CS/ENGRI 172, Fall 2002
9/13/02: Lecture Seven Handout

Topics: Pruning; non-zero-sum games.
Announcements:
• Reminder: We are seeking a volunteer to help take notes for students with registered disabilities. These students are taking their own set of notes, but the University has recognized their need for "backup". If you're a careful notetaker and regular attendee, it would be great if you could help out! If interested, please contact me by email (for confidentiality reasons).
• New office hours start next week.

Evaluation function example

Here is an example of using an evaluation function on a search-limited (i.e., pruned) version of the game tree from the previous handout. This particular evaluation function happens to be pretty good - usually relatively high scores are given to nodes corresponding to states from which player 1 could potentially win.

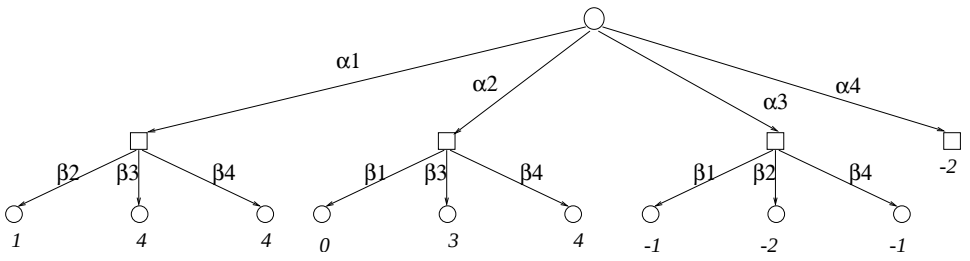

Alpha-beta pruning (The name is unrelated to the operator labels we've been using)

The fundamental principle is that nodes whose (corresponding states') values can't affect decisions "higher up" need not be examined. We will use dfs as a way to organize our examination of nodes and propagation of constraints. Two canonical situations are as follows; the shape of the node indicates whose turn it is to play in the corresponding state, and by each node our current knowledge of its minimax value (or our approximation to it based on our evaluation function as applied to the leaves) is indicated:

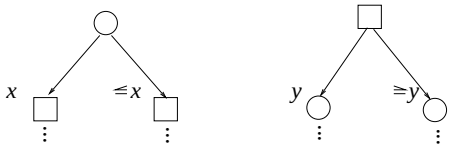

In either situation, it is not necessary to consult or compute the (pseudo-)minimax values of any heretofore unseen nodes in the right nodes' subtrees, because the leftmost operator will (or can) be chosen regardless of their values.

(over)

Prisoners' Dilemma

(A classic non-zero-sum example from game theory attributed to Merrill Flood, Melvin Dresher, and Albert Tucker.) Two criminals have been arrested for a crime that ordinarily would earn each five years in prison. But the state offers each the following deal:

> If you turn state's evidence by pleading guilty whereas your partner stubbornly maintains innocence, we will let you off with just probation and punish your partner with ten years. If both of you plead innocent, we've got enough evidence to send both of you to prison for two years each for a lesser crime.

The *payoff matrix* is as follows, where $(-y_1, -y_2)$ indicates prisoner 1 serves y_1 years and prisoner 2 serves y_2 years.

	p2: innocent	p2: guilty
p1: innocent	$(-2, -2)$	$(-10, 0)$
p1: guilty	$(0, -10)$	$(-5, -5)$