


Graph Search

Review

- Graphs


- Vertices, edges, neighbors, ...
- Dense, sparse



- Adjacency matrix implementation

	0	1	2	3	4
0		✓			✓
1	✓		✓		✓
2		✓		✓	✓
3			✓		
4	✓	✓	✓		

- Adjacency list implementation



graph.h

```

typedef unsigned int vertex;
typedef struct graph_header *graph_t;

graph_t graph_new(unsigned int numvert);
//@ensures \result != NULL;

void graph_free(graph_t G);
//@requires G != NULL;

unsigned int graph_size(graph_t G);
//@requires G != NULL;

bool graph_hasedge(graph_t G, vertex v, vertex w);
//@requires G != NULL;
//@requires v < graph_size(G) && w < graph_size(G);

void graph_addedge(graph_t G, vertex v, vertex w);
//@requires G != NULL;
//@requires v < graph_size(G) && w < graph_size(G);
//@requires v != w && !graph_hasedge(G, v, w);

-----

typedef struct neighbor_header *neighbors_t;

neighbors_t graph_get_neighbors(graph_t G, vertex v);
//@requires G != NULL && v < graph_size(G);
//@ensures \result != NULL;

bool graph_ismore_neighbors(neighbors_t nbors);
//@requires nbors != NULL;

vertex graph_next_neighbor(neighbors_t nbors);
//@requires nbors != NULL;
//@requires graph_ismore_neighbors(nbors);
//@ensures is_vertex(\result);

void graph_free_neighbors(neighbors_t nbors);
//@requires nbors != NULL;
    
```

Review

- Costs are similar for dense graphs
- AL is **more space-efficient** for sparse graphs
 - very common graphs
 - $e \in O(v)$ is typical

	Adjacency list	Adjacency matrix
Space	$O(v + e)$	$O(v^2)$
graph_new	$O(v)$	$O(v^2)$
graph_free	$O(v + e)$	$O(1)$
graph_size	$O(1)$	$O(1)$
graph_hasedge	$O(\min(v,e))$	$O(1)$
graph_addedge	$O(1)$	$O(1)$
graph_get_neighbors	$O(1)$	$O(v)$
graph_ismore_neighbors	$O(1)$	$O(1)$
graph_next_neighbor	$O(1)$	$O(1)$
graph_free_neighbors	$O(1)$	$O(\min(v,e))$

Assuming the neighbors are represented as a linked list

Review

- Typical function that **traverses** a graph
 - go over most vertices and edges

	Cost	Tally
<code>void graph_print(graph_t G) {</code>		
<code>for (vertex v = 0; v < graph_size(G); v++) {</code>	v times	
<code>printf("Vertices connected to %u: ", v);</code>	O(1)	O(v)
<code>neighbors_t nbors = graph_get_neighbors(G, v);</code>	O(1)	O(v)
<code>while (graph_ismore_neighbors(nbors)) {</code>	} O(e) altogether	
<code>vertex w = graph_next_neighbor(nbors);</code>		
<code>printf(" %u,", w);</code>		O(v + e)
<code>}</code>		
<code>graph_free_neighbors(nbors);</code>	O(1)	O(v + e)
<code>printf("\n");</code>	O(1)	O(v + e)
<code>}</code>		
<code>}</code>		

- Adjacency list: $O(v + e)$
- Adjacency matrix: $O(v^2)$


Same as
space bounds

AL is much better for
sparse graphs

Graph Connectivity

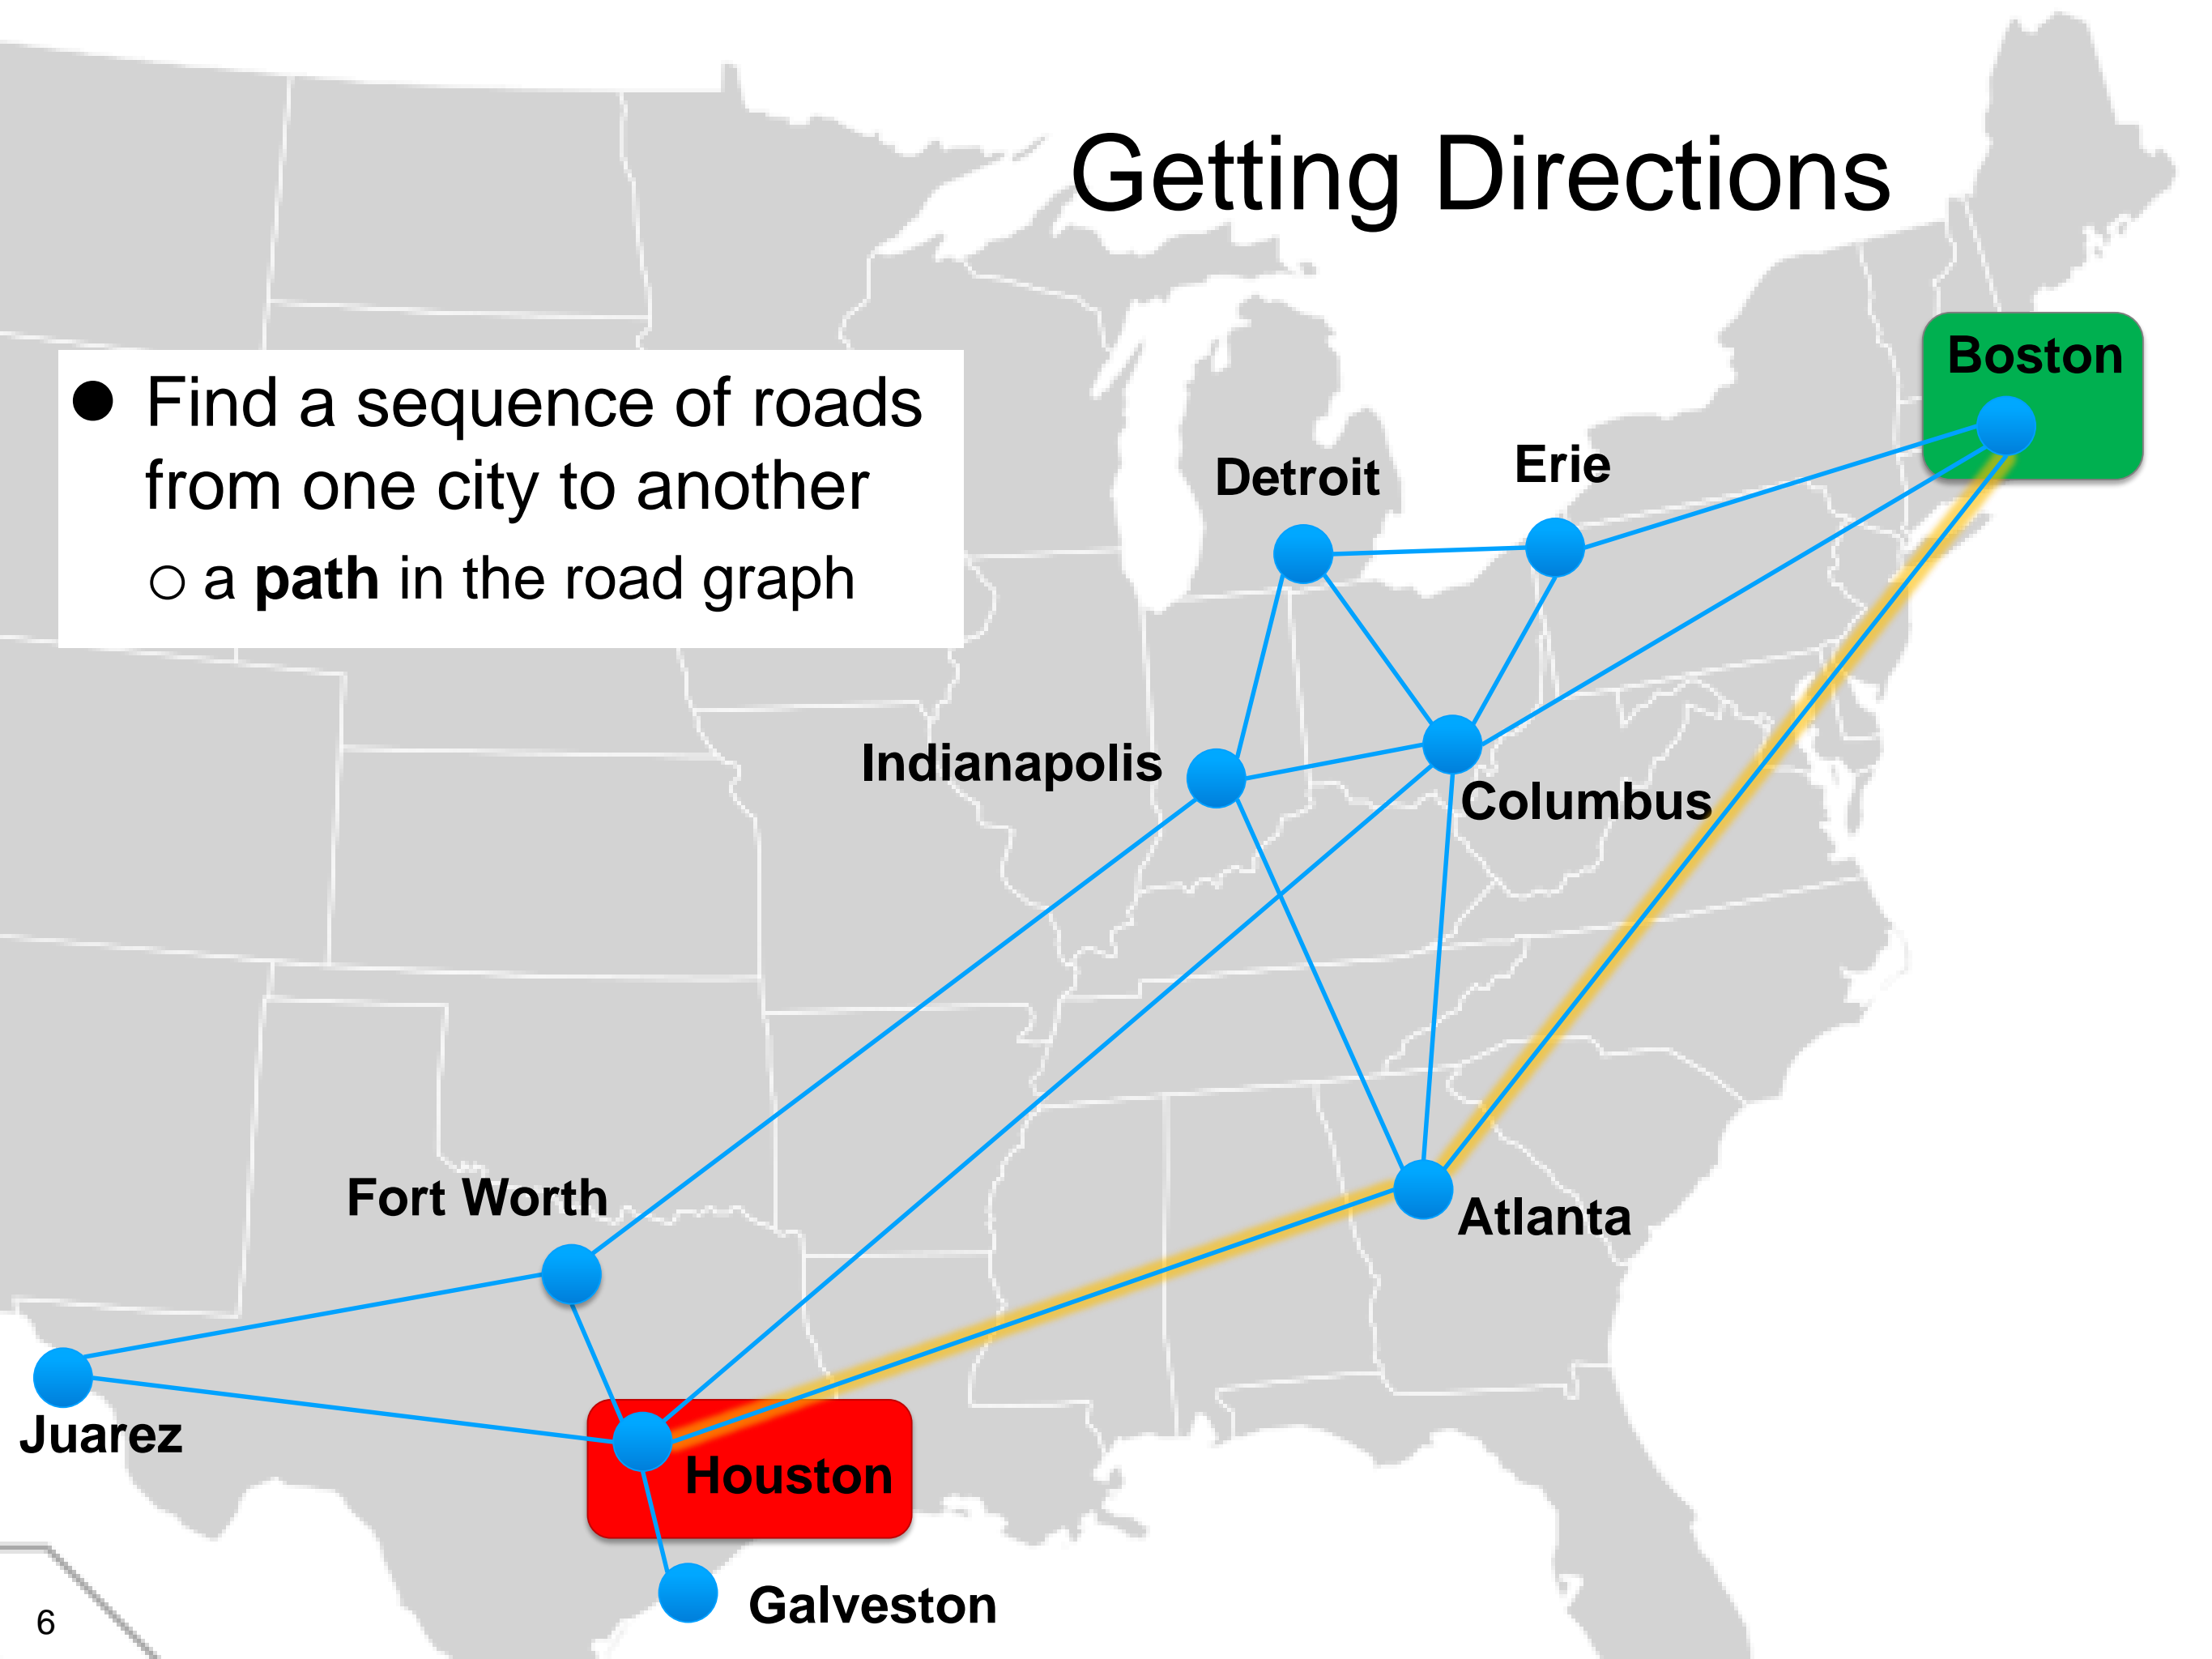
Solving Lightsout

- Find a sequence of moves from the given configuration to the solved configuration
 - a **path** in the lightsout graph



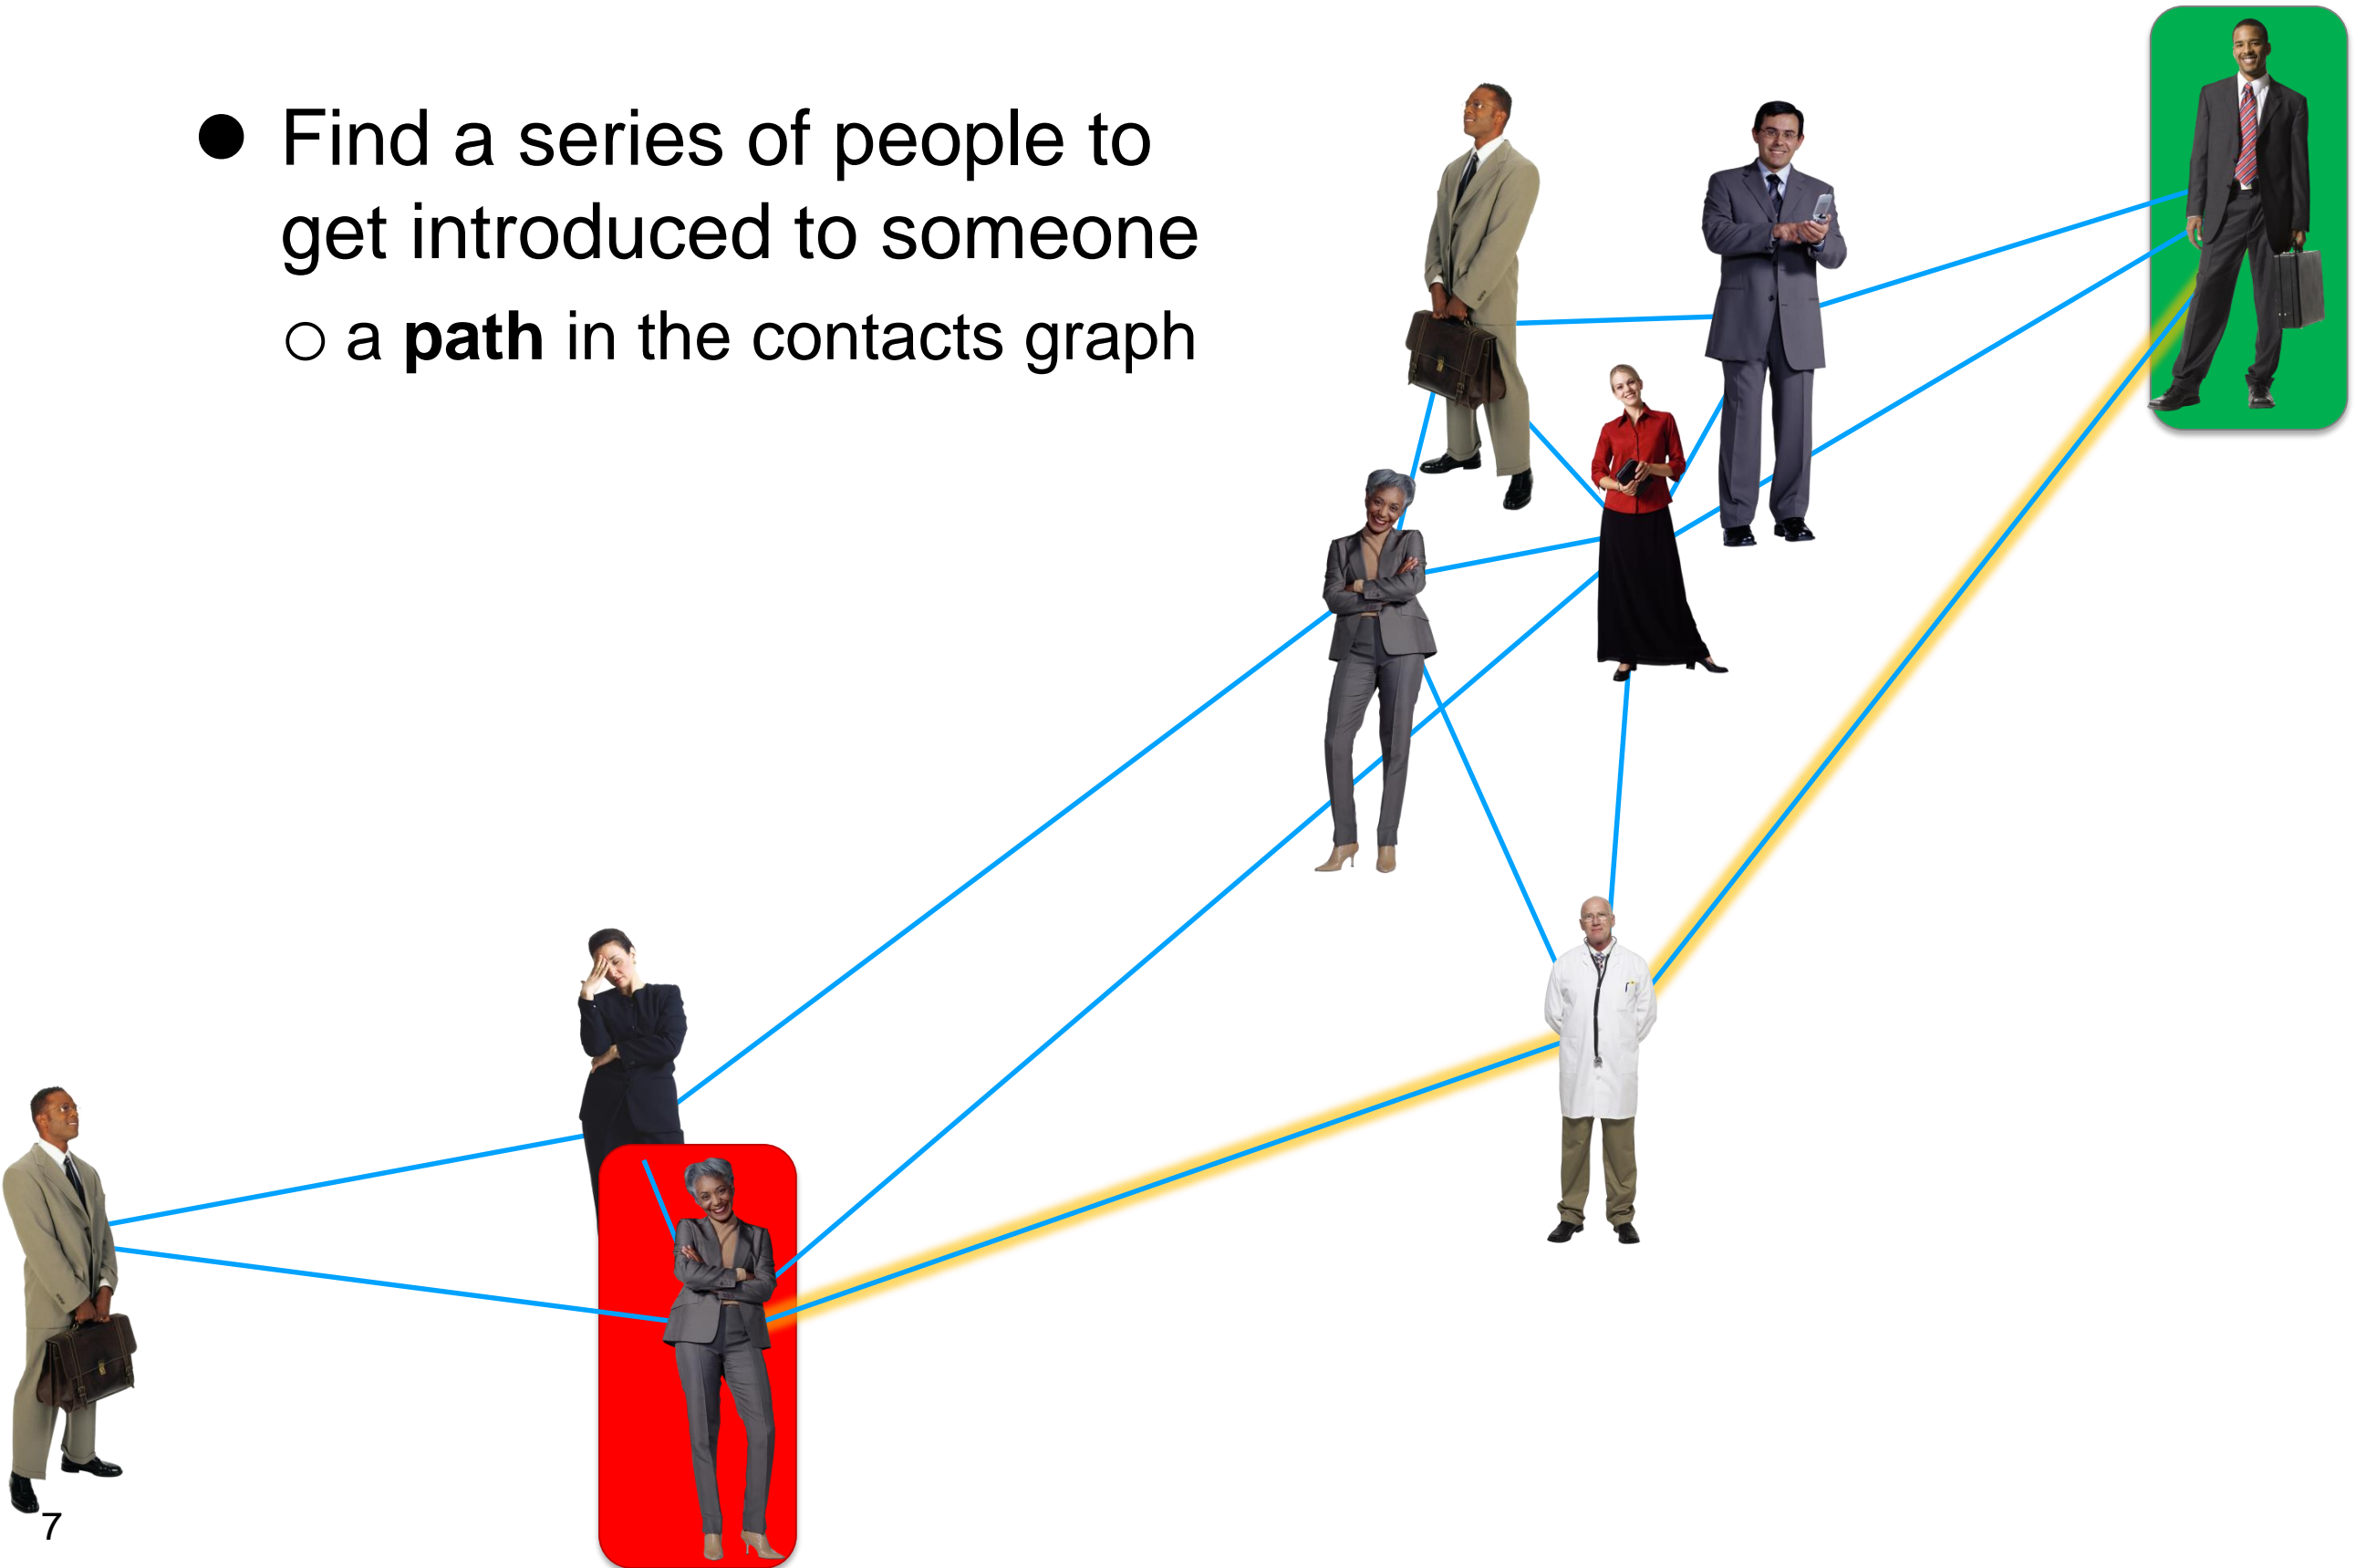
Getting Directions

- Find a sequence of roads from one city to another
- a **path** in the road graph



Getting Introduced


- Find a series of people to get introduced to someone
 - a **path** in the contacts graph



Connected Vertices

- A **path** is a sequence of vertices linked by edges

- 0-4-5-1 is a path between 0 and 1



- Two vertices are **connected** if there is a path between them


- 0 and 1 are connected

- 0 and 7 are not connected

- If v_1 and v_2 are connected, then v_2 is **reachable** from v_1

- A **connected component** is a maximal set of vertices that are connected


- this graph has two connected components



Checking Reachability

- How do we check if two vertices are connected?
 - `graph_hasedge` only tells us if they are *directly* connected
 - by an edge
 - We want to develop a general algorithm to check reachability
 - then we can use it to check reachability in any domain
 - ❑ to check if lightsout is solvable from a given board
 - ❑ to figure out if there are roads between two cities
 - ❑ to know if there is any social connection between two people

The rest of this lecture



Finding Paths

- How do we find a path between two vertices?
 - what is a solution to lightsout from a given board?
 - what roads are there between two cities?
 - what series of people can get me introduced to person X?
- an algorithm that checks reachability can be instrumented to report a path between the two vertices

We will limit ourselves to reachability

- A path is a **witness** that two vertices are connected
 - Finding a witness is called a **search problem**
 - Checking a witness is called a **verification problem**
 - checking that a witness is valid is often a lot easier than finding a witness

This is the basic principle underlying **cryptology**

Checking Reachability

- Let's define reachability mathematically


There is a path from **start** to **target** if

- start** == **target**, or
- there is an edge from **start** to some vertex **v** and there is a path from **v** to **target**


This is an inductive definition

base case

inductive case



There is a path from 0 to 0



There is a path from 0 to 3

Recursive Depth-first Search – I

Implementing the Definition

- We can immediately transcribe this inductive definition into a recursive **client-side** function

There is a path from **start** to **target** if

- **start** == **target**, or
- there is an edge from **start** to some vertex **v** and there is a path from **v** to **target**

```
bool naive_dfs(graph_t G, vertex start, vertex target) {  
    REQUIRES(G != NULL);  
    REQUIRES(start < graph_size(G) && target < graph_size(G));  
  
    // there is a path from start to target if  
    // target == start, or  
  
    // there is an edge from start to ...  
  
    // ... some vertex v ...  
  
    // ... and there is a path from v to target  
}
```

Contracts

Implementing the Definition

```
bool naive_dfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));
    printf("    Visiting %u\n", start);

    // there is a path from start to target if
    // target == start, or
    if (target == start) return true;
    // there is an edge from start to ...
    neighbors_t nbors = graph_get_neighbors(G, start);
    while (graph_ismore_neighbors(nbors)) {
        // ... some vertex v ...
        vertex v = graph_next_neighbor(nbors);
        // ... and there is a path from v to target
        if (naive_dfs(G, v, target)) {
            graph_free_neighbors(nbors);
            return true;
        }
    }
    graph_free_neighbors(nbors);
    return false;
}
```

graph.h

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

graph_t graph_new(unsigned int numvert);
//@ensures \result != NULL;

void graph_free(graph_t G);
//@requires G != NULL;

unsigned int graph_size(graph_t G);
//@requires G != NULL;

bool graph_hasedge(graph_t G, vertex v, vertex w);
//@requires G != NULL;
//@requires v < graph_size(G) && w < graph_size(G);

void graph_addege(graph_t G, vertex v, vertex w);
//@requires G != NULL;
//@requires v < graph_size(G) && w < graph_size(G);
//@requires v != w && !graph_hasedge(G, v, w);

-----

typedef struct neighbor_header *neighbors_t;

neighbors_t graph_get_neighbors(graph_t G, vertex v);
//@requires G != NULL && v < graph_size(G);
//@ensures \result != NULL;

bool graph_ismore_neighbors(neighbors_t nbors);
//@requires nbors != NULL;

vertex graph_next_neighbor(neighbors_t nbors);
//@requires nbors != NULL;
//@requires graph_ismore_neighbors(nbors);
//@ensures is_vertex(\result);

void graph_free_neighbors(neighbors_t nbors);
//@requires nbors != NULL;
```

Implementing the Definition

```
bool naive_dfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));
    printf("    Visiting %u\n", start);

    // there is a path from start to target if
    // target == start, or
    if (target == start) return true;
    // there is an edge from start to ...
    neighbors_t nbors = graph_get_neighbors(G, start);
    while (graph_ismore_neighbors(nbors)) {
        // ... some vertex v ...
        vertex v = graph_next_neighbor(nbors);
        // ... and there is a path from v to target
        if (naive_dfs(G, v, target)) {
            graph_free_neighbors(nbors);
            return true;
        }
    }
    graph_free_neighbors(nbors);
    return false;
}
```

- It has the same structure as `graph_print`

```
void graph_print(graph_t G) {
    for (vertex v = 0; v < graph_size(G); v++) {
        printf("Vertices connected to %u: ", v);
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            printf(" %u,", w);
        }
        graph_free_neighbors(nbors);
        printf("\n");
    }
}
```


- the outer loop is replaced with recursion

Does it Work?

- Let's check there is a path from 3 to 0

start	target	nbors
3	0	2
2	0	1, 3, 4
1	0	0, 2, 4
0	0	✓

Assume the neighbors are returned from smallest to biggest



- Let's run it

```
Linux Terminal
# gcc ... lib/*.c connected.c main.c
# ./a.out 3 0
  Visiting 3
  Visiting 2
  Visiting 1
  Visiting 0
  Reachable
```

... from to

Looks good


```
bool naive_dfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));
    printf("  Visiting %u\n", start);

    // there is a path from start to target if
    // target == start, or
    if (target == start) return true;
    // there is an edge from start to ...
    neighbors_t nbors = graph_get_neighbors(G, start);
    while (graph_ismore_neighbors(nbors)) {
        // ... some vertex v ...
        vertex v = graph_next_neighbor(nbors);
        // ... and there is a path from v to target
        if (naive_dfs(G, v, target)) {
            graph_free_neighbors(nbors);
            return true;
        }
    }
    graph_free_neighbors(nbors);
    return false;
}
```

Does it *Always* Work?

- Let's check there is a path from 0 to 3

start	target	nbors
0	3	1, 4
1	3	0, 2, 4
0	3	1, 4
1	3	0, 2, 4
... (this is not promising) ...		



- Let's run it

```
Linux Terminal
# gcc ... lib/* .c connected.c main.c
# ./a.out 0 3
  Visiting 0
  Visiting 1
  Visiting 0
```

runs forever!

```
bool naive_dfs(graph_t G, vertex start, vertex target) {
  REQUIRES(G != NULL);
  REQUIRES(start < graph_size(G) && target < graph_size(G));
  printf("  Visiting %u\n", start);

  // there is a path from start to target if
  // target == start, or
  if (target == start) return true;
  // there is an edge from start to ...
  neighbors_t nbors = graph_get_neighbors(G, start);
  while (graph_ismore_neighbors(nbors)) {
    // ... some vertex v ...
    vertex v = graph_next_neighbor(nbors);
    // ... and there is a path from v to target
    if (naive_dfs(G, v, target)) {
      graph_free_neighbors(nbors);
      return true;
    }
  }
  graph_free_neighbors(nbors);
  return false;
}
```

It does not Work

- Either the definition is wrong or the code is wrong

- Definition

- it magically picks the right neighbor v if there is one
 - the magic of “*there is ...*”



The definition is fine

- Code

- it must examine the neighbors in some order
 - the first v may not be the right one

There is a path from **start** to **target** if


- **start** == **target**, or
- there is an edge from **start** to some vertex v and there is a path from v to **target**

```
bool naive_dfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_si...
    printf("  Visiting %u\n", start);

    // there is a path from start to target if
    // target == start, or
    if (target == start) return true;
    // there is an edge from start to ...
    neighbors_t nbors = graph_get_neighbors(G, start);
    while (graph_ismore_neighbors(nbors)) {
        // ... some vertex v ...
        vertex v = graph_next_neighbor(nbors);
        // ... and there is a path from v to target
        if (naive_dfs(G, v, target)) {
            graph_free_neighbors(nbors);
            return true;
        }
    }
    graph_free_neighbors(nbors);
    return false;
}
```

Why doesn't it Work?

- The code examines the neighbors in some order
 - it always starts with the same v
 - the first neighbor
 - ... even if it has been examined before



- The code will never visit the second neighbor (if there is one)
 - it charges ahead with the first neighbor, always
 - if there is a path by only examining first neighbors, it will find it
 - if the path involves some other neighbor, it won't

start	target	nbors
0	3	1 4
1	3	0 2, 4
0	3	1 4
1	3	0 2, 4
...		

Recursive Depth-first Search – II

Fixing the Code

- Problems: the code examines the same neighbors over and over
- Solution: **mark** vertices that are being examined
 - only examine a vertex if it is unmarked
 - mark it right away
- How to mark vertices?
 - carry around an array of booleans
 - true = marked
 - false = unmarked



We could use any implementation of **sets**,
e.g., hash sets

Fixing the code

- Carry around an array of booleans
- Only run if start is unmarked
- Mark it right away
- Only examine a neighbor if it's unmarked
 - we need to guard the recursive call

```
bool dfs_helper(graph_t G, bool *mark, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));
    REQUIRES(!mark[start]);

    mark[start] = true;
    printf("    visiting %u\n", start);

    // there is a path from start to target if
    // target == start, or
    if (target == start) return true;
    // there is an edge from start to ...
    neighbors_t nbors = graph_get_neighbors(G, start);
    while (graph_ismore_neighbors(nbors)) {
        // ... some vertex v ...
        vertex v = graph_next_neighbor(nbors);
        // ... and there is a path from v to target
        if (!mark[v] && dfs_helper(G, mark, v, target)) {
            graph_free_neighbors(nbors);
            return true;
        }
    }
    graph_free_neighbors(nbors);
    return false;
}
```

Fixing the Code

- We have modified the prototype of the function
 - but the client should not have to deal with the added details
 - export a **wrapper** instead of `dfs_helper`

```
bool dfs (graph_t G, vertex start, vertex target) {  
    REQUIRES(G != NULL);  
    REQUIRES(start < graph_size(G) && target < graph_size(G));  
  
    bool *mark = xcalloc(graph_size(G), sizeof(bool));  
    bool connected = dfs_helper(G, mark, start, target);  
    free(mark);  
    return connected;  
}
```

Create the mark array:
`calloc` initializes all
positions to false

We must free mark
since we calloc'ated it

An Alternative Wrapper

- We can also use a *stack-allocated array* for mark

```
bool dfs(graph_t G, vertex start, vertex target) {  
    REQUIRES(G != NULL);  
    REQUIRES(start < graph_size(G) && target < graph_size(G));  
  
    bool mark[graph_size(G)];  
    for (unsigned int v = 0; v < graph_size(G); v++)  
        mark[v] = false;  
  
    return dfs_helper(G, mark, start, target);  
}
```

Create the a stack allocated array
of size graph_size(G)

We need to initialize it explicitly

But we don't need to free it

- Is this version preferable?
 - stack space is limited
 - for a large graph, the stack may not be big enough
 - **stack overflow**





Does it Work?

- Let's check there is a path from 0 to 3

start	target	nbors	marked
0	3	1, 4	0
1	3	0, 2, 4	0, 1
2	3	1, 3, 4	0, 1, 2
3	3		✓

- Let's run it

```
Linux Terminal
# gcc ... lib/*.c connected.c main.c
# ./a.out 0 3
  Visiting 0
  Visiting 1
  Visiting 2
  Visiting 3
  Reachable
```



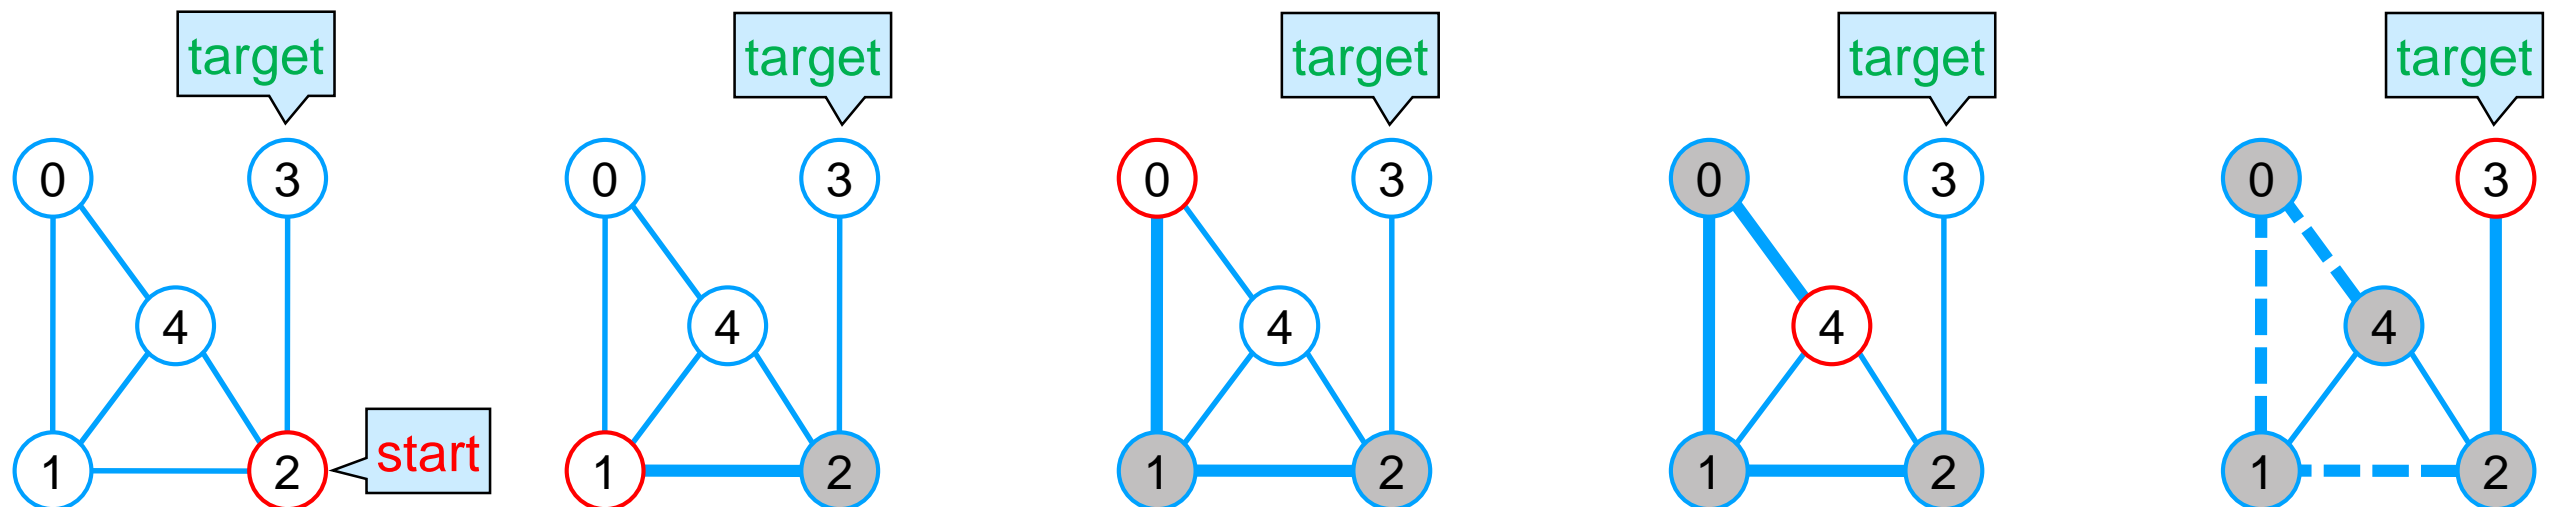
Backtracking

- Let's check there is a path from 2 to 3

start	target	nbors	marked
2	3	1, 3, 4	2
1	3	0, 2, 4	1, 2
0	3	1, 4	0, 1, 2
4	3	x 0, 1, 2	0, 1, 2, 4
3	3		✓

3 ≠ 4 and all the neighbors of 4 are marked

We **backtrack** to a vertex that has a still **unmarked neighbor** continue from it



Backtracking

- *We backtrack to a vertex that has a still **unmarked neighbor** and continue from it*

start	target	nbors	marked
2	3	① ③ 4	2
1	3	① 2, 4	1, 2
0	3	1, ④	0, 1, 2
4	3	✗ 0, 1, 2	0, 1, 2, 4
3	3		✓

- This is achieved by returning **false** from the recursive call
 - the caller will then try the next unmarked neighbor

- Let's run it

```
Linux Terminal
# gcc ... lib/*.c connected.c main.c
# ./a.out 2 3
  Visiting 2
  Visiting 1
  Visiting 0
  Visiting 4
  Visiting 3
  Reachable
```

```
...
while (graph_ismore_neighbors(nbors)) {
  // ... some vertex v ...
  vertex v = graph_next_neighbor(nbors);
  // ... and there is a path from v to target
  if (!mark[v] && dfs_helper(G, mark v, target)) {
    graph_free_neighbors(nbors);
    return true;
  }
}
graph_free_neighbors(nbors);
return false;
}
```

Complexity of dfs

- Let's call `dfs` on a graph with
 - v vertices,
 - e edges, and
 - implemented using adjacency lists

```
bool dfs (graph_t G, vertex start, vertex target) {  
    REQUIRES(G != NULL);  
    REQUIRES(start < graph_size(G) && target < graph_size(G));  
  
    bool *mark = xcalloc(graph_size(G), sizeof(bool));  
    bool connected = dfs_helper(G, mark, start, target);  
    free(mark);  
    return connected;  
}
```

O(v)

free has constant cost

- The cost of `dfs` is $O(v)$ plus the cost of `dfs_helper`

Complexity of `dfs_helper`

<code>graph_get_neighbors</code>	$O(1)$
<code>graph_ismore_neighbors</code>	$O(1)$
<code>graph_next_neighbor</code>	$O(1)$
<code>graph_free_neighbors</code>	$O(1)$

- The body of the loop runs at most $2e$ times **altogether**

- at most $2e$ calls to `graph_next_neighbors`

- e edges from either endpoint

- each endpoint is examined at most once

Just like for `graph_print`

- There are at most v recursive calls

In reality, it's more like $\min(v, e)$

- up to v vertices can be marked

- Every operation costs $O(1)$

- `dfs_helper` has cost $O(e + v)$

```

bool dfs_helper(graph_t G, bool *mark, vertex start, vertex target) {
    mark[start] = true;           O(1)           O(v)

    if (target == start) return true;           O(1)           O(v)

    neighbors_t nbors = graph_get_neighbors(G, start);           O(1)           O(v)
    while (graph_ismore_neighbors(nbors)) {           O(1)
        vertex v = graph_next_neighbor(nbors);           O(1)
        if (!mark[v] && dfs_helper(G, mark, v, target)) {           O(1)
            graph_free_neighbors(nbors);           O(1)
            return true;
        }
    }
    graph_free_neighbors(nbors);           O(1)           O(v + e)
    return false;
}
    
```

O(e) altogether

Tally

Complexity of dfs

graph_size	O(1)
graph_get_neighbors	O(1)
graph_ismore_neighbors	O(1)
graph_next_neighbor	O(1)
graph_free_neighbors	O(1)

- Let's call **dfs** on a graph with
 - v vertices,
 - e edges, and
 - implemented using adjacency lists

```
bool dfs (graph_t G, vertex start, vertex target) {  
    REQUIRES(G != NULL);  
    REQUIRES(start < graph_size(G) && target < graph_size(G));  
  
    bool *mark = xcalloc(graph_size(G), sizeof(bool));  
    bool connected = dfs_helper(G, mark, start, target);  
    free(mark);  
    return connected;  
}
```

O(v)

O(v + e)

- The cost of **dfs** is $O(v + e)$

Complexity of dfs

For a graph with v vertices and e edges

- $O(v + e)$ using the adjacency list implementation

Holds for both
sparse and dense
graphs

- $O(v^2)$ using the adjacency matrix implementation

Exercise

- AL is more efficient for sparse graphs
 - the most common kind of graphs

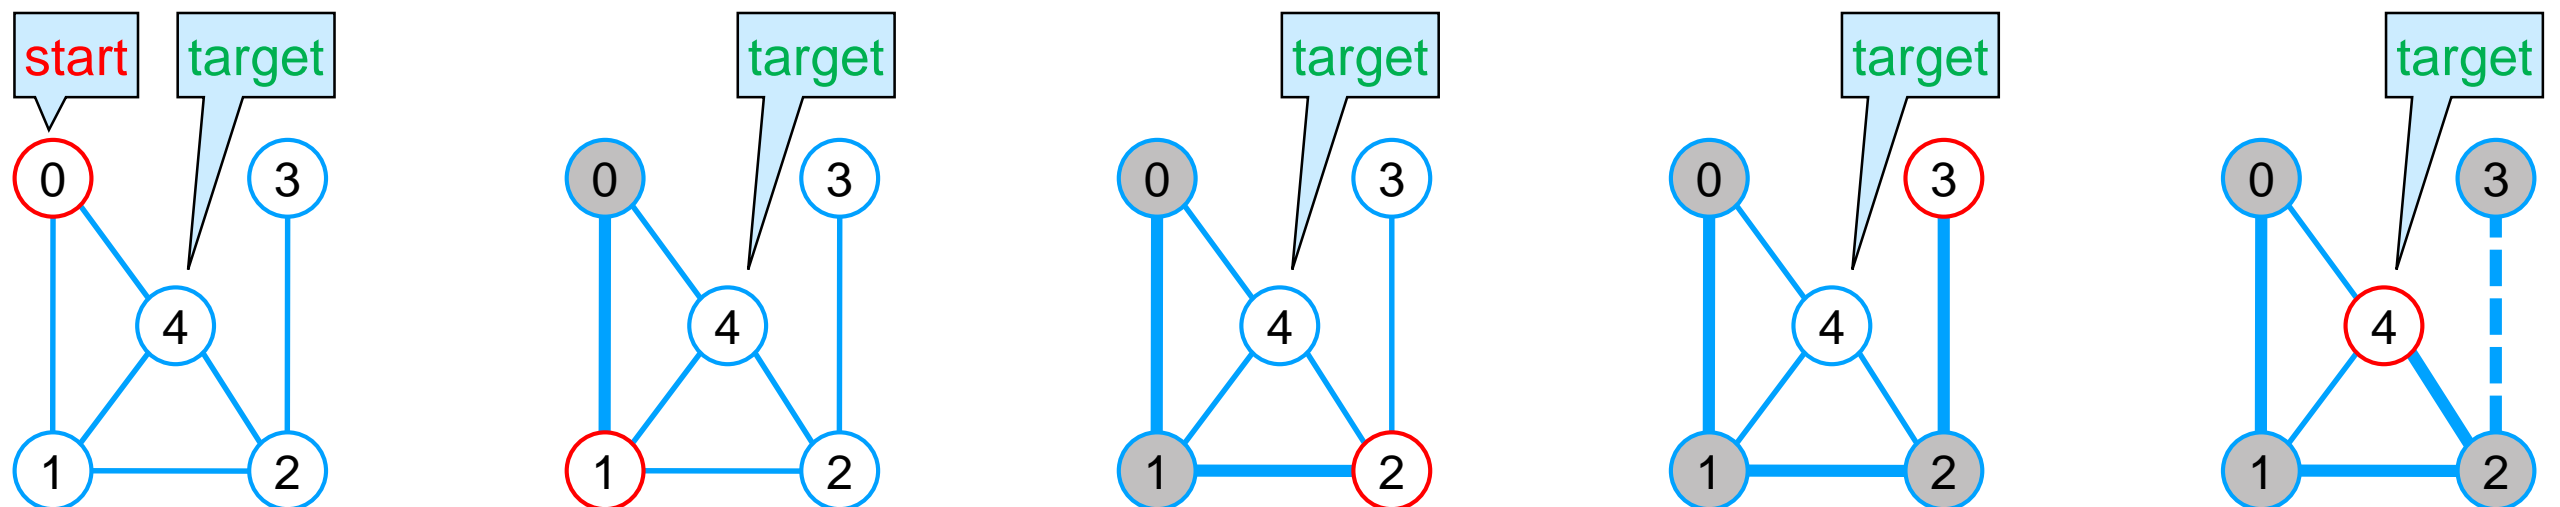
Moving forward, we will always
assume an adjacency list
implementation

Breadth-first Search

How does dfs Work?

- When calling **dfs** on 0 and 4, it finds the path 0–1–2–4
 - it also visits 3 and backtracks
- But there is a much shorter path: 0–4
 - **dfs** does more work than strictly necessary

start	target	nbors	marked
0	4	1, 4	0
1	4	0, 2, 4	0, 1
2	4	1, 3, 4	0, 1, 2
3	4	x 2	0, 1, 2, 3
4	4		








How does dfs Work?

- **dfs** charges ahead until
 - it finds the target vertex
 - or it hits a dead end
 - then it backtracks to the last choice point


start	target	nbors	marked
0	4	1, 4	0
1	4	0, 2, 4	0, 1
2	4	1, 3, 4	0, 1, 2
3	4	2	0, 1, 2, 3
4	4		✓

- This strategy is called **depth-first search** DFS



Breadth-first Search


- To find the shortest path, we need to explore the graph **level by level** from the start vertex
 - first look at the vertices 0 hops away from start,
 - if start == end
 - then look at the vertices 1 hop away from start
 - then 2 hops away
 - then 3 hops away
 - ...



- This strategy is called **breadth-first search** BFS

Breadth-first Search

- We need to traverse the graph **level by level**




- When we examine 0, we need to remember that we will have to examine 1 and 4 later
- When we examine 1, we need to remember we may have to examine 2 later
 - but first we need to look at 4

- We need a **todo list**

Breadth-first Search

- We need a **work list**


That's what we called
todo lists



- We need to traverse the graph level by level
 - finish examining the current level before starting the next level
 - we need to retrieve the vertices inserted the longest time ago
- This work list must be a **queue**
 - older nodes need to be visited before newer nodes

Breadth-first Search

- *This work list must be a queue*



- start with 0 in the queue
- at each step, retrieve the next vertex to examine

<i>next</i>	target	queue	marked
	4	0	0
0	4	1, 4	0, 1, 4
1	4	4, 2	0, 1, 4, 2
4	4		✓

- We mark the vertices so we don't put them in the queue twice
 - either because we examined them already
 - or because they are already in the queue and will be examined later

Implementing BFS

- We need
 - a **queue** where to store the vertices to examine next
 - a **mark array** where to track the vertices we know about
 - either already examined or queued up to be examined

Implementing BFS

- For as long as there are vertices still to be processed
 - retrieve the vertex v inserted in the queue the longest time ago
 - if v is **target**, we are done — **there is a path**
 - examine each neighbor w of v
 - if w is unmarked add it to the queue and mark it
 - otherwise ignore w – it was already queued up for processing
- if the queue is empty
 - there are no vertices left to process
 - and we have not found a path
 - we are done — **there is no path**

Implementing BFS – I

Initial setup

```
bool bfs(graph_t G, vertex start, vertex target) {  
    REQUIRES(G != NULL);  
    REQUIRES(start < graph_size(G) && target < graph_size(G));  
  
    if (start == target) return true;  
  
    // mark is an array containing only start  
    bool *mark = xcalloc(graph_size(G), sizeof(bool));  
    mark[start] = true;  
  
    // Q initially is a queue containing only start  
    queue_t Q = queue_new();  
    enq(Q, start);  
  
    ...  
}
```

If **start** is **target**, there is a path ✓

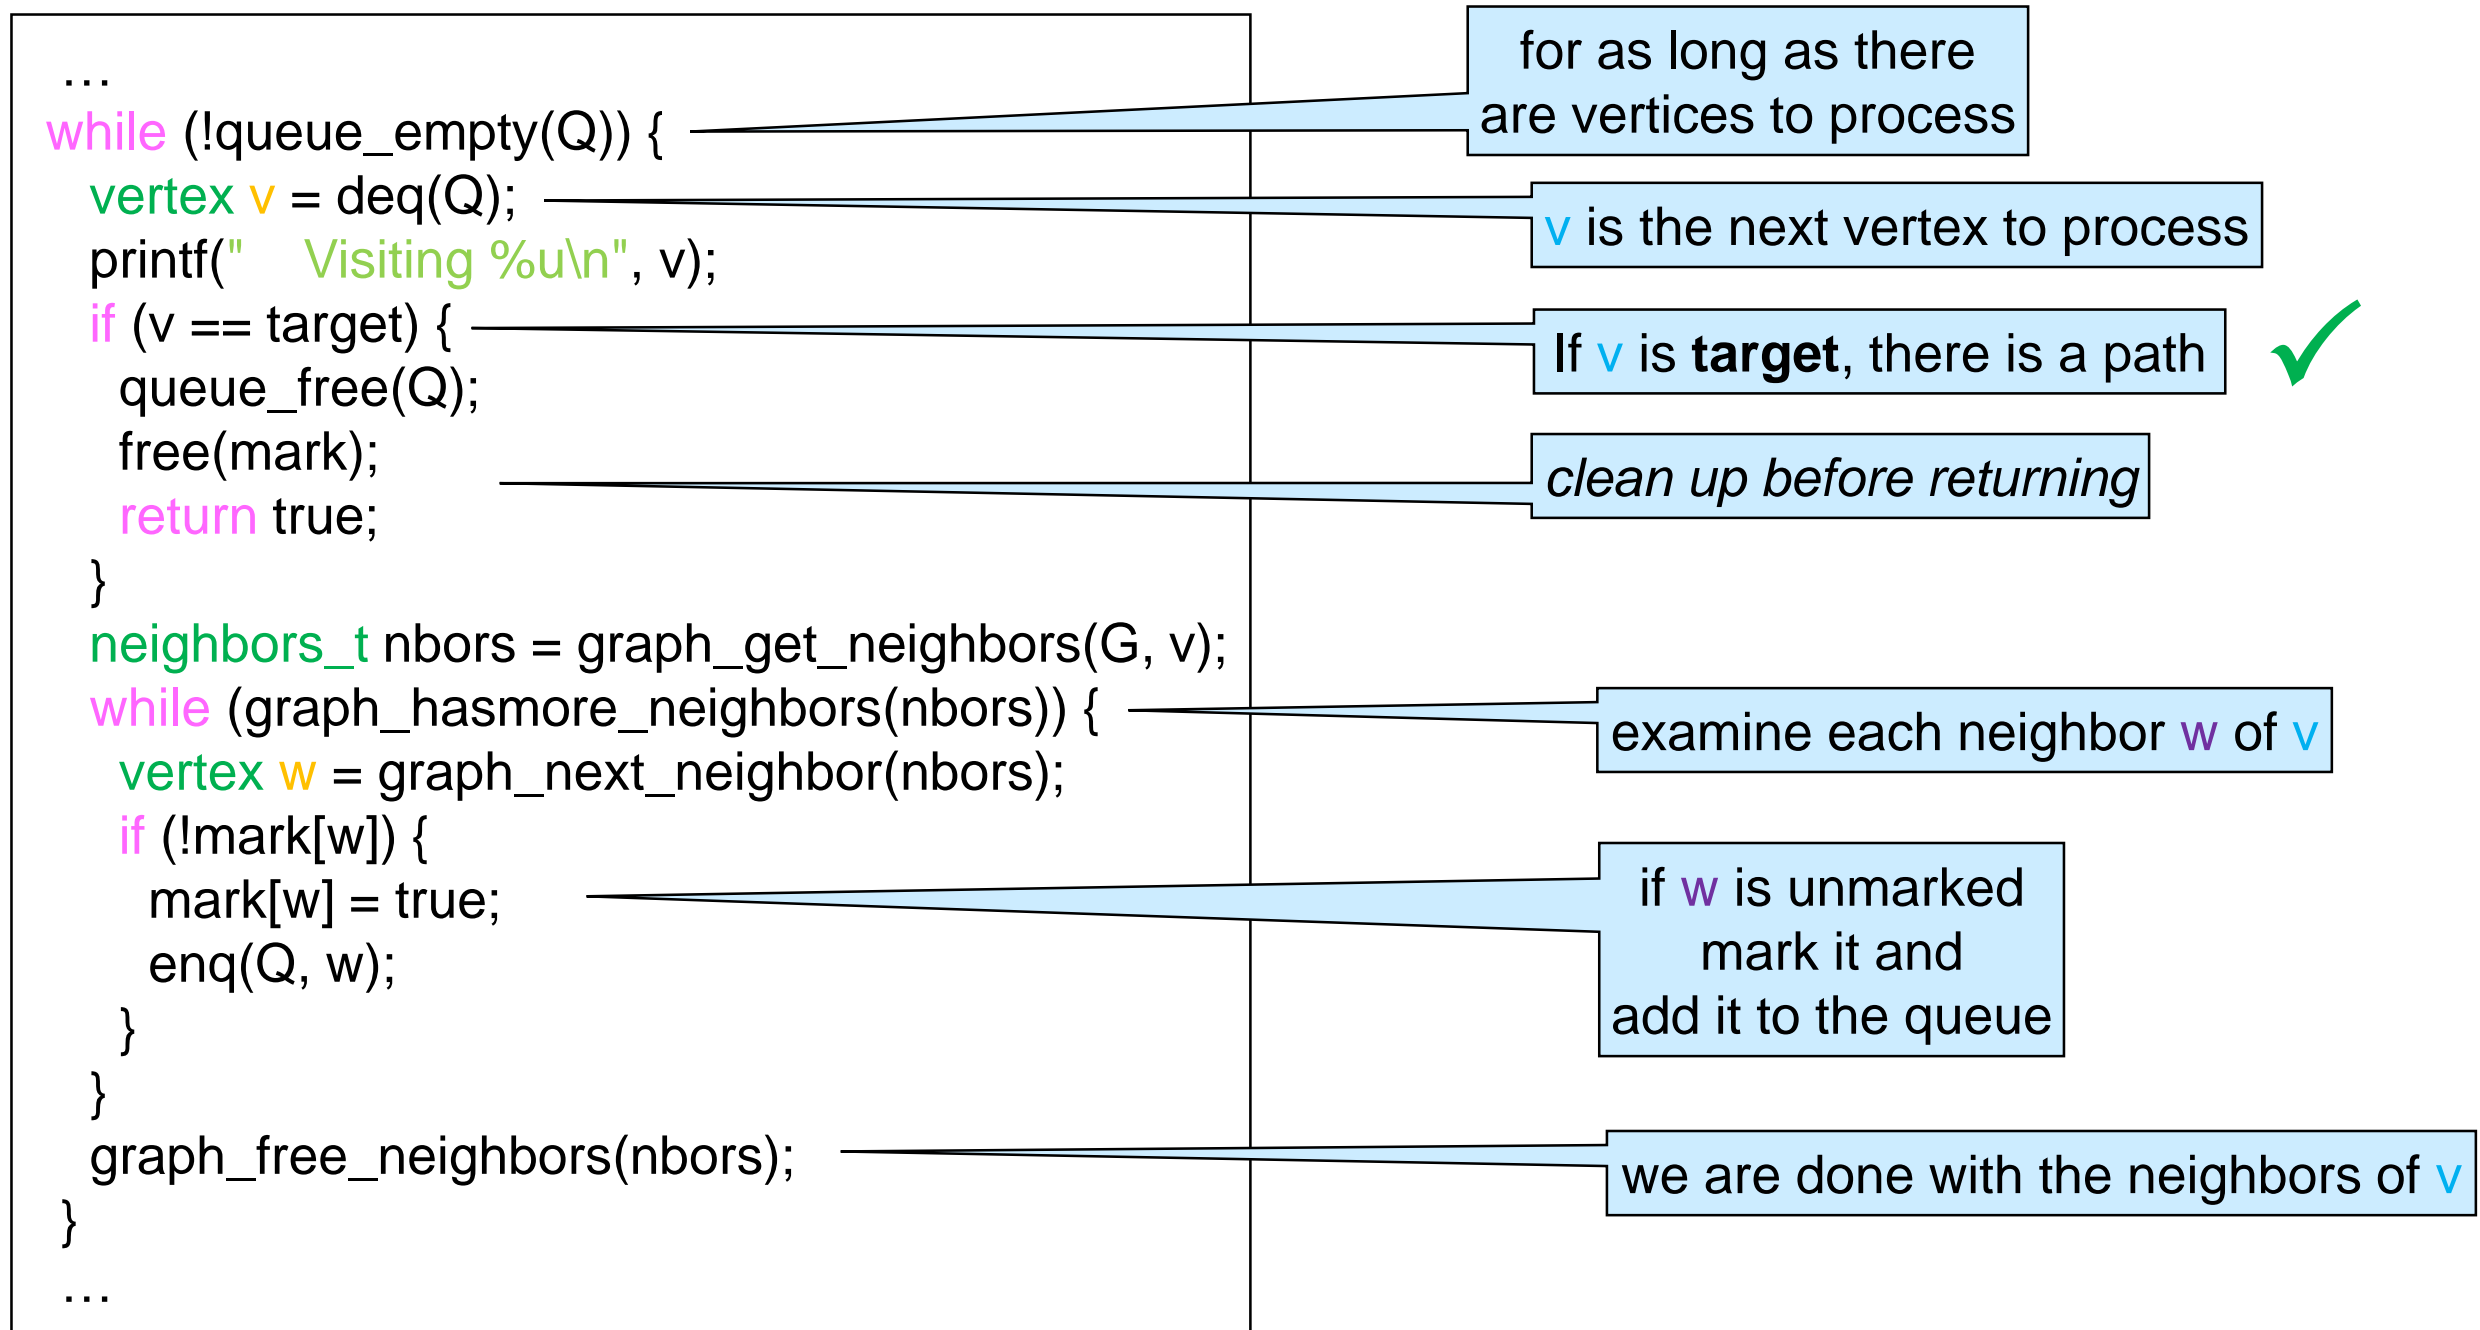
calloc initializes every vertex
as unmarked

but we want **start** to be marked

Initially only **start**
is in the queue

Implementing BFS – II

Traversing the graph



Implementing BFS – III

Giving up

```
...  
while (!queue_empty(Q)) {  
...  
}  
ASSERT(queue_empty(Q));  
queue_free(Q);  
free(mark);  
return false;  
}
```

If there are no more vertices to process

there is no path **x**

clean up before returning

Implementing BFS

- Here's the overall code

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Implementing BFS

- This code is **iterative**
 - DFS earlier was recursive
- The code structure is the same as `graph_print`

```
void graph_print(graph_t G) {
    for (vertex v = 0; v < graph_size(G); v++) {
        printf("Vertices connected to %u: ", v);
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            printf(" %u,", w);
        }
        graph_free_neighbors(nbors);
        printf("\n");
    }
}
```

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // v is the next vertex to process
        vertex v = deq(Q);
        printf(" Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Implementing BFS

- The code structure is the same as `graph_print`
 - except that we return early if we find a path
- The complexity of `bfs` is
 - $O(v + e)$ with adjacency lists
 - $O(v^2)$ with adjacency matrices
- same as `dfs`

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;           O(1)

    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool)); O(v)
    mark[start] = true;                          O(1)

    // Q is a queue containing only start initially
    queue_t Q = queue_new();                     O(1)
    enq(Q, start);                              O(1)

    while (!queue_empty(Q)) {                   v times
        // v is the next vertex to process
        vertex v = deq(Q);
        printf(" Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Complexity annotations:

- A blue bracket on the right side of the `while` loop indicates a total complexity of $O(1)$ per iteration, with a label "v times" above it.
- A red bracket on the right side of the inner `while` loop indicates a complexity of $O(e)$ altogether.
- A blue bracket on the right side of the final cleanup code (`ASSERT`, `queue_free`, `free`) indicates a complexity of $O(1)$.

Correctness

- **bfs** is **correct** if it returns
 - **true** when there is a path from **start** to **target**
 - **false** when there is no path from **start** to **target**
- It returns in three places

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```


Correctness – I

- **bfs** is correct if it returns
 - **true** when there is a path from **start** to **target**
- We need to show that there is a path in this case
 - recall the definition

There is a path from **start** to **target** if

- **start == target**, or
- there is an edge from **start** to some vertex **v** and there is a path from **v** to **target**

- we are in the first case

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Correctness – II

- **bfs** is correct if it returns
 - **true** when there is a path from **start** to **target**
- We need to show that there is a path

There is a path from **start** to **target** if

- **start** == **target**, or
- there is an edge from **start** to some vertex **v** and there is a path from **v** to **target**

- but we have nowhere to point to

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Correctness – II

There is a path from **start** to **target** if

- **start** == **target**, or
- there is an edge from **start** to some vertex **v** and there is a path from **v** to **target**

We need to show there is a path

- *but we have nowhere to point to*

● We need **loop invariants**

- What do we know about marked vertices?
 - there is a path from **start** to every marked vertex
- What do we know about vertices in the queue?
 - every vertex in the queue is marked

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_ismore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Correctness – II

● Candidate loop invariants

- LI 1: there is a path from **start** to every marked vertex
- LI 2: every vertex in the queue is marked

● INIT

- LI 1:
 - initially only **start** is marked *by I.7*
 - there is a path from **start** to **start** *by def*
- LI 2:
 - initially only **start** is in the queue *by I.10*
 - **start** is marked *by I.7*



```
1. bool bfs(graph_t G, vertex start, vertex target) {
2.     REQUIRES(G != NULL);
3.     REQUIRES(start < graph_size(G) && target < ...);
4.
5.     if (start == target) return true;
6.
7.     // mark is an array containing only start
8.     bool *mark = xcalloc(graph_size(G), sizeof(bool));
9.     mark[start] = true;
10.
11.    // Q is a queue containing only start initially
12.    queue_t Q = queue_new();
13.    enq(Q, start);
14.
15.    while (!queue_empty(Q)) {
16.        // v is the next vertex to process
17.        vertex v = deq(Q);
18.        printf("  Visiting %u\n", v);
19.        if (v == target) { // if v is target return true
20.            queue_free(Q);
21.            free(mark);
22.            return true;
23.        }
24.        // for every neighbor w of v
25.        neighbors_t nbors = graph_get_neighbors(G, v);
26.        while (graph_ismore_neighbors(nbors)) {
27.            vertex w = graph_next_neighbor(nbors);
28.            if (!mark[w]) { // if w is not already marked
29.                mark[w] = true; // mark it
30.                enq(Q, w); // enqueue it onto the queue
31.            }
32.        }
33.        graph_free_neighbors(nbors);
34.    }
35.    ASSERT(queue_empty(Q));
36.    queue_free(Q);
37.    free(mark);
38.    return false;
39. }
```

Correctness – II

● Candidate loop invariants

- LI 1: there is a path from **start** to every marked vertex
- LI 2: every vertex in the queue is marked

● PRES

- LI 1:
 - **v** is in the queue *by l.13*
 - it is marked *by LI 2*
 - there is a path from **start** to **v** *by LI 1*
 - **w** is a neighbor of **v** *by l.23*
 - there is a path from **start** to **w** *by def*
 - **w** gets marked *by l.25*
- LI 2:
 - **w** gets added to the queue *by l.26*

```
1. bool bfs(graph_t G, vertex start, vertex target) {
2.     REQUIRES(G != NULL);
3.     REQUIRES(start < graph_size(G) && target < ...);
4.
5.     if (start == target) return true;
6.
7.     // mark is an array containing only start
8.     bool *mark = xcalloc(graph_size(G), sizeof(bool));
9.     mark[start] = true;
10.
11.    // Q is a queue containing only start initially
12.    queue_t Q = queue_new();
13.    enq(Q, start);
14.
15.    while (!queue_empty(Q)) {
16.        // v is the next vertex to process
17.        vertex v = deq(Q);
18.        printf("  Visiting %u\n", v);
19.        if (v == target) { // if v is target return true
20.            queue_free(Q);
21.            free(mark);
22.            return true;
23.        }
24.        // for every neighbor w of v
25.        neighbors_t nbors = graph_get_neighbors(G, v);
26.        while (graph_hasmore_neighbors(nbors)) {
27.            vertex w = graph_next_neighbor(nbors);
28.            if (!mark[w]) { // if w is not already marked
29.                mark[w] = true; // mark it
30.                enq(Q, w); // enqueue it onto the queue
31.            }
32.        }
33.        graph_free_neighbors(nbors);
34.    }
35.    ASSERT(queue_empty(Q));
36.    queue_free(Q);
37.    free(mark);
38.    return false;
39. }
```



Correctness – II

There is a path from **start** to **target** if

- **start** == **target**, or
- there is an edge from **start** to some vertex **v** and there is a path from **v** to **target**

● We can now prove the correctness of this case

- **v** was in the queue *by l.15*
- so, **v** is marked *by LI 2*
- there is a path from **start** to **v** *by LI 1*
- **v** == **target** *by l.17*
- there is a path from **start** to **target** *by def*



```
1. bool bfs(graph_t G, vertex start, vertex target) {
2.     REQUIRES(G != NULL);
3.     REQUIRES(start < graph_size(G) && target < ...);
4.
5.     if (start == target) return true;
6.
7.     // mark is an array containing only start
8.     bool *mark = xmalloc(graph_size(G), sizeof(bool));
9.     mark[start] = true;
10.
11.    // Q is a queue containing only start initially
12.    queue_t Q = queue_new();
13.    enq(Q, start);
14.
15.    while (!queue_empty(Q)) {
16.        // @ LI 1: there is a path from start to every marked vertex
17.        // @ LI 2: every vertex in the queue is marked
18.
19.        // v is the next vertex to process
20.        vertex v = deq(Q);
21.        printf("  Visiting %u\n", v);
22.        if (v == target) { // if v is target return true
23.            queue_free(Q);
24.            free(mark);
25.            return true;
26.        }
27.        // for every neighbor w of v
28.        neighbors_t nbors = graph_get_neighbors(G, v);
29.        while (graph_ismore_neighbors(nbors)) {
30.            vertex w = graph_next_neighbor(nbors);
31.            if (!mark[w]) { // if w is not already marked
32.                mark[w] = true; // mark it
33.                enq(Q, w); // enqueue it onto the queue
34.            }
35.        }
36.        graph_free_neighbors(nbors);
37.    }
38.    ASSERT(queue_empty(Q));
39.    queue_free(Q);
40.    free(mark);
41.    return false;
42. }
```


Correctness – III

- **bfs** is correct if it returns
 - **false** when there is **no** path from **start** to **target**
- LI 1 and LI 2 are insufficient
- We need more insight into the way **bfs** works

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);


    while (!queue_empty(Q)) {
        // @ LI 1: there is a path from start to every marked vertex
        // @ LI 2: every vertex in the queue is marked

        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_hasmore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Correctness – III

- What do the elements of the queue represent?

<i>next</i>	<i>target</i>	<i>queue</i>	<i>marked</i>
	4	0	0
0	4	1, 4	0, 1, 4
1	4	4, 2	0, 1, 4, 2
4	4	<i>Success!</i>	



- The **frontier** of the search

```

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

    if (start == target) return true;


    // mark is an array containing only start
    bool *mark = xmalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // @ LI 1: there is a path from start to every marked vertex
        // @ LI 2: every vertex in the queue is marked

        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_hasmore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
    
```


Correctness – III



- All vertices behind the frontier are marked
 - they have been explored
- All vertices beyond the frontier are unmarked
 - they are still unexplored
- Every path from **start** to **target** goes through the frontier

This is a new loop invariant

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;


    // mark is an array containing only start
    bool *mark = xmalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // @ LI 1: there is a path from start to every marked vertex
        // @ LI 2: every vertex in the queue is marked

        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            free(mark);
            return true;
        }
        // for every neighbor w of v
        nbors = graph_get_neighbors(G, v);
        while (graph_hasmore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Correctness – III

- Every path from **start** to **target** goes through the frontier
- When we finally return,
 1. every path from **start** to **target** goes through the frontier
 - LI 3 hold
 2. the frontier is empty
 - negation of the loop guard
 - therefore there can't be a path from **start** to **target**
 - this is the only way (1) can hold
- bfs is correct 

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xmalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        //@ LI 1: there is a path from start to every marked vertex
        //@ LI 2: every vertex in the queue is marked
        //@ LI 3: every path from start to target goes through Q

        // v is the next vertex to process
        vertex v = deq(Q);
        printf("  Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_hasmore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Other Searches

Work List Choice

- **bfs** uses a **queue** as a work list
 - But the correctness proof does not depend on this
- We get a correct implementation of reachability whatever work list we use

```
bool bfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;

    // mark is an array containing only start
    bool *mark = xalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // Q is a queue containing only start initially
    queue_t Q = queue_new();
    enq(Q, start);

    while (!queue_empty(Q)) {
        // @ LI 1: there is a path from start to every marked vertex
        // @ LI 2: every vertex in the queue is marked
        // @ LI 3: every path from start to target goes through Q

        // v is the next vertex to process
        vertex v = deq(Q);
        printf(" Visiting %u\n", v);
        if (v == target) { // if v is target return true
            queue_free(Q);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_hasmore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                enq(Q, w); // enqueue it onto the queue
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(queue_empty(Q));
    queue_free(Q);
    free(mark);
    return false;
}
```

Work List Choice

- *We get a correct implementation of reachability whatever work list we use*
 - **Stack?**
 - The next vertex we process is the **last** we inserted
 - We get an iterative implementation of **depth-first search**
 - Complexity
 - $O(v + e)$ with adjacency lists
 - $O(v^2)$ with adjacency matrices
- because stack and queue operations have the same complexity

```
bool dfs(graph_t G, vertex start, vertex target) {
    REQUIRES(G != NULL);
    REQUIRES(start < graph_size(G) && target < graph_size(G));

    if (start == target) return true;


    // mark is an array containing only start
    bool *mark = xcalloc(graph_size(G), sizeof(bool));
    mark[start] = true;

    // S is a stack containing only start initially
    stack_t S = stack_new();
    push(S, start);

    while (!stack_empty(S)) {
        // @ LI 1: there is a path from start to every marked vertex
        // @ LI 2: every vertex in the stack is marked
        // @ LI 3: every path from start to target goes through S

        // v is the next vertex to process
        vertex v = pop(S);
        printf(" Visiting %u\n", v);
        if (v == target) { // if v is target return true
            stack_free(S);
            free(mark);
            return true;
        }
        // for every neighbor w of v
        neighbors_t nbors = graph_get_neighbors(G, v);
        while (graph_hasmore_neighbors(nbors)) {
            vertex w = graph_next_neighbor(nbors);
            if (!mark[w]) { // if w is not already marked
                mark[w] = true; // mark it
                push(S, w); // push it onto the stack
            }
        }
        graph_free_neighbors(nbors);
    }
    ASSERT(stack_empty(S));
    stack_free(S);
    free(mark);
    return false;
}
```


Work List Choice

- *We get a correct implementation of reachability whatever work list we use*
- **Priority queues?**
 - The next vertex we process is the **most promising**
 - We get artificial intelligence search algorithms like A*
 - used in planning problems, game search, ...
 - the priority function becomes a heuristic function that tells how good a vertex is
 - Complexity is higher because insertion and removal from a priority queue is not $O(1)$

pronounced "A star"

Reachability

- All these graph reachability algorithms share the same basic idea



Explore the graph by expanding the frontier

- The difference is the kind of work list they use to remember the vertices to examine next
 - DFS: a stack
 - BFS: a queue
 - A*: a priority queue