

AN EVALUATION FUNCTION FOR THE GAME OF AMAZONS

J. Lieberum

Mathematisches Institut, Univ. Basel, Rheinsprung 21, CH-4051 Basel

lieberum@math.unibas.ch, <http://www.math.unibas.ch/~lieberum/>

Abstract Amazons is a fascinating game that shares properties of chess and Go. We have written a computer program that plays Amazons. This paper reveals the secret of this program: its evaluation function. We describe it by explicit formulas, mention the ideas and goals behind these formulas, and discuss possible refinements. By analysing a tournament game of AMAZONG against the former computer world champion 8QP we illustrate how the new features of our evaluation function can lead to victory.

Keywords: Amazons, evaluation function, AMAZONG

1. Introduction

Amazons is a many-faceted game. The game set typically used to play Amazons is a draughts board of size 10×10 , four white and four black chess Queens (called amazons), and a supply of Go pieces of one colour (called arrows). The starting position and a first move by White are shown in Figure 1. Each move consists of two steps: (1) the player chooses an amazon of the own colour and moves it like a chess Queen diagonally, vertically, or horizontally; the length of the move is up to the player as long as no obstacle (another amazon or an arrow) blocks the way; (2) this amazon has to throw an arrow. Arrows also move like chess Queens. They stay at their destination square for the rest of the game and are represented by black squares (in this paper) (see Figure 1 right). The players move alternately until one player can no longer move. This happens after at most 92 moves. The player who made the last move wins the game. White's advantage of making the first move can be compensated by allowing Black to pass n times (e.g., $n = 4$).

We first heard about amazons at a workshop on combinatorial game theory at MSRI in July 2000. We were fascinated by the deepness and subtlety of 'simple' positions in Amazons that have been analysed by Berlekamp (2000), Snatzke (1996, 2002), Müller and Tegos (2002). Inspired by discus-

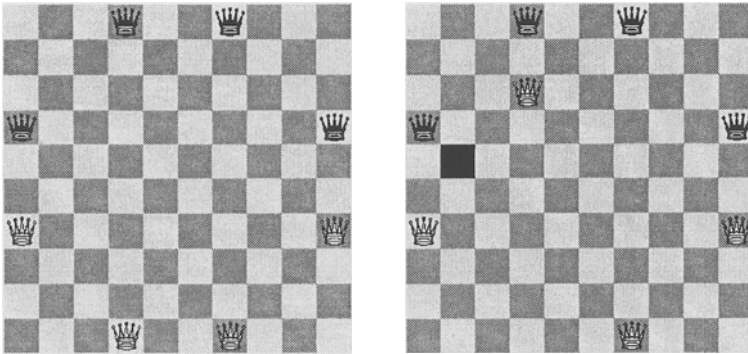


Figure 1. One good first move out of 2176 possible ones.

sions with Müller about his computer program *ARROW* and our own experiences of playing Amazons we started to write the computer program *AMAZONG*. After two years of successive improvements, *AMAZONG* has won the Amazons tournament at the seventh Computer Olympiad in Maastricht in July 2002. The reader is invited to play against the java applet *AMAZONG* at <http://www.math.unibas.ch/~lieberum/amazong/amazong.html>.

The general design of our program with a special focus on selective search has been presented in talks at the Universities of Jena and Edmonton (Lieberum, 2002). This paper complements these talks and concentrates on *AMAZONG*'s evaluation function that causes its characteristic style of play, clearly distinguishes it from other programs, and is probably its main strength.

2. The Different Phases of an Amazons Game

AMAZONG distinguishes three phases of the game: (1) the opening at the beginning of the game, (2) the filling phase at the end of the game, and (3) the main game that consists of everything else.

The opening in Amazons is the greatest challenge for computer programs. The reason is fourfold: (1) the absence of opening theory, (2) a branching factor of more than 1000, (3) many situations with more than 20 reasonable moves, and (4) the need for calculating deep variations. At this moment human play is still superior to computers in the opening. At the Computer Olympiad in Maastricht in 2002 *AMAZONG* made a random choice of the first move out of three possibilities and then started to play according to the results of a selective 5- or 6-ply search. Meanwhile, the opening book has grown to a machine generated database containing more than 30,000 moves following ideas of Lincke (2001). However, the benefit of opening books is limited in Amazons because of the enormous complexity of the game.

The filling phase consists of those positions where each empty square on the board can be reached by at most one player by some sequence of moves.

In most games this happens after approximately 50 moves. The filling phase includes positions with completely decomposed boards, meaning that amazons of different colours are separated by arrows. Then the outcome of the game can be determined by counting the number of moves left to each player. Although this problem is NP-hard (Buro, 2001), it is not difficult to play correctly in most positions that show up in real games on a board of size 10×10 . Typically, the players stop to play and agree on the outcome of the game when the filling phase starts.

Two examples of positions from the filling phase are illustrated in Figure 2. In the position on the left side of Figure 2 it seems that White has access to two empty squares, but he has to cut off one of the empty squares with his next move. Therefore this shape is called a defective territory. The position on the right side of Figure 2 is called *zugzwang* because it seems that Black has access to three empty squares, but if Black has to move before White does, then Black can only use two of the three empty squares. AMAZONG already tries to evaluate defective territory and many Zugzwang situations correctly before the filling phase starts. However, these parts of AMAZONG’s evaluation function are still far from being perfect. They will not be discussed here.




Figure 2. Defective territory and zugzwang.

Since in most Amazons games the opening book covers only the first few moves, one has to deal with many different situations in the main game until the filling phase begins and the outcome of the game becomes clear. One possible parameter which could help to choose an appropriate strategy in each situation is the number of moves played so far. AMAZONG uses a different parameter to choose its strategy. This will be discussed in the next section.

3. Territorial and Positional Evaluation

The goal of Amazons is to have access to more empty squares in the filling phase than the other player. When player j ($j \in \{1, 2\}$, player 1 is White) has exclusive access to a region of n squares, we count these squares as n secure points of territory of player j . When both players can reach a square by some sequence of moves, it is more complicated to predict which player will eventually shoot at that square. For this purpose AMAZONG uses heuristics based on the following ways to measure distances on an Amazons board.

Define the distance $d_1(a, b)$ of two squares a and b as the minimal number of chess Queen moves needed to go from a to b . When there is no path, let $d_1(a, b) = \infty$. Similarly, define the distance $d_2(a, b)$ as the minimal number of chess King moves needed to go from a to b . Obviously, we have $d_1(a, b) \leq d_2(a, b)$. The distances of player j from square a are then given by

$$D_i^j(a) = \min\{d_i(a, b) \mid \text{the square } b \text{ is occupied by an Amazon of player } j\}.$$

Figure 3 (left) is an example of $D_1^1(a)$. The upper left corners of empty squares contain the values $D_1^1(a)$ and the lower right corners contain the values $D_2^2(a)$. Figure 3 (right) is an example of $D_2^2(a)$.

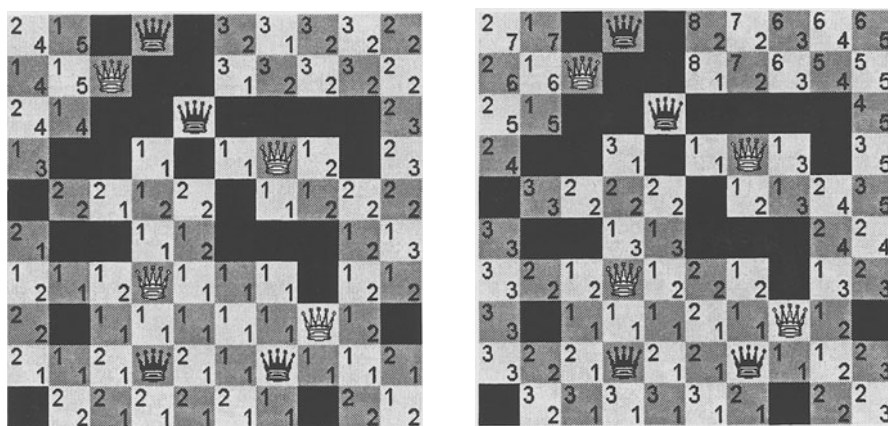


Figure 3. The minimal distances $D_i^j(a)$.

All Amazons programs seem to use D_1^j in one or another way (Hashimoto et al., 2001). The idea behind the definition of D_1^j is that $D_1^1(a) < D_1^2(a)$ indicates that player 1 has better access to the square a than player 2. One heuristic for estimating the territory of player 1 is to assume that player 1 will eventually shoot to *all* squares a with $D_1^1(a) < D_1^2(a)$. This heuristic works very well shortly before and in the filling phase. A problem of D_1^j at the beginning of the game is that a single amazon of player j in the centre can cause low values of D_1^j on the whole board, but player j cannot move the amazon into all directions at once. Here D_2^j comes in. One advantage of D_2^j is its locality: often amazons have to fulfil a certain task at their position like guarding the territory in their neighbourhood. Then a large value of $D_2^j(a)$ indicates that player j cannot move towards the square a without causing a positional damage, despite a possibly low value of $D_1^j(a)$. Another advantage of D_2^j over D_1^j is that it behaves more stable when the other player moves and shoots, especially when there are just

a few arrows. This makes D_2^j useful for long-term estimates and will stabilise the evaluation function in the beginning of the game.

We use D_i^j to assign local evaluations between -1 and 1 to each empty square. Positive values indicate an advantage of player 1. Then we sum these numbers over all empty squares in order to transform the local evaluations into global ones. One possible formula for global evaluations t_1, t_2 is given by

$$t_i = \sum_{\text{empty squares } \alpha} \Delta(D_i^1(a), D_i^2(a)),$$

where

$$\Delta(n, m) = \begin{cases} 0 & \text{if } n = m = \infty \\ \kappa & \text{if } n = m < \infty, \\ 1 & \text{if } n < m, \\ -1 & \text{if } n > m, \end{cases}$$

and $-1 < \kappa < 1$ is a constant with $(-1)^j \kappa \leq 0$ when player j is to move. The value $|\kappa|$ estimates the advantage of moving first when the distances of both players to an accessible square agree. We have made good experiences with $|\kappa| \leq 1/5$, but some fine-tuning is necessary after each modification of the evaluation function. We optimise the choice of κ frequently in order to obtain a low volatility of the evaluations during iterative deepening. This should help to avoid odd-even effects and supports aspiration search with narrow α - β -windows. (Marsland, 1986).

A program that uses the territorial evaluation t_1 as its evaluation function plays already quite reasonably, especially shortly before the filling phase. In contrast to that the value t_2 is useful in the beginning of the game but becomes less significant as the game goes on. The evaluations t_i share the drawback that they do not take into account that large values of $D_i^2(a) - D_i^1(a)$ are better for player 1 than small values. Therefore, other local evaluations than Δ seem to be important, too. The generic approach is to replace $\Delta(n, m)$ by some array of parameters and then to optimise these parameters. We have made good experiences with the choices

$$c_1 = 2 \sum_{\text{empty squares } \alpha} 2^{-D_1^1(a)} - 2^{-D_1^2(a)},$$

$$c_2 = \sum_{\text{empty squares } \alpha} \min(1, \max(-1, (D_2^2(a) - D_2^1(a))/6)).$$

Notice that in c_1 the local advantage $(D_1^1(a), D_1^2(a)) = (1, 2)$ is rewarded by 0.5 points for player 1, $(2, 3)$ by 0.25 points, $(1, 3)$ by 0.75 points, and squares a with $(D_1^1(a), D_1^2(a)) = (n, n)$ contribute 0 points. Other tuples are of minor

practical importance for c_1 . In contrast to c_1 , c_2 depends only on $D_2^2(a) - D_2^1(a)$ and only large differences indicate a clear advantage of one player.

Now we have to combine the values t_i and c_i into one evaluation function. A weighted sum with static weights does not seem to be appropriate for this because the importance of the values t_i and c_i varies during the game. Therefore, we should first try to compute the expected number W of moves needed until the filling phase starts. Instead of trying to estimate W directly we simply define

$$w = \sum_a 2^{-|D_1^1(a) - D_1^2(a)|},$$

where we sum over all empty squares a with $D_1^1(a) < \infty$ and $D_1^2(a) < \infty$. Obviously, we have $w = 0$ if and only if the position belongs to the filling phase and typically w decreases with the number of moves played. Therefore, we expect that a good estimate of W will be some function of w . For our purposes, w is just as good as W . Now define an evaluation t as

$$t = f_1(w)t_1 + f_2(w)c_1 + f_3(w)c_2 + f_4(w)t_2,$$

where $(f_i)_i$ is a partition of 1 (meaning $0 \leq f_i(w)$ and $\sum_i f_i(w) = 1$). The exact form of the functions f_i is a problem of parameter optimisation. Our choice of f_1 has been guided by the observation that t_1 becomes more and more important during the main game and gives very good estimates of the expected territory shortly before the filling phase. Hence, f_1 is monotonously decreasing and satisfies $f_1(0) = 1$. The counterpart of t_1 is t_2 . It rewards balanced distributions of the own amazons on the board or helps to hinder the other player from reaching such a distribution. This is most important at the beginning of the game. The values c_1 and c_2 allow to detect finer properties of the position than t_1 and t_2 alone, because they depend on the quality of local advantages. They support good positional play in the opening and a smooth transition between the beginning and later phases of the game. This is most evident for t_1 and c_1 : while at the end of the game only t_1 counts, c_1 rewards moves in earlier phases of the game that replace clear local disadvantages by small disadvantages and small advantages by clear advantages.

4. Mobility of Individual Amazons

AMAZONG is trying to enclose amazons of the other player inside of small regions at the beginning of the game. Compared to other computer programs, this is AMAZONG's main strength. In this section we present a modification of the evaluation function t (see Section 3) that is responsible for this behaviour.

Enclosing amazons typically does not cause an appropriate change of t (and especially of t_1) in the beginning of the game. This can be explained as follows:

when a single amazon A of player 1 is enclosed in some small region of n points, then the Amazons board is divided into two parts: the inside and the outside of that region. Player 1 has exclusive access to the territory on the inside. This contributes n points to t . On the outside, some active amazons of player 1 might overshadow the missing influence of A in D_1^1 . In addition, some amazons of player 2 that have helped to enclose A might not be in optimal positions but often have a large potential to improve their positions. The problem that A cannot reach the outside for the rest of the game is not reflected in the computation of t . The disadvantage of the enclosed amazon often starts to affect t several moves later. Then it is too late. Therefore, a correction term m is needed to take into account the mobility of individual amazons. Since active amazons can overshadow bad positions of passive amazons in the evaluation function t it seems more important to punish passive and enclosed amazons than to support active amazons in this correction term. To compute m quickly, consider first the number $N(a)$ of empty squares that can be reached from a by a single move of a chess King. The numbers $N(a)$ can be updated incrementally during the search inside of functions *doMove* and *undoMove*. For an amazon A of player j on the square a , let

$$\alpha_A = \sum_b 2^{-d_2(a,b)} N(b),$$

where we sum over all squares b with $d_1(a,b) \leq 1$ and $D_1^{3-j}(b) < \infty$. When $\alpha_A = 0$ we say that the amazon A is enclosed. Examples of the values $N(a)$, α_A and of enclosed amazons are shown in Figure 4. For example, for the white amazon A in the upper left corner of the figure on the left, we compute $\alpha_A = 7 + 6 + 5 + 3 + 3 + (5 + 4 + 7 + 4)/2 + 5/4 = 35.25$. The two white amazons in the lower right corner in this figure are enclosed.

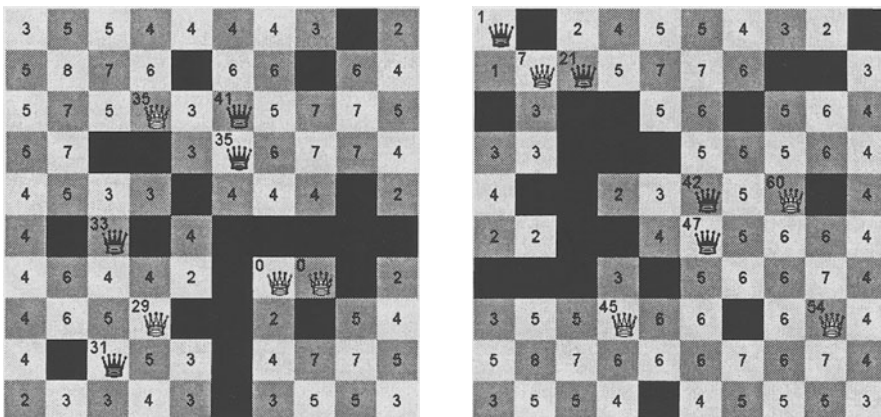


Figure 4. Neighbours $N(a)$ of empty squares a and the values α_A .

We have learned in discussions with experienced amazons players that at the beginning of a game on a board of size 10×10 enclosed amazons should be punished by a malus of at least 10 points. In general, we use w from the last section to define

$$m = \sum_{\substack{\text{amazons } B \\ \text{of player 2}}} f(w, \alpha_B) - \sum_{\substack{\text{amazons } A \\ \text{of player 1}}} f(w, \alpha_A)$$

for a suitable function $f \geq 0$. The exact choice of f is the hardest optimisation problem in our evaluation function $t + m$, so we restrict our description to the properties of f that did not change during our experiments: f satisfies $f(0, y) = 0$ and $\frac{\partial f}{\partial x}(x, y) \geq 0$ because the longer an amazon is enclosed before the filling phase starts the bigger is the disadvantage. Furthermore, f satisfies $\frac{\partial f}{\partial y}(x, y) \leq 0$ because a low value of α_A corresponds to a passive position of the amazon A . The last dependence is not linear. We have made good experiences with functions f that satisfy $2f(w, 5) < f(w, 0)$. This can be explained as follows: $\alpha_A \approx 5$ indicates that the amazon A is almost enclosed. However, there is a big difference between an enclosed and an almost enclosed amazon. The other player possibly has to move an own amazon B to an unfavourable square to prevent A from escaping. The resulting change of t then has to be compensated for by m . In addition, the task of guarding A makes the amazon B less mobile.

The big difference between enclosed and almost enclosed amazons can be seen on the right side of Figure 4: White can enclose the black amazon B with $\alpha_B = 1$ in his next move, but then Black can reply by enclosing the white amazon, too. Similarly, the task of guarding the white amazon in the upper left corner puts the black amazon B with $\alpha_B = 21$ in danger of getting enclosed.

5. Comparison between t_1 and $t + m$

In this section we compare our evaluation function $t + m$ with t_1 by using the game AMAZONG vs. 8QP played at the 7th Computer Olympiad in Maastricht. The position after 26 moves in this game is shown in Figure 3. AMAZONG won the game by 8 points, mainly due to the enclosed black amazon in the upper left corner. Figure 5 shows how t_1 and the different components of $t + m$ varied during the game.

In this diagram the values t_i are computed using $|\kappa| = 0.1$. The lines corresponding to t_1 , $t + m$ and m are clearly visible in the diagram. On move 13 White enclosed the black amazon which causes the maximum of the dashed line corresponding to m . Notice that at this point the evaluation $t + m$ predicts the outcome of the game very well and differs from t_1 by more than 18 points. Then $t + m$ and t_1 become more and more related and finally coincide when the filling phase is reached.




Figure 5. The components of the evaluation function $t + m$ during a game.

As expected, the values c_2 and t_2 are more stable than c_1 and t_1 . In addition, c_2 and t_2 are positive in almost all positions of the game. This indicates that the evaluation function of 8QP does not consider King move distances. Therefore, 8QP puts up no resistance against AMAZONG maximizing these components of $t + m$.

6. Refinements and Outlook

Consider positions with regions that are (almost) separated by arrows. How much is it worth when one player has a majority of amazons inside of such a region? Instead of looking for a general answer to this difficult question, we simply observe that the territorial evaluation t has the tendency to underestimate the advantage of the majority. A possible correction term of t could take into account the distances between each empty square and each amazon. However, the computations of these values would take almost four times longer than the computations of $D_i^j(a)$. Therefore, it seems more appropriate to compute only the numbers of amazons A_ν of player j on squares b_ν that satisfy $d_i(a, b_\nu) = D_i^j(a)$. These numbers can be computed efficiently together with $D_i^j(a)$. They are useful as additional inputs of refined definitions of c_i and t_i . In addition to these corrections, the disadvantage of having a majority of amazons in a *small* region early in the game should be reflected by m . This situation is not treated correctly by m because when amazons of both players are inside of one region the involved amazons are not considered as being enclosed.

A second refinement concerns the distribution of amazons on the board. In the opening it is desirable (especially for Black) to reach a position with exactly one amazon in each corner of the board. The distances from such a distribution can be used to improve the evaluation function in the opening phase.

In some experiments, we weighted squares in the computations of c_i and t_i . The weights depended on w and the distance of the square from the centre of the board. It is difficult to assess the importance of this third refinement.

A fourth idea for improvements is to repeat the constructions of Section 3 for other distance functions such as $d_1 + d_2$ or $2d_1 + d_2$ (or estimates of these distances that can be computed more efficiently). One has to decide very carefully how many different distance functions one should use, because each additional distance function slows down evaluations considerably.

The biggest weakness of our evaluation function seems to be the underestimation of large territorial frameworks at the beginning of the game. Our hope (potential fifth refinement) is to incorporate ideas from Lorentz (2002) to overcome this weakness. It is difficult because in many situations one has to make a choice between two plans that are often incompatible: (1) chasing and enclosing amazons or (2) building large territorial zones. The decision which plan is the more promising one in an actual position is a challenge for the next generation of Amazons programs.

References

- Berlekamp, E. R. (2000). Sums of $2 \times n$ amazons. In Bruss, F. T. and le Cam, L., editors, *Game Theory, Optimal Stopping, Probability and Statistics: Papers in honor of Thomas S. Ferguson, Institute of Mathematical Statistics Lecture Notes – Monograph Ser.*, volume 35, pages 1–34.
- Buro, M. (2001). Simple amazons endgames and their connection to hamilton circuits in cubic subgrid graphs, lecture notes in computer science. In Marsland, T.A. and Frank, I., editors, *Proceedings of the Second International Conference on Computers and Games*, volume 2063, pages 250–261.
- Hashimoto, T., Kajihara, Y., Sasaki, N., Iida, H., and Yoshimura, J., (2001). An evaluation function for amazons. In van den Herik, H.J. and Monien, B., editors, *Advances in Computers Games 9*, pages 191–202. IKAT, Universiteit Maastricht, The Netherlands.
- Lieberum, J. Selective search in amazons. Available at <http://www.math.unibas.ch/~lieberum/amazong/amazong.html>. Powerpoint presentation used in talks at the Universities of Jena and Edmonton.
- Lincke, T. R. (2001). Strategies for the automatic construction of opening books. In Marsland, T.A. and Frank, I., editors, *Proceedings of the Second International Conference on Computers and Games, CG00, Lecture Notes in Computer Science*, volume 2063, pages 74–86.
- Lorentz, R. (2002). Finding territory in amazons. In Uiterwijk, J., editor, *The Seventh Computer Olympiad Computer-Games Workshop Proceedings*, Technical Reports in Computer Science, Universiteit Maastricht.
- Marsland, T. (1986). A review of game-tree-pruning. *ICCA Journal*, 9(1):3–19.
- Müller, M. and Tegos, T. (1996). Experiments in computer amazons. In Nowakowski, R.J., editor, *More Games of No Chance*, pages 243–260.
- Snatzke, G. (1996). Exhaustive search in the game amazons. In Nowakowski, R.J., editor, *More Games of No Chance*, pages 261–278.
- Snatzke, G. (2002). New results in exhaustive search in the game amazons. University of Jena preprint.