

Theory of Computer Games: Selected Advanced Topics

Tsan-sheng Hsu

徐讚昇

tshsu@iis.sinica.edu.tw

<http://www.iis.sinica.edu.tw/~tshsu>


Abstract

- **Some advanced research issues.**
 - **The graph history interaction (GHI) problem.**
 - **Opponent models.**
 - **Searching chance nodes.**
 - **Proof-number search.**

Graph history interaction problem

- The graph history interaction (**GHI**) problem [Campbell 1985]:
 - In a game graph, a position can be visited by more than one paths.
 - The value of the position depends on **the path** visiting it.
 - ▷ *It can be win, loss or draw for Chinese chess.*
 - ▷ *It can only be draw for Western chess.*
 - ▷ *It can only be loss for Go.*
- In the transposition table, you record the value of a position, but not the path leading to it.
 - Values computed from rules on repetition cannot be used later on.
 - It takes a huge amount of storage to store all the paths visiting it.
- This is a very difficult problem to be solved in real time [Wu et al. '05].

GHI problem – example



- Assume the one causes loops loses the game.
- $A \rightarrow B \rightarrow E \rightarrow I \rightarrow J \rightarrow H \rightarrow E$ is loss because of **rules of repetition**.
 - ▷ *Memorized H as a loss position.*
- $A \rightarrow B \rightarrow D$ is a loss.
- $A \rightarrow C \rightarrow F \rightarrow H$ is loss because H is recorded as loss.
- A is loss because both branches lead to loss.
- However, $A \rightarrow C \rightarrow F \rightarrow H \rightarrow E \rightarrow G$ is a win.

Opponent models

- In a normal alpha-beta search, it is assumed that you and the opponent use the same strategy.
 - What is good to you is bad to the opponent and vice versa!
 - Hence we can reduce a minimax search to a NegaMax search.
 - This is normally true when the game ends, but may not be true in the middle of the game.
- What will happen when there are two strategies or evaluating functions f_1 and f_2 so that
 - for some positions p , $f_1(p)$ is **better** than $f_2(p)$
 - ▷ “better” means closer to the real value $f(p)$
 - for some positions q , $f_2(q)$ is **better** than $f_1(q)$
- If you are using f_1 and you know your opponent is using f_2 , what can be done to take advantage of this information.
 - This is called OM (**opponent model**) search [Carmel and Markovitch 1996].
 - ▷ In a MAX node, use f_1 .
 - ▷ In a MIN node, use f_2 .

Opponent models – comments

■ Comments:

- Need to know your opponent's model precisely or to have some knowledge about your opponent.
- How to learn the opponent model on-line or off-line?
- When there are more than 2 possible opponent strategies, use a probability model (PrOM search) to form a strategy.

Search with chance nodes

■ Chinese dark chess


- Two player, zero sum, **complete information**
- **Perfect information**
- **Stochastic**
- There is a **chance** node during searching [Ballard 1983].
 - ▷ *The value of a node is a distribution, not a fixed value.*

■ Previous work

- Alpha-beta based [Ballard 1983]
- Monte-Carlo based [Lancoto et al 2013]

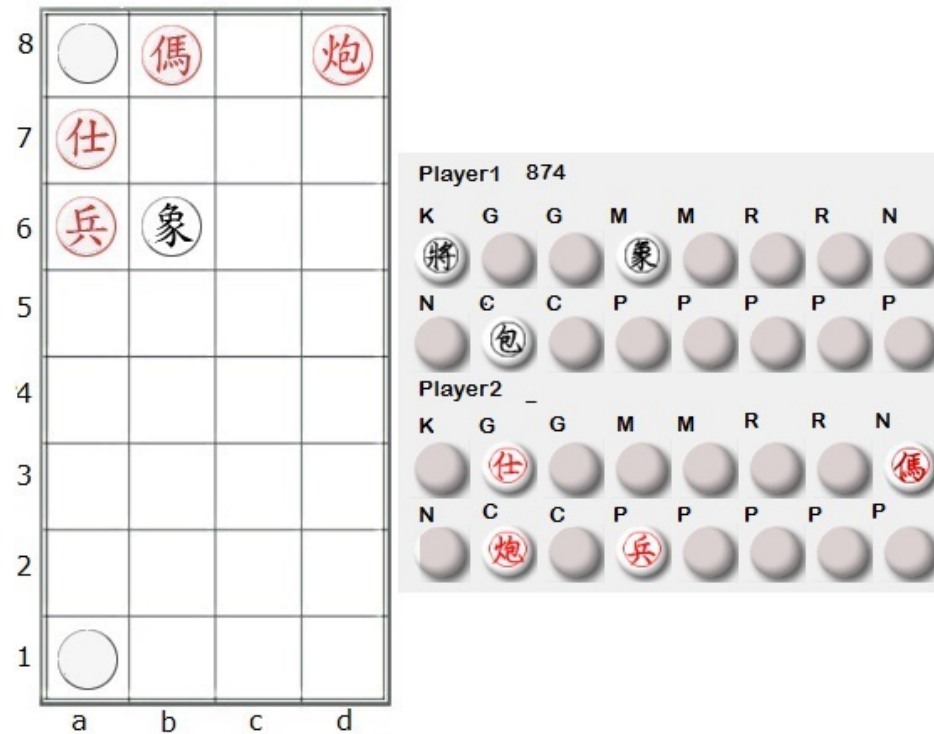
Example (1/3)

- It's black turn and black has 6 different possible legal moves including 4 of them being moving its elephant and two flipping moves at a1 or a8.
 - It is difficult for black to secure a win by moving its elephant.





Example (2/3)

- If black flips a1, then it becomes one of the 2 followings cases.
 - If a1 is black cannon, then black may win.
 - If a1 is black king, then it is difficult for black to win.



Example (3/3)


- If black flips a8, then it becomes one of the 2 followings cases.
 - If a8 is black cannon, then it is difficult for black to win.
 - If a8 is black king, then black may lose.



Basic ideas for searching chance nodes

- Assume a chance node x has a score probability distribution function $Pr(*)$ with the range of possible outcomes from 1 to N where N is a positive integer.
 - For each possible outcome i , we need to compute $score(i)$.
 - The expected value $E = \sum_{i=1}^N score(i) * Pr(x = i)$.
 - The minimum value is $m = \min_{i=1}^N \{score(i) \mid Pr(x = i) > 0\}$.
 - The maximum value is $M = \max_{i=1}^N \{score(i) \mid Pr(x = i) > 0\}$.
- Example: open game in Chinese dark chess.
 - For the first ply, $N = 14 * 32$.
 - ▷ *Using symmetry, we can reduce it to 7*8.*
 - We now consider the chance node of flipping the piece at the cell a1.
 - ▷ $N = 14$.
 - ▷ *Assume $x = 1$ means a black King is revealed and $x = 8$ means a red King is revealed.*
 - ▷ *Then $score(1) = score(8)$ since the first player owns the revealed king no matter its color is.*
 - ▷ $Pr(x = 1) = Pr(x = 8) = 1/14$.

Illustration



Bounds in a chance node

- Assume the various possibilities of a chance node is evaluated one by one in the order that at the end of phase i , $i = N$ is evaluated.
 - Assume $v_{min} \leq score(i) \leq v_{max}$.
- How do the lower and upper bounds, namely m_i and M_i , of the chance node change at the end of phase i ?
 - $i = 0$.
 - ▷ $m_0 = v_{min}$
 - ▷ $M_0 = v_{max}$
 - $i = 1$, we first compute $score(1)$, and then know
 - ▷ $m_1 \geq score(1) * Pr(x = 1) + v_{min} * (1 - Pr(x = 1))$, and
 - ▷ $M_1 \leq score(1) * Pr(x = 1) + v_{max} * (1 - Pr(x = 1))$.
 - ...
 - $i = i^*$, we have computed $score(1), \dots, score(i^*)$, and then know
 - ▷ $m_{i^*} \geq \sum_{i=1}^{i^*} score(i) * Pr(x = i) + v_{min} * (1 - \sum_{i=1}^{i^*} Pr(x = i))$, and
 - ▷ $M_{i^*} \leq \sum_{i=1}^{i^*} score(i) * Pr(x = i) + v_{max} * (1 - \sum_{i=1}^{i^*} Pr(x = i))$.

Algorithm: Chance_Search

■ Algorithm $F2.1'$ (position p , value $alpha$, value $beta$)

// max node

- determine the successor positions p_1, \dots, p_b
- if $b = 0$, then return $f(p)$
 - else begin
 - ▷ $m := alpha$
 - ▷ for $i := 1$ to b do
 - ▷ begin
 - ▷ if p_i is to play a chance node n
then $t := Star1_F2.1'(p_i, n, alpha, beta)$
 - ▷ else $t := G2.1'(p_i, m, beta)$
 - ▷ if $t > m$ then $m := t$
 - ▷ if $m \geq beta$ then return(m) **// beta cut off**
 - ▷ end
- end;
- return m

Algorithm: Chance_Search

- **Algorithm** *Star1_F2.1'*(position p , node n , value $alpha$, value $beta$)
 - **// return the expected value of a chance node n**
 - **determine the possible values of the chance node n to be k_1, \dots, k_c**
 - $m_0 = alpha$; **// current lower bound, $alpha \geq v_{min}$**
 - $M_0 = beta$; **// current upper bound, $beta \leq v_{max}$**
 - $vsum = 0$; **// current expected value**
 - **for $i = 1$ to c do**
 - **begin**
 - ▷ *let p_i be the position of assigning k_i to n in p ;*
 - ▷ $t := G2.1'(p_i, \max\{m_{i-1}, v_{min}\}, \min\{M_{i-1}, v_{max}\})$;
 - ▷ **if $t \leq m_{i-1}$ then $t := alpha$;**
 - ▷ **if $t \geq M_{i-1}$ then $t := beta$;**
 - ▷ $vsum += t * Pr_i$
 - ▷ $m_i = m_{i-1} + (t - alpha) * Pr_i$;
 - ▷ $M_i = M_{i-1} + (t - beta) * Pr_i$;
 - ▷ ...
 - **end**
- **return** $vsum$;

Example: Chinese dark chess

■ Assumption:

- The range of the scores of Chinese dark chess is $[-10, 10]$ inclusive, $\alpha = -10$ and $\beta = 10$.
- $N = 7$.
- $Pr(x = i) = 1/N = 1/7$.

■ Calculation:

- $i = 0$,
 - ▷ $m_0 = -10$.
 - ▷ $M_0 = 10$.
- $i = 1$ and **if** $score(1) = -2$, then
 - ▷ $m_1 = -2 * 1/7 + -10 * 6/7 = -62/7 \simeq -8.86$.
 - ▷ $M_1 = -2 * 1/7 + 10 * 6/7 = 58/7 \simeq 8.26$.
- $i = 1$ and **if** $score(1) = 3$, then
 - ▷ $m_1 = 3 * 1/7 + -10 * 6/7 = -57/7 \simeq -8.14$.
 - ▷ $M_1 = 3 * 1/7 + 10 * 6/7 = 63/7 = 9$.

Comments

- We illustrate the ideas using a fail hard version of the alpha-beta algorithm.
 - Fail hard version has a simple logic in maintaining the search interval.
 - The semantic of comparing an exact returning value with an expected returning value is something that needs careful thinking.
 - May want to pick a chance node with a lower value but having a hope of winning not one with a slightly higher value but having no hope of winning when you are in disadvantageous positions.
 - May want to pick a chance node with a lower value but having no chance of losing, not one with a slightly higher value but having a chance of losing when you are in advantage positions.
- Need to revise algorithms carefully when dealing with the fail sort version or the NegaScout version.
 - What does it mean to combine bounds from a fail soft version?
- Exist other improvements by considering better move orderings involving chance nodes.

How to use these bounds



- The lower and upper bounds of the expected score can be used to do alpha-beta pruning.
 - Nicely fit into the alpha-beta search algorithm.
- Can do better by not searching the DFS order.
 - It is not necessary to search completely the subtree of $x = 1$ first, and then start to look at the subtree of $x = 2$.
 - Assume it is a MAX chance node, e.g., the opponent takes a flip.
 - ▷ *Knowing some value v'_1 of a subtree for $x = 1$ gives an upper bound, i.e., $score(1) \geq v'_1$.*
 - ▷ *Knowing some value v'_2 of a subtree for $x = 2$ gives another upper bound, i.e., $score(2) \geq v'_2$.*
 - ▷ *These bounds can be used to make the search window further narrower.*
- For Monte-Carlo based algorithm, we need to use a sparse sampling algorithm to efficiently estimate the expected value of a chance node [Kearn et al 2002].

Ideas for new search methods

- Consider the case of a 2-player game tree with either 0 or 1 on the leaves.
 - win, or not win which is lose or draw;
 - lose, or not lose which is win or draw;
 - Call this a **binary valued game tree**.
- If the game tree is known as well as the values of some leaves are known, can you make use of this information to search this game tree faster?
 - The value of the root is either 0 or 1.
 - If a branch of the root returns 1, then we know for sure the value of the root is 1.
 - The value of the root is 0 only when all branches of the root returns 0.
 - An AND-OR game tree search.

Which node to search next?

- A **most proving node** for a node u : a node if its value is 1, then the value of u is 1.
- A **most disproving node** for a node u : a node if its value is 0, then the value of u is 0.



Proof or Disproof Number

- Assign a **proof number** and a **disproof number** to each node u in a binary valued game tree.
 - $proof(u)$: the minimum number of leaves needed to visited in order for the value of u to be 1.
 - $disproof(u)$: the minimum number of leaves needed to visited in order for the value of u to be 0.

Proof Number: Definition

- u is a leaf:
 - If $value(u)$ is unknown, then $proof(u)$ is the cost of evaluating u .
 - If $value(u)$ is 1, then $proof(u) = 0$.
 - If $value(u)$ is 0, then $proof(u) = \infty$.
- u is an internal node with all of the children u_1, \dots, u_b :
 - if u is a MAX node,

$$proof(u) = \min_{i=1}^{i=b} proof(u_i);$$

- if u is a MIN node,

$$proof(u) = \sum_{i=1}^{i=b} proof(u_i).$$

Disproof Number: Definition


- u is a leaf:
 - If $value(u)$ is unknown, then $disproof(u)$ is cost of evaluating u .
 - If $value(u)$ is 1, then $disproof(u) = \infty$.
 - If $value(u)$ is 0, then $disproof(u) = 0$.
- u is an internal node with all of the children u_1, \dots, u_b :
 - if u is a MAX node,

$$disproof(u) = \sum_{i=1}^{i=b} disproof(u_i);$$


- if u is a MIN node,

$$disproof(u) = \min_{i=1}^{i=b} disproof(u_i).$$

Illustrations



proof number, disproof number



proof number, disproof number

How to use these Numbers

- If the numbers are known in advance, then from the root, we search a child u with the value equals to $\min\{proof(root), disproof(root)\}$.
 - Then we find a path from the root towards a leaf recursively as follows,
 - ▷ if we try to prove it, then pick a child with the least proof number for a MAX node, and pick any node that has a chance to be proved for a MIN node.
 - ▷ if we try to disprove it, then pick a child with the least disproof number for a MIN node, and pick any node that has a chance to be disproved for a MAX node.
- Assume each leaf takes a lot of time to evaluate.
 - For example, the game tree represents an open game tree or an endgame tree.
 - Depends on the results we have so far, pick the next leaf to prove or disprove.
- Need to be able to update these numbers on the fly.


PN-search: algorithm

- *loop*: **Compute or update proof and disproof numbers for each node in a bottom up fashion.**
 - If $proof(root) = 0$ or $disproof(root) = 0$, then we are done, otherwise
 - ▷ $proof(root) \leq disproof(root)$: we try to prove it.
 - ▷ $proof(root) > disproof(root)$: we try to disprove it.
- $u \leftarrow root$; **{* find the leaf to prove or disprove *}**
 - if we try to prove, then
 - ▷ while u is not a leaf do
 - ▷ if u is a MAX node, then
 - $u \leftarrow$ leftmost child of u with the smallest non-zero proof number;
 - ▷ if current is a MIN node, then
 - $u \leftarrow$ leftmost child of u with a non-zero proof number;
 - if we try to disprove, then
 - ▷ while u is not a leaf do
 - ▷ if u is a MAX node, then
 - $u \leftarrow$ leftmost child of u with a non-zero disproof number;
 - ▷ if current is a MIN node, then
 - $u \leftarrow$ leftmost child of u with the smallest non-zero disproof number;
- **Prove or disprove u ; go to *loop*;**


Multi-Valued game Tree

- The values of the leaves may not be binary.
 - Assume the values are non-negative integers.
 - Note: it can be in any finite countable domain.
- Revision of the proof and disproof numbers.
 - $proof_v(u)$: the minimum number of leaves needed to visited in order for the value of u to $\geq v$.
 - ▷ $proof(u) \equiv proof_1(u)$.
 - $disproof_v(u)$: the minimum number of leaves needed to visited in order for the value of u to $< v$.
 - ▷ $disproof(u) \equiv disproof_1(u)$.

Illustration



Illustration



Multi-Valued Proof Number

- u is a leaf:
 - If $value(u)$ is unknown, then $proof_v(u)$ is cost of evaluating u .
 - If $value(u) \geq v$, then $proof_v(u) = 0$.
 - If $value(u) < v$, then $proof_v(u) = \infty$.
- u is an internal node with all of the children u_1, \dots, u_b :
 - if u is a MAX node,

$$proof_v(u) = \min_{i=1}^{i=b} proof_v(u_i);$$

- if u is a MIN node,

$$proof_v(u) = \sum_{i=1}^{i=b} proof_v(u_i).$$

Multi-valued Disproof Number

- u is a leaf:
 - If $value(u)$ is unknown, then $disproof_v(u)$ is cost of evaluating u .
 - If $value(u) \geq v$, then $disproof_v(u) = \infty$.
 - If $value(u) < v$, then $disproof_v(u) = 0$.
- u is an internal node with all of the children u_1, \dots, u_b :
 - if u is a MAX node,

$$disproof_v(u) = \sum_{i=1}^{i=b} disproof_v(u_i);$$

- if u is a MIN node,

$$disproof_v(u) = \min_{i=1}^{i=b} disproof_v(u_i).$$

Revised PN-search(v): algorithm

- **loop:** Compute or update proof_v and disproof_v numbers for each node in a bottom up fashion.
 - If $\text{proof}_v(\text{root}) = 0$ or $\text{disproof}_v(\text{root}) = 0$, then we are done, otherwise
 - ▷ $\text{proof}_v(\text{root}) \leq \text{disproof}_v(\text{root})$: we try to prove it.
 - ▷ $\text{proof}_v(\text{root}) > \text{disproof}_v(\text{root})$: we try to disprove it.
- $u \leftarrow \text{root}$; **{ * find the leaf to prove or disprove * }**
 - if we try to prove, then
 - ▷ while u is not a leaf do
 - ▷ if u is a MAX node, then
 - $u \leftarrow$ leftmost child of u with the smallest non-zero proof_v number;
 - ▷ if current is a MIN node, then
 - $u \leftarrow$ leftmost child of u with a non-zero proof_v number;
 - if we try to disprove, then
 - ▷ while u is not a leaf do
 - ▷ if u is a MAX node, then
 - $u \leftarrow$ leftmost child of u with a non-zero disproof_v number ;
 - ▷ if current is a MIN node, then
 - $u \leftarrow$ leftmost child of u with the smallest non-zero disproof_v number;
- **Prove or disprove u ; go to loop;**

Multi-valued PN-search: algorithm

- When the values of the leaves are not binary, use an open value binary search to find an upper bound of the value.
 - Set the initial value of v to be 1.
 - *loop*: $\text{PN-search}(v)$
 - ▷ *Prove the value of the search tree is $\geq v$ or disprove it by showing it is $< v$.*
 - If it is proved, then double the value of v and go to *loop* again.
 - If it is disproved, then the true value of the tree is between $\lfloor v/2 \rfloor$ and $v - 1$.
 - **{* Use a binary search to find the exact returned value of the tree. *}**
 - $low \leftarrow \lfloor v/2 \rfloor$; $high \leftarrow v - 1$;
 - **while** $low \leq high$ **do**
 - ▷ *if $low = high$, then return low as the tree value*
 - ▷ $mid \leftarrow \lfloor (low + high)/2 \rfloor$
 - ▷ $\text{PN-search}(mid)$
 - ▷ *if it is disproved, then $high \leftarrow mid - 1$*
 - ▷ *else if it is proved, then $low \leftarrow mid$*

Comments

- Can be used to construct opening books.
- Appears to be good for searching certain types of game trees.
 - Find the easiest way to prove or disprove a conjecture.
 - A dynamic strategy depends on work has been done so far.
- Performance has nothing to do with move ordering.
 - Performance of most previous algorithms depends heavily on whether a good move ordering can be found.
- Searching the “easiest” branch may not give you the best performance.
 - Performance depends on the value of each internal nodes.
- Commonly used in verifying conjectures, e.g., first-player win.
 - Partition the opening moves in a tree-like fashion.
 - Try to the “easiest” way to prove or disprove the given conjecture.
- Take into consideration the fact that some nodes may need more time to process than the other nodes.

References and further readings (1/2)

- L. V. Allis, M. van der Meulen, and H. J. van den Herik. Proof-number search. *Artificial Intelligence*, 66(1):91–124, 1994.
- David Carmel and Shaul Markovitch. Learning and using opponent models in adversary search. Technical Report CIS9609, Technion, 1996.
- M. Campbell. The graph-history interaction: on ignoring position history. In *Proceedings of the 1985 ACM annual conference on the range of computing : mid-80's perspective*, pages 278–280. ACM Press, 1985.

References and further readings (2/2)

- **Bruce W. Ballard** The α -minimax search procedure for trees containing chance nodes **Artificial Intelligence, Volume 21, Issue 3, September 1983, Pages 327-350**
- **Marc Lanctot, Abdallah Saffidine, Joel Veness, Chris Archibald, Mark H. M. Winands** Monte-Carlo α -MiniMax Search Proceedings **IJCAI, pages 580–586, 2013.**
- **Kearns, Michael; Mansour, Yishay; Ng, Andrew Y.** A sparse sampling algorithm for near-optimal planning in large Markov decision processes. **Machine Learning, 2002, 49.2-3: 193-208.**
- **Kuang-che Wu, Shun-Chin Hsu and Tsan-sheng Hsu** "The Graph History Interaction Problem in Chinese Chess," Proceedings of the 11th Advances in Computer Games Conference, (ACG), Springer-Verlag LNCS# 4250, pages 165–179, 2005.