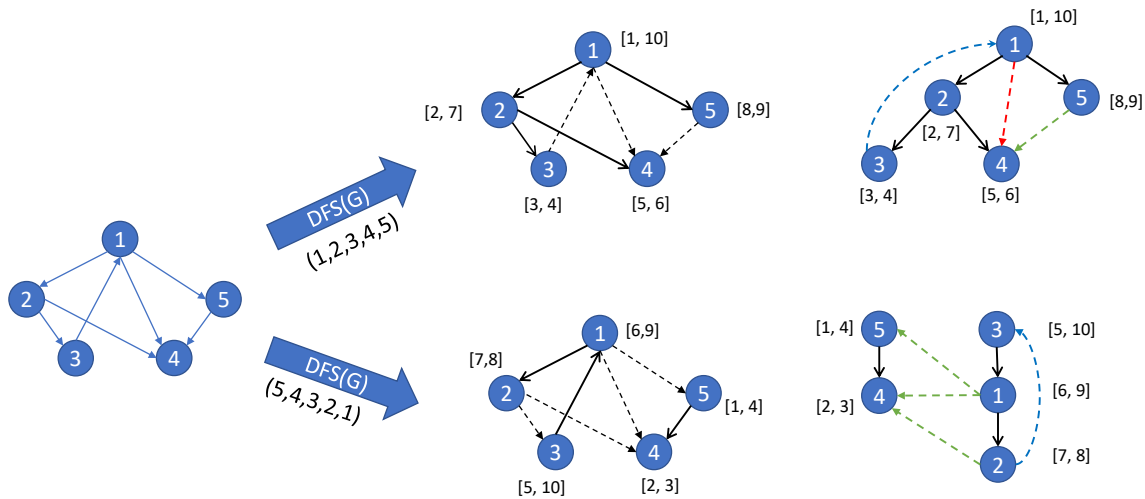

Figure 2: The edges that appear in the forest are marked in solid, while the remaining edges are dotted. The first and last are noted near the vertices. In the third figure on the right, the interval is the $[\text{first}[v], \text{last}[v]]$ interval.

Remark: *Different permutations can lead to different outcomes: see Figure 1 and Figure 2. This will be **critically** used in one application of DFS.*

Edge Classification. Running DFS on a graph G with any ordering σ leads to four kinds of edges.

4

- (Forest Edges.) These are the edges present in F. These are marked black and solid in the Figures.

- (Back Edges.) These edges go from a descendant to an ancestor. These are marked blue and dotted.

- (Forward Edges.) These edges go from an ancestor to a descendant. These are marked red and dotted. For undirected graphs the forward edges are all back edges (there is no direction).

- (Cross Edges.) All the rest. They can be among pairs in the same component, or not. These are marked green and dotted.

Properties. Next, we state and prove **three** properties of the output of the DFS algorithm. Before reading the proofs, it will be useful to see their illustrations in the examples shown in Figure 1 and Figure 2 (or any other figures you have privately made). For all the properties below, we assume we have run $\mathrm{DFS}(G, \sigma)$ for some **arbitrary** ordering σ.

> **Lemma 1** (**Nested Interval Property**). For any two vertices u and v, with $\mathrm{first}[u] < \mathrm{first}[v]$, exactly one of the following two properties hold.
>
> - $\mathrm{first}[u] < \mathrm{first}[v] < \mathrm{last}[v] < \mathrm{last}[u]$ *and v is a descendant of u in F.*
> - $\mathrm{first}[u] < \mathrm{last}[u] < \mathrm{first}[v] < \mathrm{last}[v]$ *and neither is a descendant of the other.*

This shows that the n intervals of the form $[\mathrm{first}[v], \mathrm{last}[v]]$ don't "criss-cross" (although one may be contained in the other). This property is called the *nested* property (also called laminar property).

Proof. Since $\mathrm{first}[u] < \mathrm{first}[v]$, we call $\mathrm{DFS}(G, u)$ before we call $\mathrm{DFS}(G, v)$. If v is ever discovered in the run of $\mathrm{DFS}(G, u)$, then (a) it will be a descendant of u in F (by Theorem 1), and (b) we must have $\mathrm{last}[v] < \mathrm{last}[u]$ since this recursive call must end before u's recursive call ends. This is case 1.

The only other case is v has not been discovered in the run of $\mathrm{DFS}(G, u)$. That is, $\mathrm{DFS}(G, u)$ ends before $\mathrm{DFS}(G, v)$ begins. Therefore, $\mathrm{last}[u] < \mathrm{first}[v]$ by definition of first and last. We also get v is not a descendant of u since all descendants of u are the vertices x on which we call DFS before $\mathrm{DFS}(G, u)$ ends. To end, u cannot be a descendant of v, because every descendant x of v must start DFS after v starts, that is, $\mathrm{first}[x] > \mathrm{first}[v]$. But, $\mathrm{first}[u] < \mathrm{first}[v]$. $\qquad\square$

The above property is useful, and will be useful in proving some other properties below. But it also allows us to classify the edges (not in the forest F) just looking at the first and the last values.

- (Back Edges.) Edges $(u, v) \in E \setminus F$ with $\mathrm{first}[v] < \mathrm{first}[u] < \mathrm{last}[u] < \mathrm{last}[v]$.
- (Forward Edges.) Edges $(u, v) \in E \setminus F$ with $\mathrm{first}[u] < \mathrm{first}[v] < \mathrm{last}[v] < \mathrm{last}[u]$.
- (Cross Edges.) Edges $(u, v) \in E$ such that the intervals $[\mathrm{first}[u], \mathrm{last}[u]]$ and $[\mathrm{first}[v], \mathrm{last}[v]]$ are disjoint.

> **Lemma 2** (**Edge Property**). Let (u, v) be any edge in G with $\mathrm{first}[u] < \mathrm{first}[v]$. Then, we must have $\mathrm{last}[v] < \mathrm{last}[u]$.

Proof. Suppose not. By the Nested Interval Property, we must have $\mathrm{first}[u] < \mathrm{last}[u] < \mathrm{first}[v] < \mathrm{last}[v]$. That happens when $\mathrm{DFS}(G, u)$ terminates before $\mathrm{visited}[v]$ is set to 1. But the Line 16 would discover v contradicting the above. $\qquad\square$

5

We are now ready for the first application of DFS – we can solve CONNECTED COMPONENTS of an Undirected Graph using the following lemma.

Lemma 3. Let $G = (V, E)$ be any undirected graph and consider the forest F returned by DFS(G, σ) with any permutation σ. The components of F are precisely the connected components of G.

Proof. Let $V_1, \ldots, V_k$ be the vertices in the various trees of the forest F. Clearly $G[V_i]$ is connected since they are connected in the forest. We claim that there is no edge of the form (u, v) with $u \in V_i$ and $v \in V_j$. Suppose there is, and without loss of generality assume first$[u] <$ first$[v]$ (this is where we are using the undirectedness of G). By the edge property, we have first$[u] <$ first$[v] <$ last$[v] <$ last$[u]$. But this means v is a descendant of u in F contradicting the fact they exist in different connected components of F. $\quad\square$

Theorem 3. CONNECTED COMPONENTS of an undirected graph can be found in $O(n + m)$ time by running DFS(G, σ) for any ordering σ.

Moving on to more properties.

Lemma 4 (Path Property). If $(u = v_1, v_2, \ldots, v_k = v)$ is a path in G from u to v such that first$[u] <$ first$[v_i]$ for all $2 \leq i \leq k$, then, last$[v_i] <$ last$[u]$ for all $2 \leq i \leq k$.

In English, if there is a path from a vertex u to a vertex v such that u is the first vertex to be discovered among them, then all the vertices in the path are descendants of u in the DFS forest.

Proof. Suppose not. Choose the smallest $2 \leq i \leq k$ for which last$[u] <$ last$[v_i]$. By the choice of i, we get last$[v_{i-1}] <$ last$[u]$. Also note (v_{i-1}, v_i) is an edge.

Case 1: first$[v_{i-1}] <$ first$[v_i]$. In this case, the Edge Property would imply last$[v_i] <$ last$[v_{i-1}]$, and thus last$[v_i] <$ last$[u]$. Which is what we supposed wasn't true.

Case 2: first$[v_i] <$ first$[v_{i-1}]$. In that case, we see first$[u] <$ first$[v_i] <$ last$[u] <$ last$[v_i]$ which violates the Nested Interval Property. $\quad\square$

The above property allow us immediately to solve the CYCLE? problem. The following theorem implies the algorithm: run DFS(G, σ) and check if any of the edges is a *back edge* (which is one linear time scan over all the edges and checking the first and the last).

Lemma 5. A graph G is **acyclic** if and only if there are no back edges.

Proof. One direction is trivial – if G has a back edge, then there is clearly a cycle. If the back-edge is (u, v), then by definition there is a path from v to u using F-edges, and then take the (u, v) edge back.

The other direction is more interesting. If G has a cycle C with k vertices $(v_1, \ldots, v_k, v_1)$, then without loss of generality let v_1 be the vertex with the smallest first$[v_i]$ in this cycle. Since there is a path from v_1 to v_k, using the Path property and the fact that first$[v_1]$ is the smallest, we get first$[v_1] <$ first$[v_k] <$ last$[v_k] <$ last$[v_1]$. But this implies (v_k, v_1) is a back-edge. $\quad\square$

Theorem 4. CYCLE? can be solved in $O(n + m)$ time using DFS.