Exam Performance Feedback Form

CS3191

2003/2004

General remarks: A higher exam average than last year, but still not as good as I had hoped.

Many students concentrated on Questions 1 and 2 and did not have much time left to answer their third chosen question, typically one of Questions 4 and 5. I think this plays a big role in the relatively low average marks for those questions. The second big factor is that many students clearly did not read the last two questions carefully enough and lost marks which would have been very easy to gain by putting down irrelevant detail rather than what they were asked about.

Question 1. About 95% decided to try this question. All in all this question was well answered and achieved an average mark of 61%.

Reasons why marks were lost typically were:

(a) Not giving the correct information sets in the game tree (what does a player know when he or she is making a move)? Some game trees showed probabilities although there were no chance elements in the game.

(b) The most frequent mistake was people assuming that Player 2 only had 2 strategies, to question or to accept under all circumstances. The correct answer is that this choice may depend on whether Player 1 announced 'Ace, King' or 'Two Kings', leading to four strategies.

(c) People who only had two strategies for Player 2 had a much smaller matrix to deal with, so they lost some marks here.

(d) Because a simple dominance argument turns the correct (4×4) matrix into a (4×2)-one, most answers here were of roughly the same difficulty. Most students were able to correctly calculate the sole equilibrium point. A lot forgot to translate this into how the players should behave in the game, or to comment on fairness.

Question 2. About 98% decided to answer this question. Very few of them did not know what they were supposed to do and got very low marks. The vast majority did extremely well with this question. It had an average mark of almost 60%.

Reasons why some marks were lost typically were

(a) Some students forgot to give the value, losing a mark. A few forgot that pure strategy equilibrium points lead to mixed ones.

(b) A lot of people got the mixed strategy equilibrium point but didn't spot the pure one, losing a mark.

(c) Quite a few people failed to stop reducing the matrix when dominance arguments could no longer be applied (when the matrix left is (3×3)). They were bent on getting the given equilibrium point using these methods, despite the fact that 'verifying a given equilibrium point' was something I demonstrated in the revision lecture at some length!

(d) Most people spotted the sole equilibrium point, but very few got full marks because the discussions of the sensibility of the solution in general wasn't very good.

Question 3. 15% of all students decided to answer this question, which was the hardest in the exam (and which had been announced as such a number of times). About half of those were students heading for a fail, who seemed to have decided to just answer the first 3 questions no matter what those might be. There were very few sensible discussions of winning strategies in Chomp. This question had an extremely low average mark of 26%. Most people who tried it would have been far better off choosing one of the questions 4 and 5, which were almost entirely about bookwork.

Reasons why marks were lost were typically

(a) reducing the given matrix and then solving it incorrectly, or not answering this part at all;

(b) not being able to state any sensible argument for why Player 2 can't have a winning strategy;

(c) making arguments about *plays* that would win rather than strategies, and so failing into account that in order to demonstrate the presence of a winning strategy for Player 1 it is necessary to show that *all* plays can be forced to result in a win for that player.

Question 4. This question was attempted by 57%. It had an average mark of 51%, which was rather lower than I had hoped.

By and large, this question reveals a general problem with exam technique. In many answers I got the impression that the students addressed their answers to somebody who knew the answer already, and therefore failed to explain what they were talking about on a basic level. What are these algorithms calculating? How do they do that? What is the use of alpha-beta search in game-playing programs? It would be useful for the students to think of themselves as trying to explain this to somebody who does not know anything about the topic.

Another issue with exam technique was obvious from the answers to part (c). The question clearly stated '*...give variations of the algorithm which might be applied*', yet many students only named one such. It went on to say '*Pick one of these variations and describe it in detail...*'—many students described several variations, where marks were available for only one such description.

Reasons why marks were lost were typically

(a) not stating what the algorithm is calculating; only using an example, but not stating in general how the algorithm works; using tiny examples without explaining how the tree is traversed;

(b) similar to (a);

(c) not stating what alpha-beta search actually does in a game-playing program (there were 3 marks available for that information), only listing one variant used (3 marks were available for naming 3 different ones—these were really easy marks to get), writing a confused description of one variant (3 marks for this bit).

Question 5. About 38% chose this question. It had a low average mark of only 43%. Many people got fairly low marks across the subquestions. This question is all bookwork, so the wrong answers indicate that people hadn't really internalized the main points of Sections 5 and 6 of the notes. I got the impression that many did not read this question properly as becomes obvious when looking at the reasons for losing marks:

(a) Discussing the Prisoner's Dilemma game for two rather than for several players, not giving a matrix, giving no, or only one example. There were two marks for at least two examples, one for the matrix, and the remaining two for an explanation.

(b) Most people could give a proper definition for 'invasion', but many had at best a confused argument for why the AlwaysD strategy cannot be invaded.

(c) Almost nobody got full marks for this part. Many confused the model (the indefinitely repeated PD game as an evolutionary system) with the situation to be modelled.

(d) There was a lot of confusion about the meaning of 'territorial', and the problem from (c) persisted. Similarly, a lot of the answers why it can be easier to invade a territorial system (strategies only compete with their neighbours, and two nice strategies may give each other sufficiently many points to do better than a resident 'mean' strategy) were confused at best, and plain wrong at worst. Many people brought up Nydegger as a strategy that cannot be invaded merely because it features as a relatively successful strategy in one example.