CS 121, Summer 2009
Homework #5

Out: July 22, 2009
Due: August 3, 2009

How to complete this Homework: Your answers can be typed or carefully hand-written. Please begin each problem on a new page and make sure your name is written on each page of your assignment. Print out this cover sheet and fill in your name and email address, as well as any people you collaborated with. Staple all of your work together and turn in at the beginning of class, August 3, 2009. If you will not be in class you can turn it in prior to class under the door of Gates 132 (with the time of submission written on the homework), or email it to the course staff. Late homeworks will not be accepted.

Your Name: ....

Your email address: ....

Note on Honor Code: You must not look at previously published solutions of any of these problems in preparing your answers. You may discuss these problems with other students in the class (in fact, you are encouraged to do so) and/or look into other documents (books, web sites), with the exception of published solutions, so long as your final submission is prepared on your own without referring to any notes taken during such collaboration. If you have discussed any of the problems with other students, indicate their name(s) here:

....

Any intentional transgression of these rules will be considered an honor code violation.

General Information: Justify your answers, but keep explanations short and to the point. Excessive verbosity will be penalized. If you have any doubt on how to interpret a question, tell us in advance, so that we can help you understand the question, or tell us how you understand it in your returned solution.

Grading:

Problem #	Max. Grade	Your grade
1	20	
2	25	
3	25	
4	10	
5	20	
Total	100	

---

1. **Two-player game [20 points]**

   Consider two players, player 1 and player 2, playing the following **non**-zero-sum game.

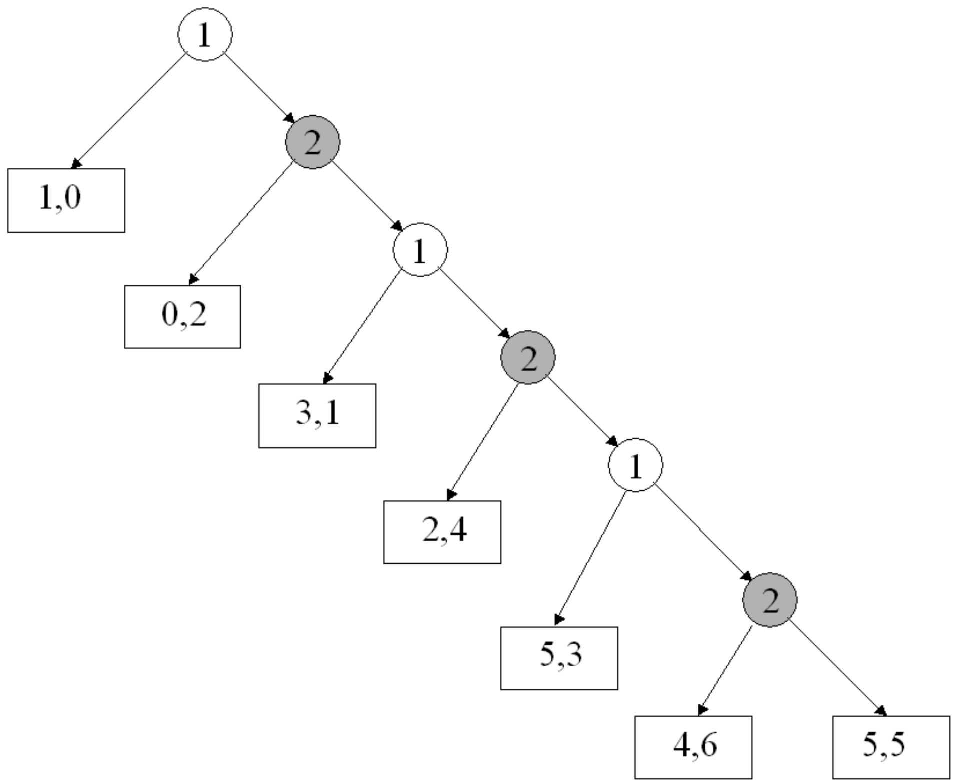

   Each circle corresponds to a decision node for one of the players, where the number in the circle corresponds to the player who gets to make the decision in that node. The rectangles correspond to the payoffs that the players get if the game ends up in that rectangle. The first number in each rectangle is the payoff for player 1, the second is for player 2.

   We assume that each player only cares about the size of their own payoff, and that both players are rational (that is they will each play perfectly), and that they have common knowledge about their rationality (they each know the other is rational, and that the other knows they are rational, etc).

   (a) **[10 points]** What will each player's strategy be? Each strategy should specify what decision each player will make at each of their decision nodes.

   (b) **[5 points]** If both players play according to these strategies, what outcome will be reached?

   (c) **[5 points]** Comment on this outcome. Is it the best that could have been achieved? Do you think you could convince either player to change their strategy to achieve a better outcome?

2. **Alpha-beta pruning [25 points]**

   Consider two players, MAX and MIN, playing a zero-sum, turn-taking game. This game is represented by the following game tree. The scores at the leaf nodes of the tree are from MAX's point of view. So, MAX is trying to get the highest score possible and MIN is trying to get the lowest score possible. In this tree the MAX player gets to move at nodes A, D, E, F, G, H and I, and the MIN player gets to move at nodes B and C.

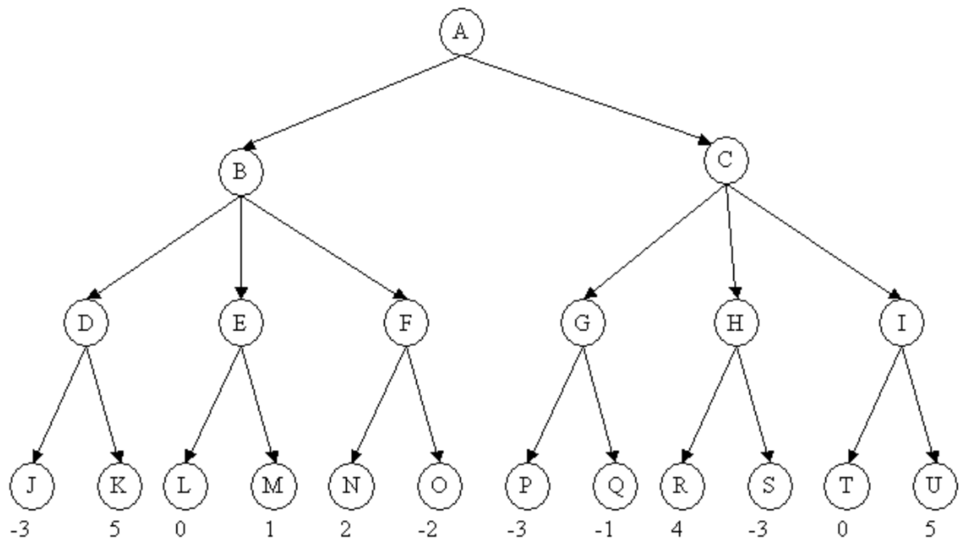

   Play begins at node A, where MAX gets to choose between going to B and C. MIN then gets to choose the next edge taken. This continues until a leaf node is reached, at which point MAX gets the score listed, and MIN gets a score of ($0-$ the score listed). The score at leaf node J is -3, at K it is 5, etc.

   (a) **[10 points]** Suppose MAX is using alpha-beta pruning to determine what move to make. Write down the final α or β value computed at each node by the alpha-beta pruning procedure. [If no value has been computed at a node, report this as well.] What move will MAX make?

   (b) **[5 points]** What nodes would not be examined using the alpha-beta pruning technique? (Assume that at each level the nodes are considered in the order listed above from left to right.)

   (c) **[5 points]** Assume that alpha-beta technique keeps track of the tree path that is responsible for the final α value of the root of the tree. This path connects the root of the tree to some leaf node, which we call x. What is this node (e.g., $x = K$)?

   (d) **[5 points]** Assume that MAX has some extra time and generates a sub-tree (singular extension) from the leaf node x that you have identified in Question 2c, and gets the following sub-tree:

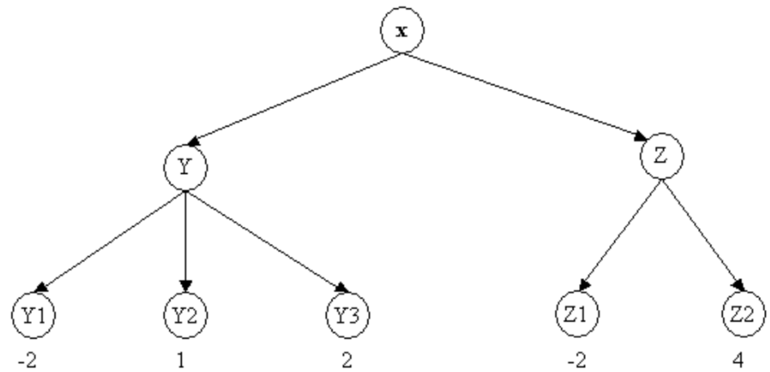

In this tree MIN gets to move at node x, and MAX gets to move at nodes Y and Z. The scores listed are again from MAX's perspective.

- Should MAX change decision at A [Answer by just yes or no]?
- If there is a new leaf node that is responsible for the final α value of the root of the tree, what is this node? [Dont provide any explanation.]

3. **Non-deterministic Uncertainty: Robot Navigation [25 points]**

A robot represented by a point moves in a regular grid placed over a two-dimensional plane. The coordinates of a grid node are (i, j), where i and j are two integers in $\mathcal{Z} = \{\ldots, -2, -1, 0, +1, +2, \ldots\}$. Figure 1 shows this environment.

The state of the robot is its current location (i, j). The belief state of the robot is the set of all locations where the robot thinks it may be with non-zero probability. We assume that the robot is always in one of the states listed in its belief state. So, for example, a belief state can be $\{(-1, 0), (0, 0), (0, +1)\}$, meaning that the robot thinks it may be at $(-1, 0)$ or at $(0, 0)$, or at $(0, +1)$. Its actual location is then any of these three locations.

For this problem, we assume that the robot can only move right by one increment, but robot control is not perfect, so that there is non-deterministic uncertainty in the outcome of this action. The uncertainty model of the action **Right** is the following: From its current location (i, j), the robot may either move to $(i + 1, j)$, $(i + 1, j + 1)$, or $(i + 1, j - 1)$, or stay at (i, j). Each outcome occurs with non-zero probability.

In addition, the robot is equipped with a position sensor that measures its location in the plane. But this sensor is also imperfect. The model of sensing uncertainty is the following: If the robot is at (i, j), then the sensor non-deterministically returns any one of the 5 positions (i, j), $(i - 1, j)$, $(i + 1, j)$, $(i, j - 1)$, and $(i, j + 1)$,

(a) **[5 points]** Let the initial belief state of the robot be $\{(-3, 0)\}$. So, the robot is actually at $(-3, 0)$. Let the robot execute **Right** once and then use its sensor to measure its location. The sensor returns $(-2, -1)$. What is the robots new belief state?

(b) **[10 points]** Assume again that the initial belief state of the robot is $\{(-3, 0)\}$. The robots goal is now to reach any grid node (i, j), such that $i \in \{0, 1, 2, 3\}$ and $-5 \leq j \leq +5$ (hence, the goal nodes form a rectangle in the grid, as shown in Figure 1). To do this, the robot plans to repeatedly execute **Right** until the data (I, J) returned by the position sensor satisfies a certain termination condition $TC(I, J)$; then, the robot

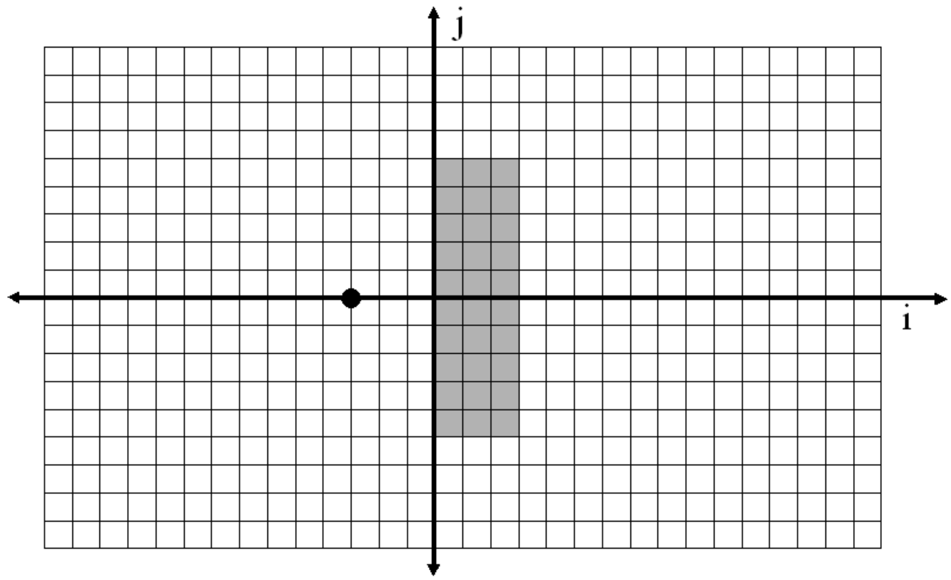

Figure 1: Grid for Problem 3. The region shown gray is the goal region for Questions 3b and 3c. The bold dot at $(-3, 0)$ is the robots initial position in Questions 3a and 3b.

stops. What should $TC(I, J)$ be in order to guarantee that the robot will eventually stop in the goal? Give a brief proof that this condition is guaranteed to make the robot stop in the goal. [Hint: To prove that the robot will stop in the goal, you must prove that: (1) if it did not stop, it would eventually reach the goal; (2) that $TC(I, J)$ cant make the robot stop before reaching the goal; and (3) that $TC(I, J)$ cant let the robot leave the goal.]

(c) **[10 points]** Let now the initial belief state of the robot be $\{(-5, 0)\}$. Is there a condition $TC(I, J)$ that is guaranteed to make the robot stop in the goal? If yes, give this condition. If not, give a possible sequence of states and sensing data that does not allow the robot to safely stop in the goal, no matter what $TC(I, J)$ it is using.

4. **Probabilistic Reasoning [10 points]**

After a recent visit to the doctor, you are given some bad news and some good news. The bad news is that your lab results came back with a positive test for a serious disease, and that the test is 99% accurate (i.e. the probability of testing positive when you do have the disease is 0.99, as is the probability of testing negative when you don't have the disease). The good news is that this is a rare disease, striking only 1 in 10,000 people of your age. Why is it good news that the disease is rare? What is the probability that you actually have the disease?

5. **Bayesian Networks [20 points]**

   You are given the following Bayesian network

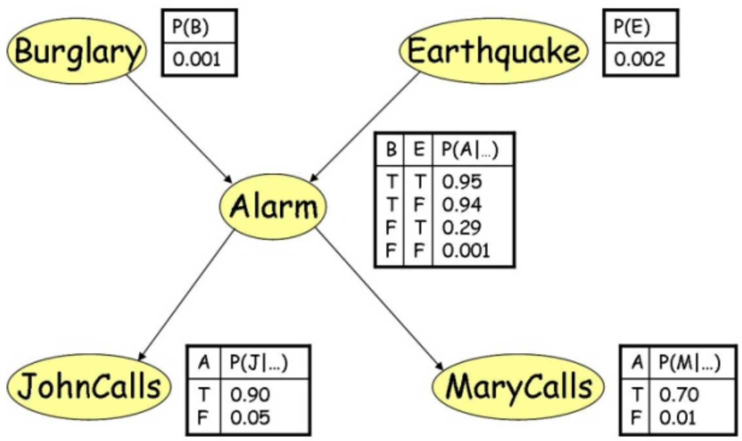

   (a) **[10 points]** Before you make any observation, how would you compute the prior probability of $MaryCalls$? [We don't ask you to do the numerical computation. We ask you to express the probability of $MaryCalls$, $P(M)$, using quantities that are given in the Bayes' net. To simplify notations, replace $MaryCalls$ by M, *Earthquake* by E, etc and let $P(M)$, $P(\neg M)$, $P(E)$, $\ldots$ denote the probability of M, not M, E, etc $\ldots$]

   (b) **[10 points]** Now, assume that you observe that $JohnCalls$ is True. How would you compute the new (posterior) probability of $MaryCalls$? [Again, we don't ask you for numerical values. Explain how you would do the computation. In this explanation, you will need to use the prior probability of $JohnCalls$, $P(J)$. No need to explain how you compute it, since its computation is very similar to that of $P(M)$ in Question 5a]