

# Searching: Deterministic single-agent

**Andrew W. Moore**  
Professor  
School of Computer Science  
Carnegie Mellon University  
[www.cs.cmu.edu/~awm](http://www.cs.cmu.edu/~awm)  
[awm@cs.cmu.edu](mailto:awm@cs.cmu.edu)  
412-268-7599

Note to other teachers and users of these slides. Andrew would be delighted if you found this source material useful in giving your own lectures. Feel free to use these slides verbatim, or to modify them to fit your own needs. PowerPoint originals are available. If you make use of a significant portion of these slides in your own lecture, please include this message, or the following link to the source repository of Andrew's tutorials: <http://www.cs.cmu.edu/~awm/tutorials>. Comments and corrections gratefully received.


Slide 1

## Overview

- Deterministic, single-agent, search problems
- Breadth First Search
- Optimality, Completeness, Time and Space complexity
- Search Trees
- Depth First Search
- Iterative Deepening
- Best First “Greedy” Search

Slide 2

## A search problem



How do we get from S to G? And what's the smallest possible number of transitions?

Slide 3

## Formalizing a search problem


A search problem has five components:

$Q$ ,  $S$ ,  $G$ , **succs**, **cost**

- $Q$  is a finite set of states.
- $S \subseteq Q$  is a non-empty set of start states.
- $G \subseteq Q$  is a non-empty set of goal states.
- **succs** :  $Q \rightarrow P(Q)$  is a function which takes a state as input and returns a set of states as output. **succs**( $s$ ) means "the set of states you can reach from  $s$  in one step".
- **cost** :  $Q, Q \rightarrow \text{Positive Number}$  is a function which takes two states,  $s$  and  $s'$ , as input. It returns the one-step cost of traveling from  $s$  to  $s'$ . The cost function is only defined when  $s'$  is a successor state of  $s$ .

Slide 4

## Our Search Problem



$Q = \{ \text{START}, a, b, c, d, e, f, h, p, q, r, \text{GOAL} \}$

$S = \{ \text{START} \}$

$G = \{ \text{GOAL} \}$

$\text{succs}(b) = \{ a \}$


$\text{succs}(e) = \{ h, r \}$

$\text{succs}(a) = \text{NULL} \dots \text{etc.}$

$\text{cost}(s,s') = 1$  for all transitions

Slide 5

## Our Search Problem



$Q = \{ \text{START}, a, b, c, d, e, f, h, p, q, r, \text{GOAL} \}$

$S = \{ \text{START} \}$

$G = \{ \text{GOAL} \}$

$\text{succs}(b) = \{ a \}$

$\text{succs}(e) = \{ h, r \}$

$\text{succs}(a) = \text{NULL} \dots \text{etc.}$

$\text{cost}(s,s') = 1$  for all transitions

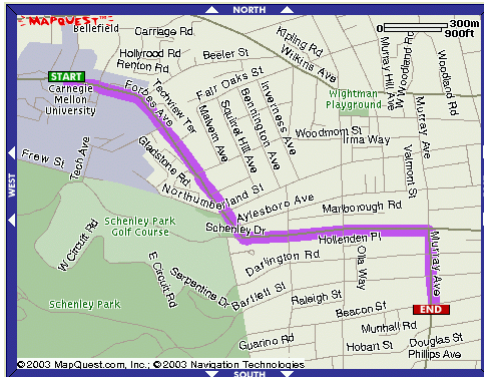
Why do we care? What problems are like this?

Slide 6

# Search Problems

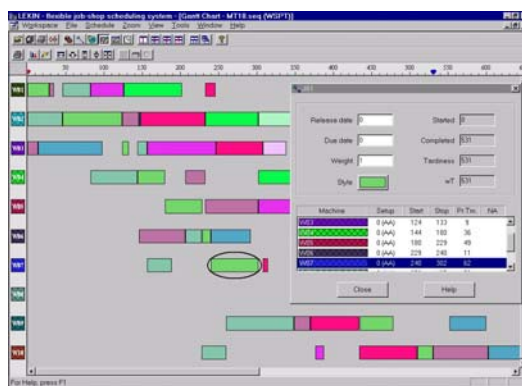


1	2	3
6	7	
8	5	4




Slide 7

# More Search Problems




Scheduling


8-Queens



What next?



Slide 8





# More Search Problems

But there are plenty of things which we'd normally call search problems that don't fit our rigid definition...



- A search problem has five components:
- $Q$ ,  $S$ ,  $G$ , **succs**, **cost**
- $Q$  is a finite set of states.
- $S \subseteq Q$  is a non-empty set of start states.
- $G \subseteq Q$  is a non-empty set of goal states.
- **succs** :  $Q \rightarrow P(Q)$  is a function which takes a state as input and returns a set of states as output. **succs**( $s$ ) means "the set of states you can reach from  $s$  in one step".
- **cost** :  $Q, Q \rightarrow \text{Positive Number}$  is a function which takes two states,  $s$  and  $s'$ , as input. It returns the one-step cost of traveling from  $s$  to  $s'$ . The cost function is only defined when  $s'$  is a successor state of  $s$ .

Can you think of examples?



Slide 9

## Our definition excludes...



Slide 10

## Our definition excludes

Game against adversary



Chance



Hidden State



Continuum (infinite number) of states




All of the above, plus distributed team control



Slide 11

## Breadth First Search



Label all states that are reachable from S in 1 step but aren't reachable in less than 1 step.


Then label all states that are reachable from S in 2 steps but aren't reachable in less than 2 steps.

Then label all states that are reachable from S in 3 steps but aren't reachable in less than 3 steps.

Etc... until Goal state reached.

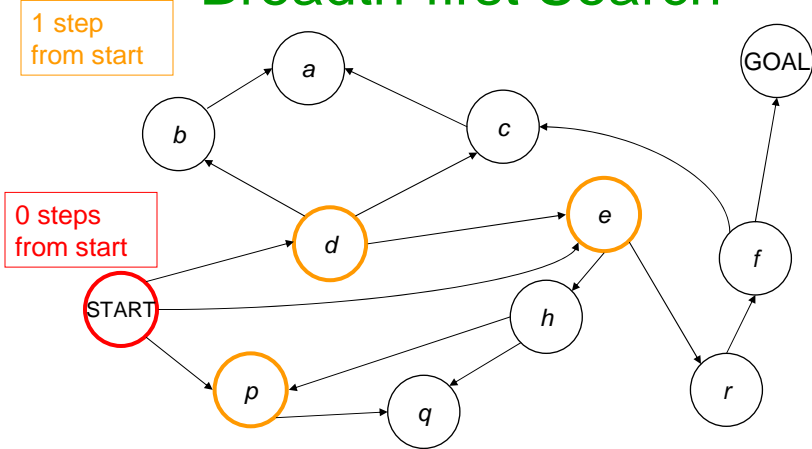
Slide 12

# Breadth-first Search

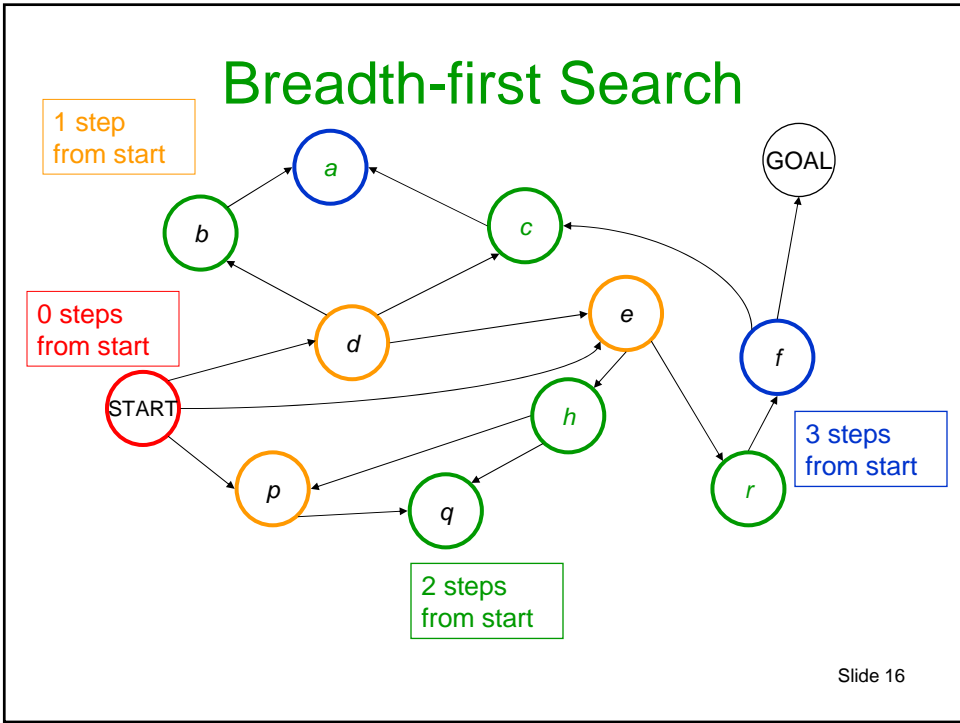
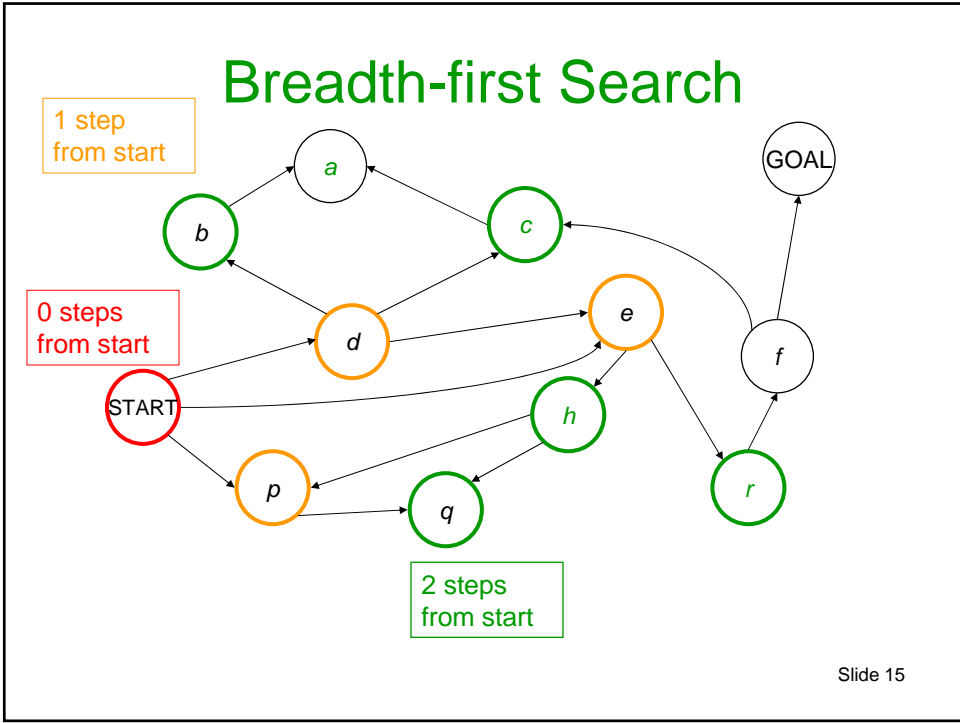


Slide 13

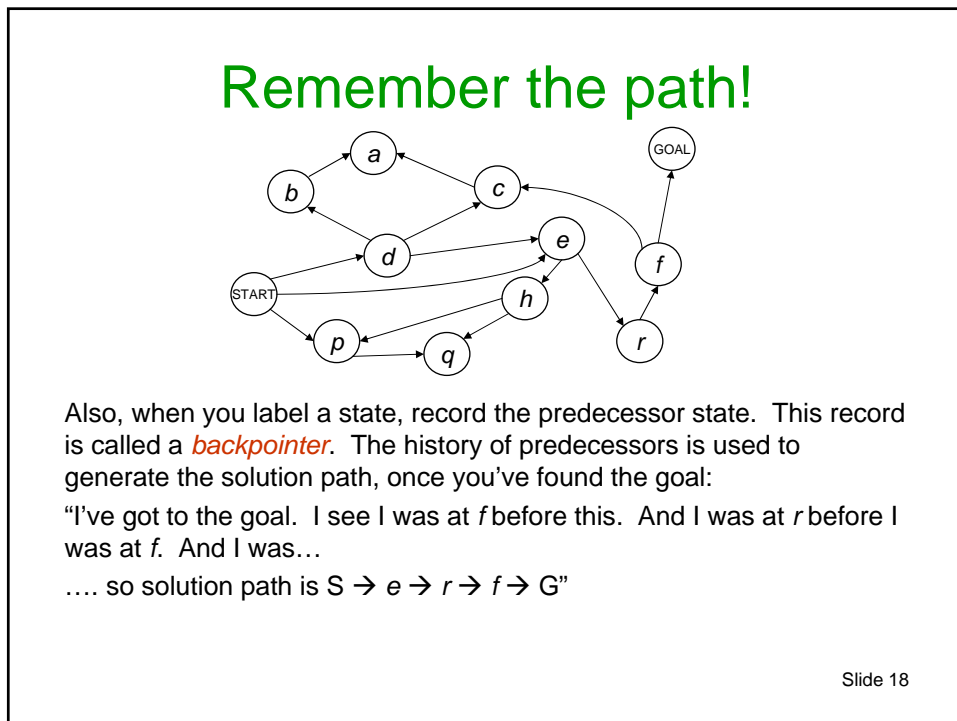
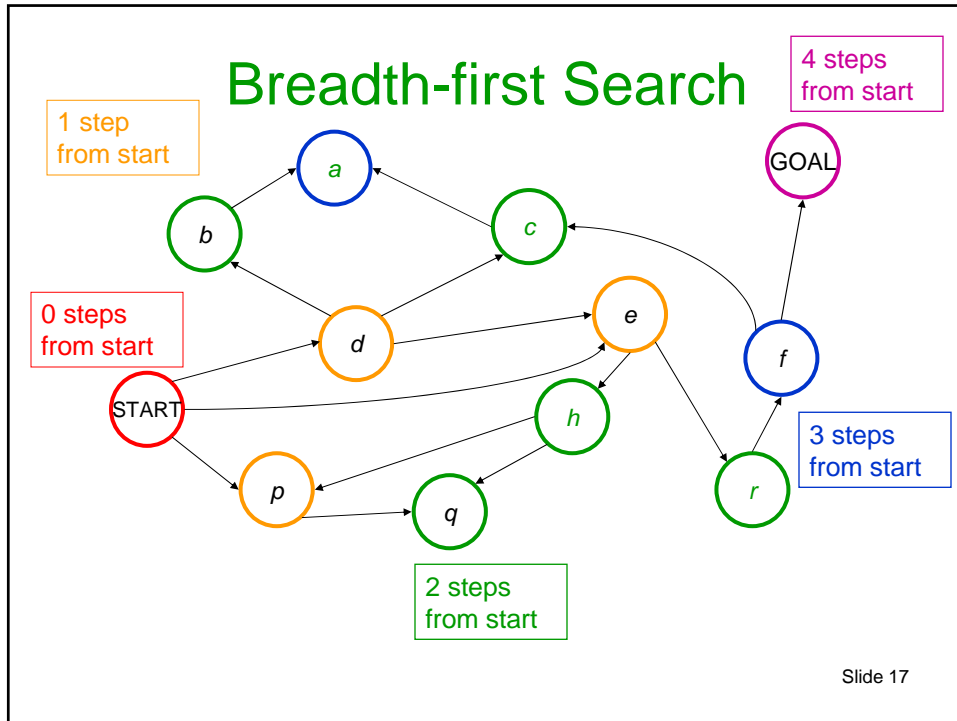
# Breadth-first Search

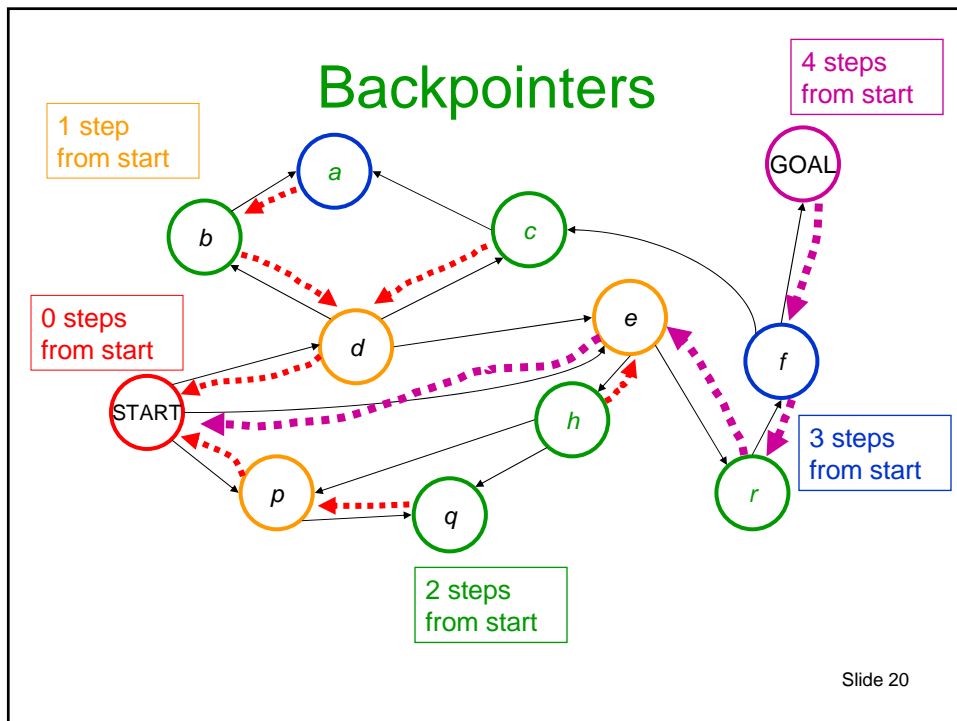
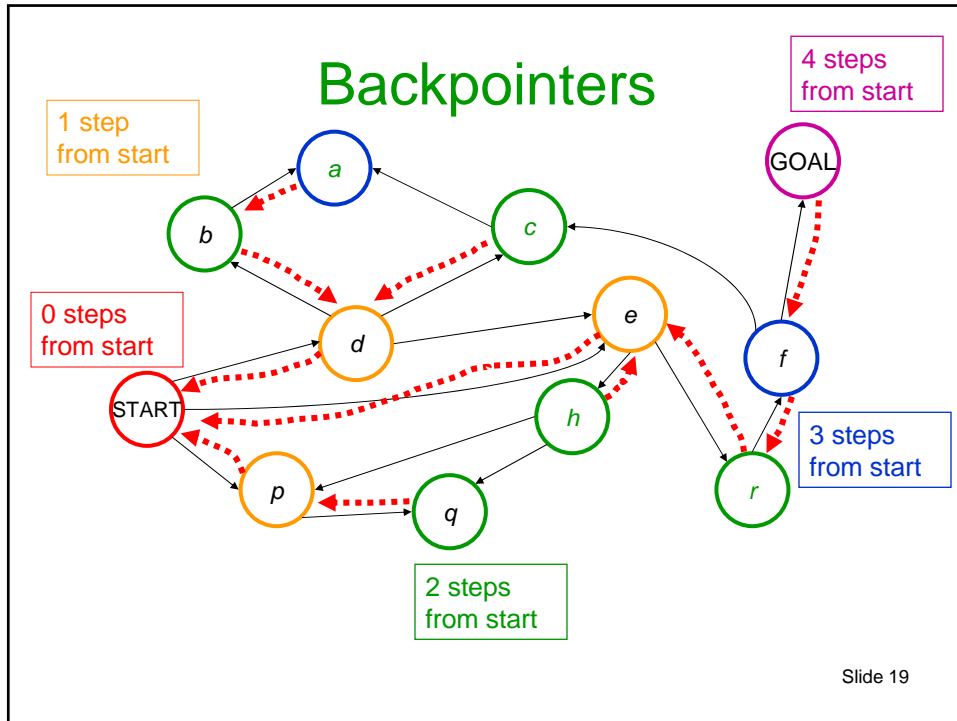


Slide 14









## Starting Breadth First Search

For any state  $s$  that we've labeled, we'll remember:

- $previous(s)$  as the previous state on a shortest path from START state to  $s$ .

On the  $k$ th iteration of the algorithm we'll begin with  $V_k$  defined as the set of those states for which the shortest path from the start costs exactly  $k$  steps


Then, during that iteration, we'll compute  $V_{k+1}$ , defined as the set of those states for which the shortest path from the start costs exactly  $k+1$  steps

We begin with  $k = 0$ ,  $V_0 = \{START\}$  and we'll define,  $previous(START) = NULL$

Then we'll add in things one step from the START into  $V_1$ . And we'll keep going.


Slide 21

## BFS




Slide 22

# BFS

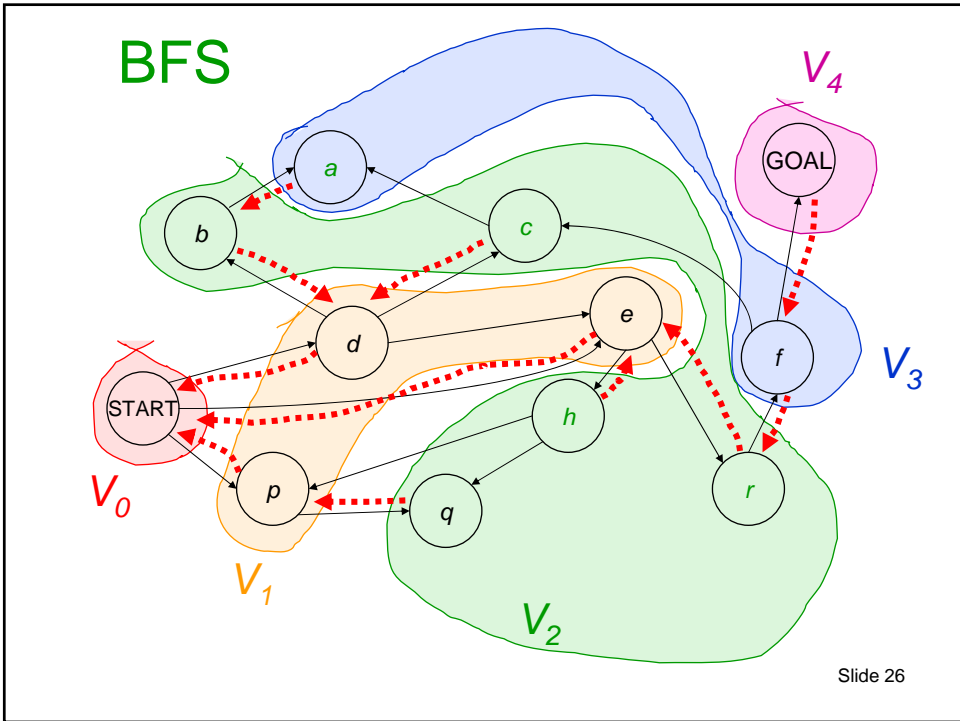
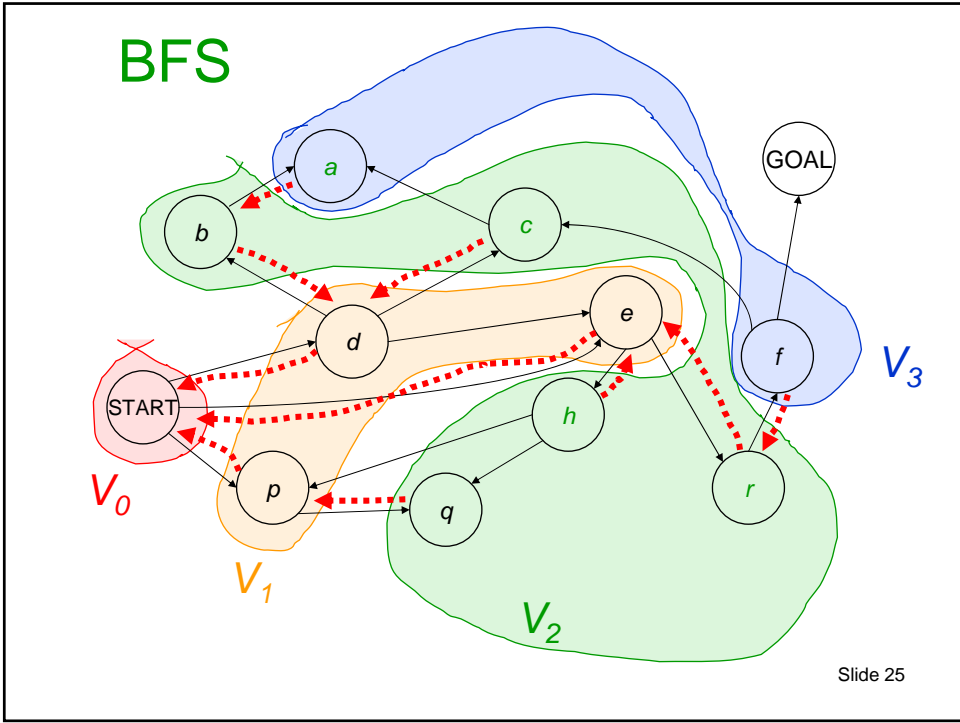


Slide 23

# BFS



Slide 24



# Breadth First Search

$V_0 := S$  (the set of start states)

$previous(START) := NIL$

$k := 0$

**while** (no goal state is in  $V_k$  and  $V_k$  is not empty) **do**

$V_{k+1} :=$  empty set

    For each state  $s$  in  $V_k$

        For each state  $s'$  in **succs**( $s$ )

            If  $s'$  has not already been labeled

                Set  $previous(s') := s$

                Add  $s'$  into  $V_{k+1}$

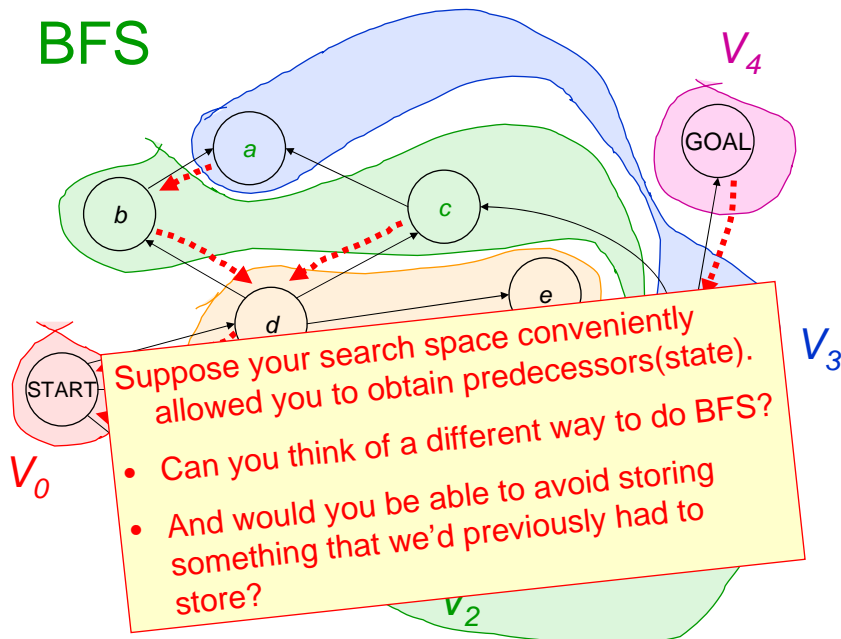
$k := k+1$

**If**  $V_k$  is empty signal FAILURE

**Else** build the solution path thus: Let  $S_i$  be the  $i$ th state in the shortest path. Define  $S_k = GOAL$ , and for all  $i < k$ , define  $S_{i-1} = previous(S_i)$ .


Slide 27

## BFS



Slide 28

## Another way: Work back



Label all states that can reach G in 1 step but can't reach it in less than 1 step.

Label all states that can reach G in 2 steps but can't reach it in less than 2 steps.

Etc. ... until start is reached.

"number of steps to goal" labels determine the shortest path. Don't need extra bookkeeping info.


Slide 29

## Breadth First Details

- It is fine for there to be more than one goal state.
- It is fine for there to be more than one start state.
- This algorithm works forwards from the start. Any algorithm which works forwards from the start is said to be *forward chaining*.
- You can also work backwards from the goal. This algorithm is very similar to Dijkstra's algorithm.
- Any algorithm which works backwards from the goal is said to be *backward chaining*.
- Backward versus forward. Which is better?

Slide 30

## Costs on transitions



Notice that BFS finds the shortest path in terms of number of transitions. It does not find the least-cost path.

We will quickly review an algorithm which does find the least-cost path. On the  $k$ th iteration, for any state  $S$ , write  $g(s)$  as the least-cost path to  $S$  in  $k$  or fewer steps.

Slide 31

## Least Cost Breadth First

$V_k$  = the set of states which can be reached in exactly  $k$  steps, and for which the least-cost  $k$ -step path is less cost than any path of length less than  $k$ . In other words,  $V_k$  = the set of states whose values changed on the previous iteration.

$V_0 := S$  (the set of start states)

$previous(START) := NIL$

$g(START) = 0$

$k := 0$

**while** ( $V_k$  is not empty) **do**

$V_{k+1} :=$  empty set

For each state  $s$  in  $V_k$

For each state  $s'$  in **succs**( $s$ )

If  $s'$  has not already been labeled

OR if  $g(s) + Cost(s, s') < g(s')$

Set  $previous(s') := s$

Set  $g(s') := g(s) + Cost(s, s')$

Add  $s'$  into  $V_{k+1}$

$k := k+1$

**If** GOAL not labeled, exit signaling FAILURE

**Else** build the solution path thus: Let  $S_k$  be the  $k$ th state in the shortest path. Define  $S_k = GOAL$ , and for all  $i < k$ , define  $S_{i-1} = previous(S_i)$ .

Slide 32



## Uniform-Cost Search

- A conceptually simple BFS approach when there are costs on transitions
- It uses priority queues

Slide 33



## Priority Queue Refresher

A priority queue is a data structure in which you can insert and retrieve (thing, value) pairs with the following operations:

Init-PriQueue(PQ)	initializes the PQ to be empty.
Insert-PriQueue(PQ, thing, value)	inserts (thing, value) into the queue.
Pop-least(PQ)	returns the (thing, value) pair with the lowest value, and removes it from the queue.

Slide 34



## Priority Queue Refresher

A priority queue is a data structure in which you can insert and retrieve *(thing, value)* pairs with the following operations:

For more details, see Knuth or Sedgwick or basically any book with the word "algorithms" prominently appearing in the title.

Init-PriQueue(PQ)	initializes the PQ to be empty.
Insert-PriQueue(PQ, thing, value)	inserts <i>(thing, value)</i> into the queue.
Pop-least(PQ)	returns the <i>(thing, value)</i> pair with the lowest value, and removes it from the queue.

Priority Queues can be implemented in such a way that the cost of the insert and pop operations are

Very cheap (though not absolutely, incredibly cheap!)

$O(\log(\text{number of things in priority queue}))$

Slide 35

## Uniform-Cost Search


- A conceptually simple BFS approach when there are costs on transitions
- It uses a priority queue

PQ = Set of states that have been expanded or are awaiting expansion

Priority of state  $s = g(s) =$  cost of getting to  $s$  using path implied by backpointers.

Slide 36


## Starting UCS



$PQ = \{(S,0)\}$

Slide 37

## UCS Iterations




$PQ = \{(S,0)\}$

Iteration:  
 1. Pop least-cost state from PQ  
 2. Add successors

Slide 38

## UCS Iterations




Iteration:  
 1. Pop least-cost state from PQ  
 2. Add successors

$PQ = \{ (p,1), (d,3), (e,9) \}$

Slide 39

## UCS Iterations




$PQ = \{ (d,3), (e,9), (q,16) \}$

Iteration:  
 1. Pop least-cost state from PQ  
 2. Add successors

Slide 40


## UCS Iterations



$PQ = \{ (b,4) , (e,5) , (c,11) , (q,16) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors
- Slide 41


## UCS Iterations



$PQ = \{ (b,4) , (e,5) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors
- Slide 42


## UCS Iterations



$PQ = \{ (e,5), (a,6), (c,11), (q,16) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors
- Slide 43


## UCS Iterations



$PQ = \{ (a,6), (h,6), (c,11), (r,14), (q,16) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors
- Slide 44

## UCS Iterations




$PQ = \{ (h,6), (c,11), (r,14), (q,16) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors

Slide 45

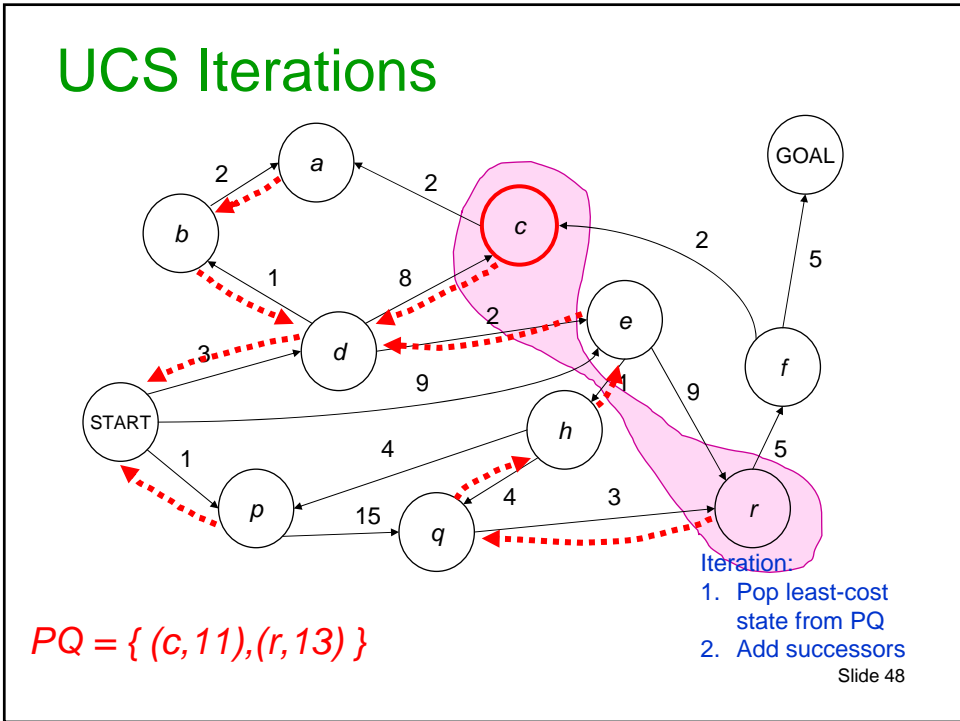
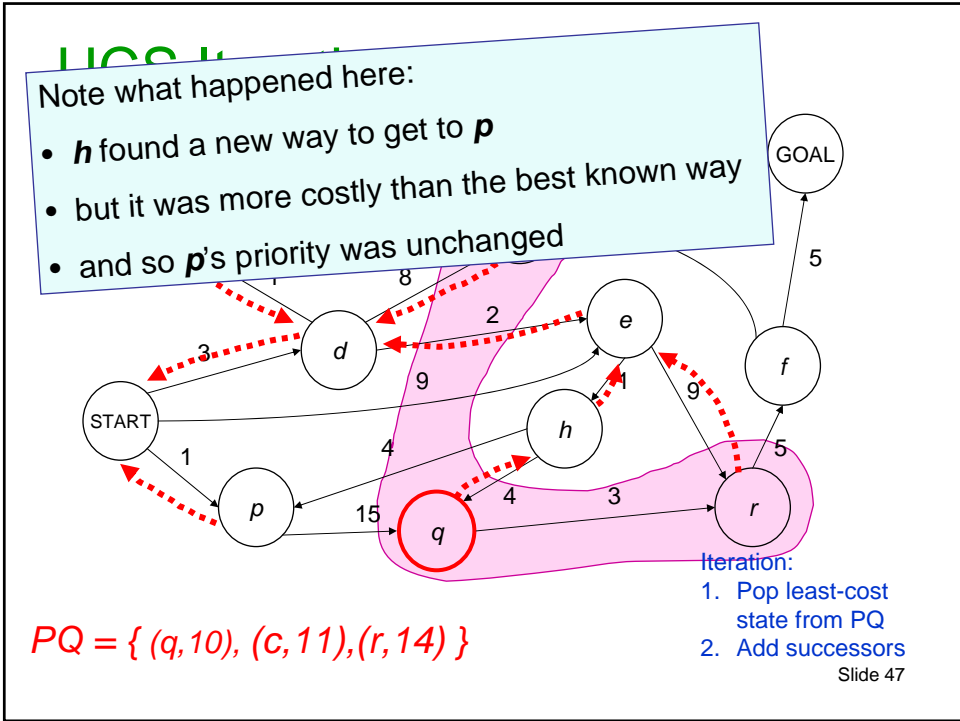
## UCS Iterations



$PQ = \{ (q,10), (c,11), (r,14) \}$


- Iteration:
1. Pop least-cost state from PQ
  2. Add successors

Slide 46





## UCS Iterations




$PQ = \{ (r, 13) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors

Slide 49

## UCS Iterations




$PQ = \{ (f, 18) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors

Slide 50

# UCS Iterations

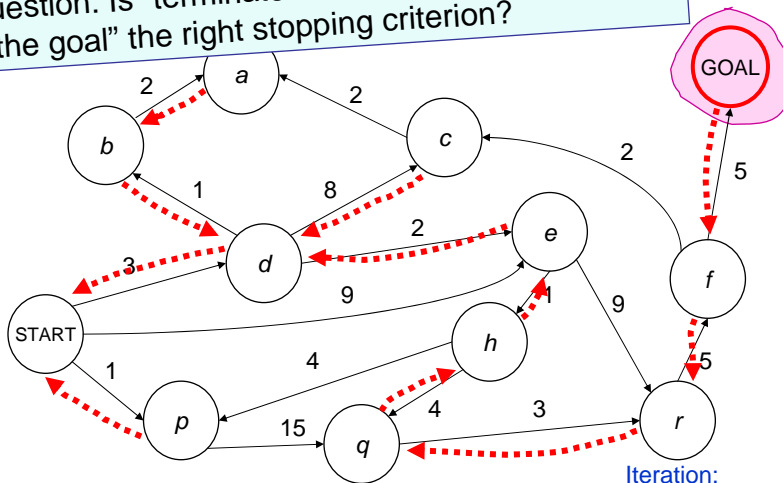


$PQ = \{ (G,23) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors

Slide 51

Question: Is "terminate as soon as you discover the goal" the right stopping criterion?

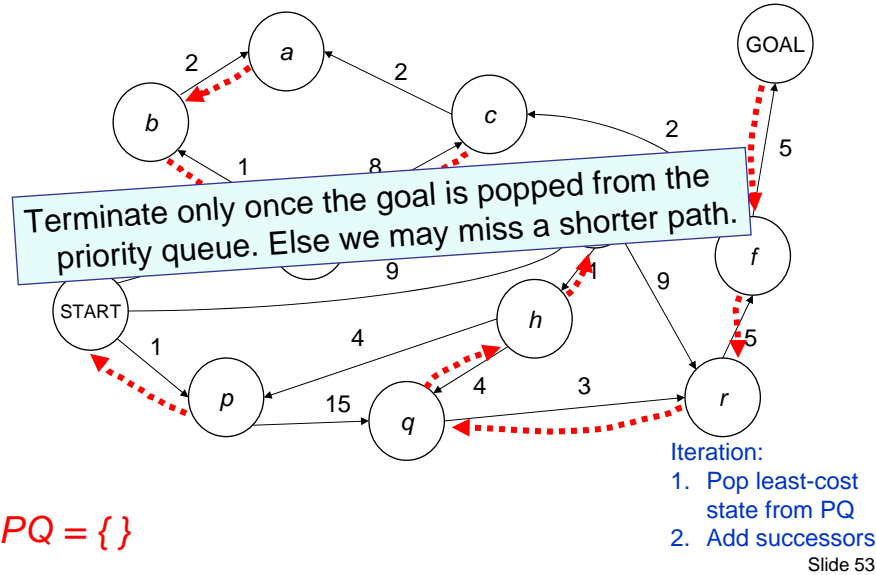


$PQ = \{ (G,23) \}$

- Iteration:
1. Pop least-cost state from PQ
  2. Add successors

Slide 52

## UCS terminates



## Judging a search algorithm

- **Completeness**: is the algorithm guaranteed to find a solution if a solution exists?
- Guaranteed to find **optimal**? (will it find the least cost path?)
- Algorithmic **time complexity**
- **Space complexity** (memory use)

Variables:

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps

*How would we judge our algorithms?*

Slide 54

## Judging a search algorithm

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search				
LCBFS	Least Cost BFS				
UCS	Uniform Cost Search				

Slide 55



## Judging a search algorithm

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$

Slide 56


## Search Tree Representation



What order do we go through the search tree with BFS?

Slide 57

## Depth First Search




An alternative to BFS. Always expand from the most-recently-expanded node, if it has any untried successors. Else backup to the previous node on the current path.

Slide 58

## DFS in action


START  
 START *d*  
 START *db*  
 START *dba*  
 START *dc*  
 START *dca*  
 START *de*  
 START *der*  
 START *derf*  
 START *derfc*  
 START *derfca*  
 START *derf* GOAL



Slide 59

## DFS Search tree traversal

Can you draw in the order in which the search-tree nodes are visited?



Slide 60



## Judging a search algorithm

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
DFS	Depth First Search				

Slide 63

## Judging a search algorithm

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
DFS	Depth First Search	N	N	N/A	N/A

Slide 64



## Judging a search algorithm

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
DFS**	Depth First Search				

Assuming Acyclic Search Space

Slide 65

## Judging a search algorithm

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
LMAX	Length of longest path from start to anywhere
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
DFS**	Depth First Search	Y	N	$O(B^{LMAX})$	$O(LMAX)$

Assuming Acyclic Search Space

Slide 66

## Questions to ponder

- How would you prevent DFS from looping?
- How could you force it to give an optimal solution?

Slide 67

## Questions to ponder

- How would you prevent DFS from looping?
- How could you force it to give an optimal solution?

Answer 1:

PC-DFS (Path Checking DFS):

Answer 2:

MEMDFS (Memoizing DFS):

Slide 68

## Questions to ponder

- How would you prevent DFS from looping?
- How could you force it to give an optimal solution?

Answer 1:

PC-DFS (Path Checking DFS):

Don't recurse on a state if that state is already in the current path

Answer 2:

MEMDFS (Memoizing DFS):

Remember all states expanded so far. Never expand anything twice.

Slide 69

## Questions to ponder

- How would you prevent DFS from looping?
- How could you force it to give an optimal solution?

Are there occasions when PCDFS is better than MEMDFS?

Are there occasions when MEMDFS is better than PCDFS?

Answer 1:

PC-DFS (Path Checking DFS):

Don't recurse on a state if that state is already in the current path

Answer 2:

MEMDFS (Memoizing DFS):

Remember all states expanded so far. Never expand anything twice.

Slide 70

## Judging a search algorithm

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
LMAX	Length of longest <b>cycle-free</b> path from start to anywhere
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
PCDFS	Path Check DFS				
MEMDFS	Memoizing DFS				

Slide 71

## Judging a search algorithm

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
LMAX	Length of longest <b>cycle-free</b> path from start to anywhere
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
PCDFS	Path Check DFS	Y	N	$O(B^{LMAX})$	$O(LMAX)$
MEMDFS	Memoizing DFS	Y	N	$O(\min(N, B^{LMAX}))$	$O(\min(N, B^{LMAX}))$

Slide 72

## Judging a search algorithm


N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
LMAX	Length of longest <b>cycle-free</b> path from start to anywhere
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
PCDFS	Path Check DFS	Y	N	$O(B^{LMAX})$	$O(LMAX)$
MEMDFS	Memoizing DFS	Y	N	$O(\min(N, B^{LMAX}))$	$O(\min(N, B^{LMAX}))$

Slide 73

## Maze example

Imagine states are cells in a maze, you can move N, E, S, W. What would **plain DFS** do, assuming it always expanded the E successor first, then N, then W, then S?

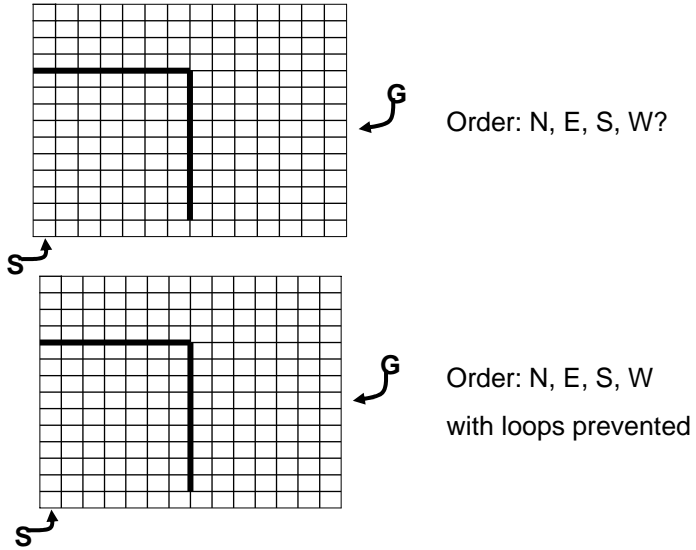


Other questions:

- What would BFS do?
- What would PCDFS do?
- What would MEMDFS do?

Slide 74

## Two other DFS examples



Slide 75

## Forward DFSearch or Backward DFSearch

If you have a predecessors() function as well as a successors() function you can begin at the goal and depth-first-search backwards until you hit a start.

Why/When might this be a good idea?

Slide 76

## Invent An Algorithm Time!

Here's a way to dramatically decrease costs sometimes. Bidirectional Search. Can you guess what this algorithm is, and why it can be a huge cost-saver?

Slide 77

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
LMAX	Length of longest <b>cycle-free</b> path from start to anywhere
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
PCDFS	Path Check DFS	Y	N	$O(B^{LMAX})$	$O(LMAX)$
MEMDFS	Memoizing DFS	Y	N	$O(\min(N, B^{LMAX}))$	$O(\min(N, B^{LMAX}))$
BIBFS	Bidirection BF Search				

Slide 78

N	number of states in the problem				
B	the average branching factor (the average number of successors) ( $B > 1$ )				
L	the length of the path from start to goal with the shortest number of steps				
LMAX	Length of longest <b>cycle-free</b> path from start to anywhere				
Q	the average size of the priority queue				

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
PCDFS	Path Check DFS	Y	N	$O(B^{LMAX})$	$O(LMAX)$
MEMDFS	Memoizing DFS	Y	N	$O(\min(N, B^{LMAX}))$	$O(\min(N, B^{LMAX}))$
BIBFS	Bidirection BF Search	Y	All trans same cost	$O(\min(N, 2B^{L/2}))$	$O(\min(N, 2B^{L/2}))$

Slide 79

## Iterative Deepening

Iterative deepening is a simple algorithm which uses DFS as a subroutine:

1. Do a DFS which only searches for paths of length 1 or less. (DFS gives up any path of length 2)
2. If "1" failed, do a DFS which only searches paths of length 2 or less.
3. If "2" failed, do a DFS which only searches paths of length 3 or less.  
....and so on until success

Cost is

$$O(b^1 + b^2 + b^3 + b^4 \dots + b^L) = O(b^L)$$


Can be much better than regular DFS. But cost can be much greater than the number of states.

Slide 80



## Maze example

Imagine states are cells in a maze, you can move N, E, S, W. What would **Iterative Deepening** do, assuming it always expanded the E successor first, then N, then W, then S?



Expansion order E, N, W, S

Slide 81

N	number of states in the problem
B	the average branching factor (the average number of successors) ( $B > 1$ )
L	the length of the path from start to goal with the shortest number of steps
LMAX	Length of longest <b>cycle-free</b> path from start to anywhere
Q	the average size of the priority queue

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
PCDFS	Path Check DFS	Y	N	$O(B^{LMAX})$	$O(LMAX)$
MEMDFS	Memoizing DFS	Y	N	$O(\min(N, B^{LMAX}))$	$O(\min(N, B^{LMAX}))$
BIBFS	Bidirection BF Search	Y	All trans same cost	$O(\min(N, 2B^{L/2}))$	$O(\min(N, 2B^{L/2}))$
ID	Iterative Deepening				

Slide 82

N	number of states in the problem				
B	the average branching factor (the average number of successors) ( $B > 1$ )				
L	the length of the path from start to goal with the shortest number of steps				
LMAX	Length of longest <b>cycle-free</b> path from start to anywhere				
Q	the average size of the priority queue				

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	if all transitions same cost	$O(\min(N, B^L))$	$O(\min(N, B^L))$
LCBFS	Least Cost BFS	Y	Y	$O(\min(N, B^L))$	$O(\min(N, B^L))$
UCS	Uniform Cost Search	Y	Y	$O(\log(Q) * \min(N, B^L))$	$O(\min(N, B^L))$
PCDFS	Path Check DFS	Y	N	$O(B^{LMAX})$	$O(LMAX)$
MEMDFS	Memoizing DFS	Y	N	$O(\min(N, B^{LMAX}))$	$O(\min(N, B^{LMAX}))$
BIBFS	Bidirection BF Search	Y	All trans same cost	$O(\min(N, 2B^{L/2}))$	$O(\min(N, 2B^{L/2}))$
ID	Iterative Deepening	Y	if all transitions same cost	$O(B^L)$	$O(L)$

Slide 83

## Best First “Greedy” Search

Needs a *heuristic*. A heuristic function maps a state onto an estimate of the cost to the goal from that state.

Can you think of examples of heuristics?

E.G. for the 8-puzzle?

E.G. for planning a path through a maze?

Denote the heuristic by a function  $h(s)$  from states to a cost value.

Slide 84

## Heuristic Search

Suppose in addition to the standard search specification we also have a *heuristic*.

***A heuristic function maps a state onto an estimate of the cost to the goal from that state.***


Can you think of examples of heuristics?

- E.G. for the 8-puzzle?
- E.G. for planning a path through a maze?

Denote the heuristic by a function  $h(s)$  from states to a cost value.

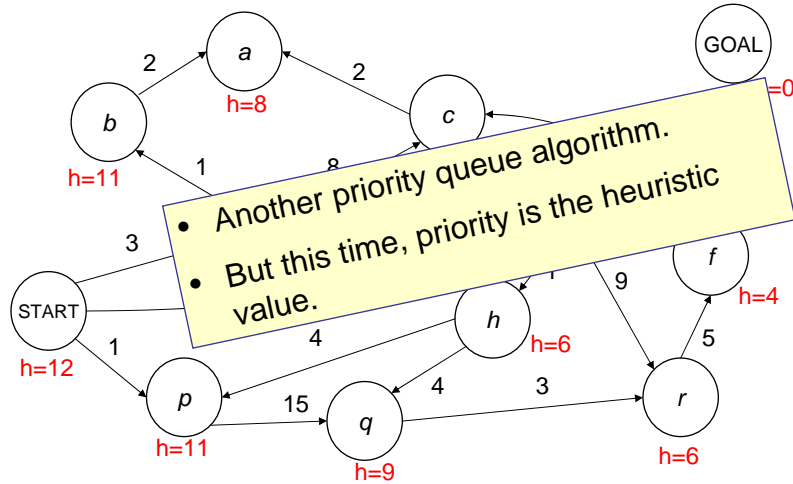
Slide 85

## Euclidian Heuristic



Slide 86

## Euclidian Heuristic



Slide 87

## Best First “Greedy” Search

```

Init-PriQueue(PQ)
Insert-PriQueue(PQ,START,h(START))
while (PQ is not empty and PQ does not contain a goal state)
    (s , h ) := Pop-least(PQ)
    foreach s' in succs(s)
        if s' is not already in PQ and s' never previously been visited
            Insert-PriQueue(PQ,s',h(s'))
    
```

Algorithm	Complete	Optimal	Time	Space
BestFS Best First Search	Y	<b>N</b>	$O(\min(N, B^{LMAX}))$	$O(\min(N, B^{LMAX}))$

A few improvements to this algorithm can make things much better. It's a little thing we like to call: A\*....

*...to be continued!*

Slide 88

## What you should know

- Thorough understanding of BFS, LCBFS, UCS, PCDFS, MEMDFS
- Understand the concepts of whether a search is complete, optimal, its time and space complexity
- Understand the ideas behind iterative deepening and bidirectional search
- Be able to discuss at cocktail parties the pros and cons of the above searches

Slide 89