# Exam Performance Feedback Form

# COMP30191 (was CS3191)

2007/2008

It should be pointed out that for the first time, the exam mark is not the final mark for this course unit. The final mark is calculated by applying a factor of .8 to the exam mark (taken out of 100) and adding it to the course work mark. I decided not to scale the course work mark since it was a good deal lower than last year's. The comments below all refer to results in the exam.

General remarks: 69 students sat the exam. The result was extremely disappointing, with an all time low average of 49%. There was an unprecedented proportion of one third of students failing the exam. There were many very poor answers, and quite a few students didn't seem to really know what they were supposed to be doing. Eleven students got a first class mark, twelve students an upper second. This is a marked contrast to the previous development in recent years. I was worried about low attendance in both, the examples classes and the question sessions, and it seems my concern was justified. I will later try to work out whether there is a correlation between attendance and marks, but I can only do that once the marks are no longer anonymous. Check the course webpage for an updated version of this note.

**Question 1.** This was the most disappointing set of answers for a Question 1 that I've had in a while. 64 students attempted this question. The average mark was 57%, which is the lowest I've had. Fourteen students got a failing mark for this question, and nine students achieved 18 or more out of 20 marks.

Admittedly, the rules of the game were slightly more complicated than usual, but that's why there were 5 marks for the correct game tree and 5 marks for the strategies for both players—I do realize that it took a little bit of time to think through what the rules said.

Reasons why marks were lost typically were:

(a) I was surprised by the fact that a few students could not come up with a meaningful way of representing the state of the game. Nim does appear variously in the notes, after all (Example 1 (a)), and sample game trees for this game were available, only without the extra complications of the first two moves. Most students had an incorrect information set: The text said that Player 1 tells Player 2 how many matches are left, hence only two of the possible three positions after the first move are in the same information set. A lot of people also got the pay-offs wrong somehow (some people had pay-offs of 2 and $-2$ throughout, and others had pay-offs of 3 and $-3$ appearing for some reason. Clearly many did not read the text carefully enough. The game tree is not that complicated: Three opening moves, two of which are in the same information set, and the game is over after at most two further moves. (For some reason some people truncated the game tree and had no moves after Player 2's first move.)

(b) People who had a game tree with just one information set had an easier job here, but still, many mistakes were made. It does not take very much to see that for Player 1, the first move is really all that he has to choose, since if he is ever asked to make another choice he can force a win on that move. Hence there are three strategies for Player 1, which correspond to the possible first moves. Player 2 should always take the one match remaining if Player 1 announces that there's only one match left, and this combines with the possibility of wanting to take 0, 1 or 2 matches if Player 1 announces there are two matches left.

Some people lost marks because they did not say what simplifications they had made, and some used such a weird notation without any explanation that I could not work out which strategies they meant. A few people do not seem to have read the text which said that one should leave out strategies that choose a losing move when an immediately winning move is on offer.

(c) A fair number of people came up with a sensible matrix (subject to their errors in parts (a) and (b)), but I did see some things that did not make any sense at all, such as having pay-offs of 0 (although that was not an outcome in the game tree), or probabilities in the matrix in a game where no moves of chance are involved.

(d) The correct matrix can be reduced using pure strategy dominance to a $2 \times 2$ matrix, which can then be solved with the usual measures. The people whose only mistake was to have the wrong information set ended up with a matrix that wasn't reducible, I did give points for people who wrote that they thought that this had to be wrong. Curiously, it didn't seem to occur to any of them to go back and check their work.

(e) By and large the calculations of the value were correct (subject to previous mistakes, which I didn't penalize again). A lot people made a hash of arguing about fairness. A surprising number wanted to play a particular role based on their interpretation of the rules of the game rather than the calculations they had just made, and a lot wouldn't have picked the equilibrium point strategy! I do wonder why they think we do all this.

**Question 2.** All but two students tackled this question. Again, the answers were fairly disappointing.

The average mark for this question was 54%, the lowest ever (I've never had an average below 60% for this question). 25(!) students failed to get a pass mark for it, and 5 achieved a mark of 18 or higher. Many students could not even do parts (a) and (c) correctly, which have appeared in every exam so far and guarantee 9 marks for less than 10 minutes' work.

Reasons why some marks were lost typically were the following.

(a) I was surprised that some students clearly had no idea how to get these very quick 5 marks. Some students forgot to give the value, losing a mark. A few missed out one or more of the four (pure strategy) equilibrium points for some reason, even fewer tried to solve this using dominance arguments—and as a result didn't find all pure strategy ones. Some students didn't give the actual strategy pairs but only circled some entries in the matrix. Four marks were available for the four equilibrium points, one for the correct value.

(b) A quick (pure dominance) full reduction of the matrix brings no new equilibrium points (some student stopped half-way and found one more), but only two students remembered that mixing pure optimal strategies for either player leads to more optimal strategies. Given that this question has been asked before, that was disappointing.

(c) Most students knew how to solve a $(2 \times 2)$-matrix. In a few cases marks were lost because

- Player 1 and Player 2 were confused;
- students made mistakes in reading off the equations—a few of them obtained non-sensical results (a probability of 0, or a negative one) and still did not go back to check their work;
- a very few students inserted the value they got from solving the equations into the wrong position in the tuple describing the equilibrium point strategy;
- a few students don't seem to have caught on to the fact that the easiest way of calculating the value of the game is to take the value calculated by solving the equations and inserting it into either of the equations—the resulting number is the value of the game.

(d) The given matrix can be reduced in several ways to a $(2 \times 2)$ matrix. There's no guessing involved if one bears in mind the rules I taught. Quite a few students seemed to try to randomly pick strategies to be removed, rather than using those rules. Most students did very poorly here, removing strategies although they had derived a contradiction(!), or without giving any reason at all. Some students had difficulties with the inequality sign for Player 2 *versus* Player 1, and for some reason some students didn't write down an inequality for every entry of the strategy they were trying to remove. I gave 8 marks for this part because I knew that there were a lot of calculations involved (but then, part (a) gives 5 marks for something that can be done in a minute), so there should have been time to finish the job.

**Question 3.** Only 7 students decided to answer this question, which was the hardest in the exam (and which had been announced as such a number of times). Most of these didn't make a serious attempt—they appeared as if they couldn't think of anything better to do than to try Questions 1, 2 and 3. Most answers either didn't make any sense at all, or restricted themselves to calculating one equilibrium point, for which only 5 marks were available, and only one student managed a passing grade for the question. Hence the average mark was a poor 18%. Only one student made an attempt to prove that there couldn't be any other equilibrium points, but he or she tried to use an incorrect assumption to do so, and ultimately could only find one further equilibrium point. As given, the game has in fact an infinite number of equilibrium points.

**Question 4.** This question was attempted by 22 students. It had an average mark of 53%. Seven students failed to get a passing mark for the question, and one student got a mark of 18 or higher.

Many of the answers were somewhat confused and contained details that were wrong. Reasons why marks were typically lost are the following:

(a) Not correctly identifying the components (see the first three sections in Chapter 5 of the notes), or giving a partially wrong account of what they do.

(b) A few students described potential components for an evaluation function here. What was asked was how a game playing program works, which is mainly about how the components identified in (a) work together. Using the move generation function the program recursively generates a tree from the current position which is searched by the alpha-beta search algorithm to a given depth; the evaluation function is then applied to nodes at that depth. In the usual way this will generate an estimated value for each position reachable by the next move by treating these values as if they were the real things. The program then picks the move to the best position so identified.

(c) This was a question that is not fully answered in the notes. I was looking for five sensible aspects of comparing how computers and human beings play, and I was looking for a bit of history (How good is a program on your typical computer compared with your average player? What happens at the top end? Does it matter which game is played?).

**Question 5.** Forty-seven students attempted this question, and 24 of them did not manage to get a passing mark (that is, they had a mark below 8). Many of these same students had achieved quite high marks in Questions 1 and/or 2. Some people seem to have run out of time for this question, but many seemed ill-prepared. Nobody managed a mark of 18 or better. The overall average for this question was a failing grad of 39%, the worst ever.

I was disappointed with the general quality of the answers to this question. For the most part, were not low because people ran out of time (although some answers clearly had been written in haste) but because many did not give me the information I had asked for. All the questions had been asked in past papers.

Marks were typically lost for the following reasons.

(a) A few people did not write about the general game, but used specific pay-offs. Many people who described the general game forgot to specify that $(T + S)/2 < R$ must hold. Some people got their pay-offs mixed up and gave a wrong matrix.

(b) Most people remembered that the only equilibrium point for the simple game is for both players to defect. The situation is more complicated for the five round game, and Exercise 19 (a) states how other equilibrium points can be identified. Most people did not know this, and some even claimed that both players always cooperating is an equilibrium point.

(c) Playing the game for a fixed number of rounds is a slight generalization of playing 5 rounds. Mentioning that always defecting for both players is the only sub-game equilibrium point would have been nice. For the indefinitely repeated game, what can be said is described in Proposition 3.2. Very few people knew that, and accordingly I got many completely wrong answers, or completely irrelevant facts about the indefinitely repeated game which did not address the question of equilibrium points at all.

(d) I asked you to give the proof of Proposition 4.1, but very few people seemed to realize this. Some managed to rescue the situation by realizing that if one strategy cannot be invaded by the AlwaysD strategy then that means that that strategy plays better against itself than AlwaysD does. Again, there were a lot of completely irrelevant discussions here, or people who compared the wrong things.