Searching Game Trees

Course: CS40002

Instructor: Dr. Pallab Dasgupta

Department of Computer Science & Engineering

Indian Institute of Technology Kharagpur

Searching Game Trees

- Consider an OR tree with two types of OR nodes, namely Min nodes and Max nodes
- In Min nodes, select the min cost successor
- In Max nodes, select the max cost successor

- Terminal nodes are winning or loosing states
  - It is often infeasible to search up to the terminal nodes
  - We use heuristic costs to compare non-terminal nodes

Shallow and Deep Pruning

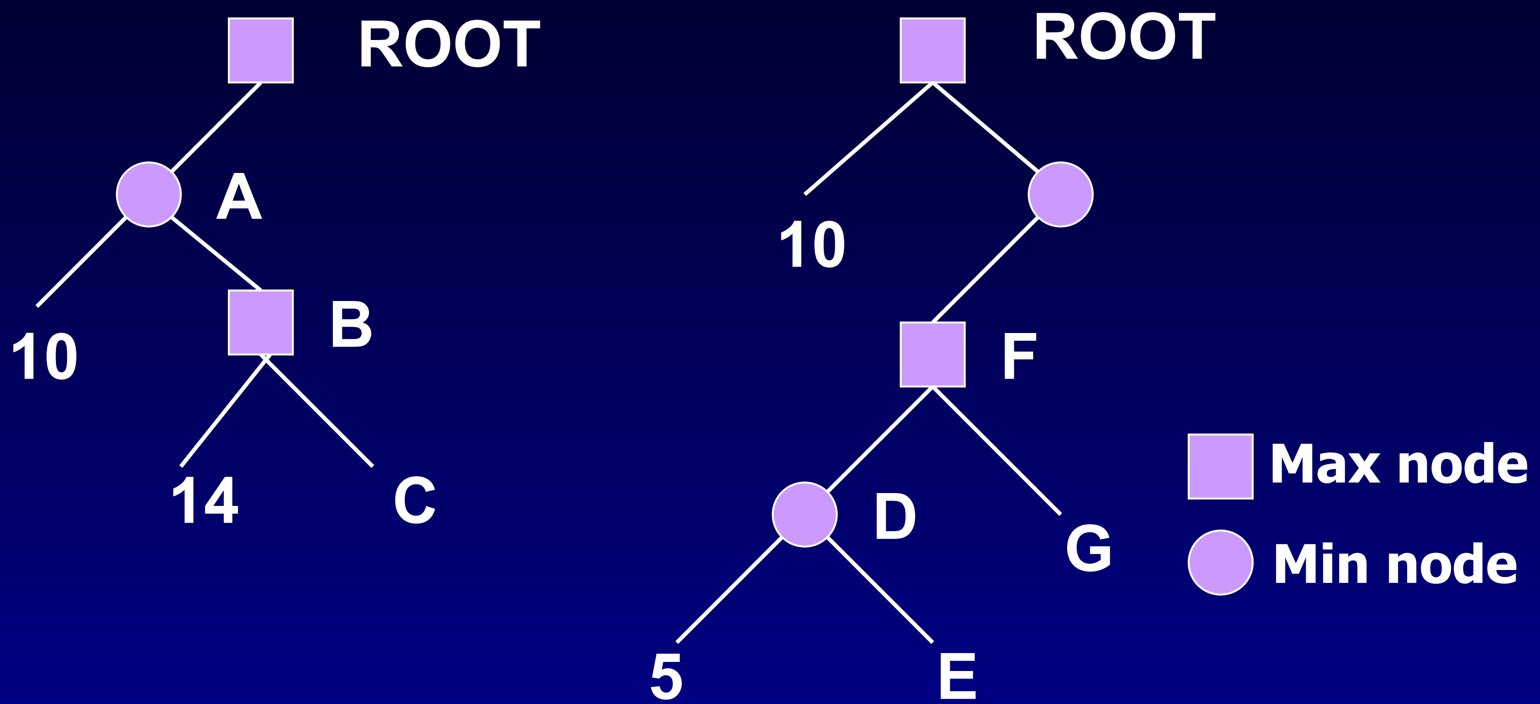

Shallow Cut-off

Deep Cut-off

Alpha-Beta Pruning

- Alpha Bound of J:
  - **The max current val of all MAX ancestors of J**
  - **Exploration of a min node, J, is stopped when its value equals or falls below alpha.**
  - **In a min node, we update beta**
- Beta Bound of J:
  - **The min current val of all MIN ancestors of J**
  - **Exploration of a max node, J, is stopped when its value equals or exceeds beta**
  - **In a max node, we update alpha**
- In both min and max nodes, we return when $\alpha \geq \beta$

Alpha-Beta Procedure: $V(J;\alpha,\beta)$

1. If J is a terminal, return V(J) = h(J).
2. If J is a max node:

   For each successor J_k of J in succession:

   Set $\alpha = \max \{ \alpha, V(J_k; \alpha, \beta) \}$

   If $\alpha \geq \beta$ then return β, else continue

   Return α
3. If J is a min node:

   For each successor J_k of J in succession:

   Set $\beta = \min \{ \beta, V(J_k; \alpha, \beta) \}$

   If $\alpha \geq \beta$ then return α, else continue

   Return β

Search: Additional Topics

Course: CS40002
Instructor: Dr. Pallab Dasgupta

Department of Computer Science & Engineering
Indian Institute of Technology Kharagpur

Iterative Refinement Search

- We iteratively try to improve the solution
  - Consider all states laid out on the surface of a landscape
  - The notion of local and global optima
- Two main approaches
  - Hill climbing / Gradient descent
  - Simulated annealing

Hill Climbing / Gradient Descent

- Makes moves which monotonically improve the quality of solution
- Can settle in a local optima

- Random-restart hill climbing

Simulated Annealing

- Let T denote the temperature. Initially T is high. During iterative refinement, T is gradually reduced to zero.

1. Initialize T
2. If T=0 return current state
3. Set next = a randomly selected succ of current
4. ΔE = Val[next] – Val[current]
5. If $\Delta E > 0$ then Set current = next
6. Otherwise Set current = next with prob $e^{\Delta E/T}$
7. Update T as per schedule and Go To Step 2.

Memory bounded A*: MA*

- Whenever |OPEN ∪ CLOSED| approaches M, some of the least promising states are removed
- To guarantee that the algorithm terminates, we need to back up the cost of the most promising leaf of the subtree being deleted at the root of that subtree
- Many variants of this algorithm have been studied. Recursive Best-First Search (RBFS) is a linear space version of this algorithm

Multi-Objective A*: MOA*

- Adaptation of A* for solving multi-criteria optimization problems
  - Traditional approaches combine the objectives into a single one
  - In multi-objective state space search, the dimensions are retained

- Main concepts:
  - Vector valued state space
  - Vector valued cost and heuristic functions
  - Non-dominated solutions