Homework 2: Chapters 6 - 9

The following exercises are due on **Friday, November 6** at 4pm. Some of these problems may take a while to solve, so I recommend that you work on this assignment over the course of multiple days.

1. ***[10 pts.]*** The prisoner's dilemma is often interpreted as being proof that somehow cooperation between self interested agents is impossible. One "solution" to the prisoner's dilemma is to consider *program equilibria*, in which players submit strategies that may be conditioned on the programs submitted by others. The following is an example of such a program:
```
IF OtherProgram == ThisProgram THEN
     COOPERATE
ELSE
     DEFECT
END-IF.
```
   With reference to this program, explain how the framework of program equilibria permits cooperation as a rational outcome in the prisoner's dilemma.

2. ***[20 pts.]*** Consider the following two payoff matrices:

	i defects	i cooperates
j defects	3 3	4 2
j cooperates	1 1	2 4

	i defects	i cooperates
j defects	-1 -1	2 1
j cooperates	1 2	-1 -1

   a) For each scenario, what is each agent's preference ordering for the outcomes?
   b) For each scenario, which strategies (if any) are weakly or strongly dominated?
   c) Does either scenario have any Nash equilibria? If so, what are they?

3. ***[10 pts.]*** Consider a task-oriented negotiation domain in which there are two agents, the set of tasks T={a,b,c,d,e}, and the cost function is defined as follows: c({a})=4, c({b})=3, c({c})=5, c({d})=3, c({e})=2. The cost of performing a set of tasks together is simply the sum of doing the tasks individually, except that if a and c are performed together, they sum to 6 (not 9), and if b and d are performed together, they sum to 4 (not 6). Assume that in the encounter <{a,b},{c,d,e}> agent A_1 has proposed deal δ_1=<{b,d},{a,c,e}> and agent A_2 has proposed deal δ_2=<{b,d,e},{a,c}>. If both agents are using the Zeuthen strategy, which agent should concede on the next round? Show your work, including the cost and utilities of both deals to the agents, as well as the measure of each agent's willingness to risk conflict.

4. ***[10 pts.]*** Under what conditions will the auctioneer receive a higher price for a good if they use a Vickrey auction as opposed to a first-price sealed-bid auction? For this problem, assume that the auctioneer does not lie and the bidders do not collude,

5. ***[15 pts.]*** Consider the following deals that are under consideration by three self-interested agents (a_1, a_2, and a_3). Each deal is specified in terms of three numbers, each of which represents the utility of the deal to one of the agents. That is, deal $d=(d_1,d_2,d_3)$ has utility d_1 for agent a_1, d_2 for agent a_2, and d_3 for agent a_3. The fallback position (the one that occurs if the agents cannot reach a deal), has the following utilities for the agents (13, 8, 2).

   A=(20, 30, 20)
   B=(15, 10, 75)
   C=(15, 15, 15)
   D=(20, 7, 5)
   E=(50, 25, 25)

   a) Are all of these deals individually rational for all agents? If not, which ones are irrational for which agents?
   b) Which of these deals maximize social welfare? There may be more than one.
   c) Which of these deals are Pareto optimal? There may be more than one.

6. ***[10 pts.]*** Using the KQML language, describe how you would implement the Contract Net protocol. Be specific about the syntax of each message and give examples. For details on the language, you may wish to refer to: http://www.cs.umbc.edu/kqml/kqmlspec/spec.html

7. ***[10 pts.]*** In "Towards Flexible Teamwork," Milind Tambe uses decision trees to calculate when an agent should communicate. Show the derivation of his formulas for communicating both under the conditions of Fig. 7 (p. 102) and under those of Fig. 8 (p. 103). Hint: Start with EU(**C**) > EU(**NC**), substitute and simplify.

8. ***[15 pts.]*** Consider the pursuit task, in which four predators attempt to surround and capture a prey. Assume that the predators and prey live on a 100x100 grid, and that each can only move one square in one of the four compass directions on each round. Only one animal can occupy a square at any given time. Each predator can only see up to three squares away from it (but can see in all directions). The predators can broadcast messages to each other, regardless of distance. A prey is captured when a predator moves into its square. Describe an appropriate coordination mechanism for the predators. Explain the rationale for your choices. You do not need to implement your solution, but it should be described at a sufficient level of detail that it could then be easily implemented by a competent programmer who is familiar with multi-agent systems.