

CSE 326: Data Structures

Graph Algorithms Part 1: Graph Search

Winter Quarter 2005

1

Topic Outline

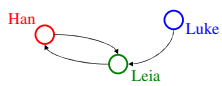
- ❑ Graph Data Structures
- ❑ Graph Properties
- ❑ Topological Sort
- ❑ Graph Search
 - Depth-first, Breadth-first, Iterated Depth-first
 - Dijkstra's Algorithm for Weighted Graphs
 - Heuristic Best-First Search
 - A* Search
- ❑ All-Pairs Shortest Paths
 - Floyd-Warshall Algorithm
- ❑ Connected Component Algorithms
 - Union/Find Algorithm using Up-trees
 - Kruskal's Minimum Spanning Tree Algorithm

2

Graph ADT

Graphs are a formalism for representing **relationships** between objects

- a graph G is represented as $G = (V, E)$
 - V is a set of vertices
 - E is a set of edges
- operations include:
 - iterating over vertices
 - iterating over edges
 - iterating over vertices adjacent to a specific vertex
 - asking whether an edge exists connected two vertices

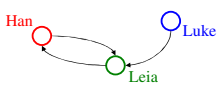


$V = \{\text{Han, Leia, Luke}\}$
 $E = \{(\text{Luke, Leia}), (\text{Han, Leia}), (\text{Leia, Han})\}$

3

Graph Representation 1: Adjacency Matrix

A $|V| \times |V|$ array in which an element (u, v) is true if and only if there is an edge from u to v



	Han	Luke	Leia
Han			
Luke			
Leia			

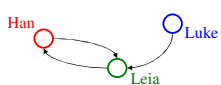
Runtime:
 iterate over vertices
 iterate over edges
 iterate edges adj. to vertex
 edge exists?

Space required:

4

Graph Representation 1: Adjacency Matrix

A $|V| \times |V|$ array in which an element (u, v) is true if and only if there is an edge from u to v



	Han	Luke	Leia
Han			
Luke			
Leia			

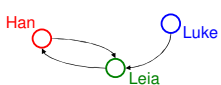
Runtime:
 iterate over vertices $O(|V|)$
 iterate over edges $O(|V|^2)$
 iterate edges adj. to vertex $O(|V|)$
 edge exists? $O(1)$

Space required: $O(|V|^2)$

5

Graph Representation 2: Adjacency List

A $|V|$ -ary list (array) in which each entry stores a list (linked list) of all adjacent vertices



Han	
Luke	
Leia	


Runtime:
 iterate over vertices
 iterate over edges
 iterate edges adj. to vertex
 edge exists?

Space required:

6

Graph Representation 2: Adjacency List

A $|V|$ -ary list (array) in which each entry stores a list (linked list) of all adjacent vertices



Runtime:

iterate over vertices $O(V)$

iterate over edges $O(|E|)$

iterate edges adj. to vertex $O(d)$ (d is number of adj. vertices)

edge exists? $O(d)$

Space required: $O(|V| + |E|)$

7


Terminology

- In *directed* graphs, edges have a specific direction
- In *undirected* graphs, edges are two-way
- Vertices u and v are *adjacent* if $(u, v) \in E$
- A *sparse* graph has $O(|V|)$ edges (upper bound)
- A *dense* graph has $\Omega(|V|^2)$ edges (lower bound)
- A *complete* graph has an edge between every pair of vertices
- An undirected graph is *connected* if there is a path between any two vertices

8

Weighted Graphs

Each edge has an associated weight or cost.




9

Paths and Cycles

A *path* is a list of vertices $\{v_1, v_2, \dots, v_n\}$ such that $(v_i, v_{i+1}) \in E$ for all $0 \leq i < n$.

A *cycle* is a path that begins and ends at the same node.




$p = \{\text{Seattle, Salt Lake City, Chicago, Dallas, San Francisco, Seattle}\}$ 10

Path Length and Cost

Path length: the number of edges in the path

Path cost: the sum of the costs of each edge




$\text{length}(p) = 5$

$\text{cost}(p) = 11.5$ 11

Trees as Graphs

Every tree is a graph with some restrictions:


- the tree is *directed*
- there are *no cycles* (directed or undirected)
- there is a *directed path from the root to every node*



12

Directed Acyclic Graphs (DAGs)

DAGs are directed graphs with no cycles.




Trees \subset DAGs \subset Graphs

13

Topological Sort

Given a directed graph, $G = (V, E)$, output all the vertices in V such that no vertex is output before any other vertex with an edge to it.



14

Topological Sort

Label each vertex's in-degree
Initialize a **queue** to contain all in-degree zero vertices
While there are vertices remaining in the queue
 Remove a vertex v with in-degree of zero and output it
 Reduce the in-degree of all vertices adjacent to v
 Put any of these with new in-degree zero on the queue

Runtime:

15

Topological Sort

Label each vertex's in-degree
Initialize a **queue** to contain all in-degree zero vertices
While there are vertices remaining in the queue
 Remove a vertex v with in-degree of zero and output it
 Reduce the in-degree of all vertices adjacent to v
 Put any of these with new in-degree zero on the queue

Runtime: $O(|V| + |E|)$

16

Topic Outline

- Graph Data Structures
- Graph Properties
- Topological Sort
- Graph Search
 - Depth-first, Breadth-first, Iterated Depth-first
 - Dijkstra's Algorithm for Weighted Graphs
 - Heuristic Best-First Search
 - A* Search
- All-Pairs Shortest Paths
 - Floyd-Warshall Algorithm
- Connected Component Algorithms
 - Union/Find Algorithm using Up-trees
 - Kruskal's Minimum Spanning Tree Algorithm

17

Graph Search

Many problems in computer science correspond to searching for a **path** in a graph, given a **start node** and **goal criteria**

- Route planning – Mapquest
- Packet-switching
- VLSI layout
- 6-degrees of Kevin Bacon
- Program synthesis
- Speech recognition
 - We'll discuss these last two later...

18

General Graph Search Algorithm

Open – some data structure (e.g., stack, queue, heap)
Criteria – some method for removing an element from Open

```
Search( Start, Goal_test, Criteria)
  insert(Start, Open);
  repeat
    if (empty(Open)) then return fail;
    select Node from Open using Criteria;
    if (Goal_test(Node)) then return Node;
    for each Child of node do
      if (Child not already visited) then Insert( Child, Open );
    Mark Node as visited;
  end
```

19

Depth-First Graph Search

Open – Stack
Criteria – Pop

```
DFS( Start, Goal_test)
  push(Start, Open);
  repeat
    if (empty(Open)) then return fail;
    Node := pop(Open);
    if (Goal_test(Node)) then return Node;
    for each Child of node do
      if (Child not already visited) then push(Child, Open);
    Mark Node as visited;
  end
```

20

Breadth-First Graph Search

Open – Queue
Criteria – Dequeue (FIFO)

```
BFS( Start, Goal_test)
  enqueue(Start, Open);
  repeat
    if (empty(Open)) then return fail;
    Node := dequeue(Open);
    if (Goal_test(Node)) then return Node;
    for each Child of node do
      if (Child not already visited) then enqueue(Child, Open);
    Mark Node as visited;
  end
```

21

Comparison: DFS versus BFS

Depth-first search

- Does not always find shortest paths
- Must be careful to mark visited vertices, or you could go into an infinite loop if there is a cycle

Breadth-first search

- Always finds shortest paths – **optimal solutions**
- Marking visited nodes can improve efficiency, but even without doing so search is guaranteed to terminate

Is BFS always preferable?

22

DFS Space Requirements

Assume:

- Longest path in graph is length d
- Highest number of out-edges is k

DFS stack grows at most to size dk

- For $k=10$, $d=15$, size is 150

23

BFS Space Requirements

Assume

- Distance from start to a goal is d
- Highest number of out edges is k BFS

Queue could grow to size k^d

- For $k=10$, $d=15$, size is **1,000,000,000,000,000**

24

Conclusion

For large graphs, DFS is hugely more memory efficient, *if we can limit the maximum path length to some fixed d .*

- If we *knew* the distance from the start to the goal in advance, we can just *not add any children to stack after level d*
- But what if we don't know d in advance?

25

Iterative-Deepening DFS (I)

```
Bounded_DFS(Start, Goal_test, Limit)
Start.dist = 0;
push(Start, Open);
repeat
  if (empty(Open)) then return fail;
  Node := pop(Open);
  if (Goal_test(Node)) then return Node;
  if (Node.dist > Limit) then return fail;
  for each Child of node do
    if (Child not already i-visited) then
      Child.dist := Node.dist + 1;
      push(Child, Open);
  Mark Node as i-visited;
end
```

26

Iterative-Deepening DFS (II)

```
IDFS_Search(Start, Goal_test)
i := 1;
repeat
  answer := Bounded_DFS(Start, Goal_test, i);
  if (answer != fail) then return answer;
  i := i+1;
end
```

27

Analysis of IDFS

Work performed with limit < actual distance to G is wasted – but the wasted work is usually small compared to amount of work done during the *last* iteration

$$\sum_{i=1}^d k^i = O(k^d) \quad \text{Ignore low order terms!}$$

Same time complexity as BFS

Same space complexity as (bounded) DFS

28

Saving the Path


Our pseudocode returns the goal node found, but not the path to it

How can we remember the path?

- Add a field to each node, that points to the previous node along the path
- Follow pointers from goal back to start to recover path


29

Example




30

Example (Unweighted Graph)



31

Example (Unweighted Graph)



32

Graph Search, Saving Path

```

Search( Start, Goal_test, Criteria)
  insert(Start, Open);
  repeat
    if (empty(Open)) then return fail;
    select Node from Open using Criteria;
    if (Goal_test(Node)) then return Node;
    for each Child of node do
      if (Child not already visited) then
        Child.previous := Node;
        Insert( Child, Open );
    Mark Node as visited;
  end
  
```

33

Shortest Path for Weighted Graphs

Given a graph $G = (V, E)$ with edge costs $c(e)$, and a vertex $s \in V$, find the shortest (lowest cost) path from s to every vertex in V

Assume: only *positive* edge costs

34

Edsger Wybe Dijkstra (1930-2002)



- Invented concepts of structured programming, synchronization, weakest precondition, and "semaphores" for controlling computer processes. The Oxford English Dictionary cites his use of the words "vector" and "stack" in a computing context.
- Believed programming should be taught without computers
- 1972 Turing Award
- "In their capacity as a tool, computers will be but a ripple on the surface of our culture. In their capacity as intellectual challenge, they are without precedent in the cultural history of mankind."

35

Dijkstra's Algorithm for Single Source Shortest Path

Similar to breadth-first search, but uses a **heap** instead of a queue:

- Always select (expand) the vertex that has a lowest-cost path to the start vertex

Correctly handles the case where the lowest-cost (shortest) path to a vertex is **not** the one with fewest edges

36

Pseudocode for Dijkstra

```

Initialize the cost of each node to  $\infty$ 
s.cost := 0
insert(s, 0, heap);
While (! empty(heap))
  n := deleteMin(heap);
  For each edge e=(n,a) do
    if (n.cost + e.cost < a.cost) then
      a.cost = n.cost + e.cost;
      a.previous = n;
    if (a is in the heap) then
      decreaseKey(a, a.cost, heap)
    else insert(a, a.cost, heap)
  end
end
end
  
```

37

Important Features

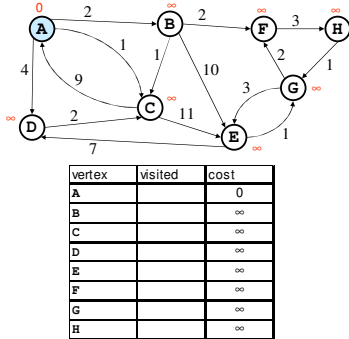
Once a vertex is **removed** from the head, the cost of the shortest path to that node is known

While a vertex is still in the heap, **another shorter path** to it might still be found

The shortest path itself can be found by following the backward pointers stored in **node.previous**

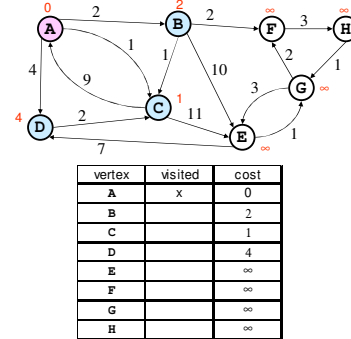
38

Dijkstra's Algorithm in Action




39

Dijkstra's Algorithm in Action



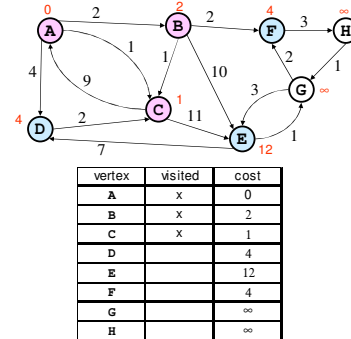
40

Dijkstra's Algorithm in Action




41

Dijkstra's Algorithm in Action



42


Dijkstra's Algorithm in Action



vertex	visited	cost
A	x	0
B	x	2
C	x	1
D	x	4
E		12
F	x	4
G		∞
H		7

43


Dijkstra's Algorithm in Action



vertex	visited	cost
A	x	0
B	x	2
C	x	1
D	x	4
E		12
F	x	4
G		8
H	x	7

44


Dijkstra's Algorithm in Action



vertex	visited	cost
A	x	0
B	x	2
C	x	1
D	x	4
E		11
F	x	4
G	x	8
H	x	7

45

Dijkstra's Algorithm in Action



vertex	visited	cost
A	x	0
B	x	2
C	x	1
D	x	4
E	x	11
F	x	4
G	x	8
H	x	7

46

Data Structures for Dijkstra's Algorithm

$|V|$ times:
 Select the unknown node with the lowest cost
 findMin/deleteMin $O(\log |V|)$

$|E|$ times:
 a 's cost = $\min(a$'s old cost, ...)
 decreaseKey or insert $O(\log |V|)$

runtime: $O(|E| \log |V|)$

47

Topic Outline

- Graph Data Structures
- Graph Properties
- Topological Sort
- Graph Search
 - Depth-first, Breadth-first, Iterated Depth-first
 - Dijkstra's Algorithm for Weighted Graphs
 - Heuristic Best-First Search
 - A* Search
- All-Pairs Shortest Paths
 - Floyd-Warshall Algorithm
- Connected Component Algorithms
 - Union/Find Algorithm using Up-trees
 - Kruskal's Minimum Spanning Tree Algorithm


48

Problem: Large Graphs

- ❑ It is expensive to find optimal paths in large graphs, using BFS, IDFS, or Dijkstra's algorithm (for weighted graphs)
- ❑ How can we search large graphs efficiently by using "commonsense" about which direction looks most promising?

49


Example



Plan a route from 9th & 50th to 3rd & 51st

50

Example



Plan a route from 9th & 50th to 3rd & 51st

51

Best-First Search

The *Manhattan distance* ($\Delta x + \Delta y$) is an **estimate** of the distance to the goal

- It is a *search heuristic*

❑ Best-First Search

- Order nodes in priority to **minimize estimated distance to the goal**

❑ Compare: BFS / Dijkstra

- Order nodes in priority to minimize distance from the start

52

Best-First Search

Open – Heap (priority queue)

Criteria – Smallest key (highest priority)

$h(n)$ – heuristic estimate of distance from n to closest goal


```

Best_First_Search( Start, Goal_test)
  insert(Start, h(Start), heap);
  repeat
    if (empty(heap)) then return fail;
    Node := deleteMin(heap);
    if (Goal_test(Node)) then return Node;
    for each Child of node do
      if (Child not already visited) then
        insert(Child, h(Child), heap);
    end
  Mark Node as visited;
end
  
```

53


Obstacles

Best-FS eventually will expand vertex to get back on the right track



54

Non-Optimality of Best-First



55

Improving Best-First

- ❑ Best-first is often tremendously faster than BFS/Dijkstra, but might stop with a non-optimal solution
- ❑ How can it be modified to be (almost) as fast, but guaranteed to find optimal solutions?
- ❑ A* - Hart, Nilsson, Raphael 1968
 - One of the first significant algorithms developed in AI
 - Widely used in many applications

56

A*

Exactly like Best-first search, but using a different criteria for the priority queue:

minimize (distance from start) + (estimated distance to goal)

priority $f(n) = g(n) + h(n)$

$f(n)$ = priority of a node

$g(n)$ = true distance from start

$h(n)$ = heuristic distance to goal

57

Optimality of A*


Suppose the estimated distance is *always* less than or equal to the true distance to the goal

- heuristic is a lower bound

Then: when the goal is removed from the priority queue, we are **guaranteed** to have found a shortest path!

58

A* in Action




59

Applications of A*: Planning

A huge graph may be **implicitly specified** by rules for generating it on-the-fly

Blocks world:


- vertex = relative positions of all blocks
- edge = robot arm stacks one block



Blocks World

Blocks world:

- distance = number of stacks to perform
- heuristic lower bound = number of blocks out of place



out of place = 2, true distance to goal = 3

61

Application of A*: Speech Recognition

(Simplified) Problem:

- System hears a sequence of 3 words
- It is unsure about what it heard
 - For each word, it has a set of possible “guesses”
 - E.g.: Word 1 is one of { “hi”, “high”, “I” }
- What is the most likely sentence it heard?

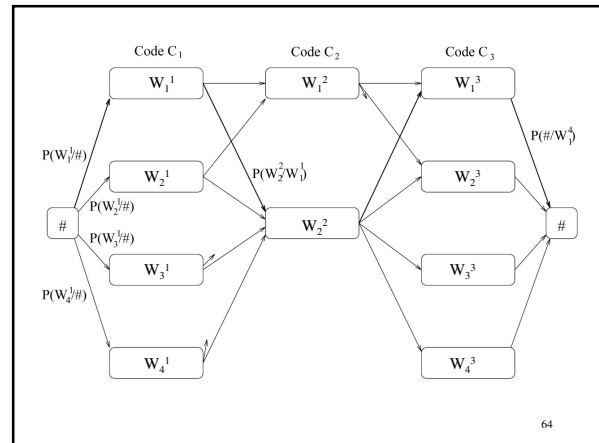
62

Speech Recognition as Shortest Path

Convert to a shortest-path problem:

- Utterance is a “layered” DAG
- Begins with a special dummy “start” node
- Next: A layer of nodes for each word position, one node for each word choice
- Edges between every node in layer i to every node in layer $i+1$
 - Cost of an edge is smaller if the pair of words frequently occur together in real speech
 - + Technically: $-\log$ probability of co-occurrence
- Finally: a dummy “end” node
- Find shortest path from start to end node

63



64

Summary: Graph Search

Depth First

- Little memory required
- Might find non-optimal path

Breadth First

- Much memory required
- Always finds optimal path

Dijkstra's Short Path Algorithm

- Like BFS for weighted graphs

Best First

- Can visit fewer nodes
- Might find non-optimal path

A*

- Can visit fewer nodes than BFS or Dijkstra
- Optimal if heuristic estimate is a lower-bound

65