

Vehicle Routing and Scheduling

Tutorial Outline

INTRODUCTION

A Service Delivery Example: Meals-for-ME

OBJECTIVES OF ROUTING AND SCHEDULING PROBLEMS

CHARACTERISTICS OF ROUTING AND SCHEDULING PROBLEMS

Classifying Routing and Scheduling Problems

Solving Routing and Scheduling Problems

ROUTING SERVICE VEHICLES

The Traveling Salesman Problem

Multiple Traveling Salesman Problem

The Vehicle Routing Problem

Cluster First, Route Second Approach

SCHEDULING SERVICE VEHICLES

The Concurrent Scheduler Approach

OTHER ROUTING AND SCHEDULING PROBLEMS

SUMMARY

KEY TERMS

DISCUSSION QUESTIONS

PROBLEMS

CASE STUDY: ROUTING AND SCHEDULING OF PHLEBOTOMISTS

BIBLIOGRAPHY

INTRODUCTION

The scheduling of customer service and the routing of service vehicles are at the heart of many service operations. For some services, such as school buses, public health nursing, and many installation or repair businesses, service delivery is critical to the performance of the service. For other services, such as mass transit, taxis, trucking firms, and the U.S. Postal Service, timely delivery *is* the service. In either case, the routing and scheduling of service vehicles has a major impact on the *quality* of the service provided.

This tutorial introduces some routing and scheduling terminology, classifies different types of routing and scheduling problems, and presents various solution methodologies. Although every effort has been made to present the topic of vehicle routing and scheduling as simply and as straightforward as possible, it should be noted that this is a technical subject and one of the more mathematical topics in this text. The tutorial begins with an example of service delivery to illustrate some of the practical issues in vehicle routing and scheduling.

A Service Delivery Example: Meals-for-ME

A private, nonprofit meal delivery program for the elderly called Meals-for-ME has been operating in the state of Maine since the mid-1970s. The program offers home delivery of hot meals, Monday through Friday, to “home-bound” individuals who are over 60 years of age. For those individuals who are eligible (and able), the program also supports a “congregate” program that provides daily transportation to group-meal sites. On a typical day within a single county, hundreds of individuals receive this service. In addition, individuals may be referred for short-term service because of a temporary illness or recuperation. Thus, on any given day, the demand for the service can be highly unpredictable. Scheduling of volunteer delivery personnel and vehicles as well as construction of routes is done on a weekly to monthly basis by regional site managers. It is the task of these individuals to coordinate the preparation of meals and to determine the sequence in which customers are to be visited. In addition, site managers must arrange for rides to the “group meals” for participating individuals.

Although these tasks may seem straightforward, there are many practical problems in routing and scheduling meal delivery. First, the delivery vehicles (and pickup vehicles) are driven by volunteers, many of whom are students who are not available during some high-demand periods (Christmas, for example). Thus, the variability in available personnel requires that delivery routes be changed frequently. Second, because the program delivers hot meals, a typical route must be less than 90 minutes. Generally, 20 to 25 meals are delivered on a route, depending on the proximity of customers. Third, all meals must be delivered within a limited time period, between 11:30 A.M. and 1:00 P.M. daily. Similar difficulties exist for personnel who pick up individuals served by the congregate program. Given the existence of these very real problems, the solution no longer seems as simple. It is obvious that solution approaches and techniques are needed that allow the decision maker to consider a multitude of variables and adapt to changes quickly and efficiently.

OBJECTIVES OF ROUTING AND SCHEDULING PROBLEMS

The objective of most routing and scheduling problems is to minimize the total cost of providing the service. This includes vehicle capital costs, mileage, and personnel costs. But other objectives also may come into play, particularly in the public sector. For example, in school bus routing and scheduling, a typical objective is to minimize the total number of student-minutes on the bus. This criterion is highly correlated with safety and with parents’ approval of the school system. For dial-a-ride services for the handicapped or elderly, an important objective is to minimize the inconvenience for all customers. For the Meals-for-ME program, the meals must be delivered at certain times of the day. For emergency services, such as ambulance, police, and fire, minimizing response time to an incident is of primary importance. Some companies promise package delivery by 10:30 A.M. the next morning. Thus, in the case of both public and private services, an appropriate objective function should consider more than the dollar cost of delivering a service. The “subjective” costs associated with failing to provide adequate service to the customer must be considered as well.

CHARACTERISTICS OF ROUTING AND SCHEDULING PROBLEMS

Routing and scheduling problems are often presented as graphical **networks**. The use of networks to describe these problems has the advantage of allowing the decision maker to visualize the problem under consideration. As an example, refer to Figure T5.1. The figure consists of five circles called **nodes**. Four of the nodes (nodes 2 through 5) represent pickup and/or delivery points, and a fifth (node 1) represents a **depot node**, from which the vehicle's trip originates and ends. The depot node is the "home base" for the vehicle or provider.

Connecting these nodes are line segments referred to as **arcs**. Arcs describe the time, cost, or distance required to travel from one node to another. The numbers along the arcs in Figure T5.1 are distances in miles. Given an average speed of travel or a distribution of travel times, distance can be easily converted to time. However, this conversion ignores physical barriers, such as mountains, lack of access, or traffic congestion. If minimizing time is the primary goal in a routing and scheduling problem, then historical data on travel times are preferable to calculations based on distances.

Arcs may be directed or undirected. **Undirected arcs** are represented by simple line segments. **Directed arcs** are indicated by arrows. These arrows represent the direction of travel in the case of routing problems (e.g., one-way streets) or precedence relationships in the case of scheduling problems (where one pickup or delivery task must precede another).

The small network in Figure T5.1 can be viewed as a route for a single vehicle. The route for the vehicle, also called a **tour**, is $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 1$ or, because the arcs are undirected, $1 \rightarrow 5 \rightarrow 4 \rightarrow 3 \rightarrow 2 \rightarrow 1$. The total distance for either tour is 51 miles.

The tour described in Figure T5.1 is a solution to a simple routing problem where the objective is to find the route that minimizes cost or any other criterion that may be appropriate (such as distance or travel time). The minimum-cost solution, however, is subject to the tour being **feasible**. Feasibility depends on the type of problem, but, in general, implies that:

1. A tour must include all nodes.
2. A node must be visited only once.
3. A tour must begin and end at a depot.

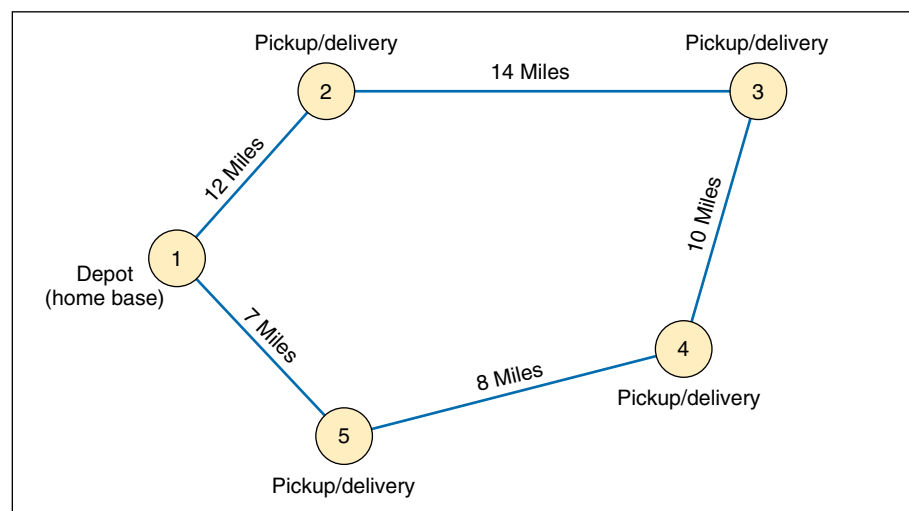
The output of all routing and scheduling systems is essentially the same. That is, for each vehicle or provider, a route and/or a schedule is provided. Generally, the **route** specifies the sequence in which the nodes (or arcs) are to be visited, and a **schedule** identifies when each node is to be visited.

Classifying Routing and Scheduling Problems

The classification of routing and scheduling problems depends on certain characteristics of the service delivery system, such as size of the delivery fleet, where the fleet is housed, capacities of the vehicles, and routing and scheduling objectives. In the simplest case, we begin with a set of nodes

FIGURE T5.1 ■

Routing Network Example



to be visited by a single vehicle. The nodes may be visited in any order, there are no precedence relationships, the travel costs between two nodes are the same regardless of the direction traveled, and there are no delivery-time restrictions. In addition, vehicle capacity is not considered. The output for the single-vehicle problem is a route or a tour where each node is visited only once and the route begins and ends at the depot node (see Figure T5.1, for example). The tour is formed with the goal of minimizing the total tour cost. This simplest case is referred to as a **traveling salesman problem (TSP)**.

An extension of the traveling salesman problem, referred to as the **multiple traveling salesman problem (MTSP)**, occurs when a fleet of vehicles must be routed from a single depot. The goal is to generate a set of routes, one for each vehicle in the fleet. The characteristics of this problem are that a node may be assigned to only one vehicle, but a vehicle will have more than one node assigned to it. There are no restrictions on the size of the load or number of passengers a vehicle may carry. The solution to this problem will give the order in which each vehicle is to visit its assigned nodes. As in the single-vehicle case, the objective is to develop the set of minimum-cost routes, where “cost” may be represented by a dollar amount, distance, or travel time.

If we now restrict the capacity of the multiple vehicles and couple with it the possibility of having varying demands at each node, the problem is classified as a **vehicle routing problem (VRP)**.

Alternatively, if the demand for the service occurs on the arcs, rather than at the nodes, or if demand is so high that individual demand nodes become too numerous to specify, we have a **Chinese postman problem (CRP)**. Examples of these types of problems include street sweeping, snow removal, refuse collection, postal delivery, and paper delivery. The Chinese postman problem is very difficult to solve, and the solution procedures are beyond the scope of this text. Table T5.1 summarizes the characteristics of these four types of routing problems.

Finally, let us distinguish between **routing** problems and **scheduling** problems. If the customers being serviced have no time restrictions and no precedence relationships exist, then the problem is a pure routing problem. If there is a specified time for the service to take place, then a scheduling problem exists. Otherwise, we are dealing with a combined routing and scheduling problem.

Solving Routing and Scheduling Problems

Another important issue in routing and scheduling involves the practical aspects of solving these types of problems. Consider, for example, the delivery of bundles of newspapers from a printing site to dropoff points in a geographic area. These dropoff points supply papers to newspaper carriers for local deliveries. The dropoff points have different demands, and the vehicles have different capacities. Each vehicle is assigned a route beginning and ending at the printing site (the depot). For a newspaper with only 10 dropoff points there are 2^{10} or 1,024 possible routings. For 50 dropoff points, there are 2^{50} or over 1 trillion possible routings. Realistic problems of this type may have over 1,000 drop points! It is evident that problems of any size quickly become too expensive to solve optimally. Fortunately, some very elegant heuristics or “rule of thumb” solution techniques have been developed that yield “good,” if not optimal, solutions to these problems. Some of the more well known of these heuristic approaches are presented in this tutorial.

TABLE T5.1

Characteristics of Four Routing Problems

TYPE	DEMAND	ARCS	NO. OF DEPOTS	NO. OF VEHICLES	VEHICLE CAPACITY
Traveling salesman problem (TSP)	At the nodes	Directed or undirected	1	=1	Unlimited
Multiple traveling salesman problem (MTSP)	At the nodes	Directed or undirected	1	>1	Unlimited
Vehicle routing problem (VRP)	At the nodes	Directed or undirected	1	>1	Limited
Chinese postman problem (CPP)	On the arcs	Directed or undirected	1	≥ 1	Limited or unlimited

TABLE T5.2Symmetric Distance
Matrix

FROM NODE	TO NODE (DISTANCES IN MILES)					
	1	2	3	4	5	6
1	—	5.4	2.8	10.5	8.2	4.1
2	5.4	—	5.0	9.5	5.0	8.5
3	2.8	5.0	—	7.8	6.0	3.6
4	10.5	9.5	7.8	—	5.0	9.5
5	8.2	5.0	6.0	5.0	—	9.2
6	4.1	8.5	3.6	9.5	9.2	—

ROUTING SERVICE VEHICLES

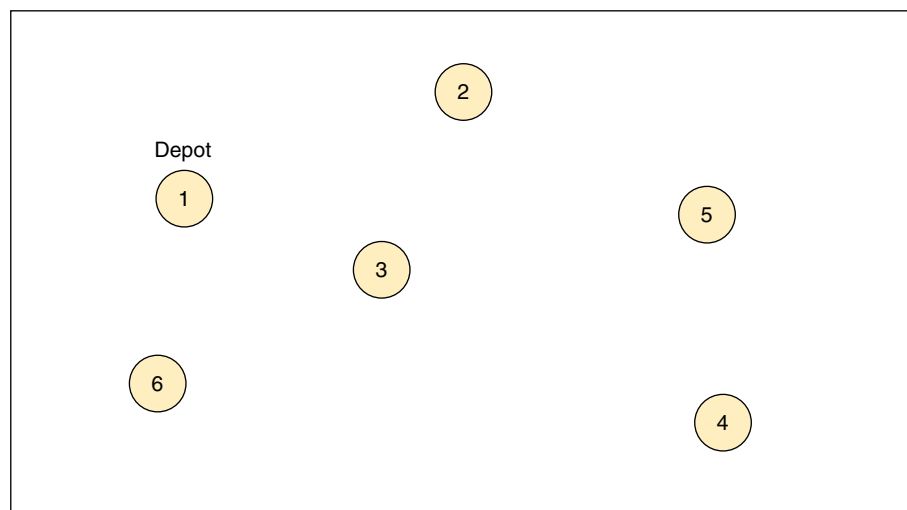
The Traveling Salesman Problem

The traveling salesman problem (TSP) is one of the most studied problems in management science. Optimal approaches to solving traveling salesman problems are based on mathematical programming. But in reality, most TSP problems are not solved optimally. When the problem is so large that an optimal solution is impossible to obtain, or when approximate solutions are good enough, heuristics are applied. Two commonly used heuristics for the traveling salesman problem are the **nearest neighbor procedure** and the **Clark and Wright savings heuristic**.

The Nearest Neighbor Procedure The nearest neighbor procedure (NNP) builds a tour based only on the cost or distance of traveling from the last-visited node to the closest node in the network. As such, the heuristic is simple, but it has the disadvantage of being rather shortsighted, as we shall see in an example. The heuristic does, however, generate an “approximately” optimal solution from a distance matrix. The procedure is outlined as follows:

1. Start with a node at the beginning of the tour (the depot node).
2. Find the node closest to the last node added to the tour.
3. Go back to step 2 until all nodes have been added.
4. Connect the first and the last nodes to form a complete tour.

Example of the Nearest Neighbor Procedure We begin the nearest neighbor procedure with data on the distance or cost of traveling from every node in the network to every other node in the network. In the case where the arcs are undirected, the distance from i to j will be the same as the distance from j to i . Such a network with undirected arcs is said to be *symmetrical*. Table T5.2 gives the complete distance matrix for the symmetrical six-node network shown in Figure T5.2.

FIGURE T5.2 ■Traveling Salesman
Problem

Referring to Figure T5.3, the solution is determined as follows:

1. Start with the depot node (node 1). Examine the distances between node 1 and every other node. The closest node to node 1 is node 3, so designate the **partial tour** or **path** as 1 → 3. (See Figure T5.3[a]. Note that the → means that the nodes are connected, not that the arc is directed.)
2. Find the closest node to the last node added (node 3) that is not currently in the path. Node 6 is 3.6 miles from node 3, so connect it to the path. The result is the three-node path 1 → 3 → 6. (See Figure T5.3[b].)
3. Find the node closest to node 6 that has not yet been connected. This is node 2, which is 8.5 miles from node 6. Connect it to yield 1 → 3 → 6 → 2. (See Figure T5.3[c].)
4. The node closest to node 2 is node 5. The partial tour is now 1 → 3 → 6 → 2 → 5. (See Figure T5.3[d].)
5. Connect the last node (node 4) to the path and complete the tour by connecting node 4 to the depot. The complete tour formed is 1 → 3 → 6 → 2 → 5 → 4 → 1. The length of the tour is 35.4 miles. (See Figure T5.3[e].)

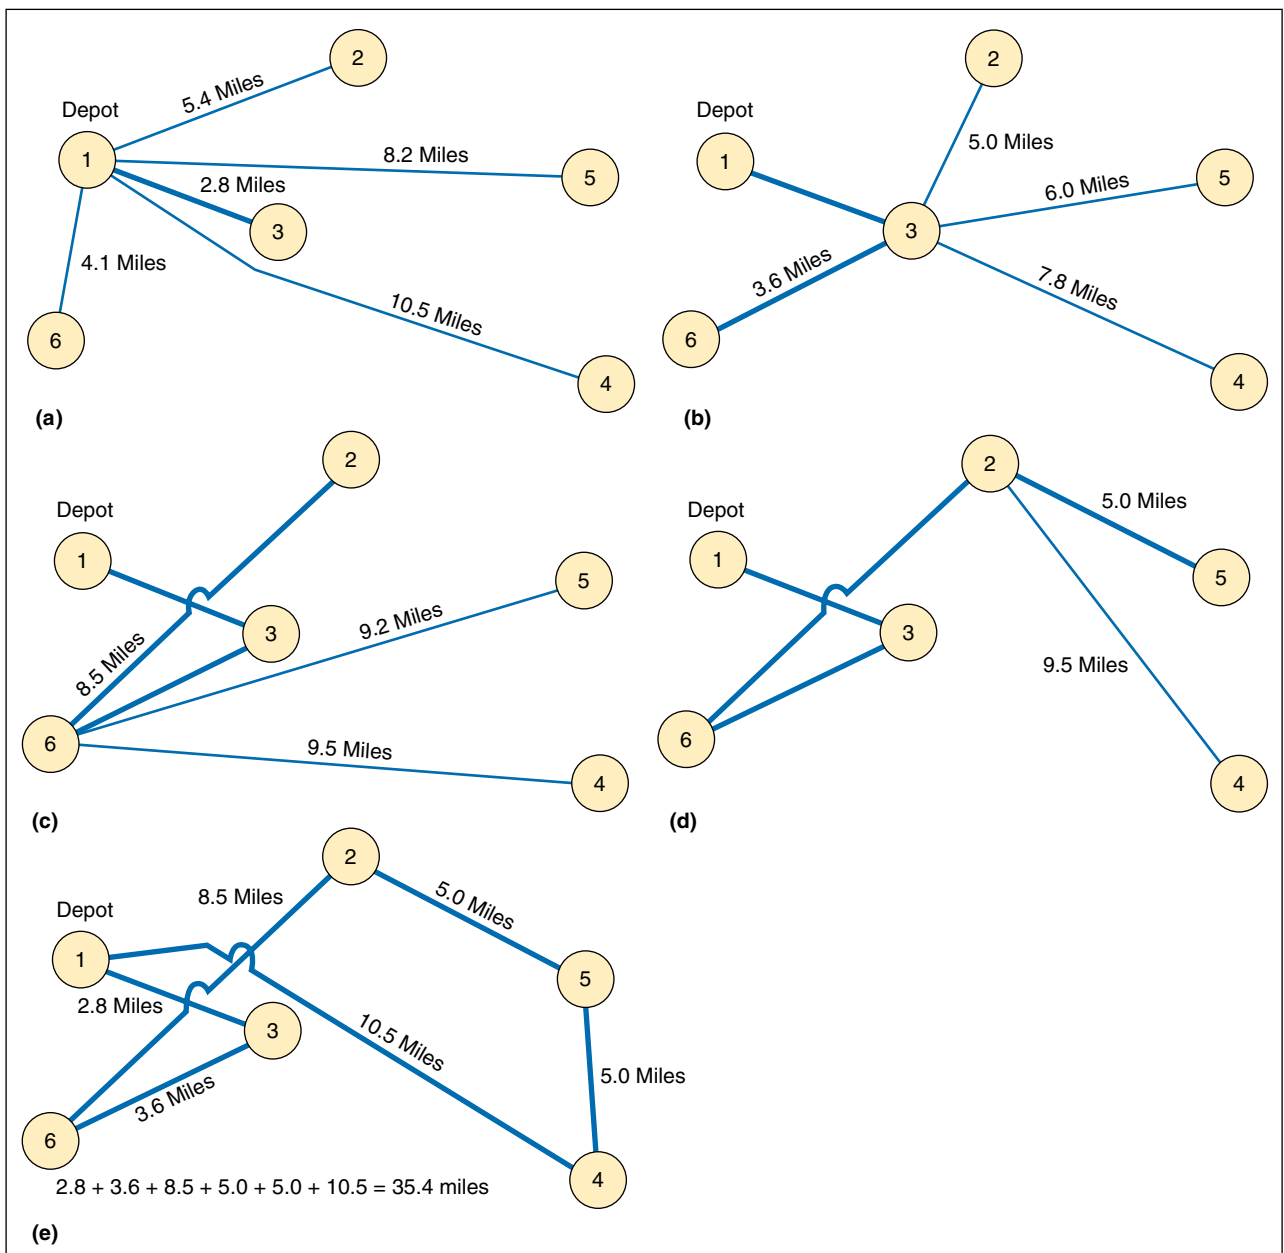
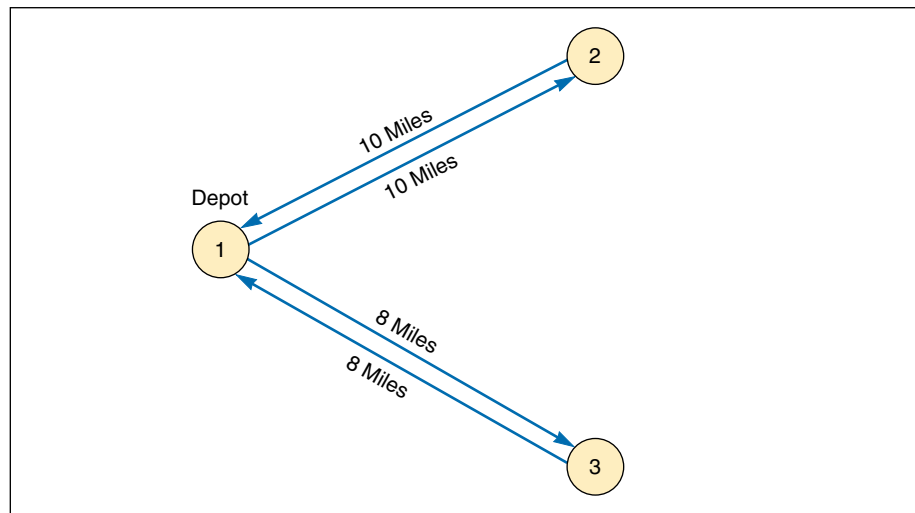


FIGURE T5.3 ■ Nearest Neighbor Procedure

FIGURE T5.4 ■

Initial C&W Network Configuration: Three-Node Problem



But is this the best-possible route? Examine the network again and try to come up with a better tour. How about $1 \rightarrow 2 \rightarrow 5 \rightarrow 4 \rightarrow 3 \rightarrow 6 \rightarrow 1$? The total distance of this tour is 30.9 miles versus 34.5 miles for the nearest neighbor-constructed tour. This result points to the limitation of heuristics; they cannot guarantee optimality. For this small a network, it would be possible to enumerate every possible tour. However, for large problems with 100 to 200 nodes, enumerating every combination would be impossible.

Before leaving the nearest neighbor heuristic, it should be noted that, in practice, the heuristic is applied repeatedly by assigning every node to be the depot node, resolving the problem, and then selecting the lowest-cost tour as the final solution. For example, if we repeat the procedure using node 6 as the depot node, the tour that results is $6 \rightarrow 3 \rightarrow 1 \rightarrow 2 \rightarrow 5 \rightarrow 4 \rightarrow 6$ with a total length of 31.3 miles.

Clark and Wright Savings Heuristic The Clark and Wright savings heuristic (C&W) is one of the most well-known techniques for solving traveling salesman problems. The heuristic begins by selecting a node as the depot node and labeling it node 1. We then assume, for the moment, that there are $n - 1$ vehicles available, where n is the number of nodes. In other words, if we have six nodes in the network, then there are five vehicles available. Each vehicle travels from the depot directly to a node and returns to the depot. Figure T5.4 shows this for a three-node network where the miles are shown on the arcs and the arcs are undirected. The distance from node 2 to node 3 is 5 miles. The total distance covered by the two vehicles in Figure T5.4 is 36 miles: 20 miles for the trip from the depot to node 2 and return, and 16 miles for the trip from the depot to node 3 and return.

But this is not a feasible solution because the objective of a traveling salesman problem is to find a tour in which all nodes are visited by *one* vehicle, rather than by two vehicles, as shown in Figure T5.4. To reduce the number of vehicles needed, we now need to combine the $n - 1$ tours originally specified.

The key to the C&W heuristic is the computation of savings. *Savings* is a measure of how much the trip length or cost can be reduced by “hooking up” a pair of nodes (in the case of Figure T5.4, nodes 2 and 3) and creating the tour $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$, which can then be assigned to a single vehicle. The savings is computed as follows. By linking nodes 2 and 3, we *add* 5 miles (the distance from node 2 to node 3), but we *save* 10 miles for the trip from node 2 to node 1 and 8 miles from the trip from 3 to 1. The total tour length for the complete tour, $1 \rightarrow 2 \rightarrow 3 \rightarrow 1$, is 23 miles. The savings obtained, over the configuration shown in Figure T5.4, is 13 miles. For a network with n nodes, we compute the savings for every possible pair of nodes, rank the savings gains from largest to smallest, and construct a tour by linking pairs of nodes until a complete route is obtained.

A statement of the C&W savings heuristic is as follows:

1. Select any node as the depot node (node 1).
2. Compute the savings, S_{ij} for linking nodes i and j :

$$S_{ij} = c_{1i} + c_{1j} - c_{ij} \text{ for } i \text{ and } j = \text{nodes } 2, 3, \dots, n \quad (\text{T5-1})$$

where

c_{ij} = the cost of traveling from node i to node j .

3. Rank the savings from largest to smallest.
4. Starting at the top of the list, form larger **subtours** by linking appropriate nodes i and j . Stop when a complete tour is formed.

Example Using the C&W Savings Heuristic To demonstrate how the C&W heuristic is used to solve a TSP problem, consider the network shown in Figure T5.5. Here, as in Figure T5.4, we assume that there is one vehicle for every node (excluding the depot) in the network. The solid lines show arcs that are in use as we begin the C&W procedure. The dashed lines show arcs that *may* be used but are not in use currently. Distances, in miles, are shown on the arcs. The savings obtained from linking nodes 2 and 3 is 13 miles. This is computed as (10 miles + 8 miles) – (5 miles). The 10- and 8-mile distances are the lengths of the return trip from nodes 2 and 3, respectively, to the depot; 5 miles is the distance from node 2 to node 3. Similarly, the savings of linking nodes 2 and 4 is 12 miles: (5 miles + 10 miles) – (3 miles). The last pair of nodes to be considered for linking is [4, 3], which yields a savings of 6 miles: (5 miles + 8 miles) – (7 miles).

We next rank the savings for every pair of nodes not yet linked. In order of savings, the pairs are [2, 3], [2, 4], and [3, 4]. The first step in specifying a tour is to link the nodes with the highest savings; nodes 2 and 3. The resulting path is shown in Figure T5.6(a). Proceeding to the next highest savings, nodes 2 and 4 are linked as shown in Figure T5.6(b). The tour is now complete—the last pair, nodes 3 and 4, cannot be linked without “breaking” the tour. The complete tour is 1 → 4 → 2 → 3 → 1, which has a total tour length of 21 miles. The total savings obtained over the “one vehicle per node” configuration shown in Figure T5.5 is 25 miles.

In general, because C&W considers cost when constructing a tour, it yields better quality solutions than the nearest neighbor procedure. Both the Clark and Wright savings heuristic and the nearest neighbor procedure can be easily adjusted to accommodate problems with directed arcs.

Multiple Traveling Salesman Problem

The MTSP is a generalization of the traveling salesman problem where there are multiple vehicles and a single depot. In this problem, instead of determining a route for a single vehicle, we wish to construct tours for all M vehicles. The characteristics of the tours are that they begin and end at the depot node. Solution procedures begin by “copying” the depot node M times. The problem is thus reduced to M single-vehicle TSPs, and it can be solved using either the nearest neighbor or Clark and Wright heuristics.

FIGURE T5.5 ■
Initial C&W Network:
Four-Node Problem

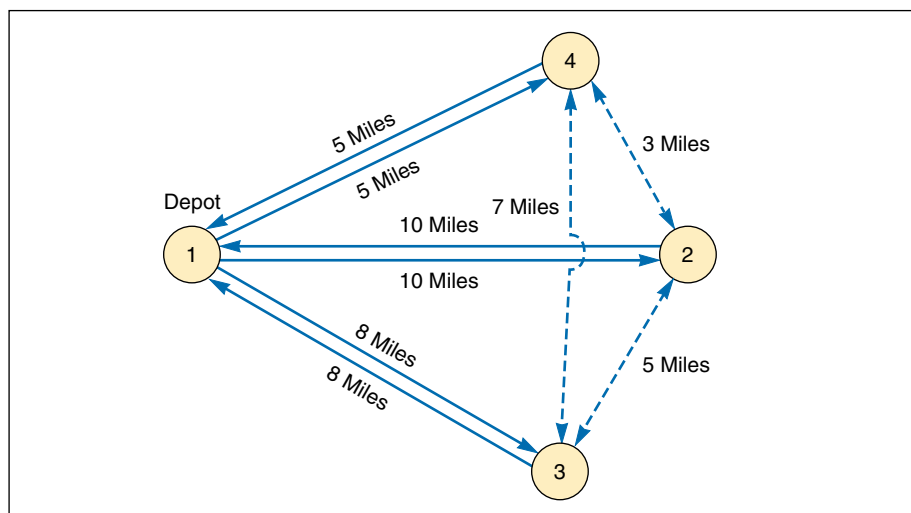
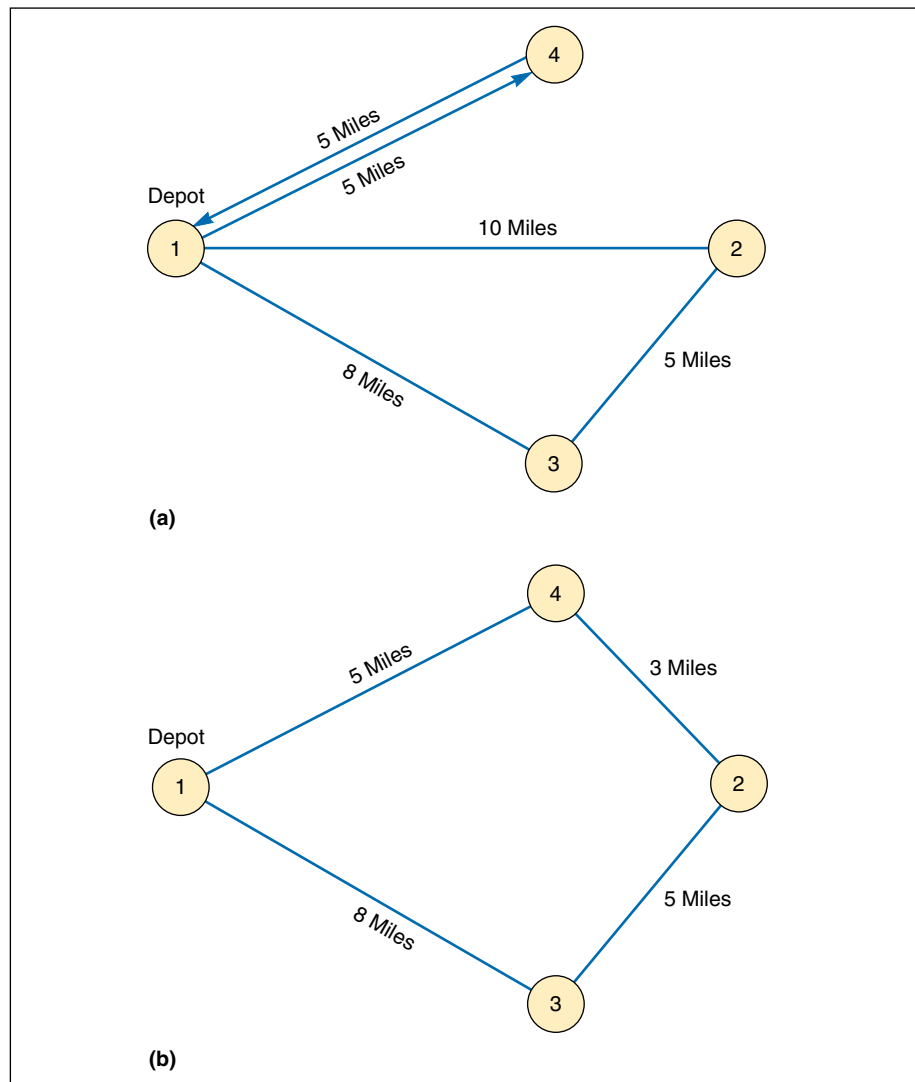


FIGURE T5.6 ■

First and Second Node
Hookups: C&W Heuristic



The Vehicle Routing Problem

The classic VRP expands the multiple traveling salesman problem to include different service requirements at each node and different capacities for vehicles in the fleet. The objective of these problems is to minimize total cost or distance across all routes. Examples of services that show the characteristics of vehicle routing problems include United Parcel Service deliveries, public transportation “pickups” for the handicapped, and the newspaper delivery problem described earlier.

The vehicle routing problem cannot be fully solved with the same procedures as the multiple traveling salesman problem. Consider the simple example illustrated in Figure T5.7. Suppose we have a single depot and two buses, 1 and 2. Vehicle 1 has a capacity of 20 people and vehicle 2 a capacity of 10. There are three nodes where travelers are to be picked up. The number of travelers to be picked up is shown in brackets beside each node.

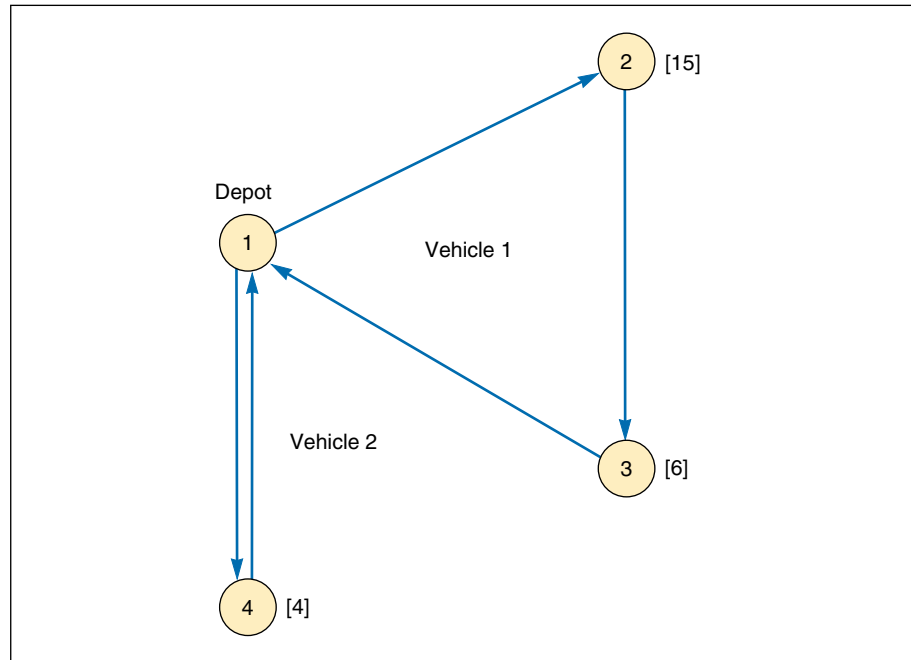
Ignoring for the moment the capacity of the buses and the demand at each node, the Clark and Wright heuristic would construct a tour for each vehicle as follows:

- Bus 1’s tour: 1 → 2 → 3 → 1
- Bus 2’s tour: 1 → 4 → 1

This assignment, however, sends 21 passengers on bus 1, which violates the capacity constraints of bus 1. Thus, this type of problem cannot be solved as a multiple traveling salesman problem. The characteristics of the vehicle routing problem also make it a difficult problem to solve optimally. However, a good heuristic solution can be obtained with the cluster first, route second approach.

FIGURE T5.7 ■

Four-Node Vehicle Routing Problem



Cluster First, Route Second Approach

The **cluster first, route second approach** is best illustrated by an example. Figure T5.8 shows a 12-node problem in which two vehicles must deliver cargo to 11 stations and return to the depot. Cargo demand is bracketed at each node, and distances, in miles, are shown on the arcs. The 12 nodes have been clustered initially into two groups, one for each vehicle. Nodes 2 through 6 are assigned to vehicle 1 and nodes 7 through 12 to vehicle 2. Node 1 is the depot node. In practice, clustering takes into account physical barriers such as rivers, mountains, or interstate highways, as well as geographic areas such as towns and cities that form a natural cluster. Capacity restrictions are also taken into account when developing the clusters. For this example, the capacities of vehicles 1 and 2 are 45 and 35 tons, respectively.

From the initial clustering, vehicle 1 must carry 40 tons and vehicle 2 must carry 34 tons. Both assignments are feasible (i.e., the demands do not exceed either vehicle's capacity). Using the C&W heuristic, a tour is constructed for vehicle 1 (tour 1), $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 6 \rightarrow 1$, with a total tour length of 330 miles. Vehicle 2's tour (tour 2) is $1 \rightarrow 7 \rightarrow 8 \rightarrow 9 \rightarrow 10 \rightarrow 11 \rightarrow 12 \rightarrow 1$. Its length is 410 miles.

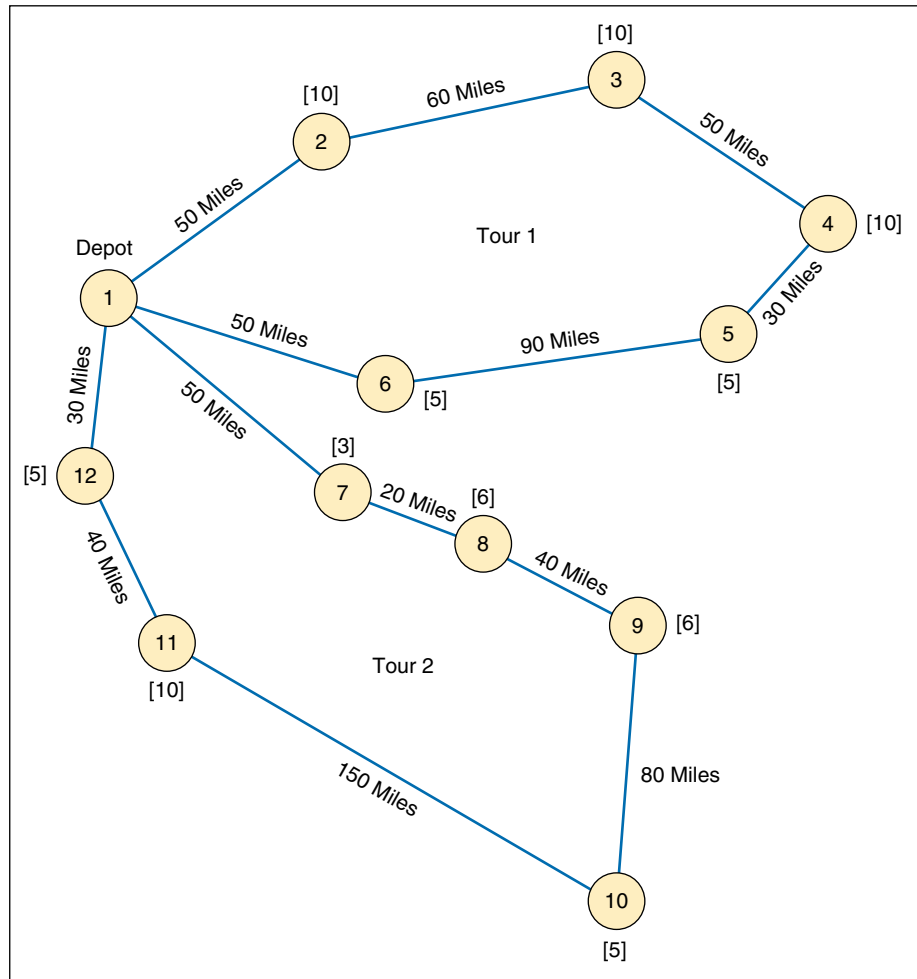
The next phase of the procedure is to determine whether a node or nodes can be switched from the longest tour (tour 2) to tour 1 such that the capacity of vehicle 1 is not exceeded and the sum of the two tour lengths is reduced. This step is referred to as *tour improvement*. We first identify the nodes in tour 2 that are closest to tour 1. These are nodes 7 and 8. Node 8 has a demand of 6 tons and cannot be switched to tour 1 without exceeding vehicle 1's capacity. Node 7, however, has a demand of 3 tons and is eligible to switch. Given that we wish to consider a switch of node 7, how can we evaluate where the node should be inserted into tour 1 and whether it will reduce the distance traveled? Both these questions can be answered by means of the **minimum cost of insertion technique**.

The minimum cost of insertion is calculated in the same way as the Clark and Wright heuristic. If all distances are symmetrical, then the cost of insertion, I_{ij} , can be calculated as follows:

$$I_{ij} = c_{i,k} + c_{j,k} - c_{ij} \text{ for all } i \text{ and } j, i \neq j \quad (\text{T5-2})$$

where c_{ij} = the cost of traveling from node i to node j . Nodes i and j are already in the tour, and node k is the node we are trying to insert. Referring to Figure T5.8, node 7 is a candidate for insertion because it is near tour 1. Node 7 could be inserted between nodes 6 and 1 or between nodes 5 and 6. Both alternatives will be evaluated. In order to calculate the cost of inserting node 7 into tour 1, we

FIGURE T5.8 ■
 Vehicle Routing
 Problem: Initial Solution



require the additional distance information provided in the following table. In practice, this information would be available for all pairs of nodes.

FROM NODE	TO NODE	DISTANCE
1	7	50 miles
6	7	30 miles
5	7	60 miles
1	5	130 miles
1	8	60 miles

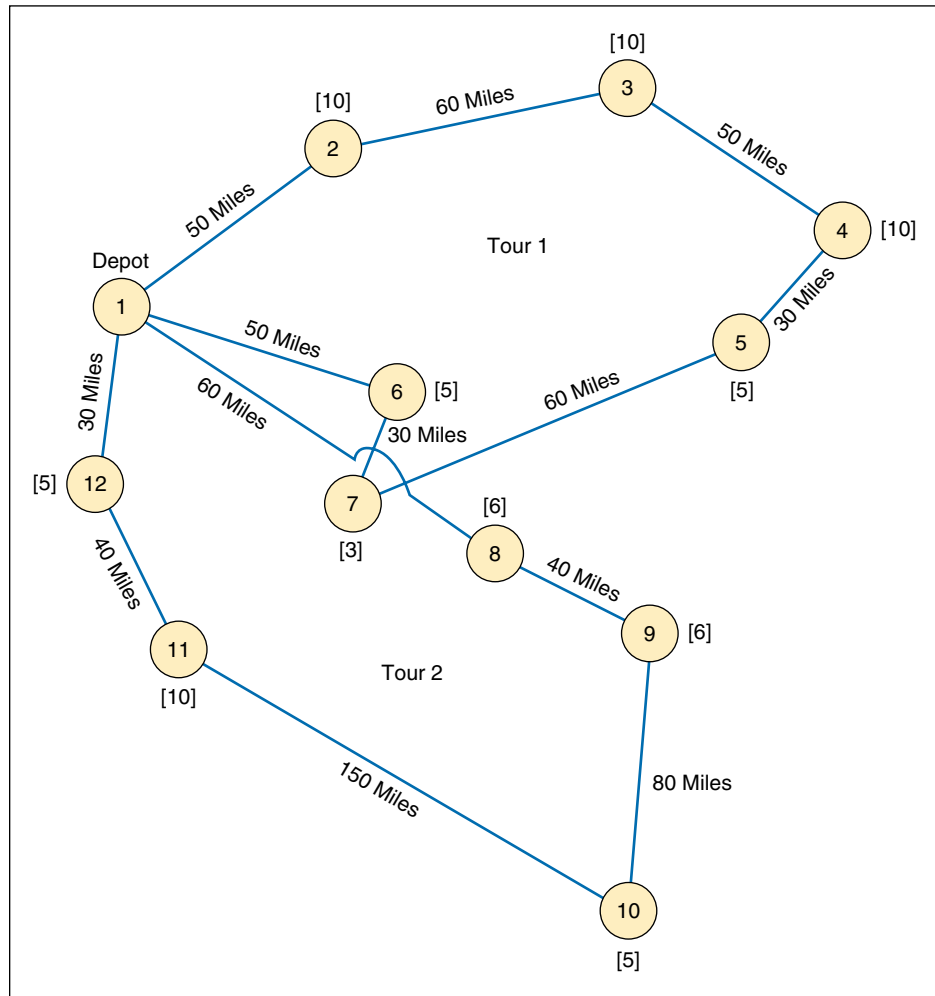
The cost of inserting node 7 between nodes 1 and 6 is 30 miles: $(30 + 50 - 50)$. The cost of inserting the node between nodes 5 and 6 is 0: $(60 + 30 - 90)$. The lowest cost is found by inserting node 7 between nodes 5 and 6, resulting in a completed tour for vehicle 1 of $1 \rightarrow 2 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 7 \rightarrow 6 \rightarrow 1$. Figure T5.9 shows the revised solution. The total length of tour 1 is now 330 miles, and the length of tour 2 is 400 miles. The distance traveled by the two vehicles has decreased from 740 to 730 miles.

SCHEDULING SERVICE VEHICLES

Scheduling problems are characterized by delivery-time restrictions. The starting and ending times for a service may be specified in advance. Subway schedules fall into this category in that the arrival times at each stop are known in advance and the train must meet the schedule. Time windows bracket the service time to within a specified interval. Recall that in the Meals-for-ME program described earlier, meals had to be delivered between 11:30 A.M. and 1:00 P.M.. This is an example of

FIGURE T5.9 ■

Vehicle Routing
Problem: Revised
Solution



a **two-sided window**. A **one-sided time window** either specifies that a service precede a given time or follow a given time. For example, most newspapers attempt to have papers delivered before 7:00 A.M. Furniture delivery is usually scheduled after 9:00 A.M. and before 4:30 P.M. Other characteristics that further complicate these problems include multiple deliveries to the same customer during a week's schedule.

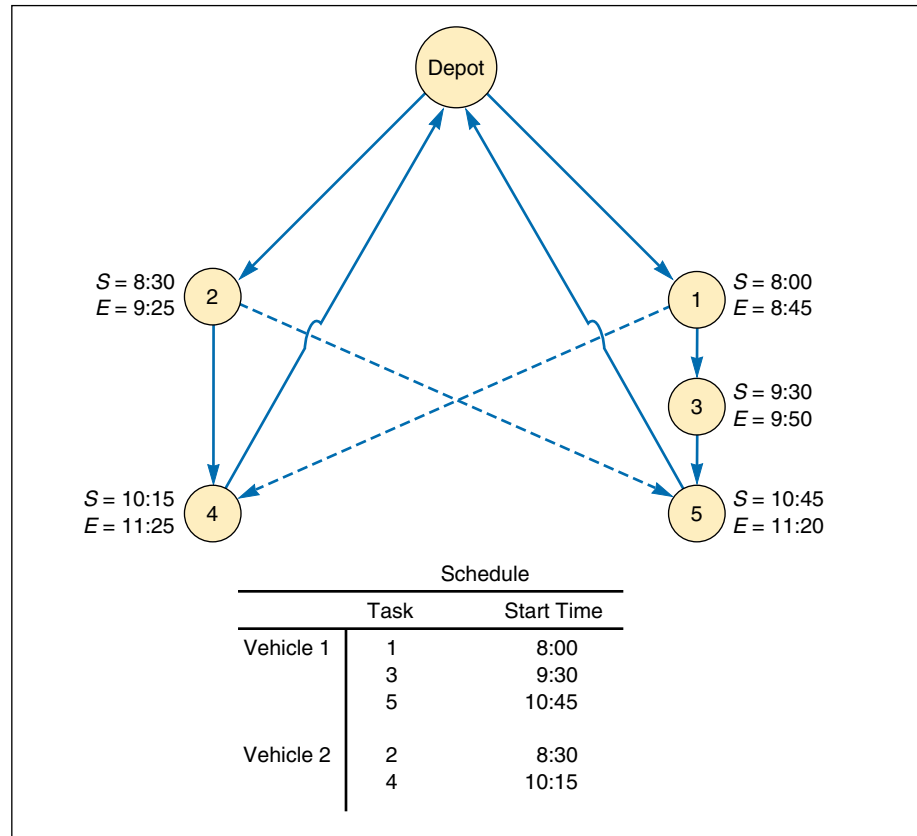
The general input for a scheduling problem consists of a set of tasks, each with a starting and ending time, and a set of directed arcs, each with a starting and ending location. The set of vehicles may be housed at one or more depots.

The network in Figure T5.10 shows a five-task scheduling problem with a single depot. The nodes identify the tasks. Each task has a start and an end time associated with it. The directed arcs mean that two tasks are assigned to the same vehicle. The dashed arcs show other feasible connections that were not used in the schedule. An arc may join node i to node j if the start time of task j is greater than the end time of task i . An additional restriction is that the start time of task j must include a user-specified period of time longer than the end time of task i . In this example, the time is 45 minutes. This is referred to as **deadhead time** and is the nonproductive time required for the vehicle to travel from one task location to another or return to the depot empty. Also, the paths are not restricted in length. Finally, each vehicle must start and end at the depot.

To solve this problem, the nodes in the network must be partitioned into a set of paths and a vehicle assigned to each path. If we can identify the minimum number of paths, we can minimize the number of vehicles required and thus the vehicle capital costs. Next, if we can associate a weight to each arc that is proportional or equal to the travel time for each arc (i.e., the deadhead time), we can minimize personnel and vehicle operating costs as well as time.

FIGURE T5.10 ■

Schedule for a Five-Task Network (S = Start Time, E = End Time)



The Concurrent Scheduler Approach

This problem may be formulated as a special type of network problem called a **minimal-cost-flow problem**. Alternatively, a heuristic approach may be used. One that is simple to use is the **concurrent scheduler approach**. The concurrent scheduler proceeds as follows:

1. Order all tasks by starting times. Assign the first task to vehicle 1.
2. For the remaining number of tasks, do the following. If it is feasible to assign the next task to an existing vehicle, assign it to the vehicle that has the minimum deadhead time to that task. Otherwise, create a new vehicle and assign the task to the new vehicle.

Table T5.3 presents start and end times for 12 tasks. The deadhead time is 15 minutes. The problem is solved using the concurrent scheduler approach. Initially, vehicle 1 is assigned to task 1. Because task 2 begins before vehicle 1 is available, a second vehicle is assigned to this task. Vehicle 2 finishes task 2 in time to take care of task 3 also. In the meantime, vehicle 1 completes task 1 and is available for task 4. A third vehicle is not required until task 5, when vehicles 1 and 2 are busy with tasks 4 and 3, respectively. Continuing in a similar fashion, the schedule for vehicle 1 is 1 → 4 → 7 → 10 → 12, for vehicle 2 the schedule is 2 → 3 → 6 → 9, and for vehicle 3 the schedule is 5 → 8 → 11.

OTHER ROUTING AND SCHEDULING PROBLEMS

Scheduling workers is often concerned with staffing desired vehicle movements. The two are of necessity related in that vehicle schedules restrict staffing options, and vice versa. In general, vehicle scheduling is done first, followed by staff scheduling. This approach is appropriate for services such as airlines, where the cost of personnel is small in comparison to the cost of operating an airplane. It is less appropriate, however, for services such as mass transit systems, where personnel costs may account for up to 80% of operating costs. For such systems it is more appropriate to either schedule personnel first, then schedule vehicles, or to do both at the same time.

TABLE T5.3

Task Times and Schedule
for the Concurrent
Scheduler Example

TASK	START	END	ASSIGN TO VEHICLE
1	8:10 A.M.	9:30 A.M.	1
2	8:15 A.M.	9:15 A.M.	2
3	9:30 A.M.	10:40 A.M.	2
4	9:45 A.M.	10:45 A.M.	1
5	10:00 A.M.	11:30 A.M.	3
6	11:00 A.M.	11:45 A.M.	2
7	1:00 P.M.	1:45 P.M.	1
8	1:15 P.M.	2:45 P.M.	3
9	1:45 P.M.	3:00 P.M.	2
10	2:00 P.M.	2:45 P.M.	1
11	3:00 P.M.	3:40 P.M.	3
12	3:30 P.M.	4:00 P.M.	1

SCHEDULE		
	TASK	START TIME
Vehicle 1	1	8:10 A.M.
	4	9:45 A.M.
	7	1:00 P.M.
Vehicle 2	2	8:15 A.M.
	3	9:30 A.M.
	6	11:00 A.M.
Vehicle 3	5	10:00 A.M.
	8	1:15 P.M.

Problems that have elements of both routing and scheduling are numerous. Examples include school bus routing and scheduling, dial-a-ride services, municipal bus transportation, and the Meals-for-ME program and other meals-on-wheels programs. Certain routing problems also may take on the characteristics of a combined problem. For example, snow plows must clear busier streets prior to clearing less-traveled streets. In addition, there are usually repeated visits depending on the rate of snowfall. These components introduce a scheduling aspect to the routing problem. Considering the fact that there may be literally thousands of variables involved in the formulation of such problems, it becomes apparent that an optimal solution is impossible to obtain. In order to solve real-world problems of this type, management scientists have developed some elegant solution procedures. With rare exception, the procedures use heuristic approaches to obtain “good” but not optimal routes and schedules.

The delivery of emergency services, such as ambulance, police, and fire, is not usually considered a routing or scheduling problem. Rather, emergency services are more concerned with resource allocation (how many units are needed) and facility location (where the units should be located).

SUMMARY

Effective routing and scheduling of service vehicles are two important and difficult problems for managers of services. The consequences of poor planning are costly, and a decision maker must frequently fine-tune the system to ensure that the needs of the customer are being met in a timely and cost-effective fashion. The criterion used to measure the effectiveness of service delivery depends on the type of service. Although minimizing total cost is an important criterion, for some services, criteria such as minimizing customer inconvenience and minimizing response time may be equally if not more important.

Solution of routing and scheduling problems begins with a careful description of the characteristics of the service under study. Characteristics, such as whether demand occurs on the nodes or the arcs, whether there are delivery-time constraints, and whether the capacity of the service vehicles is a concern, determine the type of problem being considered. The type of problem then determines the solution techniques available to the decision maker.

This tutorial discussed the characteristics of routing problems, scheduling problems, and combined routing and scheduling problems. Optimal solution techniques for these types of problems are generally based on mathematical programming. However, in practice, a good but perhaps

nonoptimal solution is usually sufficient. To obtain a good solution, several heuristic solution approaches have been developed. Two well-known heuristics for solving the traveling salesman problem were presented, the nearest neighbor procedure and the Clark and Wright savings heuristic. Also presented was the minimum cost of insertion technique for use in solving the vehicle routing problem.

KEY TERMS

- Networks (*p. T5-3*)
- Nodes (*p. T5-3*)
- Depot node (*p. T5-3*)
- Arcs (*p. T5-3*)
- Undirected arcs (*p. T5-3*)
- Directed arcs (*p. T5-3*)
- Tour (*p. T5-3*)
- Feasible (*p. T5-3*)
- Route (*p. T5-3*)
- Schedule (*p. T5-3*)
- Traveling salesman problem (TSP) (*p. T5-4*)
- Multiple traveling salesman problem (MTSP) (*p. T5-4*)
- Vehicle routing problem (VRP) (*p. T5-4*)
- Chinese postman problem (CRP) (*p. T5-4*)
- Routing (*p. T5-4*)
- Scheduling (*p. T5-4*)
- Nearest neighbor procedure (*p. T5-5*)
- Clark and Wright savings heuristic (*p. T5-5*)
- Partial tour (path) (*p. T5-6*)
- Path (*p. T5-6*)
- Subtours (*p. T5-8*)
- Cluster first, route second approach (*p. T5-10*)
- Minimum cost of insertion technique (*p. T5-10*)
- Two-sided window (*p. T5-12*)
- One-sided window (*p. T5-12*)
- Deadhead time (*p. T5-12*)
- Minimal-cost-flow problem (*p. T5-13*)
- Concurrent scheduler approach (*p. T5-13*)

DISCUSSION QUESTIONS

1. Compare the characteristics of the following types of problems:
 - (a) Routing problems
 - (b) Scheduling problems
 - (c) Combined routing and scheduling problems
2. Describe the differences between and give an example of:
 - (a) A traveling salesman problem
 - (b) The Chinese postman problem
 - (c) A vehicle routing problem
3. A mail carrier delivers mail to 300 houses in Blacksburg. The carrier also must pick up mail from five drop boxes along the route. Mail boxes have specified pickup times of 10:00 A.M., 12:00 noon, 1:00 P.M., 1:30 P.M., and 3:00 P.M. daily. Describe the characteristics of this problem using the information provided in Figure T5.1. What types of service-time restrictions apply?
4. Define each of the following:
 - (a) Deadhead time
 - (b) Depot node
 - (c) Undirected arc
5. Describe what is meant by
 - (a) A feasible tour for a vehicle routing problem
 - (b) A feasible tour for a traveling salesman problem
 - (c) A two-sided time window
 - (d) A node precedence relationship
6. Discuss the differences between the nearest neighbor procedure and the Clark and Wright savings heuristic procedure for constructing a tour.
7. Discuss under what circumstances a distance or cost matrix in a routing problem would be asymmetrical.
8. What are some objectives that might be used to evaluate routes and schedules developed for
 - (a) School buses
 - (b) Furniture delivery trucks
 - (c) Ambulances
9. What are some practical problems that might affect the routing and scheduling of
 - (a) A city's mass transit system
 - (b) A national trucking fleet
 - (c) Snow plows
10. What is the "savings" in the Clark and Wright savings heuristic?

PROBLEMS

- T5.1** Use the Clark and Wright savings heuristic procedure, and the data that follow, to compute the savings obtained by connecting
- a) 2 with 3
 - b) 3 with 4
 - c) 2 with 5

FROM NODE	TO NODE (DISTANCES IN MILES)			
	2	3	4	5
1	10	14	12	16
2	—	5	—	18
3	5	—	6	—

T5-16 ONLINE TUTORIAL 5 VEHICLE ROUTING AND SCHEDULING

T5.2 Assume that a tour $1 \rightarrow 3 \rightarrow 5 \rightarrow 1$ exists and has a total length of 23 miles. Given the distance information that follows and using the minimum cost of insertion technique, determine where node 2 should be inserted.

FROM NODE	TO NODE	DISTANCE
1	3	6
1	5	9
3	5	8
1	2	5
2	3	7
2	6	5
2	5	8

T5.3 A vehicle-routing problem has 20 nodes and two vehicles. How many different routes could be constructed for this problem?

T5.4 Given the distance matrix for a traveling salesman problem shown in Table T5.4,
 a) Assume node 1 is the depot node, and construct a tour using the nearest neighbor procedure.
 b) Assume the depot is node 4, and construct a tour using the nearest neighbor procedure.

TABLE T5.4

FROM NODE	DISTANCE TO NODE (IN MILES)							
	1	2	3	4	5	6	7	8
1	—	2.2	5.8	4.0	5.0	8.5	3.6	3.6
2	2.2	—	4.1	3.6	5.8	9.4	5.0	5.8
3	5.8	4.1	—	3.2	6.1	9.0	6.7	9.2
4	4.0	3.6	3.2	—	3.0	6.3	3.6	6.7
5	5.0	5.8	6.1	3.0	—	3.6	2.0	6.0
6	8.5	9.4	9.0	6.3	3.6	—	3.6	8.5
7	3.6	5.0	6.7	3.6	2.0	3.6	—	4.0
8	3.6	5.8	9.2	6.7	6.0	8.5	4.0	—

T5.5 Using the Clark and Wright savings heuristic, construct a tour for the data given in the distance matrix for Problem T5.4. Assume node 1 is the depot node.

T5.6 You have been asked to route two vehicles through a 10-node network. Node 1 is the depot node; nodes 2 through 5 have been assigned to vehicle 1 and nodes 6 through 10 to vehicle 2. The cost matrix for the network is given in Table T5.5.

- a) Construct the two tours using the nearest neighbor procedure and state the total cost of the tour.
- b) Construct the two tours using the Clark and Wright savings heuristic and state the total cost of the tour.

TABLE T5.5

FROM NODE	COST TO NODE (\$)									
	1	2	3	4	5	6	7	8	9	10
1	—	22	22	32	32	14	45	56	51	35
2	22	—	32	22	54	36	67	78	67	41
3	22	32	—	22	36	41	42	67	70	64
4	32	22	22	—	56	51	71	86	83	63
5	32	54	36	56	—	32	10	32	45	54
6	14	36	41	51	32	—	40	45	32	32
7	45	67	42	71	10	40	—	20	42	71
8	56	78	67	86	32	45	20	—	32	71
9	51	67	70	83	45	32	42	32	—	45
10	35	41	64	63	54	32	71	71	45	—

- : **T5.7** Referring to Problem T5.6, assume vehicle 1 has a capacity of 35 passengers and vehicle 2 a capacity of 55 passengers. The number of passengers to be picked up at each node is

NODE	NUMBER OF PASSENGERS
2	10
3	10
4	5
5	5
6	5
7	5
8	20
9	10
10	5

Using the tours constructed in Problem T5.6, attempt to improve the total cost of the two tours using the minimum cost of insertion technique.

- : **T5.8** Convert the distance matrix given in Problem T5.4 to a cost matrix using the following information. The cost of routing a vehicle from any node *i* to any node *j* is \$100. This is a fixed cost of including a link in a tour. The variable cost of using a link (or arc) is \$3.30 per mile for the first 5 miles and \$2.00 for the remainder of the arc distance. After computing the cost matrix, resolve the problem using the Clark and Wright savings heuristic.

- : **T5.9** Using the task times provided below, determine the number of vehicles required and the task sequence for each vehicle using the concurrent scheduler approach. The deadhead time is 30 minutes.

TASK	START	END
1	8:00 A.M.	8:30 A.M.
2	8:15 A.M.	9:15 A.M.
3	9:00 A.M.	9:30 A.M.
4	9:40 A.M.	10:20 A.M.
5	10:10 A.M.	11:00 A.M.
6	10:45 A.M.	11:30 A.M.
7	12:15 P.M.	12:40 P.M.
8	1:30 P.M.	1:50 P.M.
9	2:00 P.M.	2:40 P.M.
10	2:15 P.M.	3:30 P.M.

CASE STUDY

Routing and Scheduling of Phlebotomists

Phlebotomists are clinical laboratory technicians who are responsible for drawing blood specimens from patients in the hospital. Their routine responsibilities include drawing samples for laboratory tests ordered that are to be completed on that day by the day crew. A 500-bed medical center usually employs five to seven technicians in this capacity. The morning pickups are made between 6:30 A.M. and 8:00 A.M. On a given morning there may be requests to draw blood samples from 120 to 150 patients. The time required to draw the blood necessary to complete a physician's order varies depending on age, physical condition, and the number of different types of tests required of a patient. For example, a single phlebotomist may be able to draw samples for 20 maternity patients in 90 minutes, since most of these women are healthy and do not require "unusual" types of blood work. However, that same phlebotomist may only be able to draw blood from eight critically ill patients, who usually require more varied tests and may, because of their physical condition,

require more time. The same limitation is true for infants and small children who require special collection techniques due to their size.

In addition to their routine pickups, which must be completed within the 90-minute preshift interval, there are routine specimens that must be drawn at a specified time. These timed specimens include fasting specimens (such as blood glucose tests), which must be collected before the patient eats, and *blood gases*, which are collected 30 minutes after a patient has received a respiratory treatment. With either of these tests, there is a margin for "error" of 15 minutes. Generally, more routine tests are also collected along with the timed specimens.

The medical center has five floors, each of which specializes in a particular type of patient. For example, one floor may handle surgical patients and another orthopedic patients. In addition, there are special sections, including the nursery, the pediatric floor, and the intensive care unit. Because of the location of the respiratory equipment and monitors, all patients requiring daily blood gases are located in intensive care.

It is the task of the chief phlebotomist to estimate the number of phlebotomists needed on a given day and to assign patients to technicians such that all deliveries are made before the start of the day shift

(continued)

and timed specimens are collected within a 15-minute window of the specified time.

Discussion Questions

1. What characteristics of routing and scheduling are exhibited in this problem?
2. What type of data would you need to collect in order to most effectively schedule technicians?
3. Does a deadhead time exist in this situation? If so, where?
4. If you were to view this as a cluster first, route second situation, based on what criteria would you form clusters?
5. Suggest how you would solve the problem if the timed specimens and routine pickups were considered separately.



BIBLIOGRAPHY

- Bodin, Lawrence, Bruce Golden, Arjang Assad, and Michael Ball. "Routing and Scheduling of Vehicles and Crews: The State of the Art." *Computers and Operations Research* 10, no. 2 (1983): 63–211.
- Fitzsimmons, J. A., and R. S. Sullivan. "Service Vehicle Scheduling and Routing." In *Service Operations Management*. New York: McGraw-Hill (1982): 312–336.
- Hall, R. W. "Change of Direction." *OR/MS Today* 29, no. 4 (February 2002): 38–41, 46.
- Partyka, J. G., and R. W. Hall. "On the Road to Service." *OR/MS Today* 27, no. 4 (August 2000): 26–30.
- "Vehicle Routing Software Survey." *OR/MS Today* 29, no. 4 (February 2002): 42–45, 47.