Evaluation of Zoning Design with Transfers for Paratransit Services

Chung-Wei Shen and Luca Quadrifoglio

This paper evaluates the effects of including transfers between service zones on overall service performance in a paratransit system. Transfers were included to improve the operational efficiency of a system when maintenance of a desirable zoning structure was obligatory. This proposed innovative service design was compared with more traditional cases of no transfer zoning and no zoning. A set of instances was generated from demand data obtained from the Metropolitan Transit Authority of Harris County, Texas, and evaluated through simulation analyses. The results demonstrated that under a zoning structure, this transfer design (in comparison with a nontransfer design) provided noticeable improvements in efficiency measures and better passenger trips per vehicle revenue hour while maintaining a minimum customer service standard; however, the overall performance of the no-zoning strategy used by the Houston, Texas, Metropolitan Transit Authority of Harris County performed the best, on average.

This paper investigates the different organizational structures of paratransit services that are used in large regions. A paratransit service is a demand-responsive shared-ride transit service that uses vans or small buses. It is characterized by the use of vehicles that do not operate on a fixed route or on a fixed schedule. The paratransit route and schedule are arranged according to a user-specified origin and to a user-specified destination at a user-specified time.

The size of the service area is one of the key factors that affects the productivity of demand-responsive transit. Americans with Disabilities Act (ADA) paratransit service is a type of demand-responsive transit that provides transportation to people with disabilities. In general, the larger the service area is, the longer the trip length will be, and thus, demand-responsive transit will not always be able to serve a given number of passengers consistently in a specified amount of time (*1*). The impact of areas of different sizes on the productivity of transit was first studied by Wilson et al. (*2*). They demonstrated that the number of vehicles used is proportional to the size of the service area. Chira-Chavala and Venter adopted the data provided by the Outreach Paratransit Service in Santa Clara County, California, and observed that longer trip lengths contributed to an increase in empty trip miles within an expanding service area (*3*).

In addition, large areas usually necessitate trips more dispersed than those enjoyed by more compact service areas. This pattern of dispersed trips, which translates into a lower demand density, makes it difficult to achieve the most beneficial effects of ridesharing. It has been shown that in low-density areas, demand-responsive transit systems have a lower level of productivity than systems that function in higher-density municipal areas (*4*).

To retain productivity by focusing on shorter trips within denser areas, some larger systems have outsourced operations to more than one contractor, with each contractor responsible for the service zone to which its vehicles have been assigned. This service design is called a "zonal structure" or a "zoning approach." Adjacent zones have no overlapping or shared buffer areas. The zoning approach is attractive not only because it creates more manageable pieces of work but also, more importantly, because it establishes an ongoing spirit of competition throughout the contract term (*5*). Zonal demand-responsive service is also used for dispatching, as well as for fare determination purposes (*6*).

Developers of zoning strategies, however, need to decide how to accommodate trips that cross into different zones. The zonal approach can be divided into two variations: (*a*) zoning without transfer, such as the service provided in Los Angeles County, California, and (*b*) zoning with transfer, such as the service provided in the Twin Cities metropolitan area in Minnesota. In zoning without transfer, interzonal customers may not need to switch vehicles during their trips. Alternatively, a system of zoning with transfer may require interzonal customers to switch vehicles. Quadrifoglio et al. performed a simulation study to test the productivity of zoning without transfer, comparing the performance of that strategy with a centralized, no-zoning case based on data obtained from Los Angeles, California (*7*). Shen and Quadrifoglio investigated the zoning-without-transfer designs used by the ADA paratransit system in Houston, Texas (*8*). They concluded that centralized cases perform better than zoning without transfer according to the number of passenger trips per vehicle revenue hour. The decrease in the number of passenger trips per vehicle revenue hour is probably the result of the higher number of empty trip miles that tend to occur in the system with a zoning-without-transfer design. Introduction of transfers to interzonal customers would be a promising method of decreasing these empty trip miles.

The zoning-with-transfer system coordinates vehicle schedules at various transfer locations. The schedule coordination of interzonal mechanisms of transportation likely reduces trip costs because of an increase in the rideshare rate and a lowering of the number of empty return miles (*9*). The proper coordination of paratransit services would increase not only efficiency and productivity but also mobility.

Although the operational consolidation of providers appears to achieve economies of scale, the following may impede their coordination: (*a*) a user may have some concern that the current service level will decrease, (*b*) the sponsoring agency may have doubts over whether the cost savings is significant, and (*c*) the different jurisdictions within which component transportation systems operate

---

Zachry Department of Civil Engineering, Texas A&M University, College Station, TX 77843-3136. Corresponding author: L. Quadrifoglio, quadrifo@tamu.edu.

Transportation Research Record: Journal of the Transportation Research Board, No. 2277, Transportation Research Board of the National Academies, Washington, D.C., 2012, pp. 82–89.
DOI: 10.3141/2277-10

may have different operational standards particularly designed to meet local riders' needs (*5*). To the best of the authors' knowledge, no quantitative evidence exists to demonstrate the benefits and concurrent costs that occur from adapting a zoning-with-transfer design to a large-scale paratransit system.

The scheduling and routing of classic paratransit systems is known as the "dial-a-ride problem," in the terminology used for research on vehicle routing problems. A dial-a-ride problem without ride time constraints is denoted by the term "pickup-and-delivery problem." The most recent surveys published on the dial-a-ride problem and pickup-and-delivery problem were presented by Cordeau and Laporte (*10*) and Berbeglia et al. (*11*), respectively.

A paratransit service that uses a transfer system is a generalization of the dial-a-ride problem. The transfer of passengers will always require more than one vehicle to fulfill a trip; therefore, the spatial and temporal synchronization constraints will, by necessity, be imposed on more than one vehicle. A schedule delay in one vehicle route may necessitate a change to all other routes. Therefore, solutions to such problems are computationally difficult, even when one is simply trying to develop a heuristic algorithm.

Shang and Cuff have provided a concurrent heuristic approach to solve the issue of the pickup-and-delivery problem with transfer, using as an example a health maintenance organization (*12*). They showed that their proposed heuristic performed better than the health maintenance organization's scheduling heuristic, according to the overall lower number of delays, total travel time (in hours), and total number of vehicles. Cortes et al. studied a pickup-and-delivery problem with transfers through a process of mixed-integer programming (*13*). They found that the transfers permitted a higher level of efficiency in the total vehicle travel time. Because of the complexity of the problem, this method could be applied only to a very small number of customers, which was maximized at six customers.

In this paper, a heuristic-based simulation was used as a study method to better understand the effects of zoning and zoning with transfers on paratransit services. In an experiment performed with data from Houston's demand-responsive service, the productivity and service levels of three organizational structures were compared: zoning with transfer, zoning without transfer, and no zoning. In the zoning-without-transfer structure, each zonal service provider can pick up only those customers whose pickup location is within their service area; however, the provider is allowed to drop off customers outside their service area. Each provider is unaware of the state of the system in the other zones. Alternatively, the no-zoning control system is totally centralized. In this scenario, basic paratransit service vehicles are, in general, allowed to move within the entire service area.

The main contribution of this report is the quantitative evidence that shows the effects of transfer design on zoning policy.

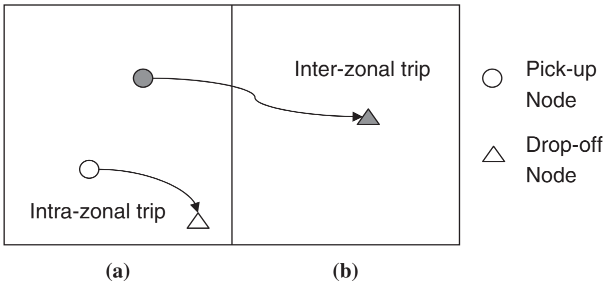

FIGURE 1 Categories of trip by zonal structure.

The rest of this paper is organized into four sections. The paratransit services of the zoning-with-transfer system are first defined. The demand data used are then described, and the computation of the algorithm is outlined. Finally, the results of the simulation are summarized.

SYSTEM DEFINITIONS

This section provides descriptions of two zoning strategies, those with and those without transfers, and a summary of the scheduling procedures used.

Within a demand-response service area, the service provider may subdivide the service area into zones. A "zone" is a geographical boundary. A list of customers will request a certain number of trips, and each trip has a specific pair of scheduled pickup and drop-off locations, as well as a desired pickup (or drop-off) time attached to each pickup (or drop-off) location. Each pickup and drop-off is considered a "node" in the system. Each trip can be categorized as either an interzonal or an intrazonal trip, as determined by the pickup and drop-off locations. Trips with pickup and drop-off locations in different zones are "interzonal trips," and trips with pickup and drop-off locations within the same zone are "intrazonal trips" (Figure 1).

Under the zoning-without-transfer policy, vehicles in each zone are served and independently operated by different carriers. Figure 2 illustrates the characteristics of this policy. The pickup location of each customer determines which carrier is eligible to serve that customer. Vehicles are, however, allowed to traverse zone boundaries to drop off interzonal customers.

In zoning-with-transfer control, interzonal passengers must transfer from one vehicle to another at given transfer locations to reach their final destinations. Conversely, intrazonal passengers do

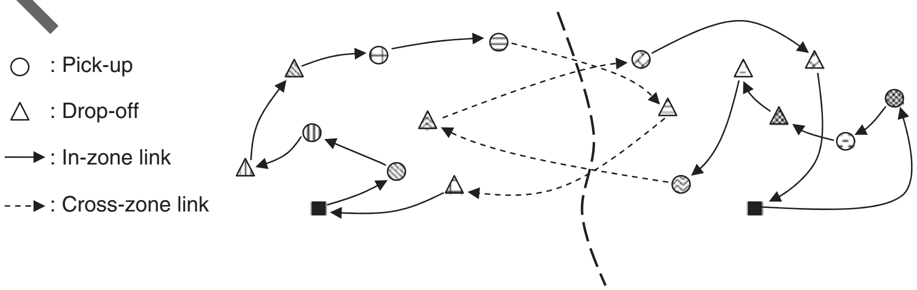

FIGURE 2 Zoning-without-transfer policy (each symbol represents a different customer).

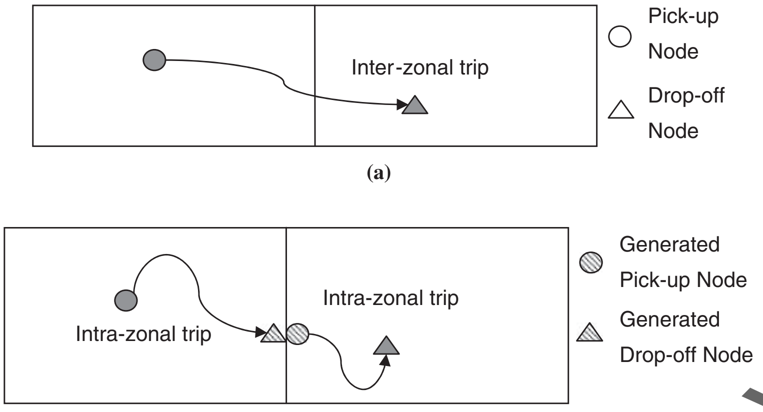

(a)

(b)

FIGURE 3 Example of generating an intrazonal trip.

not have to switch vehicles to complete their trips. To highlight the process of loading and unloading at various transfer locations, two corresponding nodes (a load node and an unload node) were generated at each transfer location for each interzonal trip (i.e., when a vehicle visits the transfer node, it either loads or unloads passengers according to the node's characteristics). Thus, an interzonal trip (Figure 3*a*) could be treated as two intrazonal trips when the schedules are coordinated, so that interzonal customers could switch vehicles at specific transfer locations (Figure 3*b*).

It is assumed that for each interzonal trip, vehicles can be switched only at particular transfer locations and only once per trip. The transfer locations at which a vehicle might stop are typically found on the boundaries of subzones. If passengers need to travel between zones that do not border one another, a transfer can be arranged at a suitable location between the two zones. Because of customer discomfort, more than one transfer might be undesirable, and under certain circumstances, it might be quite unreasonable. In practice, the passengers of the paratransit systems in both Chicago, Illinois, and Boston, Massachusetts, are assured that they will need to make, at most, one transfer (both systems use coordinated zoning systems).

For this study, the hard time windows were set as follows: the earliest arrival time is ET_i, and the latest departure time is LT_i (where $i = 1, 2, \ldots, N$) for both the pickup and the drop-off nodes. In the following context, $+i$ and $-i$ denote the points of pickup and drop-off for customer i, respectively. The earliest vehicle arrival time is denoted AT_i, and the earliest vehicle departure time is denoted DT_i. At pickup nodes, the time gap between ET_i and LT_i denotes the width of a predefined pickup time window. For example, one node may be a pickup home address scheduled within a half-hour window of time between 6:45 and 7:15 a.m.

In many demand-response scheduling systems' insertion algorithms, the objective is to minimize the vehicle travel distance while maintaining an acceptable level of service. To maintain such a service level, the ratio of the maximum ride time (MRT_i) to the direct ride time (DRT_i) needs to be within a specified value, R, called the "maximum ride time factor," for every customer. Therefore, ET_{-i} and LT_{-i} of the drop-off node would be decided by the corresponding ET_{+i} and LT_{+i} of the pickup node and R:

$$\mathrm{ET}_{-i} = \mathrm{ET}_{+i} + \mathrm{DRT}_i$$

$$\mathrm{LT}_{-i} = \mathrm{ET}_{+i} + R \times \mathrm{DRT}_i$$

if $\mathrm{LT}_{-i} < \mathrm{LT}_{+i}$, then

$$\mathrm{LT}_{-i} = \mathrm{LT}_{+i} + \mathrm{DRT}_i$$

R can be a constant (as is the case in Los Angeles County) or an inverse function of the direct trip length (as is the case in Houston), to avoid extremely long maximum trips for already long direct journeys.

Except in the case in which the pickup and drop-off vehicles arrive at the transfer location at exactly the same time, the vehicle arriving earlier must wait until the other vehicle arrives (i.e., customers would not be allowed to wait alone at transfer locations). The distances in Manhattan (New York City) are used to calculate the symmetrical travel distances between any two pairs of nodes. Quadrifoglio et al. verified that these estimated travel distances are close to the actual travel distances (*7*). The distance calculations imply that the network is arranged in a rectilinear grid pattern. It is also assumed that the system has no traffic jams and that the travel time between any two points is a matter of only the travel distance and the vehicle speed. This assumption might not allow calculation of the precise travel time between two points, but it does not alter the results of the following performance comparison. The link distances and speeds were input into the model and can easily be updated with more accurate values if and when those values become available.

SCHEDULING ALGORITHM

The new insertion-based heuristic makes use of the generic insertion framework of Solomon's sequential approach (*14*). This algorithm processes ride requests sequentially, inserting one customer into the vehicle schedule at a time until all requests have been serviced.

After sorting of all customers by requested pickup times, one empty route is generated in each service zone. Each empty route starts from and ends at the same depot. Every interzonal trip generates a drop-off and pickup node at a transfer location. According to the designated zone of each trip, the possible insertions of unassigned trips are searched for sequentially by their earliest pickup times. In this study, the insertion procedure from the first to the last unassigned trip is called one "round." A more detailed description of the procedure used for insertion review is described below. Those trips that cannot be inserted into the schedule during a round are copied to the

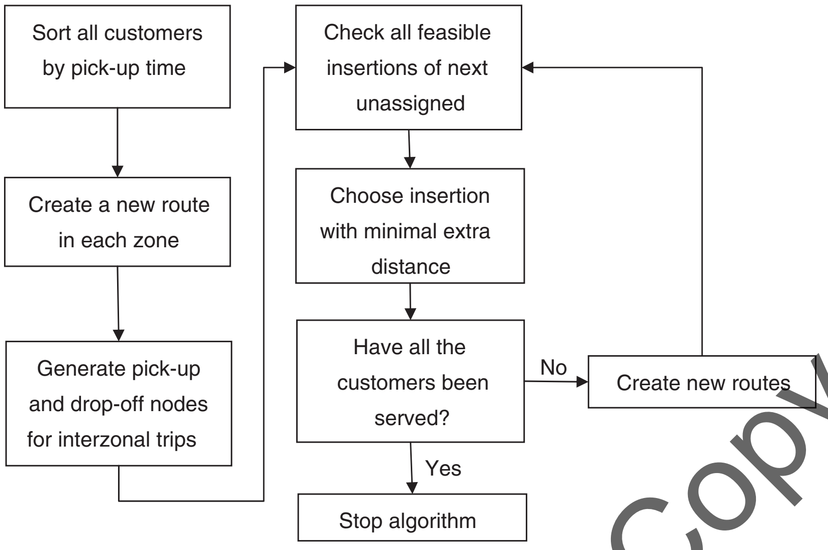

FIGURE 4 Flowchart of algorithm.

unassigned list of trips. This insertion procedure requires that one route be maintained in each zone during each round.

During the search procedure, four constraints are taken into consideration. First, the arrival time (AT$_i$) of a vehicle at the pickup (or drop-off) location can be no later than LT$_{+i}$ (LT$_{-i}$). Second, for each passenger, the drop-off time needs to be later than the pickup time; this is also known as a "precedence" constraint. Third, after insertion of the new trip, a check is performed to determine whether the insertion will violate the assigned customers' successive time windows. Finally, the capacity of each vehicle is also necessary to consider the proper process for insertion of the unassigned trips. Figure 4 illustrates the algorithm procedure in a diagram.

COMPUTATIONAL EXPERIMENT

To demonstrate the productivity and level of service provided by the proposed zoning-with-transfer paratransit system, the results of zoning without transfer and no zoning obtained by use of the same sequential insertion algorithm proposed in the previous section were compared. The real demand data provided by the Metropolitan Transit Authority of Harris County, Texas, which were used to generate the random samples, are presented. The configurations of three organizational structures are then described. Finally, an analysis of the simulation results is provided.

Demand Data Description

METROLift is a paratransit service in Harris County, Texas, currently in compliance with the ADA. On average, more than 5,000 trips are made daily from 3:45 a.m. to 1:30 a.m. on the following day. The fare for a single ticket is $1.15 per ride. All trips need to be scheduled 1 day in advance. Once customers make a reservation, the schedule operator gives them an estimated scheduled pickup time. The time can change plus or minus 20 min, which results in a 40-min time window. (Other U.S. cities have adopted 20- or 30-min time windows.) Comparisons with other systems are provided in Table 1.

Test samples were generated according to the locations (pickup and drop-off) and time distributions. The number of pickup and drop-off locations for every square mile was counted with geographic information system software (Figure 5). The actual pickup time distribution is shown in Figure 6. Because the pickup and drop-off locations were independently generated, the pickup and drop-off points were occasionally unrealistically generated within the same square mile. In these rare cases, new drop-off locations were generated.

Zoning Configurations

The configuration of a zoning structure is defined by its boundaries; transfer locations are often located at a zone boundary. The following

TABLE 1 Operating Characteristics and Populations Served by Different Systems

					Boarding Time (min)		Disembarking Time (min)	
City	Service Area (mi^2)	Service Area Population (millions)	Number of ADA Customers	Service Hours	Lift Required	No Lift Required	Lift Required	No Lift Required
Houston	751	3.2	17,695	3:45 a.m.–1:30 a.m.	6	1	4	1
Chicago	3,750	8	42,516	24 h	7	3	6	2
Boston	729	2.5	67,329	6 a.m.–1 a.m.	5	2	3	2
Washington, D.C.	1,500	3.4	25,575	5 a.m.–12 a.m.	7	2	6	2

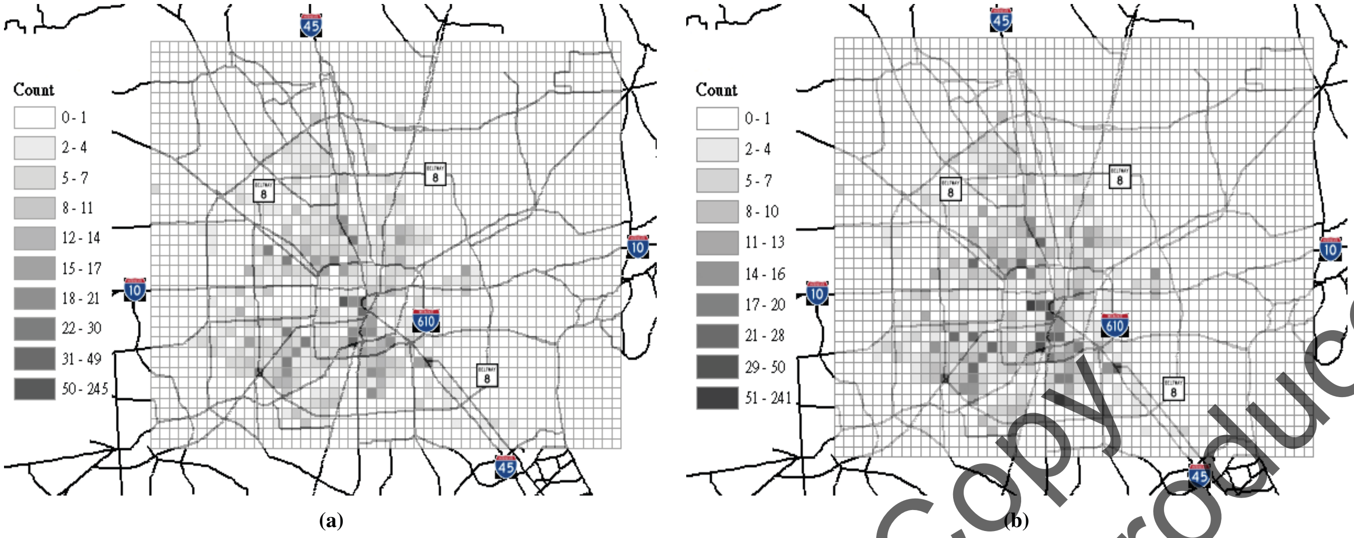

FIGURE 5 Distribution of (*a*) pickup locations and (*b*) drop-off locations.

four rules were used to build the subzones in Houston, as shown in Figure 7:

1. It is better not to situate a popular destination or an area with a high demand density in one exclusive zone.
2. Each zone should accommodate a certain number of trips originating from it.
3. The percentage and number of interzonal trips attached to each zone should be close.
4. Zones should be mutually adjacent so that more than one transfer can be avoided.

From a review of the distribution of pickup and drop-off locations, a 1-mi^2 area with an extremely high demand density (250 pickups per day) was located. This spot sits roughly in the lower center section of the service area. Both the origins of the trips leading to this spot and the destinations desired from this spot are scattered throughout the area. Therefore, this made an ideal center from which to form zones. If this spot was included in one specific zone, trips for other zones would have had to make more interzonal trips, which in turn would have decreased the overall service quality. On the basis of the selection of this spot, the service area was administratively divided into four geographical quadrants of unequal size: the northwest, the northeast, the southeast, and the southwest quadrants. Trips within each zone were observed to be large enough to maintain a minimum level of operational scale, although individual trips from each zone were not equal in length. In practice, passengers do not usually require transfers if their destinations are just one or two blocks beyond a particular zone boundary. Therefore, a 1-mi-wide buffer area was set along each zone boundary.

For the zones generated, five locations were sufficient to provide for all transfer needs. The center of the four quadrants was selected to be the transfer location for all interzonal trips traveling between the northwest and southeast or the northeast and southwest. The research found that transfer locations are best located at the edges of zones, nearest the major interzonal corridors.

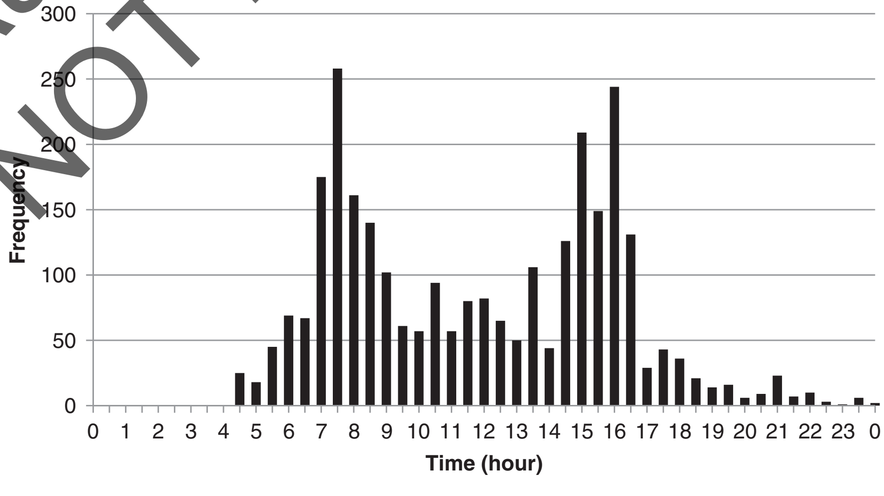

FIGURE 6 Distribution of requested pickup times.

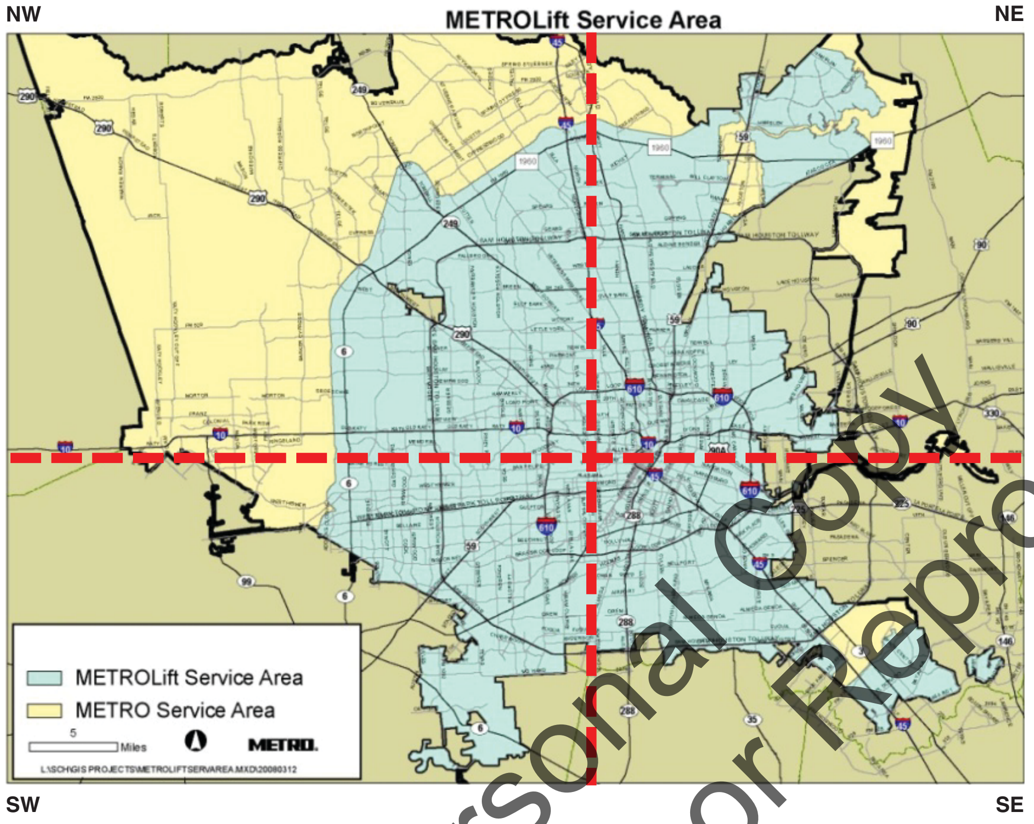

FIGURE 7 Zones built in Houston region.

The default parameters used in the simulation are as follows:

- Vehicle speed: 20 mph;
- Boarding time: ambulatory passenger = 1 min; wheelchair passenger = 6 min;
- Disembarkation time: ambulatory passenger = 1 min; wheelchair passenger = 4 min;
- Maximum ride time parameters: different parameters according to the direct travel distances of the customers; and
- Allowable deviation from estimated scheduled pickup time: 20 min plus or minus the scheduled pickup time.

The three scenarios listed below were tested on randomly generated instances, and 10 replications were run to deal with the randomness of the simulation.

1. Zoning without transfer (Scenario 1). The region was divided into four service zones, and each zone had its own carrier. Customers were zoned by their pickup locations and served by their designated service carrier. Vehicles in each zone could cross boundaries only to drop off interzonal customers.
2. Zoning with transfer (Scenario 2). The zoning-with-transfer scenario respected the same geographical zones and carrier design described for Scenario 1. Vehicles in this system, however, always remained within a single zone. Customers needed to transfer at zone boundaries.
3. No zoning (Scenario 3). The region was served by a single carrier. The current Houston paratransit service adheres to this scenario.

The statistics reported are the averages of 10 replications. The heuristic was implemented via the computer program C and was run on a computer with a 2.33-GHz Core2 Duo processor and 2 GB of memory.

Performance Measurements and Analysis of Results

The performance characteristics of the various scenarios were investigated according to system efficiency and service quality. For system efficiency, the number of vehicles used was the most straightforward indicator for a comparison of alternative scenarios. "Deadhead miles" were the number of miles that a vehicle traveled from its home depot to its first pickup node and from its last drop-off node to its home depot. "Vehicle revenue miles" were defined for all vehicles as the total number of miles traveled from the first pickup location to the last drop-off location. Vehicle revenue miles with no passengers on board were defined as "empty miles." "Total miles" included revenue miles and deadhead miles.

The number of passenger trips per vehicle revenue hour served as an important performance measure for capturing the productivity of a particular demand-responsive system. A higher number of passenger trips per vehicle hour usually means that more trips can be scheduled within a given time period.

The number of passenger miles traveled was calculated as the sum of the number of miles traveled multiplied by the number of customers on board for each travel segment. The number of passenger miles per vehicle revenue mile was another performance measurement used to calculate the productivity of the demand-responsive system. This measurement captured the difference in travel demand patterns between the systems that averaged longer or shorter trips. "Vehicle idle time" is the time gap between the vehicle arrival time and the earliest pickup time at the pickup location.

The service quality of the various different strategies was thoroughly analyzed according to several characteristics except efficiency. From the service quality point of view, deviation from the desired pickup time and passenger ride time were the major passenger concerns (besides the fare). Passenger wait time was calculated as the difference in time between the requested pickup time and the scheduled pickup time. Passenger ride time was the actual drop-off time minus the actual pickup time. Again, the passenger ride time could not exceed the maximum ride time factor for both intrazonal and interzonal requests.

The results generated by the three test scenarios are shown in Table 2. It was observed that the no-zone system had the smallest number of vehicles, whereas the zoning-with-transfer and zoning-without-transfer policies had larger numbers. This may be attributed to the following two reasons. The no-zoning system had no restrictions on the choice of the next unassigned trip; thus, the probability that a better insertion would be found was higher. In addition, in favor of the sequential insertion method, the number of trips in each of the routes created earlier was higher than that in the route created later. Therefore, if the route created later included only one or two interzonal trips, it could be served by one vehicle in a no-zoning system or in a zoning-without-transfer system. The route created later would have to be served by two vehicles in a zoning-with-transfer case.

When transfers were allowed in the zoning policy, the numbers of deadhead miles and empty miles were decreased compared with the numbers for the zoning-without-transfer policy. For the operator, a smaller number of empty miles is a better result because the number of passenger miles per vehicle revenue mile increases as the number of empty miles decreases. The zoning-with-transfer policy showed a significant improvement in the number of passenger miles over the numbers for both the no-zoning and zoning-without-transfer policies. The higher number of passenger miles could contribute to the longer travel time or the higher rideshare rate. Because the same data set was used to run each simulation, it was concluded that the zoning with transfer had a higher rideshare rate.

Although zoning limits the likelihood that a better insertion would be found, the results indicate that the transfer policy not only recovered the deficit from the no-transfer case but also improved that number compared with that for the no-zoning strategy. Because it had the highest number of passenger miles among the three cases, the zoning-with-transfer policy showed the highest number of passenger miles per total number of miles.

Zoning with transfer significantly improved the number of passenger trips per revenue hour, especially when it was compared with the number obtained in situations involving a zoning-without-transfer design. For interzonal customers, a zonal service that acts as a feeder and a distributor and that has a coordinated schedule and routes around a particular transfer point has increased productivity. This improvement is mainly due to the decrease in the number of empty miles from the last unloading point to the point where the new customers are scheduled for pickup. For interzonal trips with long travel distances, it was found that simultaneous two-way passenger exchanges at particular transfer points largely decrease empty backhaul miles.

Such a transfer policy would increase vehicle idle time, partially because of the vehicle's time spent idling at transfer points while it waits to pick up interzonal customers for a route created later. Analysis for the level of service showed that coordination at the transfer locations slightly increased passenger wait times compared with those experienced in a zoning-without-transfer system. However, passenger wait times were lower than those in the no-zoning case. The zoning-with-transfer policy showed the longest passenger ride times among the three scenarios. This was because passengers on interzonal trips had to switch vehicles at various transfer locations; thus, the system required some extra travel distances and additional wait times. However, it has been found that passengers can usually better endure longer travel times than longer wait times.

In general, the results show that the zoning-with-transfer design is suitable for a large service area in which the majority of trips are short and the number of long trips is determinable, similar to the demand pattern in Houston. Hence, productivity is retained with a focus on shorter trips within a denser area. For a small community, it is unlikely that longer trip lengths that contribute to an increase in empty trip miles will exist. As a result, transfer design will not be able to increase productivity and, instead, will only downgrade the service level. However, the extent to which the transfer design would benefit from an increase in the service area should be further investigated.

CONCLUSION

The effects of inclusion of transfers between service zones were examined in depth by use of the ADA paratransit system design. The results indicate that such systems can provide significant benefits to paratransit operations that are managed within a zoning structure. The results were obtained with the demand data for the paratransit system in Houston, Texas (a relatively low-density region), and it was concluded that the transfer method provides a productive

TABLE 2 Comparison of Performance for Three Zoning Scenarios

Scenario	Number of Vehicles	Number of Miles: Total Revenue	Number of Miles: Deadhead	Number of Miles: Empty	Number of Passenger Miles	Number of Passenger Miles/ Total Miles	Number of Passenger Trips/ Revenue Hour	Vehicle Idle Time (min)	Average Passenger Waiting Time (min)	Average Passenger Ride Time (min)
Zoning with transfer	379	49,828	4,668	10,043	64,026	1.17	1.60	77,041	22.5	44.3
Zoning without transfer	363	51,569	7,013	14,655	57,477	0.98	1.30	60,607	21.9	36.2
No zoning	295	43,289	8,337	7,824	57,376	1.11	1.53	51,438	23.4	36.2

organizational structure if the system operates under an obligation to maintain a zoning design. It was found that the transfer design described in this paper especially excelled at significantly enabling the zoning system to increase the number of passenger trips per revenue hour without an excessive increase in in-vehicle ride times for passengers but with maintenance of service within promised pickup and drop-off time windows. The no-zoning cases adopted by the Metropolitan Transit Authority of Harris County still perform better than zoning cases, on average, according to efficiency.

Furthermore, the comparisons of the simulations of the two zoning scenarios are generally considered to be indicative of their relative performances. Although the exact level of benefit will vary according to the different demand types and operational standards, this simulation methodology is easily and quickly adaptable to any large-scale paratransit system. Future work should combine searches for optimal transfer locations and optimal numbers of transfer locations to improve the performance of the transfer system proposed here. Finally, other means of improvement of this zonal transfer design for multiple paratransit operators would be a promising body for further study.

REFERENCES

1. KFH Group, Inc. *TCRP Report 124: Guidebook for Measuring, Assessing, and Improving Performance of Demand-Response Transportation.* Transportation Research Board of the National Academies, Washington, D.C., 2008.
2. Wilson, N. H. M., J. M. Sussman, L. A. Goodman, and B. T. Higonnet. Simulation of a Computer Aided Routing System (CARS). *Proc., Third Conference on Applications of Simulation*, 1970, pp. 171–183.
3. Chira-Chavala, T., and C. Venter. Cost and Productivity Impacts of a “Smart” Paratransit System. In *Transportation Research Record 1571,* TRB, National Research Council, Washington, D.C., 1997, pp. 81–87.
4. Ellis, E., and B. McCollom. *TCRP Report 136: Guidebook for Rural Demand-Response Transportation: Measuring, Assessing, and Improving Performance.* Transportation Research Board of the National Academies, Washington, D.C., 2009.
5. Lave, R., and R. Mathias. State of the Art of Paratransit. In *Transportation in the New Millennium* (CD-ROM), TRB, National Research Council, Washington, D.C., 2000.
6. Burkhardt, J. E., B. Hamby, and A. T. McGavock. *TCRP Report 6: Users’ Manual for Assessing Service-Delivery Systems for Rural Passenger Transportation.* TRB, National Research Council, Washington, D.C., 1995.
7. Quadrifoglio, L., M. M. Dessouky, and F. Ordonez. A Simulation Study of Demand Responsive Transit System Design. *Transportation Research Part A: Policy and Practice,* Vol. 42, No. 4, 2008, pp. 718–737.
8. Shen, C.-W., and L. Quadrifoglio. Centralize vs. Decentralize Zoning Strategies for Metropolitan Paratransit Systems. Presented at 89th Annual Meeting of the Transportation Research Board, Washington, D.C., 2010.
9. Cook, T. J., J. J. Lawrie, and A. J. Henry. From Rural Single-County to Multicounty Regional Transit Systems: Benefits of Consolidation. In *Transportation Research Record: Journal of the Transportation Research Board, No. 1841,* Transportation Research Board of the National Academies, Washington, D.C., 2003, pp. 54–61.
10. Cordeau, J. F., and G. Laporte. The Dial-a-Ride Problem: Models and Algorithms. *Annals of Operations Research,* Vol. 153, 2007, pp. 29–46.
11. Berbeglia, G., J. F. Cordeau, I. Gribkovskaia, and G. Laporte. Static Pickup and Delivery Problems: A Classification Scheme and Survey. *TOP,* Vol. 15, 2007, pp. 1–31.
12. Shang, J. S., and C. K. Cuff. Multicriteria Pickup and Delivery Problem with Transfer Opportunity. *Computers and Industrial Engineering,* Vol. 30, No. 4, 1996, pp. 631–645.
13. Cortes, C. E., M. Matamala, and C. Contardo. The Pickup and Delivery Problem with Transfers: Formulation and a Branch-and-Cut Solution Method. *European Journal of Operational Research,* Vol. 200, No. 3, 2010, pp. 711–724.
14. Solomon, M. M. Algorithms for the Vehicle Routing and Scheduling Problems with Time Window Constraints. *Operations Research,* Vol. 35, No. 2, 1987, pp. 254–265.

The Public Transportation Group peer-reviewed this paper.