CSE 4351/5351 Notes 10: Obstacles to Parallelism

Overview

Inherently sequential problems

Discrete event simulation

Two-person games

NP-Completeness (Limit to sequential computation)

Traditionally, these problems are assumed to require exponential time (intractable).

Problems that can be solved in polynomial time are considered ‘‘practical’ (tractable).

Precise definition requires concept of non-determinism (and Turing machines).

Most situations allow use of polynomial-time verification as substitute concept.

NP is closed under ‘‘polynomial time’’ reductions.

Boolean satisfiability was the ‘‘first’’ NP-Complete problem (Stephen Cook, U. Toronto).

Practical description: Is there a way to set the inputs to a digital circuit such that an output LED glows?

Early work by Richard Karp (Berkeley) provides many NP-Complete problems that are useful for reductions.

Open question: P = NP?

Book: Garey & Johnson

P-Completeness (Limit to massively parallel computation)

PRAM theory suggests that problems solvable in polylogarithmic time on a polynomial number of processors (class NC) have the best potential (‘‘massively parallel’’) for parallel solution.

P-Complete problems are in P, but are assumed to not be in NC.

Theoretical definition: A problem in P is P-complete if all other problems in P are logspace-reducible to this problem.

The ‘‘path systems’’ problem was the first P-complete problem (S. Cook).

A less esoteric P-Complete problem - Boolean circuit evaluation: Given input values for a Boolean circuit, does the output LED glow?

Other P-Complete problems:

- Variations on Boolean circuit evaluation.
- Many variations on breadth-first, depth-first search, and problems with optimal greedy solutions.
- Linear programming.
- Maximum network flow.
- Gaussian elimination with partial pivoting

Open question: P = NC?

Book: Greenlaw, Hoover, & Ruzzo

Discrete Event Simulation

Sequential Simulation

State of simulation

Events

Priority queue (event list)

Causality errors are avoided

Termination is obvious

Applications: communication networks, circuits, military

Conservative approach (Chandy & Misra, Bryant)

Causality errors are still avoided

Distributed model - limited communication channels between logical processes

Null messages to accelerate progress - ‘‘promises’’ that all future messages from this logical process will have a timestamp past some time (abundance of these can be a problem)

Optimistic approach (Jefferson & Fujimoto - Time Warp/Virtual Time)

Local time vs. global time

Small number of potential causality errors in real simulations

Window for time ahead of global time

Causality errors lead to (expensive) rollbacks and anti-messages

Exceptional storage issues (garbage collection) for input, output, and previous messages

Worst-case: If you can tune optimistic approach, then you already understand simulated world.

Two-Person Games

Min-max α-β tree

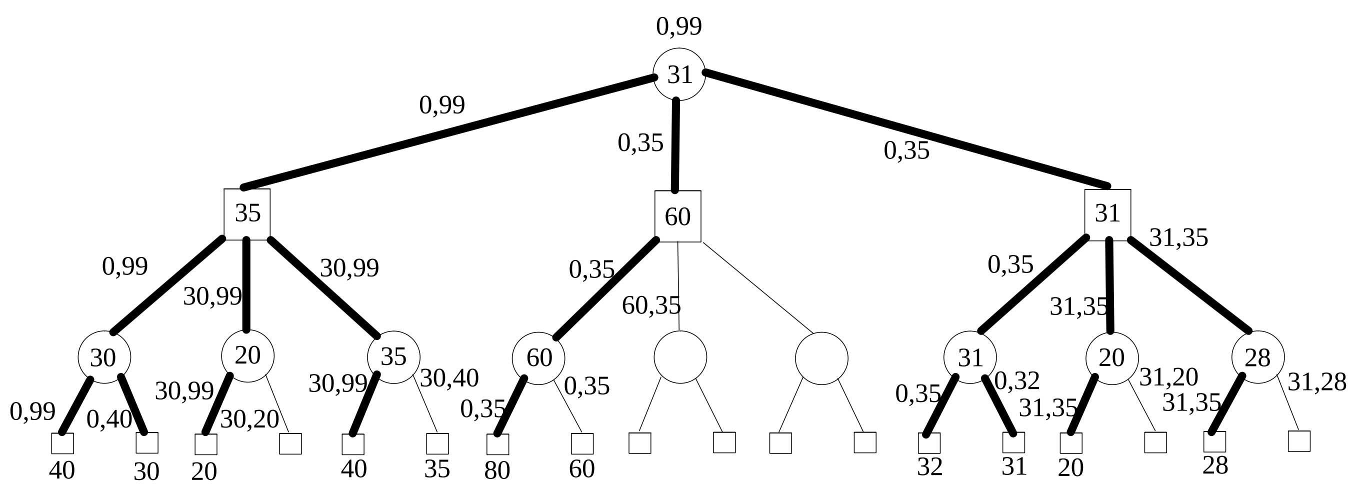

Minimizing nodes modify the upper bound in the following way:

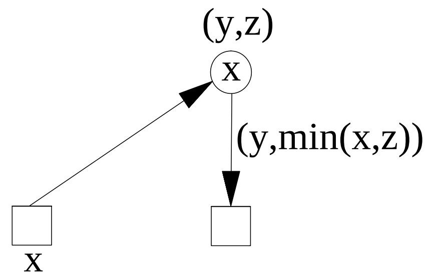

Maximizing nodes modify the lower bound in the following way:

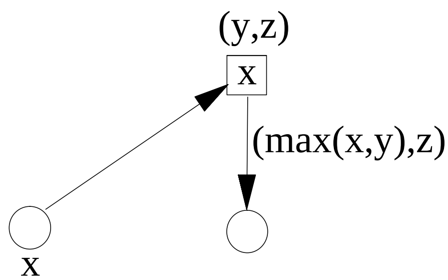

In either case, if (a,b) does not have a < b, then the branch is pruned.

Parallelizing

Move generation

Iterative deepening - adding each successive level improves search bounds and helps with move ordering

Window (aspiration) search

Each window has a narrow range for search bounds

Window with final result within its bounds is the one with the best move

Significant effort to tune to avoid load imbalance

Ordering - minimum-to-maximum, maximum-to minimum

Concurrent transposition table