

A - Informed Search*

15-491, Fall 2008

Manuela Veloso

Carnegie Mellon

(Thanks to past instructors and several book authors)

Uninformed Search Complexity


- N = Total number of states
- B = Average number of successors (branching factor)
- L = Length for start to goal with smallest number of steps
- Q = Average size of the priority queue
- L_{max} = Length of longest path from *START* to any state

Algorithm		Complete	Optimal	Time	Space
BFS	Breadth First Search	Y	Y, If all trans. have same cost	$O(\text{Min}(N, B^L))$	$O(\text{Min}(N, B^L))$
BIBFS	Bi- Direction. BFS	Y	Y, If all trans. have same cost	$O(\text{Min}(N, 2B^{L/2}))$	$O(\text{Min}(N, 2B^{L/2}))$
PCDFS	Path Check DFS	Y	N	$O(B^{L_{max}})$	$O(BL_{max})$
MEMDFS	Memorizing DFS	Y	N	$O(\text{Min}(N, B^{L_{max}}))$	$O(\text{Min}(N, B^{L_{max}}))$
IDS	Iterative Deepening	Y	Y, If all trans. have same cost	$O(B^L)$	$O(BL)$

Uninformed vs Informed

- Uninformed – only guided by
 - *successor* relationships
 - topological structure (leftmost,...)
 - length as number of nodes
- Informed
 - assume *cost* of edges
 - more knowledge?

Search Revisited

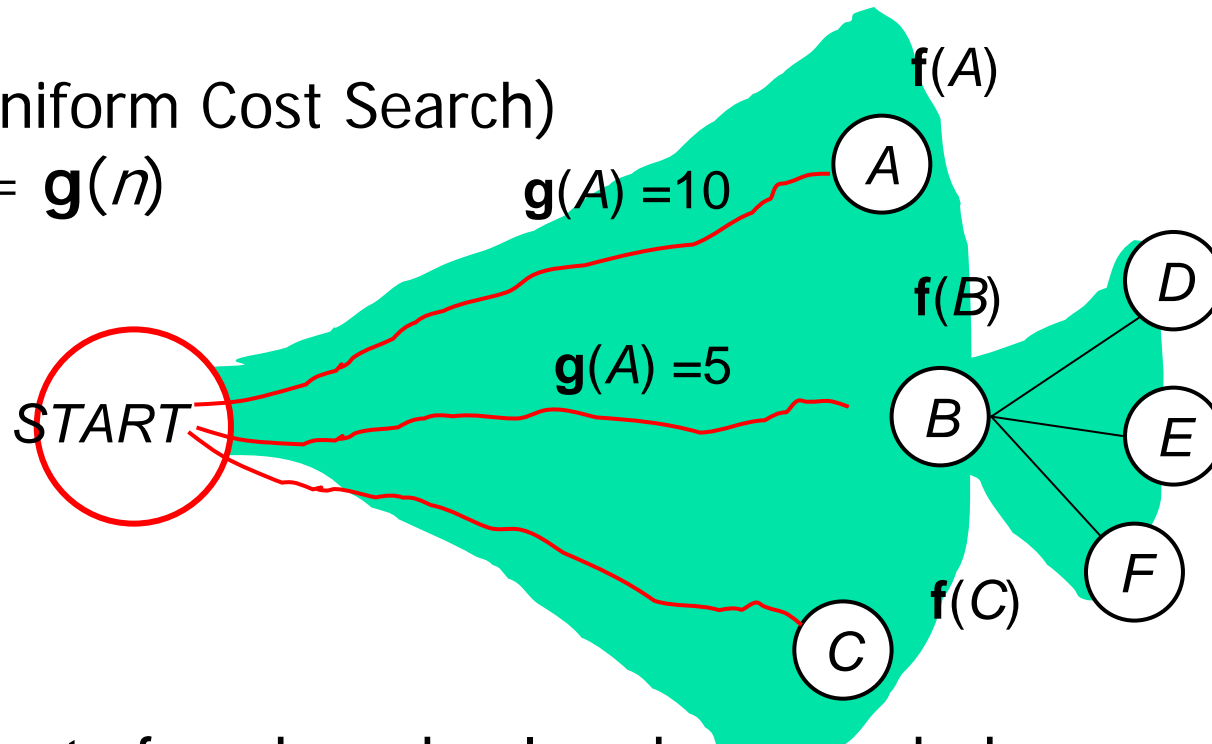


1. Store a value $f(s)$ at each state s
2. Choose the state with lowest f to expand next
3. Insert its successors

If $f(.)$ is chosen carefully, we will eventually find the lowest-cost sequence

- UCS (Uniform Cost Search)

$$f(n) = g(n)$$

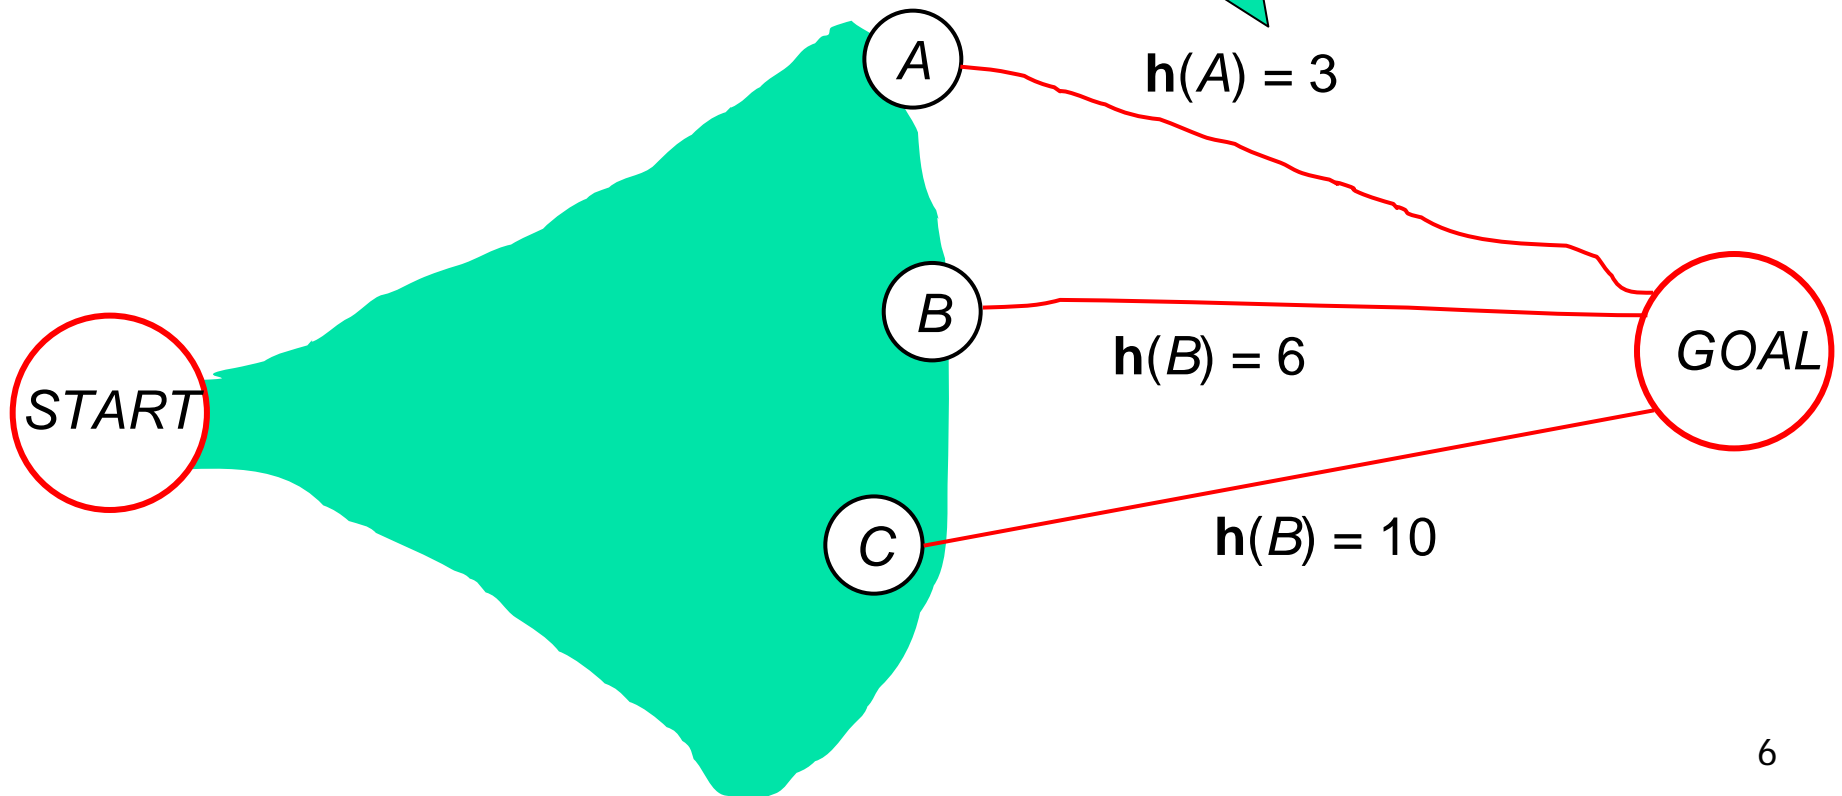


- $g(n)$ - cost of each node already expanded
length of shortest path from START to n
- Implementation – Store open successor states (waiting to be expanded) in a *priority queue* for efficient retrieval of minimum f
- Optimal → Guaranteed to find lowest cost sequence, *but guidance is about known path...*

Estimate “Cost” to Goal

- Introduce a function $h(s)$ to estimate the unknown distance from state s to the goal

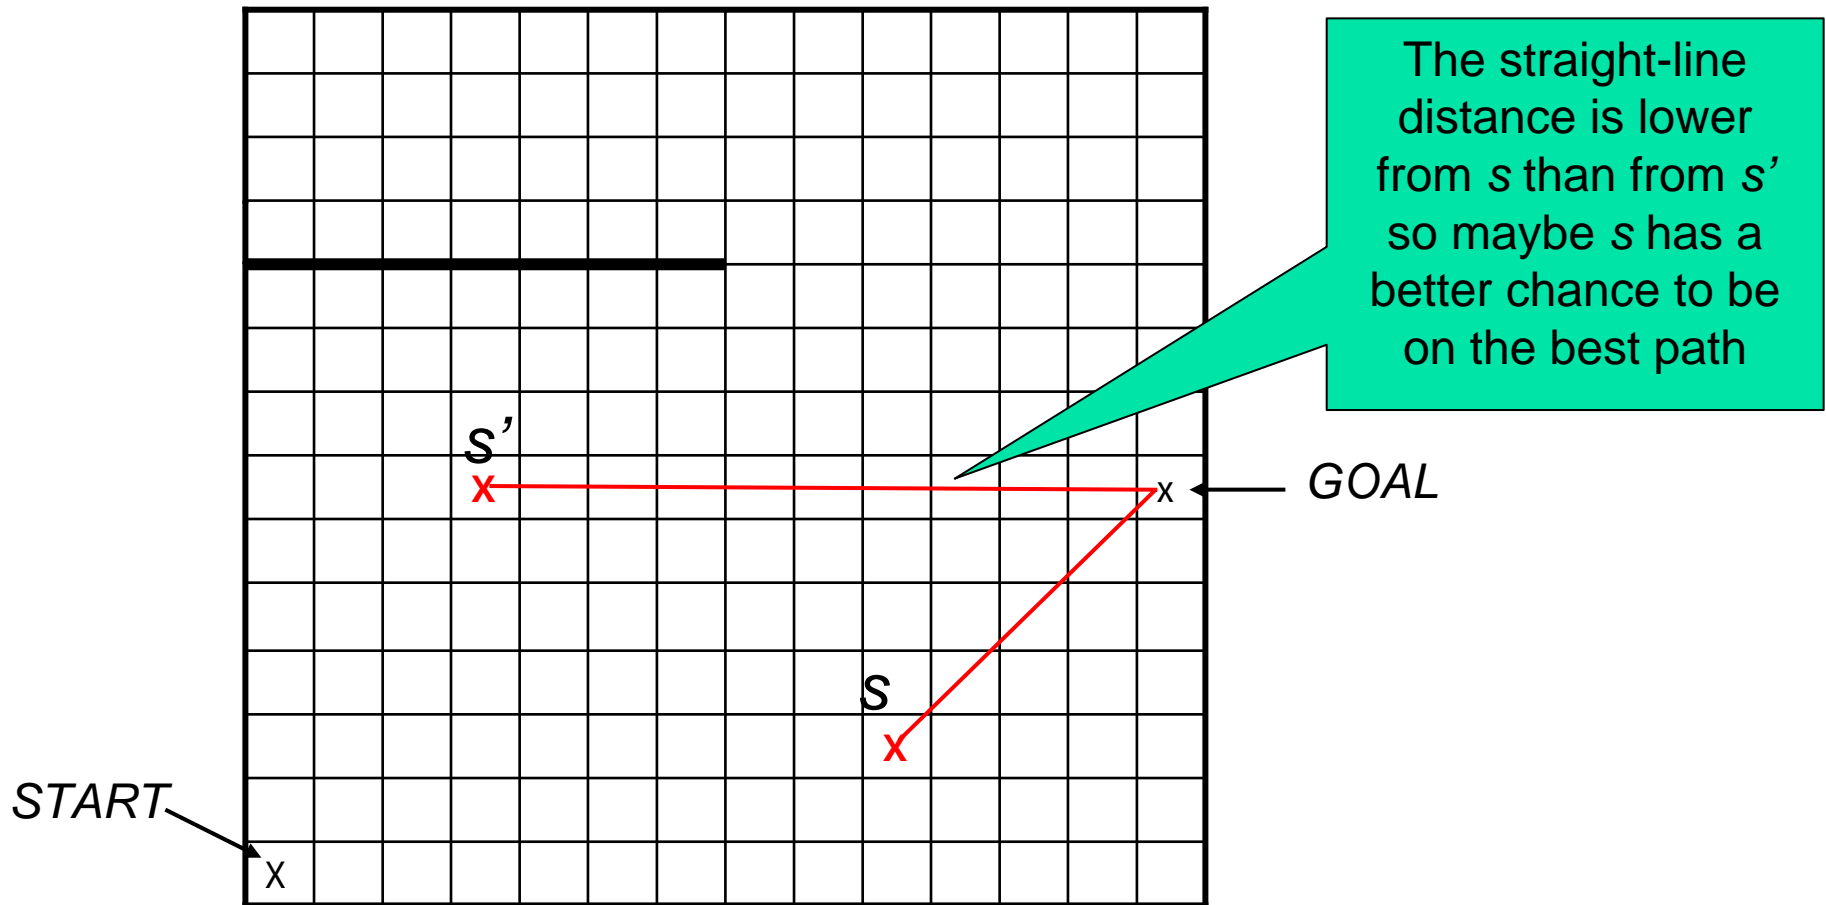
Our best **guess** is that A is closer to $GOAL$ than B so maybe it is a more promising state to expand



Heuristic Functions

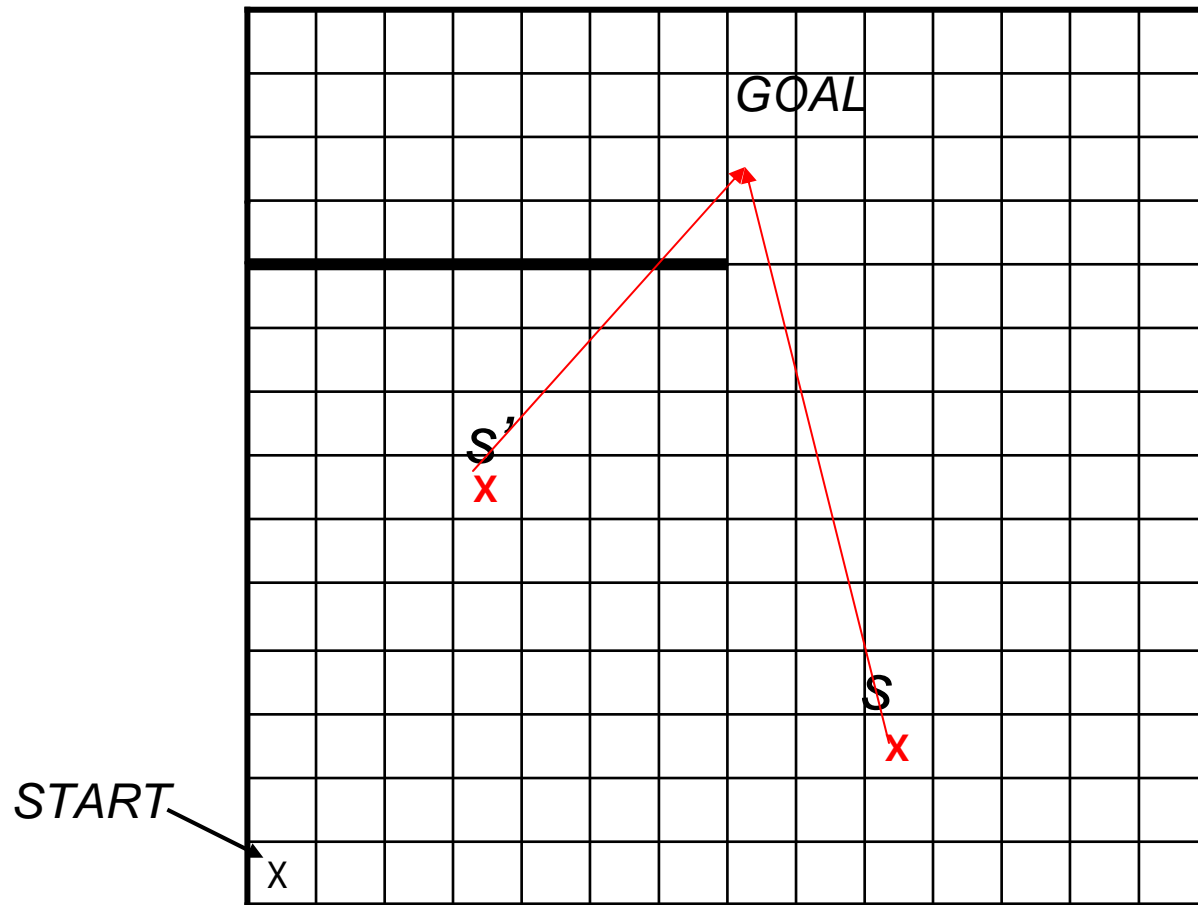
- h is a *heuristic* function for the search problem
- $h(s)$ = estimate of the cost of the shortest path from s to *GOAL*
- h cannot be computed solely from the states and transitions in the current problem → If we could, we would already know the optimal path!
- $h(.)$ is based on external knowledge about the problem → *informed* search
- Questions:
 1. Typical examples of h ?
 2. How to use h ?
 3. What are desirable/necessary properties of h ?

Heuristic Functions Example




- $h(s) = \text{Euclidean distance to } GOAL$

Heuristic Functions Example



- $h(s) = \text{Euclidean distance to } GOAL$
- *Euclidean distance is an heuristic.*

Heuristic Functions Example



- How could we define $h(s)$?



s



GOAL

Misplaced tiles:

$$h_1(s) = 7$$

Manhattan distance:

$$h_2(s) = 2 + 3 + 3 + 2 + 4 + 2 + 0 + 2 = 18$$

First Attempt: Greedy Best First Search

- Simplest use of heuristic function: Always select the node with smallest $\mathbf{h}(\cdot)$ for expansion (i.e., $\mathbf{f}(s) = \mathbf{h}(s)$)

Initialize PQ

Insert $START$ with value $\mathbf{h}(START)$ in PQ

While (PQ not empty and no goal state is in PQ)

 Pop the state s with the minimum value of \mathbf{h} from PQ


 For all s' in $\mathbf{succs}(s)$

 If s' is not already in PQ and has not already been visited

 Insert s' in PQ with value $\mathbf{h}(s')$

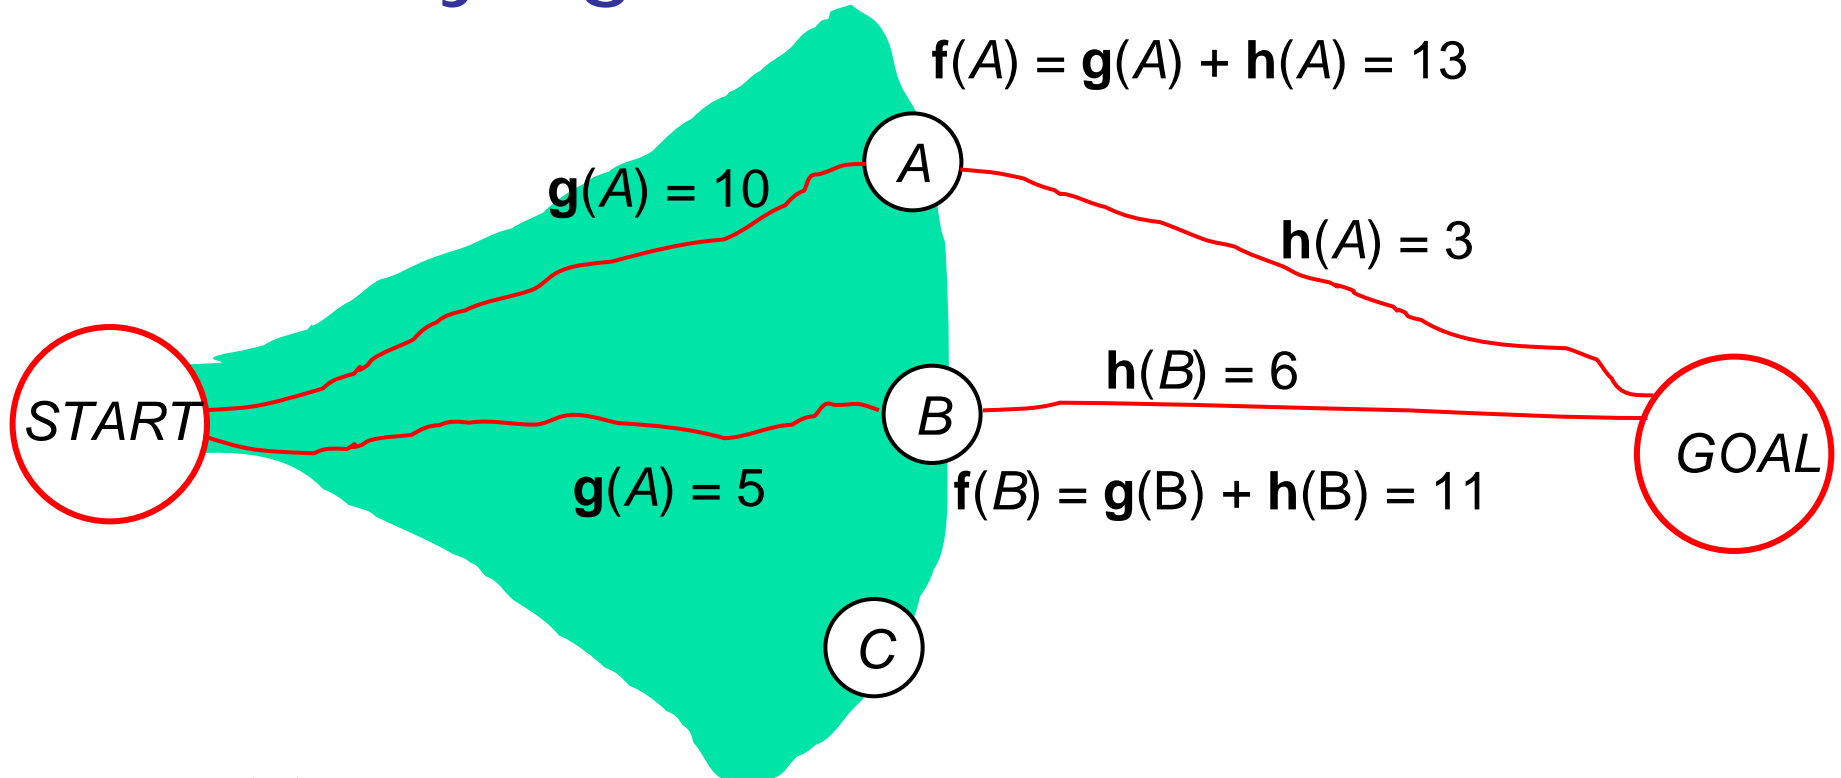
Problem

4



- What solution do we find in this case?
- Greedy search clearly not optimal, even though the heuristic function is non-stupid

Trying to Fix the Problem




- $g(s)$ is the cost from *START* to s only
- $h(s)$ estimates the cost from s to *GOAL*
- Key insight: $g(s) + h(s)$ estimates the **total** cost of the cheapest path from *START* to *GOAL* going through s
- → **A*** algorithm

A* Algorithm

- $f(s) = g(s) + h(s)$
- heuristics
 - good, less good..., alternative, efficiency
 - “easy” to define...
- efficiency

Can A* Fix the Problem?



$\{(START, 4)\}$

$\{(A, 5)\}$

$$(f(A) = h(A) + g(A) = 3 + g(START) + \text{cost}(START, A) = 3 + 0 + 2)$$

$\{(B, 5) (C, 7)\}$


$$(f(C) = h(C) + g(C) = 1 + g(A) + \text{cost}(A, C) = 1 + 2 + 4)$$

$\{(C, 5)\}$

$$(f(C) = h(C) + g(C) = 1 + g(B) + \text{cost}(B, C) = 1 + 3 + 1)$$

$\{(GOAL, 6)\}$

Can A* Fix the Problem?



C is placed in the queue with backpointers {A, START}

$\{(START, 4)\}$
 $\{(A, 5)\}$

$$(f(A) = h(A) + g(A) = 3 + g(START) + \text{cost}(START, A) = 3 + 0 + 2)$$

$\{(B, 5) (C, 7)\}$

$$(f(C) = h(C) + g(C) = 1 + g(A) + \text{cost}(A, C))$$

A lower value of $f(C)$ is found with backpointers {B, A, START}

$\{(C, 5)\}$


$$(f(C) = h(C) + g(C) = 1 + g(B) + \text{cost}(B, C) = 1 + 3 + 1)$$

$\{(GOAL, 6)\}$

A* Core Issues

- Termination condition
- Revisiting states
- Algorithm
- Optimality
- Avoiding revisiting states
- Choosing good heuristics
- Reducing memory usage

A* Termination Condition



Queue:

{(B,4) (A,8)}

{(C,4) (A,8)}


{(D,4) (A,8)}

{(A,8) (G,10)}

We have encountered G before we have a chance to visit the branch going through A. The problem is that at each step we use only an estimate of the path cost to the goal.

- Stop when GOAL is popped from the queue.

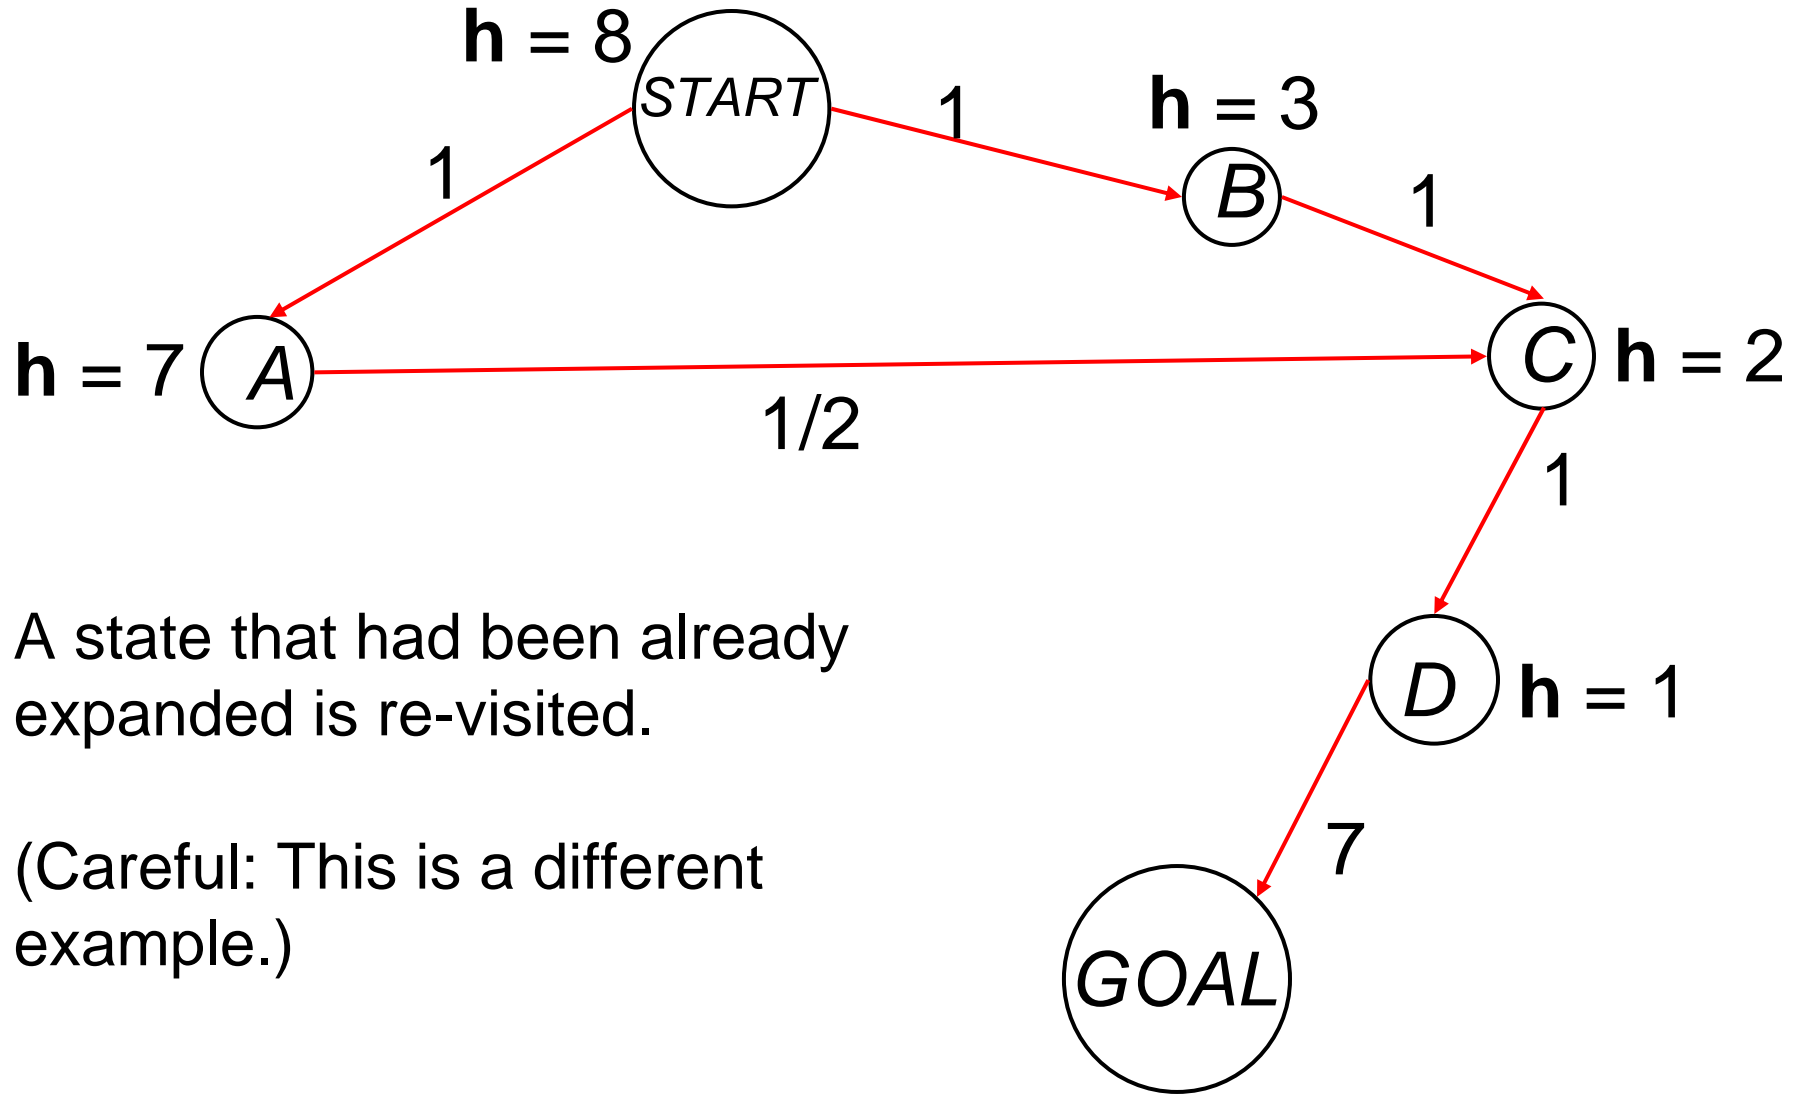
Revisiting States



A state that was already in the queue is re-visited.

How is its priority updated?

Revisiting States



Pop state s with lowest $f(s)$ in queue

If $s = GOAL$

return *SUCCESS*

Else expand s :

For all s' in **succs** (s):

$$f(s') = g(s') + h(s') = g(s) + \mathbf{cost}(s, s') + h(s')$$

If (s' not seen before OR

s' previously expanded with $f(s') > f'$ OR

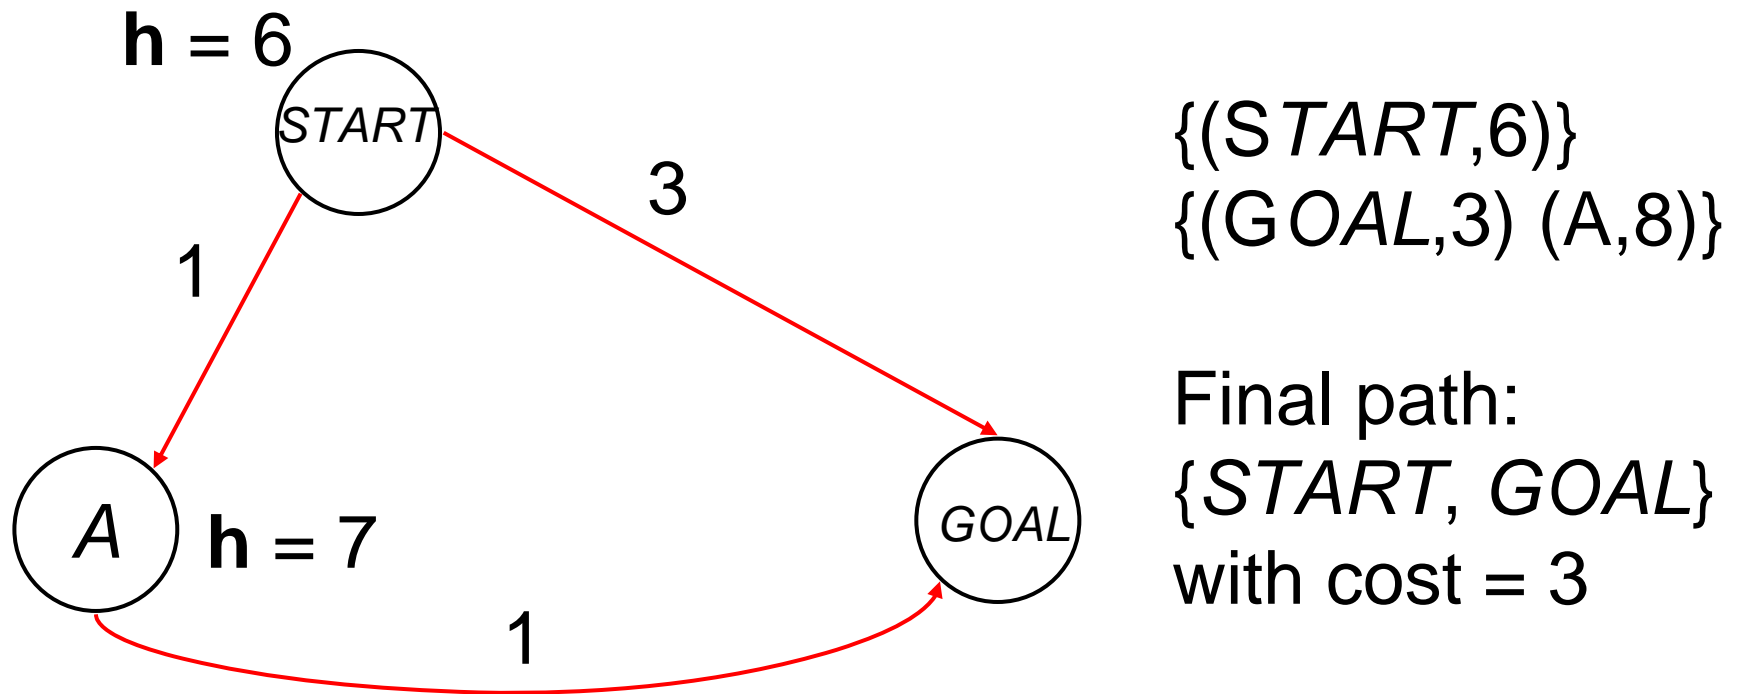
s' in PQ with $f(s') > f'$)

Promote/Insert s' with new value f' in PQ

previous(s') $\leftarrow s$

Else Ignore s' (because it has been visited and its current path cost $f(s')$ is still the lowest path cost from *START* to s')

Under what Conditions is A* Optimal?



- Problem: $h(\cdot)$ is a *poor* estimate of path cost to the goal state

Admissible Heuristics


- Define $h^*(s)$ = the true minimal cost to the goal from s
- h is admissible if

$$h(s) \leq h^*(s) \text{ for all states } s$$

- I.e., an admissible heuristic never overestimates the cost to the goal.
"Optimistic" estimate of cost to goal.


A^* is guaranteed to find the optimal path if h is admissible. (proof in chapter 4)

Examples



For the navigation problem:
The length of the shortest path is at least the distance between s and $GOAL \rightarrow$
Euclidean distance is an admissible heuristic

What about the puzzle?



s



GOAL

Misplaced tiles:

$$h_1(s) = 7$$

Manhattan distance:


$$h_2(s) = 2 + 3 + 3 + 2 + 4 + 2 + 0 + 2 = 18$$

Comparing Heuristics

	$L = 4$ steps	$L = 8$ steps	$L = 12$ steps	
$h_1 =$ misplaced tiles				
	Iterative Deepening	112	6,300	3.6×10^6
$h_2 =$ Manhattan distance				
	A* with heuristic h_1	13	39	227
	A* with heuristic h_2	12	25	73

- Overestimates A* performance because of the tendency of IDS to expand states repeatedly
- Number of states expanded does not include $\log()$ time access to queue

Comparing Heuristics



$$h_1(s) = 7$$

$$h_2(s) = 2 + 3 + 3 + 2 + 4 + 2 + 0 + 2 = 18$$

h_2 is larger than h_1 and, at same time, A^* seems to be more efficient with h_2 .

h_2 dominates h_1 , if $h_2(s) \geq h_1(s)$ for all s

For any two heuristics h_2 and h_1 :

If h_2 dominates h_1 then A^* is more efficient (expands fewer states) with h_2

Intuition: since $h \leq h^*$, a larger h is a better approximation of the true path cost

Limitations

- Computation: In the worst case, we may have to explore all the states
- The good news: A^* is optimally efficient
→ For a given $h(\cdot)$, no other optimal algorithm will expand fewer nodes
- The bad news: Storage is also potentially exponential (all states)

IDS (Iterative Deepening Search)

- Need to make DFS optimal
- IDS (Iterative Deepening Search):
 - Run DFS by searching only path of length 1 (DFS stops if length of path is greater than 1)
 - If that doesn't find a solution, try again by running DFS on paths of length 2 or less
 - If that doesn't find a solution, try again by running DFS on paths of length 3 or less
 -
 - Continue until a solution is found

Example: IDA* (Iterative Deepening A*)

- Same idea as Iterative Deepening DFS except use $f(s)$ to control depth of search instead of the number of transitions
- Example, assuming integer costs:
 1. Run DFS, stopping at states s such that $f(s) > 0$
Stop if goal reached
 2. Run DFS, stopping at states s such that $f(s) > 1$
Stop if goal reached
 3. Run DFS, stopping at states s such that $f(s) > 2$
Stop if goal reached.....Keep going by increasing the limit on f by 1 every time
- Complete (assuming we use loop-avoiding DFS)
- Optimal
- More expensive in computation cost than A*
- Memory order L as in DFS

Summary

- Informed search and heuristics
- First attempt: Best-First Greedy search
- A* algorithm
 - Optimality
 - Condition on heuristic functions
 - Completeness, efficiency
- IDA*

Nils Nilsson. Problem Solving Methods in Artificial Intelligence. McGraw Hill (1971)

Judea Pearl. Heuristics: Intelligent Search Strategies for Computer Problem Solving (1984)

Chapters 3&4 Russell & Norvig