15-381 Spring 06 Midterm

03/09/06

Name:_________________ AndrewID:______________

- This is an open-book, open-notes examination. You have 80 minutes to complete this examination.
- Write your answers legibly in the space provided on the examination sheet. If you use the back of a sheet, indicate clearly that you have done so on the front.
- Write your name and Andrew ID on this page and your andrew id on the top of each successive page in the space provided.
- Calculators are allowed but laptops and PDAs are not allowed.
- Good luck!

Problem	Points
Problem 1	/18
Problem 2	/6
Problem 3	/9
Problem 4	/9
Problem 5	/25
Problem 6	/17
Problem 7	/16
Total	/100

1 Informed Search

1.1 (8 points)

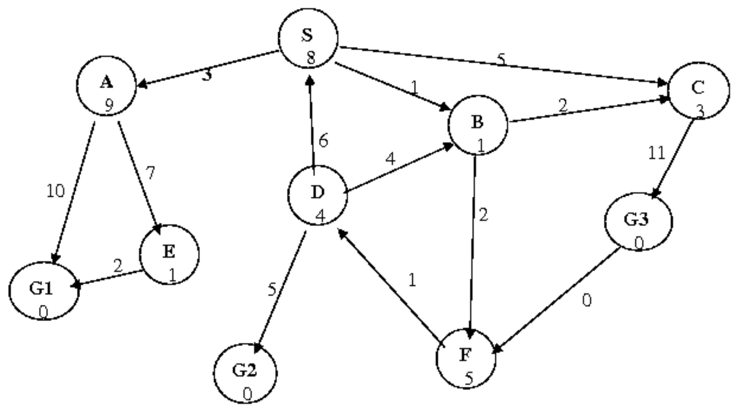

Consider the above search space, where S is the start node and G1, G2, and G3 are goal states. Search with multiple goals proceeds the same way as search with a single goal, except that the algorithm terminates when any of the goal states in the goal set is reached by an optimal path. The links are labeled with the cost of traversing them and the heuristic cost to a goal is reported inside the nodes.

For each of the following search strategies, indicate which goal state is reached (if any) and list, in order, all the states visited during the search. In the case of ties, the algorithm should expand nodes in lexicographical order (A before B, etc.).

- **Best-First Search**

  Goal State Reached: *G3* States visited: *S,B,C,G3 (Note: It does not matter if S was listed or not)*

- **A***

  Goal State Reached: *G2* States visited: *B,C,F,D,G2*

1.2 (5 points)

Suppose that we have a search problem for which it is not possible to compute directly the cost of each node (say s) to the goal (or the start). However, in this problem, we can assume that we have a procedure $P(s, s')$ which returns $+1$ if s is better (estimated as having lower cost to the goal) than s' and -1 otherwise.

Is is always possible to implement a version of best-first search by using this type of procedure instead of a numerical heuristic function? If not, state the condition(s) that must be met for best-first search to be implementable with such a procedure.

There are several ways of implementing best first search with such a function. Choosing the node in the queue that is better (according to P) than any of the other nodes is the most popular one (with different implementations suggested). Any reasonable variation was accepted. There is one condition, however, for best-first search to make sense, which is transitivity: $P(s, s") = +1$ if $P(s, s') = +1$ and $P(s', s") = +1$. If this condition is not satisfied, there is no notion of "bestänymore. One point was deducted for missing that one.

1.3 (5 points)

The relaxed version of a search problem P is another search problem P' with the same states, such that any solution of P is also a solution of P'. More precisely, P' uses the same states as P, and if s' is a successor of s in P, it is also a successor in P' with the same cost. If we view P and P' as graphs of states, the set of arcs in P is a subset of the set of arcs in P' with the same costs.

Prove that for any state s, the cost $c(s)$ of the path between s and the goal found by solving P' is an admissible heuristics for P.

Since P' contains all the arcs from P with the same costs: The lowest-cost path in P (of cost $h^*(s)$) is a valid path in P', therefore its cost must be greater than the cost of the minimal path in P': $c(s) \leq h^*(s)$, since, by definition, $c(s)$ is the lowest cost path in P'.

This turned out to be a more difficult question than anticipated; the situation being not helped by a typo in the original question. As a result, this ended being an almost guaranteed 5 points.

2 Local Search (6 points)

For each of the following algorithms, state whether it is guaranteed to converge to the *global* maximum. Assume that the state space is finite. Justify your answer in one sentence for each algorithm:

- Hill-climbing from a randomly chosen initial condition.

- Simulated annealing

- Genetic algorithm search

Answer: The answer is NO to all three. Although this is obvious for hill-climbing, it is not so obvious for the other two. Simulated annealing is only guaranteed to converge to the global maximum only in probability and with infinite iterations and temperature decrease. GA is not guaranteed either. Full marks for any reasonable statement of the type "yes, but in probability only", or "yes, if inifinite iterations".

3 Uninformed Search (9 points)

Given a maze with the start state S and the goal state G, number the cells in the order they are visited by each search algorithm using the following guidelines:

- Gray cells contain obstacles and can not be used.
- Moves are considered in the following order: Up, Right, Down, Left. Diagonal moves are not allowed.
- Loops are detected automatically and no state is visited twice.
- For Bi-Directional BFS, the start state is evaluated before the goal state.

Begin by labeling the start state as '1', the next state as '2', etc.

A. Breadth-First Search

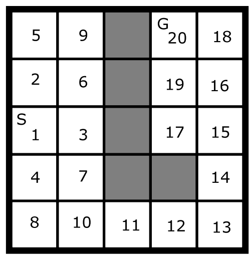

B. Depth-First Search

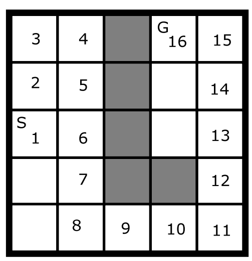

C. Bi-Directional Breadth-First Search
Either of the following was accepted:

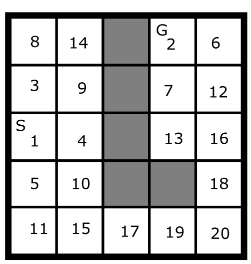

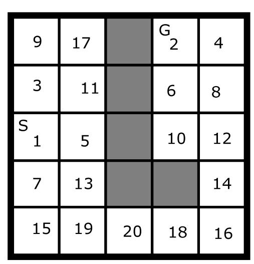

4 Game Trees (9 points)

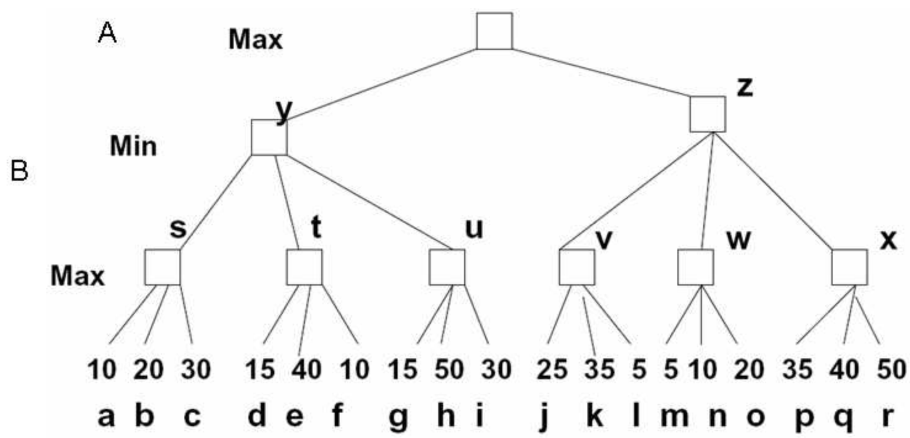

Consider the above game tree.

4.1 (2 points)

What is Player A's next move?

Y

4.2 (7 points)

Which nodes would *not* be visited if alpha-beta search is used?

f,i,x,p,q,r

5 Two-player mixed strategy Nash Equilibrium

Two competing grocery stores, Tiny Eagle and Partfood, are planning to renovate their Shadyside stores to sell more fruits in this neighborhood. But as part of their strategic plans, they have to simultaneously decide what kinds of customers they want to target and therefore what kinds of fruits they should buy from the suppliers. There are two type of customers in the Shadyside neighborhood. One kind prefers to buy organic fruits which are more expensive than the normal fruits. The other kind prefers to buy fruits with lower price and do not care if they are organic or not. Let us call the former type of customers organic customers(O) and the latter type non-organic customers(N).

Suppose there are 15,000 organic customers and 20,000 non-organic customers in Shadyside. If these two stores favor different types of customers, they will get the maximum amount of each type of customers to visit their stores. But if both stores favor the organic customers, Partfood will get 3/5 of the organic customers and Tiny Eagle will get 2/5 of the organic customers. In this case, non-organic customers will not go to any of these two stores. On the other hand, if both stores favor the non-organic customers, Tiny Eagle will get 3/5 of the non-organic customers and partfood will get 2/5. Similarly, organic customers will not go to any of these stores in this case. Assume the profit of each customer is 1.

5.1 (5 points)

Write down the matrix form of the game.

Answer: When Tiny Eagle's strategies are rows, Partfood's are columns,

	N	O
N	12, 8	20, 15
O	15, 20	6, 9

When Tiny Eagle's strategies are columns, Partfood's are rows,

	N	O
N	8, 12	20, 15
O	15, 20	9, 6

5.2 (5 points)

Are there any pure dominant strategies? Are there any pure strategy solutions? If yes, show what they are. If no, briefly explain why.

Answer: No pure dominant strategies. Pure strategy solutions are (N,O) and (O,N), since none of the stores would like to change their strategies in these cases.

5.3 (5 points)

Let us assume Tiny Eagle favors the organic customers with a probability $p = 0.2$ and Partfood already knows this. Which pure strategy will Partfood pick?

Answer: Organic customers.

5.4 (5 points)

Suppose Tiny Eagle favors the organic customers with a probability p and Partfood favors the organic customers with a probability q, compute the expected profit of Tiny Eagle and Partfood respectively in terms of p, q.

Answer: Payoff for Tiny Eagle is A1 = 6pq + 20(1-p)q + 15p(1-q) + 12(1-p)(1-q) = -17pq + 3p + 8q + 12. Payoff for Partfood is A2 = 9pq + 15(1-p)q + 20p(1-q) + 8(1-p)(1-q) = -18pq + 12p + 7q + 8.

5.5 (10 points)

Are there any mixed strategy solutions that are not pure strategy solutions? If yes, show what these solutions are. If no, briefly explain why.

Answer: Yes. Calculate the derivative dA1/dp = -17q + 3 = 0, dA2/dq = -18p + 7 = 0, so p = 3/17, and q* = 7/18.*

6 Robot Motion Planning

6.1 Voronoi Diagrams

Below is a map of an environment with walls and boundaries.

1. Roughly draw the Voronoi roadmap of this environment, taking into account both its walls as well as its boundaries.

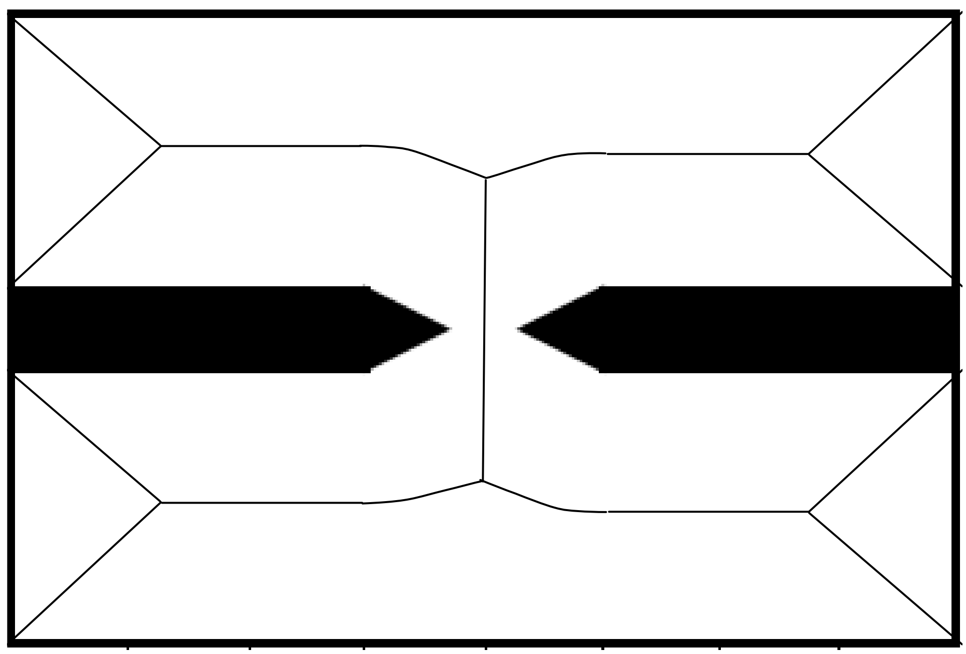

Any Voronoi diagrams which looked roughly like this one got full points.

6.2 Approximate Cell Decomposition

Consider the same environment with an initial approximate cell decomposition as shown. Cells are free space if they are *entirely* free of obstacles, else they are considered blocked. A point-robot (no shape or mass) starts at *Start* and wishes to reach *Goal*.

1. Is there a path from *Start* to *Goal* in this decomposition? `No`
2. **If so**, draw one such path in the diagram below. **If not**, modify the decomposition in a way that allows a valid path. draw your modifications as well as a valid path.

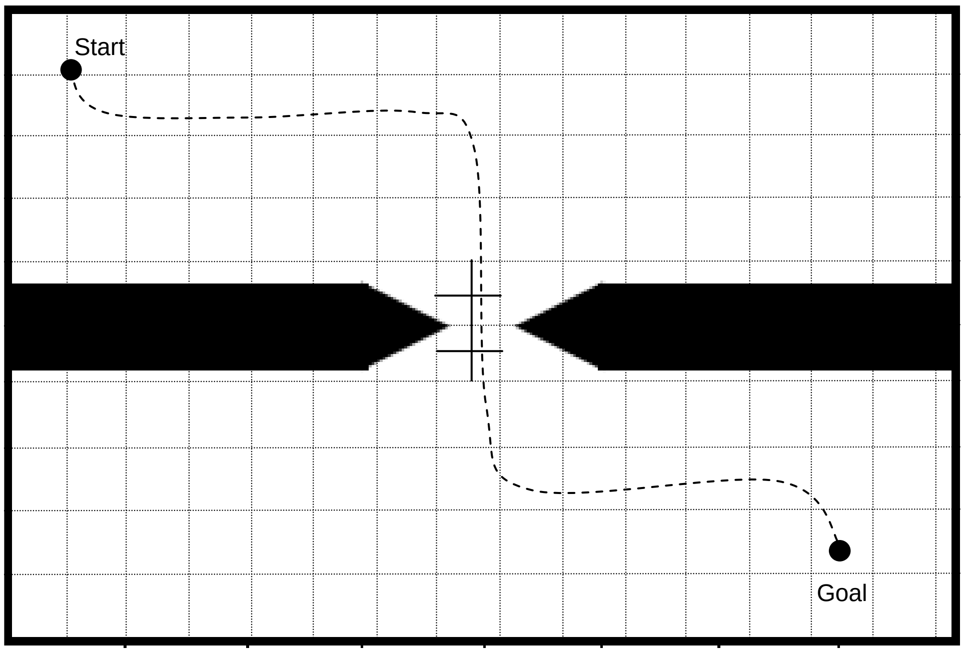

There are many possible solutions. The main point is that cells are considered 'blocked' if there is even the slightest overlap with an obstacle.

6.3 Sampling-based Methods

We now tackle the path-planning task with a sampling approach. Assume that we've sampled the points shown below and we wish to plan a path using the method for obtaining a *Probabilistic Road Map* (PRM) discussed in class.

1. Draw a PRM with $K = 2$ on these points, considering the walls and boundaries to be "forbidden regions".

2. Is there a path from *Start* to *Goal* in this PRM? `No`

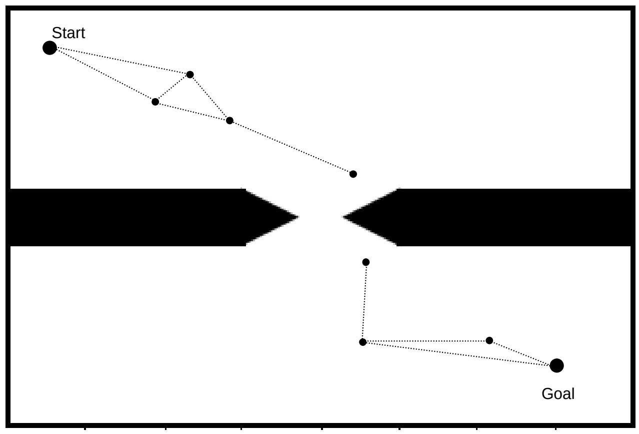

A valid PRM contains lines from 'Start' and 'Goal' to the other points, something which many people overlooked. Also, note that every point connects to its K nearest neighbors, so it is possible for a point to end up having more than K edges touching it. Valid PRMs with slightly different edges were not penalized since it is difficult to visually evaluate distances, but it should be clear that the resulting PRM has no path.

3. **If the above PRM has no path**, add points at any desired positions below to induce a PRM that contains a path. Draw the resulting PRM using the same procedure as above.

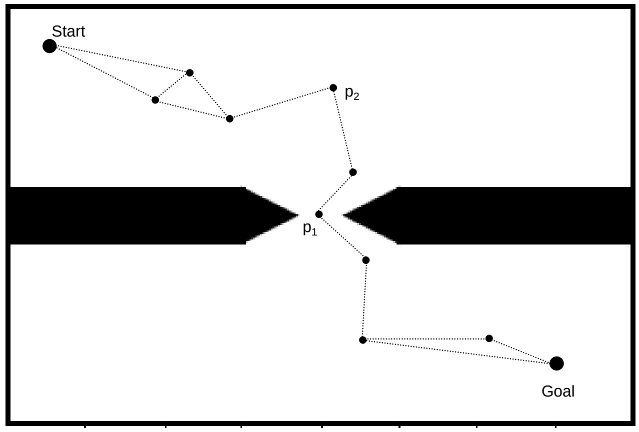

Note: Adding point p_1 by itself does not induce a PRM with a path, as you can check and verify. However, adding just a single point was not penalized because it is difficult to visually evaluate distances. Again, a valid PRM must contain edges from 'Start' and 'Goal' to other points.

4. If the two walls in this environment were to be extended until they left a much smaller gap (as shown below), which of the following path planning algorithms' running time would be affected? Assume that all algorithms are executed until a valid path is found.

Algorithm	**Affected? (Yes/No)**
Voronoi diagrams	No
Approximate Cell Decomposition	Yes
PRMs with $K = 2$	Yes
Visibility Graphs	No
Rapidly Expanding Random Trees (RRTs)	Yes

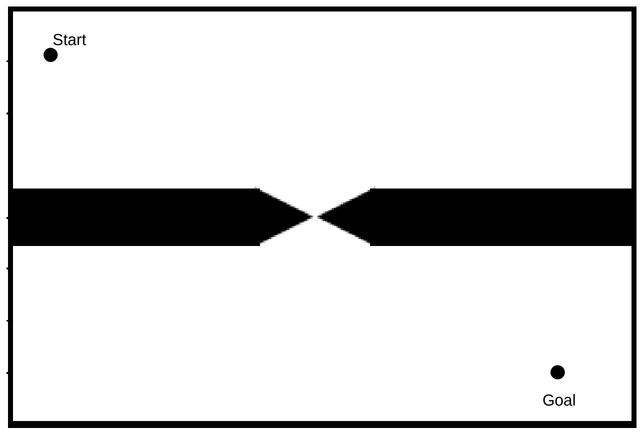

Sampling-based algorithms (PRMs,RRTs) are affected because the chances of randomly sampling just the right set of points to induce a path decrease with the size of the gap. Approximate cell-decomposition is affected because finer granularity decompositions would be needed until the path is found. Voronoi diagrams and Visibility graphs are unaffected because they are 'geometric' algorithms dependent only on the number of edges and vertices in the map.

7 Airport Counter Allocation Problem

The Pittsburgh International Airport has a fixed number of check-in counters, grouped into islands. Assume that there are n islands and m counters in total. The goal is to allocate enough counters to each flight (the number depends on the aircraft type). The counters for a particular flight have to be in the same island. A counter can not be shared by two flights running at the same time. You can assume the number of flights to be k and a function $isoverlap(i, j)$ which returns true if flights i and j overlap, and false otherwise. Also assume the number of counters needed by flight number i as c_i, and the number of counters in island i is m_i.

7.1 (4 points)

Suppose each counter needed by each flight is represented by the pair (island number, counter number). Define the counter allocation problem as a CSP.

- *Variables :* (I_{ij}, C_{ij}),
  where i *denotes flight number varying from 1 to* k
  and j *denotes the* j^{th} *counter assigned to the* i^{th} *flight and varies from 1 to* c_i.
- *Domain : domain of* I_{ij} *is* $[1, n]$
  domain of C_{ij} *is* $[1, m]$
- *Constraints :* $\forall i, j, k,\ I_{ij} = I_{ik}$ *(unique island for a flight)*
  $\forall i, j, k,\ C_{ij} \neq C_{ik}$ *(all counters for a flight are different)*
  $\forall i, j, k, l,\ i \neq j \wedge isoverlap(i, j) \rightarrow C_{ik} \neq C_{jl}$ *(two overlapping flights never share a counter)*

7.2 (6 points)

Now, suppose the counters needed by each flight are represented by the pair (island number, first counter number). Now, define the counter allocation problem as a CSP using this representation. What is the effect of the changed representation of the CSP problem?

- *Variables :* (I_i, C_i),
  where i *denotes flight number varying from 1 to* k
- *Domain : domain of* I_i *is* $[1, n]$
  domain of C_i *is* $[1, m]$
- *Constraints :* $\forall i,\ C_i + c_i - 1 <= m_{I_i}$ *(ensures that enough counters are available in the island assigned to a flight)*
  $\forall i, j,\ isoverlap(i, j) \rightarrow [C_i, C_i + c_i - 1]$ *and* $[C_j, C_j + c_j - 1]$ *dont intersect (two overlapping flights never share a counter)*

With the changed representation, the number of variables is reduced, hence the running time decreases.

7.3 (3 points)

Consider a portion of the Pittsburgh International Airport, where there are 5 counters. Counters 1,2, and 3 are in island 1, and counters 4,5 are in island 2. Let there be three flights A, B and C departing at the same time. Flights A, B, and C require 1, 1, and 3 counters respectively. Consider the formulation of the problem as a CSP discussed in the previous problem based on first counter number. Please answer True of False for the following questions:

- If the variables of flights are assigned in the order A, B, and C, the CSP search does encounter a conflict. *True*
- If the variables are assigned in the order of most restricted variable, the CSP search does encounter a conflict. *False*
- If the variables of flights are assigned in the order A, B, and C, forward checking helps in reducing the number of steps in search. *True*