

UNIT V - COPING WITH THE LIMITATIONS OF ALGORITHM POWER

5.1 LIMITATIONS OF ALGORITHM POWER

There are many algorithms for solving a variety of different problems. They are very powerful instruments, especially when they are executed by modern computers.

The power of algorithms is limited because of the following reasons:

- There are some problems cannot be solved by any algorithm.
- There are some problems can be solved algorithmically but not in polynomial time.
- There are some problems can be solved in polynomial time by some algorithms, but they are usually lower bounds on their efficiency.

Algorithms limits are identified by the following:

- Lower-Bound Arguments
- Decision Trees
- P, NP and NP-Complete Problems

5.2 LOWER-BOUND ARGUMENTS

We can look at the efficiency of an algorithm two ways. We can establish its **asymptotic efficiency class** (say, for the worst case) and see where this class stands with respect to the **hierarchy of efficiency classes**.

For example, selection sort, whose efficiency is quadratic, is a reasonably fast algorithm, whereas the algorithm for the Tower of Hanoi problem is very slow because its efficiency is exponential.

Lower bounds means estimating the minimum amount of work needed to solve the problem. We present several methods for establishing lower bounds and illustrate them with specific examples.

1. Trivial Lower Bounds
2. Information-Theoretic Arguments
3. Adversary Arguments
4. Problem Reduction

In analyzing the efficiency of specific algorithms in the preceding, we should distinguish between a lower-bound class and a minimum number of times a particular operation needs to be executed.

Trivial Lower Bounds

The simplest method of obtaining a lower-bound class is based on counting the number of items in the problem's **input** that must be **processed** and the number of **output** items that need to be **produced**.

Since any algorithm must at least “read” all the items it needs to process and “write” all its outputs, such a count yields a **trivial lower bound**.

For example, any algorithm for generating all permutations of n distinct items must be in $\Omega(n!)$ because the size of the output is $n!$. And this bound is **tight** because good algorithms for generating permutations spend a constant time on each of them except the initial one.

Consider the problem of **evaluating a polynomial of degree n** at a given point x , given its coefficients a_n, a_{n-1}, \dots, a_0 . $p(x) = a_n x^n + a_{n-1} x^{n-1} + \dots + a_0$. All the coefficients have to be processed by any polynomial-evaluation algorithm. i.e **$\Omega(n)$** . This is tight lower bound.

Similarly, a trivial lower bound for computing **the product of two $n \times n$ matrices** is **$\Omega(n^2)$** because any such algorithm has to process $2n^2$ elements in the input matrices and generate n^2 elements of the product. It is still unknown, however, whether this bound is tight.

The trivial bound for the **traveling salesman problem** is **$\Omega(n^2)$** , because its input is $n(n-1)/2$ intercity distances and its output is a list of $n + 1$ cities making up an optimal tour. But this bound is useless because there is no known algorithm with the running time being a polynomial function.

Determining the lower bound lies in **which part of an input must be processed** by any algorithm solving the problem. For example, searching for an element of a given value in a sorted array does not require processing all its elements.

Information-Theoretic Arguments

The information-theoretical approach seeks to establish a lower bound based on **the amount of information it has to produce** by algorithm.

Consider an example “**Game of guessing number**”, the well-known game of deducing a positive integer between 1 and n selected by somebody by asking that person questions with yes/no answers. The amount of uncertainty that any algorithm solving this problem has to resolve can be measured by $\lceil \log_2 n \rceil$.

The number of bits needed to specify a particular number among the n possibilities. Each answer to the question gives information about each bit.

1. Is the first bit zero? → No → first bit is 1
2. Is the second bit zero? → Yes → second bit is 0
3. Is the third bit zero? → Yes → third bit is 0
4. Is the fourth bit zero? → Yes → fourth bit is 0

The number in binary is 1000, i.e. 8 in decimal value.

The above approach is called the *information-theoretic argument* because of its connection to information theory. This is useful for finding *information-theoretic lower bounds* for many problems involving comparisons, including sorting and searching.

Its underlying idea can be realized the mechanism of *decision trees*. Because

Adversary Arguments

Adversary Argument is a method of proving by **playing a role of adversary (opponent)** in which algorithm has to work more for **adjusting input** consistently.

Consider the Game of guessing number between positive integer 1 and n by asking a person (Adversary) with yes/no type answers for questions. After each question at least one-half of the numbers reduced. If an algorithm stops before the size of the set is reduced to 1, the adversary can exhibit a number.

Any algorithm needs $\lceil \log_2 n \rceil$ iterations to shrink an n -element set to a one-element set by halving and rounding up the size of the remaining set. Hence, at least $\lceil \log_2 n \rceil$ questions need to be asked by any algorithm in the worst case. This example illustrates the *adversary method* for establishing lower bounds.

Consider the problem of **merging two sorted lists** of size n $a_1 < a_2 < \dots < a_n$ and $b_1 < b_2 < \dots < b_n$ into a single sorted list of size $2n$. For simplicity, we assume that all the a 's and b 's are distinct, which gives the problem a unique solution.

Merging is done by repeatedly comparing the first elements in the remaining lists and outputting the smaller among them. The number of key comparisons (lower bound) in the worst case for this algorithm for merging is $2n - 1$.

Problem Reduction

Problem reduction is a method in which a difficult unsolvable problem P is reduced to another solvable problem B which can be solved by a known algorithm.

A similar reduction idea can be used for finding a lower bound. To show that problem P is at least as hard as another problem Q with a known lower bound, we need to reduce Q to P (not P to Q !). In other words, we should show that an arbitrary instance of problem Q can be transformed to an instance of problem P , so any algorithm solving P would solve Q as well. Then a lower bound for Q will be a lower bound for P . Table 5.1 lists several important problems that are often used for this purpose.

TABLE 5.1 Problems often used for establishing lower bounds by problem reduction

Problem	Lower bound	Tightness
Sorting	$\Omega(n \log n)$	yes
searching in a sorted array	$\Omega(\log n)$	yes
element uniqueness problem	$\Omega(n \log n)$	yes
multiplication of n-digit integers	$\Omega(n)$	unknown
multiplication of $n \times n$ matrices	$\Omega(n^2)$	unknown

Consider the Euclidean minimum spanning tree problem as an example of establishing a lower bound by reduction:

Given n points in the Cartesian plane, construct a tree of minimum total length whose vertices are the given points. As a problem with a known lower bound, we use the element uniqueness problem.

We can transform any set x_1, x_2, \dots, x_n of n real numbers into a set of n points in the Cartesian plane by simply adding 0 as the points' y coordinate: $(x_1, 0), (x_2, 0), \dots, (x_n, 0)$. Let T be a minimum spanning tree found for this set of points. Since T must contain a shortest edge, checking whether T contains a zero length edge will answer the question about uniqueness of the given numbers. This reduction implies that $\Omega(n \log n)$ is a lower bound for the Euclidean minimum spanning tree problem,

Note: Limitations of algorithm can be studied by obtaining lower bound efficiency.

5.3 DECISION TREES

Important algorithms like sorting and searching are based on comparing items of their inputs. The study of the performance of such algorithm is called a **decision tree**. As an example, Figure 5.1 presents a decision tree of an algorithm for finding a minimum of three numbers. Each internal node of a binary decision tree represents a key comparison indicated in the node.

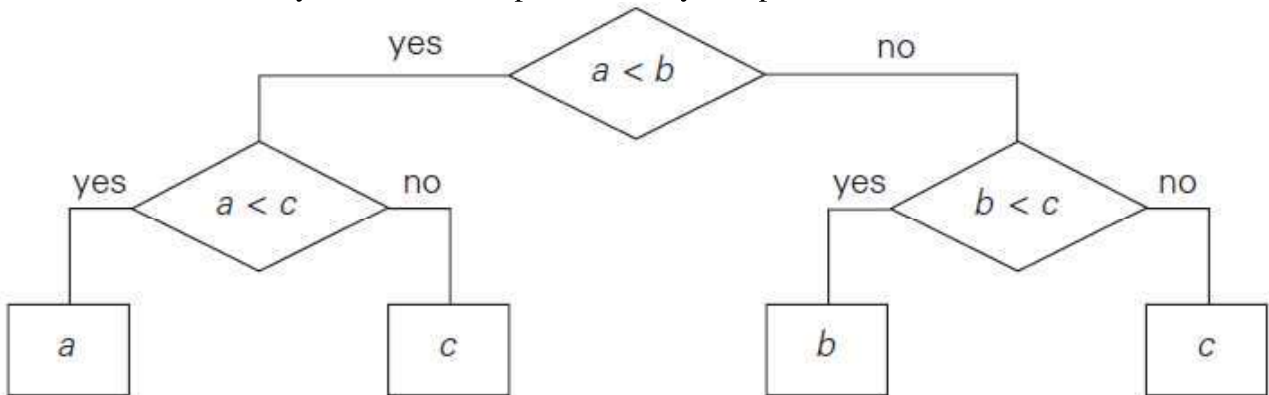
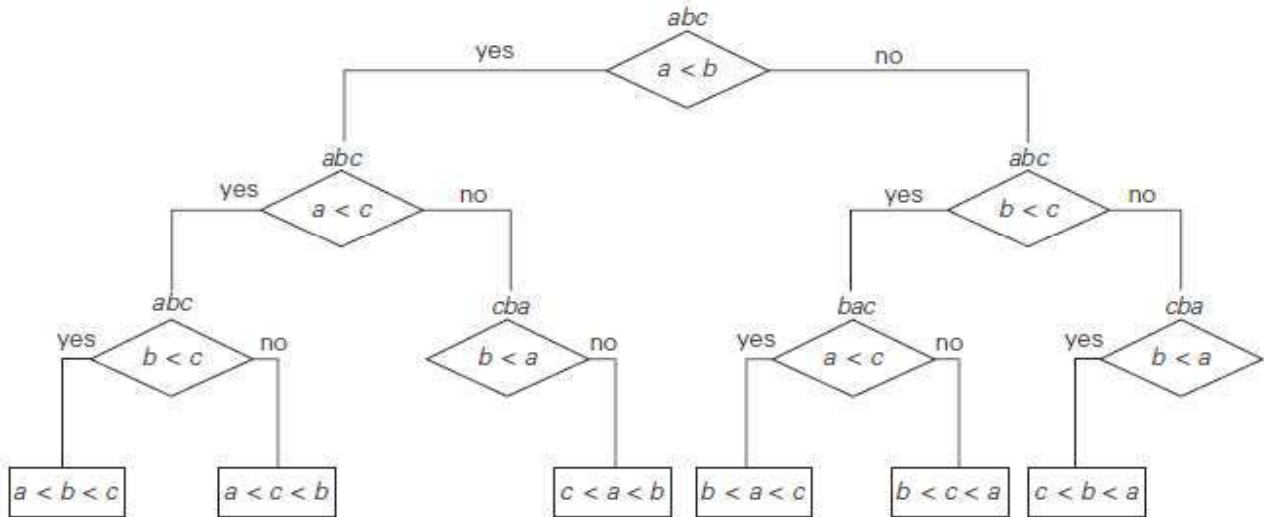


FIGURE 5.1 Decision tree for finding a minimum of three numbers.

Consider a binary decision tree with height h and leaves n . and height h , then $h \geq \lceil \log_2 n \rceil$. A binary tree of height h with the largest number of leaves on the last level is 2^h . In other words, $2^h \geq n$, which puts a lower bound on the heights of binary decision trees. Hence the worst-case number of comparisons made by any comparison-based algorithm for the problem is called the information theoretic lower bound.

Decision Trees for Sorting



C b a
1 2 3

FIGURE 5.2 Decision tree for the tree-element selection sort.

A triple above a node indicates the state of the array being sorted. Note two redundant comparisons $b < a$ with a single possible outcome because of the results of some previously made comparisons.

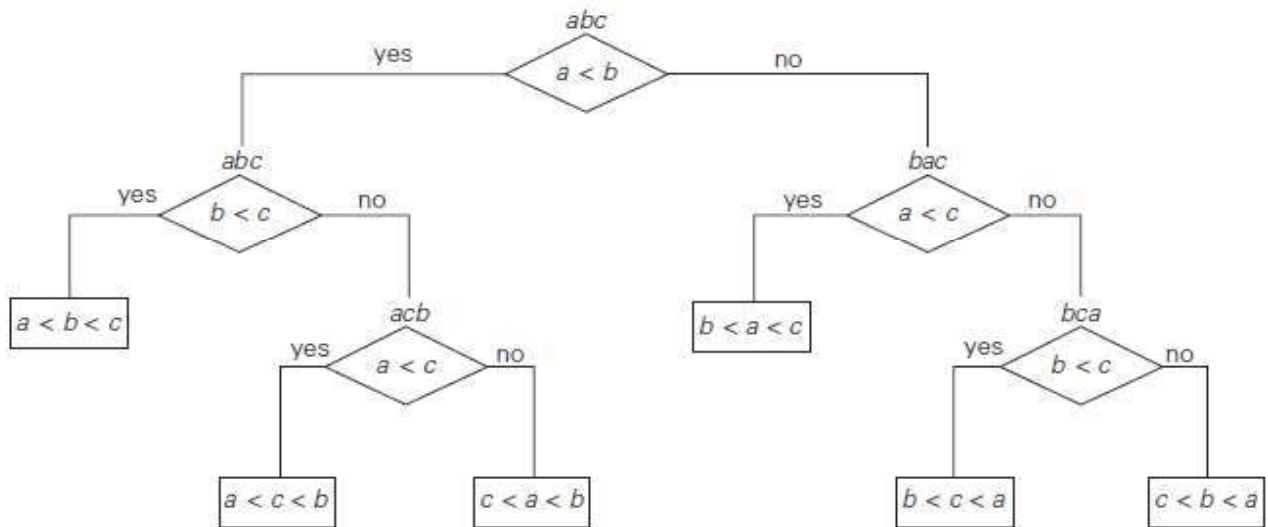


FIGURE 5.3 Decision tree for the three-element insertion sort.

The three-element insertion sort whose decision tree is given in Figure 5.3, this number is $(2 + 3 + 3 + 2 + 3 + 3)/6 = 2.66$. Under the standard assumption that all $n!$ outcomes of sorting are equally likely, the following lower bound on the average number of comparisons C_{avg} made by any comparison-based algorithm in sorting an n -element list has been proved:

$$C_{avg}(n) \geq \log_2 n!$$

Decision tree is a convenient model of algorithms involving comparisons in which

- internal nodes represent comparisons
- leaves represent outcomes (or input cases)

Decision Trees and Sorting Algorithms

- Any comparison-based sorting algorithm can be represented by a decision tree (for each fixed n)
- Number of leaves (outcomes) $\geq n!$

- Height of binary tree with $n!$ leaves $\geq \lceil \log_2 n! \rceil$
- Minimum number of comparisons in the worst case $\geq \lceil \log_2 n! \rceil$ for any comparison-based sorting algorithm, since the longest path represents the worst case and its length is the height
- $\lceil \log_2 n! \rceil \approx n \log_2 n$ (by Sterling approximation)
- This lower bound is tight (mergesort or heapsort)

Decision Trees for Searching a Sorted Array

Decision trees can be used for establishing lower bounds on the number of key comparisons in searching a sorted array of n keys: $A[0] < A[1] < \dots < A[n - 1]$.

The principal algorithm for this problem is binary search. The number of comparisons made by binary search in the worst case, $C_{\text{worst}}(n)$, is given by the formula

$$C_{\text{worst}}(n) = \lfloor \log_2 n \rfloor + 1 = \lceil \log_2(n + 1) \rceil$$

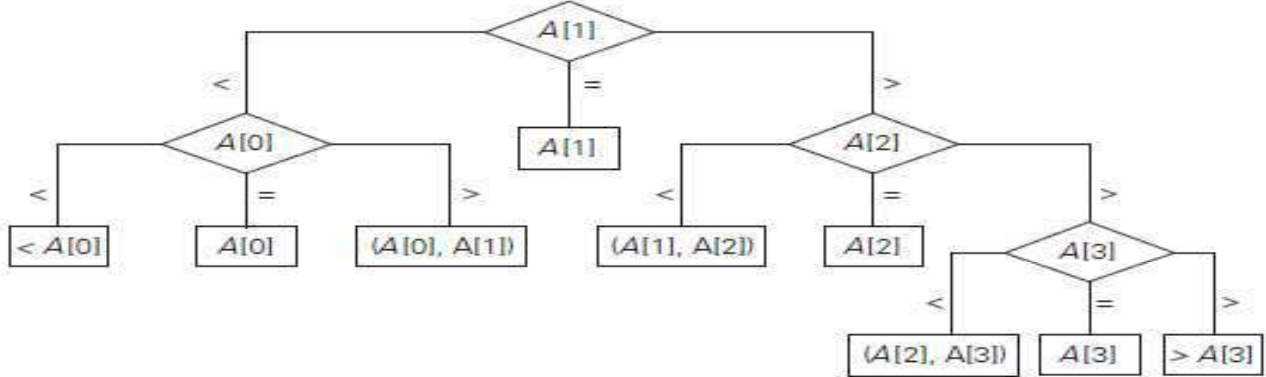


FIGURE 5.4 Ternary decision tree for binary search in a four-element array.

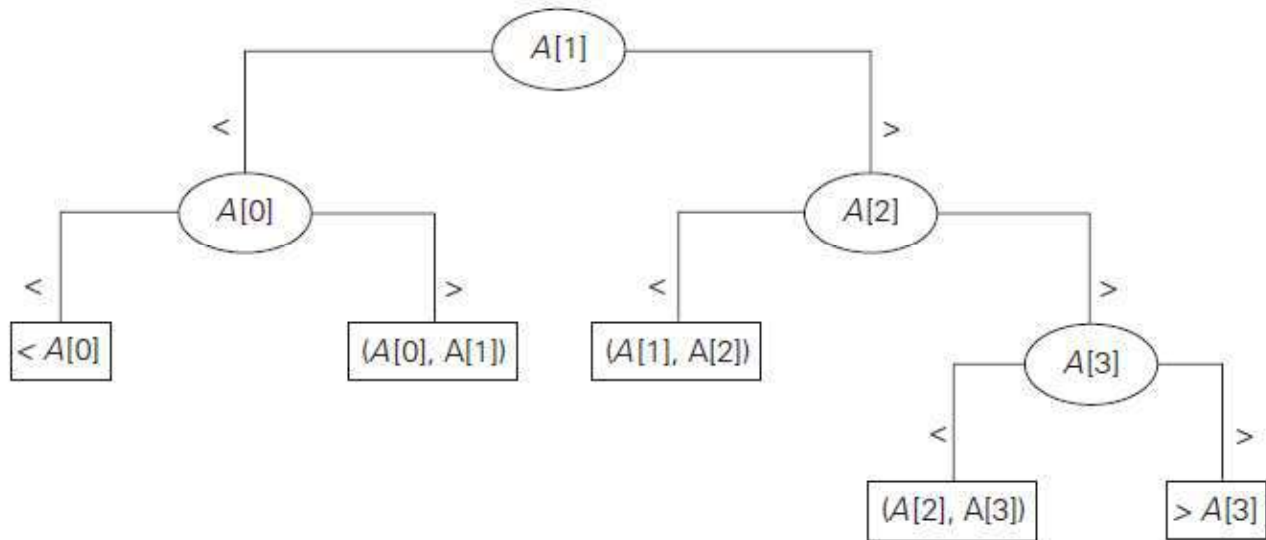


FIGURE 5.5 Binary decision tree for binary search in a four-element array.

As comparison of the decision trees in the above illustrates, the binary decision tree is simply the ternary decision tree with all the middle subtrees eliminated. Applying inequality to such binary decision trees immediately yields $C_{\text{worst}}(n) \geq \lceil \log_2(n + 1) \rceil$

5.4 P, NP AND NP-COMPLETE PROBLEMS

Problems that can be solved in polynomial time are called *tractable*, and problems that cannot be solved in polynomial time are called *intractable*.

There are several **reasons for intractability**.

- **First**, we **cannot solve** arbitrary instances of intractable problems in a reasonable amount of time unless such **instances are very small**.
- **Second**, although there might be a huge difference between the running times in $O(p(n))$ for polynomials of **drastically different degrees**, where $p(n)$ is a polynomial of the problem's input size n .
- **Third**, polynomial functions possess many convenient properties; in particular, both the sum and composition of two polynomials are **always polynomials too**.
- **Fourth**, the choice of this class has led to a development of an extensive theory called *computational complexity*.

Definition: Class **P** is a class of decision problems that can be solved in polynomial time by deterministic algorithms. This class of problems is called *polynomial class*.

- Problems that can be solved in polynomial time as the set that computer science theoreticians call **P**. A more formal definition includes in **P** only **decision problems**, which are problems with **yes/no** answers.
- The class of decision problems that are solvable in $O(p(n))$ **polynomial time**, where $p(n)$ is a polynomial of problem's input size n

Examples:

- Searching
- Element uniqueness
- Graph connectivity
- Graph acyclicity
- Primality testing (finally proved in 2002)
- **The restriction of P** to decision problems can be justified by the following reasons.
 - First, it is sensible to **exclude problems not solvable in polynomial time** because of their exponentially large output. e.g., generating subsets of a given set or all the permutations of n distinct items.
 - Second, **many important problems that are not decision problems** in their most natural formulation can be reduced to a series of decision problems that are easier to study. For example, instead of asking about the minimum number of colors needed to color the vertices of a graph so that no two adjacent vertices are colored the same color. Coloring of the graph's vertices with no more than m colors for $m = 1, 2, \dots$ (The latter is called the **m-coloring problem**.)
 - So, every decision problem can not be solved in polynomial time. Some **decision problems** cannot be solved at all by any algorithm. Such problems are called **undecidable**, as opposed to **decidable** problems that can be solved by an algorithm (**Halting problem**).
- **Non polynomial-time algorithm:** There are many important problems, however, for which no polynomial-time algorithm has been found.
 - **Hamiltonian circuit problem:** Determine whether a given graph has a Hamiltonian circuit—a path that starts and ends at the same vertex and passes through all the other vertices exactly once.
 - **Traveling salesman problem:** Find the shortest tour through n cities with known positive integer distances between them (find the shortest Hamiltonian circuit in a complete graph with positive integer weights).

- **Knapsack problem:** Find the most valuable subset of n items of given positive integer weights and values that fit into a knapsack of a given positive integer capacity.
- **Partition problem:** Given n positive integers, determine whether it is possible to partition them into two disjoint subsets with the same sum.
- **Bin-packing problem:** Given n items whose sizes are positive rational numbers not larger than 1, put them into the smallest number of bins of size 1.
- **Graph-coloring problem:** For a given graph, find its chromatic number, which is the smallest number of colors that need to be assigned to the graph's vertices so that no two adjacent vertices are assigned the same color.
- **Integer linear programming problem:** Find the maximum (or minimum) value of a linear function of several integer-valued variables subject to a finite set of constraints in the form of linear equalities and inequalities.

Definition: A **nondeterministic algorithm** is a two-stage procedure that takes as its input an instance I of a decision problem and does the following.

1. **Nondeterministic (“guessing”) stage:** An arbitrary string S is generated that can be thought of as a candidate solution to the given instance.
2. **Deterministic (“verification”) stage:** A deterministic algorithm takes both I and S as its input and outputs yes if S represents a solution to instance I . (If S is not a solution to instance I , the algorithm either returns no or is allowed not to halt at all.)

Finally, a nondeterministic algorithm is said to be **nondeterministic polynomial** if the time efficiency of its verification stage is polynomial.

Definition: Class **NP** is the class of decision problems that can be solved by nondeterministic polynomial algorithms. This class of problems is called **nondeterministic polynomial**.

Most decision problems are in NP. First of all, this class includes all the problems in P:

$$P \subseteq NP$$

This is true because, if a problem is in P, we can use the deterministic polynomial time algorithm that solves it in the verification-stage of a nondeterministic algorithm that simply ignores string S generated in its nondeterministic (“guessing”) stage. But NP also contains the Hamiltonian circuit problem, the partition problem, decision versions of the traveling salesman, the knapsack, graph coloring, and many hundreds of other difficult combinatorial optimization. The halting problem, on the other hand, is among the rare examples of decision problems that are known not to be in NP.

Note that $P = NP$ would imply that each of many hundreds of difficult combinatorial decision problems can be solved by a polynomial-time algorithm.

Definition: A decision problem $D1$ is said to be **polynomially reducible** to a decision problem $D2$, if there exists a function t that transforms instances of $D1$ to instances of $D2$ such that:

1. t maps all yes instances of $D1$ to yes instances of $D2$ and all no instances of $D1$ to no instances of $D2$.
2. t is computable by a polynomial time algorithm.

This definition immediately implies that if a problem $D1$ is polynomially reducible to some problem $D2$ that can be solved in polynomial time, then problem $D1$ can also be solved in polynomial time

Definition: A decision problem D is said to be **NP-complete** if it is hard as any problem in NP.

1. It belongs to class **NP**
2. Every problem in **NP** is polynomially reducible to D

The fact that closely related decision problems are polynomially reducible to each other is not very surprising. For example, let us prove that the Hamiltonian circuit problem is polynomially reducible to the decision version of the traveling salesman problem.




FIGURE 5.6 Polynomial-time reductions of *NP* problems to an *NP*-complete problem

Theorem: A decision problem is said to be *NP-complete* if it is hard as any problem in *NP*.

Proof: Let us prove that the Hamiltonian circuit problem is polynomially reducible to the decision version of the traveling salesman problem.

We can map a graph *G* of a given instance of the Hamiltonian circuit problem to a complete weighted graph *G'* representing an instance of the traveling salesman problem by assigning 1 as the weight to each edge in *G* and adding an edge of weight 2 between any pair of nonadjacent vertices in *G*. As the upper bound *m* on the Hamiltonian circuit length, we take $m = n$, where *n* is the number of vertices in *G* (and *G'*). Obviously, this transformation can be done in polynomial time.

Let *G* be a yes instance of the Hamiltonian circuit problem. Then *G* has a Hamiltonian circuit, and its image in *G'* will have length *n*, making the image a yes instance of the decision traveling salesman problem.

Conversely, if we have a Hamiltonian circuit of the length not larger than *n* in *G'*, then its length must be exactly *n* and hence the circuit must be made up of edges present in *G*, making the inverse image of the yes instance of the decision traveling salesman problem be a yes instance of the Hamiltonian circuit problem.

This completes the proof.

Theorem: State and prove Cook's theorem.

Prove that CNF-sat is *NP*-complete.

Satisfiability of boolean formula for three conjunctive normal form is *NP*-Complete.

NP problems obtained by polynomial-time reductions from a *NP*-complete problem

Proof: The notion of *NP*-completeness requires, however, polynomial reducibility of *all* problems in *NP*, both known and unknown, to the problem in question. Given the bewildering variety of decision problems, it is nothing short of amazing that specific examples of *NP*-complete problems have been actually found.

Nevertheless, this mathematical feat was accomplished independently by Stephen Cook in the United States and Leonid Levin in the former Soviet Union. In his 1971 paper, Cook [Coo71] showed that the so-called **CNF-satisfiability problem** is *NP*complete.

x_1	x_2	x_3	\bar{x}_1	\bar{x}_2	\bar{x}_3	$x_1 \vee \bar{x}_2 \vee \bar{x}_3$	$\bar{x}_1 \vee x_2$	$\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3$	$(x_1 \vee \bar{x}_2 \vee \bar{x}_3) \wedge (\bar{x}_1 \vee x_2) \wedge (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3)$
T	T	T	F	F	F	T	T	F	F
T	T	F	F	F	T	T	T	T	T
T	F	T	F	T	F	T	F	T	F
T	F	F	F	T	T	T	F	T	F

F	T	T	T	F	F	F	T	T	F
F	T	F	T	F	T	T	T	T	T
F	F	T	T	T	F	T	T	T	T
F	F	F	T	T	T	T	T	T	T

The CNF-satisfiability problem deals with boolean expressions. Each boolean expression can be represented in conjunctive normal form, such as the following expression involving three boolean variables x_1 , x_2 , and x_3 and their negations denoted \bar{x}_1 , \bar{x}_2 , and \bar{x}_3 respectively:

$$(x_1 \vee \bar{x}_2 \vee \bar{x}_3) \& (\bar{x}_1 \vee x_2) \& (\bar{x}_1 \vee \bar{x}_2 \vee \bar{x}_3)$$

The CNF-satisfiability problem asks whether or not one can assign values *true* and *false* to variables of a given boolean expression in its CNF form to make the entire expression *true*. (It is easy to see that this can be done for the above formula: if $x_1 = \text{true}$, $x_2 = \text{true}$, and $x_3 = \text{false}$, the entire expression is *true*.)

Since the Cook-Levin discovery of the first known *NP*-complete problems, computer scientists have found many hundreds, if not thousands, of other examples. In particular, the well-known problems (or their decision versions) mentioned above—Hamiltonian circuit, traveling salesman, partition, bin packing, and graph coloring—are all *NP*-complete. It is known, however, that if $P \neq NP$ there must exist *NP* problems that neither are in *P* nor are *NP*-complete.

Showing that a decision problem is *NP*-complete can be done in two steps.

1. First, one needs to show that the problem in question is in *NP*; i.e., a randomly generated string can be checked in polynomial time to determine whether or not it represents a solution to the problem. Typically, this step is easy.
2. The second step is to show that every problem in *NP* is reducible to the problem in question in polynomial time. Because of the transitivity of polynomial reduction, this step can be done by showing that a known *NP*-complete problem can be transformed to the problem in question in polynomial time.

The definition of *NP*-completeness immediately implies that if there exists a deterministic polynomial-time algorithm for just one *NP*-complete problem, then every problem in *NP* can be solved in polynomial time by a deterministic algorithm, and hence $P = NP$.




FIGURE 5.7 NP-completeness by reduction

Examples: TSP, knapsack, partition, graph-coloring and hundreds of other problems of combinatorial nature $P = NP$ would imply that every problem in *NP*, including all *NP*-complete problems, could be solved in polynomial time. If a polynomial-time algorithm for just one *NP*-complete problem is discovered, then every problem in *NP* can be solved in polynomial time, i.e. $P = NP$. Most but not all researchers believe that $P \neq NP$, i.e. *P* is a proper subset of *NP*. If $P \neq NP$, then the *NP*-complete problems are not in *P*, although many of them are very useful in practice.

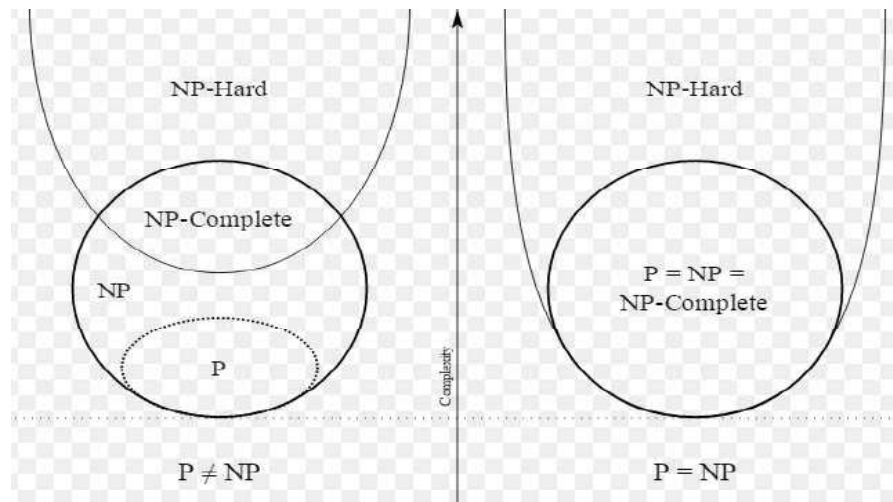


FIGURE 5.8 Relation among P, NP, NP-hard and NP Complete problems

5.5 COPING WITH THE LIMITATIONS OF ALGORITHM POWER

There are some problems that are difficult to solve algorithmically. At the same time, few of them are so important, we must solve by some other technique. Two algorithm design techniques ***backtracking*** and ***branch-and-bound*** that often make it possible to solve at least some large instances of difficult combinatorial problems.

Both backtracking and branch-and-bound are based on the construction of a state-space tree whose nodes reflect specific choices made for a solution's components. Both techniques terminate a node as soon as it can be guaranteed that no solution to the problem can be obtained by considering choices that correspond to the node's descendants

We consider a few approximation algorithms for solving the Assignment Problem, traveling salesman and knapsack problems. There are three classic methods like the bisection method, the method of false position, and Newton's method for approximate root finding.

Exact Solution Strategies are given below:

Exhaustive search (brute force)-

- useful only for small instances

Dynamic programming

- applicable to some problems (e.g., the knapsack problem)

Backtracking

- eliminates some unnecessary cases from consideration
- yields solutions in reasonable time for many instances but worst case is still exponential

Branch-and-bound

- further refines the backtracking idea for optimization problems

Coping with the Limitations of Algorithm Power are given below:

Backtracking

- *n*-Queens Problem
- Hamiltonian Circuit Problem
- Subset-Sum Problem

Branch-and-Bound

- Assignment Problem
- Knapsack Problem
- Traveling Salesman Problem

Approximation Algorithms for *NP*-Hard Problems

- Approximation Algorithms for the Traveling Salesman Problem
- Approximation Algorithms for the Knapsack Problem

Algorithms for Solving Nonlinear Equations

- Bisection Method
- False Position Method
- Newton's Method

5.6 BACKTRACKING

- Backtracking is a more intelligent variation approach.
- The principal idea is to construct solutions one component at a time and evaluate such partially constructed candidates as follows.
- If a partially constructed solution can be developed further without violating the problem's constraints, it is done by taking the first remaining legitimate option for the next component.
- If there is no legitimate option for the next component, no alternatives for any remaining component need to be considered. In this case, the algorithm backtracks to replace the last component of the partially constructed solution with its next option.
- It is convenient to implement this kind of processing by constructing a tree of choices being made, called **the state-space tree**.
- Its root represents an initial state before the search for a solution begins.
- The nodes of the first level in the tree represent the choices made for the first component of a solution, the nodes of the second level represent the choices for the second component, and so on.
- A node in a state-space tree is said to be promising if it corresponds to a partially constructed solution that may still lead to a complete solution. otherwise, it is called **nonpromising**.
- Leaves represent either nonpromising dead ends or complete solutions found by the algorithm. In the majority of cases, a statespace tree for a backtracking algorithm is constructed in the manner of depthfirst search.
- If the current node is promising, its child is generated by adding the first remaining legitimate option for the next component of a solution, and the processing moves to this child. If the current node turns out to be nonpromising, the algorithm backtracks to the node's parent to consider the next possible option for its last component; if there is no such option, it backtracks one more level up the tree, and so on.
- Finally, if the algorithm reaches a complete solution to the problem, it either stops (if just one solution is required) or continues searching for other possible solutions.
- Backtracking techniques are applied to solve the following problems
 - *n*-Queens Problem
 - Hamiltonian Circuit Problem
 - Subset-Sum Problem

5.7 N-QUEENS PROBLEM

The problem is to place n queens on an $n \times n$ chessboard so that no two queens attack each other by being in the same row or in the same column or on the same diagonal.

For $n = 1$, the problem has a **trivial solution**.

Q

For $n = 2$, it is easy to see that there is **no solution** to place 2 queens in 2×2 chessboard.

Q	

For $n = 3$, it is easy to see that there is **no solution** to place 3 queens in 3×3 chessboard.

	1	2	3	
1	Q			← queen 1
2			Q	← queen 2
3				

Or

	1	2	3	
1	Q			← queen 1
2				
3		Q		← queen 2

Or

	1	2	3	
1		Q		← queen 1
2				
3	Q			← queen 2

For $n = 4$, There is **solution** to place 4 queens in 4×4 chessboard. the four-queens problem solved by the backtracking technique.

Step 1: Start with the empty board

	1	2	3	4	
1					← queen 1
2					← queen 2
3					← queen 3
4					← queen 4

Step 2: Place queen 1 in the first possible position of its row, which is in column 1 of row 1.

	1	2	3	4	
1	Q				
2					
3					
4					

Step 3: place queen 2, after trying unsuccessfully columns 1 and 2, in the first acceptable position for it, which is square (2, 3), the square in row 2 and column 3.

	1	2	3	4	
1	Q				
2			Q		
3					
4					

Step 4: This proves to be a dead end because there is no acceptable position for queen 3. So, the algorithm backtracks and puts queen 2 in the next possible position at (2, 4).

	1	2	3	4
1	Q			
2				Q
3				
4				

Step 5: Then queen 3 is placed at (3, 2), which proves to be another dead end.

	1	2	3	4
1	Q			
2				Q
3		Q		
4				

Step 6: The algorithm then backtracks all the way to queen 1 and moves it to (1, 2).

	1	2	3	4
1		Q		
2				
3				
4				

Step 7: The queen 2 goes to (2, 4).

	1	2	3	4
1		Q		
2				Q
3				
4				

Step 8: The queen 3 goes to (3, 1).

	1	2	3	4
1		Q		
2				Q
3	Q			
4				

Step 9: The queen 3 goes to (4, 3). This is a solution to the problem.

	1	2	3	4
1		Q		
2				Q
3	Q			
4			Q	

FIGURE 5.9 Solution four-queens problem in 4x4 Board.

The state-space tree of this search is shown in Figure 12.2




FIGURE 5.10 State-space tree of solving the four-queens problem by backtracking. × denotes an unsuccessful attempt to place a queen.

For $n = 8$, There is **solution** to place 8 queens in 8×8 chessboard.




FIGURE 5.11 Solution 8-queens problem in 8×8 Board.

5.8 HAMILTONIAN CIRCUIT PROBLEM

A **Hamiltonian circuit** (also called a **Hamiltonian cycle**, **Hamilton cycle**, or **Hamilton circuit**) is a graph cycle (i.e., closed loop) through a graph that visits each node exactly once. A graph possessing a **Hamiltonian cycle** is said to be a **Hamiltonian graph**.




FIGURE 5.12 Graph contains Hamiltonian circuit

Let us consider the problem of finding a Hamiltonian circuit in the graph in Figure 5.13.

Example: Find Hamiltonian circuit starts at vertex a .




FIGURE 5.13 Graph.

Solution:

- Assume that if a Hamiltonian circuit exists, it starts at vertex a . accordingly, we make vertex a the root of the state-space tree as in Figure 5.14.
- In a Graph G , Hamiltonian cycle begins at some vertex $V_1 \in G$, and the vertices are visited only once in the order V_1, V_2, \dots, V_n . (V_i are distinct except for V_1 and V_{n+1} which are equal).
- The first component of our future solution, if it exists, is a first intermediate vertex of a Hamiltonian circuit to be constructed. Using the alphabet order to break the three-way tie among the vertices adjacent to a , we
- Select vertex b . From b , the algorithm proceeds to c , then to d , then to e , and finally to f , which proves to be a dead end.
- So the algorithm backtracks from f to e , then to d , and then to c , which provides the first alternative for the algorithm to pursue.
- Going from c to e eventually proves useless, and the algorithm has to backtrack from e to c and then to b . From there, it goes to the vertices f , e , c , and d , from which it can legitimately return to a , yielding the Hamiltonian circuit a, b, f, e, c, d, a . If we wanted to find another Hamiltonian circuit, we could continue this process by backtracking from the leaf of the solution found.

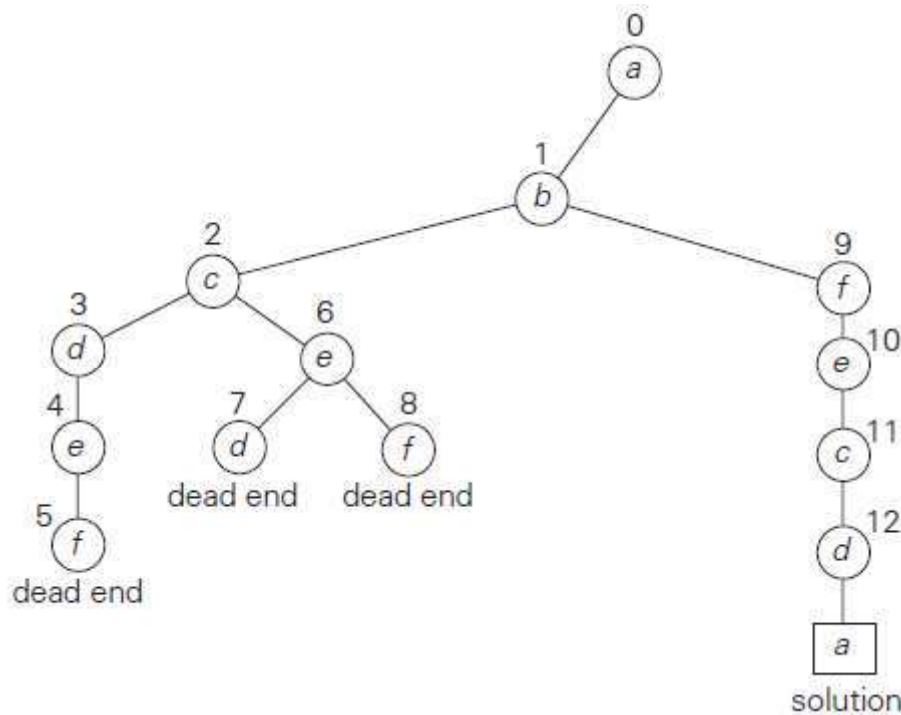


FIGURE 5.14 State-space tree for finding a Hamiltonian circuit.

5.9 SUBSET SUM PROBLEM

The *subset-sum problem* finds a subset of a given set $A = \{a_1, \dots, a_n\}$ of n positive integers whose sum is equal to a given positive integer d . For example, for $A = \{1, 2, 5, 6, 8\}$ and $d = 9$, there are two solutions: $\{1, 2, 6\}$ and $\{1, 8\}$. Of course, some instances of this problem may have no solutions.

It is convenient to sort the set's elements in increasing order. So, we will assume that $a_1 < a_2 < \dots < a_n$.

$A = \{3, 5, 6, 7\}$ and $d = 15$ of the subset-sum problem. The number inside a node is the sum of the elements already included in the subsets represented by the node. The inequality below a leaf indicates the reason for its termination.

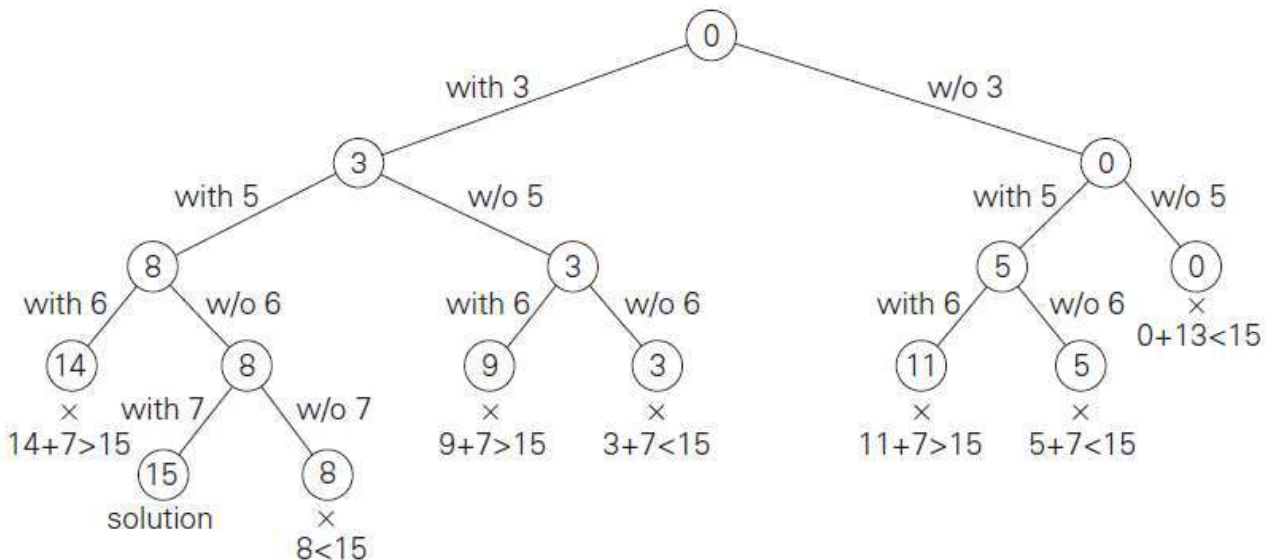


FIGURE 5.15 Complete state-space tree of the backtracking algorithm applied to the instance

Example:

- The state-space tree can be constructed as a binary tree like that in Figure 5.15 for the instance $A = \{3, 5, 6, 7\}$ and $d = 15$.
- The root of the tree represents the starting point, with no decisions about the given elements made as yet.
- Its left and right children represent, respectively, inclusion and exclusion of a_1 in a set being sought. Similarly, going to the left from a node of the first level corresponds to inclusion of a_2 while going to the right corresponds to its exclusion, and so on.
- Thus, a path from the root to a node on the i th level of the tree indicates which of the first I numbers have been included in the subsets represented by that node.
- We record the value of s , the sum of these numbers, in the node.
- If s is equal to d , we have a solution to the problem. We can either report this result and stop or, if all the solutions need to be found, continue by backtracking to the node's parent.
- If s is not equal to d , we can terminate the node as nonpromising if either of the following two inequalities holds:

$$s + a_{i+1} > d \text{ (the sum } s \text{ is too large),}$$

$$s + \sum_{j=i+1}^n a_j < d \text{ (the sum } s \text{ is too small).}$$

General Remarks

From a more general perspective, most backtracking algorithms fit the following description. An output of a backtracking algorithm can be thought of as an n -tuple (x_1, x_2, \dots, x_n) where each coordinate x_i is an element of some finite linearly ordered set S_i . For example, for the n -queens problem, each S_i is the set of integers (column numbers) 1 through n .

A backtracking algorithm generates, explicitly or implicitly, a state-space tree; its nodes represent partially constructed tuples with the first i coordinates defined by the earlier actions of the algorithm. If such a tuple (x_1, x_2, \dots, x_i) is not a solution, the algorithm finds the next element in S_{i+1} that is consistent with the values of $((x_1, x_2, \dots, x_i))$ and the problem's constraints, and adds it to the tuple as its $(i + 1)$ st coordinate. If such an element does not exist, the algorithm backtracks to consider the next value of x_i , and so on.

ALGORITHM *Backtrack*($X[1..i]$)

```
//Gives a template of a generic backtracking algorithm
//Input:  $X[1..i]$  specifies first  $i$  promising components of a solution
//Output: All the tuples representing the problem's solutions
if  $X[1..i]$  is a solution write  $X[1..i]$ 
else //see Problem this section
    for each element  $x \in S_{i+1}$  consistent with  $X[1..i]$  and the constraints do
         $X[i + 1] \leftarrow x$ 
        Backtrack( $X[1..i + 1]$ )
```

5.10 BRANCH AND BOUND

An optimization problem seeks to minimize or maximize some objective function, usually subject to some constraints. Note that in the standard terminology of optimization problems, a *feasible solution* is a point in the problem's search space that satisfies all the problem's constraints (e.g., a Hamiltonian circuit in the travelling salesman problem or a subset of items whose total weight does not exceed the knapsack's capacity in the knapsack problem), whereas an *optimal solution* is a feasible solution with the best value of the objective function (e.g., the shortest Hamiltonian circuit or the most valuable subset of items that fit the knapsack).

Compared to backtracking, branch-and-bound requires two additional items:

1. a way to provide, for every node of a state-space tree, a bound on the best value of the objective function on any solution that can be obtained by adding further components to the partially constructed solution represented by the node
2. the value of the best solution seen so far

If this information is available, we can compare a node's bound value with the value of the best solution seen so far. If the bound value is not better than the value of the best solution seen so far—i.e., not smaller for a minimization problem and not larger for a maximization problem—the node is nonpromising and can be terminated (some people say the branch is “pruned”). Indeed, no solution obtained from it can yield a better solution than the one already available. This is the principal idea of the branch-and-bound technique.

In general, we terminate a search path at the current node in a state-space tree of a branch-and-bound algorithm for any one of the following three reasons:

1. The value of the node's bound is not better than the value of the best solution seen so far.
2. The node represents no feasible solutions because the constraints of the problem are already violated.
3. The subset of feasible solutions represented by the node consists of a single point (and hence no further choices can be made)—in this case, we compare the value of the objective function for this feasible solution with that of the best solution seen so far and update the latter with the former if the new solution is better.

Some problems can be solved by Branch-and-Bound are:

1. Assignment Problem
2. Knapsack Problem
3. Traveling Salesman Problem

5.11 ASSIGNMENT PROBLEM

Let us illustrate the branch-and-bound approach by applying it to the problem of assigning n people to n jobs so that the total cost of the assignment is as small as possible. An instance of the assignment problem is specified by an $n \times n$ cost matrix C .

$$C = \begin{matrix} & \begin{matrix} \text{job 1} & \text{job 2} & \text{job 3} & \text{job 4} \end{matrix} \\ \begin{matrix} \text{person } a \\ \text{person } b \\ \text{person } c \\ \text{person } d \end{matrix} & \begin{bmatrix} 9 & 2 & 7 & 8 \\ 6 & 4 & 3 & 7 \\ 5 & 8 & 1 & 8 \\ 7 & 6 & 9 & 4 \end{bmatrix} \end{matrix}$$

We have to find a lower bound on the cost of an optimal selection without actually solving the problem. We can do this by several methods. For example, it is clear that the cost of any solution, including an optimal one, cannot be smaller than the sum of the smallest elements in each of the matrix's rows. For the instance here, this sum is $2 + 3 + 1 + 4 = 10$. It is important to stress that this is not the cost of any legitimate selection (3 and 1 came from the same column of the matrix); it is just a lower bound on the cost of any legitimate selection. We can and will apply the same thinking to partially constructed solutions. For example, for any legitimate selection that selects 9 from the first row, the lower bound will be $9 + 3 + 1 + 4 = 17$.

It is sensible to consider a node with the best bound as most promising, although this does not, of course, preclude the possibility that an optimal solution will ultimately belong to a different branch of the state-space tree. This variation of the strategy is called the **best-first branch-and-bound**.

The lower-bound value for the root, denoted lb , is 10. The nodes on the first level of the tree correspond to selections of an element in the first row of the matrix, i.e., a job for person a as shown in Figure 5.15.

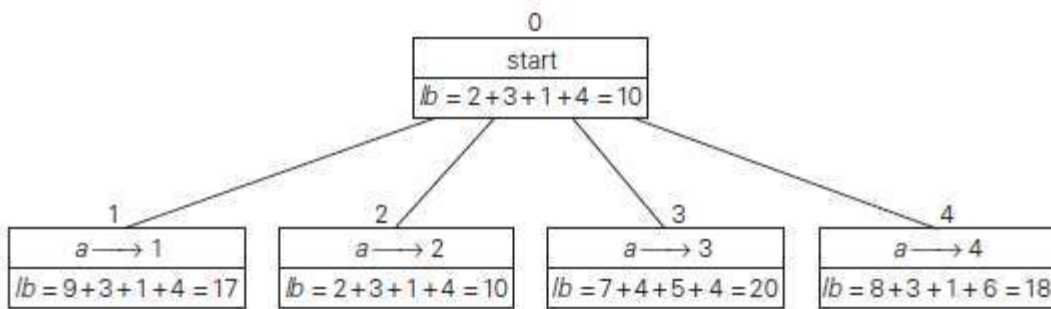


FIGURE 5.15 Levels 0 and 1 of the state-space tree for the instance of the assignment problem being solved with the best-first branch-and-bound algorithm. The number above a node shows the order in which the node was generated. A node's fields indicate the job number assigned to person a and the lower bound value, lb , for this node.

So we have four live leaves (promising leaves are also called **live**)—nodes 1 through 4—that may contain an optimal solution. The most promising of them is node 2 because it has the smallest lowerbound value. Following our best-first search strategy, we branch out from that node first by considering the three different ways of selecting an element from the second row and not in the second column—the three different jobs that can be assigned to person b (Figure 5.16).

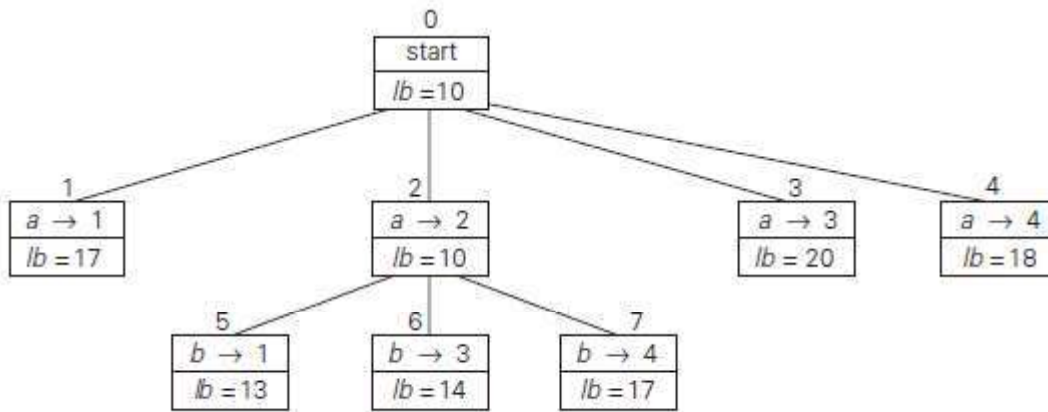


FIGURE 5.16 Levels 0, 1, and 2 of the state-space tree for the instance of the assignment problem being solved with the best-first branch-and-bound algorithm.

Of the six live leaves—nodes 1, 3, 4, 5, 6, and 7—that may contain an optimal solution, we again choose the one with the smallest lower bound, node 5. First, we consider selecting the third column’s element from *c*’s row (i.e., assigning person *c* to job 3); this leaves us with no choice but to select the element from the fourth column of *d*’s row (assigning person *d* to job 4). This yields leaf 8 (Figure 5.17), which corresponds to the feasible solution $\{a \rightarrow 2, b \rightarrow 1, c \rightarrow 3, d \rightarrow 4\}$ with the total cost of 13. Its sibling, node 9, corresponds to the feasible solution $\{a \rightarrow 2, b \rightarrow 1, c \rightarrow 4, d \rightarrow 3\}$ with the total cost of 25. Since its cost is larger than the cost of the solution represented by leaf 8, node 9 is simply terminated. (Of course, if its cost were smaller than 13, we would have to replace the information about the best solution seen so far with the data provided by this node.)

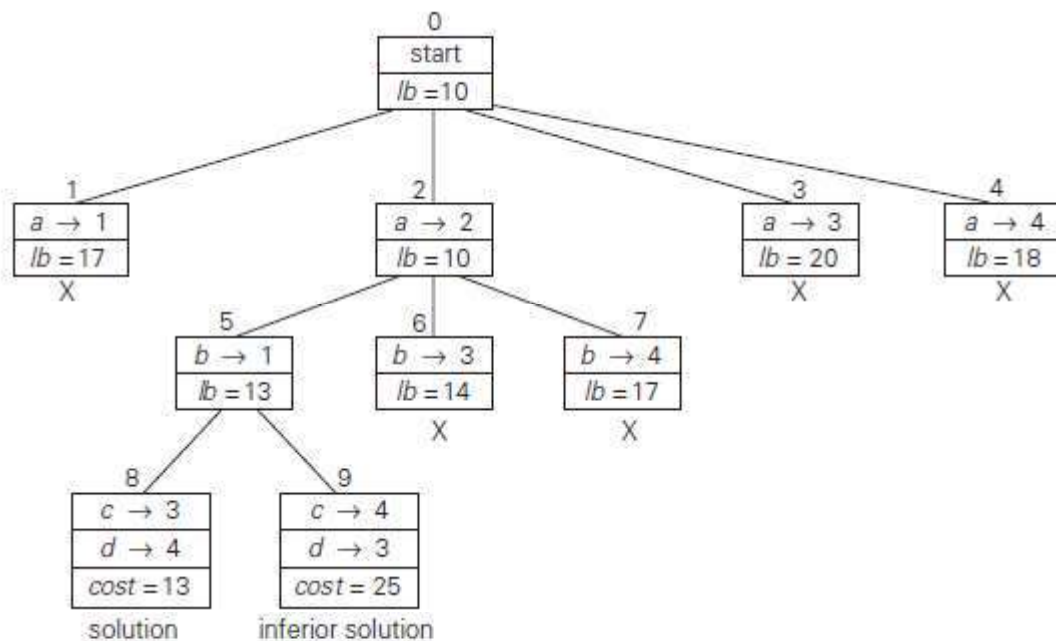


FIGURE 5.17 Complete state-space tree for the instance of the assignment problem solved with the best-first branch-and-bound algorithm.

Now, as we inspect each of the live leaves of the last state-space tree—nodes 1, 3, 4, 6, and 7 in Figure 5.17—we discover that their lower-bound values are not smaller than 13, the value of the best selection seen so far (leaf 8). Hence, we terminate all of them and recognize the solution represented by leaf 8 as the optimal solution to the problem.

5.12 KNAPSACK PROBLEM

Let us now discuss how we can apply the branch-and-bound technique to solving the knapsack problem. Given n items of known weights w_i and values v_i , $i = 1, 2, \dots, n$, and a knapsack of capacity W , find the most valuable subset of the items that fit in the knapsack. It is convenient to order the items of a given instance in descending order by their value-to-weight ratios. Then the first item gives the best payoff per weight unit and the last one gives the worst payoff per weight unit, with ties resolved arbitrarily:

$$v_1/w_1 \geq v_2/w_2 \geq \dots \geq v_n/w_n.$$

It is natural to structure the state-space tree for this problem as a binary tree constructed as follows. Each node on the i th level of this tree, $0 \leq i \leq n$, represents all the subsets of n items that include a particular selection made from the first i ordered items. This particular selection is uniquely determined by the path from the root to the node: a branch going to the left indicates the inclusion of the next item, and a branch going to the right indicates its exclusion. We record the total weight w and the total value v of this selection in the node, along with some upper bound ub on the value of any subset that can be obtained by adding zero or more items to this selection.

Item	Weight	value	value / weight	capacity
1	4	\$40	10	W = 10
2	7	\$42	6	
3	5	\$25	5	
4	3	\$12	4	
	w=19	v=119	$v_{i+1}/w_{i+1}=25$	

A simple way to compute the upper bound ub is to add to v , the total value of the items already selected, the product of the remaining capacity of the knapsack $W - w$ and the best per unit payoff among the remaining items, which is v_{i+1}/w_{i+1} :

$$\begin{aligned} ub &= v + (W - w)(v_{i+1}/w_{i+1}). \\ &= 0 + (10 - 0)(10) \\ &= 100 \end{aligned}$$

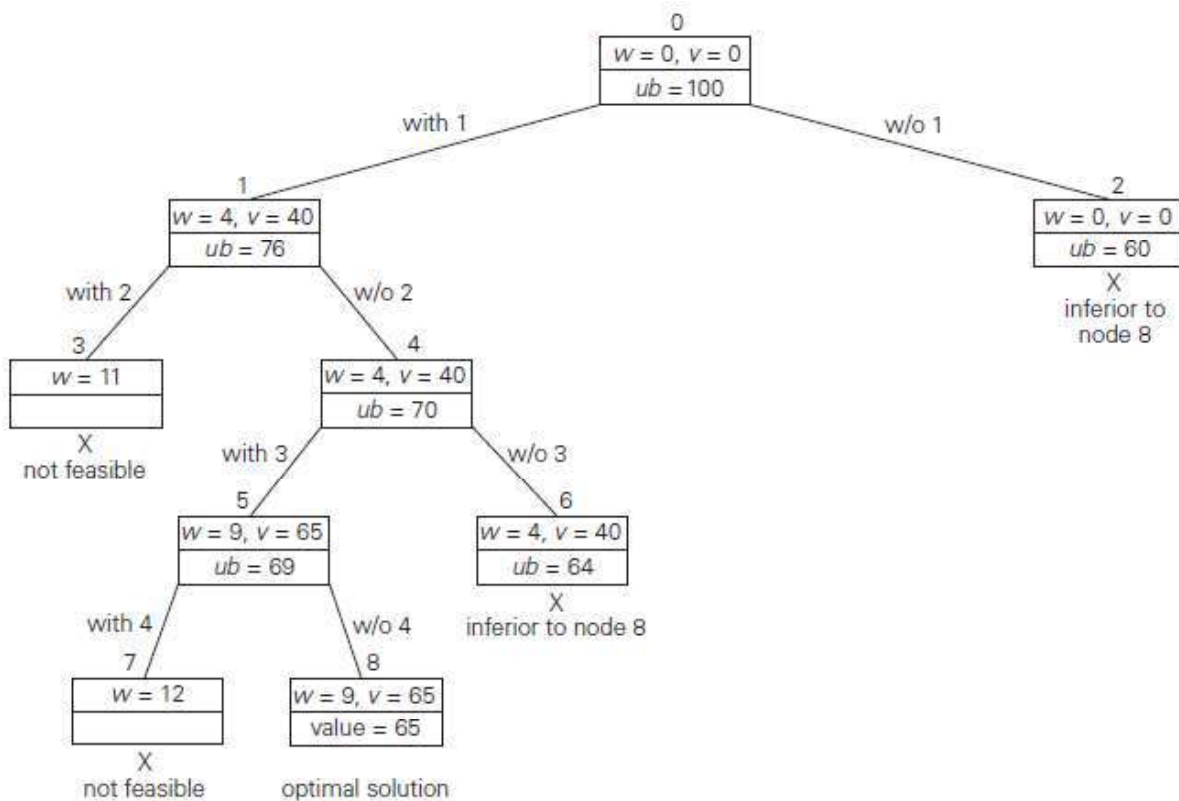


FIGURE 5.18 State-space tree of the best-first branch-and-bound algorithm for the instance of the knapsack problem.

At the root of the state-space tree (see Figure 5.18), no items have been selected as yet. Hence, both the total weight of the items already selected w and their total value v are equal to 0. The value of the upper bound computed by formula (12.1) is \$100. Node 1, the left child of the root, represents the subsets that include item 1. The total weight and value of the items already included are 4 and \$40, respectively; the value of the upper bound is $40 + (10 - 4) * 6 = \$76$. Node 2 represents the subsets that do not include item 1. Accordingly, $w = 0$, $v = \$0$, and $ub = 0 + (10 - 0) * 6 = \60 . Since node 1 has a larger upper bound than the upper bound of node 2, it is more promising for this maximization problem, and we branch from node 1 first. Its children—nodes 3 and 4—represent subsets with item 1 and with and without item 2, respectively.

Since the total weight w of every subset represented by node 3 exceeds the knapsack's capacity, node 3 can be terminated immediately. Node 4 has the same values of w and v as its parent; the upper bound ub is equal to $40 + (10 - 4) * 5 = \$70$. Selecting node 4 over node 2 for the next branching (why?), we get nodes 5 and 6 by respectively including and excluding item 3. The total weights and values as well as the upper bounds for these nodes are computed in the same way as for the preceding nodes. Branching from node 5 yields node 7, which represents no feasible solutions, and node 8, which represents just a single subset $\{1, 3\}$ of value \$65. The remaining live nodes 2 and 6 have smaller upper-bound values than the value of the solution represented by node 8. Hence, both can be terminated making the subset $\{1, 3\}$ of node 8 the optimal solution to the problem.

Solving the knapsack problem by a branch-and-bound algorithm has a rather unusual characteristic. Typically, internal nodes of a state-space tree do not define a point of the problem's search space, because some of the solution's components remain undefined. If we had done this for the instance investigated above, we could have terminated nodes 2 and 6 before node 8 was generated because they both are inferior to the subset of value \$65 of node 5.

5.13 TRAVELING SALESMAN PROBLEM

We will be able to apply the branch-and-bound technique to instances of the travelling salesman problem if we come up with a reasonable lower bound on tour lengths. One very simple lower bound can be obtained by finding the smallest element in the intercity distance matrix D and multiplying it by the number of cities n . But there is a less obvious and more informative lower bound for instances with symmetric matrix D , which does not require a lot of work to compute. It is not difficult to show (Problem 8 in this section's exercises) that we can compute a lower bound on the length l of any tour as follows. For each city i , $1 \leq i \leq n$, find the sum s_i of the distances from city i to the two nearest cities; compute the sum s of these n numbers, divide the result by 2, and, if all the distances are integers, round up the result to the nearest integer:

$$lb = \lceil s/2 \rceil$$

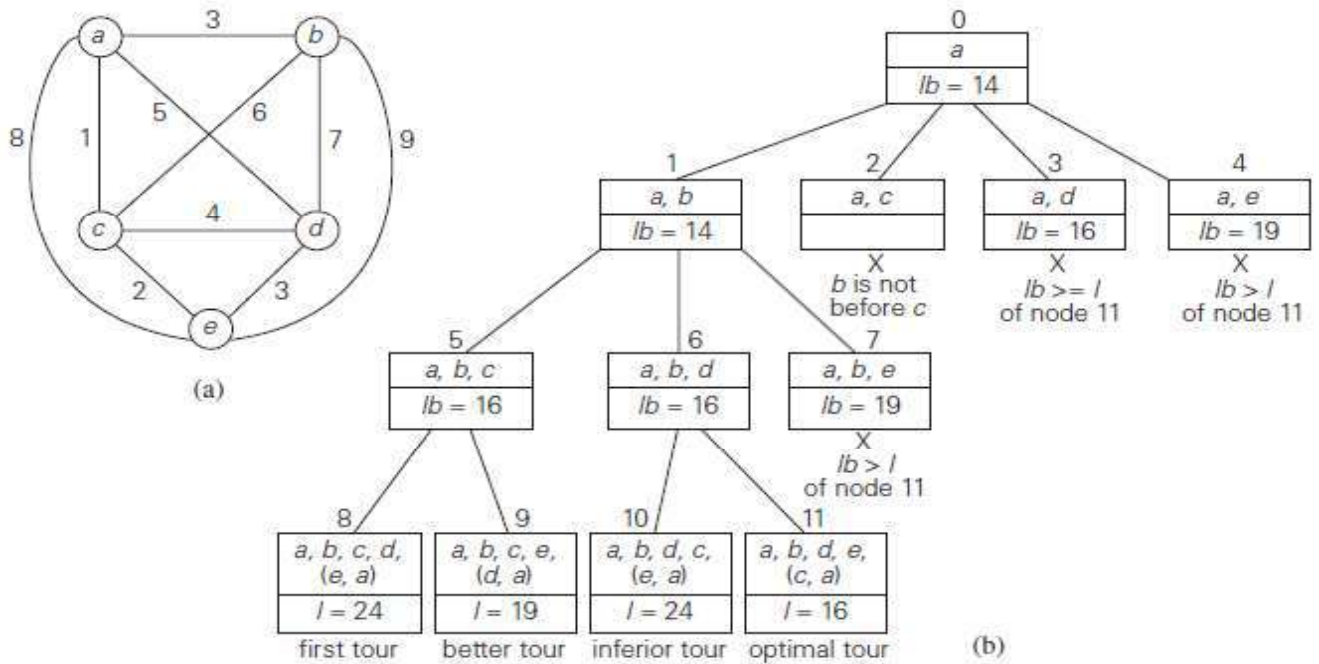


FIGURE 5.19 (a)Weighted graph. (b) State-space tree of the branch-and-bound algorithm to find a shortest Hamiltonian circuit in this graph. The list of vertices in a node specifies a beginning part of the Hamiltonian circuits represented by the node.

For example, for the instance in Figure and above formula yields

$$lb = \lceil [(1 + 3) + (3 + 6) + (1 + 2) + (3 + 4) + (2 + 3)]/2 \rceil = 14.$$

Moreover, for any subset of tours that must include particular edges of a given graph, we can modify lower bound accordingly. For example, for all the Hamiltonian circuits of the graph in Figure that must include edge (a, d) , we get the following lower bound by summing up the lengths of the two shortest edges incident with each of the vertices, with the required inclusion of edges (a, d) and (d, a) :

$$\lceil [(1 + 5) + (3 + 6) + (1 + 2) + (3 + 5) + (2 + 3)]/2 \rceil = 16.$$

We now apply the branch-and-bound algorithm, with the bounding function given by formula, to find the shortest Hamiltonian circuit for the graph in Figure 5.19a. To reduce the amount of potential work. First, without loss of generality, we can consider only tours that start at a . Second, because our graph is undirected, we can generate only tours in which b is visited before c . In addition, after visiting $n - 1 = 4$ cities, a tour has no choice but to visit the remaining unvisited city and return to the starting one. The state-space tree tracing the algorithm's application is given in Figure 5.19b.

5.14 APPROXIMATION ALGORITHMS FOR NP HARD PROBLEMS

Now we are going to discuss a different approach to handling difficult problems of combinatorial optimization, such as the **travelling salesman problem and the knapsack problem**. The decision versions of these problems are *NP*-complete. Their optimization versions fall in the class of *NP-hard problems*—problems that are at least as hard as *NP*-complete problems. Hence, there are no known polynomial-time algorithms for these problems, and there are serious theoretical reasons to believe that such algorithms do not exist.

Approximation algorithms run a gamut in level of sophistication; most of them are based on some problem-specific heuristic. A **heuristic** is a common-sense rule drawn from experience rather than from a mathematically proved assertion. For example, going to the nearest unvisited city in the travelling salesman problem is a good illustration of this notion.

Of course, if we use an algorithm whose output is just an approximation of the actual optimal solution, we would like to know how accurate this approximation is. We can quantify the accuracy of an approximate solution s_a to a problem of **minimizing** some function f by the size of the relative error (*re*) of this approximation,

$$re(s_a) = \frac{f(s_a) - f(s^*)}{f(s^*)}$$

where s^* is an exact solution to the problem. Alternatively, $re(s_a) = f(s_a)/f(s^*) - 1$, we can simply use the **accuracy ratio**

$$r(s_a) = \frac{f(s_a)}{f(s^*)}$$

as a measure of accuracy of s_a . Note that for the sake of scale uniformity, the accuracy ratio of approximate solutions to **maximization** problems is usually computed as

$$r(s_a) = \frac{f(s^*)}{f(s_a)}$$

to make this ratio greater than or equal to 1, as it is for minimization problems. Obviously, the closer $r(s_a)$ is to 1, the better the approximate solution is. For most instances, however, we cannot compute the accuracy ratio, because we typically do not know $f(s^*)$, the true optimal value of the objective function. Therefore, our hope should lie in obtaining a good upper bound on the values of $r(s_a)$. This leads to the following definitions.

A polynomial-time approximation algorithm is said to be a ***c* approximation algorithm**, where $c \geq 1$, if the accuracy ratio of the approximation it produces does not exceed c for any instance of the problem in question: $r(s_a) \leq c$.

The best (i.e., the smallest) value of c for which inequality holds for all instances of the problem is called the **performance ratio** of the algorithm and denoted R_A .

The performance ratio serves as the principal metric indicating the quality of the approximation algorithm. We would like to have approximation algorithms with R_A as close to 1 as possible. Unfortunately, as we shall see, some approximation algorithms have infinitely large performance ratios ($R_A = \infty$). This does not necessarily rule out using such algorithms, but it does call for a cautious treatment of their outputs.

Approximation Algorithms for NP Hard Problems are:

- Traveling salesman problem (tsp)
- Knapsack problem

5.15 TRAVELING SALESMAN PROBLEM (APPROXIMATION ALGORITHM)

Greedy Algorithms for the TSP The simplest approximation algorithms for the traveling salesman problem are based on the greedy technique. We will discuss here two such algorithms.

1. Nearest-neighbor algorithm
2. Minimum-Spanning-Tree-Based Algorithms

NEAREST-NEIGHBOR ALGORITHM

The following well-known greedy algorithm is based on the *nearest-neighbor* heuristic: always go next to the nearest unvisited city.

Step 1 Choose an arbitrary city as the start.

Step 2 Repeat the following operation until all the cities have been visited: go to the unvisited city nearest the one visited last (ties can be broken arbitrarily).

Step 3 Return to the starting city.

EXAMPLE 1 For the instance represented by the graph in Figure 5.20, with a as the starting vertex, the nearest-neighbor algorithm yields the tour (Hamiltonian circuit) $s_a: a - b - c - d - a$ of length 10.




FIGURE 5.20 Instance of the traveling salesman problem.

The optimal solution, as can be easily checked by exhaustive search, is the tour $s^*: a - b - d - c - a$ of length 8. Thus, the accuracy ratio of this approximation is

$$r(s_a) = \frac{f(s_a)}{f(s^*)} = \frac{10}{8} = 1.25$$

(i.e., tour s_a is 25% longer than the optimal tour s^*).

Multifragment-heuristic algorithm

Another natural greedy algorithm for the traveling salesman problem considers it as the problem of finding a minimum-weight collection of edges in a given complete weighted graph so that all the vertices have **degree 2**.

Step 1 Sort the edges in increasing order of their weights. (Ties can be broken arbitrarily.) Initialize the set of tour edges to be constructed to the empty set.

Step 2 Repeat this step n times, where n is the number of cities in the instance being solved: add the next edge on the sorted edge list to the set of tour edges, provided this addition does not create a vertex of degree 3 or a cycle of length less than n ; otherwise, skip the edge.

Step 3 Return the set of tour edges.

As an example, applying the algorithm to the graph in Figure 5.20 yields $\{(a, b), (c, d), (b, c), (a, d)\}$. This set of edges forms the same tour as the one produced by the nearest-neighbor algorithm. In general, the multifragment-heuristic algorithm tends to produce significantly better

tours than the nearest-neighbor algorithm, as we are going to see from the experimental data quoted at the end of this section. But the performance ratio of the multifragment-heuristic algorithm is also unbounded, of course.

There is, however, a very important subset of instances, called *Euclidean*, for which we can make a nontrivial assertion about the accuracy of both the nearest-neighbor and multifragment-heuristic algorithms. These are the instances in which intercity distances satisfy the following natural conditions:

- **triangle inequality** $d[i, j] \leq d[i, k] + d[k, j]$ for any triple of cities i, j , and k (the distance between cities i and j cannot exceed the length of a two-leg path from i to some intermediate city k to j)
- **symmetry** $d[i, j] = d[j, i]$ for any pair of cities i and j (the distance from i to j is the same as the distance from j to i)

MINIMUM-SPANNING-TREE-BASED ALGORITHMS

There are approximation algorithms for the travelling salesman problem that exploit a connection between Hamiltonian circuits and spanning trees of the same graph. Since removing an edge from a Hamiltonian circuit yields a spanning tree, we can expect that the structure of a minimum spanning tree provides a good basis for constructing a shortest tour approximation. Here is an algorithm that implements this idea in a rather straightforward fashion.

Twice-around-the-tree algorithm

- Step 1** Construct a minimum spanning tree of the graph corresponding to a given instance of the traveling salesman problem.
- Step 2** Starting at an arbitrary vertex, perform a walk around the minimum spanning tree recording all the vertices passed by. (This can be done by a DFS traversal.)
- Step 3** Scan the vertex list obtained in Step 2 and eliminate from it all repeated occurrences of the same vertex except the starting one at the end of the list. (This step is equivalent to making shortcuts in the walk.) The vertices remaining on the list will form a Hamiltonian circuit, which is the output of the algorithm.

EXAMPLE 2 Let us apply this algorithm to the graph in Figure 5.21a. The minimum spanning tree of this graph is made up of edges (a, b) , (b, c) , (b, d) , and (d, e) (Figure 5.21b). A twice-around-the-tree walk that starts and ends at a is $a, b, c, b, d, e, d, b, a$. Eliminating the second b (a shortcut from c to d), the second d , and the third b (a shortcut from e to a) yields the Hamiltonian circuit a, b, c, d, e, a of length 39.

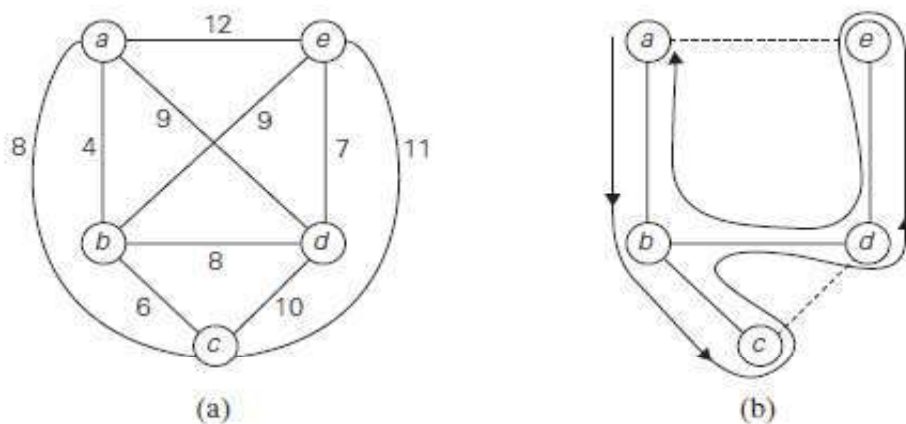


FIGURE 5.21 Illustration of the twice-around-the-tree algorithm. (a) Graph. (b) Walk around the minimum spanning tree with the shortcuts.

5.16 KNAPSACK PROBLEM (APPROXIMATION ALGORITHM)

The knapsack problem is one well-known *NP*-hard problem. Given n items of known weights w_1, \dots, w_n and values v_1, \dots, v_n and a knapsack of weight capacity W , find the most valuable subset of the items that fits into the knapsack.

GREEDY ALGORITHMS FOR THE KNAPSACK PROBLEM

We can think of several greedy approaches to this problem. One is to select the items in decreasing order of their weights; however, heavier items may not be the most valuable in the set. Alternatively, if we pick up the items in decreasing order of their value, there is no guarantee that the knapsack's capacity will be used efficiently. We find a greedy strategy that takes into account both the weights and values by computing the value-to-weight ratios v_i/w_i , $i = 1, 2, \dots, n$, and selecting the items in decreasing order of these ratios. Here is the algorithm based on this greedy heuristic.

Greedy algorithm for the discrete knapsack problem

Step 1 Compute the value-to-weight ratios $r_i = v_i/w_i$, $i = 1, \dots, n$, for the items given.

Step 2 Sort the items in nonincreasing order of the ratios computed in Step 1. (Ties can be broken arbitrarily.)

Step 3 Repeat the following operation until no item is left in the sorted list: if the current item on the list fits into the knapsack, place it in the knapsack and proceed to the next item; otherwise, just proceed to the next item.

EXAMPLE 1 Let us consider the instance of the knapsack problem with the knapsack capacity 10 and the item information as follows:

Item	weight	value
1	4	\$40
2	7	\$42
3	5	\$25
4	3	\$12

Computing the value-to-weight ratios and sorting the items in non increasing order of these efficiency ratios yields

Item	weight	value	value / weight	capacity
1	4	\$40	10	$W = 10$
2	7	\$42	6	
3	5	\$25	5	
4	3	\$12	4	

The greedy algorithm will select the first item of weight **4**, skip the next item of weight 7, select the next item of weight **5**, and skip the last item of weight 3. The solution obtained happens to be optimal for this instance. So the total items value in knapsack is **\$65**.

GREEDY ALGORITHM FOR THE CONTINUOUS KNAPSACK PROBLEM

Step 1 Compute the value-to-weight ratios v_i/w_i , $i = 1, \dots, n$, for the items given.

Step 2 Sort the items in nonincreasing order of the ratios computed in Step 1. (Ties can be broken arbitrarily.)

Step 3 Repeat the following operation until the knapsack is filled to its full capacity or no item is left in the sorted list: if the current item on the list fits into the knapsack in its

entirety, take it and proceed to the next item; otherwise, take its largest fraction to fill the knapsack to its full capacity and stop.

EXAMPLE 2 A small example of an approximation scheme with $k = 2$ is provided. The algorithm yields $\{1, 3, 4\}$, which is the optimal solution for this instance.

Item	weight	value	value / weight	capacity
1	4	\$40	10	W = 10
2	7	\$42	6	
3	5	\$25	5	
4	1	\$4	4	

subset	Added items	value
{}	1, 3, 4	\$69
{1}	3, 4	\$69
{2}	4	\$46
{3}	1, 4	\$69
{4}	1, 3	\$69
{1, 2}	Not feasible	
{1, 4}	4	\$69
{1, 4}	3	\$69
{2, 3}	Not feasible	
{2, 4}	-	\$46
{3, 4}	1	\$69

For each of those subsets, it needs $O(n)$ time to determine the subset's possible extension. Thus, the algorithm's efficiency is in $O(kn^{k+1})$. Note that although it is polynomial in n , the time efficiency of Sahni's scheme is exponential in k . More sophisticated approximation schemes, called *fully polynomial schemes*, do not have this shortcoming.

All the best - There is no substitute for hard work