automatically eliminated from the active order file or displayed with the appropriate vehicle number if a unit has already been sent. Other special and emergency calls are given order file priority. Also, as an option to the order sender, the system processes and displays cab status, which is manually input by the order sender. Vehicle status retains an automatic drop time and also is removed automatically from status screens when the vehicle number is used on an order. All disk access times have been accelerated with advance programming technology to ensure that operators are not "waiting" for the system to respond to a command.

Several business-oriented reports monitor the total communications operation and the individual performance of the operators. The computer hardware consists of 2 Data General Nova series minicomputers of 32K each, 2 dual disk driver units, 1 line printer, 1 teletype, 14 Hazeltime CRTs, and appropriate switching gear to enable the system to be fully backed up in case of computer hardware failure.

Some conclusions may be drawn from this operation:

1. The use of EDP equipment in dispatching demand-responsive vehicles is technically feasible;
2. It is economically feasible for an operation in which a minimum of 2,700 orders per day are handled;
3. It gives management greater flexibility in the utilization of personnel;
4. It improves service to the public; and
5. When the day of economically feasible AVM arrives, the circle of control of the historically independent taxicab driver will be more nearly complete.

Nigel H. M. Wilson, Massachusetts Institute of Technology

From 1967 to 1971, much research at M.I.T. and elsewhere was devoted to the potential use of computers in the control of demand-responsive transportation systems. Two of the most tangible outputs of these efforts were

1. A computer simulation model to test alternative computer control algorithms and to predict system performance; and
2. A recommended set of computer control procedures in which (a) the immediate assignment of each request was made to the current "tour" of the best vehicle, (b) the assignment was based on feasibility conditions, under which each user receives service within specified bounds, and (c) the determination of the best assignment was based on the minimization of total service times for current and future passengers.

These control procedures were tested by a simulation model and were found to perform well on intuitive grounds (i.e., an examination of individual assignments and their comparison with judgment) and relative to other proposed algorithms. However, since no optimal-solution algorithm has been developed, absolute statements about their performance were impossible.

One result of this research program was the decision to mount a demonstration project of the concept in Haddonfield, New Jersey, to obtain a market test of the service concept and to obtain data on the potential of computer dispatching. The system (which has been extensively described elsewhere) has just terminated; its demonstration project phase provided valuable data in both of these areas. In particular the computer control system used in the latter stages of Haddonfield was developed by the Mitre Corporation using the control algorithms previously developed at M.I.T.

M.I.T. is now the recipient of a university research and training grant from the Urban Mass Transportation Administration to develop advanced DRT control procedures based on the experience gained in Haddonfield and to look explicitly at the problem of controlling integrated DRT and fixed-route transit services. This presents a rare opportunity to evaluate academic research in light of subsequent operational experience and specifically to validate the simulation model and to analyze and improve on the operation of the total system. An additional benefit of the Haddonfield experiment has been the collection of extensive data on a similar manual system (the characteristics to the user are identical) that will permit evaluation of the quality of computer assignment. This paper presents preliminary results of this research and concentrates on the single DRT system.

ASSUMPTIONS IN DESIGN OF SIMULATION MODEL

Numerous assumptions and simplifications of the real-world system were required in the design of the simulation model. This model was designed to provide the analyst with the ability to simulate a wide range of systems. The input parameters include area dimensions, demand rate, demand pattern, number of vehicles, vehicle size, and vehicle speed. However, as the model was originally designed, 2 major assumptions warranted further investigation in light of Haddonfield operating experience:

1. A constant number of vehicles are in service continuously throughout the simulated period, and
2. The demand rate is constant during the simulated period, although the time between successive demands is selected from a user specified distribution.

To investigate the validity of these assumptions, 2 new options that relax these 2 simplifications have now been implemented in the model. The first option allows the analyst to use either completely random demand inputs or a fully specified set of demands that occur at known times between known origins and destinations. This allows the simulation of an actual set of demands from a day's operation at Haddonfield, for example. The second option allows vehicles to enter or leave service at any times specified by the analyst or to use a constant, continuous supply of vehicles. These options provide a great deal of flexibility and power in validating the simulation model. Simulation experiments were then run of the Haddonfield system; real and approximate demand and vehicle input were used.

FINDINGS ON MODELING ASSUMPTIONS

Comparing an actual demand stream simulation as obtained from Haddonfield transaction tapes with random demands based on approximations of the Haddonfield demand showed that approximate and random demands are quite satisfactory for the prediction of system performance. This implies that estimating the approximate spatial distribution of demand and level of demand is sufficient to predict future performance in a demand-responsive transportation system. This is fortunate, for if this assumption were not valid, prediction of future systems performance would have been infeasible.

However, the assumption of a constant and continuous supply of vehicles was found to result in significant overestimation of vehicle productivity or overestimation of the quality of service that can be provided or both. The reason for this is that, when a vehicle enters (leaves) service, it is significantly underused in the hour immediately following (preceding) the change. The greater the number of changes are in vehicle status, the greater the overall impact is; and, since fully demand-responsive operations occur in the base period of the schedule, vehicle status changes are frequent because of shift changes and driver lunch breaks.

To approximate Haddonfield results by using the basic unmodified simulation model with the constant number of vehicles equal to the average number of vehicles actually operating was impossible. However, by using actual vehicle in-service times, we were able to closely approximate actual Haddonfield quality of service. The operations from 9 a.m. to 3 p.m. on September 19, 1974 were as follows:

Item	Number
Passengers	262
Vehicle productivity (passengers/vehicle/hour)	5
Vehicles in service	10 to 12
Distinct vehicle shifts	34

The statistical analysis of actual and simulated quality of service is given in Table 1. The constant number of vehicles in continuous service demonstrates that similar service can be provided with about 30 percent fewer vehicles if they provide continuous service. The results from the third assumption reflect actual vehicle in-service times and show close correspondence with the actual operation.

The conclusion must be that, although the simulation model was sophisticated by any standard, it was not, as originally designed, realistic enough to provide reliable estimates of productivity and service quality. At the time the simulation model was developed, not enough was known about the transient behavior of the system to recognize this as a significant factor. The implications of this behavior are

1. The new model should be used in planning new systems in conjunction with expected (and realistic) vehicle in-service times (indeed the model can be an important factor in planning driver shifts), and
2. From a control procedure and operation viewpoint, more attention must be given to system performance under transient supply conditions.

ALGORITHM PERFORMANCE

In general the algorithm used in Haddonfield has performed well although no definitive comparison of the system performance with computer and manual assignments has yet been made. Preliminary evaluation indicates that the quality of service provided is at least as good under computer control as under manual decision making, and probably somewhat better; however, a fuller evaluation is now under way.

Based on operational experience in Haddonfield, the following are areas in which improved performance might be achieved:

1. Inflexibility of hard constraints,
2. Objective function as a true reflection of customer utility,
3. Handling of advanced and periodic requests,
4. Constraint of vehicle position at future time,
5. Restriction of certain vehicles to given zones,
6. Preassignment capability,
7. Scheduling at start and end of driver and vehicle shift, and
8. Gearing of algorithm to underused system.

Each of these areas is described, and, where appropriate, possible remedial actions are suggested.

Table 1. Statistical analysis of quality of service.

Operation	Assumption	Vehicles	Time	Mean	Standard Deviation	Max
Actual	–	10 to 12	Wait	9.5	6.0	34
			Ride	9.5	5.4	32
Simulated	1	Constant 8	Wait	6.7	5.0	22
			Ride	9.6	6.0	31
	2	Constant 7	Wait	8.9	6.8	31
			Ride	10.4	6.8	34
	3	In and out of service	Wait	7.4	6.5	34
			Ride	10.2	6.5	36

Inflexibility of Hard Constraints

The algorithm was designed to minimize total service time (for current and future passengers) within fixed constraints on wait, travel, and total service times.

Any assignment in which no constraint is violated is preferred to any assignment involving a constraint violation, independent of the value of the objective function. This constraint was developed to reduce the number of passengers experiencing unreasonably long service times; the effect of some increase in the mean service times was acknowledged and expected. To achieve this goal requires that the constraints be set about 50 percent above the mean service times. In practice, 2 problems arise with this approach.

1. Because the short-run demand rate varies widely during the course of the day and because mean service times are sensitive to the recent demand rate, a constraint set correctly for some time of the day may be incorrect for many other times of the day. The problem is that the constraints are not dynamically set as a function of the number of passengers currently on the system and the number of vehicles currently in service. This problem could be solved by using a short-memory heuristic to compute the current constraint set.
2. More basic is the problem that assignments that may be far superior from the objective function's viewpoint will be rejected if a constraint is violated. This introduces a perturbation in performance and can lead to short-sighted decisions that tend to waste system resources. This problem cannot be solved by any useful setting of the constraints, and its existence argues for a reduction in the role of constraints in future algorithm development work. This is possible only if the individual customer utility function can be equally or better represented by some other construct.

Objective Function as a True Reflection of Customer Utility

The objective function implies that users of the system associate with the service a utility function that is linear in service time. This seems to be an inaccurate and simplistic representation of actual passenger satisfaction, and hence its use can result in customer dissatisfaction. Although the actual utility function associated with DRT service has not yet been identified, clearly measures of the distribution of service time, other than the mean, are also important, e.g., standard deviation. It is also probable that the uncertainty in service is also an important characteristic. One measure of this is the difference between estimated and actual pickup and delivery times. Once again the means and standard deviations of these distributions should be considered.

That actual utility functions will vary not only from customer to customer but from area to area is highly likely. For these reasons, the next generation of algorithms must incorporate a richer mix of elements in the objective function and provide the user (operator) with ways to manipulate the objective function to achieve desired service characteristics. If the objective function is more realistic, the service constraints can then be used as a means to reduce computation (by eliminating unpromising assignments early) rather than as an integral part of the algorithm.

Handling of Advanced and Periodic Requests

At present, advance requests (this term will be used to include periodic requests) are assigned a fixed period before their desired pickup times and have a special set of (tight) constraints. A modified objective function that attempts to minimize the time between expected and desired pickup time is used. All subsequent assignments to a tour, including the advance request, are made as if the tour consists of only immediate service requests. This results in service for the advance request being no better than service for immediate requests, an unsatisfactory state of affairs, for advance requests should be easier to schedule and serve than immediate requests. This is an important area for future work.

Constraint of Vehicle Position at Future Time

The system was designed for the dynamic many-to-many case in which scheduled or repetitive demands or both on the system are not a major factor. In actual operation, vehicles may frequently have to make regularly scheduled or one-time appearances at specific locations, even though no originating service requests have been made (e.g., PATCO station in Haddonfield for scatter operations). This capability is an integral part of current algorithm development work at M.I.T.

Restriction of Certain Vehicles to Given Zones

For ease of use at high-density demand generators, specifying service zones is desirable so that passengers know immediately which vehicle serves their destinations—each vehicle can then post 1 or more zone numbers. For this operational technique to be compatible with computer dispatching, the computer system must be able to restrict a vehicle to serve only limited origin-destination pairs. This capability does not exist in the Haddonfield system, but recently M.I.T. implemented a scheme whereby vehicles can be restricted in terms of the origins and destinations served in the simulation model.

Preassignment Capability

The Haddonfield computer system does not have a passenger reassignment capability except when a vehicle breaks down, in which case the tour (including both collected and uncollected passengers) is shifted to the end of the vehicle that can first reach the breakdown point empty. Passenger reassignment as an element of the algorithm was investigated previously by M.I.T. and found to be of only marginal benefit. However, it may well be worth implementing specifically just for vehicles that break down and for vehicles that suffer large delays en route.

Scheduling at Start and End of Driver and Vehicle Shifts

As discussed previously, the computer should be able to efficiently build up tours and stop further assignments at specific times so as to maximize system productivity.

Gearing of Algorithm to Underused System

The previous algorithm development research was geared heavily to system (and hence algorithm) performance at or near the point of maximum system use. This resulted in higher vehicle productivities than typically observed in Haddonfield, and so the algorithm has been operating at much lower productivities than previously studied. As it turns out, both through observations in Haddonfield and through simulation experiments, the algorithm may not perform most effectively in this situation. Specifically the increase in tour length in the objective function can lead to significant imbalances in the use of vehicles; i.e., the probability is high that a new request will be assigned to an already highly used vehicle, and once a vehicle becomes unassigned it tends to remain so. The best objective function may well depend on the current use of the system.

CONCLUSIONS

The simulation model can accurately predict system performance providing that vehicle in-service times are used; otherwise, system performance can be significantly overestimated. With this caveat, the control algorithm performed as predicted by previous simulation modeling. However, much of the previous research and performance prediction was at significantly higher demand densities than have been observed at Haddonfield or most other demand-responsive systems. The implication of these lower demand densities is that the economies of scale possible with these systems cannot yet be realized—and that productivities of 5 to 8 passenger trips per vehicle hour are more realistic than previously cited ranges of 9 to 13. Stress must now be on making the service more attractive to potential users so that economies of scale can be achieved and at the same time increasing productivity for a given quality of service. With regard to integrating DRT and fixed-route transit, the computer must be used to make the overall service more attractive and to enable larger systems to be operated. Current research at M.I.T., which is addressing all these issues, strongly suggests that it is both feasible and desirable for the computer algorithms to achieve better service and to allow the operation of large integrated DRT and conventional transit systems.

Roy E. Murphy, Phillip L. Paisley, and John N. Siersema, LEX Systems, Inc.

Few people deny that one of the major problems today is the satisfaction of demand for an attractive, practical, economic alternative to the door-to-door transportation service offered by the automobile. Although much public and private money has been spent on the conveyance aspect of transportation, this expenditure has not brought us much closer to the development of an alternative to the automobile.

Many people think that the personalized transportation service offered by demand-responsive transportation technology provides this alternative to the automobile. If this is so, why has this new technology not been adopted by professional transit people to any great extent? The fact remains that most current DRT systems have serious defects for the practical transit operator.

DEFECTS IN DRT TECHNOLOGY

What are some of these defects? We suggest that too little attention has been paid to the economic efficiency of vehicle use in DRT applications. The current pressure to maintain high DRT service levels and the labor-intensive cost structure have reduced vehicle economic efficiency to such an extent that no conventional transit operating budget can long sustain such a DRT system.

The second defect in current DRT technology is its inability to provide practical DRT services to a large geographic area where, for example, door-to-door travel times could be as long as 2 hours. Another aspect of this defect is the current lack of DRT technology to truly integrate with express bus or rail transit facilities in a large area.

The third weakness in DRT technology is the poor accuracy of current scheduling methods. Given fixed resources, promised response times grow less and less reliable as demand increases. This fault is not so much due to the inability of current scheduling methods to cope with DRT demands as to the lack of scheduling tools that can assist in carrying out the methods while keeping up with the demands. Therefore, the scheduling of increasing numbers of vehicles or passengers or both, plus the introduction of other complexities such as the integration of DRT and other forms of transit, is hard to imagine without some automated scheduling assistance.

AUTOMATED SCHEDULING ASSISTANCE

To assist the scheduling (and dispatching) functions of DRT control and to help overcome the defects, LEX has developed various levels of automated control system