

# Applying GIS and OR Techniques to Solve Sears Technician-Dispatching and Home-Delivery Problems

---

Don Weigel

*Transportation/Logistics Services  
ESRI, Inc.  
380 New York Street  
Redlands, California 92373*

Buyang Cao

*Transportation/Logistics Services  
ESRI, Inc.*

---

Sears, Roebuck and Company uses a vehicle-routing-and-scheduling system based on a geographic information system to run its delivery and home service fleets more efficiently. Although the problems to be solved can be modeled as vehicle-routing problems with time windows (VRPTW), the size of the problems and thus practical complexity make these problems of both theoretical and practical interest. We constructed a series of algorithms, including the algorithm to build the origin-and-destination matrix, the algorithm to assign resources, and algorithms to perform sequencing and route improvement. The combination of GIS and OR techniques makes the system quite efficient. The system has improved the Sears technician-dispatching and home-delivery business; resulting in over \$9 million in one-time savings and over \$42 million in annual savings. The success of this application also suggests a promising link between GIS and OR techniques.

Sears, Roebuck and Company began with an enterprise established in 1886 and was incorporated under the laws of New York in 1906. Sears and its consolidated subsidiaries conduct domestic and

international merchandising and credit operations. The company, a multiline retailer, is among the largest retailers in the world selling merchandise and services. US domestic operations include merchan-

dising and credit operations in the United States and Puerto Rico, and consist primarily of providing goods and services from the company's retail stores, home services, and direct-response marketing businesses. International operations consist of merchandising and credit operations conducted through majority-owned subsidiaries in Canada and Mexico. The company employs approximately 320,000 people worldwide.

ESRI is a privately held corporation selling software with estimated annual revenue of \$250 million. It provides a range of geographic information system (GIS) software packages and related services to clients around the world. Headquartered in Redlands, California, ESRI has regional offices throughout the United States, several subsidiary companies overseas, and distributors in 91 countries. Founded in 1969, ESRI pioneered the development of cartographic data structures and specialized GIS software tools.

### **The Problem**

Sears logistics services (SLS) manages a US fleet of over 1,000 delivery vehicles that includes contract carriers and Sears-owned vehicles. Sears provides the largest home-delivery service of furniture and appliances in the United States, with over 4 million deliveries a year of 21,000 unique items. When we began the project, SLS had 46 routing offices serving 70 percent of the US population, each office responsible for a designated delivery region. When a customer asks for a delivery, Sears determines the day and estimated time window based on customer desire and the delivery schedule in the region where the customer is located. One day before the delivery,

Sears creates the routes for the next day's deliveries based on merchandise types, the quantity of merchandise, the delivery vehicles available, customer time windows, and so forth. It may modify these routes because of delivery cancellations or new next-day deliveries. Once it has finalized the routes, SLS center personnel call the affected customers to confirm the deliveries and their time windows. The routing offices try (1) to provide customers with accurate and convenient time windows for deliveries, (2) to minimize operational costs, and (3) to give drivers consistent routes.

Sears product services (SPS) operates a US fleet of 12,500 service vehicles and the associated technicians, who repair and install appliances and provide home improvements and homeowner services. The SPS call center receives 15 million calls for on-site service annually. This is the largest home-service business in the United States (a \$3 billion business in the \$160 billion industry) and the sixth largest fleet in North America. Home service is also a key growth engine for Sears. Like SLS, SPS partitions the whole country into regions, with an SPS regional office covering each region. When a customer calls an SPS call center with a request for home service, the customer representative assigns the customer a service date and time window based on the customer's wants and the working schedule in the customer's area. One day before the service date, the regional office builds routes based on the customer requests, the available technicians, their skills, and their schedules. When it has finalized the routes, it confirms the times to provide service with the

customers. However, it may revise a route further on the day the service is provided because of emergency services or technician schedule changes. SPS tries to plan the routes for services so that (1) it maximizes the completion of service calls on first attempt, (2) it minimizes operational costs, and (3) it enhances customer service. Table 1 summarizes the Sears home-delivery and technician-dispatching businesses.

**The New System**

The purpose of the new systems ESRI developed for SLS and SPS is to enhance Sears' existing mainframe-based delivery order system (DOS) and national product services (NPS) systems to consolidate operations, improve services, and reduce costs. The new system should be able to deal with the complicated home-delivery and home-service businesses effectively and also provide seamless data communications with the existing system. We call

the SLS home delivery system the *Enhanced Home Delivery System* (EHDS) and the SPS technician dispatch system the *Computer Aided Routing System* (CARS). The EHDS/CARS system incorporates geographic information system (GIS) techniques that provide spatial data-processing capabilities, which are impossible within traditional tabular data-processing environments (Figure 1).

EHDS/CARS is UNIX based and operates on either a central server or distributed workstations. In either case, the system interfaces with Sears mainframes to receive delivery or service orders. After routing and scheduling, the system uploads vehicle and driver assignments, stop sequences, travel times, arrival times, and other data to mainframe databases. EHDS is implemented on distributed IBM RS/6000 workstations with multiple X-terminals in each routing facility. CARS is implemented on a single IBM RS/6000 SP

	Sears Logistics Services	Sears Product Services
Vehicles or personnel	1,000+ consisting of contract carriers and Sears-owned trucks	12,500 service technicians
Annual stops	4+ million	15 million
Service area	Regional delivery center based	Regional service center based
Business objective	To deliver furniture and appliances	To provide repair, installation, home improvements, and homeowner service
System objectives	Improve customer satisfaction	Increase completed service calls on first attempt
	Reduce operational costs	Improve customer service
	Consolidate delivery operations	Provide same day service
	Plan consistent routes	Reduce operational costs
Algorithm objectives	Automatically build routes that reduce travel time while honoring side constraints.	Consolidate dispatch operations
		Automatically build routes that reduce travel time while honoring side constraints.

**Table 1: Although the Sears technician-dispatching and home-delivery businesses differ in the size of their problems, they share many of the same objectives.**

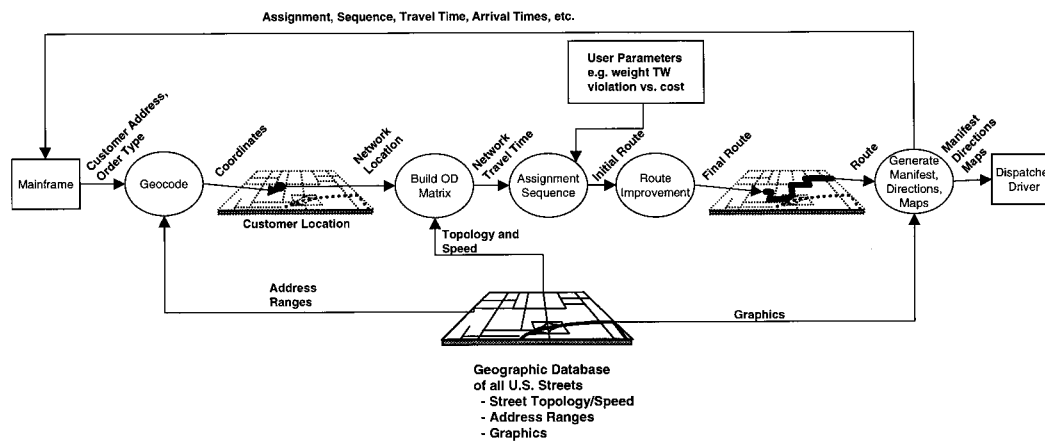


Figure 1: The EHDS or CARS system contains several components for building routes. When the customer data are downloaded from the mainframe, the geocode module will locate the customer locations ( $x$ -,  $y$ -coordinates). Based on the street information stored in a GIS, it calculates the distances between all customers, providing the basic information to the assignment and route improvement modules. Using distance and other customer information, such as imposed time windows and specialties, the assignment and route-improvement modules generate the final routes for home delivery or home service.

with multiple CPUs. Users access CARS using X-emulation over a wide-area network.

The software environment consists of ESRI's ARC/INFO GIS along with additional programs ESRI developed to implement the algorithms. ARC/INFO is interfaced with Sears mainframe and distributed databases, including DB2 and INFORMIX. The GIS integrates Sears customer data and commercially available street network data, and provides the necessary flow of control for the algorithms that solve the VRPTW. The solver was implemented in the C/C++ programming language. The system consists of the following modules:

An address matching and geocoding module matches the address of a customer requiring service, delivery, or pickup to a node in the geographic street database. A digital map layer in the ARC/INFO GIS

environment is called a *coverage*. A coverage contains both the location data and descriptive data for streets (or other map features) in a given geographic area.

A view environment module is fully integrated with the GIS and has standard GIS display capabilities. It can (1) display all roads and identify or list attribute information about them; (2) display work (service or delivery) areas and their text; (3) automatically zoom to a selected route and draw the stops on that route; (4) provide driving directions; (5) view all routes, seed points, and stops; and (6) select an individual technician or driver and assign him or her to a new seed point location. (A *seed point* is the geographic centroid of the desired working area.)

The pre-edit module allows dispatchers to manually modify any day's service or delivery-order assignments prior to batch optimization, but it allows the computer to

seek the optimal expression of those interventions. The dispatcher can preassign a service or delivery order to a particular technician or driver. The module will automatically check the feasibility of the preassignment and, if necessary, warn the dispatcher. The module can also be used to update the attributes of a stop, such as service time, service request, and special instructions.

---

### Sears employs approximately 320,000 people worldwide.

---

The assignment rules module allows dispatchers, prior to running the batch optimization, to set rules that dictate how technicians or drivers will be assigned to service or delivery orders. These rules are customizable. For example, they could (1) allow the system to assign any necessary overtime to technicians or drivers, or assign overtime only to technicians or drivers who are preauthorized to work overtime; (2) assign technicians with primary and secondary specialties to service orders considering only their primary specialty or both their primary and secondary specialties; (3) use average travel time (distance) from stop to stop in determining maximum impedance in calculating origin and destination matrix; and (4) cover additional rules required by the client.

The routing module consists of all solution methods and solves the vehicle-routing problems with time windows (VRPTW) based on the input parameters (areas, customer information, technician or driver schedules, preassigned stops, and assignment rules).

The graphic route-editing module can

be used by the dispatcher to manually fine-tune the routes resulting from the batch routing process, for instance, moving stops between routes. When the dispatcher moves a stop from one route to another, the module will resequence the affected routes.

The routing options module allows the dispatcher to review and evaluate the routes resulting from the routing procedure. Using its tools, a dispatcher can easily identify some exceptional cases. He or she can then use other modules to edit the routes, change the technician or driver profiles (schedules), change assignment rules, and resequence the routes.

The reporting module generates on-line and hard-copy reports, for example, (1) a route report containing such information as the route number, technician's or driver's name and identification number, the number of stops, starting time and transit time, total mileage, total service time, and overtime; (2) a time-window-violation report for any failure to meet the time window for a stop; (3) a stop report containing information about individual stops, such as the route number, the visiting sequence number, the service or delivery order number, the time window imposed, the estimated arrival time, the service time, the transit time from the previous stop, and the parts needed; (4) a direction report listing the street-level directions for the selected route; and (5) a routing-summary report containing information about an entire day's routes (number of routes, number of stops, total mileage, total travel time, total dead time including waiting and free times, average number of stops per route, and estimated

cost).

The GUI allows users to visually execute geographic data processes, such as dropping stops from routes, moving stops between routes, and adding new street segments; evaluating routing results; and interacting with the optimization procedures to solve VRPTW.

### Optimization Techniques

Sears' problems of dispatching technicians and scheduling deliveries can be modeled as vehicle-routing problems with time windows (VRPTW). Many researchers have worked on these problems (for instance, Ball et al. [1995] and Bodin et al. [1983]). They have developed a variety of algorithms (exact and heuristic methods) to solve these problems (for examples, see Desrochers et al. [1992] and Solomon [1987]). The basic VRPTW approach is to determine  $M$  vehicle routes, where a route is a tour run by a vehicle starting at the depot, visiting a subset of customers in a certain sequence, and returning to the depot. Each customer must be visited by exactly one of  $M$  vehicles, and the size (weight, volume) associated with customers must not exceed the capacity of the vehicle. Furthermore, the time windows customers imposed should be met. The routes should be built to minimize the total cost (travel time, distance, or other costs). For practical reasons, we need to consider factors other than those in the basic model: (1) In dispatching technicians, we must assign a technician to skills; (2) we must consider employee schedules, such factors as working hours, days off, training, and breaks; (3) we must consider technicians' starting and ending locations, which may be different; (4) in dispatching technicians,

we must include parts depots in some routes; (5) in scheduling deliveries, we must impose precedence constraints so that the merchandise is first picked up at certain predefined stores; (6) we must sometimes predefine the areas technicians or truck drivers serve to ensure, for instance, that they speak the dominant language in the area, or that they become familiar with certain areas, or that a union's request that senior employees be assigned to favored areas is honored, or that carrier contracts are assigned to their contracted areas; (7) we must assign some service orders to specific technicians because of their specialties, their knowledge, or the preference of the customer; (8) we must restrict total route time, including travel and on-site times, if the technician or driver is not authorized for overtime; (9) we must consider the number of service orders or delivery orders on a route to honor contract or union agreements that the employee services at least  $n$  customers, or no more than  $m$  customers, or both; (10) we must take account of exceptional geographic regions in which customers can be visited only within a specified time period during a day; and (11) we must ensure that the travel time between any pair of locations on a route does not exceed some predefined amount.

Based on our analysis of the problem, we developed an objective function that includes travel time (distance), the duration of routes, the penalties of time-window violation, and waiting (free) time. The time-window constraints are soft. Users can modify the weights, which provides dispatchers flexibility in building routes and allows regional managers to

achieve specific local objectives.

This particular VRPTW is a very difficult combinatorial optimization problem. Furthermore the problem size (the underlying street network can contain two million or more arcs and the number of stops can be two thousand or more) makes the problem very complex. To cope with the problem's complexities and complete the solution within a reasonable time (less than one hour), ESRI developed and implemented a series of algorithms based on heuristic strategies. The solution procedure can be viewed as a "cluster-first, route-second" method, in which the service orders or deliveries are assigned to technicians or drivers first, and then each route is improved by applying the algorithms. The core solver is composed at the following major components:

- Origin-destination (OD) matrix

- construction;

- Resource assignment;

- Sequence and route improvement.

These algorithms are able to solve the technician dispatch and home-delivery problems. We address the different requirements of these two VRPTWs through problem-specific constraint sets and weights in the objective function.

#### **Constructing the Origin and Destination (OD) Matrix**

An OD matrix containing the travel time (distance) from one customer location to another provides the primary data for the resource assignment, sequencing, and route-improvement procedures. We obtain the travel time (distance) by applying a shortest-path algorithm to a GIS street-network database. In this case, we can prevent such unrealistic scenarios as routes

crossing geographic barriers (mountains, bodies of water areas with no road access). Furthermore, accurate distances are crucial for the solution quality. For instance, although two locations in a mountain area seem close based on their distance by air (Euclidean distance), it may take a long time to drive from one to the other. Sears uses ESRI's GIS software to integrate the client's existing customer information with map data provided by GDT to obtain accurate and realistic travel times (distances) between pairs of locations. The GIS software provides the data needed to build the OD matrices, for instance, street information, including speed limits, one-way streets, and barriers; vehicle starting and ending locations; depot locations; and order locations.

For a home-delivery problem, each delivery center usually has several hundred customers to visit, and each truck driver starts and ends at the delivery center. It is not difficult to create an OD matrix that contains the distances between pairs of customers and between customers and the delivery center. However, it may be impractical to build an OD matrix that contains the distances for all technicians and service orders within a region. On the one hand, the size of a technician-dispatching problem (for example, several thousand service orders and several hundred technicians) presents the computation of an entire OD matrix in a timely fashion. On the other hand, many distances obtained in this way may be useless, for instance, a distance between a technician starting location and a service order for which the technician does not have the necessary skills. We use the information on skill sets

(specialties), desired working areas (usually we represent working areas by their centroids or seed points), maximum travel times, and so forth to create several smaller and more manageable OD matrices for use in the assignment and routing procedures. We use the following parameters to build OD matrices efficiently:

- The technician's skills in repairing or installing (for example, televisions, washers, or computers);
- Minimum impedance, the minimum time (distance) a technician must travel;
- Maximum impedance, the maximum allowed time (distance) a technician can travel from one location to another;
- Minimum candidate, the minimum number of locations the technician should reach from a starting location;
- Maximum candidate, the maximum number of locations the technician can reach from a starting location.

The purpose of setting these parameters is to provide enough distances so that the assignment and routing procedures can create reasonable routes in a reasonable amount of computational time. The OD matrix used to solve the VRPTW for the technician-dispatching problem consists of two parts: OD1 and OD2.

OD1 contains the distances from a technician's starting point, ending point, and seed point (the centroid of the technician working area) to the service orders the technician can serve (skill consideration). The matrix also includes the distances between the same technician's starting and ending locations as well as the seed point. While computing the distances, the algorithm takes the parameters listed above into account. For each point (starting

point, ending location, or seed point of a technician) the algorithm computes the distances to potential service orders calling for his or her skills until the maximum candidate or maximum impedance value is reached. If it reaches the maximum candidate value, the algorithm will check whether it has met the minimum impedance value. If it has not, it will increase the maximum candidate value by  $n$  (user defined), and the process will continue; otherwise that ends the distance calculation for one technician. Likewise, if the algorithm meets the maximum impedance, it will verify whether it has also reached the minimum candidate value. If it has, the computation for one technician is terminated; otherwise it will increase the maximum impedance by  $x$  (user defined) units, and the process continues. In both cases, the upper bounds are set for maximum impedance and maximum candidate. Once the algorithm reaches the upper bound, the distance calculation for one technician terminates. Using this algorithm, we can collect service orders within the particular technician's skill set starting at the technician's seed point until he or she travels the network for a specified distance (maximum impedance) or collects a specified number of service orders (maximum candidates).

While OD1 includes the distances between a technician's starting point, ending point, and seed point and the service orders within his skill set, OD2 contains the distances between service orders. In addition to the parameters mentioned above, the procedure to build OD2 also utilizes information on the skill set on the day the VRPTW is solved.



For any service order, let its specialty be  $s$ . Then

$S_i$  = the skill set of technician  $i$ ; and

$S = \{\cup S_i \mid s \in S_i\}$ , that is, the union of skill sets of all technicians who have the specialty  $s$ .

For a given service order that requires a particular specialty in  $S$ , the algorithm computes the distances from this order to all other service orders requiring specialties in  $S$ . The parameters and computing strategies used to create OD1 are applied here similarly. In this case, we eliminate the distances between pairs of service orders that cannot appear on the same route.

In case some stops are not covered, the required distances will be appended in the resource-assignment procedure. However, the number of distances to be computed is very small. Furthermore, the procedure to create an OD matrix can be run repeatedly so that whenever a new order comes in, the desired travel time (distance) can be appended properly.

### Assigning Resources

After the OD matrix is built, it is possible to solve the VRPTW for the technician-dispatch or home-delivery problem. In the following discussion, we call a service order or a customer needing delivery a *stop*, and a technician or a truck driver a *vehicle*. The resource-assignment algorithm (multiple insertion) assigns stops to vehicles. Sequence- and route-improvement algorithms improve the results of the resource assignment by revising the assignment decision for each stop. The multiple-insertion (MI) algorithm extends and adapts the existing (generalized assignment) algorithms [Ball et al. 1995; Solomon 1987] for

VRPTW and incorporates multiple objectives by using adjustable weights in the objective function while assigning stops to vehicles. The MI algorithm treats some of the constraints mentioned above, such as skill set, vehicle capacity (weight or volume), and precedence, as hard rules; that is, it cannot violate these constraints during the assignment procedure without producing infeasible routes. Other constraints, such as time windows and duration of a route, will be viewed as soft rules; that is, it can violate these constraints by accepting associated penalties. The MI algorithm tries to minimize these penalties and travel time (distance) when it assigns stops, even though it is essentially an initial construction procedure. The major components of the MI algorithm are building initial routes, assigning stops, and an optional post-insertion improvement step.

First, the MI algorithm builds an initial (dummy) route for each available vehicle  $r$ . The initial route contains only the starting, seed, and ending points. Then, we compute the travel time (distance)  $t_r$ , the amount of time-window violation  $v_r$ , and waiting (free) time  $w_r$  for each vehicle  $r$ , treating the overtime as the time-window violation of vehicle  $r$ . If a vehicle returns to the ending point before the shift end time, we consider the vehicle to have free time. The objective function of the current solution can be determined by

$$\alpha_1 \sum t_r + \alpha_2 \sum v_r + \alpha_3 \sum w_r,$$

where  $\alpha_1$ ,  $\alpha_2$ , and  $\alpha_3$  are the weights for travel time (distance), the amount of time-window violation, and waiting (free) time, respectively. Users can adjust these

weights based on their local or seasonal business objectives and geographic characteristics.

Second, the MI algorithm inserts any unassigned stop into a route between two consecutive vertices that represent starting, seed, or end points, or stops. Let  $U$  represent the set of all unassigned stops. For each route  $r$  and position  $k$  where a stop may be inserted, we define a corresponding insertion cost for an unassigned stop  $i$  as follows:

$$c_{irk} = \alpha_1 \Delta t_{irk} + \alpha_2 \Delta v_{irk} + \alpha_3 \Delta w_{irk},$$

where  $\Delta t_{irk}$  is the change in travel time,  $\Delta v_{irk}$  is the amount of time-window violation, and  $\Delta w_{irk}$  is the waiting time of route  $r$ , if stop  $i$  is inserted at position  $k$  of this route. If route  $r$  does not have the specialty stop  $i$  requires or the capacity of route  $r$  will be exceeded if stop  $i$  is inserted at position  $k$ , then the insertion cost is infinite. We select a stop  $j$  and insert it into route  $t$  at position  $q$  such that

$$c_{jtq} = \min \{c_{irk} \mid \text{for all } i, r, \text{ and } k\}.$$

As soon as a route receives a real stop, the seed point is eliminated from the route. The route attributes, such as travel time (distance), time-window violation, and waiting (free) time, are modified correspondingly. After unassigned stop  $j$  is inserted into some route, let  $U = U \setminus \{j\}$ . If  $U = \Phi$ , the MI algorithm either goes on to the postinsertion improvement procedure or it stops. If  $U \neq \Phi$ , it updates the insertion costs for all elements in  $U$  by considering all possible routes and positions in those routes and repeats this insertion step.

In the optional postinsertion improve-

ment step, we modify routes (clusters) created by the procedure described above that are not satisfactory because of an unacceptable imbalance in workload, such as number of stops or overtime. In this case, we will transfer stops from route to route to create balanced workloads for all routes. In making such transfers, we always try to ensure that the increment of the weighted objective function is minimal.

Once we solve the technician-dispatch problem, we address other issues. A technician may have primary and secondary specialties. The insertion cost of an unassigned stop will be adjusted based on the technician's primary specialty or secondary specialty. The adjustment favors a technician getting stops that require his or her primary specialty. Technicians are classed as full-time, part-time, or flexible-time. The MI algorithm adjusts the insertion cost during the insertion procedure so that the routes corresponding to full-time technicians get as many stops as possible (considering the duration of the route and the technician's schedule). The MI algorithm generates a set of routes, assigns all stops to these routes, and determines the sequence (position) for each stop within a route.

### Improving Sequences and Routes

The sequence-and-route-improvement procedure improves the initial routes generated by the MI procedure. Let  $R = (V, A)$  represent a route, where  $V$  is the set of vertices representing the starting point, ending point, and stops, and  $A$  is the set of arcs connecting pairs of vertices in the route. Each arc  $(i, j) \in A$  has an associated cost  $t_{ij}$ , indicating the travel time (distance)

from  $i$  to  $j$ . The procedure considers all the routes built in the assignment process as a whole system, and based on that, the algorithms try to find the most suitable route for each stop and the best position for each stop within individual routes according to the objective function. The procedure consists of two major heuristic procedures—intraroute and interroute improvements.

### Intraroute Improvement

The purpose of intraroute improvement is to find the best position for each stop within the single route. Theoretically, the problem we solve is the traveling salesman problem with time windows and other side constraints (for instance, capacity and precedence). Since the underlying optimization is an NP-hard problem, we use heuristics to solve the problem. We define a *move* as an operation that transforms the current solution (route) to a new one. To solve this difficult problem efficiently, we developed moves similar to the procedure Or [1976] describes. We consider two types of moves to solve the problem—forward and backward insertions. They were proved to be very effective in solving the traveling salesman problem with time windows and other side constraints [Cao and Rinderle 1992; Carlton and Barnes 1996]. We use forward (backward) insertions to try to improve a route by relocating a stop within it. If we index the vertices representing stops in a route from 1 to  $n$  in sequence, where  $n$  is the total number of vertices in the route, removing a stop from its current position in the sequence and placing it a position later in the route is a forward-insertion move, while removing a stop from its cur-

rent position and inserting it earlier in the route is a backward-insertion move. If the forward-insertion move removes a stop from its current position  $i$  and places it at  $j$  ( $j > i$ ) position in the sequence, then we can determine the resulting change of the travel time for the route as follows:

$$\Delta t_{ij} = t_{i-1,i+1} + t_{j,i} + t_{ij} - t_{i-1,i} - t_{i,i+1} - t_{jj+1}.$$

We also compute the changes in time-window violation  $\Delta v_{ij}$  and waiting (free) time  $\Delta w_{ij}$ , since the arrival times for all stops after  $i - 1$  in the route are changed because of the move. We define

$$\Delta c_{ij} = \alpha_1 \Delta t_{ij} + \alpha_2 \Delta v_{ij} + \alpha_3 \Delta w_{ij}$$

to be the cost associated with this specific forward-insertion move, and we compute this cost for any possible forward-insertion move. Similarly, we can determine the cost of a specific backward-insertion move. At each intraroute improvement step, we choose the least costly forward or backward insertion. The insertion cost may be infinite for a forward- or backward-insertion move if it violates the precedence constraints or exceeds the capacity of a vehicle that runs this route.

We extend the moves discussed here to deal with inserting more than one stop. Quite frequently, the travel times (distances) between some stops within a route are zero, because the corresponding service orders or deliveries are located in the same building (they are referred to as combined stops). For practical reasons, the same place should not be visited more than once as long as all the stops at that location have compatible time windows. For combined stops, the forward (back-

ward) insertion considers moving these combined stops as a group using a strategy similar to that for moving a single vertex.

### **Introute Improvement**

In solving a VRPTW, we make two types of decisions, the assignment decision (which route serves which stops) and the routing decision (what is the sequence of the stops within a route). We use the intraroute-improvement procedure to improve the routing decision. Although we may have several thousand stops to assign, a route usually contains no more than 70 stops. It is impossible to obtain a significant improvement by revising only the routing decision. Therefore the intraroute-improvement procedure alone cannot solve the technician-dispatching or the home-delivery problem effectively. The interrout-improvement procedure is a heuristic that utilizes the multiple-route structure to explore additional opportunities, revising the assignment decision to obtain better solutions. It uses two types of moves to relocate stops between two routes: transferring moves and exchanging moves.

In a transferring move, the heuristic removes a stop from one route (the origin route) and inserts it in another (the destination route) at a determined insertion position. For each potential move, the heuristic computes the transferring cost based on the transferred stop, the destination route, and the insertion position. The transferring cost is infinite if (1) the destination route does not have the specialty the stop requires (for technician-dispatching problems); (2) the precedence constraint in either route is violated; (3) the capacity of

either route is exceeded; or (4) the stop was preassigned to the origin route. A transferring move will be carried out if it has the least transferring cost.

In an exchanging move, two stops from different routes are relocated simultaneously into the other route. The heuristic decides the insertion positions for these stops in their respective destination routes. We compute the exchanging cost for each potential move based on the exchanged stops, the routes involved, and the insertion positions in the routes. The exchanging move is prohibited if (1) a route does not have the specialty the stop to be inserted requires; (2) the capacity of either route is exceeded owing to the move; (3) the precedence constraint will be violated; or (4) one or both stops are preassigned. The exchanging move that has the least cost is performed.

Similar to the intraroute-improvement procedure, the transferring and exchanging moves can be extended to a group of stops (combined stops) as well. Based on our computational experiments, this kind of extension sometimes can improve the overall solution significantly.

In solving the technician-dispatch problem, we modify the transfer and exchange costs so that generally the stops served by full-time technicians will not be switched to routes run by part-time technicians. If the "exception" geographic region restriction is imposed, a stop within this "exception" region cannot be switched to a route outside of the region. After interrout-improvement ends, we apply the intraroute-improvement procedure to each route.

### **Tabu Search Application**

The moves in the intraroute- and

interroute-improvement procedures are all local heuristic searches, and these searches are easily trapped in the local optima. We use tabu search techniques to obtain better solutions. Tabu search [Glover 1986] is a metaheuristic that guides local heuristic search procedures (for example, the moves in intraroute and interroute improvements) to explore the solution space beyond local optimality. Tabu search has been widely applied to a variety of combinatorial optimization problems [Carlton and Barnes 1996; Gendreau et al. 1992; Glover 1995]. Given a function  $f(x)$  to be optimized over a finite set  $X$ , tabu search proceeds iteratively from one solution to another until the termination criterion is met. Each  $x \in X$  has an associated neighborhood  $N(x) \subset X$ , and each solution  $x' \in N(x)$  can be reached by an operation called a move. To avoid being trapped at a poor local optimum, tabu search accepts a new solution  $x'$  even if its objective function value is worse than the best one. Since cost improvement is not enforced, tabu search runs the danger of cycling—repeating the same sequence of solutions indefinitely. To resolve this problem, one can use one or more tabu search lists to classify certain attributes (for example, a variable whose value changed or a vertex that is added or dropped) of solutions recently visited as *tabu active*. A solution  $x' \in N(x)$  is classified as tabu to be regenerated if specified attributes of  $x'$  appear in the tabu lists. However, the tabu status of a solution can be overruled if certain *aspiration levels* are met. The result is the solution with the best objective-function value found in the entire solution process.

Tabu search ideas are used in the

sequence-and-route-improvement procedure. The attributes of moves that we use to define tabu status consist of stops on routes. Among these attributes, we identify several as critical attributes of the move. If the critical attributes are tabu active, then the corresponding move is tabu, that is, the solution created by this move is excluded from consideration unless it meets a certain aspiration level. The tabu-active classification is defined based directly on recency. Therefore a move is classified as tabu if it reiterates critical attributes that were changed by recent moves.

---

### Home delivery has become one of Sears' core competencies.

---

For intraroute improvement, we define the position on a route where the stop to be relocated is to be inserted as a *destination position*. Then we define the stop to be relocated and the stop at the destination position as the critical attributes of the forward- (backward-) insertion moves. When a stop is relocated, it will be put into tabu list 1 with size  $s_1$  and the stop at the destination position will be put into tabu list 2 with size  $s_2$ . A forward or backward move that contains these two critical attributes within the next  $s = \min\{s_1, s_2\}$  iterations is excluded from consideration unless it satisfies the aspiration level (defined later).

For the transferring move of the interroute-improvement procedure, we define the tabu rules as follows: A stop to be transferred from one route to another is defined as the critical attribute of the move. If a stop is transferred, it cannot be

the critical attribute of any transferring move during the next  $s$  (the size of the tabu list for the move) iterations. Tabu restriction for the exchanging move defines the two affected to be the critical attributes. Both stops will be put into tabu list 1 (size  $s_1$ ) and tabu list 2 (size  $s_2$ ), respectively. The move that uses these two stops as its critical attributes is prohibited within next  $s = \min\{s_1, s_2\}$  iterations (again, subject to aspiration level). The aspiration criterion we used is the simple aspiration by objective, that is, the tabu status of a move can be overruled if the value of the weighted objective function obtained by this move is better than the value of the best solution obtained previously. Based on our computational experiments on real problems, the solutions obtained by this tabu-search-based procedure are better than (15 percent improvement) those yielded by the heuristic without tabu-search techniques.

### System Roll Out

ESRI and Sears jointly built a project team in October 1993 after Sears asked ESRI to develop and implement the systems for SLS and SPS. This team was also responsible for training staff, explaining and analyzing the results obtained by the systems, reporting on project progress to the management of both companies, collecting criticism from the field, and revising the systems.

Because of the complexity and the scale of this project, the project team first created pilot systems for SLS and SPS. Using the pilot system, the project team found problems it had not predicted and realized some business logic it had not considered. Based on feedback from the pilot field, the

project team figured out how to improve the system. Because of their similarity, the project team developed the two systems in parallel. After about one year of development, the project team implemented the pilot system for SLS in 1995 and the pilot system for SPS in early 1996. The team chose several regional offices for implementation that would reflect different scenarios, for example, the service area covered had various geographic characteristics, including mountains, deserts, and beaches, and also offered all of Sears' services. This allowed the project team to test the system thoroughly and acquire wide experience, which greatly helped it in the nationwide rollout. The project team spent another half year improving the systems based on the pilot and finalizing the systems for the gradual nationwide rollout. By the end of 1996, both systems were working for the entire country.

As we rolled out the system, we encountered several impediments to smooth implementation. The EHDS/CARS system required Sears employees to make a transition from text-based terminals to mouse-driven GUIs. Many found this difficult. In addition, users had to learn the effects of various algorithm-input parameters on the quality of routes. SLS and SPS took different approaches to training. SLS rolled out EHDS incrementally with a small Sears project team visiting each of the then 46 delivery offices to provide intensive training. SPS, because of the large size of its operation, hired a third-party implementation firm to provide training.

During the initial rollout, various aspects of the algorithms were not complete

or parameters needed tuning. This resulted in routes that were infeasible or impractical, which reduced the dispatchers' and field personnel's confidence in the system. Also, because the system seeks a global optimal (to reduce costs across all routes on a given day), some individual routes may suffer. Usually the joint team had to analyze the results and explain them to dispatchers and field personnel. And the team had to modify the algorithms to correct problems. Confidence improved as the system provided more reasonable routes and people became familiar with it.

---

### ESRI's projects for Sears have driven the improvement of commercially available street-network data.

---

Technicians and drivers resented being tracked by the system. Mileage calculations compared with odometer readings exposed off-route activities or failure to run routes in the prescribed sequence. However, this obstacle could be overcome as the field manager gained more confidence in the system and the system matured. Because of the benefits the system brought, the field manager was able to convince the technicians and drivers to follow the routes the system suggested.

When we initially deployed the system, we set uniform algorithm parameters for all dispatch and delivery offices nationwide. We discovered that different offices would perform better with different parameter settings. Also, regional managers wanted to tune the system to accomplish specific regional objectives, for example, to

honor time windows rather than minimize travel time or mileage. We devised individual parameter settings for each region. Tuning the system based on the geography and the business objectives further improved the results.

#### **Benefits and Impact for Sears**

Since the systems for SLS and SPS were deployed, SLS and SPS have improved their overall service. Using the system, SLS and SPS have improved their honoring of time windows. SLS and SPS have reduced their operational costs because of reductions in travel time, mileage, and routing time. SLS is able to utilize its fleets more efficiently, because the mileage reduction increases the number of stops a vehicle can visit. Table 2 summarizes the improvements in some key cost factors and business criteria for SLS and SPS.

By decreasing miles per stop, SLS achieves more stops per truck, maximizing the use of delivery resources. Because of the huge number of deliveries, even a fractional saving on each stop results in a significant annual savings. Sears increased its delivery orders 9.2 percent with virtually the same fleet size. The increased number of stops per vehicle has also allowed Sears to negotiate a rate reduction with third-party carriers.

EHDS has eliminated the need for dispatchers to have local knowledge, and it has sped up the routing process. Dispatchers no longer need map hard copies to build routes. Efficient and automated routing has allowed Sears to consolidate its routing offices, going from 46 to 22 with further reductions planned, which reduces facility, equipment, and personnel costs.

Home delivery has become one of Sears'

	SLS		SPS	
	Before	After	Before	After
Geocoding accuracy	55%	95%	55%	95%
Arrival time window	4 hours	2 hours	2 hours	2 hours
On-time performance	78%	95%	84%	95%
Time spent routing	5 hours	20 minutes	8 hours	1–2 hours
Miles per stop	1.6	1.2		
Stops per vehicle	16	20		
Dispatch facilities	46	22	92	6
Completed calls	N/A	N/A	---	+ 3%
Overtime				– 15%
Drive time				– 6%

**Table 2: Sears achieved many benefits by using EHDS and CARS: It improved its accuracy in finding customer locations, it improved its fulfillment of time-window commitments, and it reduced its operational costs by reducing mileage, utilizing vehicles more efficiently, and consolidating facilities.**

core competencies distinguishing it from its competitors. Sears can now promise two-hour time windows, which makes Sears a more compelling place to shop. Since Sears can notify customers of their delivery windows early, drivers encounter fewer not-at-home customers. Because Sears can predict delivery times accurately, customer satisfaction with the delivery process has increased from 84.7 percent to 87.2 percent. Arrival within promised time windows retains customers, and excellent service attracts more customers. According to the statistical data, a typical “happy” Sears customer spends \$20,000 over his or her lifetime at Sears stores. Customer satisfaction will bring revenue.

With EHDS, Sears can meet its objectives and those of its customers, rather than be subject to the objectives of the third-party carrier. The system has eliminated off-route driver activities and the fees Sears paid to third-party dispatchers.

Altogether EHDS results in annual sav-

ings of \$30 million.

CARS improved technician productivity: (1) It reduced driving time by six percent; (2) it increased Sears’ ability to replace cancelled service orders; (3) it increased Sears’ ability to balance service orders across technicians throughout the day; (4) it increased the number of service orders each technician completed per day by three percent; and (5) it reduced overtime by 15 percent. Because of these improvements, Sears has achieved annual savings of \$9 million.

Like EHDS, CARS has eliminated the need for dispatchers to have local knowledge and to interface directly with technicians. CARS helps Sears to use its resources more efficiently and to reduce its resource requirements. The system has helped Sears to establish standard processes, to generate higher quality routes that have lower costs and higher time-window commitment, and to balance the workloads among the technicians. The system also helps to dramatically reduce the



time taken by the routing process. Using CARS, call center agents can respond to customer inquiries quickly with such information as technician-arrival time. CARS has made it possible for Sears to consolidate its dispatching facilities, which saves \$9 million.

CARS generates routes for service technicians more efficiently than the dispatchers did, reducing the transportation costs associated with home services. Based on the statistics, CARS reduced variable truck costs by 10 percent, yielding a \$3 million annual saving.

Furthermore, CARS improves Sears' ability to respond the same day to requests for service, and it increases technicians' on-time arrivals. Since CARS was implemented, customers who report being "very satisfied" have increased from 74 percent to 77 percent, and recommendation scores have been over 80 percent. Because of the improvements in service quality, customer satisfaction, and costs, revenues for paid service calls have increased by 16 percent—another significant impact.

#### **Benefits and Impact for ESRI**

ESRI has used the knowledge it gained and developed in the Sears projects to enhance and extend its product line. Specifically, we have developed new algorithms in our Arc/Info and ArcView products, developed new products, and improved data structures and memory management to solve very large network optimization problems.

We enhanced the Arc/Info Network extension to include precedence ordering in the Tour command (TSP heuristic). We enhanced geocoding functionality to im-

prove various address-matching criteria and tools to assist in reject processing (the methodology to deal with some addresses that cannot be matched by the normal address-match process). We enhanced our dynamic memory management tools so that we can manage very large networks (millions of arcs) in virtual memory.

We developed the ArcView Network Analyst product to leverage these ideas in a mass market commercial product. Many ESRI developers use ArcView to create value-added products to resell, for example, MileMaster for compiling mileage data for tax accounting and the product Miner and Miner developed for analyzing the outage of electric networks.

We developed RouteServer, a client-server routing system that combines the assignment and sequencing algorithms described here with the transaction processing (TP) monitor technology originally developed for Internet applications. This is a very promising technology for enterprise applications because it allows multiple users to share a common street database and computing resources to drastically reduce the incremental cost of client workstations for routing operations.

We developed ArcLogistics, a low-cost PC-based routing-and-scheduling application that brings high-end functionality to small organizations who were previously unable to afford this technology. We created this product based on the experience we gained in the two Sears projects. It helps small organizations to solve both the technician-dispatch problem and the delivery-routing problem.

Arc/Info Version 8.0 is a complete redesign of ESRI's GIS product, which will in-

clude network modeling tools, many of which result from our work with Sears.

Because of ESRI's success with Sears, we have expanded our transportation and logistics consulting practice and customer base. ESRI established its transportation and logistics services in January 1996 to satisfy a demand for services similar to those provided to Sears. We have increased our recruitment and hiring of OR professionals, and OR/MS is now a recognized field of expertise at ESRI.

ESRI is now well known for its capabilities in routing and scheduling, and it integrates routing-and-scheduling technology in many projects our professional services division conducts, for example, applications in Internet mapping and routing, 911 and emergency response, forestry management, information kiosks, dial-a-ride services, warehouse-picking optimization, yield management, and insurance-claims-adjuster dispatch, and in the layout design of telecommunications fiber cable. We expect to continue developing applications with this technology in these industries and many others.

Though not a specific achievement, ESRI's projects for Sears have driven the improvement of commercially available street-network data. Our data partner in these projects, Geographic Data Technologies (GDT), now offers, as a result of these projects, a continuous street-network database for the entire US, suitable for routing and scheduling applications. This greatly reduces the cost for organizations desiring to implement this technology.

We feel that our primary accomplishment in these projects is the integration of GIS data structures and processing tech-

niques with leading-edge OR theories and algorithms. We accomplished this integration at the basic levels of our products, making this much more than a loose coupling of techniques. The true integration of geoprocessing software with the algorithms we employed allows us to operate on much larger street networks with faster processing times than has previously been possible.

ESRI's association with Sears has resulted in our offering new products and services that benefit not only home-service and home-delivery industries, but also many unrelated industries. The algorithms and their technical implementation have proven to be generic enough that we can successfully apply them with little or no modification for a wide variety of customers.

## References

- Ball, M. O.; Magnati, T. L.; Monma, C. L.; and Nemhauser, G. L., eds. 1995, *Network Routing*, Institute for Operations Research and Management Sciences, Elsevier Science B. V., Amsterdam, The Netherlands.
- Bodin, L. D.; Golden, B. L.; Assad, A.; and Ball, M. 1983, "Routing and scheduling of vehicles and crews: The state of the art," *Computers and Operations Research*, Vol. 10, No. 1, pp. 69-211.
- Cao, B. and Rinderle, K. 1992, "Traveling salesman problem with path dependent costs—Generalization of one-dimensional search with traveling cost," Technical report S-9201, Institute of System and Operations Research, University of the Federal Armed Forces-Munich, Neubiberg, Germany.
- Carlton, W. B. and Barnes, J. W. 1996, "Solving the traveling salesman problem with time windows using tabu search," *IEEE Transactions*, Vol. 28, pp. 617-629.
- Desrochers, M.; Desrosiers, J.; and Solomon, M. 1992, "A new optimization algorithm for the vehicle routing problem with time windows," *Operations Research*, Vol. 40, No. 2,

- pp. 342–354.
- Glover, F. 1986, “Future paths for integer programming and links to artificial intelligence,” *Computers and Operations Research*, Vol. 13, No. 6, pp. 533–549.
- Glover, F. 1992, “Tabu search fundamentals and uses,” School of Business, University of Colorado at Boulder, Colorado.
- Gendreau, M.; Hertz, A.; and Laporte, G. 1992, “A tabu search heuristic for the vehicle routing problem,” Technical report CRT-777, Center for Research on Transportation, Université de Montréal, Montréal, Canada.
- Or, I. 1976, “Traveling salesman-type combinatorial problems and their relation to the logistics of blood banking,” PhD dissertation, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, Illinois.
- Solomon, M. 1987, “Algorithms for the vehicle routing and scheduling problem with time window constraints,” *Operations Research*, Vol. 35, No. 2, pp. 254–265.