Fleet Size and Mix Optimization for Paratransit Services

Liping Fu and Gary Ishkhanov

Most paratransit agencies use a mix of different types of vehicles ranging from small sedans to large converted vans as a cost-effective way to meet the diverse travel needs and seating requirements of their clients. Currently, decisions on what types of vehicles and how many vehicles to use are mostly made by service managers on an ad hoc basis without much systematic analysis and optimization. The objective of this research is to address the underlying fleet size and mix problem and to develop a practical procedure that can be used to determine the optimal fleet mix for a given application. A real-life example illustrates the relationship between the performance of a paratransit service system and the size of its service vehicles. A heuristic procedure identifies the optimal fleet mix that maximizes the operating efficiency of a service system. A set of recommendations is offered for future research; the most important is the need to incorporate a life-cycle cost framework into the paratransit service planning process.

The planning of demand-responsive paratransit services requires addressing two fleet-related decision problems: what types of vehicles to use and how many vehicles to use. The underlying problem is commonly referred to as fleet size and mix (FSM) problem and is of critical importance for a paratransit agency because it has an effect on both the costs of delivering the service (capital and operating costs) and the level of service (LOS) that can be provided to the clients in regard to comfort, convenience, and enjoyment (*1*).

Most paratransit agencies use a mix of vehicles of different types or sizes, from small sedans (for ambulatories only) to vans and small buses that have more seats and that can accommodate both ambulatories and wheelchair clients. The main advantage of using a fleet of mixed vehicles is the cost-effectiveness in dealing with variation in seating requirements as well as spatial and temporal clustering of requests.

Larger vehicles can accommodate, on a single trip, more passengers with different seating needs and thus allow for more ridesharing, which, in turn, can lead to higher productivity and fewer vehicles required to deliver the service. However, there are two potential drawbacks with using larger vehicles. First, using larger vehicles does not automatically yield higher productivity, because ridesharing is also limited by the time constraints of the clients. For example, a vehicle can take more passengers with its available seats, but it may not be able to do so because of the time constraints associated with some of its already committed passengers (e.g., maximum allowable ride time or pickup and delivery time windows). Second, the use of larger vehicles usually means higher capital and operating costs (higher fuel consumption per operating mileage or hour), higher emissions, and lower maneuverability.

L. Fu, Department of Civil Engineering, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada. G. Ishkhanov, TSS Wireless, Inc., 5750 Wilshire Boulevard, Suite 530, Los Angeles, CA 90036.

Transportation Research Record: Journal of the Transportation Research Board, No. 1884, TRB, National Research Council, Washington, D.C., 2004, pp. 39–46.

In addition, in situations of low demand (either in some subareas or time periods), opportunities for ridesharing are minimal, and smaller vehicles are often sufficient to handle the trips without any loss of efficiency.

While these conceptual relationships and issues are well understood in the paratransit industry, current planning decisions on fleet mix are mostly made by program managers on an ad hoc basis without much systematic analysis and optimization, and little systematic research effort has been devoted directly to this issue. The research closest to the problem determines the minimum fleet size for a paratransit service system, assuming either a fixed fleet mix or unlimited vehicle capacity (*2, 3*).

In contrast, the problem of selecting vehicle size for fixed-route services has attracted more attention, including both analytical studies (*4*) and simulation analyses (*5, 6*). Most of these models attempt to minimize the total sum of vehicle operating costs, the generalized costs of passenger waiting and ride time, and the social costs of traffic congestion. These models, however, cannot be extended to the paratransit fleet mix problem when the routing of individual service vehicles is not fixed.

The FSM problem has also been studied by the operations research (OR) community, mostly from a theoretical perspective. Because a solution to the FSM problem requires solving the vehicle routing and scheduling problem as a subproblem, the FSM is commonly dealt with as part of a joint problem called the fleet size and mix vehicle routing problem (FSMVRP). One of the most important contributions to this problem was made by Golden et al. (*7*), who proposed several solution heuristics based on the well-known Clarke and Wright saving algorithm (*8*). These algorithms have been further extended by Renaud and Boctor (*9*). However, all existing algorithms deal exclusively with the basic vehicle routing problems with no time constraints and precedence conditions (e.g., pickup stops must precede drop-off stops), and thus they cannot be used to solve the FSM problem coupled with a dial-a-ride routing and scheduling problem arising in paratransit services. To the authors' knowledge, no previous academic research deals specifically with the FSM problem associated with paratransit services.

This research has two primary objectives: (1) to show the importance of addressing the fleet size and mix problem and (2) to develop a practical procedure that can be used to determine the optimal fleet mix for a given service system. Ideally, a comprehensive analysis framework (e.g., life-cycle costing technique) that can take a systematic account of all cost elements associated with a fleet should be incorporated into the process of determining the optimal fleet size and mix of a paratransit service system. This research, however, focuses on the most technically challenging issue that must be addressed for such a full cost analysis procedure—that is, how the composition of a fleet would influence the operations of the fleet in regard to required fleet size (relating to capital costs) and system productivity (relating to operating costs). It is hoped that by attempting to address this issue, the present work would stimulate new research interest on this important problem and contribute to the development of a practical and systematic tool for paratransit services planning.

The paper first presents a real-life example to illustrate the relationship between the performance of a paratransit service system and the size of its fleet vehicles. A heuristic procedure is then proposed for identifying the optimal vehicle mix for a given operating condition. The procedure is subsequently applied to an example case from a paratransit agency with results compared with an ideal scenario and the actual fleet mix used by the agency. Some remaining issues are highlighted in the conclusion.

EFFECTS OF VEHICLE SIZE ON PARATRANSIT OPERATIONS

The objective of this section is to demonstrate how various performance indicators of a paratransit system, such as vehicle productivity, required fleet size, and average ride time, are related to the size (or seating capacity) of the service vehicles. The methodology employed to conduct the investigation is presented, followed by a sensitivity analysis on an example case.

As indicated previously, many factors other than vehicle size influence the performance of a paratransit system. Among the most important factors are the process and the technique used to schedule the trips. For example, if the trips were scheduled manually by schedulers, the outcome would depend on the level of skill and experience of the schedulers. For the computer-aided scheduling method, the results would depend on how service policy and objectives were handled by the scheduling algorithm and the mechanism of the algorithm itself, which might lead to schedules of different performance even for the same operating conditions. As a result, to isolate the effect of the vehicle size, an experimental procedure similar to that of Fu (*2*) was used.

A computer scheduling software called FirstWin (*10*) was used to generate schedules and associated performance statistics for the given case. Trips are then scheduled with a procedure called sequential iterative scheduling (SIS) to determine the number of vehicles required for a given operating condition, vehicle productivity, and LOS as represented by average excess ride time.

Two objectives are considered in scheduling. The first objective, which is also considered the top priority, is to minimize fleet size. This is achieved heuristically by using the neighborhood-based sequential insertion algorithm (NSI) in FirstWin, which considers vehicles one at a time and uses trip-clustering knowledge in the insertion process. From numerous experiments with both simulated and real cases, this algorithm was found to perform better than other algorithms such as the parallel insertion algorithm in regard to minimizing the number of vehicles required to service a given set of trips. Minimization of total travel time was explicitly considered in selecting trips and identifying optimal insertion positions in the scheduling algorithm. The created routes are subsequently improved with a set of improvement procedures.

The maximum allowable ride time ratio and service time window are considered as hard constraints that must be satisfied for all trips. These constraints define the minimum level of service that must be guaranteed for all trips. In all tests the fleet size is assumed to be unlimited so that all trips are guaranteed to be scheduled.

Trips are scheduled with the following scheduling procedure, the SIS algorithm:

Step 0. Import network/vehicle/trip data.
Step 1. Set maximum ride ratio (MRR) and service time window.
Step 2. Schedule all trips by using the NSI algorithm.
Step 3. Remove all trips from those vehicles that have fewer than three trips assigned and try to reassign removed trips to other scheduled vehicles by using the swap and reinsertion algorithm included in FirstWin. This is an attempt to reduce the number of vehicles required for delivering the service.
Step 4. Apply the improvement procedure to improve the generated schedules and go back to Step 3. This step continues for a prespecified number of iterations.
Step 5. Record the scheduling statistics, including the number of vehicles that have been scheduled with trips, vehicle productivity, average ride time, average excess ride, and deadheading time.

To investigate the sensitivity of service performance to vehicle size, a modified real-life example consisted of a weekday service covered by a paratransit service provider in Canada. Two cases were extracted for analysis: Case I represented a low-demand scenario consisting of 460 trips with both ambulatory (59%) and wheelchair trips (41%) over the peak period 11:00 a.m. to 2:00 p.m.; Case II represented a high-demand scenario with a total of 682 trips taken from the p.m. peak period (15:00 to 18:00). A fleet of identical vehicles with capacities ranging from [1, 1] to [10, 10] was assumed (where the first number represents the available ambulatory seats and the second number represents the available wheelchair seats; the same notation convention will be used in the following discussion without further explanation). In the original trip data file, some trips were group trips consisting of a group of people from the same origin to the same destination. The authors also converted multiple passenger trips into trips with at most one passenger for each seating type. Demand distributions of these two cases are shown in Figure 1. Trips were scheduled using the SIS procedure described previously, with a fixed time window of 30 min and a set of maximum ride ratios of 1.2, 1.5, and 2.0 (a maximum of 20%, 50%, and 100% extra ride time as compared to direct ride time).

Figures 2 and 3 give the effects of vehicle size on three key performance indicators of a paratransit system, namely, vehicle productivity (Figures 2*a* and 3*a*), fleet size (Figures 2*b* and 3*b*), and passenger ride time (Figures 2*c* and 3*c*). Corresponding numerical values are given in Table 1. For vehicle productivity and fleet size, the figures give the percentage of increase relative to the scenario with vehicles of unlimited seating capacity (50 seats per vehicle). From these results, the following observations can be made:

1. Vehicle size has a significant impact on vehicle productivity and the number of vehicles required. As expected, the larger the vehicles, the higher the average vehicle productivity and the smaller the required fleet size. However, there exists a critical point beyond which additional capacity would not result in better performance. This finding confirms the authors' expectation that it may not always make economic sense to use large vehicles.
2. The optimal vehicle size of a fleet depends on the level of travel demand that the fleet will cover. In comparing Figures 2 and 3, one can see that larger vehicles should be used in high-demand cases. This makes intuitive sense, for the higher the demand is, the more opportunities there are for ridesharing, and thus the more advantageous it is to use larger vehicles.

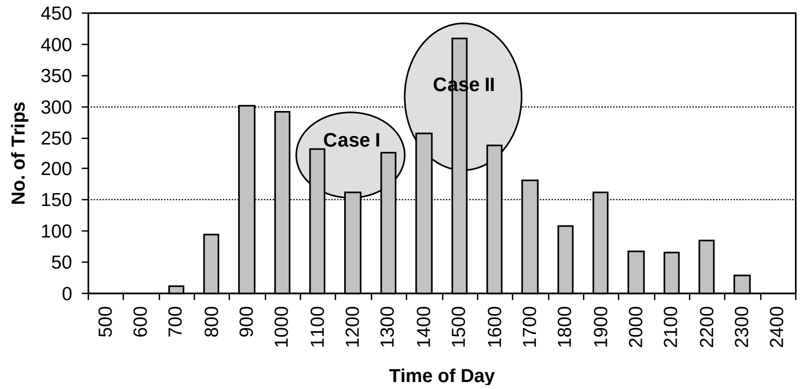

FIGURE 1 Trip distribution over time of day.

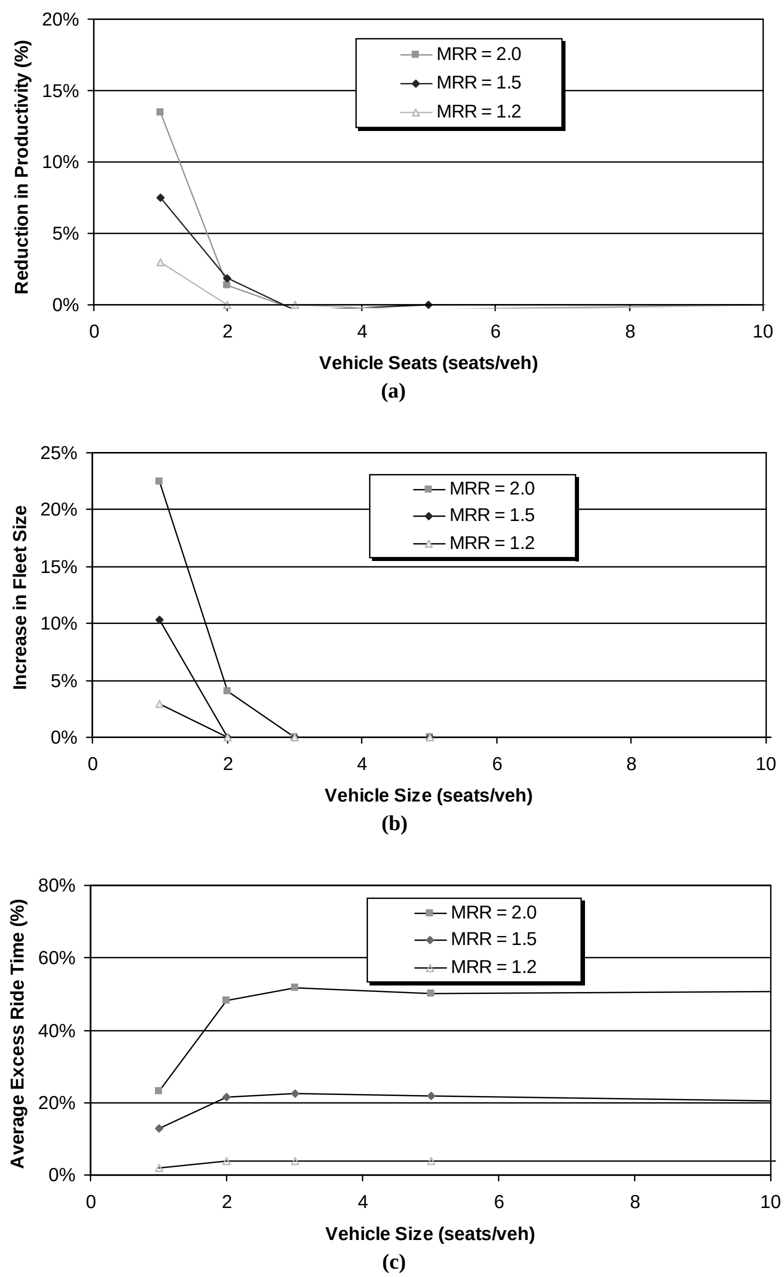

FIGURE 2 Effects of vehicle size on system performance: low-demand case showing effect on (*a*) vehicle productivity, (*b*) fleet size, and (*c*) passenger ride time.

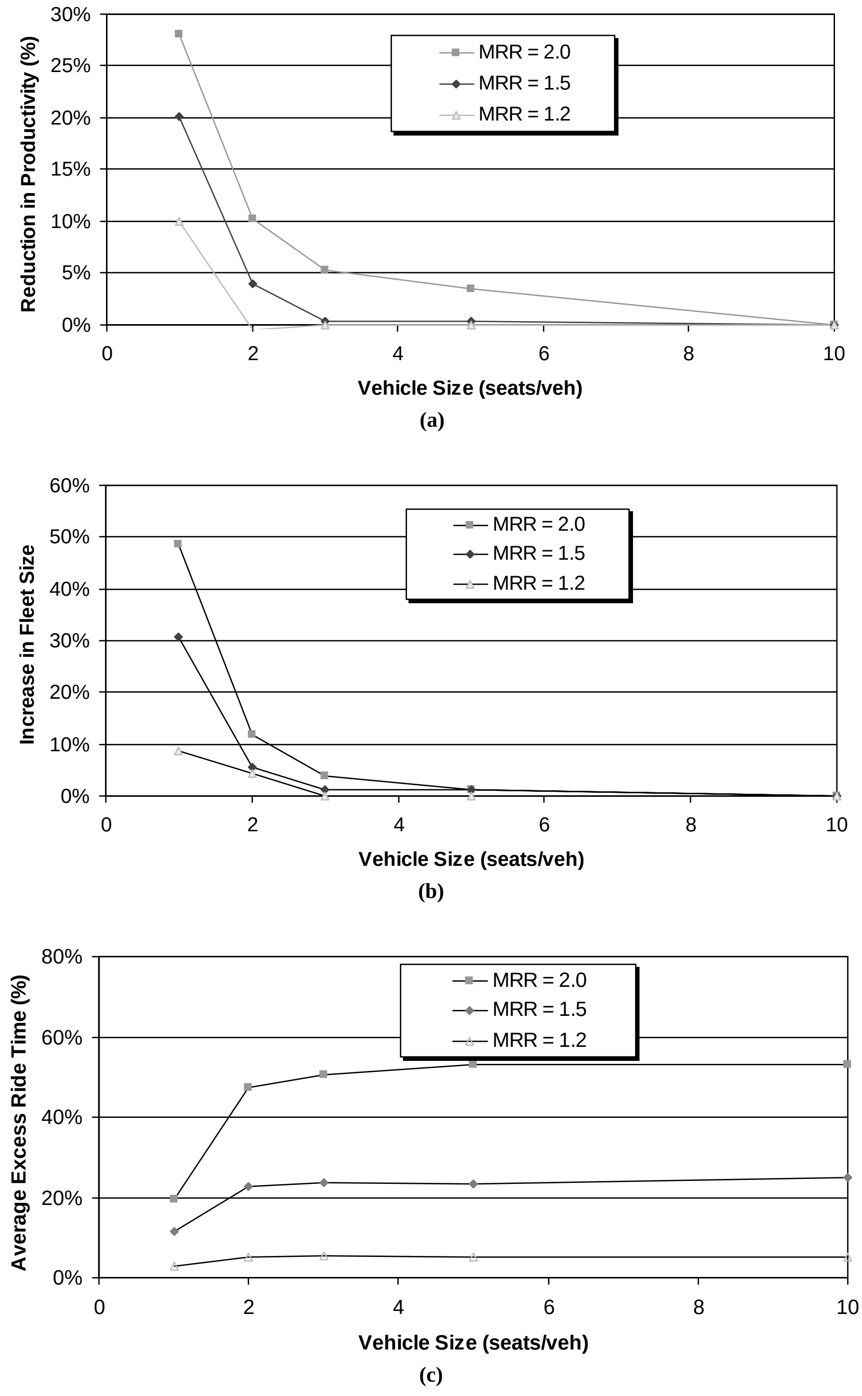

FIGURE 3 Effects of vehicle size on system performance: high-demand case showing effect on (*a*) vehicle productivity, (*b*) fleet size, and (*c*) passenger ride time.

3. In addition to travel demand, another factor that influences the optimal vehicle size is the service constraint and policy. Figures 2 and 3 show the system performance under different levels of ride deviation allowed for each trip, ranging from a high LOS scenario with an MRR of 1.2 (extra ride time should not exceed 20% of the direct ride time) to a low LOS scenario (MRR of 2.0). In the high LOS scenario, the scheduling process is mainly constrained by the ride time window, and therefore ridesharing becomes less possible: smaller vehicles should be sufficient. Conversely, larger vehicles may be more cost-effective under more relaxed service constraints (e.g., MRR of 2).

HEURISTIC PROCEDURE FOR DETERMINING OPTIMAL FLEET MIX

The performance of a paratransit fleet and the size of its service vehicles were discussed, as well as a critical point in vehicle size beyond which additional capacity could become ineffective. If the capital and operating costs of a vehicle increase as does the size of the vehicle, one should expect an optimal vehicle size that minimizes the total cost of the system. For simplicity, the analysis is limited to identifying a fleet of vehicles that maximizes the operating efficiency (vehicle pro-

TABLE 1 Effects of Vehicle Size on System Performance

MRR	Vehicle Size (seats/veh)	Travel Time (h)	Productivity (trips/h)	(%)	Excess Ride (%)	No. of Vehicles	(%)
Low-Demand Scenario							
1.2	1	332.2	1.88	10	3	124	9
	2	287.7	2.10	0	5	119	4
	3	297.2	2.09	0	5	114	0
	5	297.1	2.09	0	5	114	0
	10	297.1	2.09	0	5	114	0
1.5	1	311.9	1.99	20	12	115	31
	2	260.0	2.39	4	23	93	6
	3	251.1	2.48	0	24	89	1
	5	250.5	2.48	0	23	89	1
	10	250.2	2.49	0	25	88	0
2.0	1	288.4	2.03	28	20	113	49
	2	235.8	2.53	10	47	85	12
	3	228.4	2.67	5	51	79	4
	5	228.7	2.72	4	53	77	1
	10	220.8	2.82	0	53	76	0
High-Demand Scenario							
1.2	1	220.3	2.3	b	2	70	3
	2	212.9	2.4	0	4	68	0
	3	212.3	2.4	0	4	68	0
	5	212.0	2.4	0	4	68	0
	10	215.1	2.4	0	4	68	0
1.5	1	205.0	2.5	8	13	64	10
	2	193.5	2.6	2	21	58	0
	3	188.6	2.7	0	23	58	0
	5	189.6	2.7	0	22	58	0
	10	189.4	2.7	0	23	58	0
2.0	1	201.8	2.5	13	23	60	22
	2	177.1	2.9	1	48	51	4
	3	173.5	2.9	0	52	49	0
	5	175.3	2.9	0	50	49	0
	10	174.5	2.9	0	53	49	0

ductivity) of the service system, subject to certain service policies and constraints. The underlying problem is not trivial, because the relationship between system performance and fleet mix is not explicitly known; no simple analytical expression can be developed. Also, there could be an enormous number of combinations in both seating combination for a given vehicle and vehicle mix of different types: thus enumeration approaches are computationally intractable.

In recognition of the challenge, a heuristic algorithm, called scheduling, matching, allocation, and reduction (SMAR), is proposed to address the problem. The algorithm assumes that the types of vehicles available for selection are given; the objective is to decide the number of vehicles of each type to be used. The heuristic is fundamentally a greedy search procedure, with the idea of using as many small vehicles as possible without loss of productivity. It includes five major steps as follows (see Figure 4):

1. Prepare representative cases. Cases represent the typical operating conditions of the service system, including trip database, travel time and speed estimates, and service constraints.
2. Schedule trips with an idealized fleet. For each case, trips are scheduled with an idealized fleet composed of vehicles with unlimited seating capacities for both ambulatory passengers and those with wheelchairs. The scheduling process is constrained only by time windows and ride time limits. The fleet is assumed to be sufficiently large to handle all the trips. The scheduling process should follow a predefined procedure as in the SIS algorithm described in the previous section, with the objective of using as small a fleet as possible. The output of this scheduling process is a set of routes, each assigned a set of trips to be covered by a vehicle of unlimited capacity.
3. Determine the distribution of routes by required vehicle size. For each generated route, a loading profile can be established to show the number of passengers on board over the vehicle trip duration. The maximum number of seats (seating capacity or vehicle size) required for a vehicle to cover the route can be subsequently determined. Figure 5 shows an example of such loading profiles, from which it can be observed that the route is scheduled to carry both ambulatory and wheelchair passengers. To cover this route, the assigned vehicle must have a minimum of seven wheelchair seats and four regular seats for ambulatory passengers. Accordingly, the required vehicle size, or the minimum number of seats for each seating type, for each scheduled route can be determined. To visualize the distribution of all the scheduled routes by their minimum seating requirements, a bubble chart can be used. Figure 6 is an example of a bubble chart to display the distribution of routes by required vehicle size. In this example, 13 routes were created, including 4 routes requiring vehicles with [2, 3] seats, 3 having vehicles with [4, 2] seats, 5 having vehicles with [4, 1] seats, and 1 having a vehicle with [10, 4] seats.
4. Match and allocate vehicles. Following Figure 5 as an example, if there were four types of vehicles (V_1, V_2, V_3, and V_4) with seating combinations of V_1, [2, 3], V_2, [4, 1], V_3, [4, 0], and V_4 [10, 4], we could simply decide to use four vehicles of V_1, three vehicles of V_2, five vehicles of V_3, and one vehicle of V_4. In reality, however, the

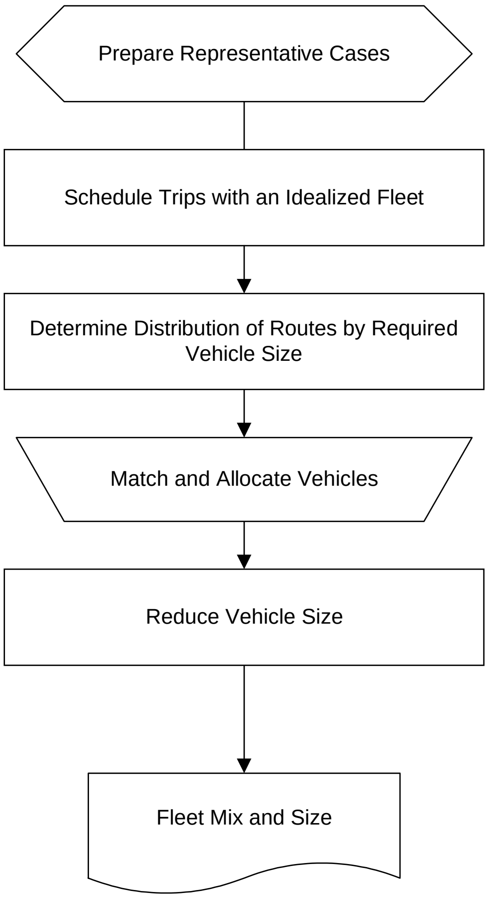

FIGURE 4 SMAR procedure for solving fleet mix and size problem.

distribution of routes by required vehicle size is usually much more spread out than what the available types of vehicles can provide. As a result, some routes need to be allocated with vehicles that do not match the exact seating needs. For the example shown in Figure 5, suppose there are only two types of vehicles, V_1 [4, 0], and V_2 [5, 5], for selection. Only one group of the routes can be assigned with an exactly matched vehicle type, and the other groups of routes need to be allocated with vehicles of unmatched capacity. Many alternatives may be used to allocate vehicles of specific types to the routes generated on the basis of an idealized fleet. One alternative would be to

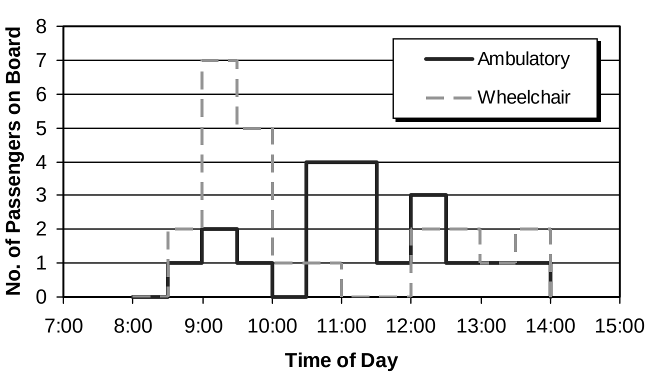

FIGURE 5 Loading profile of route.

let the planners decide the allocation so that they can incorporate their experience into the process as much as possible. An automated process is also possible, with distribution of routes on the basis of some heuristic rules. For example, one such rule would be to assign routes to the smallest vehicles that can handle routes and assign any routes that cannot be handled by just any types of vehicles to the largest vehicles.

The result of this matching and allocation step is an initial solution to the FSM problem. To make sure that the resulting fleet can handle all the trips, one needs to reschedule the trips with the new fleet (with limited capacity). If there are trips that cannot be assigned to just any of the existing vehicles, one would increase the number of the largest vehicles by one. This step continues until all the trips are scheduled.

5. Reduce vehicle size. The logic of the previous step may lead to overuse of large vehicles. This step initiates a subprocess with the objective of reducing the number of large vehicles. The following steps are involved:

a. Iterate from the largest vehicle type to the smallest vehicle type.
b. For the current vehicle type, convert one of its vehicles into one of next smaller vehicle type (one size smaller).
c. Reschedule all trips and compare the scheduling statistics between before and after the change. If there is no or negative change in schedule performance, keep the original fleet mix (before the conversion) and select the next vehicle type. Go back to Step 5*b.* Otherwise, keep the change, and go back to Step 5*b.*

It is important to emphasize that to take into account variations in passenger demand, demand scenarios of many operating days need to be evaluated with the above algorithm. The result should be a distribution of optimal fleet size and mix, from which the final solution can be obtained by incorporating other factors, such as fleet maintenance, life-cycle costs of vehicles, and maneuverability of vehicles on narrow streets. The following example illustrates the application of the SMAR heuristic using a real-life example.

EXAMPLE

The objective of this section is to demonstrate the application of the proposed heuristic for fleet mix optimization. Again, a real-life case is used from the same service provider, with the following operating conditions: a whole day service with a total of 2,992 trips (trip distribution by time of day is the same as shown in Figure 1); a service time window of 30 min; and a maximum ride ratio (MRR) of 1.5 (i.e., 50% excess ride time). Currently, a fleet of vehicles with capacities ranging from [4, 0] to [13, 8] is used to provide the service. According to their seating capacity, the fleet vehicles are grouped into four types—V_1 [4, 0]; V_2, [6, 5]; V_3, [14, 0]; and V_3 [10, 5]—which were also assumed to be the only vehicle types available for selection.

Following the MAR algorithm described in the previous section, trips were first scheduled using a fleet of vehicles with unlimited capacity. The loading profile of each scheduled route was then examined to determine the minimum number of seats required for each seating type. A manual process was then initiated to match and allocate the routes to the four vehicle types. Figure 7 schematically illustrates the matching and allocation results, which suggest that 40 vehicles of V_1, 32 vehicles of V_2, 21 vehicles of V_3, and 19 vehicles of V_4 should be used. This represents the initial solution of the optimization process discussed previously. Instead of invoking the

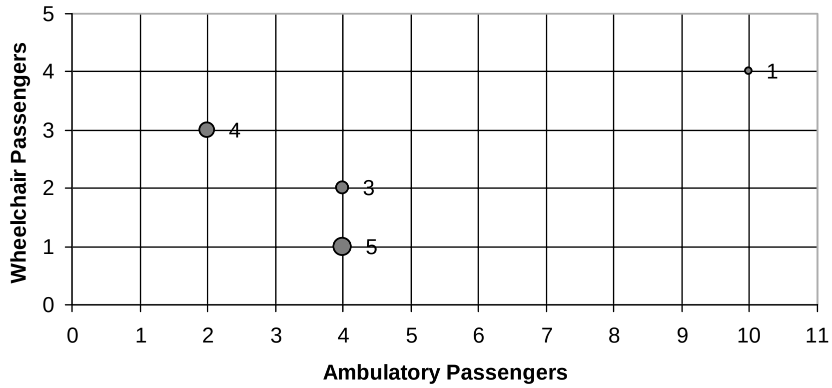

FIGURE 6 Distribution of routes by seat requirements.

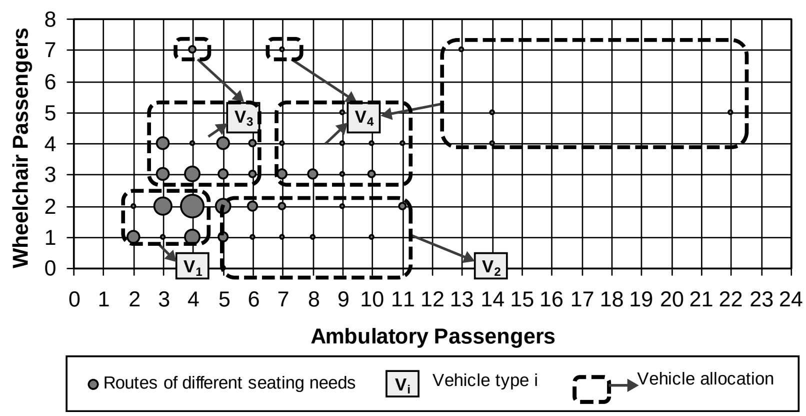

FIGURE 7 Determination of fleet mix by SMAR.

reduction step to obtain an improved solution, the authors directly used this initial solution and compared it with the ideal scenario (unlimited fleet size and vehicle capacity) and the current fleet mix. Table 2 gives the scheduling statistics from the following three scenarios: current fleet mix, fleet mix by SMAR, and ideal fleet mix. It can be observed that the number of vehicles as determined by SMAR (112 vehicles) was approximately 10% higher than in the case of ideal fleet mix (103 vehicles). However, the values of other performance indicators, such as productivity, deadheading, and average ride time, were quite close to the ideal scenario. In contrast, compared with the results from the current fleet mix, the optimization process has resulted in a smaller fleet (two vehicles less) and higher productivity

TABLE 2 Comparison of Solution Results

Scheduling Statistics			Scenarios		
			Current Fleet Mix	Optimal Fleet Mix by SMAR	Ideal Fleet Mix
Total No. of Vehicles Used			114	112	103
No. of Vehicles by Type		Seats			
	V_1	[4,0]	58	40	N/A
	V_2	[6,5]	42	32	N/A
	V_3	[14,0]	9	21	N/A
	V_4	[10,5]	5	19	N/A
Total Service Hours			761.02	642.62	621.4
Productivity (trips/veh/hour)			4.43	5.25	5.43
Average Ride Time (min/trip)			19.95	18.75	18.36
Average Excess Ride (%/trip)			18.12	18.12	20.25
Average Deadheading (hours/veh)			1.7	1.2	1.2

(18% higher). These differences have further demonstrated the potential of and need for fleet mix optimization.

CONCLUDING REMARKS

Paratransit operations planning requires determining the composition of service fleet that is most cost-effective for a given operation environment and best for their clients. While many factors influence the selection of vehicle types, this paper approached the problem from the perspective of service efficiency. Experimental results from a case analysis have indicated the dependence of the performance of a paratransit system on the size of its fleet vehicles and the existence of a critical point beyond which additional capacity becomes ineffective. The paper proposes a heuristic procedure that can be used by paratransit agencies to identify the optimal fleet mix for their specific operating conditions and environments. It should be cautioned, however, that the results from this research are still preliminary and further research is needed in the following directions:

1. A comprehensive life-cycle cost analysis framework that would incorporate the proposed model and recognize the differences in the capital and maintenance costs of different types of vehicles should be introduced.
2. Some agencies may experience wide variations in manifests from one day to the next, and typical operating conditions are difficult to identify. Under such circumstances, sufficient number of days of operations must be analyzed to identify the distribution of optimal fleet size and mix.
3. The fleet size and mix analysis should also take into account future travel demand so that the recommended fleet is flexible enough to handle future travel demand variations. In addition, a demand prediction model should be part of the analysis.
4. The proposed fleet mix optimization procedure was implemented in a semiautomatic manner. A software tool that automates the whole process could therefore be useful.
5. In practice, many factors other than service efficiency have to be taken into account in selecting the appropriate types of vehicles. It would therefore be valuable to develop a synthesis that summarizes current practice and experience related to life-cycle costs of different types of vehicles, vehicle selection, and fleet mix planning.

ACKNOWLEDGMENT

This research was supported by the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

1. Tolliver-Nigro, H. Choosing A Shuttle Bus: What Should Agencies Look for When Purchasing Vehicles for Their Fleets? *Mass Transit,* Jan.–Feb. 1999, pp. 42–48.
2. Fu, L. Analytical Model for Paratransit Capacity and Quality-of-Service Analysis. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1841*, TRB, National Research Council, Washington D.C., 2003, pp. 81–89.
3. Schofer, J. L., B. L. Nelson, R. Eash, M. Daskin, Y. Yang, H. Wan, J. Yan, and L. Medgyesy. *TCRP Report 98: Resource Requirements for Demand-Responsive Transportation Services*, TRB, National Research Council, Washington, D.C., 2003.
4. Jansson, J. O. A Simple Bus Line Model for Optimization of Service Frequency and Bus Size. *Journal of Transport Economic Policy,* Vol. 14, No. 1, 1980, pp. 53–80.
5. Glaiseter, S. Bus Deregulation, Competition and Vehicle Size. *Journal of Transport Economic Policy,* Vol. 20, No. 2, 1986, pp. 217–244.
6. Bly, P. H., and R. H. Oldfield. Competition between Minibuses and Regular Bus Services. *Journal of Transport Economic Policy,* Vol. 20, No. 1, 1986, pp. 47–68.
7. Golden, B. L., A. A. Assad, L. Levy, and F. Gheysens. The Fleet Size and Mix Vehicle Routing Problem. *Computers and Operations Research,* Vol. 11, 1984, pp. 49–66.
8. Clarke, G., and J. Wright. Scheduling of Vehicles from a Central Depot to a Number of Delivery Points. *Operations Research,* Vol. 12, 1964, pp. 568–581.
9. Renaud, J., and F. F. Boctor. A Sweep-Based Algorithm for the Fleet Size and Mix Vehicle Routing Problem. *European Journal of Operational Research,* Vol. 140, 2002, pp. 618–628.
10. TSS Wireless Inc. *User's Guide: FirstWin—A Tool for Routing and Scheduling Dial-A-Ride Paratransit Vehicles,* Los Angeles, Calif., 2003.

Publication of this paper sponsored by Paratransit Committee.