The University of Nottingham

SCHOOL OF MATHEMATICAL SCIENCES

A LEVEL 3 MODULE, AUTUMN 2006–2007

GAME THEORY

Time allowed TWO hours and THIRTY minutes

Candidates must NOT start writing their answers until told to do so.

This paper contains FIVE questions which carry equal marks.
Full marks may be obtained for FOUR complete answers.
Credit will be given for the best FOUR answers.

Only silent, self-contained calculators with a Single-line Display or Dual-line Display are permitted in this examination.

Dictionaries are not allowed with one exception. Those whose first language is not English may use a standard translation dictionary to translate between that language and English provided that neither language is the subject of this examination. Subject specific translation dictionaries are not permitted.

No electronic devices capable of storing and retrieving text, including electronic dictionaries, may be used.

DO NOT turn examination paper over until instructed to do so

1 (*a*) Suppose that complete analysis of a game tree to a given depth would return a value V. Show that with $\alpha-\beta$ pruning, analysis to the same depth will return the same value in the case $\alpha < V < \beta$. [5]
What conclusions can be drawn about the returned value in the cases $V > \beta$ and $\alpha > V$? [5]
Describe briefly TWO heuristics used in practice to speed up dynamic analysis by improving the efficiency of $\alpha-\beta$ pruning. [4]

(*b*) In an experiment, a game-playing program had its static evaluation replaced by a random-number generator. [The program still understood the rules of the game, and correctly evaluated terminal positions.] To the initial surprise of the experimenters, the program played much better than completely random moves. Discuss features of games and their trees that could explain this result. [11]

2 (*a*) If A and B are games, explain briefly what are meant by $-A$, $A+B$, $A > B$ and $A \parallel B$. [4]
For each of the following relations, either find games $A > 0$ and $B \parallel 0$ such that the relation holds, or prove that no such games exist:

(i) $A+B > 0$; (ii) $A+B = 0$; (iii) $A+B \parallel 0$; (iv) $A+2B = 0$; (v) $A+2B \parallel 0$. [9]

(*b*) Let G be the game $\{A, B, C, \ldots \mid D, E, F, \ldots\}$, and let H be the game which is the same as G but with Left's option A omitted.
(i) Show that $G \geqslant H$. [3]
(ii) Show that if $A \leqslant B$ then $G = H$. [3]
(iii) Suppose that A is strictly greater than all Left options of H. Does it follow that $G > H$? Give reasons. [6]

3 (*a*) In the game of Rayles, some dots are drawn on the page, and a move is to draw a closed loop which passes through exactly either one or two dots without crossing any other loop; that player loses who is unable to draw any such loop. Explain why this game is equivalent to Kayles, [4]
and hence or otherwise determine a winning move in the position shown below, in which four moves have already been played. [11]

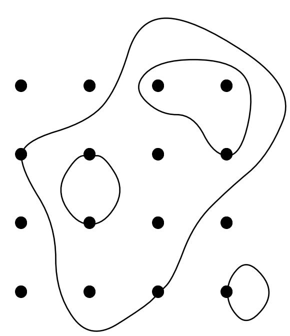

(*b*) The last move played was to draw the small loop through the dot in the bottom right corner. What move should have been played instead? [3]

(*c*) Suppose that the rule that the loop must pass through exactly one or two dots is replaced by a rule that the loop must pass through at least one dot. Show that the resulting game is equivalent to Nim played on heaps equal to the number of dots in each region, [4]
and find a winning move in the position shown. [3]

4 (*a*) Show that it is not possible on the ‘continental’ Solitaire board, as shown in the examples below, to reduce the standard start position [with the central square empty and the rest of the board occupied] to a single peg. [8]

(*b*) Show how to reduce the standard start position to the ‘The lecturer and his audience’ position shown below. [Occupied squares are indicated by ‘•’ and empty squares by ‘○’. You need not write out all the moves as long as the order in which any packs are applied is made clear.] [7]

(*c*) Show how to reduce the standard start position to the ‘Le tricolet’ position shown below. [10]

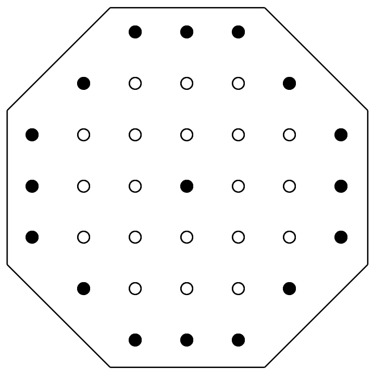

The lecturer and his audience

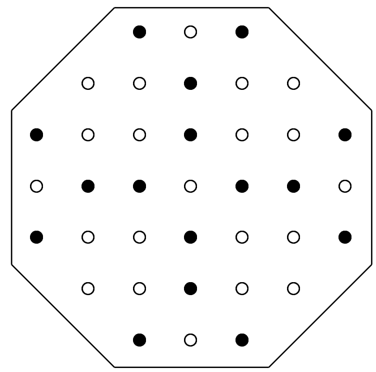

Le tricolet

5 Solve the matrix game with payoff matrix

$$\begin{pmatrix} -1 & 4 & 2 & 1 & 5 \\ 7 & 0 & 8 & 4 & 1 \\ -3 & -1 & 4 & 0 & 0 \end{pmatrix},$$

determining the lower and upper values, dominated strategies, the optimal strategy for each player, and the value of the game. [25]

1 Solution

(*a*) At any given node, we can assume recursively that analysis with pruning of sub-nodes will return values in accordance with the conclusions drawn here; note that static analysis [at the leaves] always returns the 'true' value [as determined by the program], and thus conforms to these rules. Suppose that at a non-leaf node, analysis of subnodes $1, 2, 3, \ldots$ returns values [negated as necessary in accordance with the player to move] $V_1, V_2, V_3, \ldots$. Then the true value should be the maximum of the true V_i. If this maximum, say V^*, satisfies $\alpha < V^* < \beta$, then no V_i has exceeded β, at least one has exceeded α, any that have exceeded α must be true values by 'the rules', and any that have not must have true values below α, also by 'the rules'. So V^* is the true maximum of the V_i, and $V^* = V$, as required. [5]

If V exceeds β, then at least one V_i must also exceed β, and this will cause a β-cutoff. So remaining nodes [if any] will not be searched, and it is possible that there is some $n > i$ such that $V_n > V_i$. So the returned value $V_i \leqslant V^* \leqslant V$, and is a lower bound [but $> \beta$] on the true value of the node.

If $V < \alpha$, then no V_i can exceed α. In this case, some or all of the V_i may have been determined by β-cutoffs in the child nodes, and so [after negation] may be upper bounds [but $< \alpha$] on the true value. So the returned value is less than α and is an upper bound on the true value. [5]

[Several heuristics given in lectures. These include iterative deepening and history heuristics to try to get good moves to the front of the move-list and thereby improve the chance of a β-cutoff; and aspiration search and minimal-windowing to reduce the gap between α and β, which will speed up the search in the case that $\alpha < V < \beta$, but slow it down if V is outside this range and a re-search is necessary.] [4]
[All bookwork to here.]

(*b*) There are two primary possibilities: *(i)* that there are many terminal positions in the tree, and their correct evaluation is strongly influencing play; *(ii)* that somehow the maximum of a collection of random numbers is correlated with a true static evaluation.

Case *(i)* could arise if bad moves are 'instantly' punished. In terms of a game like chess, this would mean that forced mates [within the 'horizon'] are still discovered, and so threats of forced mate result in wins unless the 'random' reply counters the threat. So the program is 'forced' to discover good replies to strong threats.

Case *(ii)* arises if 'mobility' is a significant direct or indirect component in the evaluation. The evaluation is more likely to choose lines where a player has more moves available, as the maximum of a set of random numbers is more likely to be in a large part of the set than a small part. In chess terms, this means that the program is likely to prefer positions with more moves available; this is correlated not only with 'mobility' in the sense that the pieces have plenty of room and are not cramped, but also with material, in the sense that the player with more pieces can usually choose between more moves. [11]
[Unseen. Marks to be obtained not for precise statements, but for sensible consideration of what the situation implies.]

2 Solution

(a) $-A$ is the game obtained by interchanging the roles of Left and Right; $A+B$ is the disjunctive sum of A and B—the game wherein a move consists of moving in *either A or B*; $A > B$ means that Left (whether or not on move) wins the difference game $A-B = A+(-B)$; $A \parallel B$ means that whoever is to move wins the difference game. [4]
[Bookwork.]

(a) For example, $A = 2$, $B = \pm 1$. [1]
(b) No such game; Left to play wins by playing first in B, replying in whichever game Right moves in. Alternatively, $A+B = 0$ implies $B = -A$ implies $B < 0$, contradicting $B \parallel 0$. [2]
(c) For example, $A = 1$, $B = \pm 2$. [1]
(d) For example, $A = 2$, $B = -1 \pm 2$. [2]
(e) For example, $A = 2 \pm 1$, $B = -1 \pm 2$. [3]
[Many other solutions, e.g. using games like *. Unseen. Last part harder, as all the fuzzy games B they 'know' have the property that $2B$ isn't fuzzy, so they need to realise that A can have a fuzzy component.]

(b) Part *(i)* is a matter of strategy. To show that $G \geqslant H$, we must show that $G-H \geqslant 0$, in other words that $G-H$ is a win for Left if Right moves first [we don't care who wins if Left moves first—it will be Left if $G > H$ or Right if $G = H$]. But Right moving first must play to one of $D-H$, $E-H$, $F-H$, $\ldots$, $G-B$, $G-C$, $\ldots$, and in each of these Left has a reply to the zero game $D-D$, $\ldots$, $B-B$, $\ldots$. The only move by either side to which there is no symmetric response is Left's move to A in G, and that cannot be played as Right's first move. [3]
[Unseen, but easy development of:]

For part *(ii)*, we already know that Right to play loses; we now need also to show that if $A \leqslant B$ then Left to play also loses. As in *(i)*, there is a symmetric response in $G-H$ to every move except Left's move to $A-H$, in which Right can reply to $A-B \leqslant 0$, so that Left [then] to play loses. So $G-H = 0$ or $G = H$. [3]
[Bookwork.]

For part *(iii)*, the supposition is that move A is the best move available for Left. Does it make any difference if move A can't be played? It does if the move is 'good', and it doesn't if the move is futile. Now it's easy to construct examples where $G = H$. For example, the games $\{-2, -1 \mid 1\}$ and $\{-2 \mid 1\}$ are equal [both zero]. Informally, it can't do you any harm to be given extra possible moves, and it may [but won't necessarily] do you some real good [even if the extra moves are the best available]. [6]
[Unseen.]

3 Solution

(a) Because the loops don't cross, they divide the page into regions, each of which contains some dots [possibly none]. Each move consists of 'crossing off' one or two dots, and, because the loop will create a new region, dividing the remaining dots in the same region [if any] into those inside and those outside that region. These are exactly the rules for Kayles, where there are [instead] rows of skittles, and a move is to knock down either one or two adjacent skittles, dividing the remainder of the row into two disjoint rows [possibly empty]. [4]
In the given position, the regions contain 5 dots [outside the big loop], 3 dots [between loops] and 2 dots [inside one of the small loops]. [2]
So we need to know the Grundy numbers for Kayles rows of size up to 5:

size	moves, to rows of size:	Nim equivalents	Grundy number [mex]
0			0
1	0	0	1
2	1, 0	1, 0	2
3	2, 1 and 1, 1	2, $1\oplus1=0$, 1	3
4	3, 2 and 1, 2, 1 and 1	3, $2\oplus1=3$, 2, $1\oplus1=0$	1
5	4, 3 and 1, 2 and 2, 3, 2 and 1	1, $3\oplus1=2$, $2\oplus2=0$, 3, $2\oplus1=3$	4

[3]

So the given position is equivalent to Kayles on rows of sizes 5, 3 and 2, and thus to Nim on heaps of size 4, 3 and 2. As $4\oplus3\oplus2 = 5$, the potential winning moves are to convert the Nim heap of size 4 to one of size $4\oplus5=1$, or that of size 3 to one of size $3\oplus5=6$, or that of size 2 to one of size $2\oplus5=7$; clearly, only the first of these is possible. [2]
So we win the Kayles by playing the heap of size 5 to one [or two] with Nim-equivalent value 1, and looking along the row we see this is the heap of size 4. [2]
So we should 'knock out' just one dot from the outside region. This could be a loop, like that on the bottom right dot, on any one of these five dots, or alternatively could be a large loop, again through one of the outer dots and otherwise surrounding the whole diagram. [2]

(b) Without the last move, the regions would have been 6, 3, 2. We could extend the previous table by another row [to find that a Kayles row of length 6 has Nim-equivalent 3], but the winning move to 4, 3, 2 from the previous part is already known and is still available. So a winning move is to knock out two dots from the outer region, for example by drawing a loop through the two bottom right dots [instead of through just one of them]. [3]

(c) Similarly to the above, the moves divide the page into regions, and knock out at least one dot in each region. These correspond directly to the moves in Nim, except that moves may also divide regions into two non-empty regions. But as $a\oplus b \leqslant a+b$, and we can assume inductively that the result is true for all simpler regions, the extra moves do not give Nim values that could not have been reached anyway, so the Grundy number of the resulting moves is the same as for Nim. [4]
The given position is equivalent to Nim on regions of size 5, 3 and 2; but $5\oplus3\oplus2 = 4$, so a winning move has to reduce the outer region to one worth $5\oplus4 = 1$. This can be achieved by drawing a loop through any four of the outer dots. [3]
[Kayles, its Nim equivalents, and Nim all familiar, and here applied to numeric problems; relationship of Rayles and Rims to Kayles and Nim mentioned in passing as a different way to play similar games, but not explored in detail.]

4 Solution

(*a*) Exercise in Reiss classification. [1]
After that, several ways. One of the simplest is to observe that the standard 'English' board is complementing, so the 'continental' standard start, consisting of the 'English' start plus the four pegs in $b2$, $b6$, $f2$ and $f6$, is in the same Reiss class as the 'English' finish of a single peg in $d4$ plus those same four pegs. Now use the 'Rule of Three' to move $f2$ to $c2$ and $f6$ to $c6$, jump $b2$ over $c2$ to $d2$, and move $d2$ to $d5$, and jump $b6$ over $c6$ to $d6$. Then observe that the remaining pegs in $d4$, $d5$ and $d6$ constitute a 3-pack which may be removed, leaving an empty board. So the continental' standard start is in the same Reiss class as the empty board, [3]
and so is in a different class from any single-peg position. [2]
As the Reiss class is conserved by any actual move, there is no possible sequence of moves to convert the start into a single peg. [2]
[Bookwork.]

(*b*) Many ways. For example, the 3-packs shown in the diagram below, applied in the given order, clear out all squares to be emptied except $d5$ and $d6$, leaving a final jump down $d5-d4$ to create the peg in the centre. [7]
[Unseen problem.]

(*c*) Many ways. For example, the 3-pack shown clears out $b5$, $c5$ and $d5$; then jumping $d7-d5$ restores a peg to $d5$, and jumping $b6-d6$ restores a peg to $d6$, leaving the 'tricolet' configuration in the top-left corner. Repeating this sequence in the other corners completes the solution. [10]
[Unseen problem, harder as there are no obvious packs.]

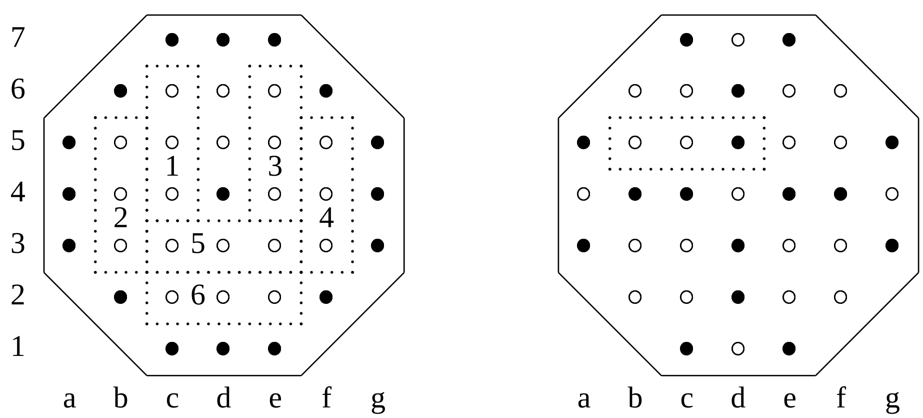

The lecturer and his audience   Le tricolet

5 Solution

Including the row minima and column maxima, and labelling players and strategies, the payoff matrix is

L\R	a	b	c	d	e	minima
a	–1	4	2	1	5	–1
b	7	0	8	4	1	0
c	–3	–1	4	0	0	–3
maxima	7	4	8	4	5	

So the lower value is the maximum of the row minima, or 0, and the upper value is the minimum of the row maxima, or 4. [2]
As these differ, the optimal strategy must be mixed [probabilistic]. [1]
By inspection, Left's strategy Lc is dominated, as each element is less than that of Lb; similarly, Rc is dominated by Rd, as each element is greater [and Right is trying to minimise the payoff], and Re by Rb. [3]
So we have an equivalent reduced matrix:

L\R	a	b	d
a	–1	4	1
b	7	0	4

[2]

As there are only two rows, a graphical approach is simplest; plot the probability [so between 0 and 1] of Left playing strategy La along the x-axis against the resulting value of each Right strategy.

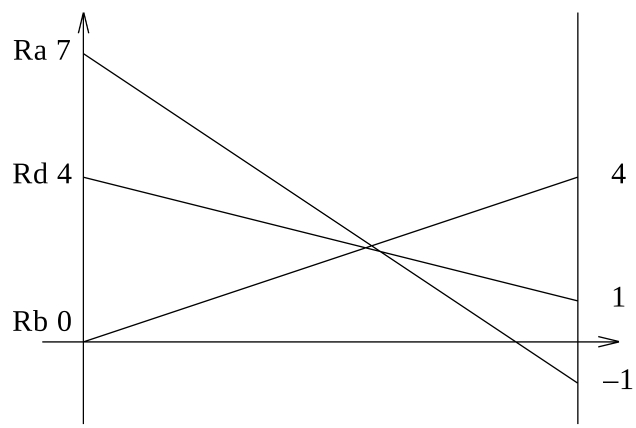

[3]

By inspection, we see that the minimax value [the highest point of the lowest value of the three lines] occurs where Rb crosses Rd, and in fact Ra is 'not worthwhile'. [If we don't trust the accuracy of the diagram, then it will be necessary to calculate all three intersection points.] [3]
Rb has equation $y = 4x$, Rd has equation $y = 4-3x$, so the required point of intersection is given by $x = \frac{4}{7}$, $y = \frac{16}{7}$. [3]
So Left should play strategy La with probability $\frac{4}{7}$, and strategy Lb with probability $\frac{3}{7}$, [3]
and the value of the game is $2\frac{2}{7}$. [2]
If Right plays Rb with probability p and therefore Rd with probability $1-p$, then Right's expected payoff against La is $4p+1-p$, and this must equal $\frac{16}{7}$, so $3p = \frac{9}{7}$, and Right should play Rb with probability $\frac{3}{7}$ and Rd with probability $\frac{4}{7}$. [3]
[Standard procedures on unseen numerical values.]

Commentary

1 This was quite poorly done. The bookwork was clearly only hazily remembered, and little understood. The last part of the question was intended to test understanding. Many of you noticed that because terminal positions were still recognised the program could still play an immediate win or avoid an immediate loss. But very few recognised that ‘played much better than random moves’ implied rather more than ‘sometimes even random moves are lucky’. As noted above, I was more concerned with whether you made sensible statements, showing that you understood the general process of tree searching, than with whether you produced mathematically sound assertions. This was a genuine experiment, back in 1994, by the way; from time to time since then it has been a ‘hot’ research topic.

2 This question was cobbled together from old questions, so those who had been diligent were rewarded. ‘Explanations’ such as ‘$A > B$ means that A is greater than B’ did not gain very much credit! Far too many students said that $B \parallel 0$ implies that $2B = 0$, despite producing examples such as 1 ± 1 or $\uparrow + *$ for part *(aiii)*. So there were a few good scores for part *(a)*, but a lot of very poor ones. Part *(b)* was better done, except by those who told me that if A is omitted, then in the difference game the players can swap moves until only A is left. Think rather about playing a game, say Noughts-and-Crosses, in which some move [say, Nought playing to a particular square] is forbidden in version H.

3 The Rayles/Kayles equivalence was quite well done, and there were many good solutions to the diagrammed position, and to finding the correct alternative move. A fair number of students slightly bizarrely decided that the five external points in the diagram were actually 2+2+1 [which I didn’t penalise very heavily if the rest of the work was correct] without realising that the position is then lost rather than won, and without applying the same rule in part *(c)*. Very few students managed to show that Rims, the ‘relaxed’ game of part *(c)*, is equivalent to Nim. Most of you either failed to notice that splitting heaps is not an option in Nim, or else noticed but then ignored the problem.

4 This question was well done. Most of you managed *(a)* and *(b)*; the main problem was students who reduced the standard start to some simple position from which no single-peg reduction was possible, and then just asserted that therefore the problem was impossible. You needed also to mention the Reiss classification, *and* to mention that the Reiss class is conserved. Sorry about the typo in *(c)*; the draft version, above, was correct, so it crept in during production. It put the position in the wrong Reiss class, so should have been manifest to anyone who was near a solution. Anyway, this was quite well done, with many good solutions, but also a fair amount of floundering. Neither my dictionary nor Babelfish tells me what ‘le tricolet’ means; but it appears to be a type of daffodil, and a word for the French flag, and a type of dress. I assume the position is meant to look a little like a daffodil; the position, name and solution date back at least to Lucas in 1882.

5 There were a lot of good solutions to this question, too. Far too many of you assumed that Right strategies with high values dominated those with low values instead of the other way around. I might have allowed this if you had then gone the other way with Left strategies as well; but if both players are looking for high values, then you might just as well agree to play the 8 in the middle of the matrix, and there is no game! The next problem was students who misdrew the diagram; sometimes even after commenting that it was only a sketch. You really do have either to draw the diagram reasonably accurately or to calculate where the lines intersect so that you know which strategies are worthwhile. After that came the students who assumed that all three remaining Right strategies were worthwhile, leaving themselves some rather complicated algebra and, unsurprisingly, a contradiction in the value of x. It is possible to recover from there—basically, we're part-way through the simplex method for linear programming—but you do have to realise at some stage that one of Right's strategies is still not worthwhile and so will have probability zero. There were a fair number of method marks available, so those who got the wrong strategies, *etc.*, often still got decent marks for knowing the general process, and for correctly solving the wrong matrix.

Summary

Overall, the marks were mostly reasonably decent and in line with your performances elsewhere; this is different from some previous years where the exam has either spread or bunched students too much. But some individual students departed a long way—two classes or so—either way from that. I assume, without proof, that that reflected students who either play a lot of games in their spare time or who avoid them. The median raw mark was slightly low, which I attributed to having the exam at the tail end of the period rather than to the exam being difficult; so all marks were raised slightly, which brought the median into the upper seconds.

In general, direct problems ['solve this situation'] were much better done than either explicit bookwork or 'please tell me your ideas about this' problems. At one level, this is a Good Thing; it bodes well for a career as a practising mathematician. On the other hand, it would be nice to think that you all knew what you were doing when practising!

A final comment. There were four 'visitors' taking the module. It's not easy to go to a foreign university, pick up a module in a style, and perhaps language, that you aren't used to, and perform well while also adjusting to the culture. So it was particularly pleasing that all four did very well, and obtained first-class marks.

ANW