

Lecture Notes on Search in Graphs (Partial Draft)

15-122: Principles of Imperative Computation
Frank Pfenning, André Platzer, Rob Simmons

Lecture 24
April 23, 2013

1 Introduction


In this lecture we introduce *graphs*. Graphs provide a uniform model for many structures, for example, maps with distances or Facebook relationships. Algorithms on graphs are therefore important to many applications. They will be a central subject in the algorithms courses later in the curriculum; here we only provide a very small sample of graph algorithms.

2 Paths in Graphs

We start with *undirected graphs* which consist of a set V of *vertices* (also called *nodes*) and a set E of *edges*, each connecting two different vertices. In particular, these graphs have no edges from a node back to itself. A graph is connected if we can reach any vertex from any other vertex by following edges in either direction. In a *directed graph* edges provide a connection from one node to another, but not necessarily in the opposite direction. More mathematically, we say that the *edge relation* between vertices is *symmetric* for undirected graphs. In this lecture we only discuss undirected graphs, although directed graphs also play an important role in many applications.

The following is a simple example of a connected, undirected graph

with 5 vertices (A, B, C, D, E) and 6 edges (AB, BC, CD, AE, BE, CE).



A *path* in a graph is a sequence of vertices where each vertex is connected to the next by an edge. That is, a path is a sequence

$$v_0, v_1, v_2, v_3, \dots, v_l$$

of some length $l \geq 0$ such that there is an edge from v_i to v_{i+1} in the graph for each $i < l$. For example, all of the following are paths in the graph above:

$$\begin{aligned}
 &A - B - E - C - D \\
 &A - B - A \\
 &E - C - D - C - B \\
 &B
 \end{aligned}$$

The last one is a special case: The length of a path is given by the number of edges in it, so a node by itself is a path of length 0 (without following any edges). Paths always have a starting vertex and an ending vertex, which coincide in a path of length 0. We also say that a path connects its endpoints.

The *graph reachability problem* is to determine if there is a path connecting two given vertices in a graph. If we know the graph is connected, this problem is easy since one can reach any node from any other node. But we might refine our specification to request that the algorithm return not just a boolean value (reachable or not), but an actual path. At that point the problem is somewhat interesting even for connected graphs. Using our earlier terminology, a path from vertex v to vertex w is a *certificate* or *explicit evidence* for the fact that vertex w is reachable from another vertex v . It is easy to check whether the certificate is valid, since it is easy to check if each node in the path is connected to the next one by an edge. It is more difficult to produce such a certificate.

For example, the path

$$A - B - E - C - D$$

is a certificate for the fact that vertex D is reachable from vertex A in the above graph. It is easy to check this certificate by following along the path and checking whether the indicated edges are in the graph.

3 Implicit Graphs

There are many, many different ways to represent graphs. In some applications they are never explicitly constructed but remain implicit in the way the problem was solved. One such example was *peg solitaire*. The vertices of the graph implicit in this problem are *board positions*. There is an edge from A to B if we can make a move in position A to reach position B . Note that this implicit graph is actually a *directed graph* since the game does not allow us to undo a move we just made. The classical reachability question here would be if from some initial position we can reach another given final position. We actually solved a related question, namely if we can reach any of a number of alternative positions (those with exactly one peg) from a given initial position. We win the game if we can reach any of those positions with a single peg.

The reason why we did not explicitly construct the full graph is that for standard boards it is unreasonably large – there are too many reachable positions. Instead, we incrementally construct it as we search for a solution in the hope we can find a solution without ever generating all nodes. In some examples (like the standard English board), this hope was justified if we were lucky enough to pick a good move strategy. To make sure that unsolvable boards had no solution, however, we still had to visit every reachable position. Just because we have 3 pegs remaining with one attempt of trying to solve the board does not mean we could not have been more successful if we had moved the pegs around in a different way.

4 Explicit Graphs & Graph Interface

Sometimes, we *do* want to represent a graph as a specific set of edges and vertices, however. In the C code, we'll refer to our vertices with unsigned integers. A minimal interface for graphs would allow us to create and free graphs, check whether an edge exists in the graph, and add a new edge to the graph.

```
typedef unsigned int vertex;
typedef struct graph_header* graph;

graph graph_new(unsigned int numvert);
void graph_free(graph G);
unsigned int graph_size(graph G);
    // number of vertices in the graph
bool graph_hasedge(graph G, vertex v, vertex w);
    //@requires v < graph_size(G) && w < graph_size(w);
void graph_addedge(graph G, vertex v, vertex w);
    //@requires v < graph_size(G) && w < graph_size(w);
    //@requires !graph_hasedge(G, v, w);
```

We use the C0 notation for contracts on the interface functions here. Even though C compilers do not recognize the `//@requires` contract and will simply discard it as a comment, the contract still serves an important role for the programmer reading the program. For the graph interface, we decide that it does not make sense to add an edge into a graph when that edge is already there, hence the second requires.

With this minimal interface, we can create a graph for our running example (letting $A = 0$, $B = 1$, and so on).

```
graph G = graph_new(5);
graph_addedge(G, 0, 1);
graph_addedge(G, 1, 2);
graph_addedge(G, 2, 3);
graph_addedge(G, 0, 4);
graph_addedge(G, 1, 4);
graph_addedge(G, 2, 4);
```

5 Adjacency Matrices

There are two simple ways to implement the graph interface. One way is to represent the graph as a two-dimensional array that represents its edge relation. We can check if there is an edge from $B (= 1)$ to $D (= 3)$ by looking for a checkmark in row 1, column 3. In an undirected graph, the top-right half of this two-dimensional array will be a mirror image of the bottom-left,


because the edge relation is symmetric.

	A	B	C	D	E
A		✓			✓
B	✓		✓		✓
C		✓		✓	✓
D			✓		
E	✓	✓	✓		

This representation of a graph is called an *adjacency matrix*, because it is a matrix that stores which nodes are neighbors.

6 Adjacency Lists

The other classic representation of a graph is as an *adjacency list*. In an adjacency list representation, we have a one-dimensional array that looks much like a hash table. Each vertex has a spot in the array, and each spot in the array contains a linked list of all the other vertices connected to that vertex. Our running example would look like this as an adjacency list:



Adjacency lists and adjacency matrices have different tradeoffs in the time and space it takes to perform operations. In a *sparse* matrix, where there are lots of vertices and not a lot of edges, it usually makes more sense to use an adjacency list. In a *dense* matrix, where there are lots of edges, it may be more efficient to use adjacency matrix.

If we do use an adjacency list representation, it will often make sense to extend the interface to graphs to give the client access to this linked list of adjacent edges.

```

typedef struct adjlist_node adjlist;

struct adjlist_node {

```

```

    vertex vert;
    adjlist *next;
};

/* Returns a linked list of the neighbors of vertex v.
 * This adjacency list is owned by the graph and should
 * not be modified by the user.
 * @requires(v < graph_size(G)) */
adjlist *graph_connections(graph G, vertex v);

```

One way to implement the adjacency list version of graphs is as a pointer to a special kind of C struct, a struct with a *flexible array member*, i.e. its last field is an array whose length is only specified at runtime, not at the time of declaring the type of the struct. This is a struct with two fields: the first is a unsigned integer representing the actual size, and the second is an array of adjacency lists.

```

struct graph_header {
    unsigned int size;
    adjlist *adj[]; // Flexible array member!
};

```

The array `adj` of adjacency lists will be contiguous in memory with the `size` field – this is quite different than, say, a hashtable chain, which is a linked list with data fields of type `elem` and `next` pointers. We allocate this adjacency list using `xcalloc` to make sure that the adjacency list is initialized to an empty array. Behind the scenes `xcalloc` just multiplies its two arguments; because we are allocating a struct with a flexible array member, we pass in 1 for the first argument and explicitly figure out the desired size of the array for the second argument:

```

graph graph_new(unsigned int size) {
    size_t adj_size = sizeof(adjlist*) * size;
    graph G = xcalloc(1, sizeof(struct graph_header) + adj_size);
    G->size = size;
    ENSURES(is_graph(G));
    return G;
}


```

7 Depth-First Search

The first algorithm we consider for determining if one vertex is reachable from another is called *depth-first search*.

Let's try to work our way up to this algorithm. Assume we are trying to find a path from u to w . We start at u . If it is equal to w we are done, because w is reachable by a path of length 0. If not we pick an arbitrary edge leaving u to get us to some node v . Now we have "reduced" the original problem to the one of finding a path from v to w .

The problem here is of course that we may never arrive at w even if there is a path. For example, say we want to find a path from A to D in our earlier example graph.




We can go $A - B - E - A - B - E - \dots$ without ever reaching D (or we can go just $A - B - A - B - \dots$), even though there exists a path.

We need to avoid repeating nodes in the path we are exploring. A *cycle* is a path of length 1 or greater that has the same starting and ending point. So another way to say we need to avoid repeating nodes is to say that we need to avoid cycles in the path. We accomplish this by *marking* the nodes we have already visited so when we visit them again we know not to consider them again.

Let's go back to the earlier example and play through this idea while trying to find a path from A to D . We start by marking A (indicated by hollowing the circle) and go to B . We indicate the path we have been following


by drawing a double-line along the edges contained in it.



When we are at B we mark B and have three choices for the next step.

1. We could go back to A , but A is already marked and therefore ruled out.
2. We could go to E .
3. We could go to C .

Say we pick E . At this point have again three choices. We might consider A as a next node on the path, but it is ruled out because A has already been marked. We show this by dashed the edge from A to E to indicate it was considered, but ineligible. The only possibility now is to go to C , because we have been at B as well (we just came from B).




From C we consider the link to D (before considering the link to B) and we arrive at D , declaring success with the path

$$A - B - E - C - D$$


which, by construction, has no cycles.

There is one more consideration to make, namely what we do when we get stuck. Let's reconsider the original graph



and the goal to find a path from E to B . Let's say we start $E - C$ and then $C - D$. At this point, all the vertices we could go to (which is only C) have already been marked! So we have to *backtrack* to the most recent choice point and pursue alternatives. In this case, this could be C , where the only remaining alternative would be B , completing the path $E - C - B$. Notice that when backtracking we have to go back to C even though it is already marked.

Depth-first search is characterized not only by the marking, but also that we always return to our most recent choice and follow a different path. When no other alternatives are available, we backtrack further. Let's consider the following slightly more larger graph, where we explore the outgoing edges using the alphabetically last label first.



We write the current node we are visiting on the left and on the right a *stack* of nodes we have to return to when we backtrack. For each of these we also remember which choices remain (in parentheses). We annotate marked nodes with an asterisk, which means that we never pick them as the next node to visit. In particular, when our current node has been marked, then we have been there and did not find the goal (yet), so we do not explore again. Hence, we do not care what their neighbors are now, because those had already been put on the stack earlier and will be or will have been considered at some point.

For example, at step 5 we do not consider E^* but go to D instead. We backtrack, when no unmarked neighbors remain to the current node.

Step	Current	Neighbors	Stack	Remark
1	A	(E, B)		
2	E	(C, B, A*)	A* (B)	
3	C	(G, E*, D)	E* (B, A*) A* (B)	
4	G	(C*)	C* (E*, D) E* (B, A*) A* (B)	
5	C*	don't care	G* () C* (E*, D) E* (B, A*) A* (B)	Backtrack
6	E*	don't care	C* (D) E* (B, A*) A* (B)	
7	D	(F, C*)	C* () E* (B, A*) A* (B)	
8	F	(D*)	D* (C*) C* () E* (B, A*) A* (B)	
9	D*	don't care	F* () D* (C*) C* () E* (B, A*) A* (B)	Backtrack
10	C*	don't care	D* () C* () E* (B, A*) A* (B)	Backtrack ×2
11	B	(A*)	E* (B, A*) A* (B)	Goal Reached

We are keeping the visited nodes on a stack so we can easily return to the most recent one.

```
bool dfs(graph G, bool *mark, vertex start, vertex target) {
    REQUIRES(G != NULL && mark != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));
    if (mark[start]) return false;
    mark[start] = true;

    printf("Visiting %d\n", start);
    if(start == target) return true;

    for(adjlist *L = graph_connections(G, start); L != NULL; L = L->next) {
        if(dfs(G, mark, L->vert, target)) return true;
    }

    return false;
}
```

Note that this recursive implementation of DFS uses the (implicit) stack of the function calls to `dfs` and each `dfs` function body has its own linked list for the adjacency list. In effect, that gives the search management data structure the form of a stack of queues (see Clac lab) as indicated in the example above. The stack elements are separated by `|` and the elements of the queue are wrapped in parentheses (B, A^*) etc.

8 Depth-First Search with a single stack

When scrutinizing the above example, we notice that the sophisticated data structure of a stack of queues was really quite unnecessary for DFS. The recursive implementation is simple and elegant, but its effect on the data management is nontrivial.

This can all be simplified by making the stack explicit. In that case there is a single stack with all the nodes on it that we still need to look at.

Step	Current	Neighbors	+ Old stack	= New stack
1	A	(E, B)	()	(E, B)
2	E	(C, B, A*)	(B)	(C, B, A*, B)
3	C	(G, E*, D)	(B, A*, B)	(G, E*, D, B, A*, B)
4	G	(C*)	(E*, D, B, A*, B)	(C*, E*, D, B, A*, B)
5	C*	don't care	(E*, D, B, A*, B)	(E*, D, B, A*, B)
6	E*	don't care	(D, B, A*, B)	(D, B, A*, B)
7	D	(F, C*)	(B, A*, B)	(F, C*, B, A*, B)
8	F	(D*)	(C*, B, A*, B)	(D*, C*, B, A*, B)
9	D*	don't care	(C*, B, A*, B)	(C*, B, A*, B)
10	C*	don't care	(B, A*, B)	(B, A*, B)
11	B	goal reached	(A*, B)	

```

bool dfs_iter1(graph G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    stack S = stack_new();
    push(S, (void*)(uintptr_t)start);

    bool mark[graph_size(G)];
    for(unsigned int i = 0; i < graph_size(G); i++)
        mark[i] = false;

    while(!stack_empty(S)) {
        vertex v = (vertex)(uintptr_t) pop(S);
        if(!mark[v]) {
            printf("Visiting %d\n", v);
            mark[v] = true;
            if (v == target) {
                stack_free(S, NULL);
                return true;
            }
        }

        for(adjlist *L = graph_connections(G, v); L != NULL; L = L->next)
            push(S, (void*)(uintptr_t)L->vert);
    }
}

```

```
    stack_free(S, NULL);  
    return false;  
}
```

We explicitly put the starting node on the stack. Then, every time we pop an element of the stack, we check to see whether it has been marked already (like we did in the iterative implementation). If it wasn't, we visit the node by marking it and comparing it to the target. Ultimately, we push all neighbors on the stack to make sure we look at them later.

9 Breadth-First Search

The iterative DFS algorithm managed his agenda, i.e. the list of nodes it still had to look at using a stack. But there's no reason to insist on a stack for the purposes of that. What happens if we replace the data management by a queue instead? All of a sudden, we will no longer explore the most recently found neighbor first as in depth-first search, but, instead, we will look at the oldest neighbor first. This corresponds to a breadth-first search where you explore the graph layer by layer. So BFS completes a layer of the graph before proceeding to the next layer. The code for that and many other interesting variations of graph search can be found on the web page.