9. A.K. Sen and E.K. Morlok. On Estimating Effective Frequencies and Average Waiting Times for Indirect Connections. Transportation Planning and Technology, Vol. 3, No. 3, 1976, pp. 175-184.
10. L.D. Bodin and R.B. Dial. Hierarchical Procedures for Determining Vehicle and Crew Requirements for Mass Transit Systems. TRB, Transportation Research Record 746, 1980, pp. 58-64.
11. L.R. Ford, Jr., and D.R. Fulkerson. Flows in Networks. Princeton Univ. Press, Princeton, NJ, 1962.
12. A.B. Sosslau and others. Quick-Response Urban Travel Estimation Techniques and Transferable Parameters, User's Guide. NCHRP, Rept. 187, 1978.
13. R.H. Pratt; DTM, Inc. Development and Calibration of Mode Choice Models for the Twin Cities Area. Metropolitan Council, Minneapolis, 1976.
14. G.J. Fielding, R.E. Glauthier, and C.A. Lave. Applying Performance Indicators in Transit Management. Presented at the National Conference on Transit Performance, Norfolk, VA, Sept. 1977.

Transportation, Final Rept., in preparation.
5. J.J. Bakker. Transit Operating Strategies and Levels of Service. TRB, Transportation Research Record 606, 1976, pp. 1-5.
6. M.A. Turnquist and L.A. Bowman. Control of Service Reliability in Urban Networks. U.S. Department of Transportation, Final Rept. DOT-OS-80018, 1979.
7. W.C. Jordan and M.A. Turnquist. Zone Scheduling of Bus Routes to Improve Service Reliability. Transportation Science, Vol. 13, No. 3, 1979, pp. 242-268.
8. L.A. Bowman and M.A. Turnquist. Service Frequency, Schedule Reliability, and Passenger Wait Time at Transit Stops. Transportation Research, Vol. 15A, No. 6, 1981, pp. 465-472.

Publication of this paper sponsored by Committee on Bus Transit Systems.

State of the Art of Current Bus Transfer Practices

MICHAEL NELSON, DANIEL BRAND, AND MICHAEL MANDEL

The use and consequences of alternative bus transfer policies are examined. A bus transfer policy consists of a set of operator actions that involves vehicle routing and scheduling, transfer charges, information for passengers, and terminal facilities that affects the movement of passengers between buses as part of a continuing trip. In this paper, the bus transfer policies currently in use on U.S. transit properties are described and summarized. Reasons why properties use or do not use particular transfer policies are identified, and the specific consequences of alternative transfer policies in different settings on cost, ridership, revenue, and user satisfaction are assessed. Situations or settings in which particular transfer policies can be applied beneficially are then identified and analyzed.

Abstract goes here

Under ideal circumstances, transit would carry all users directly from their orgins to their destinations without requiring a change of vehicles. However, given the geographic and temporal distribution of trips, such direct service is of course uneconomical for transit to provide. Therefore, operators must undertake some set of actions that involves such factors as vehicle routing, scheduling, transfer charges, and/or information for passengers (a transfer policy) to accommodate transferring riders.

This paper examines the use and impacts of the following 11 bus transfer-policy components, which are listed under four main components:

1. Routing components--distance between routes at transfer points and through-routing;
2. Scheduling components--schedule coordination, dynamic control of departure times at transfer points, timed transfers, schedule adherence on connecting routes, and service frequency on connecting routes;
3. Pricing components--transfer charge and use of transfer slips; and
4. Information components--provision of schedule information and marketing initiatives.

Note that the 11 transfer-policy components examined here do not exhaust the list of possible operator actions that affect transfers. However, most of the remaining ones (such as transit shelters, terminal facilities, and temporal or directional restrictions) are reviewed at least briefly in conjunction with one or more of the above components.

The material presented in this paper is drawn from the results of a recently completed study conducted under the Service and Methods Demonstration program of the Urban Mass Transportation Administration (UMTA) for the Transportation Systems Center (1,2). Data for that study were drawn primarily from a series of telephone and on-site discussions with experienced transit professionals on 39 different properties.

On any particular transit property, the demand for transferring clearly influences the type of transfer policy adopted. Relevant transfer demand characteristics include the following:

1. The percentage of riders who transfer (i.e., transfer rate),
2. Their socioeconomic and trip purpose characteristics,
3. Transfer-point locations, and
4. Directional and temporal characteristics.

The transfer rate is the percentage of transit person trips that involve transfers between transit vehicles. Often, the transfer rate cannot be calculated directly from available data but rather must be estimated from transfer slip data, passenger counts, or special surveys. Data problems include transit pass users who do not use transfer slips or riders who transfer more than once in the course of a trip. In general, however, it is possible to obtain reasonable estimates of transfer rates on most properties.

For bus-to-bus transfers, the average transfer rate on the properties examined is approximately 21 percent. However, several bus properties have a transfer rate on the order of 5 percent, while transfer rates as high as 50 percent have been observed. The size of the property has a large effect on the overall transfer rate. Not including properties that currently use timed transfers extensively, large bus properties have a much higher average transfer rate than small properties (20 versus 12 percent). Bus properties that currently use timed transfers extensively are uniformly small properties and have a much higher transfer rate (28 percent) than either large or small non-timed-transfer cities. Also, bus properties that do not charge for transfers have a higher average transfer rate than those that do charge (22 versus 18 percent). In many of these cases, the causal relation is not clear. That is, rather than the action causing a higher transfer rate, the presence of a high transfer rate may cause a property to institute options such as timed transfers or no transfer charge.

Riders who transfer vary by socioeconomic and demographic groups. Low-income riders transfer more often than higher-income riders, and young people also have above average transfer rates. Elderly people, on the other hand, tend to have a lower transfer rate than other riders, perhaps because of the effort involved in changing vehicles.

In the following sections, the current practices of transit operators regarding each of the 11 transfer-policy components listed above are described. Reasons offered by operators for the use or lack of use of the policy component are examined, and the cost, user satisfaction, ridership, and revenue consequences of the component when used in different settings are provided. In this way, the types of properties and settings for which the transfer-policy component is most beneficial are identified, along with other components that work well when combined with the component in question.

TRANSFER-POLICY OPTIONS

Distance between Routes at Transfer Points

A basic attribute of transferring is the walk required between vehicles. There may be only a few feet or, alternatively, passengers may have to walk several blocks to transfer. The greater the distance, the less useful the transfer is for the passenger. A significant number of bus properties separate by 500 ft or more some routes between which transfers are expected to occur. On at least one property, passengers must walk up to 1500 ft to transfer. In contrast, an ideal transfer arrangement is one where the transfer walk distance is less than one or two blocks, with clear lines of sight between buses.

There are two major factors that determine how closely a given transfer point can approach this ideal situation--the number of vehicles that meet at the transfer point and the size of the central business district (CBD). Available information strongly suggests that the upper bound on the number of buses than can be present simultaneously in the area that surrounds a single bus transfer point is approximately 20. Above this number, even if all buses formally meet in the same area, there will necessarily be a significant transfer walk distance and obstructed lines of sight between at least some pairs of routes.

The size of the CBD is important because line-haul transit often serves as a downtown distributor if the CBD is large enough so that no single terminal area is within walking distance of all of it. Concentrating the termination points of all routes in one spot in such a CBD may cut down overall coverage and greatly increase transferring unless costly detours are made. Therefore, except where the layout of the CBD is well suited to single-point termination, transit systems with large CBDs cannot generally obtain the most desirable walking distance between all vehicles at a single transfer point.

It is therefore often necessary to consider alternatives to the most desirable transfer arrangement. Specific alternatives available to the operator that relate to spatial separation at transfer points include the following:

1. Building an off-street terminal facility. A terminal facility increases the number of buses that can meet at one point and reduces pedestrian obstacles. However, it may be necessary to use a noncentral location, and there may be significant capital costs. This option is best implemented when there are non-transfer-related benefits as well (e.g., reduction of street congestion).
2. Establishing a bus transit mall or street. With a bus transit mall or street, the number of routes is not a constraint and transfer walk distance is effectively very low, since vehicles from different routes pass the same points. Depending on the shape of the CBD, transit malls may reduce CBD coverage and require more walking and more transfers. This option is most feasible when the CBD is narrow (e.g., four blocks), and can often be implemented without significant capital costs.
3. Collecting route termini into several subfoci such that all or most routes intersect, although not all at the same point. With the option of collecting route termini, the overall number of routes is not a limitation, transfers between routes within subfoci are easy, and the CBD is well covered. Operating costs increase due to the extra vehicle miles of travel (VMT) in the congested downtown area. These costs tend to restrict the use of subfoci in larger cities.
4. Laying out routes in a grid network. With a grid system, the distance between routes at transfer points is minimized. Unfortunately, dispersed transfer points may be less understandable and less safe in the evening than single or multiple route foci. This option is most often employed in cities with large central areas of high-density employment and population, where a grid system may be the only appropriate route structure.

The principal cost consequences of any of these strategies typically arise from the changes in bus VMT needed to move the routes closer together. On the demand side, key aspects of transfer distance are walk time, comprehensibility of the transfer system, and potential pedestrian obstacles. These factors are particularly important to the elderly, shoppers, and infrequent users.

One additional important factor to consider when examining the trade-offs involved in reducing spatial separation is that some transfer-policy options discussed below, such a through-routing, schedule coordination, and timed transfers, require or are greatly facilitated by the physical proximity of connecting vehicles. In these cases, the trade-offs connected with spatial separation cannot be treated independently of the trade-offs connected with consideration of the other options. Thus, although reducing spatial separation has its costs, it may also have benefits that go beyond those of the single transfer-policy component standing alone.

Through-Routing

Through-routing, also known as interlining, involves linking two routes so that the same vehicle travels on both routes. It eliminates transfers between the two routes, since a passenger can board a vehicle at a stop on one route and get off at a stop on the other without having to change vehicles. A number of U.S. transit properties use through-routing; some properties use it quite extensively.

Five types of bus through-routing are currently in use:

1. Interlining, or classic through-routing, is when two separately identified routes share the same vehicles;
2. Single route through-routing differs from classic through-routing only in that the two halves of the route are joined on a permanent basis and are formally treated as a single route;
3. Variable through-routing differs from classic through-routing in that buses are exchanged among multiple routes rather than just between pairs of routes;
4. Trippers are when buses are through-routed at particular times of the day, usually during the rush hour or to meet shift or school times; and
5. Overlap involves terminating a radial route on the opposite side of the CBD from which it came in.

Through-routing can be used for two distinctly different reasons: operations and ridership. Both types of through-routing are considered in detail below.

Through-Routing for Operations

Through-routing can produce significant cost savings through elimination of turnaround time and distance, opportunities for logical scheduling, and potential gains in service reliability (if layovers are preserved). Although headway matching may add costs and extra scheduling effort may be needed, the net effect of through-routing is generally to reduce costs. These cost savings are most likely to occur in cities with a congested CBD where routes enter from more than one direction. Through-routing is most applicable as an aid to logical scheduling when properties are constrained by service-area boundaries or when operators seek to maintain clock-face or pulse scheduling. The presence of clock-face or pulse scheduling will also tend to minimize the need for further headway matching to implement through-routing, thereby avoiding potentially adverse cost impacts. Nevertheless, it must be emphasized that the operational and cost consequences of through-routing are heavily dependent on the street layout and other conditions. For instance, on some properties the dominant reason for implementing through-routing might be the elimination of dangerous left turns.

Through-Routing for Ridership

Through-routing eliminates transfers, which thus eliminates waiting and walking time and produces significant benefits for riders. Through-routing for ridership is often profitably employed where there is a high volume of transferring passengers between two routes with a common terminus. For instance, properties with strong and definite flows to outlying shopping malls may want to interline the mall route with a route running through a densely populated residential area. The groups that tend to benefit from this would be shoppers and the elderly--groups whose user satisfaction is most increased by through-routing. Properties with periodic peak flows to particular points, on the other hand, might profitably run trippers. If there is a relatively dispersed flow of transferring passengers, variable through-routing is a possible option. This will principally benefit the elderly or others who are made aware of and can afford to wait until a particular time of day for service. Properties that have a large amount of transferring to reach destinations within the CBD may consider overlap.

A reasonable range for the increase in ridership resulting from the connection of a pair of routes that serve logical origins and destinations is between 4 and 7 percent of the original ridership on the two routes. Conversely, pairing routes that do not connect logical origins and destinations may not increase ridership at all. Overall, through-routing for ridership is not necessarily incompatible with through-routing for operational reasons, although passengers who transfer between the two routes may be only a small portion of the total ridership on the routes. However, route pairings for maximum user satisfaction may not be the same as route pairings for maximum cost savings and operations benefit.

Schedule Coordination

Schedule coordination involves the adjustment of schedules on routes to change the relative times of arrival of vehicles at transfer points to reduce average transfer wait time. Schedule coordination used alone benefits passengers who transfer in one direction more than passengers who transfer in the other, since to ensure that one vehicle arrives before another without disruption of the regular schedule, an offset must be used (rather than attempting to have two vehicles meet at exactly the same time). Hence, schedule coordination used alone is applied most beneficially to route pairs where the majority of transfers are in a single direction at any one time.

Schedule coordination generally takes one of three forms:

1. CBD schedule coordination is used in situations where there is a strong directional flow of transfers through the CBD during peak hours. This option can improve the level of service for transferring passengers, although it requires some scheduling effort and sometimes changes in headways.
2. Trunk-crosstown coordination involves transfers between trunk lines and crosstowns where the schedule coordination is imposed outbound on the low-frequency crosstown lines. This option generally costs little and has minor negative effects on people who transfer in the opposite direction due to the generally high level of trunk-line frequencies. Trunk-crosstown coordination is more widely applicable than CBD schedule coordination, since its effects are less sensitive to the directionality of the transfer flow.
3. Minor schedule coordination characteristically is implemented in response to the complaints of passengers on a particular run who are unable to make a connection to another route. This option is easy to implement, although it typically does not lead to large ridership gains. It can be implemented on any transit system, even the largest and most complex, at any time of day.

If transfers are strongly directional between two lines at any time of the day and if a reasonable degree of schedule reliability exists, schedule coordination may be a very productive action for the operator to undertake. It typically costs little, can involve only minor headway changes, and demands almost no real-time operator attention. Because user satisfaction for the (assumed) large proportion of people who transfer in the correct direction is increased as their average transfer time is decreased, ridership will be induced in most cases. (Overall ridership gains are likely to be on the order of 3-4 percent for CBD schedule coordination, 1-2 percent for trunk-crosstown coordination, and minimal for minor schedule coordination.) Schedule coordination therefore may be a very cost-effective way to improve service.

However, there are definite limitations on the opportunities for application of most types of schedule coordination. The major restriction is the need for strong directionality of transfers. People who transfer in the "wrong" direction will have a transfer wait time equal to the entire headway (minus the advance) of their connecting bus. From the point of view of ridership, equity, and public relations, this may be unacceptable if a sizable number of people are affected. The result is that schedule coordination is inapplicable in many situations. For instance, it is largely inappropriate for off-peak use, since shopping traffic is inherently two-way. More important, it cannot be used in the CBD unless there is a skewed distribution of origins and destinations by time of day. Because this is a condition that is much more likely to occur in small cities than large ones, city size is a determinant of the applicability of at least CBD schedule coordination.

Dynamic Control of Departure Times at Transfer Points

Dynamic control involves holding a vehicle beyond its scheduled departure time from a transfer point if it is known that a vehicle on another route is approaching that is likely to have transferring passengers on board. Such information can be conveyed by radio or by some other signaling device (e.g., headlights). Dynamic control of bus transfers is found in many different settings. Several small properties use it extensively, either to control meetings between trunk and crosstown routes or to facilitate transfers in CBDs where the schedule permits. Some larger properties use dynamic control marginally, on only a few routes or only in the evening. However, on smaller properties that use timed transfers, dynamic control is regularly used to ensure the meeting of buses at the transfer point.

By definition, dynamic control perturbs the schedule. On a simple system this disturbance may not have widespread effects. On a more complex network of routes, use of extensive dynamic control may produce harmful schedule disruptions. There is also a limit on the number of dynamic-control messages that a radio system can handle. The major constraints on the use of dynamic control thus tend to be the size and complexity of the system.

This does not mean that larger properties cannot use dynamic control. It does mean, however, that it should be used sparingly and substituted for as appropriate, particularly on larger properties. For example, if dynamic control is used regularly at a particular transfer point, then an option such as schedule coordination might be more appropriate than a regular real-time adjustment in operations.

There are many situations where dynamic control is a low-cost method for obtaining large gains in user satisfaction for some riders and for improving overall public relations. Dynamic control is applicable whenever two low-frequency routes intersect, and it is productive to guarantee that transferring passengers will make their bus. Dynamic control is also applicable in cases where a low-frequency route receives a significant volume of transferring passengers from a higher-frequency route. By holding the vehicle on the low-frequency route to ensure that it meets an approaching vehicle on the other route, wait time is reduced.

More generally, dynamic control as a separate option is appropriate either when transfer flows are intermittent or when schedule unreliability is common. In the first case, dynamic control provides a way of making adjustments in operations only when they are needed to accommodate transferring passengers and is thus a substitute for schedule coordination. In the second case, dynamic control can cause buses that would not have met otherwise to meet, thus mitigating the effects of schedule unreliability on transferring passengers. This is particularly important on timed-transfer properties, where a guarantee that buses that will meet is necessary to attract new riders and ensure the satisfaction of old riders. Dynamic control with timed transfers may require some additional layover time, although not as much as if layovers alone were used to overcome reliability problems. In this situation, dynamic control is generally a workable compromise on both cost and user-satisfaction grounds between no alleviation of schedule uncertainty and the addition of layover time sufficient to absorb all schedule variance.

Timed Transfers

A timed transfer is defined as a set of operator actions that provides some degree of certainty that vehicles on different routes will meet at regular intervals to exchange transferring passengers. Timed transfers for buses can be divided into four distinct types:

1. Simple timed transfers: Two routes are scheduled and operated to guarantee that some or all buses on the routes will meet at the transfer point;
2. Pulse scheduling: Buses on all (or most) routes that meet at the major transfer point are scheduled to arrive nearly simultaneously, hold until all buses have come in, and then leave together;
3. Line-ups: In larger cities, buses on all (or most) routes that meet at the major transfer point are scheduled to allow a pulse-type exchange of passengers, typically in the context of low service frequencies and possibly long layovers at the transfer point, and most often in the evening; and
4. Neighborhood pulse: The schedules of neighborhood circulator routes are coordinated to make travel within a sector of a city easier.

Simple timed transfers are used on many properties, from the smallest to the largest, and are most commonly employed in the evening when both routes have low frequencies. Pulse scheduling is currently found in smaller cities (service-area population up to 300 000) and is used all day at central transfer points with a normal pulse frequency of approximately 30 min. Line-ups are used by many larger, nonpulse properties (service-area population of 500 000 or more), typically with headways of 1 h. Neighborhood pulse is currently being implemented on at least two large properties (Portland, Oregon, and Denver, Colorado) as part of their conversion to the bus transit-center concept.

In general, implementation of timed transfers requires several operator actions. Headways on different routes must be synchronized by altering route length and/or modifying layovers. Extra layover times may be needed to improve schedule adherence. This may also be accomplished by providing dynamic control of buses at the transfer point in case any are late. The operator must provide suitable space and facilities to permit easy simultaneous interchange of passengers between buses and must make important decisions concerning user information.

Property size is the principal criterion for timed-transfer applicability, serving as a proxy for headway reliability, service frequency, and the number of buses meeting at one time. Properties with less than 400 000 people in their service area are generally able to use pulse scheduling at their main transfer point. Larger properties often have line-ups at night. Simple timed transfers are usable on any property, while neighborhood pulse is applicable on any system with subcenters that serve as logical pulse points.

Service frequency and reliability appear to be the major determinants of whether user satisfaction is greatly increased by timed transfers. Ridership gains on the order of 5-12 percent appear reasonable with a highly reliable pulse-type timed transfer. Because increasing reliability generally costs money, the operator can implement timed transfers in different ways, depending on local objectives (e.g., increasing layover times and shortening routes versus adding equipment). Overall, however, timed transfers appear to be a cost-effective way of increasing service and ridership under certain circumstances without necessarily increasing costs (3).

Schedule Adherence on Connecting Routes

Schedule adherence is an important aspect of the overall level of service on transit properties that affects all (transferring and nontransferring) riders on the system. Major causes of bus schedule-adherence problems include traffic congestion, bunching, and, somewhat surprisingly, interference from trains and breakdowns of new buses. Remedies include skip-stopping, use of electronic and manual monitoring systems to control bus bunching, passing the first bus, and insertion of extra buses.

The principal consequences of increased schedule unreliability are an increase in the variance in transfer wait time and the expected (average) transfer wait time. These, of course, lead to decreases in user satisfaction, ridership, and revenue. However, the direct tranfer-related consequences of schedule unreliability are usually dominated by the nontransfer effects. Nevertheless, there are indirect transfer-related benefits that are nontrivial. If unreliability is too high, other operator actions regarding transfers (e.g., pulse scheduling or schedule coordination) are likely to meet limited success. Because the consequences of these other transfer-policy options can be very significant, minimizing unreliability for the purpose of aiding transfers can be an important objective.

Service Frequency on Connecting Routes

Service frequency, like schedule adherence, is an important component of transit level of service that has consequences beyond its impact on transfers. Given good schedule adherence, increasing the frequency of service on a connecting route should decrease the transfer wait time. Typically, however, operators raise or lower service frequency in response to non-transfer-related factors. The exceptions to this rule arise when other transfer components such as timed transfers, through-routing, and schedule coordination are implemented, since headways must be synchronized between routes. Even in these cases, however, the headway adjustments currently made are usually not large.

In general, user satisfaction, ridership, and revenue will rise due to the reduction in transfer wait time associated with service-frequency increases. Furthermore, there is a threshold headway of 10-15 min below which other transfer-related actions regarding scheduling may not be worthwhile due to limits on how great a reduction in transfer wait time they can produce. However, because of the significant cost of increasing service frequency, large changes in service frequency are typically only made in response to changes in overall demand or other factors and not simply transfer demand.

Transfer Charge

The transfer charge is the amount of money, over and above the basic fare, that a passenger pays to transfer to a second bus. Most bus properties currently charge nothing or $0.05 for a transfer. Other transfer charge levels are comparatively rare. Very few properties have full-fare transfers.

There is no consistent trend in transfer charges over recent years. Some properties have raised the charge slightly, e.g., from $0.02 to $0.05. Other properties have made transfers free. On average, nominal (and certainly real) transfer charges have tended to drift downward, although this tendency is neither pronounced nor universal.

A variety of reasons exists for setting the transfer charge at a particular level. In approximate order of importance, these include the following:

1. Historical precedent: Many properties have not recently given serious consideration to the level of their transfer charge;
2. Transfer abuse: A nonzero transfer charge may reduce the resale of transfers at a price below that of a full fare or the giving away of the transfers to friends and relatives, since fewer people would take transfers with the intent of later distribution if they had to pay something for them; and
3. Political or equity considerations: A particular transfer charge may be justified on the basis of a desire not to penalize transfers, not to subsidize long trips, etc.

Revenue, public relations, bus running times, and other considerations may also affect the selection of the transfer charge.

The cost consequences of a particular transfer charge result from the possible slowdown in bus passenger entrance and the minor cost of counting and handling additional revenue. The cost of slowing down the bus to process the transfer charge may be significant and results from both the need to pay a charge and from disputes that may develop between drivers and passengers over transfer abuse.

User satisfaction is decreased as the transfer charge goes up. However, the magnitude of this effect is determined by the disutility associated with charges by different user groups and the justifiability of the charge (i.e., the feeling among riders that the charge is fair and has a purpose, such as to make longer trips cost more). Both total bus ridership and the bus-to-bus transfer rate are sensitive to the level of transfer charge, although different types of trips will be affected differently by a given change in transfer charge. Captive riders have their riding patterns altered least by an increase in transfer charge, while shopping and other discretionary trips would be most discouraged. These ridership changes translate directly into effects on revenue. Because of the generally inelastic demand for transit, total revenue will typically increase (sometimes by a substantial amount) as the transfer charge goes up.

Clearly, operator goals and policies play a major role in determining the best transfer charge in a particular setting. For example, maintaining a low base fare to encourage total ridership may call for relatively high transfer charges for revenue reasons. A large deficit may also necessitate raising transfer charges to raise more revenue. Of course, transferring and overall ridership may be discouraged by a high transfer charge. Furthermore, it seems unlikely that a full transfer charge could be imposed on a system that currently has a high transfer rate without substantial political or equity problems. Overall, each of the three levels of transfer charge--zero, small but not nonzero, and full--seems to be stable and viable. The selection of one over the others is based on the operator's priorities and various other site-specific factors.

Transfer Slips

Transfer slips are the principal method for offering reduced-fare bus transfers that entitle riders to board subsequent vehicles at a reduced fare. Most properties use conventional transfer slips, but some use daily (or longer-term) passes or even no transfer slips at all to grant reduced-fare transfers. (No transfer slips at all is relatively rare but can be used when a small number of buses meet on a pulse schedule or when there is a restricted-access facility at the single transfer point on a system.)

The cost of administration of transfer slips is low. User satisfaction, ridership, and revenue consequences of transfer slips follow primarily from their use in setting fare policy and not from any characteristic intrinsic to the transfer slips themselves.

Schedule Information

Schedule information useful for transferring can be provided either at the transfer point or prior to the start of the transit trip. At the transfer point, transit properties can supply or post printed schedules or maps and/or disseminate information about whether the connecting vehicle is late. Prior to the trip, sources of information include printed schedules (which may include information on transfer points, time points, and best connecting vehicles) and telephone information systems. Schedules can also provide information on the other components of the transfer policy (e.g., through-routing, dynamic control, schedule coordination, timed transfers). In general, most properties only indicate the transfer charge and procedure for transferring on their schedule and almost never indicate the use of dynamic control or schedule coordination.

The direct costs of providing schedule information include printing schedules, manning phone banks, etc. The indirect costs are a type of opportunity cost, which occurs when a operator publicly states a transfer policy and then feels committed to it even when it becomes unproductive to do so. Provision of schedule information at the transfer point may raise the satisfaction of the rider by reducing uncertainty and by creating opportunities for other productive activities (which may be equivalent to a significant reduction in transfer wait time). Awareness of schedule information prior to the start of a trip will raise user satisfaction by enabling the passenger to make beneficial changes in trip-making behavior.

If the schedule and routes of a transit system never changed, provision of schedule information of all types would almost universally be the preferred action. However, transit routes and schedules frequently change, typically requiring an information or cost trade-off to be made. Each operator must determine whether the benefits produced by providing information are offset by direct costs and the need to make periodic adjustments in schedules and routes. This trade-off may be considered separately for each aspect of the transfer system that might be publicized. Route structure, for instance, is usually more stable than the schedule, so listing transfer points in a schedule poses fewer potential problems than listing the schedules of connecting routes. Also, the best connecting vehicle may change as a result of a small change in the schedule. Therefore, few properties put best-connecting-vehicle information in their printed schedules. Each component of the transfer policy may or may not be ruled out by the need to make periodic adjustments in schedules and routes on any particular property.

Marketing

Transfer-related marketing initiatives by the transit operator can focus completely on transfers, be part of a broader marketing effort, or use transfers to market other aspects of the transit system. When the transfer policy has some important distinguishing feature that can potentially affect a significant number of riders (e.g., pulse scheduling or universal transfer valid between carriers), it can be the focal point of a marketing effort. In addition, it is possible to market responsiveness by such means as a limited use of schedule coordination in which individual runs are adjusted to promote user satisfaction in response to user complaints.

Many properties promote transfers as part of broader marketing efforts. For instance, properties often produce brochures that describe their special services, including brochures on how to transfer. Transit fare-prepayment plans are another important example of the marketing of transfers as part of a larger effort, since the transferring rider usually receives free transfers. Transfer slips can also be used as part of marketing efforts that have nothing to do with transfers, such as the use of transfer slips as daily passes or special promotions in which retail establishments offer their customers return fares in exchange for transfer slips.

The cost of transfer-related marketing can be low, especially if there are few transfer points. The user-satisfaction consequences of marketing are related to the changes it can cause in awareness of and attitudes toward transit. Transfer-related marketing can change people's perceptions that transfers are onerous by promoting aspects of the transfer policy that make transfers easier. Marketing has been used to raise the awareness of different market segments concerning the existence of various transfer-policy components, as well as coverage and services provided by the system as a whole. It remains uncertain, however, whether marketing directed specifically toward transfers on properties whose transfer system has no special attributes is appropriate or productive.

CONCLUSIONS

Each of the 11 transfer-policy components described can be cost effective in various situations. The decision to use a particular policy component must address the trade-offs among that component's various consequences. For instance, the introduction of a transfer fee on a property where there were previously free transfers involves balancing equity, revenue, and user-satisfaction considerations. Trade-offs of this type are important from the point of view of the operator in determining how well a particular transfer policy meets the goals and objectives of the system.

It is important to reiterate that combinations of transfer-policy components may produce consequences that are not simply additive. For example, instituting timed transfers will, in general, have a positive effect on user satisfaction, as will increasing schedule reliability. However, the magnitude of the consequences of timed transfers will usually depend on the reliability of the connection, which would be enhanced by increasing schedule reliability. Hence, the increases in user satisfaction caused by implementing timed transfers and increasing schedule reliability may exceed the sum of the benefits derived from using those two components individually. Furthermore, some options have more widespread applicability than others; through-routing, for instance, can probably be implemented on a wider range of property types than pulse scheduling, although pulse scheduling has more far-reaching effects. Each operator must evaluate the service, cost, and demand conditions on the property and the consequences of alternative policies to determine which actions would be the most productive.

ACKNOWLEDGMENT

The material presented in this paper is drawn from the results of a recently completed study funded by the Transportation Systems Center (TSC) under the Service and Methods Demonstration program, UMTA. We would like to thank Robert Casey of TSC and Thomas Parody of Charles River Associates for their many helpful comments and observations. Also, the contributions of the transit professionals who participated in this study are gratefully acknowledged. However, the opinions and conclusions expressed or implied in this paper are ours, and we are solely responsible for any errors of content or ommission that remain.

REFERENCES

1. Charles River Associates. State of the Art of Current Practices for Transit Transfers. Transportation Systems Center, U.S. Department of Transportation, Cambridge, MA, June 1980.
2. Charles River Associates. Transit Operator Guidelines for Transfer Policy Design. Transportation Systems Center, U.S. Department of Transportation, Cambridge, MA, June 1980.
3. M. Nelson, D. Brand, and M. Mandel. Use and Consequences of Timed Transfers on U.S. Transit Properties. TRB, Transportation Research Record 798, 1981, pp. 50-55.

Publication of this paper sponsored by Committee on Bus Transit Systems.

Short-Term Ridership-Projection Model

CY ULBERG

A statistical model has been developed for use by a transit agency in making short-term forecasts of transit ridership. These factors have been used successfully to plan service changes and to forecast revenues. A by-product of the use of the model is an increasing understanding by staff members of determinants of ridership changes and a corresponding reduction in the emphasis on ridership as a performance indicator by the agency. The model uses a combination of multiple-regression and time-series analyses to produce monthly projections of ridership. The variables included in the model were chosen for simplicity, ease of collection, and explanatory power. The validity and reliability of the model are quite strong, given its simplicity. During a two-year validation period, the average monthly error was 2 percent. Errors in annual totals were 0.9 and 1.7 percent, respectively. One objective in the development of the model was to make it a useful tool for planners and managers within the agency. A monthly report has been developed that has become a part of the decisionmaking process in the agency. Even though experience with a model has been limited, it has been demonstrated that a transit agency can make use of a relatively sophisticated (although simple) statistical technique to develop ridership forecasts.

In the past several decades, transit ridership has varied dramatically. Long-term trends have been influenced by phenomena such as the rising popularity of the automobile, world wars, and population shifts from farms into cities and suburbs. In contrast to these long-term trends, short-term ridership gains and losses occur due to more rapidly varying factors such as seasonal effects, service levels and quality, fares, gasoline prices and supply, parking rates, employment, and population. This paper describes one transit agency's experience with producing useful short-term forecasts.

Transit agencies use a variety of nonstatistical and quasi-statistical methods to produce forecasts of ridership. Generally, these methods use interpretations of past trends modified by management objectives for increasing ridership. Most agencies try to predict the impact of fare changes and service changes on ridership. In the Seattle metropolitan area, Metro Transit traditionally has projected ridership by using a modified Delphi technique. Objectives for productivity (passengers per hour) were set by using qualitative assessments of the environment, particularly the impact of fare and service changes. Service hours were projected by using budget constraints and perceived ridership demand. Total ridership projections were determined by multiplying productivity and service hours.

When ridership changes were relatively stable (such as between 1975 and 1979), these methods worked fairly well. However, in 1980 ridership trends changed abruptly. A gasoline crisis and rising employment were followed by a drop in gasoline price and declining employment trends. A major fare increase was implemented. Rapid increases in ridership changed to a leveling-off period. The extent of the change was unanticipated and resulted in major adjustments in service planning and budgeting.

In order to anticipate similar short-term changes in the future, Seattle's Metro Transit has developed a short-term ridership-projection model. It has been used during the past year to assist in the preparation of revenue projections and in planning service changes. It has also been used to anticipate the impact of a fare increase implemented in February 1982. Because the model uses variables extraneous to Metro's control, such as gasoline price and supply and employment, it has helped develop a new perspective on the use of ridership data for evaluating the effectiveness of the transit agency and its components.

BASIC STRUCTURE OF MODEL

A major objective in the development of the model