Automated Paratransit Routing and Scheduling Using a Highway Network Model

DANIEL J. BOWER

The quality of vehicle routing and scheduling has a great impact on the operating cost and service quality of any paratransit service. The difficulties associated with many-to-many trip requests, vehicle availability and capacity constraints, traffic patterns, and geographical obstacles are such that the process, when done manually, is error prone, labor intensive, and difficult to optimize. The development of computer software that creates paratransit schedules can, in addition to addressing these problems, significantly reduce the cost of generating schedules and improve the overall efficiency of providing service. COMSIS Corporation has developed the COMSIS Routing and Scheduling System (CRSS) to address these issues. It uses a highway network to model point-to-point travel times that recognizes geographical obstacles. CRSS takes a set of trip requests entered by means of a paratransit management information system and generates vehicle manifests. The program also tracks the estimated vehicle locations and remaining vehicle capacity throughout the day. The software is processed in batch after all requests have been received. CRSS has been implemented at a number of paratransit agencies, and comparisons have been made. Manually generated schedules have been compared with CRSS schedules in the areas of service quality and operating cost and the results are favorable. Areas for further development have been identified.

Paratransit scheduling can be considered an expansion of the classic traveling salesman problem, which involves the calculation of an optimum itinerary for visiting a number of predetermined nodes on a network. The use of a network to model a roadway system provides a realistic representation of geographical obstacles and travel costs, thus enabling the development of a feasible itinerary. In its simplest forms, this problem has a closed solution using linear programming techniques.

In the paratransit scenario, the following constraints beyond the traveling salesman problem exist:

1. There are two types of stops—pickup and drop-off. All are paired and each pair must be serviced in sequence. A pickup–drop-off pair can be considered a trip request.
2. Trip requests may have unique pickup and drop-off locations. Thus the problem presents a many-to-many instead of a one-to-many scenario.
3. Many vehicles may be available to service trip requests. Each vehicle has a capacity limit, and the capacities of vehicles may vary. Also, some vehicles may be physically unable (too wide, too high) to access certain locations. Furthermore, some vehicles may be used preferentially, such as agency-operated vehicles versus supplemental contract carrier vehicles.
4. Types of passengers may vary. At a minimum, one must distinguish between wheelchair and ambulatory passengers for the use of vehicle capacity and the variations in load and unload time.
5. Constraints must be placed upon the amount of time a passenger is kept on a vehicle (i.e., a limit on ride time).
6. Limits must be placed upon the amount of variation between a requested pickup or drop-off and a scheduled pickup or drop-off to ensure timely service.
7. Some trips may have a greater priority than others. This may involve the division between promised, recurring service (standing orders) and ad hoc (demand-response) trip requests. In the case where not all vehicles can transport passengers in wheelchairs, a priority may be placed upon the wheelchair trips.

These constraints upon the traveling salesman problem, in the author's opinion, render the problem intractable. Even if a closed solution was possible, the computing resources necessary to produce solutions in a timely fashion would be prohibitively expensive. The only practical methods of paratransit scheduling are to perform the task manually, or to develop a computerized heuristic to solve the problem.

MANUAL PARATRANSIT SCHEDULING TECHNIQUES

Manual techniques for paratransit scheduling vary in application from agency to agency; however, the substance of various methods is largely the same. Simply stated, the process for a service with no preference by trip type involves

1. Sorting trips chronologically. This is a simple sort throughout the trip requests for a service day.
2. Grouping trips geographically. Trip grouping involves collecting parallel trips at similar times of the day. This is done with consideration of geographical obstacles and traffic conditions.
3. Assembling groups into manifests. Once trips are collected into groups, the groups are sequenced chronologically into manifests.
4. Resolving exception trips. Trips that do not fit into groups, as well as ones that are handled outside the rules, are inserted into the schedule as best the scheduler can.

COMSIS Corporation, 2275 Swallow Hill Road, Building 500, Suite 3, Pittsburgh, Pa. 15220.

A common variation involves assigning priority to trips by trip type. In some cases, standing order (prearranged) trip requests take precedence over ad hoc requests. If a vehicle fleet is only partially equipped to handle wheelchairs, priority would commonly be given to the wheelchair trips. To handle priorities, one would make multiple passes through the procedure, one for each of the priorities.

Throughout the second and third steps, the constraints should not be violated. Specifically, this involves constructing groups that do not exceed vehicle capacities. Manifests must be arranged to provide service without keeping any passengers on the vehicles for inordinate amounts of time and without changing requested pickup or drop-off times too severely. Above all, the resulting schedule must be feasible. The schedule must be written considering expected traffic conditions and the existing connectivity of highways so it can be driven on time.

ADVANTAGES AND DISADVANTAGES OF MANUAL PARATRANSIT SCHEDULING

There are many advantages and drawbacks to manual scheduling. Foremost among the advantages is versatility. A good scheduler can be adept at handling exceptions. In cases where service demand starts to exceed capacity, the scheduler can make judgments about where to bend service policy rules. Even under an automated system, the resolution of exceptions is still best left as a manual process.

Another advantage is the low capital cost. The tools needed for manual paratransit scheduling can be as simple as a large table and a set of trip requests, each on a standard form.

It should be noted that, in some cases, the day-to-day effort involved in a manual procedure can be small, even for a relatively large number of daily trip requests. This typically happens when most trip requests are standing orders. In such a case, the schedule for the standing trips is written once for each day of the week. It is refined periodically as standing requests change, but it changes little from day to day. For each service day, relatively few demand-response requests need to be added to the largely standing-order schedule.

The disadvantages of manual scheduling include the amount of labor necessary to schedule a large number of trips, when most are not standing requests. Here, the schedule must be, to a great extent, written from scratch every day. A typical case is that of Capital Metropolitan Transit Authority (Austin, Texas), where 18 to 22 person hours are required to prepare a weekday schedule of 1,600 trips. Sixty percent or more are not standing orders.

In agencies serving a large number of trips every day, it is difficult to maintain a consistent level of service quality. Often, as deadlines approach, the scheduler hurries to make the last vehicle assignments, sometimes with less regard to efficiency and passenger service quality than were the case with the first trips scheduled. Furthermore, like any manual process, the element of human error exists, resulting in conflicts and infeasibilities in manifests.

For paratransit services where the amount of work presses the limits of the human ability, a computer program to create paratransit schedules can be very useful. The output of such a program should meet the following goals:

1. Schedules must be feasible. There must be sufficient time between scheduled stops to allow for on-time service. Physical and operational constraints (e.g., vehicle capacity, ride-time limits, etc.) must not be violated.
2. Schedules must be created quickly. Ideally, the scheduling process should take considerably less time than with manual processing, to allow for review and revision as necessary.
3. Schedules must conform to service quality criteria. Like trips must be dealt with equally and within the limits of a set operating policy.
4. Schedules must be cost-effective, while still providing service quality within set limits.

SERVICE QUALITY MEASURES

A set of criteria is necessary to judge schedules. To measure the passenger's perception of service quality, ride time can be considered. Here, the passenger could be said to enjoy adequate quality service if the scheduled ride time was no greater than some factor of the direct ride time. The closer the scheduled ride time is to the direct ride time, the better the service quality.

The difference between the requested pickup and the scheduled pickup should also be considered. Here, the closer the schedule adheres to the passengers requested time, the better the quality of service. (Note that, if the passenger is concerned more with the drop-off time than the pickup, such as in the case of on-time arrival at the passenger's workplace, the difference between requested and scheduled drop-off times would be the criteria.)

SERVICE PRODUCTIVITY MEASURES

Paratransit productivity can be measured by a number of criteria. Two that can be useful are passengers served per vehicle hour, and total vehicle hours. The first measures the overall rate of service provision; the higher the rate, the more productive the service. Total vehicle hours have a direct effect on the overall cost of operation. Here, if the annual budget allows for a set average cost per day, this almost directly translates to a limit on total vehicle hours.

Note that some measures commonly used for fixed-route services can be misleading when applied to paratransit. For example, the ratio of live to dead vehicle time is of little value when it disregards the overall number of vehicle hours. It is entirely possible to arrange vehicle tours so there is little or no idle or deadhead time by simply retaining passengers longer than necessary, by making circuitous routings, or by both. In these cases, service quality can suffer, and the overall number of vehicle hours (and therefore overall cost) is higher.

Depending on the ratio of rate of pay to drivers to the cost of vehicle operation (fuel and maintenance), the issue of total mileage can be significant. Where driver pay is a relatively high percentage of total vehicle operating cost, it may not be necessary to consider reduced vehicle mileage separately from reduced vehicle hours. If, however, vehicle operating costs are a large portion of overall cost, the overall vehicle mileage can be a significant concern.

TRADEOFFS BETWEEN SERVICE QUALITY AND QUANTITY

For any given level of service supply (determined by driver and vehicle availability, and ultimately, an operating budget) and a given level of scheduling ability, there is an intrinsic tradeoff between quality and quantity of service. To improve service quality to the passenger (by reducing ride time or improving service timeliness), productivity (passengers per vehicle hour) must be reduced. In essence, to improve service to some passengers, one would be unable to take some others. Conversely, to improve productivity under these constraints, service quality to more passengers is reduced.

A computerized scheduler might be able to improve productivity without degrading quality, if the computer program is more capable than the person performing the task. Once the level of scheduling ability is established, the ability to control the tradeoff between quality and quantity is needed.

COMSIS ROUTING AND SCHEDULING SYSTEM

COMSIS Corporation has developed the COMSIS Routing and Scheduling System (CRSS), a computer program designed on the basis of the rules and goals outlined previously, to address the paratransit scheduling problem. This program includes a scheduling function, which is surrounded by functions to maintain the data bases needed to support the scheduler. It is not a full-functioned paratransit management information system (MIS); rather, it is a tool intended for use with an MIS.

The client list, standing order, and daily trip entry functions are managed by the MIS. Once all trips for a service date have been input to the MIS, the trip information for that date is transferred to CRSS. CRSS performs the scheduling functions, then reports the schedule back to the MIS. The MIS is then used to print driver manifests. CRSS can be integrated with any MIS that can export trip information and import schedule information by means of formatted ASCII files.

In addition to the trip request information imported from the MIS, CRSS maintains a number of data bases to accommodate variations in operating policies and conditions among paratransit agencies. One file describes the loading characteristics and restrictions on each vehicle type. A run file identifies, on a daily basis, a run's availability, preference type, starting location, and assigned vehicle type. A third data base describes a highway network of the service area, recording travel distances and speeds. From this network file, minimum path travel times are automatically derived, as well as a data base of geographical proximity. (See Figures 1–4 for data file layouts.) The use of this highway network is described at length later in this paper.

CRSS addresses the goals for the scheduling process as follows:

1. *Processing speed.* CRSS is written entirely in C, which is perhaps the fastest programming language that can be ported between computer systems. On a 20 MHz 80386-based PC

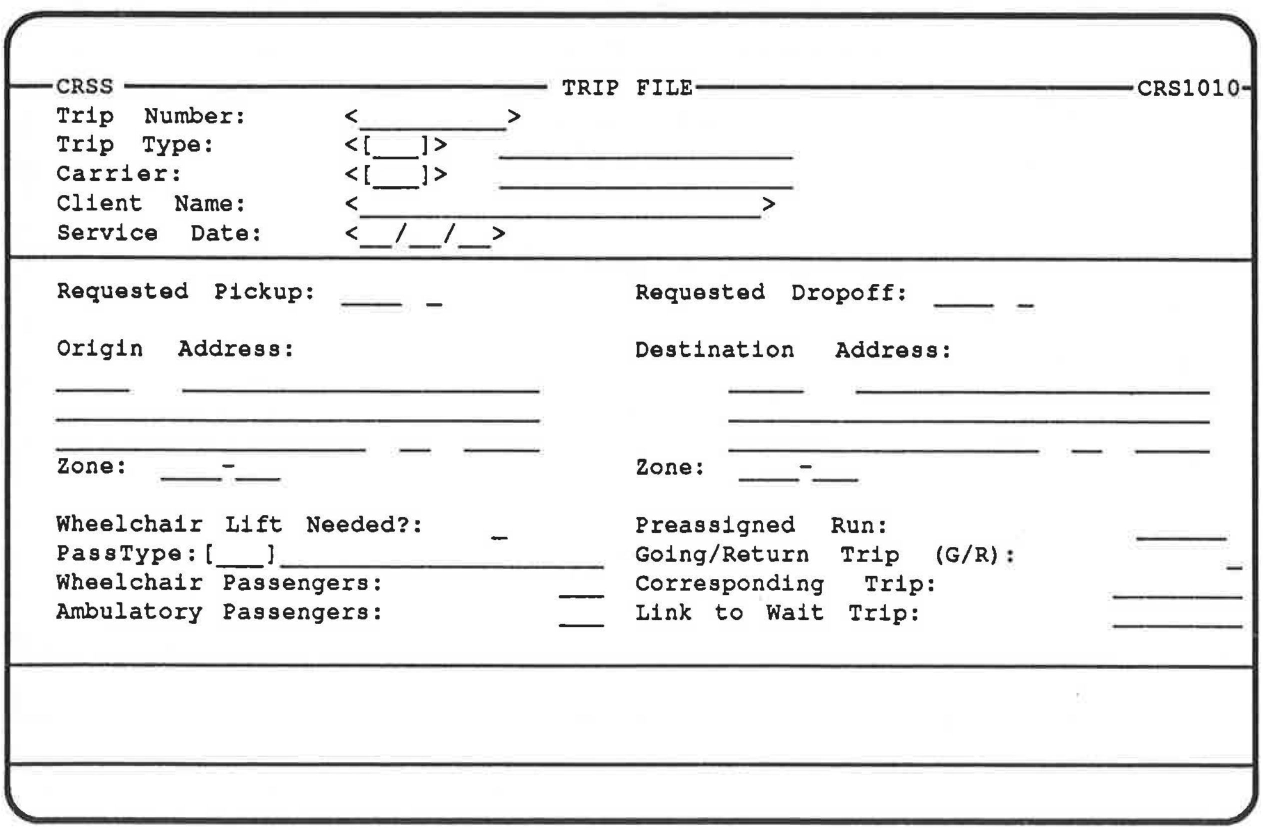

FIGURE 1 CRSS trip file screen displays select information about trip requests.

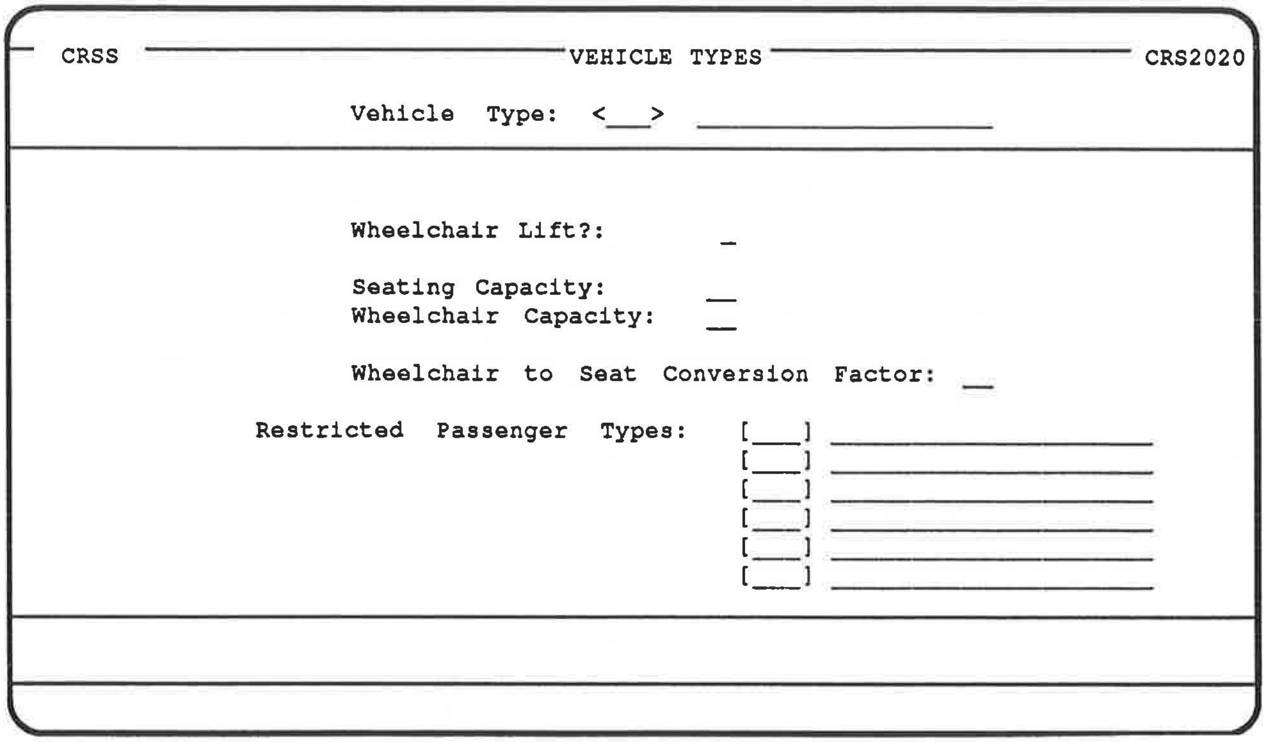

FIGURE 2 CRSS vehicle type file tracks the capacity and passenger type limitations of each type of vehicle.

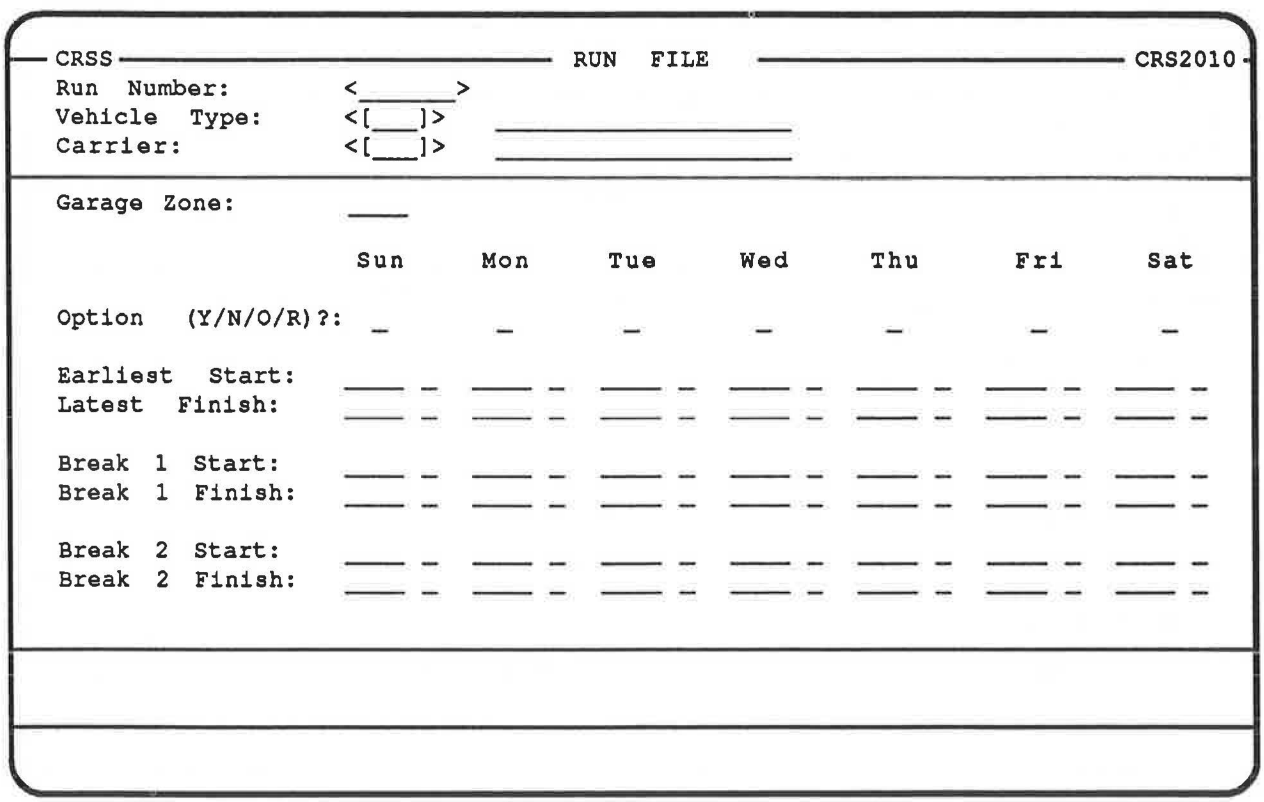

FIGURE 3 CRSS run file describes vehicle availability by time of day and starting location and identifies a vehicle type for each run.

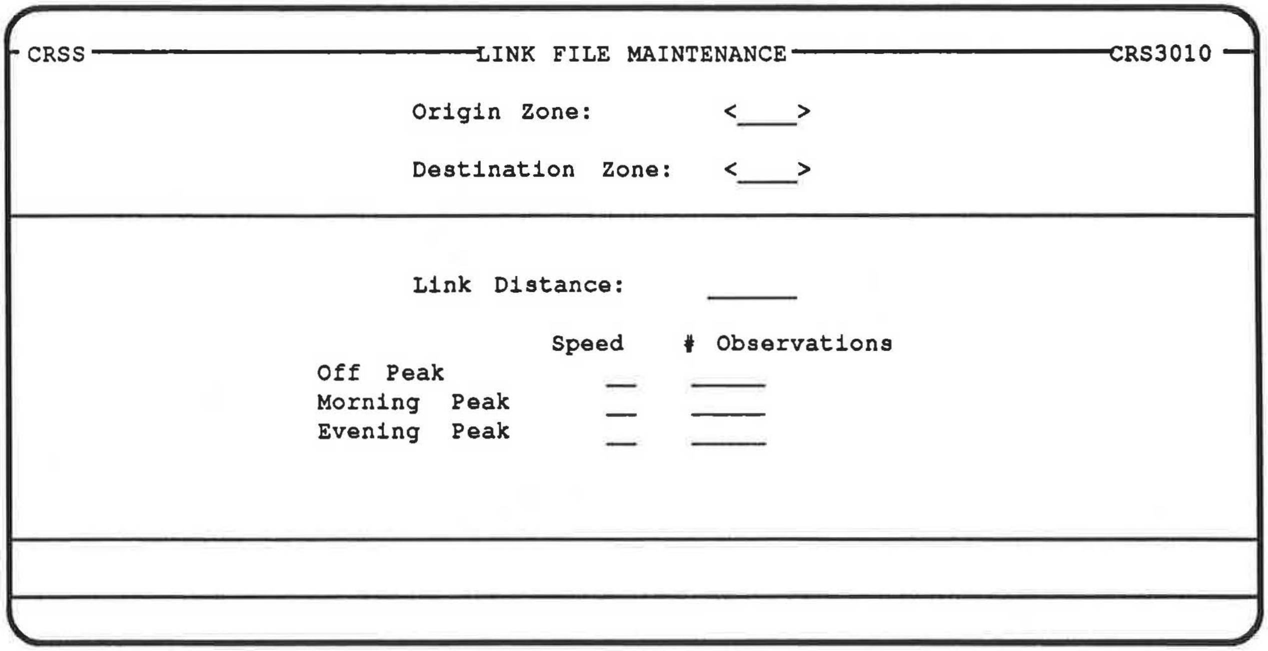

FIGURE 4 CRSS highway network link file computes minimum travel times and distances as part of scheduling process.

running Xenix, schedules of 300 trips are calculated in 15 min or less. Running time does not increase linearly; a schedule of 1,300 trips can take as long 90 to 100 min. (Performance varies with the hardware configuration and operating system. Running times are typically up to 50 percent greater with identical hardware under MS DOS.)

2. *Consistency*. CRSS allows the user to identify any number of trip types (for example, combinations of standing, demand response, ambulatory, and wheelchair are common). Service quality limits can be set independently for each trip type. Any number of passenger types can be defined, and load and unload times are independently set for each. Once these constraints are set, CRSS does not violate them. Thus, a consistent level of service is maintained, regardless of the person using the software, the number of trip requests, vehicle availability, and so on.

3. *Control of quality-quantity tradeoff*. Beyond the setting of absolute of passenger service quality, CRSS provides a mechanism to control this tradeoff. This is done during the scheduling process by means of a disutility function. When inserting a group of trips on any given run, CRSS calculates the following factors:

a. Additional minutes for all passengers in the difference between requested and scheduled trip times.
b. Additional minutes of ride time for all passengers.
c. Change in total live vehicle time.
d. Change in total idle vehicle time, where idle time is when the vehicle is "on the road," but sitting still with no passengers aboard.
e. Change in total deadhead vehicle time, where deadhead is any time the vehicle is moving with no passengers on board. This is not limited to the time to and from the garage at the beginning and end of the day.

The user of CRSS may assign a different weight to each of these factors, signifying relative importance of minimizing each independently. By manipulating these factors, the user can bias the schedule program's decisions toward vehicle productivity or passenger convenience.

4. *Feasibility*. As used here, feasibility of a schedule refers to the degree to which a driver can make all pickups and dropoffs on time and in proper sequence. The means by which CRSS addresses this issue are through the use of a highway network model of the service area. This approach was chosen primarily to avoid errors in schedule sequences and travel times inherent in Cartesian ("crow-flies") distance estimations and average travel time methodologies.

Illustrated in Figures 5–7 is a simple example of two trip requests. The points labeled Pickup 1, Dropoff 1, Pickup 2,

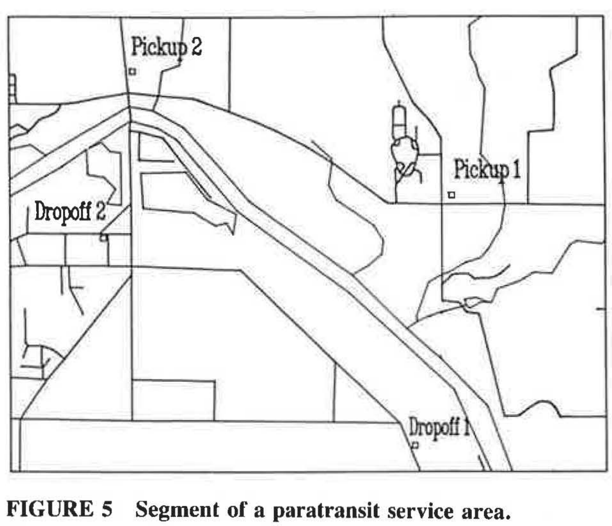

FIGURE 5 Segment of a paratransit service area.

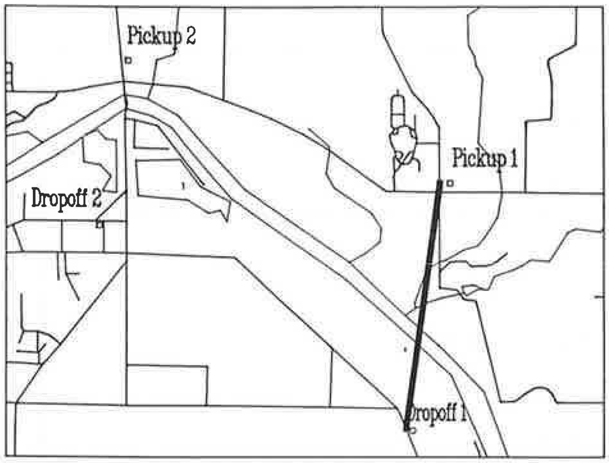

FIGURE 6 Intermediate step in scheduling process—line-of-sight travel estimations.

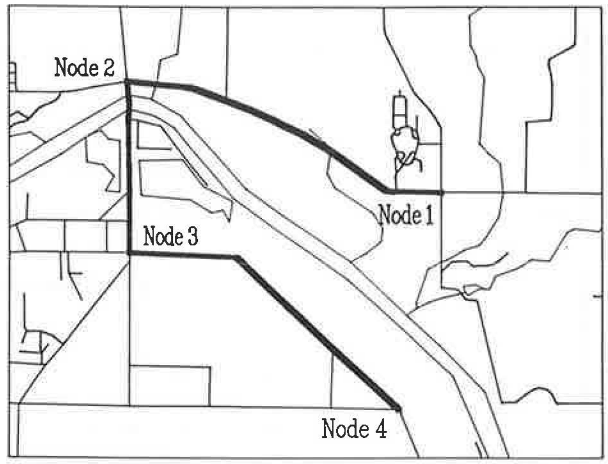

FIGURE 8 Highway network for sample paratransit service area segment.

and Dropoff 2 represent the origins and destinations of two trip requests at the same time. The trips are separated by a river with only one bridge. As each trip is entered, the origin and destination addresses are assigned coordinates. (This may be done many ways, ranging from simply looking up the coordinates on a map, to automatically matching the addresses against a list of all streets in the service area.) The travel times between each stop in a manifest are calculated from the distance, and a typical travel speed is determined.

Using a grid-based system, with constant travel speeds, a scheduling program might schedule the first trip (P1–D1) in correct sequence, but with insufficient travel time. Using the straight-line distance between the pickup and drop-off, there is no accounting for the lack of a bridge at that point on the river. Thus, the travel time calculated by this method is inadequate.

If the second trip is added to the vehicle's itinerary, the errors compound. Not only is the time between the first two stops too short, but the sequence would be P1–D1–P2–D2. Given the location of the bridge, the driver would have to pass the pickup and drop-off of the second trip on the way to the drop-off of the first, then backtrack to serve the second request.

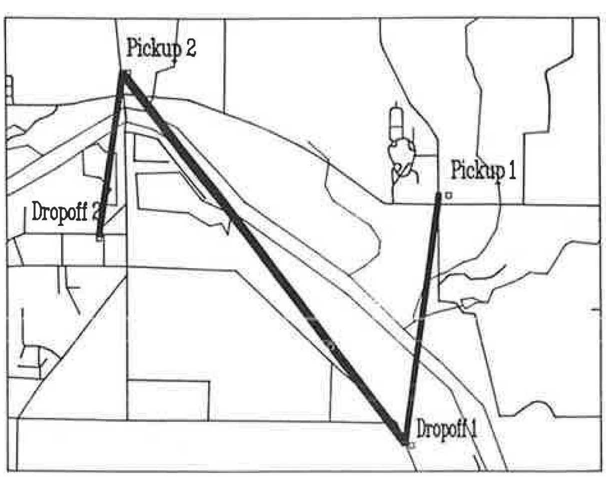

FIGURE 7 Final step of the scheduling process—line-of-sight estimations.

The way that one considers this problem, using CRSS, is to first describe the highway network available for use. Networks appropriate for operational use are much simpler than typical transportation planning networks. CRSS only requires the distance of each link, and the typical speeds (accounting for stops and turns) during the morning-peak, evening-peak, and off-peak travel. This information may be adapted from a transportation planning oriented network, depending on how closely the general travel demand follows that of paratransit travel. As illustrated, it is not necessary to code every last street and alley. However, it is necessary to define the network in sufficient detail so that no two commonly visited locations are associated with the same node.

In the example seen before, four nodes would be defined as shown in Figure 8. Six records would be entered into the link data base as presented in Table 1. Two records are associated with each link (one for each direction). This allows proper representation of directional congestion by time of day; typically, inbound congestion during the morning peak, outbound congestion during the evening. Note that each adjacent pair of nodes has two, one-way links between them.

As each trip request is accepted, the origin and destination are geocoded as before, but only one number is assigned: that of the nearest node in the network. P1 would be coded to Node 1, and so on. Assuming that the trip requests are for the morning peak, when the first trip is scheduled (Figure 9),

TABLE 1 CRSS HIGHWAY NETWORK LINK DATA

Orig Node	Dest Node	Distance	AM Speed	Off Peak Speed	PM Speed
1	2	.5	15	20	20
2	1	.5	20	20	15
2	3	.2	15	30	15
3	2	.2	15	30	15
3	4	.6	25	25	25
4	3	.6	25	25	25

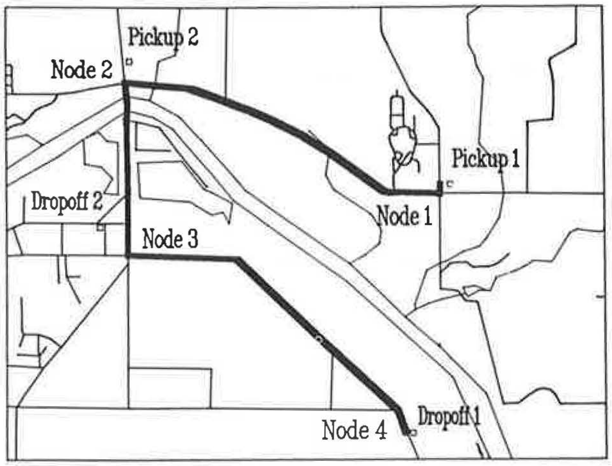

FIGURE 9 Vehicle manifest produced by CRSS after inserting the first trip.

CRSS uses the morning peak travel time along the minimum travel path defined by the network. In this case, the minimum path is by means of the bridge so the travel time used is 5 min. (CRSS works with integer time values; here, the fractional value of 4.5 min is rounded to 5 min.) Since CRSS will only route vehicles by the links defined, a line-of-sight time and distance are not considered for the first trip. An adequate amount of time is allotted.

CRSS considers all the feasible sequences when it inserts the second trip into the manifest (Figure 10). However, since travel times are derived from the highway network, CRSS uses a more representative time for each possibility. The vehicle time required by the solution from the crow flies' method (allowing one extra minute for loading and unloading) totals 12 min. A better solution would be P1–P2–D2–D1, which takes 9 min.

When the second trip was inserted, its proximity to the vehicle's path along the network was determined to be close enough to allow it to be scheduled on the vehicle as it served the first trip. The resulting manifest is correctly sequenced, and using travel times derived from the highway network, the manifest is feasible.

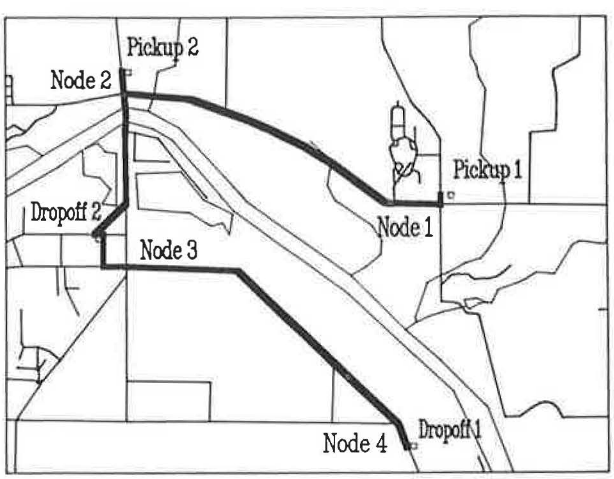

FIGURE 10 Final vehicle manifest produced by CRSS.

Note that, with any method, if the only way to accommodate the second trip is to cause the first passenger to be late for an appointment, it would be necessary to schedule the second trip on a different vehicle, or in a different sequence.

Maintenance of the Highway Network

During day-to-day use of CRSS, the data composing a highway network remain relatively static. Changes are necessary only when links on the network change. This occurs when bridges are closed, lanes are closed on major highways for maintenance, new roads are opened, and so on. Typically only a few records need to be changed. To change the speed on the bridge in the four node example, one would alter two records. If a new bridge was built between Nodes 1 and 4, two new records would be added. (Note that doing so could change nearly every point-to-point travel time in the service area.) Once this maintenance is done, one would run a function to automatically recompute the minimum paths, and scheduling may resume.

CRSS provides a separate method to address across-the-board travel time changes, such as those resulting from inclement weather. Instead of slowing down travel times on each link of the network, the user can enter a factor which is multiplied by each travel time as it is accessed. Thus, there is no need for long-range planning for weather adjustments. If the schedule needs to be slower on the next day, then one simply changes one entry and runs the schedule function.

Other Uses of Highway Network Information

In addition to computing travel times during the scheduling process, CRSS uses these travel times to group trips. CRSS defines a group as a set of trips with common or nearby origins and destinations, occurring at approximately the same time of day. To determine geographical proximity, CRSS uses the travel time between origins of trips. If the travel time is less than a configurable parameter, the trip origins are considered sufficiently close. The same test is applied to destinations. If both tests pass, and the requested time of day is within a configurable window, CRSS will attempt to schedule the trips onto the same run. (Considering vehicle capacity, passenger appointment times, and so on, all trips in a group might not end up on the same vehicle.)

CRSS uses the network travel times after a schedule is complete to assist in computing passenger hours and vehicle hours in a statistical summary. (See Figure 11 for an example.) Combined with the total number of passengers scheduled, these data can be used to compute average ride time (passengers per passenger hours), and vehicle productivity (passengers per vehicle hours). These measures can be useful in judging the effect of variations on service constraints, service quantity and quality weights, and so on.

Recent Implementations of CRSS

The task of quantitatively evaluating the effectiveness of a computerized paratransit scheduling program is often diffi-

Date: 03/14/90 CRSS Quick Schedule Report Page: 1
Time: 03:13 PM

Service Date: 03/15/90 Total Trips: 54 Scheduled: 49
Unscheduled: 5

Estimated Vehicle Statistics Carrier: COL Colonial Taxi

Run Number	Veh. hours	Revenue hours	Deadhead hours	Idle hours	Passengers Amb.	Passengers WC	Passengers Tot	Passenger Hours Amb.	Passenger Hours WC	Passenger Hours Total
1	10.1	2.4	0.8	6.9	11	4	15	4	1	5
2	12.6	3.9	1.9	6.8	22	10	32	5	3	8
Total	22.7	6.3	2.7	13.7	33	14	47	9	4	13

Runs Printed: 2

Date: 03/14/90 CRSS Quick Schedule Report
Time: 03:13 PM

Service Date: 03/15/90 Total Trips: 54 Scheduled: 49
Unscheduled: 5

Estimated Vehicle Statistics Carrier: YEL Yellow Cab

Run. Number	Veh. hours	Revenue hours	Deadhead hours	Idle hours	Passengers Amb.	Passengers WC	Passengers Tot	Passenger Hours Amb.	Passenger Hours WC	Passenger Hours Total
3	4.7	1.7	1.1	1.9	4	5	9	1	1	2
4	5.2	0.3	0.5	4.4	9	2	11	1	0	1
Total	9.9	2.0	1.6	6.3	13	7	20	2	1	3

Vehicles Printed: 2

FIGURE 11 Sample CRSS schedule statistical summary.

cult. This is not because of the lack of information provided by the software, rather it results from the difficulty in obtaining an historical data base for comparison purposes. It is often the case that a paratransit agency is very tightly funded, and the efforts of the entire staff are involved in providing service. Little time is left for collection of operational data beyond the requirements imposed by funding agencies. Under such circumstances, it is difficult make a quantitative judgment on the improvement between manual and automated scheduling.

Data that typically are available for comparison involve the total number of passengers served with the number of vehicles available, and the total effort required to produce schedules. Information on such variables as average passenger ride time and on-time service, among other things, is at best only collected according to UMTA Section 15 reporting criteria, which are oriented towards fixed-route service, not paratransit. A test using similar data under similar conditions is often not possible.

An agency where some good data has been available is Central New York Regional Transportation Authority, of Syracuse, New York (Centro). They did not computerize an entirely manual system; rather COMSIS Paratransit Information System software replaced an older paratransit MIS. Shortly thereafter, CRSS was installed; and the old system was run parallel with the new one.

At Centro, significant benefits were realized in improving productivity as well as reducing the staff time required to prepare schedules. Part of this resulted from the new MIS software: the average telephone request was handled in approximately 3 min, which was a 300 percent improvement over the old system. The average number of requests per hour increased from 50 to 125.

Improvements strictly relative to CRSS were in the area of staff effort and vehicle productivity. Vehicle productivity improved from 2.1 passengers per vehicle hour to over 2.8 passengers per vehicle hour. This translates to an increase from 375 trips per day to 450 trips per day. The level of staff effort is significantly reduced. Under previous methods, 2 to 4 overtime hours were necessary every day to finalize schedules. Under current conditions, no overtime is typically needed.

It is necessary to reiterate that the level of improvement at any given agency will vary. An example might be Miami Valley Regional Transit Authority, of Dayton, Ohio. In this case, the agency serves overall a small number of trips; typically 250 per weekday. Under this relatively light load, the scheduling staff produces good schedules in less than 4 hr a day. CRSS schedules have no noticeable improvement in productivity, although the CRSS scheduling process takes approximately 15 min.

Future Development

From experience at these sites, as well as others, COMSIS is implementing a number of improvements to CRSS. Among these are two dealing specifically with the way that the highway network is used.

The first is a method of better identifying trips "along the way." Here, CRSS will consider all the nodes in the minimum path between the origin and destination of a trip. All trips at that time of day will be checked against the list of nodes in the minimum path. If there is a trip going in the same direction along any part of the minimum path, it will be considered part of the group.

Many agencies have a significant variation in the geographical density of service requests. At any given time of day, there are a number of trips in one area. A half hour later, the activity moves elsewhere. If a vehicle "runs out of work" in an area, CRSS tends to send it to where the work is instead of letting it sit idle, even though there will be work in the area 15 to 30 min later. Under such circumstances, vehicles migrate throughout the service area.

For agencies where this is not desirable, it is necessary to place penalties on live and deadhead travel distance independent from travel time. To lessen migration, one would emphasize a penalty of deadhead miles, but not necessarily deadhead and idle time. The overall productivity may be less, but vehicles would tend to work in the same areas.

CONCLUSION

The development of CRSS has been ongoing since 1987. A great deal of experimental effort has been directed toward providing a practical paratransit scheduling package. The author would like to emphasize the practical implications of the paratransit scheduling process upon the development of CRSS. The scheduling problem is very difficult and complex. There is a calibration process, parallel to that encountered in urban transportation planning. Users of the package must be trained to enter realistic trip requests. (One cannot, for example, expect the software to schedule ten widely scattered pickups at the same time to the same vehicle, even though it is what was being done via manual scheduling.) It has also been learned that the highway network information can be used effectively to determine proximity in the process of grouping.

Results have been favorable as of early 1990. It is believed that CRSS is a viable response to scheduling needs in the paratransit industry. During this effort, it was believed that a great deal has been learned about automating the scheduling process. As progress continues, it is expected that further refinements to the package will be made.