

CS 188: Artificial Intelligence

Spring 2007

Lecture 3: Queue-Based Search

1/23/2007

Srini Narayanan – UC Berkeley

Many slides over the course adapted from Dan Klein, Stuart Russell or Andrew Moore

Announcements

- § Assignment 1 due 1/30 11:59 PM
 - § You can do most of it after today
- § Sections start this week
- § Stay tuned for Python Lab

Summary

- § Agents interact with environments through actuators and sensors
 - § The agent function describes what the agent does in all circumstances
 - § The agent program calculates the agent function
 - § The performance measure evaluates the environment sequence

- § A perfectly rational agent maximizes expected performance

- § PEAS descriptions define task environments

- § Environments are categorized along several dimensions:
 - § Observable? Deterministic? Episodic? Static? Discrete? Single-agent?

- § Problem-solving agents make a plan, then execute it

- § State space encodings of problems

Problem-Solving Agents

```
function SIMPLE-PROBLEM-SOLVING-AGENT(percept) returns an action
  static: seq, an action sequence, initially empty
           state, some description of the current world state
           goal, a goal, initially null
           problem, a problem formulation

  state ← UPDATE-STATE(state, percept)
  if seq is empty then
    goal ← FORMULATE-GOAL(state)
    problem ← FORMULATE-PROBLEM(state, goal)
    seq ← SEARCH(problem)
  action ← FIRST(seq); seq ← REST(seq)
  return action
```

This is the hard part!

§ This offline problem solving!

§ Solution is executed “eyes closed.”

Tree Search

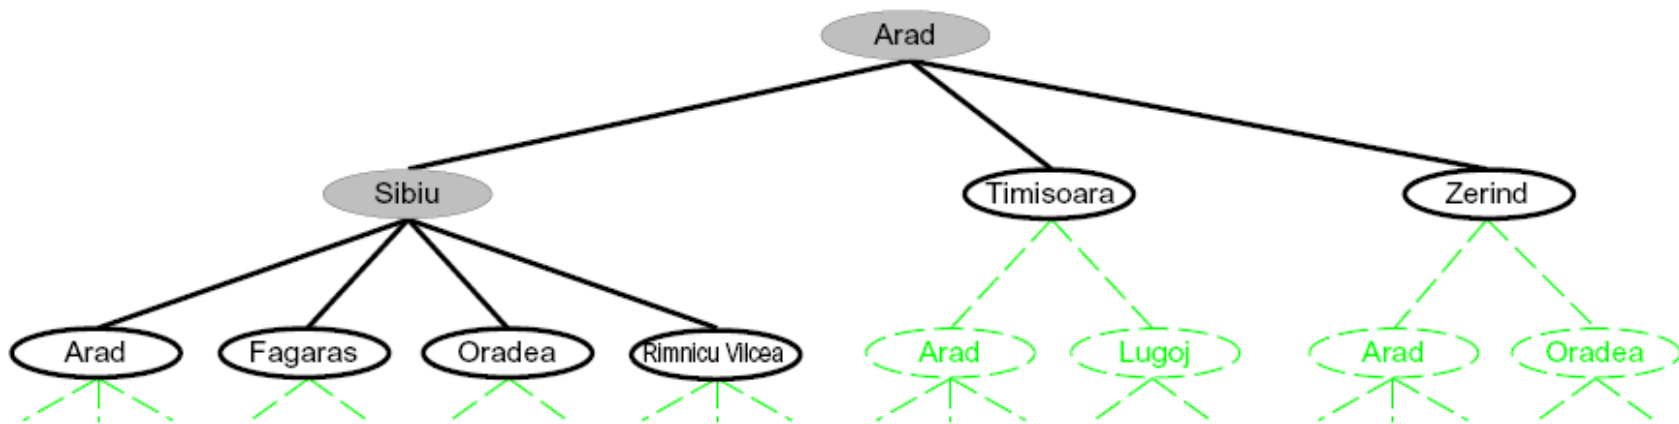
§ Basic solution method for graph problems

§ Offline simulated exploration of state space

§ Searching a model of the space, not the real world

```
function TREE-SEARCH(problem, strategy) returns a solution, or failure
  initialize the search tree using the initial state of problem
  loop do
    if there are no candidates for expansion then return failure
    choose a leaf node for expansion according to strategy
    if the node contains a goal state then return the corresponding solution
    else expand the node and add the resulting nodes to the search tree
  end
```

A Search Tree



§ Search:

§ Expand out possible plans

§ Maintain a **fringe** of unexpanded plans

§ Try to expand as few tree nodes as possible

Tree Search

```
function TREE-SEARCH(problem, fringe) returns a solution, or failure
  fringe ← INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)
  loop do
    if fringe is empty then return failure
    node ← REMOVE-FRONT(fringe)
    if GOAL-TEST(problem, STATE(node)) then return SOLUTION(node)
    fringe ← INSERTALL(EXPAND(node, problem), fringe)
```

```
function EXPAND(node, problem) returns a set of nodes
  successors ← the empty set; state ← STATE[node]
  for each action, result in SUCCESSOR-FN(problem, state) do
    s ← a new NODE
    PARENT-NODE[s] ← node; ACTION[s] ← action; STATE[s] ← result
    PATH-COST[s] ← PATH-COST[node] + STEP-COST(state, action, result)
    DEPTH[s] ← DEPTH[node] + 1
    add s to successors
  return successors
```

General Tree Search

§ Important ideas:

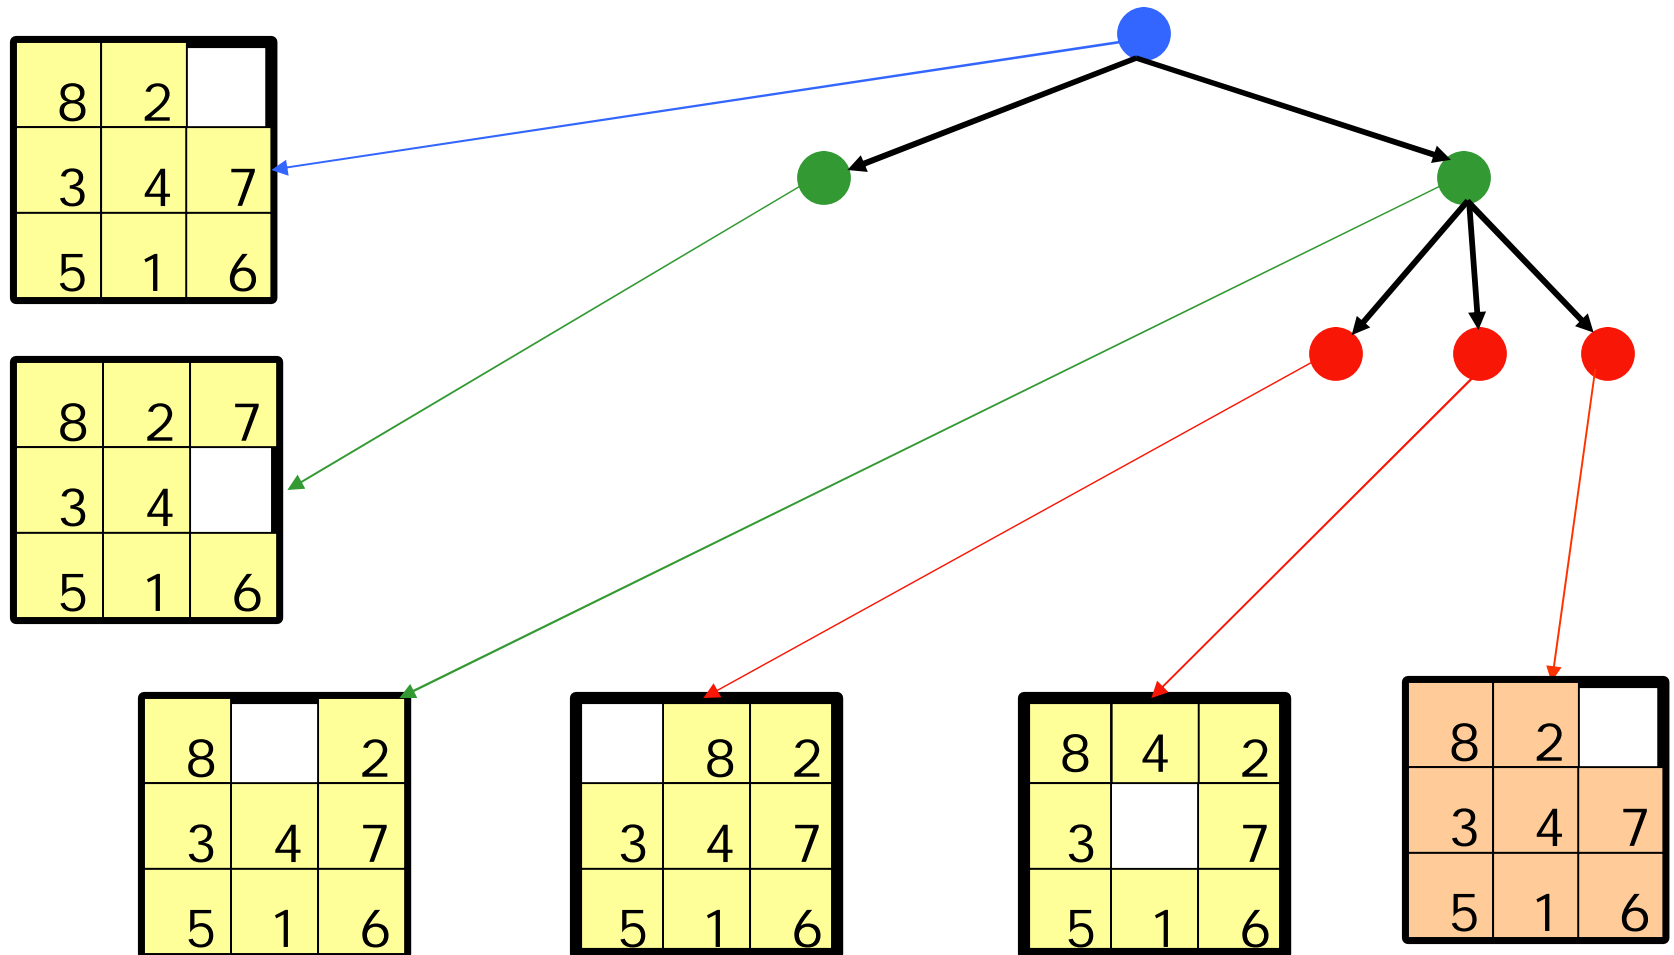
- § Fringe

- § Expansion

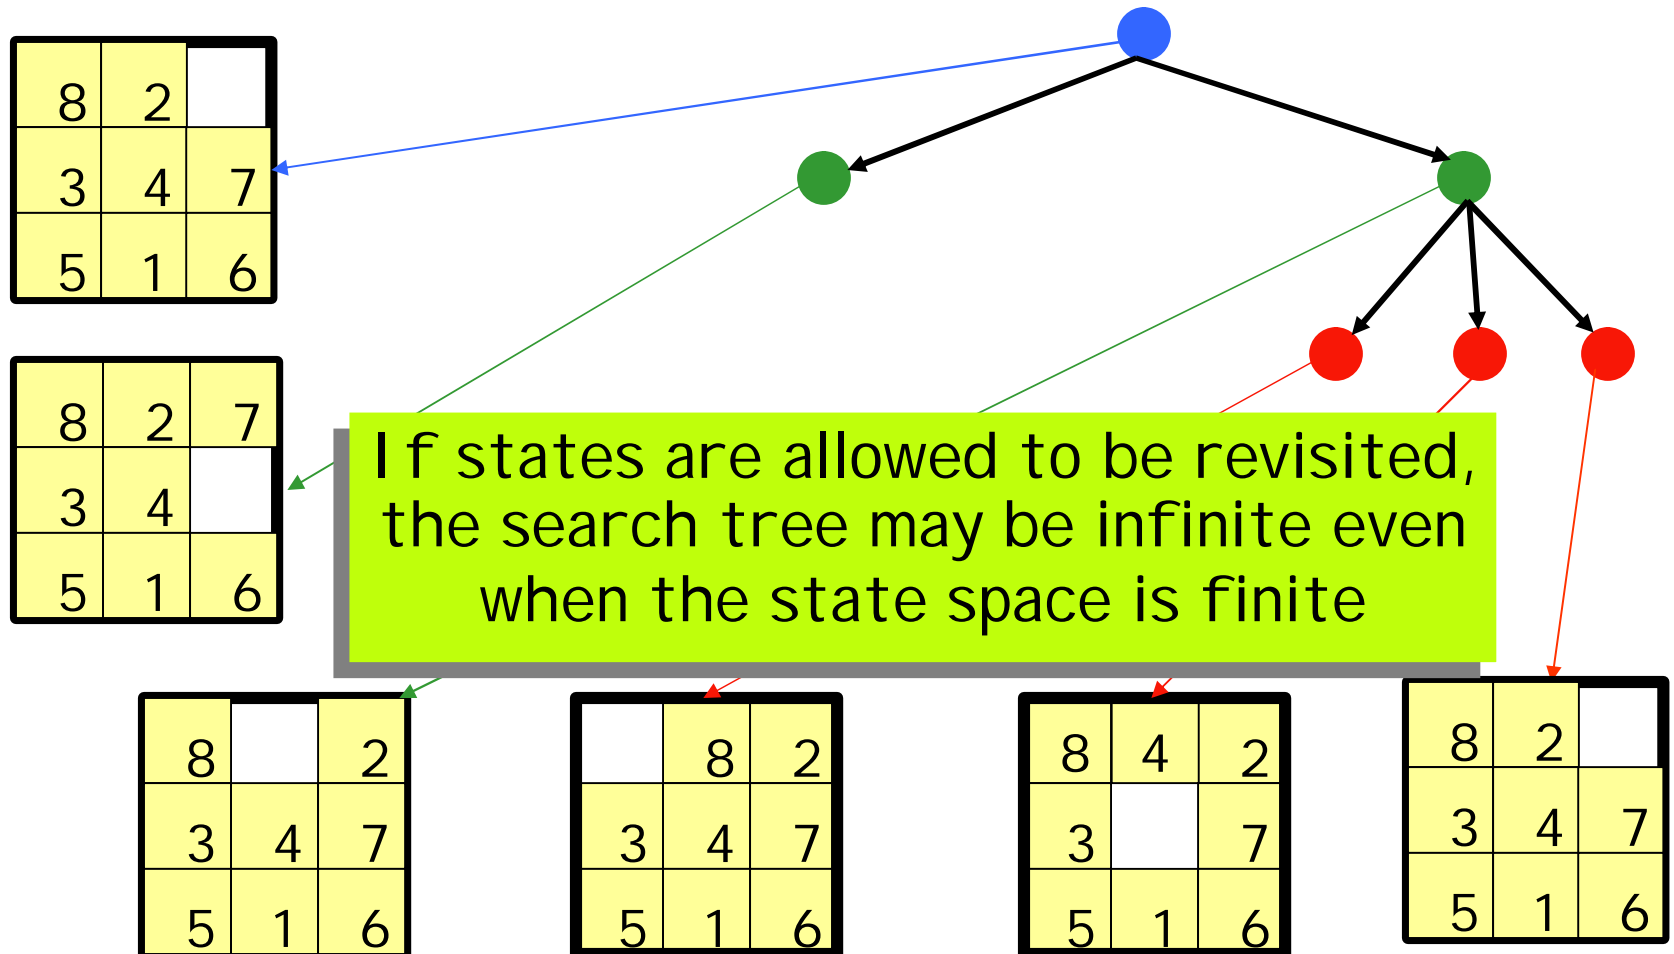
- § Exploration strategy

§ Main question: which fringe nodes to explore?

Search Nodes vs. States




Search Nodes vs. States



States vs. Nodes

- § Problem graphs have problem states
 - § Have successors
- § Search trees have search nodes
 - § Have parents, children, depth, path cost, etc.
 - § Expand uses successor function to create new search tree nodes
 - § The same problem state may be in multiple search tree nodes



Uninformed search strategies

§ (a.k.a. blind search) = use only information available in problem definition.

§ When strategies can determine whether one non-goal state is better than another → *informed search*.

§ Categories defined by expansion algorithm:

§ Breadth-first search

§ Depth-first search

§ (Depth-limited search)

§ Iterative deepening search

§ Uniform-cost search

§ Bidirectional search

State Space Graphs


§ There's some big graph in which

§ Each state is a node

§ Each successor is an outgoing arc

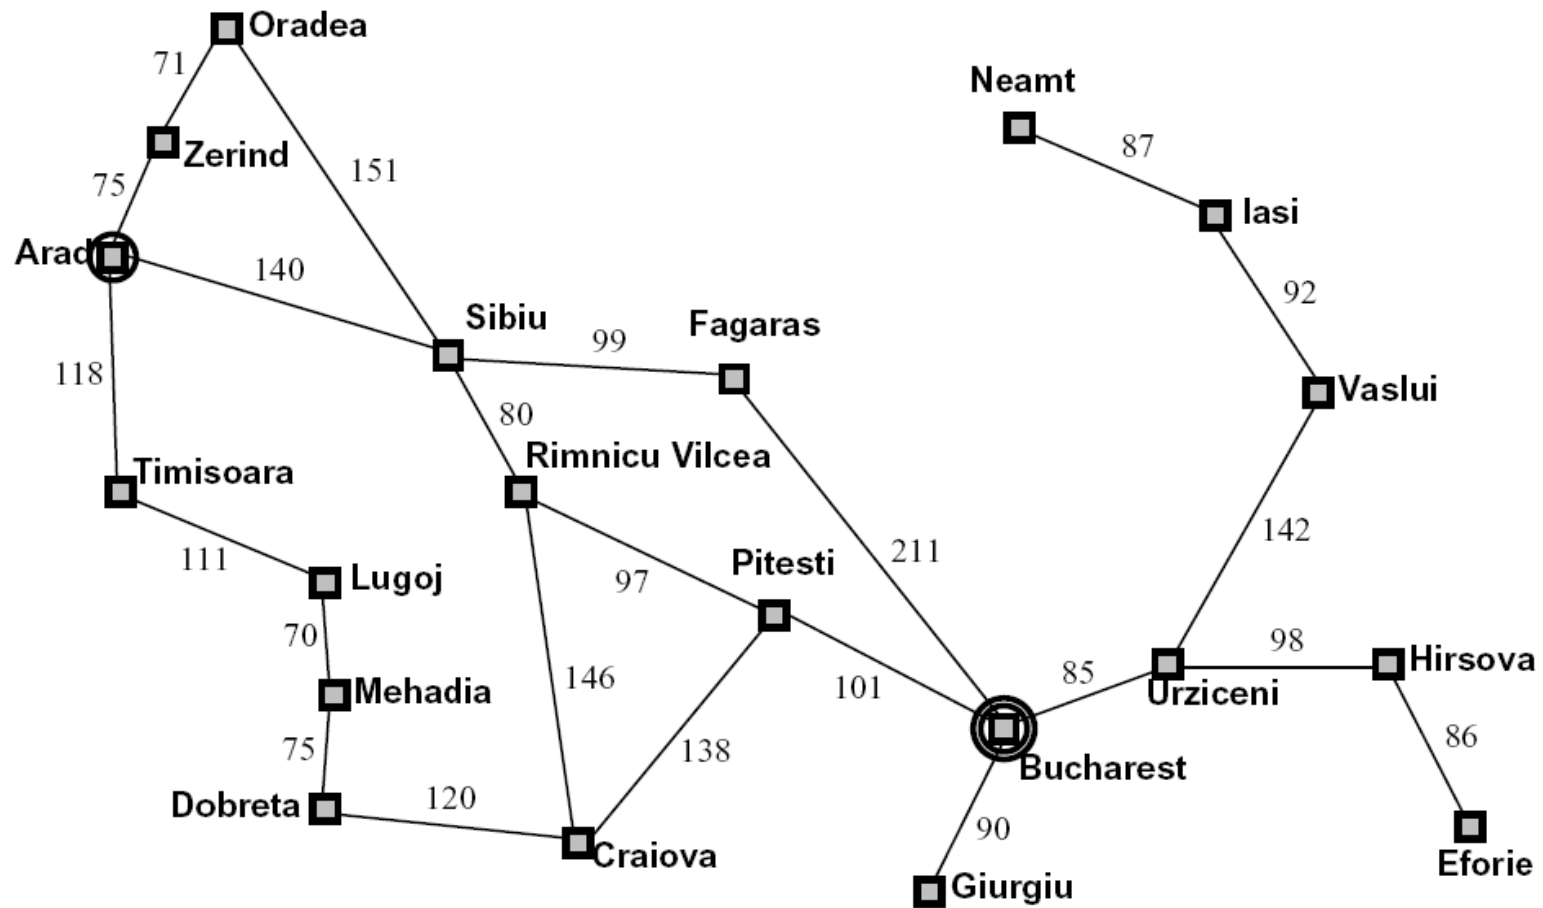
§ Important: For most problems we could never actually build this graph

§ How many states in 8-puzzle?




Laughably tiny search graph for a tiny search problem


Example: Romania



Example: Tree Search




State Graphs vs Search Trees



Each NODE in in the search tree is an entire PATH in the problem graph.



We almost always construct both on demand – and we construct as little as possible.



Review: Depth First Search

Strategy: expand deepest node first

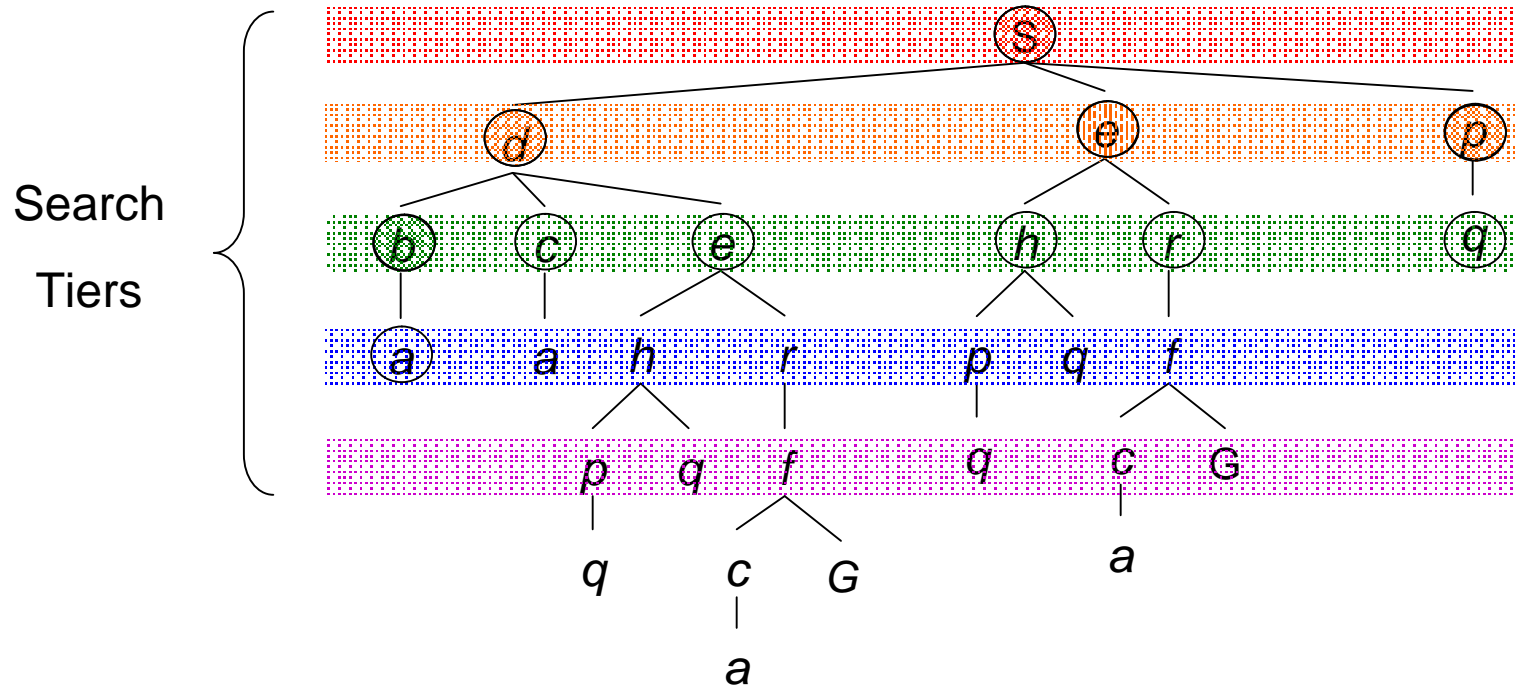

Implementation:
Fringe is a LIFO stack



Review: Breadth First Search

Strategy: expand shallowest node first

Implementation: Fringe is a FIFO queue



Search Algorithm Properties


- § **Complete?** Guaranteed to find a solution if one exists?
- § **Optimal?** Guaranteed to find the least cost path?
- § **Time complexity?**
- § **Space complexity?**

Variables:

| | |
|-------|--|
| n | Number of states in the problem |
| b | The average branching factor B (the average number of successors) |
| C^* | Cost of least cost solution |
| s | Depth of the shallowest solution |
| m | Max depth of the search tree |

DFS

| Algorithm | | Complete | Optimal | Time | Space |
|-----------|--------------------|----------|---------|----------|----------|
| DFS | Depth First Search | N | N | Infinite | Infinite |

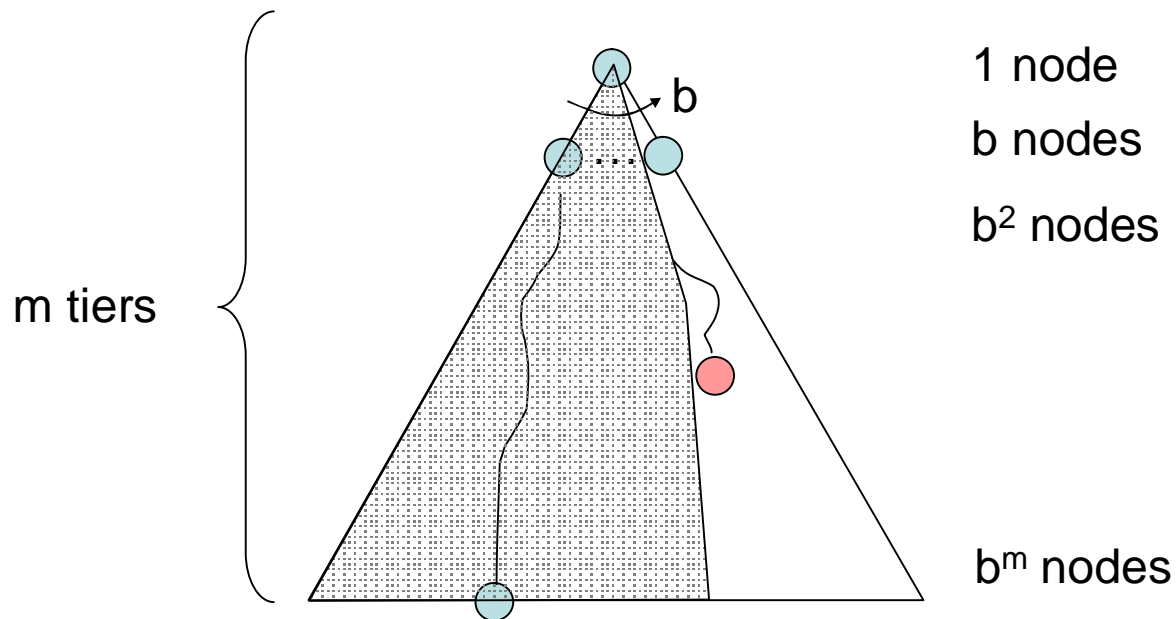


§ Infinite paths make DFS incomplete...

§ How can we fix this?

DFS

§ With cycle checking, DFS is complete.




| Algorithm | | Complete | Optimal | Time | Space |
|-----------|------------------|----------|---------|----------|---------|
| DFS | w/ Path Checking | Y | N | $O(b^m)$ | $O(bm)$ |

§ When is DFS optimal?

BFS

| Algorithm | | Complete | Optimal | Time | Space |
|-----------|------------------|----------|---------|--------------|--------------|
| DFS | w/ Path Checking | Y | N | $O(b^m)$ | $O(bm)$ |
| BFS | | Y | N* | $O(b^{s+1})$ | $O(b^{s+1})$ |




§ When is BFS optimal?

Comparisons

§ When will BFS outperform DFS?

§ When will DFS outperform BFS?

Costs on Actions





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

We will quickly cover an algorithm which does find the least-cost path.

Uniform Cost Search

*Expand cheapest node first:
Fringe is a priority queue*





Priority Queue Refresher

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

| | |
|----------------------------------|---|
| <code>pq.push(key, value)</code> | inserts <i>(key, value)</i> into the queue. |
| <code>pq.pop()</code> | returns the key with the lowest value, and removes it from the queue. |


§ You can promote or demote keys by resetting their priorities

§ Unlike a regular queue, insertions into a priority queue are not constant time, usually $O(\log n)$

§ We'll need priority queues for most cost-sensitive search methods.

Uniform Cost Search


§ What will UCS do for this graph?



§ What does this mean for completeness?

Uniform Cost Search

| Algorithm | | Complete | Optimal | Time | Space |
|-----------|------------------|----------|---------|----------------|----------------|
| DFS | w/ Path Checking | Y | N | $O(b^m)$ | $O(bm)$ |
| BFS | | Y | N | $O(b^{s+1})$ | $O(b^{s+1})$ |
| UCS | | Y* | Y | $O(b^{C^*/e})$ | $O(b^{C^*/e})$ |



We'll talk more about uniform cost search's failure cases later...

Uniform Cost Problems



§ Remember: explores increasing cost contours

§ The good: UCS is complete and optimal!

§ The bad:

§ Explores options in every “direction”

§ No information about goal location



Depth-limited search

depth-first search with depth limit l ,
i.e., nodes at depth l have no successors

§ Recursive implementation:

```
function DEPTH-LIMITED-SEARCH(problem, limit) returns soln/fail/cutoff
  RECURSIVE-DLS(MAKE-NODE(INITIAL-STATE[problem]), problem, limit)

function RECURSIVE-DLS(node, problem, limit) returns soln/fail/cutoff
  cutoff-occurred? ← false
  if GOAL-TEST[problem](STATE[node]) then return SOLUTION(node)
  else if DEPTH[node] = limit then return cutoff
  else for each successor in EXPAND(node, problem) do
    result ← RECURSIVE-DLS(successor, problem, limit)
    if result = cutoff then cutoff-occurred? ← true
    else if result ≠ failure then return result
  if cutoff-occurred? then return cutoff else return failure
```

Iterative deepening search

```
function ITERATIVE-DEEPENING-SEARCH(problem) returns a solution, or fail-  
ure  
  inputs: problem, a problem  
  for depth  $\leftarrow$  0 to  $\infty$  do  
    result  $\leftarrow$  DEPTH-LIMITED-SEARCH(problem, depth)  
    if result  $\neq$  cutoff then return result
```

Iterative deepening search $l = 0$

Limit = 0




Iterative deepening search $l=1$

Limit = 1




Iterative deepening search $l=2$

Limit = 2



Iterative deepening search $l=3$

Limit = 3



Iterative deepening search

§ Number of nodes generated in a depth-limited search to depth d with branching factor b :

$$N_{DLS} = b^0 + b^1 + b^2 + \dots + b^{d-2} + b^{d-1} + b^d$$

§ Number of nodes generated in an iterative deepening search to depth d with branching factor b :

$$N_{IDS} = (d+1)b^0 + d b^1 + (d-1)b^2 + \dots + 3b^{d-2} + 2b^{d-1} + 1b^d$$

§ For $b = 10$, $d = 5$,

§ $N_{DLS} = 1 + 10 + 100 + 1,000 + 10,000 + 100,000 = 111,111$


§ $N_{IDS} = 6 + 50 + 400 + 3,000 + 20,000 + 100,000 = 123,456$

§ Overhead = $(123,456 - 111,111)/111,111 = 11\%$

Iterative Deepening

Iterative deepening uses DFS as a subroutine:

1. Do a DFS which only searches for paths of length 1 or less. (DFS gives up on 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.





| Algorithm | | Complete | Optimal | Time | Space |
|-----------|------------------|----------|---------|--------------|--------------|
| DFS | w/ Path Checking | Y | N | $O(b^m)$ | $O(bm)$ |
| BFS | | Y | N^* | $O(b^{s+1})$ | $O(b^{s+1})$ |
| ID | | Y | N^* | $O(b^d)$ | $O(bd)$ |

Summary of algorithms

| Criterion | Breadth-First | Uniform-Cost | Depth-First | Depth-Limited | Iterative Deepening |
|-----------|---------------|-------------------------------------|-------------|---------------|---------------------|
| Complete? | Yes | Yes | No | No | Yes |
| Time | $O(b^{d+1})$ | $O(b^{\lceil C^*/\epsilon \rceil})$ | $O(b^m)$ | $O(b^l)$ | $O(b^d)$ |
| Space | $O(b^{d+1})$ | $O(b^{\lceil C^*/\epsilon \rceil})$ | $O(bm)$ | $O(bl)$ | $O(bd)$ |
| Optimal? | Yes | Yes | No | No | Yes |


Extra Work?

§ Failure to detect repeated states can cause exponentially more work. Why?



Graph Search


§ In BFS, for example, we shouldn't bother expanding the circled nodes (why?)



Graph Search

§ Very simple fix: never expand a node twice

```
function GRAPH-SEARCH(problem, fringe) returns a solution, or failure
  closed ← an empty set
  fringe ← INSERT(MAKE-NODE(INITIAL-STATE[problem]), fringe)
  loop do
    if fringe is empty then return failure
    node ← REMOVE-FRONT(fringe)
    if GOAL-TEST(problem, STATE[node]) then return node
    if STATE[node] is not in closed then
      add STATE[node] to closed
      fringe ← INSERTALL(EXPAND(node, problem), fringe)
    end
  end
```



§ Can this wreck correctness? Why or why not?

Search Gone Wrong?

