



CS 540 Introduction to Artificial Intelligence

Search II: Informed Search

Yudong Chen
University of Wisconsin-Madison

Nov 16, 2021

Announcements

- HW8 released today, due next Tuesday
- Annotated slides
- Grading Info
- Class roadmap

Today = informed search

Thursday = advanced search.

Next week = Games.

Outline

- Uninformed vs Informed Search
 - Review of uninformed strategies, adding heuristics
- A* Search
 - Heuristic properties, stopping rules, analysis
- Extensions: Beyond A*
 - Iterative deepening, beam search

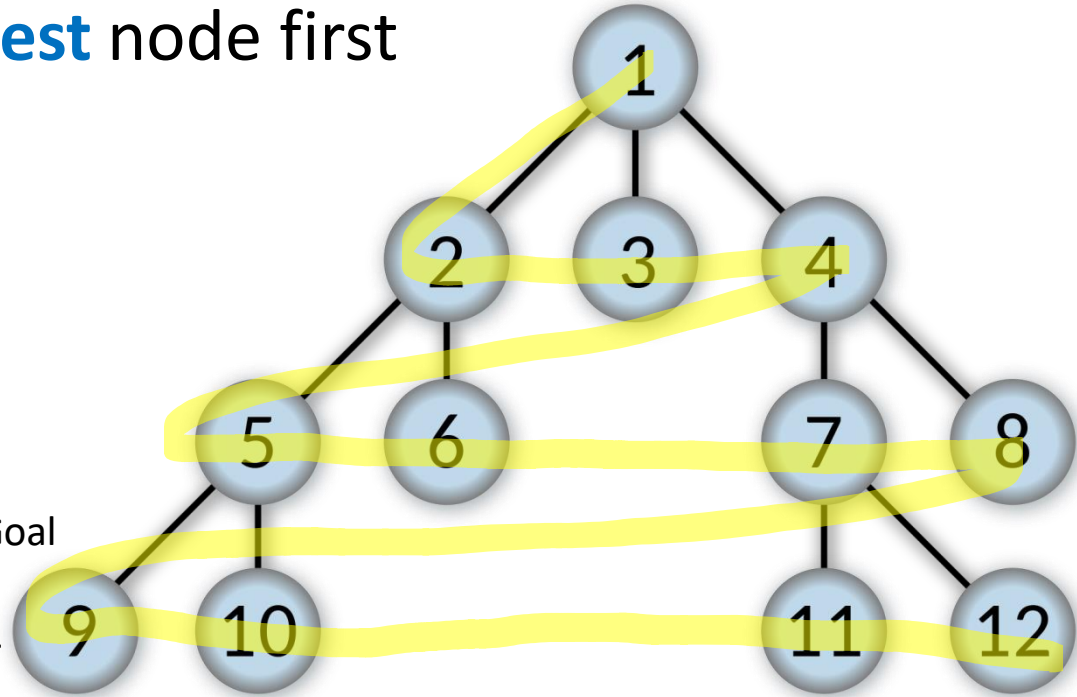
Breadth-First Search

Recall: expand **shallowest** node first

- Data structure: queue
- **Properties:**
 - Complete
 - Optimal (if edge cost 1)
 - Time $O(b^d)$
 - Space $O(b^d)$

Depth of Goal

Branching Factor



Uniform Cost Search

Recall: expand **least-cost** node first

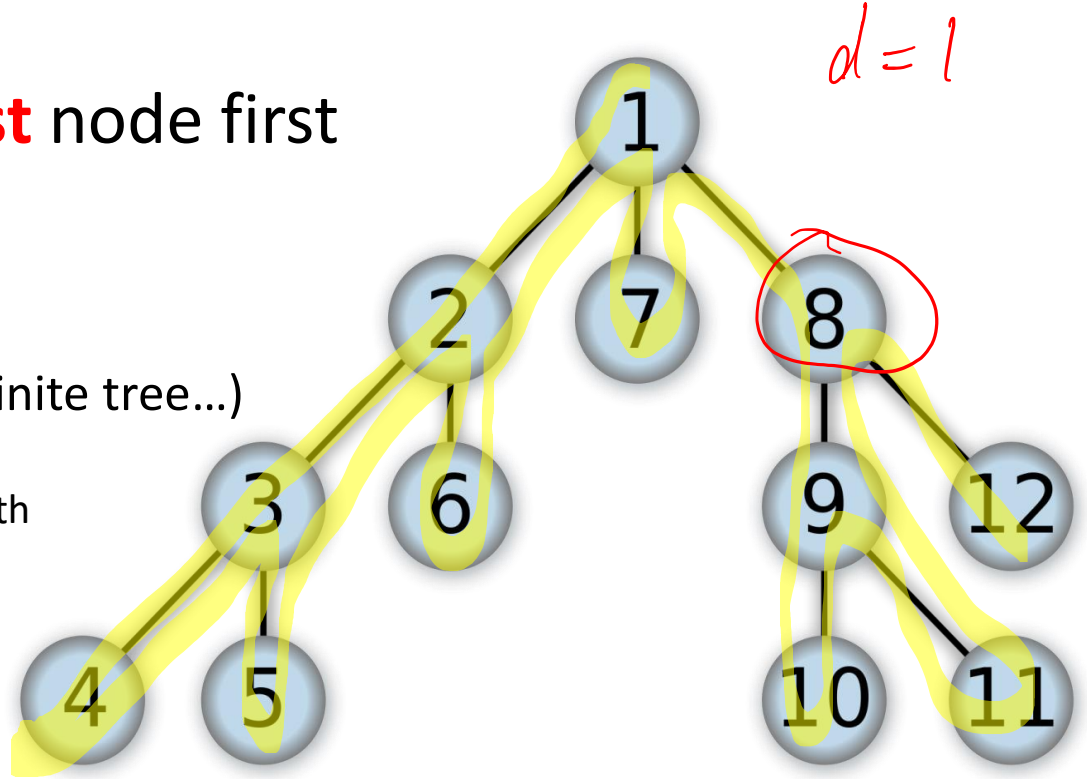
- Generalization of BFS
- Data structure: priority queue
- **Properties:**
 - Complete
 - Optimal (if weight lower bounded by ϵ)
 - Time $O(b^{C^*/\epsilon})$
 - Space $O(b^{C^*/\epsilon})$

← Optimal goal path cost

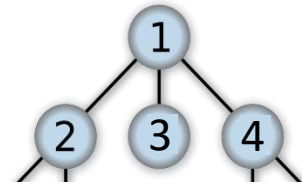
Depth-First Search

Recall: expand **deepest** node first

- Data structure: stack
- **Properties:**
 - Incomplete (stuck in infinite tree...)
 - Suboptimal
 - Time $O(b^m)$ ← Max Depth
b ~~4~~ *3*
 - Space $O(bm)$

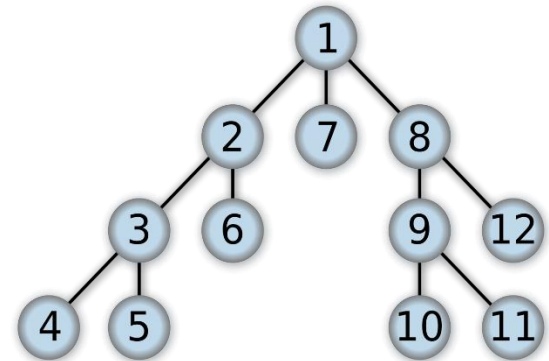
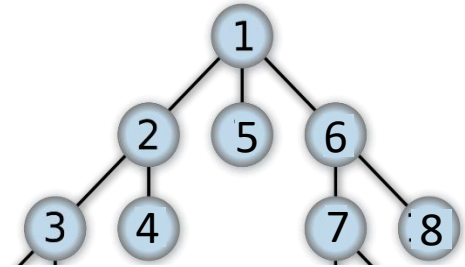


Iterative Deepening DFS



Repeated limited DFS

- Search like BFS, fringe like DFS
- **Properties:**
 - Complete
 - Optimal (if edge cost 1)
 - Time $O(b^d)$
 - Space $O(bd)$



A good option!

Uninformed vs Informed Search

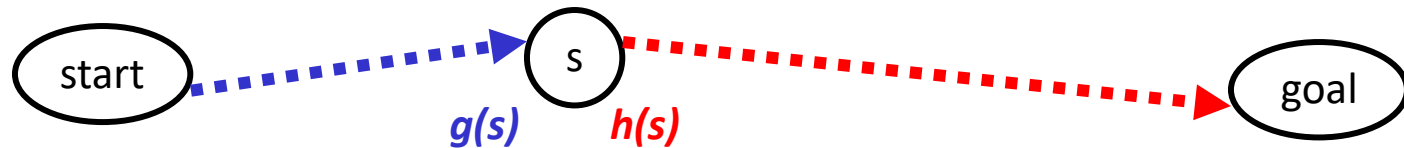
Uninformed search (all of what we saw). Know:

- Path cost $g(s)$ from start to node s
- Successors.



Informed search. Know:

- All uninformed search properties, plus
- Heuristic $h(s)$ from s to goal

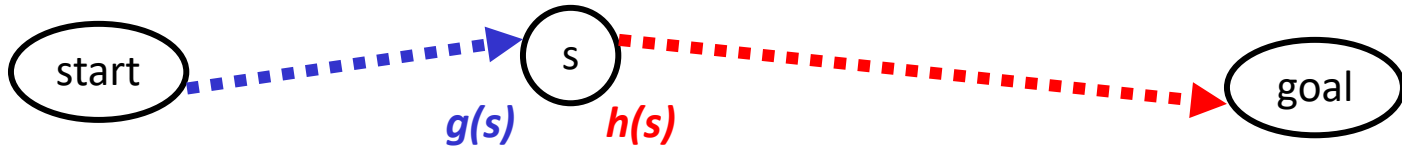


Informed Search

Informed search. Know:

- All uninformed search properties, plus
- Heuristic $h(s)$ from s to goal

hope: $h(s)$ is true cost from s to goal.

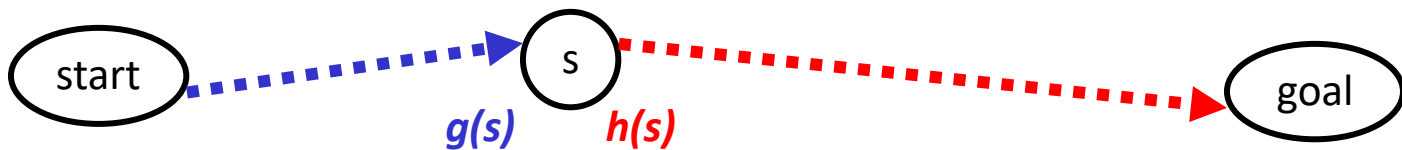


- Use side information to **speed up search**.

Using the Heuristic

Back to uniform-cost search

- We had the priority queue
- Expand the node with the smallest $g(s)$
 - $g(s)$ “first-half-cost”



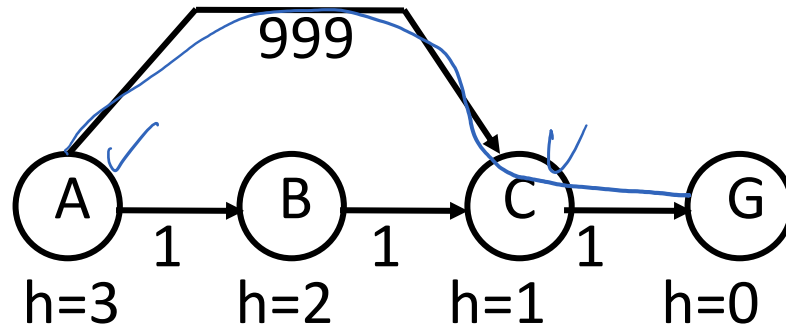
- Now let's use the heuristic (“second-half-cost”)
 - Several possible approaches: let's see what works

Attempt 1: Best-First Greedy

One approach: just use $h(s)$ alone

- Specifically, expand node with smallest $h(s)$
- This isn't a good idea. Why?

B: 2
C: 1



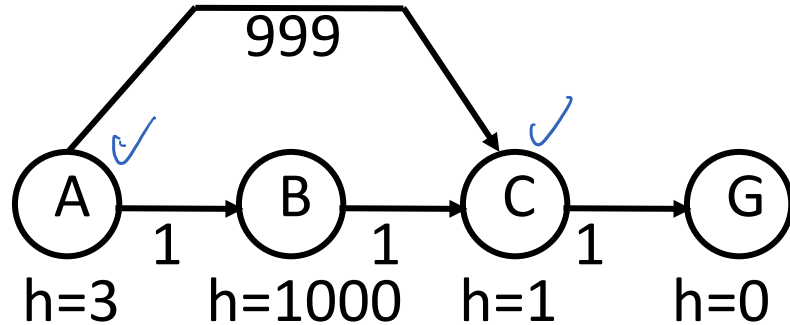
- Not optimal! **Get** $A \rightarrow C \rightarrow G$. **Want:** $A \rightarrow B \rightarrow C \rightarrow G$

Attempt 2: A Search

$g(s) + h(s)$
B 1 + 1000
~~C 999 + 1~~

Next approach: use both $g(s) + h(s)$ alone

- Specifically, expand node with smallest $g(s) + h(s)$
- Again, use a priority queue
- Called “A” search



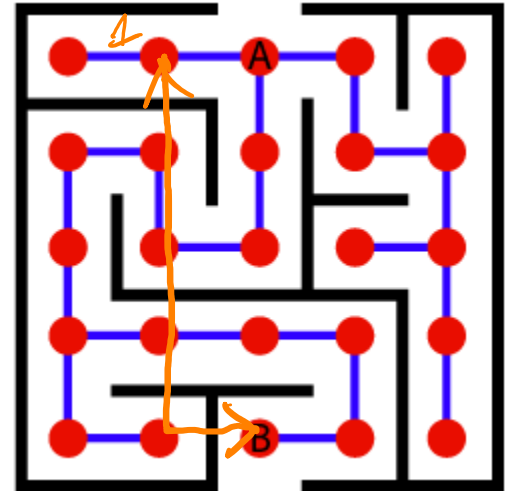
- **Still not optimal!** (Does work for former example).

Attempt 3: A* Search

$$h(s) = 5$$

Same idea, use $g(s) + h(s)$, with one requirement

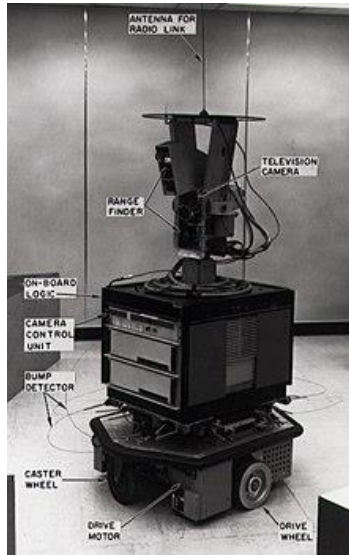
- Demand that $h(s) \leq h^*(s)$
- If heuristic has this property, “admissible”
 - Optimistic! Never over-estimates
- Still need $h(s) \geq 0$
 - Negative heuristics can lead to strange behavior
- This is **A*** search



Attempt 3: A* Search

$0 = \text{fringe}$
 $\text{color} = h(s)$

Origins: robots and planning



Shakey the Robot,
1960's

Credit: Wiki



Animation: finding a path
around obstacle

Credit: Wiki

Admissible Heuristic Functions

Have to be careful to ensure admissibility (**optimism!**)

- Example: **8 Game**

Example
State

1		5
2	6	3
7	4	8

Goal
State

1	2	3
4	5	6
7	8	

- One useful approach: **relax constraints**

– $h(s)$ = number of tiles in wrong position

- allows tiles to fly to destination in a single step

$$h(s) = 5.$$

Break & Quiz

Q 1.1: Consider finding the fastest driving route from one US city to another. Measure cost as the number of hours driven when driving at the speed limit. Let $h(s)$ be the number of hours needed to ride a bike from city s to your destination. $h(s)$ is

- A. An admissible heuristic
- B. Not an admissible heuristic

Break & Quiz

Q 1.1: Consider finding the fastest driving route from one US city to another. Measure cost as the number of hours driven when driving at the speed limit. Let $h(s)$ be the number of hours needed to ride a bike from city s to your destination. $h(s)$ is

- A. An admissible heuristic
- **B. Not an admissible heuristic**

Break & Quiz

Q 1.1: Consider finding the fastest driving route from one US city to another. Measure cost as the number of hours driven when driving at the speed limit. Let $h(s)$ be the number of hours needed to ride a bike from city s to your destination. $h(s)$ is

- A. An admissible heuristic **No: riding your bike take longer.**
- **B. Not an admissible heuristic**

Break & Quiz

Q 1.2: Which of the following are admissible heuristics?

- ✓ (i) $h(s) = h^*(s)$
- ✗ (ii) $h(s) = \max(2, h^*(s))$
- ✓ (iii) $h(s) = \min(2, h^*(s)) \leq h^*(s)$
- ✗ (iv) $h(s) = h^*(s) - 2$ *may be negative*
- (v) $h(s) = \text{sqrt}(h^*(s))$

- A. All of the above
- B. (i), (iii), (iv)
- C. (i), (iii)
- D. (i), (iii), (v)

Break & Quiz

Q 1.2: Which of the following are admissible heuristics?

(i) $h(s) = h^*(s)$

(ii) $h(s) = \max(2, h^*(s))$

(iii) $h(s) = \min(2, h^*(s))$

(iv) $h(s) = h^*(s) - 2$

(v) $h(s) = \text{sqrt}(h^*(s))$

- A. All of the above
- B. (i), (iii), (iv)
- **C. (i), (iii)**
- D. (i), (iii), (v)

Break & Quiz

Q 1.2: Which of the following are admissible heuristics?

(i) $h(s) = h^*(s)$

(ii) $h(s) = \max(2, h^*(s))$ No: $h(s)$ might be too big

(iii) $h(s) = \min(2, h^*(s))$

(iv) $h(s) = h^*(s) - 2$ No: $h(s)$ might be negative

(v) $h(s) = \text{sqrt}(h^*(s))$ No: if $h^*(s) < 1$ then $h(s)$ is bigger

- A. All of the above
- B. (i), (iii), (iv)
- **C. (i), (iii)**
- D. (i), (iii), (v)

Heuristic Function Tradeoffs

Dominance: h_2 dominates h_1 if for all states s ,

$$h_1(s) \leq h_2(s) \leq h^*(s)$$

- **Idea:** we want to be as close to h^* as possible
 - But not over!
- **Tradeoff:** being very close might require a very complex heuristic, expensive computation
 - Might be better off with cheaper heuristic & expand more nodes.

Heuristic Function Tradeoffs

h_2 dominates h_1 .

- Example: **8 Game**

Example State

1	5	5
2	6	3
7	4	8

Goal State

1	2	3
4	5	6
7	8	

- Previous heuristic: $h_1(s)$ = number of tiles in wrong position
- Better heuristic?

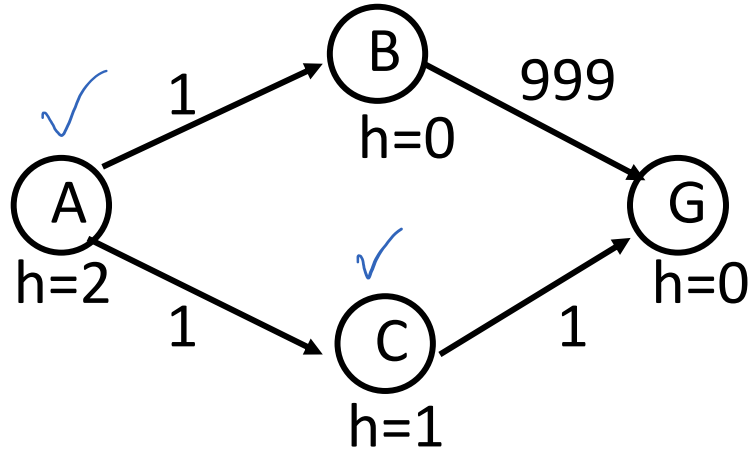
$$h_2(s) = \sum_{\text{tile: wrong}} \text{Manhattan distance (tile, destination)}$$

A* Termination

$g(s) + h(s)$
~~B | 1 + 0~~
~~C | 1 + 1~~
A | 1000 + 0

When should A* **stop**?

- One idea: as soon as we reach goal state?

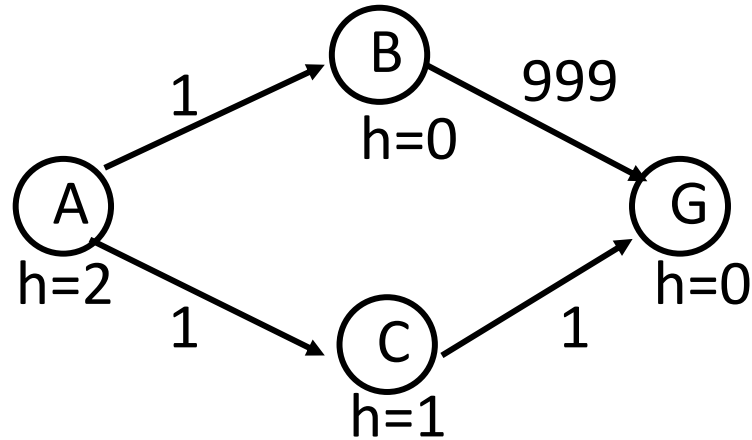


- **h** admissible, but note that we get $A \rightarrow B \rightarrow G$ (**cost 1000**)!

A* Termination

When should A* stop?

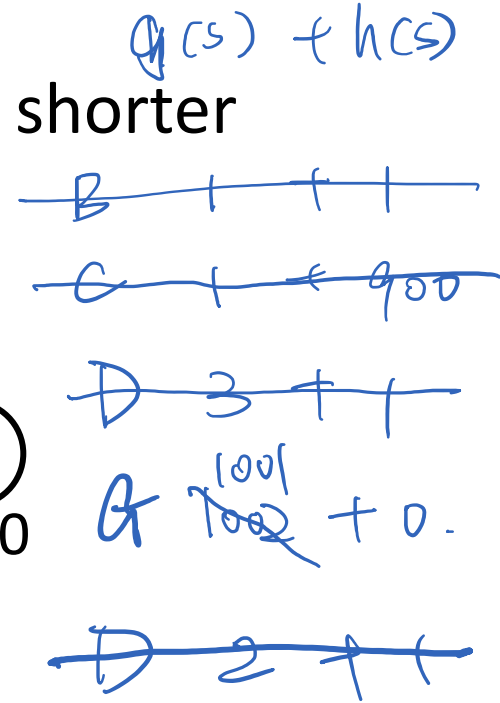
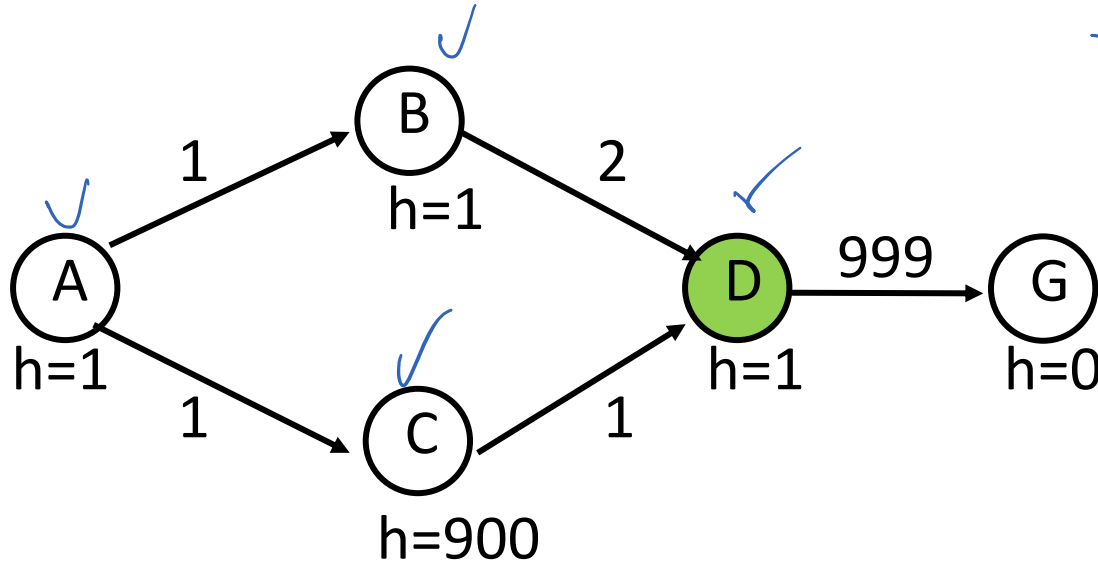
- **Rule:** terminate **when a goal is popped** from queue.



- Note: taking $h = 0$ reduces to uniform cost search rule.

A* Revisiting Expanded States

Possible to revisit an expanded state, get a shorter path:



- Put D back into priority queue, smaller $g+h$

A* Full Algorithm

1. Put the start node S on the priority queue, called $OPEN$
2. If $OPEN$ is empty, exit with failure
3. Remove from $OPEN$ and place on $CLOSED$ a node n for which $f(n)$ is minimum (note that $f(n)=g(n)+h(n)$)
4. If n is a goal node, exit (trace back pointers from n to S)
5. Expand n , generating all successors and attach to pointers back to n . For each successor n' of n
 1. If n' is not already on $OPEN$ or $CLOSED$ estimate $h(n')$, $g(n')=g(n)+c(n,n')$, $f(n')=g(n')+h(n')$, and place it on $OPEN$.
 2. If n' is already on $OPEN$ or $CLOSED$, then check if $g(n')$ is lower for the new version of n' . If so, then:
 1. Redirect pointers backward from n' along path yielding lower $g(n')$.
 2. Put n' on $OPEN$.
 3. If $g(n')$ is not lower for the new version, do nothing.
6. Goto 2.

A* Analysis

Some properties:

- Terminates!
- A* can use **lots of memory**: $O(\# \text{ states})$.
- Will run out on large problems.
- Next, we will consider some alternatives to deal with this.

The image shows a terminal window with two parts. The top part displays the progress of an A* search algorithm over 24 iterations. Each iteration is represented by a row of progress bars for 'Solved' (green) and 'Pruned' (red) states, along with numerical counts for 'Solved', 'Pruned', and 'Total' states. The bottom part shows system resource usage with a table of running processes and a summary of system statistics.

```
ubuntu@ip-172-31-46-218: ~
File Edit View Search Terminal Help

1 [||||| 93.5%] 25 [||||| 89.6%] 49 [||||| 100.0%] 73 [||||| 100.0%]
2 [||||| 89.6%] 26 [||||| 87.7%] 50 [||||| 100.0%] 74 [||||| 100.0%]
3 [||||| 66.4%] 27 [||||| 89.5%] 51 [||||| 49.0%] 75 [||||| 100.0%]
4 [||||| 70.4%] 28 [||||| 89.5%] 52 [||||| 100.0%] 76 [||||| 100.0%]
5 [||||| 93.5%] 29 [||||| 86.3%] 53 [||||| 98.2%] 77 [||||| 100.0%]
6 [||||| 83.0%] 30 [||||| 94.1%] 54 [||||| 100.0%] 78 [||||| 100.0%]
7 [||||| 76.4%] 31 [||||| 90.2%] 55 [||||| 77.2%] 79 [||||| 100.0%]
8 [||||| 69.7%] 32 [||||| 83.0%] 56 [||||| 100.0%] 80 [||||| 100.0%]
9 [||||| 62.7%] 33 [||||| 79.1%] 57 [||||| 100.0%] 81 [||||| 100.0%]
10 [||||| 86.3%] 34 [||||| 86.9%] 58 [||||| 89.9%] 82 [||||| 100.0%]
11 [||||| 86.9%] 35 [||||| 86.9%] 59 [||||| 100.0%] 83 [||||| 100.0%]
12 [||||| 62.3%] 36 [||||| 83.7%] 60 [||||| 81.2%] 84 [||||| 100.0%]
13 [||||| 84.2%] 37 [||||| 81.9%] 61 [||||| 80.3%] 85 [||||| 100.0%]
14 [||||| 70.4%] 38 [||||| 76.0%] 62 [||||| 100.0%] 86 [||||| 100.0%]
15 [||||| 70.6%] 39 [||||| 79.7%] 63 [||||| 96.7%] 87 [||||| 92.2%]
16 [||||| 85.8%] 40 [||||| 77.9%] 64 [||||| 100.0%] 88 [||||| 100.0%]
17 [||||| 85.7%] 41 [||||| 80.5%] 65 [||||| 100.0%] 89 [||||| 100.0%]
18 [||||| 65.8%] 42 [||||| 85.1%] 66 [||||| 100.0%] 90 [||||| 100.0%]
19 [||||| 80.5%] 43 [||||| 81.7%] 67 [||||| 100.0%] 91 [||||| 100.0%]
20 [||||| 76.6%] 44 [||||| 90.1%] 68 [||||| 100.0%] 92 [||||| 100.0%]
21 [||||| 85.7%] 45 [||||| 75.8%] 69 [||||| 100.0%] 93 [||||| 100.0%]
22 [||||| 83.8%] 46 [||||| 83.0%] 70 [||||| 83.8%] 94 [||||| 100.0%]
23 [||||| 88.9%] 47 [||||| 65.1%] 71 [||||| 100.0%] 95 [||||| 29.8%]
24 [||||| 84.3%] 48 [||||| 86.9%] 72 [||||| 96.8%] 96 [||||| 100.0%]
Mem [||||| 177G/376G] Tasks: 332, 49 thr, 749 kthr; 89 running
Swp [||||| 0K/0K] Load average: 157.55 143.11 124.84
Uptime: 13:52:44

PID USER      PRI  NI  VIRT   RES   SHR  S  CPU%  MEM%   IN#E#  Command
3231 ubuntu    20   0 2384M 1925M 5820  R 102.0  0.5  10h48:36 /usr/lib/R/bin/exec/R
3211 ubuntu    20   0 2189M 1664M 5800  R 102.0  0.4  11h21:04 /usr/lib/R/bin/exec/R
3232 ubuntu    20   0 2384M 1925M 5820  R 102.0  0.5  10h38:29 /usr/lib/R/bin/exec/R
3176 ubuntu    20   0 2792M 2252M 5788  R 102.0  0.6  11h12:51 /usr/lib/R/bin/exec/R
3179 ubuntu    20   0 2384M 1925M 5800  R 102.0  0.5  11h37:07 /usr/lib/R/bin/exec/R
3154 ubuntu    20   0 1878M 1361M 5800  R 102.0  0.4  11h19:01 /usr/lib/R/bin/exec/R
3146 ubuntu    20   0 2588M 1943M 5788  R 102.0  0.5  11h18:23 /usr/lib/R/bin/exec/R
3208 ubuntu    20   0 2402M 1644M 5788  R 102.0  0.4  11h03:11 /usr/lib/R/bin/exec/R
3148 ubuntu    20   0 2922M 2367M 5788  R 102.0  0.6  11h38:07 /usr/lib/R/bin/exec/R
3230 ubuntu    20   0 1980M 1522M 5788  R 102.0  0.4  10h48:57 /usr/lib/R/bin/exec/R
3150 ubuntu    20   0 2588M 1985M 5788  R 102.0  0.5  11h42:22 /usr/lib/R/bin/exec/R
3207 ubuntu    20   0 2615M 2157M 5788  R 102.0  0.6  11h13:52 /usr/lib/R/bin/exec/R
3200 ubuntu    20   0 2996M 2337M 5788  R 101.0  0.6  11h46:18 /usr/lib/R/bin/exec/R
3157 ubuntu    20   0 2792M 2263M 5800  R 101.0  0.6  11h52:37 /usr/lib/R/bin/exec/R
3152 ubuntu    20   0 2591M 2068M 5788  R 101.0  0.5  11h37:38 /usr/lib/R/bin/exec/R
F1Help F2Setup F3Search F4Filter F5Tree F6SortBy F7Nice F8NICE F9Kill F10Quit
```

Break & Quiz

Q 2.1: Consider two heuristics for the 8 puzzle problem. h_1 is the number of tiles in wrong position. h_2 is the l_1 /Manhattan distance between the tiles and the goal location. How do h_1 and h_2 relate?

- A. h_2 dominates h_1
- B. h_1 dominates h_2
- C. Neither dominates the other

Break & Quiz

Q 2.1: Consider two heuristics for the 8 puzzle problem. h_1 is the number of tiles in wrong position. h_2 is the l_1 /Manhattan distance between the tiles and the goal location. How do h_1 and h_2 relate?

- **A. h_2 dominates h_1**
- B. h_1 dominates h_2
- C. Neither dominates the other

Break & Quiz

Q 2.1: Consider two heuristics for the 8 puzzle problem. h_1 is the number of tiles in wrong position. h_2 is the l_1 /Manhattan distance between the tiles and the goal location. How do h_1 and h_2 relate?

- **A. h_2 dominates h_1**
- **B. h_1 dominates h_2 (No: h_1 is a distance where each entry is at most 1, h_2 can be greater)**
- **C. Neither dominates the other**

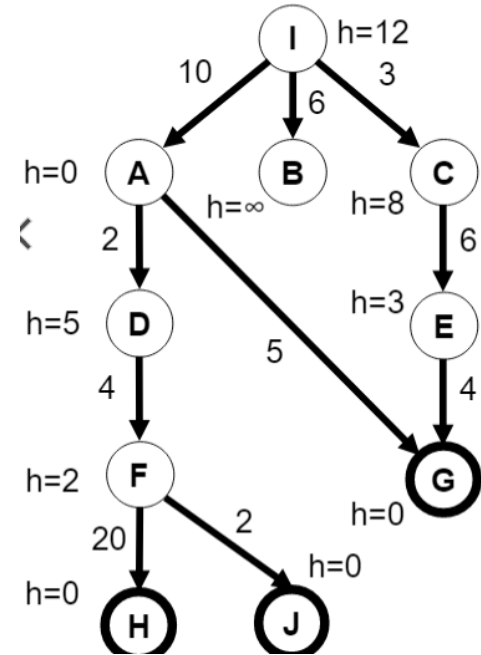
Break & Quiz

Q 2.2: Consider the state space graph below. Goal states have bold borders. $h(s)$ is show next to each node. What node will be expanded by A* after the initial state I?

- A. A
- B. B
- C. C

$g(s) + h(s)$

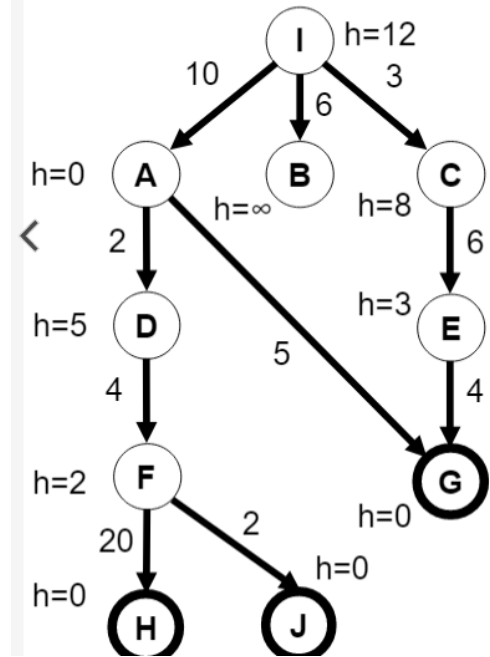
A	10	+	0
B	6	+	∞
C	3	+	8 8



Break & Quiz

Q 2.2: Consider the state space graph below. Goal states have bold borders. $h(s)$ is show next to each node. What node will be expanded by A* after the initial state I?

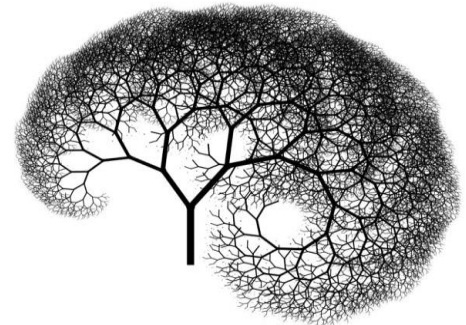
- **A. A**
- B. B
- C. C



IDA*: Iterative Deepening A*

Similar idea to our earlier iterative deepening.

- Bound the memory in search.
- At each phase, don't expand any node with $g(s) + h(s) > k$,
 - Assuming integer costs, do this for $k=0$, then $k=1$, then $k=2$, and so on
- Complete + optimal, might be costly time-wise
 - Revisit many nodes
- Lower memory use than A*



IDA*: Properties

How many restarts do we expect?

- With integer costs, optimal solution C^* , at most C^*

What about non-integer costs?

- Initial threshold k . Use the same rule for non-expansion
- Set new k to be the $\min g(s) + h(s)$ for non-expanded nodes
- Worst case: restarted for each state

Beam Search

General approach (beyond A* too)

- Priority queue with fixed size k ; beyond k nodes, **discard!**
- **Upside**: good memory efficiency
- **Downside**: not complete or optimal

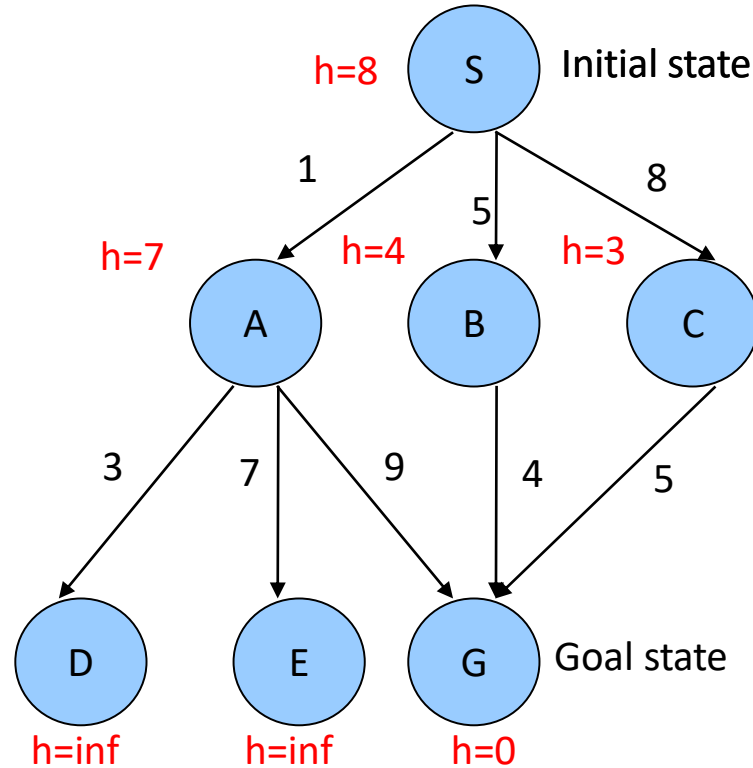
Variation:

- Priority queue with nodes that **are at most ϵ worse** than best node.



Recap and Examples

Example for A*:

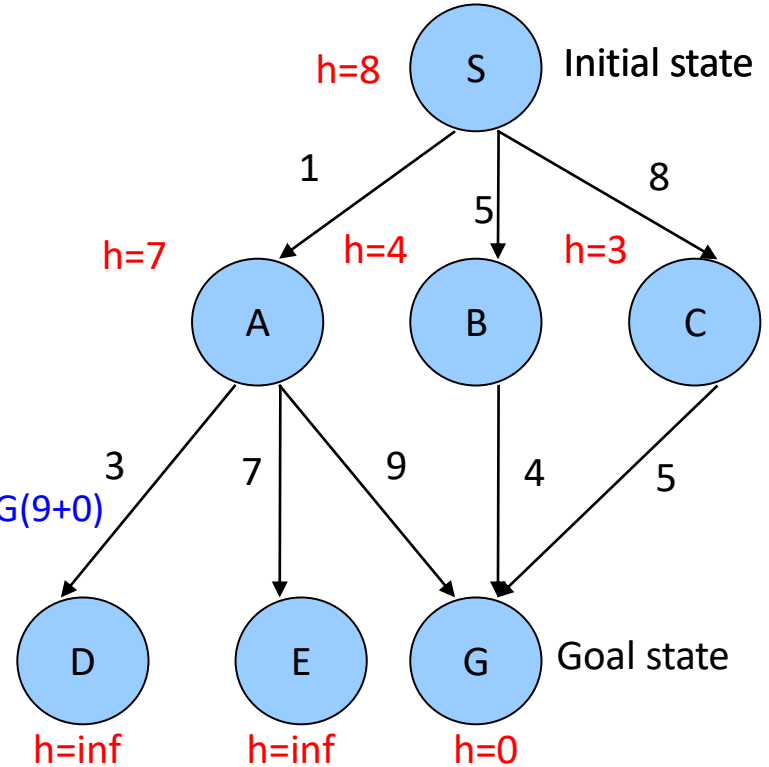


Recap and Examples

Example for A*:

OPEN	CLOSED
S(0+8)	-
A(1+7) B(5+4) C(8+3)	S(0+8)
B(5+4) C(8+3) D(4+inf) E(8+inf) G(10+0)	S(0+8) A(1+7)
C(8+3) D(4+inf) E(8+inf) G(9+0)	S(0+8) A(1+7) B(5+4)
C(8+3) D(4+inf) E(8+inf)	S(0+8) A(1+7) B(5+4) G(9+0)

G → B → S



Recap and Examples

Example for IDA*:

Threshold = 8

PREFIX

OPEN

-

S(0+8)

S

A(1+7)

SA

H(2+2) D(4+4)

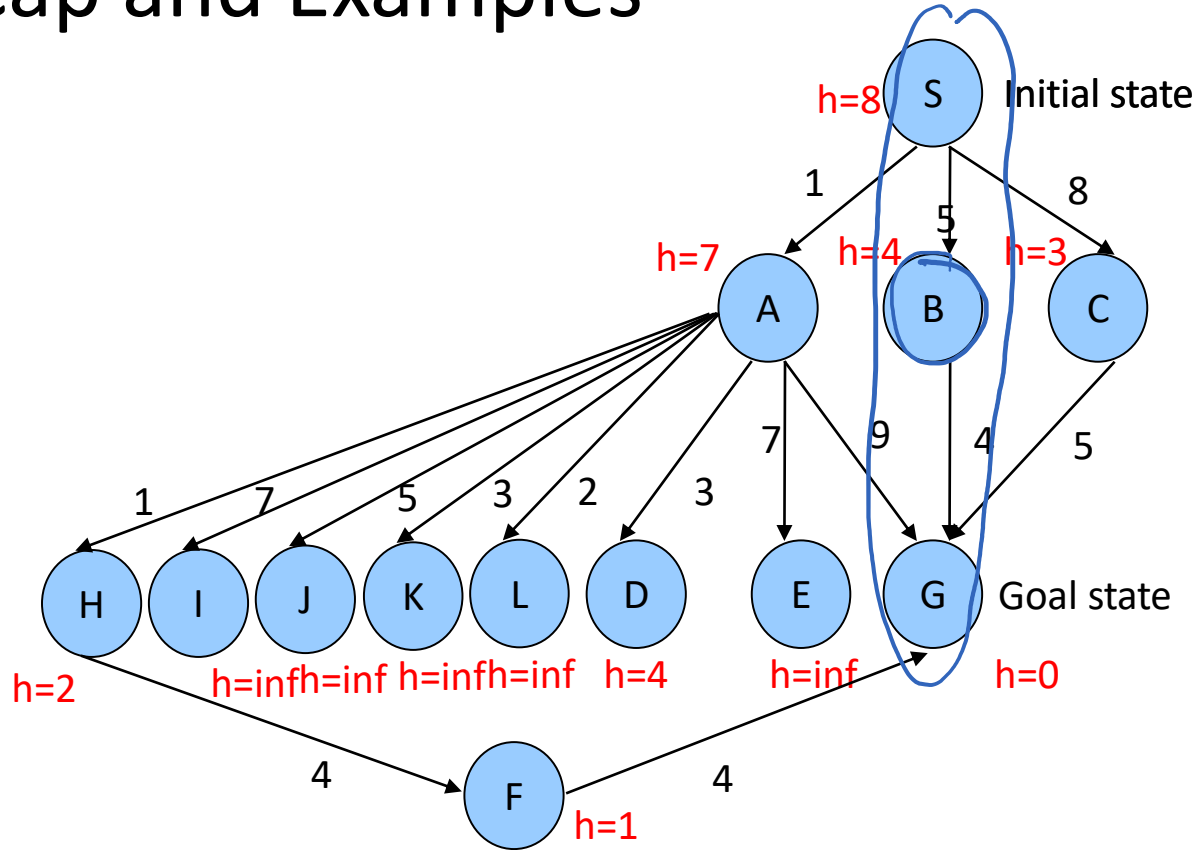
SAH

D(4+4) F(6+1)

SAHF

D(4+4)

SAD

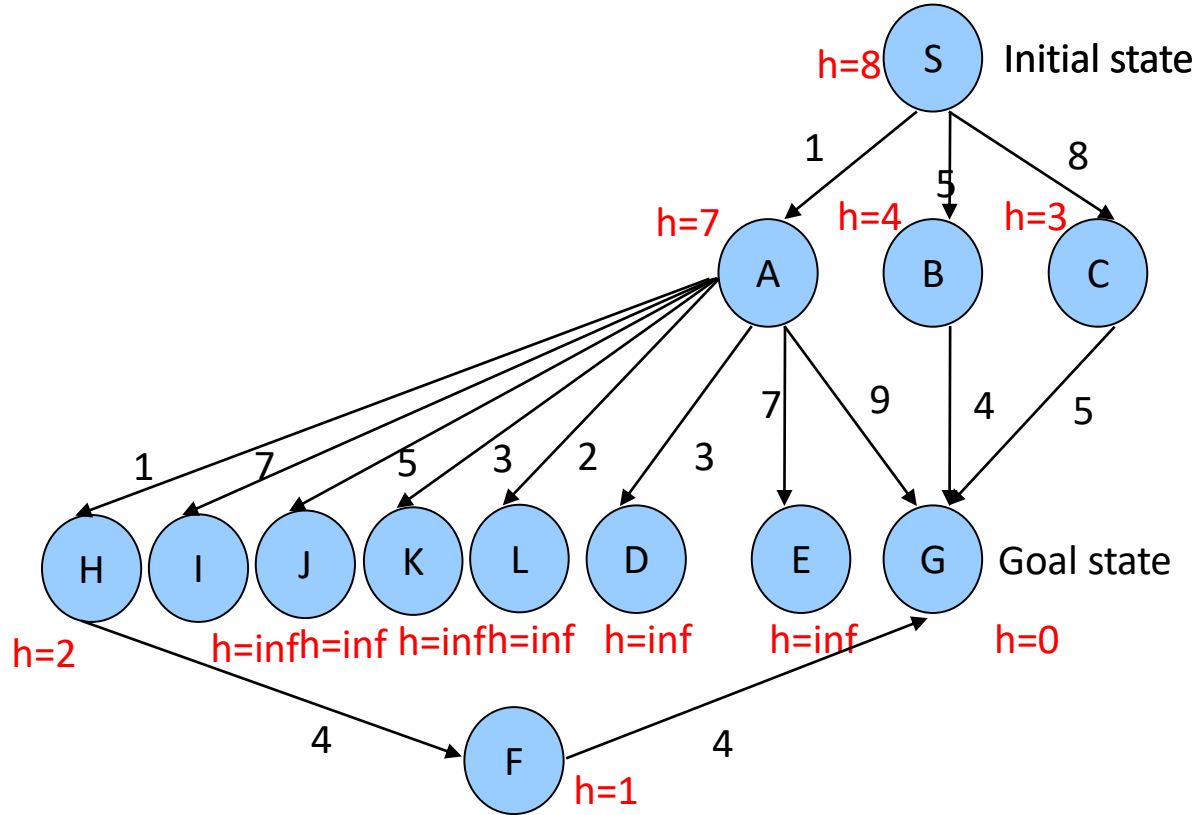


Recap and Examples

Example for IDA*:

Threshold = 9

PREFIX	OPEN
-	S(0+8)
S	A(1+7) B(5+4)
SA	B(5+4) H(2+2) D(4+4)
SAH	B(5+4) D(4+4) F(6+1)
SAHF	B(5+4) D(4+4)
SAD	B(5+4)
SB	G(9+0)
SBG	



Recap and Examples

Example for Beam Search: $k=2$

CURRENT

-

S

A

H

F

D

G

OPEN

S(0+8)

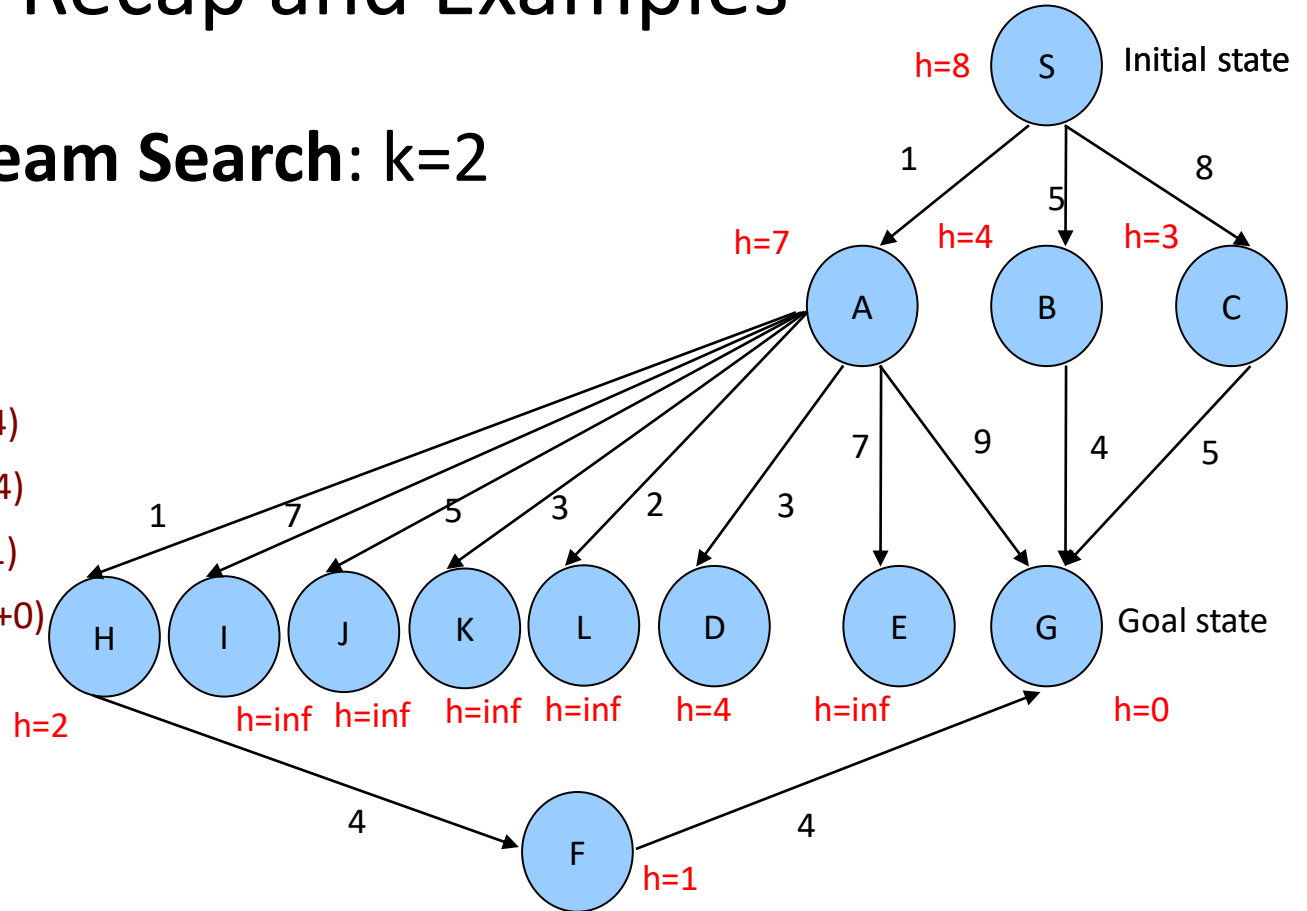
A(1+7) B(5+4)

H(2+2) D(4+4)

D(4+4) F(6+1)

D(4+4) G(10+0)

G(10+0)



Summary

- Informed search: introduce heuristics
 - Not all approaches work: best-first greedy is bad
- A* algorithm
 - Properties of A*, idea of admissible heuristics
- Beyond A*
 - IDA*, beam search. Ways to deal with space requirements.



Acknowledgements: Adapted from materials by Jerry Zhu and Fred Sala (University of Wisconsin-Madison).