Name: ____ CSE 473, Spring 2012

Please answer clearly and succinctly. Show your work clearly for full credit. If an explanation is requested, think carefully before writing. Points will be removed for rambling answers with irrelevant information (and may be removed in cases of messy and hard to read answers). If a question is unclear or ambiguous, feel free to make the additional assumptions necessary to produce the answer. State these assumptions clearly; you will be graded on the basis of the assumption as well as subsequent reasoning.

There are 10+0 problems worth 78 points on 7 pages

Problem 0 (1 point) Write your name on the top of each page.

Problem 1 (11 points) Suppose that an agent is in a 3x3 maze environment similar to the ones shown in the illustration below. The agent knows the size of the maze, that the initial state is (1,1) in the lower left and that the goal is (3,3). There are four actions, **Left**, **Right**, **Up** and **Down**, which have their usual effects, except when blocked by a wall or by the maze boundaries; in this case, the agent does not move. The agent does ***not*** (initially) know if there are any walls present in the maze nor where such walls may be. After each action, the agent receives a single percept, which tells it if it has failed to move; by using this percept, the agent may deduce the presence of walls.

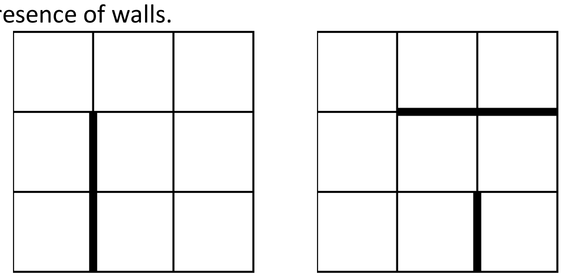

a) (5 points) Describe the problem of finding a path from the initial state to the goal as a search problem in belief state space.

b) (3 points) How large is the *initial* belief state?

c) (3 points) Approximately how many distinct belief states are there? (Some of the possible belief states may not be reachable from the initial state, but include them anyway).

Problem 2 (10 points) The figure below shows a problem-space graph, where A is the initial state and G denotes the goal. Edges are labeled with their true cost. We have a heuristic function, f(), written in the standard form: f(n)=g(n)+h(n) where g(n) is the cost to get from A to n and h(n) is an estimate of the remaining distance to G.

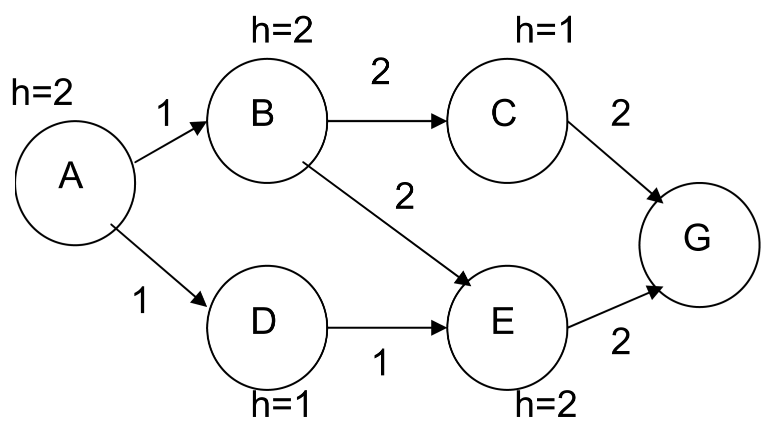

a) (2 points) In the graph above, is f() admissible? Why or why not?

b) (2 points) Is f() monotonic? Why or why not?

c) (6 points) Suppose we use IDA* to search the graph and that states having the same f values are visited in alphabetical order. In what order does the algorithm consider f-limits and visit states? Fill out the table below:

f-limit	**States visited (in order)**
2	A

Problem 3 (8 points) Label the intermediate nodes in the search tree, below, with values and draw a line across edges to denote any pruning done by alpha-beta search. Which action should the agent choose?

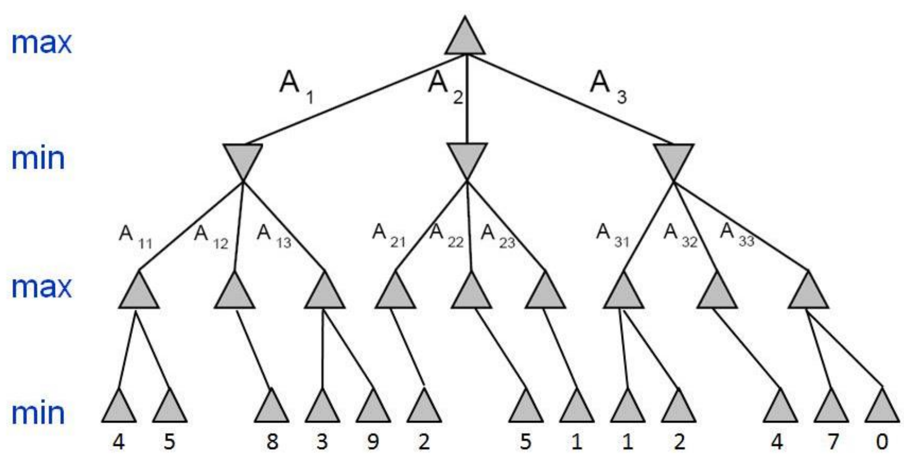

Problem 8 (13 points) Suppose an agent inhabits a world with two states, S and $\neg$S, can do exactly one of two actions, a and b. Action a does nothing and action b flips from one state to the other. Suppose that S is initially true.

b) (6 points) Consider this world as an MDP (with deterministic actions) and let R(S)=3, R($\neg$S)=2 and γ=0.5. Complete the columns of the following table. (Note we are assuming that reward is a function of the destination state and is independent of the starting state or action executed.)

Time	Q(S,a)	Q(S,b)	V(S)	Q($\neg$S,a)	Q($\neg$S,b)	V($\neg$S)
0			0			0
1						
2						
3						

c) (1 point) Based on the V3 values you have computed, what is the policy an agent should choose?