



Lemmas on Partial Observation, with Application to Phantom Games

Fabien Teytaud, Olivier Teytaud

► **To cite this version:**

Fabien Teytaud, Olivier Teytaud. Lemmas on Partial Observation, with Application to Phantom Games. Computational Intelligence and Games, Aug 2011, Seoul, North Korea. 2011. <inria-00625794>

HAL Id: inria-00625794

<https://hal.inria.fr/inria-00625794>

Submitted on 22 Sep 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Lemmas on Partial Observation, with Application to Phantom Games

F. Teytaud and O. Teytaud

Abstract—Solving games is usual in the fully observable case. The partially observable case is much more difficult; whenever the number of strategies is finite (which is not necessarily the case, even when the state space is finite), the main tool for the exact solving is the construction of the full matrix game and its solving by linear programming. We here propose tools for approximating the value of partially observable games. The lemmas are relatively general, and we apply them for deriving rigorous bounds on the Nash equilibrium of phantom-tic-tac-toe and phantom-Go.

I. INTRODUCTION

Solving games is a common artificial intelligence exercise. One of the simplest case is the 3x3 tic-tac-toe, that many children solve manually by exhaustive analysis. Some games involve deep mathematics; for example, the standard Nim is solved exactly without exhaustive search [Bouton, 1902]; many mathematical developments exist on variants of Nim. Some games involve massive computer-based analysis [Schaeffer et al., 2007]; partial solutions are often given with restricted numbers of pieces [Kryukov, 2006]. These big successes of artificial intelligence involve both human expertise (through value functions) and big searches.

There are far fewer results in the partially observable case. If there were finitely many strategies per player, then we can rewrite the game under matrix form: $M_{i,j}$ is the result of the game (1 if player 1 wins, 0 if player 2 wins, $\frac{1}{2}$ in case of draw), and by linear programming we can find the Nash equilibrium of matrix game M . The case of Rock-Paper-Scissors is easily solved by this method, but there is no examples of exact solving of big partially observable games (except some restricted forms of Poker). There are consistent approaches [Littman et al., 1995], but they also do not scale to real games and approximate tools are classically used [Parr and Russell, 1995]. A real progress has been provided by probabilistic bounds in e.g. [Grigoriadis and Khachiyan, 1995], [Audibert and Bubeck, 2009]; these algorithms (bandit-type algorithms) provide results of the form “with probability at least p , the value of the game is in the confidence interval $[a, b]$ ”; these algorithms provide a precision ϵ (i.e. $b - a \leq \epsilon$) in time $O(K \log(K)/\epsilon^2)$ where K is the number of strategies. This is impressive in particular as it is sublinear in the number of entries in the matrix. However, it does not allow the use of human expertise for focusing on important parts.

F. Teytaud and O. Teytaud are with the TAO team. This team is a member of the INRIA, the Cnrs Umr 8623 and the university of Paris South.

We here provide simple intuitive mathematical tools for providing rigorous bounds on the Nash value of a game, and strategies realizing these bounds; for example, strategies ensuring a probability of winning $2/3$ for the first player in 4x4 Phantom-Ponnuki-Go. These results are qualitatively different from existing results:

- A difference with bandit tools is that here the bounds are not ensured with a given confidence rate but with certainty, i.e. (the constant p above is $p = 1$); a second difference is that we can use human expertise for improving the bound, without losing the rigor of the lower bound.
- A difference with classical alpha-beta based analysis is that we work on a partially observable game.
- A difference with exhaustive search is that we do not rely on a big computational effort; our methodology can be used with a big computer-based search, but this is not necessary and our examples below are built without using computers.

Phantom-games

Partially observable board games have been designed as a better approximation of war (for training of military officers) than classical board games. The ancestor of these games is probably “L’attaque” [Boutin, 2010], and then “stratego”; a classical challenge in AI is phantom-Go [Cazenave, 2006], [Cazenave and Borsboom, 2007], which is part of the annual computer Olympiads. Consider fully observable games with reward 0, $\frac{1}{2}$, or 1 (loss, draw, win respectively). Phantom-games are built from fully observable games by making all opponent’s stones/pieces/moves invisible; the only source of information is that whenever a move is illegal (due to unknown moves of the opponent), the move is cancelled and the player is informed that his move was illegal. It is therefore always a good piece of news to play an illegal move, as it provides information on the state of the game, without loosing one’s turn (the move is then cancelled and the player chooses another move, until he finds a legal move; the player is not allowed to play twice the same illegal move so that the game remains finite).

II. REMARKS ON PARTIALLY OBSERVABLE GAMES

This section presents some simple lemmas useful for analysis; we use them in later sections and therefore decided to show the lemmas that we first implicitly used; the key point is the relevant use of these lemma. The first lemma discusses symmetries, and the second lemma shows how a game can be replaced by a simpler version without increasing

the Nash value (i.e. in Lemma 2 we do not ensure that the Nash value is preserved, but only that it does not increase). Nash equilibria are well defined because we consider cases in which finitely many pure¹ strategies exist (by finiteness of the horizons). Some lemmas are then given specifically for phantom-games, which are an important special case of partially observable games.

Here a game is a tree G :

- with finitely many nodes;
- partitioned into nodes in which player 1 plays and nodes in which player 2 plays;
- with each node equipped with observations for player 1 and observations for player 2;
- with each leaf equipped with a reward (to be maximized by player 1 and minimized by player 2).

The edges are oriented and labelled with actions. Players are possibly randomized functions from sequences of observations to actions. A is the set of actions, supposed to be the same for player 1 and player 2. $\sigma(E)$ denotes the set of permutations of the set E . $R(\pi_1, \pi_2)$ is the expected reward associated to a game in which player 1 has strategy π_1 and player 2 has strategy π_2 .

Lemmas 1 and 3 are aimed at providing (through symmetries and dominating moves) very concise representations of strategies; we will use it for our human analysis of phantom-games, and believe that it is also helpful for computer-based analysis of partially observable games. Lemmas 2 and 4 are tools for proving bounds on the value of phantom-games.

Lemma 1: Consider a game G and a set $S \subset \sigma(A)$. Let's assume that if $\{s^{-1}; s \in S\} \subset S$ and if for any pure strategies π_1, π_2 for player 1 and 2 respectively, and $\forall s \in S$,

$$R(s \circ \pi_1, s \circ \pi_2) = R(\pi_1, \pi_2), \quad (1)$$

$$\text{then } \sup_{\pi_1} \inf_{\pi_2} R(p \circ \pi_1, \pi_2) = \sup_{\pi_1} \inf_{\pi_2} R(\pi_1, \pi_2),$$

where p is uniformly distributed in S (p is a randomly drawn permutation and p^{-1} is the inverse of the permutation p), and where \sup and \inf are on mixed strategies ("mixed strategies" are all strategies, including stochastic ones).

Remarks:

- Eq. 1 is the formal statement of the invariances of G w.r.t. S .
- This implies that we can consider only strategies which are invariant with respect to the symmetries of the game when evaluating the Nash equilibria.

Proof: Let S_1 (resp. S_2) be the set of mixed strategies for player 1 (resp. player 2). Assume Eq. 1 and consider p uniform on S . First, the inequality:

$$\sup_{\pi_1} \inf_{\pi_2} R(p \circ \pi_1, \pi_2) \leq \sup_{\pi_1} \inf_{\pi_2} R(\pi_1, \pi_2)$$

is clear; we just have to show

$$\sup_{\pi_1} \inf_{\pi_2} R(p \circ \pi_1, \pi_2) \geq \sup_{\pi_1} \inf_{\pi_2} R(\pi_1, \pi_2).$$

¹Pure strategies are deterministic strategies (as opposed to mixed strategies).

The proof is as follows:

$$\begin{aligned} & \sup_{\pi_1} \inf_{\pi_2} R(p \circ \pi_1, \pi_2) \\ &= \sup_{\pi_1} \inf_{\pi_2} R(\pi_1, p^{-1} \circ \pi_2) \text{ by assumption 1,} \\ & \geq \sup_{\pi_1} \inf_{\pi_2} R(\pi_1, \pi_2) \end{aligned} \quad (2)$$

because $\{p \circ \pi_2; \pi_2 \in S_2\} \subset S_2$. Eq. 2 concludes the proof. \square

Lemma 2: Consider a partially observable game G , and a subset T of the nodes of G , such that player 2 is to play in each node of T .

Then, consider the game G' defined as follows:

- The nodes are the same as those of G ;
- The edges from G are preserved;
- The observation for player 2 in each node z of T is unique, so that player 2 knows in which state he is when he reaches z .
- We add an edge from each node z of T to any node such that player 1 has the same sequence of observations as those from the root to z .

Then, the value of the game G for Player 1 is at least the value of the game G' .

Remark: This lemma shows that, when analyzing the value of the game for player 1, we can replace the unknown part of the state by the worst possible distribution on it. This means that, if we take the point of view of player 1, the value of the game will not increase if in some states, the game is stopped with its reward equal to the value of the Matrix game in which the opponent chooses the unknown part of the state (but remains consistent with the first player's observation) and we choose our strategy. In phantom-Go, this means that we can allow, without increasing the value of the game, the white player to change (privately) the position of the stones that black does not know.

The key point here is that, at first view, this lemma is too conservative. In fact, it will provide an efficient way of lower bounding the value of phantom-Ponnuki in 4x4.

Proof: The game allows the same actions for player 1, and more actions and more precise observations for player 2; therefore, the game is easier for player 2 (more precisely: the set of strategies for player 2 is extended, therefore the Nash value of the game becomes better for player 2). \square

The third lemma, given without proof, is about phantom-games specifically; it means that moves which are either a forced win or are illegal can be inserted into a strategy without decreasing its expected rewards. This very simple lemma provides concise representations: a strategy can be represented without specifying such moves, with the convention that all dominating moves are inserted.

Lemma 3: Consider a phantom-game G . If a move m is either illegal or a win for any state associated to a sequence $o = (o_1, \dots, o_k)$ of observations for player 1, and if π is a strategy for player 1 which does not play m after observing o , then π' dominates π , where π' plays equivalently to π in all cases except that

- $\pi'(o) = m$;
- $\pi'(o_1, \dots, o_k, o_{k+1}, o'_1, \dots, o'_l) = \pi(o_1, \dots, o_k, o'_1, \dots, o'_l)$.

Remark and definition: we will term such moves “dominating moves”, in the sense that inserting such moves will make strategies better (in the classical domination sense). This lemma is simple but allows a very short writing of sophisticated strategies: implicitly, untested dominating moves are played whenever possible. As detecting such moves is usual much faster than playing optimally in the general case, this is also a good tool in a code provided that the non-phantom-version can be solved at least in some cases.

Finally, we will use results for the fully observable game for guessing the value of the phantom version of a game, as follows:

Lemma 4: If game G is fully observable and G has at most N possible sequences of actions for player 1 and G is a win for player 1 (i.e. player 1 can win independently of the strategy of player 2), then the value of the phantom version of G is at least $1/N$ for player 1.

Proof: $1/N$ is the minimum probability for player 1 to play perfectly (for the non-phantom version of the game) if playing randomly. Therefore, player 1 can ensure a probability $1/N$ of winning just by playing randomly and uniformly. \square

Remark: The bound is tight, as one can see with the following game: in the non-phantom version, player 1 plays $i \in [[1, N]]$, and then player 2 plays $j \in [[1, N]]$. Player 2 wins if $i = j$. This (stupid) game is a clear win for player 2; in the phantom version player 2 wins with probability $1/N$.


III. APPLICATION TO PHANTOM-TIC-TAC-TOE

In 3x3 boards we do not write board coordinates; A, B and C are in abscissa, 1, 2, and 3 are in ordinates.

3x3 Tic-tac-toe is an easy case in the standard version of the game; but in the phantom-case it’s non trivial. We will here use lemmas 1, 3, and 4. Let’s term black the first player. We here consider six families of strategies for white, and two families of strategies for black, all of the black strategies starting with B2. White strategies (informed of the first black move) are as follows

- A. If no illegal move occurs, the two first white moves are contiguous (e.g. A3 B3).
- B. If no illegal move occurs, the two first white moves are adjacent corners (e.g. A3 C3).
- C. If no illegal move occurs, the two first white moves are in “knight move” (e.g. A3 C2).
- D. If no illegal move occurs, the two first white moves are opposite and in corners (e.g. A3 C1).
- E. If no illegal move occurs, the two first white moves are in “kosumi” (e.g. A2 B3).
- F. If no illegal move occurs, the two first white moves are opposite and not in corners (e.g. A2 C2).

We show below these 6 strategies, from the point of view of white (i.e. we show the initial black move in the center, plus the two white moves played if no white move is illegal):



We consider the following black strategy. Using lemmas above, we specify only a few moves, and the strategy must then be adapted as follows: as long as there are untested moves which are either illegal or a forced win, such moves are played; and the strategy is randomly symmetrized (any of the 8 natural symmetries of a board). This gives a very concise description of the strategy:

- Play B2;
- Then, play B1 (and then associated dominating moves); if B1 is illegal play a symmetry (B3, C2, or A2) of B1;
- Then, play C3 (and then associated dominating moves); if C3 is illegal play a symmetry (C1) of C3;
- Then, play dominating moves if possible, and moves ensuring a draw otherwise, until the game is over.

This black strategy ensures a probability of winning: 75% against strategy A; 100% against strategy B; 75% against strategy C; 100% against strategy D; 13/16=81.25% against strategy E; 7/8 against strategy F. We develop the most difficult case, the case of strategy E: in this case we point out that if black B2, white B1 or B3 or C2 or A2 (equilikelly), then black B1 leads to two cases:

- (75% of cases) black B1 legal, the black strategy ensures 75% of win; If B1 is legal (75% of cases), Black will win against strategies B, D and F in all the cases. For strategies B and D, it is simply because White has two stones in the corners and then can not block the line B1-B2-B3 from Black. Black will win against white strategy F because in that case, the two white stones are in A2 and C2 and then can not block neither the line B1-B2-B3 for Black. Against strategies A,C and E, Black can ensure at least a draw. Then the resulting win rate is $\frac{3}{6} * 1 + \frac{3}{6} * \frac{1}{2}$ which is equal to 75%
- (25% of cases) black B1 illegal, then black knows completely the state of the board and has a forced win as in the fully observable case by playing C1, forcing white A3, black C3.

As the 6 strategies A-F cover all possible strategies for white, the black strategy ensures an average reward 75%.




Fig. 1. The two possible situations in Phantom 2x2-Ponnuki.

The game is a draw in the non-phantom version; therefore the second player can ensure a draw with probability at least $(8 \times 6 \times 4 \times 2)^{-1}$ by playing randomly (Lemma 4, adapted to draws).

Combining these bounds (upper and lower), we conclude that the value of phantom-tic-tac-toe for the first player is in $[\frac{3}{4}, 1 - 0.5/384]$.

IV. APPLICATION TO PHANTOM-PONNUKI

Ponnuki-Go is a simpler version of the game of Go. The rules are the same, but the goal consists in capturing first a stone of the opponent.

This section is devoted to an application to phantom-Go. More precisely, we focus on phantom-Ponnuki, which gets rid of parameters (like Ponnuki).

The non-phantom versions of Ponnuki are solved until 6x6 [van der Werf et al., 2002], [Boissac and Cazenave, 2006]. In particular, 2x2 and 4x4 are wins for white, whereas 1x1, 3x3 and 5x5 are wins for black. This implies the following for phantom-Ponnuki:

- With N locations on the board, in 3x3 and 5x5, black can ensure a probability at least $1 - 1/((N - 2) \times (N - 4) \times \dots \times 1)$ of winning (by playing a first perfect move, and then by random play).
- With N locations on the board, in 2x2 and 4x4, white can ensure a probability at least $1 - 1/((N - 1) \times (N - 3) \times \dots \times 1)$ of winning (by playing a first perfect move, and then by random play).


In the 4x4 phantom-Ponnuki, we can therefore claim that white can win with probability at least $1/2027025$, by application of Lemma 4. We will also use Lemmas 1, 3; Lemma 2 will be used for the 4x4 case.

A. Phantom 2x2-Ponnuki


Phantom 2x2-Ponnuki is a draw on average: on Figure 1 (representing all possible cases up to a permutation), white wins (left) and black wins (right). It's clear that both situations are equally likely if any of the two players want them to be equally likely, and therefore the Nash equilibrium is a draw on average.

B. Phantom 3x3-Ponnuki


Phantom 3-3 Ponnuki is a win for black. Black tries to build a line B2 A2 C2, leading to:



If black succeeds it is a win. If black fails then black can ensure B2 B3 C2:



Then A1 wins if it is legal, otherwise B1 wins if it is legal, otherwise A2 wins if it is legal, otherwise C3 leads to the following situation which is a win (because it is White turn):




C. Phantom 4x4-Ponnuki

The 4x4 case is much more difficult, and in fact we don't have the complete solving; we just know that the value of the game (with 1 for a win for black, 0 for a loss) is in $[\frac{2}{3}, 1 - 1/2027025]$ (the upper bound has been shown in the beginning of section IV).


Let's now show that the probability of winning for black is lower bounded by $\frac{2}{3}$.

Black plays C3 or B3 or C2 or B2 (randomly with probability $\frac{1}{4}$); let's consider the C3 case without loss of generality. Then at his next turn, black plays B2 if possible:


- If B2 is legal, then black wins by trying B3:
 - if B3 is legal, then black reaches the following state which is a win (see section IV-C1):



- if B3 is illegal, then black tries C2; if C2 is legal, then black reaches the same state as above, within symmetry. If C2 is not legal, then black wins by an atari in B4, as here:




- If B2 is not possible, then black plays B3 and wins with probability $\frac{1}{3}$ with the following state (if white does not play C2, then black can play C2 and goes to the situation above):



(proof in section IV-C2).


These two situations are equally likely, so the value of the game is at least $\frac{1}{2} + \frac{1}{2} \times \frac{1}{3} = \frac{2}{3}$.

1) *Case in which black wins.*: Let's consider the following case:




Black plays C2. If it is legal, black has the central square and wins easily. Otherwise, black plays B1 (if B1 is illegal, black tries C3 which is equivalent by symmetry; if C3 is also illegal, black knows all the state and has an easy forced win by C4 D4 D2 B2 and immediate consequences of this):


- If B1 is legal, the situation is:




then black tries to win by D3; if D3 is legal the game is an easy win. If D3 is not legal, black plays C1; if legal, the game is an easy win. If C1 is not legal, black plays C4, as below:



which is another easy win. If C4 is not legal then B4 is necessary legal, leading to the following easy win:




2) *Case in which black wins with probability $\frac{1}{3}$.*: We now have the most difficult part: showing that black wins with probability $\frac{1}{3}$ (at least) in the following case:




Black plays D2 (which is necessary legal), and then D3 (which is legal, unless it provides a quick win for black by D4). Black then plays A3:

- if A3 is legal, then this move A3 leads to the following situation:




which is an easy win by D1 (if legal; otherwise C1 if legal);

- B4 if legal (then, if D4 is legal it makes a sufficient territory for ensuring a win as the seven white stones have no room without suicide; if D4 is not legal there is a win by C4).
- C4 otherwise (then A4 if legal, B4 otherwise), which is equivalent to the case above.
- If A3 is not legal, then black plays B4:
 - if B4 is not legal then blacks completely knows the situation:



and this is an easy win by C4 which is an atari.

- if B4 is legal, then we get the situation below, where B4 has been played by black and two white stones are unknown:



We have to analyse the case in the figure above. We use the main lemma from Section II:

We therefore allow white to choose (privately) the distribution of its 2 missing white stones, and will show that black wins anyway with probability at least $\frac{1}{3}$.

We will see two strategies for black, two families of strategies for white (covering all white possibilities); we'll then build the matrix of this 2x2 game (black choosing between his two strategies, white choosing between his two families of strategies). The purpose of the following lines is to fill the diagonal of the matrix in Table I.

We distinguish two strategies for black:


- First strategy, black plays D4, if D4 is illegal, then C4; the rest of the strategy does not matter.
- Second strategy: black plays A4, and then A2; the rest of the strategy will be detailed below.

We distinguish two families of strategies for white:

- First family of strategies for white: white has one stone in either C4 or D4, plus one stone in D1. In this case, the strategy "black D4" is a win (if D4 is illegal, C4 is a win). This fills the lower right part of Table I.
- Second family of strategies for white: if white has no stone in C4 or D4 or no stone in D1², then black has three liberties, which allows three moves before any trouble. The following strategy for black (which is the second strategy in the list above) wins with probability at least $\frac{1}{2}$:

- Black plays A4. If legal, then black plays A2 (legal or not). Then black plays D1 or D4 (with probability $\frac{1}{2}$ each), and we show that this ensures a win with probability $\frac{1}{2}$ whatever may be the choice of white in this family of strategies:

- * D4 is a win if there is a white stone in C4 or in D4³.
- * D1 is a win if there is no white stone in C4 or in D4; this is because if there's no white stone here, the lower left part has necessarily filled its liberties, as shown below:




| | White has no stone in C4/D4 or no stone in D1 | White has a stone in C4 or D4 and a stone in D1 |
|----------------|---|---|
| Black A4 (...) | $\frac{1}{2}$ | |
| Black D4 (...) | | 1 |

TABLE I

MATRIX GAME IN SECTION IV-C2; THIS GAME IS EASIER FOR WHITE, AND BLACK ALREADY GETS A WIN WITH PROBABILITY $\frac{1}{3}$; THIS IS ENOUGH FOR THE EXPECTED RESULT. WE DON'T NEED THE RESULT OUT OF THE DIAGONAL HERE, AS WE ONLY WANT A LOWER BOUND FOR BLACK.

- * If A4 is not legal, then black plays A2; if it is not a win, A2 is illegal; we get the following figure, with white to play:




White can either play in black's eye (C4D4) or in its own eye (A1B1). Black does not know where white has played. Then, there are two possibilities:

- There's a white stone in D4 or C4; then black wins by D4 (or C4 if D4 is illegal).
- There's a white stone in A1 or B1; then black wins by A1 (or B1 if A1 is illegal).

This is exactly a matching penny game. Playing each of these two strategies with probability $\frac{1}{2}$ each ensures a win with probability $\frac{1}{2}$ for black.

So we get a game with black choosing its strategy and white choosing his two hidden stones, from the following situation:



The matrix (probability of winning for black, row player) is as given in Table I: The value of this matrix game is $\frac{1}{3}$, hence the expected result. Black first wins with probability $\frac{1}{2}$, and in remaining cases black wins with probability $\frac{1}{3}$ at least, which leads to an overall probability $\frac{1}{2} + \frac{1}{2} \times \frac{1}{3} = \frac{2}{3}$.

²Please note that white can not have two stones in C4-D4.

³If there's a white stone in C4, black D4 is an immediate win, and if there's a white stone in D4, then black plays again, and wins with E4.

V. CONCLUSION

We provided lemmas for analyzing partially observable games. These lemmas provide concise representations of strategies, without loss of performance, by pruning dominated strategies (Lemma 3) and by implicitly symmetrizing (Lemma 1). This concise representation is helpful for human analysis (in this paper) and our main further work is its use for programs as well. Two other lemmas are useful for deriving upper and lower bounds on values of partially observable games.

Using just pen and paper, we could show that, on average, 4x4 phantom-Ponnuki is a win for Black (probability of winning at least $\frac{2}{3}$). The non-phantom-version is a win for white, and the only size which is solved (from the empty board) in the standard (full information) setting and not solved by our analysis in the phantom case is the 5x5 case.

The main further work is the use of these lemmas within an implementation, in order to find bounds on bigger partially observable games. We conjecture that phantom-Ponnuki in 5x5 is a win for Black.

REFERENCES

- [Audibert and Bubeck, 2009] Audibert, J.-Y. and Bubeck, S. (2009). Minimax policies for adversarial and stochastic bandits. In *22th annual conference on learning theory*, Montreal.
- [Boissac and Cazenave, 2006] Boissac, F. and Cazenave, T. (2006). De nouvelles heuristiques de recherche appliques la resolution d'Atarigo. In *Intelligence Artificielle et Jeux*, pages 127–141. Hermes Science Lavoisier.
- [Boutin, 2010] Boutin, M. (2010). Les jeux de pions en france dans les années 1900 et leurs liens avec les jeux étrangers. Invention dun jeu singulier : lattaque. In *Proceedings of BGA'2010*.
- [Bouton, 1902] Bouton, C. (1901-1902). Nim: A game with a complete mathematical theory. *The Annals of Mathematics*, 3(1/4):35–39.
- [Cazenave, 2006] Cazenave, T. (2006). A phantom-go program. In *ACG*, pages 120–125.
- [Cazenave and Borsboom, 2007] Cazenave, T. and Borsboom, J. (2007). Golois wins phantom go tournament. *ICGA Journal*, 30(3):165–166.
- [Grigoriadis and Khachiyan, 1995] Grigoriadis, M. D. and Khachiyan, L. G. (1995). A sublinear-time randomized approximation algorithm for matrix games. *Operations Research Letters*, 18(2):53–58.
- [Kryukov, 2006] Kryukov, K. (2006). EGTs online.
- [Littman et al., 1995] Littman, M. L., Cassandra, A. R., and Kaelbling, L. P. (1995). An efficient algorithm for dynamic programming in partially observable markov decision processes. *Technical Report CS95 -19, Brown University, Providence, Rhode Island*.
- [Parr and Russell, 1995] Parr, R. and Russell, S. (1995). Approximating optimal policies for partially observable stochastic domains. In *Proceedings of the International Joint Conference on Artificial Intelligence*.
- [Schaeffer et al., 2007] Schaeffer, J., Burch, N., Bjornsson, Y., Kishimoto, A., Muller, M., Lake, R., Lu, P., and Sutphen, S. (2007). Checkers is solved. *Science*, pages 1144079+.
- [van der Werf et al., 2002] van der Werf, E., Uiterwijk, J., and van den Herik, H. (2002). Solving Ponnuki-Go on small boards. In *Proceedings of 14th Belgium-Netherlands Conference on Artificial Intelligence (BNAIC'02)*, pages 347–354.