4511W, Spring-2019
ASSIGNMENT 3 :
Assigned: 03/11/19 Due: 03/24/19 at 11:55 PM (submit via Canvas, you may take a picture of handwritten solutions, but you must put them in a pdf) Submit only pdf or txt files

Written/drawn:

Problem 1. (10 points)
Run minimax on the following tree, then show which action the root should take (and state the expected outcome). (Show work for minimax computation.)

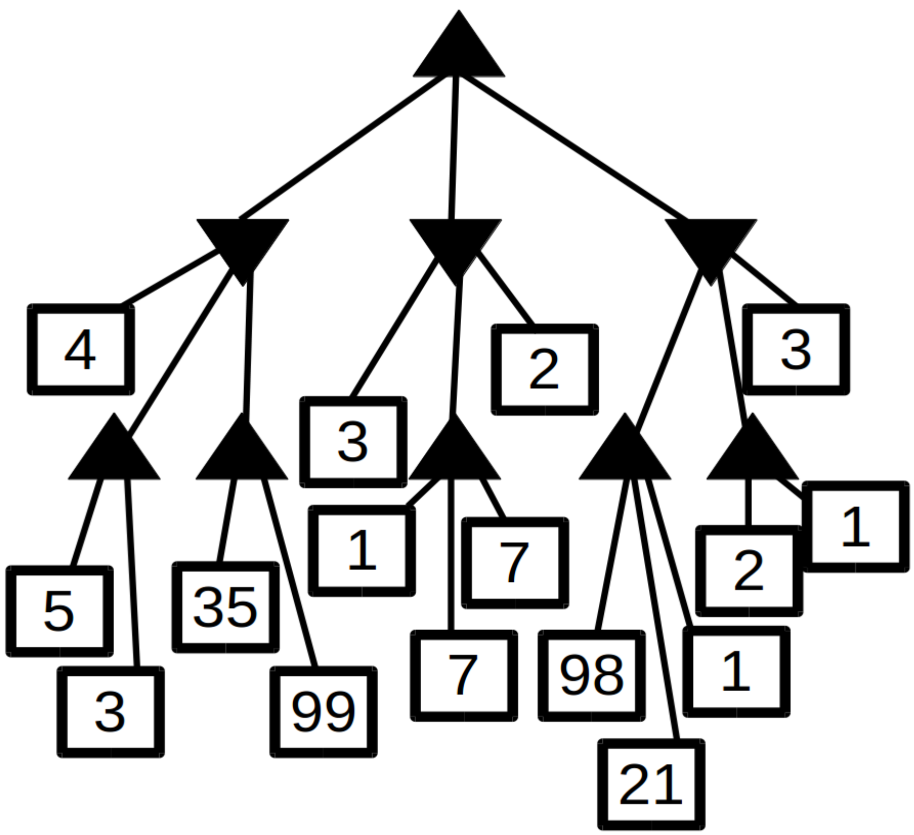

Problem 2. (15 points)
(1) What is/are the pure strategy Nash equilibrium of the game shown below? (Show work.)
(2) What is/are the pure strategy Pareto optimal point(s)? (Show work.)
(3) What single number could you change to simultaneously change both the set of Nash equilibrium and set of Pareto optimal point(s)? (Note: this is one number, not one cell (a pair of numbers).)

4, 2	4, 5	1, 11
3, 0	7, 8	3, 9
0, 3	8, 2	5, 5

Problem 3. (15 points)
Suppose you were running Monte-Carlo tree search on a binary tree (branching factor=2 always). Assume the "random rollouts/playouts" resulted in the following sequence: win, loss, win, win, win, loss, loss, win, win. Show this tree after every 3 rollouts/playouts, so do: win, loss, win, (show tree), win, win, loss, (show tree), loss, win, win (show final tree). Show the UCB values of all nodes when showing the tree.

Problem 4. (20 points)

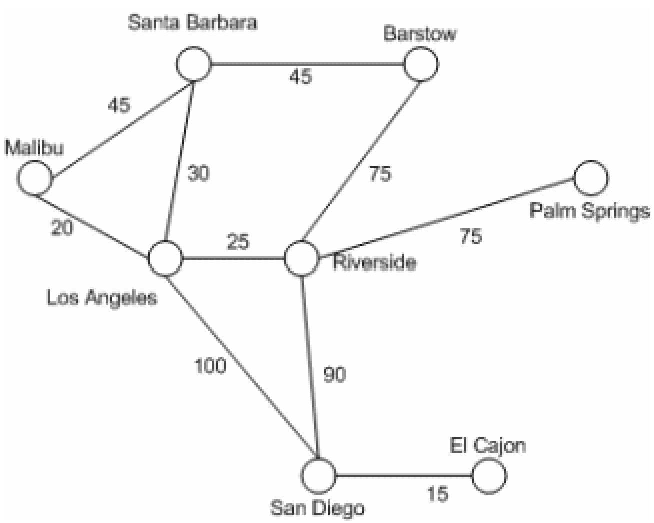

h(city) table:
h(Barstow) = 50
h(El Cajon) = 70
h(Los Angeles) = 90
h(Malibu) = 100
h(Palm Springs) = 0
h(Riverside) = 60
h(San Diego) = 80
h(Santa Barbara) = 80

Given the graph above, assume you are trying to get to Palm Springs (i.e. that is the goal).
(1) Run the basic hill-climbing algorithm from Malibu.
(2) Run stochastic hill-climbing starting from Malibu, where your probability of traveling to a neighbor is defined as: (the equations below have been slightly modified to clarify the corner case when the heuristics are equal to the current state)

Let $total$ be defined as: $total = \sum_{x \in h(x) < h(current)} \frac{1}{h(x)}$

$$P(current \to next) = \begin{cases} 0, \text{if } h(next) > h(current) \\ \frac{1/h(next)}{total}, \text{otherwise} \end{cases}$$

Assume you have the following random numbers: (and you try to use this probability to travel to cities in alphabetical order. So if you have a 30% chance to go to city Ant, 50% to go to Bat and 20% chance to go to Cup... if your random number is 0.42 you would first try to go to "Ant" and fail, then go to "Bat" and succeed (and never even try "Cup").) (Note: more random numbers than you need...)

0.154 0.758 0.724 0.099 0.421 0.303 0.932 0.527 0.457 0.727

(3) Rerun stochastic hill-climbing from Malibu except with the following random numbers:

0.714 0.540 0.991 0.911 0.857 0.340 0.493 0.134 0.741 0.292

(4) Using this stochastic hill-climbing, what is the probability that you find the goal starting at Malibu with the given heuristics?

Problem 5. (30 points)
Assume you are running alpha-beta pruning with a root "max" node (and it always alternates max/min on each depth). "Max" nodes have a branching factor of 2, while "min" nodes have a branching factor of 3.

(1) At what depth are you able to prune more than 75% of the leaf nodes?
(2) At what depth are you able to prune more than 90% of the leaf nodes?
(3) Can you prune more or less on this tree if you start with a "min" node instead? Justify your answer.

Programming (python/lisp):

For this assignment, you will need to know how to use the implemented simulated annealing algorithm. Of note are:
/root/search.py
/root/tests/test_search.py

Problem 6. (10 points)
Assume you have a robot that operates in a grid world and can only travel Up, Down, Left or Right. On each square the robot can see what value it will get (and that of the neighboring four squares). Run simulated annealing to try and end up with the highest value square

Consider the following 4x4 grid world (start in the top-left square):

0	5	10	20
-1	7	8	14
3	4	11	5
23	2	8	25

(1) Run simulated annealing at least 10 times and state where you would end up when the algorithm finishes. How many times did you reach the global maximum? How many times did you reach a local maximum?

(2) Change the bottom left number in the 4x4 grid from a "23" to a "33" and run simulated annealing 10 times. How many times did you reach the global maximum? How many times did you reach a local maximum?

(3) Give your impressions about whether or not the "cooling schedule" for determining where to go in simulated annealing is good or bad.