CS472 Sample Midterm

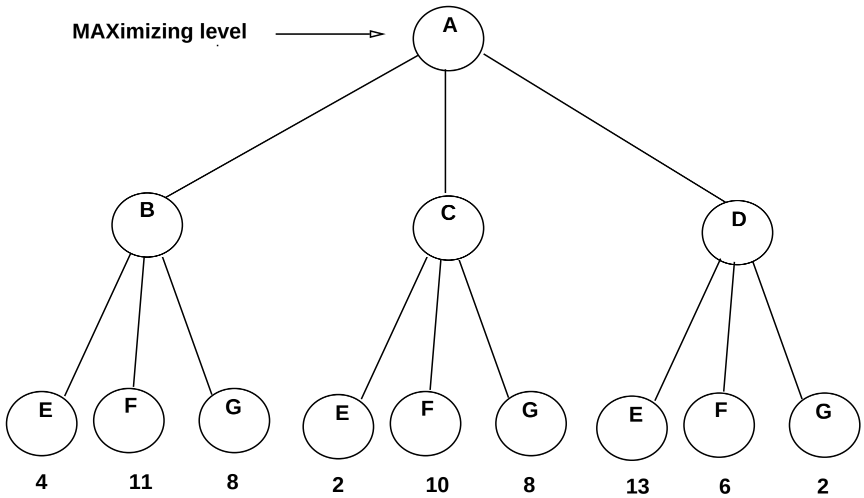

Figure 1: A minimax game tree.

Adversarial Search (15 points total)
Use the game tree from Figure 1 for the following questions. Leaves of the tree are labeled with the value returned by a static evaluation function.

1. (5 pts.) What should the next move be for the maximizing player?
2. (5 pts.) Indicate **one** example of $\alpha - \beta$ pruning in this tree. Assume left-to-right evaluation of nodes. Assume that there is no depth cutoff for the search, i.e., the $\alpha - \beta$ algorithm simply searches until it reaches a leaf.
3. (Choose one, 2 pts.) The horizon effect results from
   (a) ____ large branching factors
   (b) ____ depth bounds on search
   (c) ____ inadequate game-tree search heuristics
   (d) ____ inadequate state representations
4. (3 pts.) Name or describe one technique for limiting the horizon effect.

State Space Search: (15 points total)

Consider the search tree below. In this tree, node **A** represents the start state. Goal states are represented by double circles. In addition, each arc is labeled with the actual cost of traversing it. The number inside each node indicates the heuristic estimate of the remaining distance to the closest goal node.

For each of the search strategies named below, list the nodes in the order they would be expanded. Assume that the search operators return each node's successors according to the left-to-right ordering shown in the picture. (In the case of ties, choose the node that has been on the list of nodes to be expanded the longest.

1. (5 pts.) **Greedy Best-First Search**:
2. (5 pts.) **Iterative Deepening**:

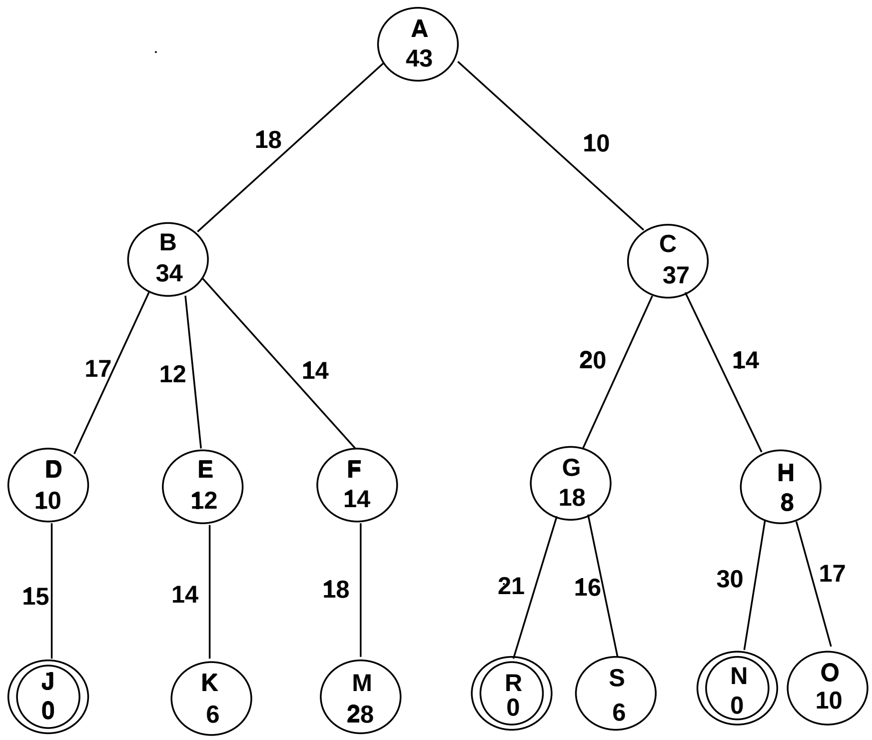

3. (5 pts.) **A* search**:

Resolution theorem-proving (10 points total)

If a course is easy, some students are happy. If a course has a final, no students are happy.

1. (3 pts.) Represent the above statements using *first-order logic*.
2. (7 pts.) Prove by resolution that, *if a course has a final, the course isn't easy.*

Heuristic Search and CSP: (20 points total)

You and some friends are going to a Chinese restaurant. You want to order four dishes from the menu to share. they must satisfy various constraints. Examples of such constraints are:

- One of your friends is a quasi-vegetarian, so at most one dish can use beef or pork.
- One of your friends doesn't like hot foods, so at most one dish can be spicy.
- No two dishes can have the same main ingredient.

1. (5 pts) Translate this task into a state-space search problem. What is the initial state? The goal state? The operators that transform one state into another?
2. (3 pts) How does the **state space** change if you are to find a legal combination of four dishes with the least total cost? How must the **search** change?
3. (3 pts) Describe a heuristic evaluation function for the problem of part 2, suitable for use in the A^* algorithm, that can improve performance over blind search.

4. (4 pts) Suppose you and a friend whom you are taking to dinner instead play a game where you alternate choosing dishes, always satisfying the constraints, until four dishes have been chosen. You try to make the total price as low as possible (since you are picking up the tab), while your friend tries to make it as high as possible. Explain why the A^* algorithm does not apply to finding an optimal strategy for this problem. Mention an algorithm that does.

5. (5 pts) Specify the problem as a constraint satisfaction problem.