Use and Consequences of Timed Transfers on U.S. Transit Properties

MICHAEL NELSON, DANIEL BRAND, AND MICHAEL MANDEL

A recently completed study conducted for the Transportation Systems Center examined the current use and impacts of a total of 11 transfer policy options (including timed transfers) on U.S. transit properties to identify the situations or settings in which particular transfer policies can most beneficially be applied. Data for the study were drawn from a series of telephone and on-site discussions with experienced transit professionals on 39 different properties. The information resulting from these discussions has been supplemented with a limited amount of site-specific quantitative data and references to the literature as appropriate. The findings of that study regarding timed transfers are presented. Implementation of timed transfers can involve adjustments of headways, route lengths, and/or layover times as well as provision of suitable space, facilities, and information to permit the easy interchange of passengers between buses. Transit-property size is the principal criterion for the applicability of timed transfers, serving as a proxy for headway reliability, service frequency, and the number of buses meeting at one time. Small properties are generally able to use timed transfers at their main transfer point, whereas larger properties may only be able to use this option on a relatively more limited scale. Ridership gains on the order of 5-12 percent may be realized under some circumstances solely from the implementation of timed transfers. Overall, timed transfers appear to be a cost-effective way of increasing service and ridership in many settings without necessarily increasing costs.

A timed transfer is defined as any set of operator actions that provides for vehicles on different routes to meet at regular intervals to exchange transferring passengers. The simplest form of timed transfers involves only two routes. At the other end of the complexity scale is the extensive use of timed transfers (also known as "pulse scheduling"), where vehicles on all (or most) routes are scheduled to meet at the major transfer point nearly simultaneously, hold until all the vehicles have come in, and then leave together. When this occurs at regular intervals, the effect is as if the vehicles were pulsing.

In between these extremes are two other types of timed transfers. When pulse scheduling of buses is used only in the evening or off-peak hours, with low service frequencies and possibly long layovers at the transfer point, it is commonly called a "lineup". Unlike pulse scheduling, lineups are found in larger cities. Another variant of timed transfers, "neighborhood pulse", is also found on large properties. It involves coordinating the schedules of neighborhood bus circulator routes to make travel within a section of a city easier.

This paper examines situations in which these variants of bus timed transfers are used and the operator actions associated with implementing them. (Examples of simple timed transfers in which rail is the connecting mode were also found, but this practice is quite uncommon and is not addressed further in this paper.) The effects of timed transfers on operator costs, user satisfaction, ridership, and revenue are then described, and conclusions are drawn concerning the applicability of timed transfers in different settings.

The findings in this paper are drawn from a recently completed study of 11 transfer policy options (including timed transfers) on U.S. transit properties (1). Data for the study were drawn largely from an extensive series of telephone and on-site discussions with experienced transit professionals on 39 different properties, 16 of which used timed transfers of some sort, as indicated below:

Type of Timed Transfer	Property
Simple	Albany, New York
	Washington, D.C.
Pulse scheduling	Fresno, California
	Lafayette, Indiana
	Brockton, Massachusetts
	Westport, Connecticut
	Lewiston, Maine
	Haverhill, Massachusetts
Lineup	Knoxville, Tennessee
	Portland, Oregon
	Columbus, Ohio
	Memphis, Tennessee
	Toledo, Ohio
	Albany, New York
Neighborhood pulse	Denver, Colorado
	Portland, Oregon

This sample clearly does not include all properties that use some form of timed transfers, nor was it a random sample designed to yield statistically representative results. Rather, the survey was designed to yield the greatest possible amount of information on different operating environments and practices as was feasible with a limited sample size.

CURRENT PRACTICES

As the following table indicates, the demand for transferring on a transit property is clearly related to the type of transfer policy adopted:

Use of Timed Transfers	Transfer Rate (%) Average	Low	High
Extensive	28	18	50
Not extensive	18	5	33

As the table shows, bus properties that currently use timed transfers extensively have an average transfer rate of 28 percent, whereas properties that do not use timed transfers extensively have an average transfer rate of approximately 18 percent. Furthermore, when the properties that use timed transfers extensively (all of which are small) are separated from the remaining small and large properties, the suggested relation between ridership and the use of timed transfers becomes even more pronounced:

Type of Transit Property	Transfer Rate (%) Average	Low	High
Large	20.3	10	33
Small			
No use of timed transfers	11.8	5	20
Use of timed transfers	28.0	18	50

It should be noted, though, that the causal relation here is not clear. Timed transfers may increase transfer ridership through a reduction in transfer time but, conversely, the existence of travel patterns that result in a high transfer rate may make it more likely that a property will institute timed transfers. Thus, it cannot be concluded with certainty that the use of timed transfers will always cause substantial increases in transfer ridership. Rather, it is necessary to consider carefully the circumstances surrounding each possible application of timed transfers before ridership and other impacts can reliably be assessed.

Simple Timed Transfers

Simple timed transfers, where buses on two routes are guaranteed to meet regularly, illustrate the basic principles of timed transfers. Simple timed transfers are used on many properties, from the smallest to the largest. They are most commonly used in the evening when both routes have low frequencies. Simple timed transfers, almost by definition, are more likely to be found at outlying transfer points where few routes may meet. However, their use is not restricted to any particular setting.

In order to implement simple timed transfers, schedules must be adjusted so that buses arrive at the transfer point at the same time. There are differences, though, in the way operators handle the unavoidable problems of schedule unreliability. Some operators have the buses lay over for 2-5 min at the transfer point, assuming that such a layover provides enough of a cushion to ensure that the buses will meet. Other operators use "dynamic control" to hold the first bus until the second bus arrives, if the second bus has transferring passengers. (This real-time modification of schedules is usually accomplished through verbal communication by radio, although other communications media are sometimes used.) The problem with "static control", where each bus is simply scheduled to hold until the other arrives, is that, if one bus breaks down or is extremely late, the schedule of the second bus is needlessly disrupted. Therefore, true static control is rarely used and a limit is typically placed on the length of time spent waiting. All of these operator actions, however, have the common objective of guaranteeing transfers with a low wait time between two routes.

Pulse Scheduling

Pulse scheduling, or extensive timed transfers, is the type of timed transfer that has the most far-reaching operational consequences. The transit properties that currently use this option are extremely diverse, serving a wide variety of communities all over the United States, including college towns, industrial cities, and bedroom suburbs. Table 1 gives service, route, and scheduling data (based on operator interviews, timetables and maps from each property, and data from the American Public Transit Association) for several of these properties that participated in the study sample. Annual ridership among these properties ranges from the tens of thousands to more than 4 million. Fleet sizes range from 3 to about 100 buses and service-area populations from less than 30 000 to more than 300 000. It is interesting to note, however, that virtually all of the pulse properties offer free transfers. This may reflect both a "philosophy" on pulse properties of simplifying and reducing the burden of transfers and an effective marketing approach.

Important aspects of pulse transfers include service frequency, routing, schedule adherence, space for buses to meet, and operator information policies. Since all buses are meeting, it is possible to speak of a "pulse frequency" of which all route frequencies are a multiple. The most common pulse frequency is 30 min. Other pulse frequencies such as 35 and 45 min (with some buses meeting in between the major pulses) are also in use, as Table 1 indicates. These frequencies typically do not change much between the peak and off-peak periods, although some properties halve the frequency in the evenings and on weekends.

Given the ranges of headways that are found on different properties and the different ways of making them compatible, it is somewhat surprising to find that 30 min is almost uniformly perceived as the preferred pulse frequency. Many operators believe that a 30-min headway makes the transit system easier to understand. In addition, 30-min headways are quite compatible with clock-face scheduling. These reasons are consistent with the design of timed transfers as a popularly supported, easily understood, and not easily disrupted public transit system. Virtually all of the properties that use a pulse other than 30 min originally implemented pulse at a 30-min frequency and later modified it because of schedule unreliability, increases in ridership that led to longer running times, or other site-specific reasons.

Because of the need for a uniform frequency, implementation of timed transfers may involve reducing frequency on some lines, which would reduce level of service, or increasing frequency on others, which is costly. Both possibilities require making "artificial" changes to the schedule that may be wholly unsatisfactory in some settings. Forcing a wide variety of routes to meet in time and space may be essentially infeasible, especially in large cities. In the opinion of one experienced transit professional, "In large cities, crosstowns are better and cheaper, too."

The need for all or most routes to have the same headway in turn constrains the routing of buses. When implementing pulse scheduling, many properties find that their natural routes are too long and that pulse limits their route miles. A typical remedy is to cut short the ends of the routes. In addition,

Table 1. Characteristics of sample U.S. transit properties that use pulse scheduling.

City	Population of Service Area (000s)	Transfer Charge (cents)	Total Annual Ridership (000s)	Approximate No. of Buses in Fleet	Approximate No. of Routes in System[a]	Pulse Frequency[b] (min)	No. of Buses Meeting Each Pulse
Eugene, Oregon	210	0	2860	67	20	30	12
Lafayette, Indiana	110	5	1100	29	14	30	14
Brockton, Massachusetts	100	0	2790	32	16	45[c]	15
Westport, Connecticut	30	0	640	25	7	35	7
Fresno, California	310	0	4400	99	21	30	12[d]
Lewiston, Maine	70	0	460	18	9	30	9
Everett, Washington	50	0	1020	18	12	30	11
Haverhill, Massachusetts	50	0	90	3	6	30[e]	3

[a]Not including express routes.
[b]Generally constant throughout the day.
[c]Some routes meet every 22.5 min at an intermediate pulse.
[d]With "syncopated pulse", the buses are divided (seven and five) between two adjacent pulse points.
[e]Because three buses run six routes, each route has a 60-min headway.

through routing can be used to achieve headways that do not divide evenly into route run times. On the other hand, several properties have routes that are too short. The operator response to this problem is typically to increase layovers to equalize running times or to extend the routes by loops or other means, thus adding area coverage.

In practice, the choice of a pulse frequency can never be made independently of routing decisions. A major influence in the balance between frequency and routing is the size and shape of the relevant transit district. Lafayette, Indiana, and Brockton, Massachusetts, are two pulse cities that have relatively compact service areas and thus have no trouble operating a 30-min pulse with good loop area coverage. Another pulse property, Everett, Washington, had difficulty expanding the length of its routes because the service area is long and thin and the central business district (CBD) is not in its geographic center. In general, properties whose CBDs are in the geographic center of the relevant area find it easier to pick an appropriate pulse frequency and then equalize running times on different routes based on the size of the area.

Ensuring schedule adherence is a major problem for pulse properties. The reasons for schedule unreliability tend to be the same as those on nonpulse properties: traffic congestion, breakdowns of new buses, and interference from trains. However, since the essence of timed transfers is to ensure that transferring passengers make connections, maintaining schedule adherence is more important on pulse properties.

Two strategies are available for coping with problems of schedule reliability on a pulse system. The first strategy is to build extra layover time into the schedule. Most pulse systems use layovers of 5 min or less out of each half-hour. Use of additional layover time is limited if the same schedule is to be used for both peak and off-peak periods. That is, if long enough layovers are added to absorb peak-hour unreliability, there will be costly unused layovers during the off-peak period. However, layovers of 5 min or less are usually not sufficient to handle all schedule-adherence problems.

Therefore, almost all pulse properties also use the second possible strategy, dynamic control, to mitigate problems with schedule reliability at the pulse point. In general, buses will hold for a maximum of 3-6 extra min for late buses before leaving the pulse point. If the late bus is radio equipped, the driver can inform the dispatcher or starter which routes will be receiving transferring passengers so that buses can be selectively released.

Typically, "lengthy" detention of buses through dynamic control is used most effectively during off-peak hours, during the last pulse of the day, and toward the end of the peak. Its use is avoided at the beginning of peak hours because during the peak buses have difficulty catching up to the schedule if they have been held any length of time. It is generally thought to be better to let one or two peak-hour buses miss the pulse than to disturb the rest of the system. On the other hand, it is very important on the last trips of the day to ensure that no one is stranded.

Some properties use short layovers and static control (holding "blindly" for up to 5 min) to deal with schedule uncertainty. These operators, who do not have radios, sometimes encounter a situation that might be called "disintegrating pulse". Because layovers are short and buses may be detained inefficiently, routes on which traffic congestion is bad may not be able to stay on schedule and are simply dropped from the pulse.

Another important requirement of pulse scheduling is the provision of space for buses to meet at the pulse point. Most pulse properties have a single pulse point that is located in the CBD. Typically, 9-12 buses occupy the pulse point at each pulse, although as many as 15 or as few as 3 buses have been observed in practice, depending on the size of the system. These numbers refer only to the buses meeting the pulse; pulse points may have unsynchronized routes that terminate there as well.

There is a need to keep all pulsing buses close together for the benefit of riders and for better control of the pulse. Buses are most often distributed over one or two blocks along a street. This may create problems for some passengers, since it is far enough to cause them to miss their buses. In general, though, the use of on-street stops is not viewed as intolerable by operators. Of the pulse properties participating in this study, only Brockton has an off-street facility and that was only opened in March 1979.

Two other properties have adopted atypical solutions to the problem of arranging buses at the pulse point. Because of space limitations, Fresno, California, had to adopt what might be called syncopated pulse, where the buses at one pulse point are routed so that they pass by the other pulse point both coming in and going out. The buses that terminate at the first pulse point drop off passengers at the second pulse point just before it pulses and pick up passengers at the second pulse point just after it pulses. In this way, passengers can make their transfers within a reasonably short time without all of the routes having to terminate at the same spot.

The second property that has used an unusual pulse-point arrangement is Lafayette. Until January 1979, Lafayette had two pulse points, one in the CBD and a second at Purdue University. The two pulse points were approximately 1 mile apart across a river and connected by a shuttle route that met both pulses. This arrangement was originally instituted to increase coverage to the west side of town and to keep large numbers of buses off the single major bridge over the river. However, Lafayette went back to using a single pulse point in the CBD in January 1979 because of problems in adhering to schedules.

This experience raises the problem of conflict between pulsing buses--both parked and moving in platoons--and automobile traffic. In many cases, some traffic engineering work and cooperation from the police are necessary to ensure smooth operations. These aids, and the possible tendency of automobiles to avoid "pulse streets", tend to keep traffic-congestion problems to a minimum.

Pulse properties vary considerably in the degree to which they publicize their use of pulse scheduling. Several properties make it clear from their schedules that pulse scheduling is a keystone of their system. Other properties place some emphasis on pulse scheduling without making it the dominant feature of the system. Finally, there are some properties that do not highlight their use of pulse scheduling at all. This last group includes systems that historically have had some sort of timed transfers or clock-face scheduling and do not regard it as an especially distinctive feature.

Other Types of Timed Transfers

The other variants of timed transfers--lineups and neighborhood pulse--are basically pulse scheduling applied in different situations. A lineup is pulse scheduling used in the evening and in off-peak hours. A neighborhood pulse is pulse scheduling used on only a portion of the system. Most of the operator actions associated with these variants are similar to those for pulse scheduling. The major differences that do exist are pointed out below.

Lineups

Lineups are used by many nonpulse transit properties in the evening or on weekends. The populations served by the sample of lineup properties that participated in this study ranged from 190 000 to 1 800 000, and all but one served more than 500 000 people. Most of these properties use a headway of 1 h for their lineups, which is the same headway often used in the evening by pulse properties.

Given that the term lineup conjures up an image of a row of buses sitting in a line for long stretches of time, it is important to note that most lineups have no more than 5- to 10-min layovers. Again, there may be some adjustments made in routing to accommodate the schedule. For instance, one property reduces some coverage of outlying suburbs, while another adds a "night loop" to some routes. Most of the other actions taken by properties are the same for lineups as for pulse scheduling. In addition, emphasis may be placed on the fact that lineups tend to guarantee that no one gets stranded after the last trip of the day.

Neighborhood Pulse

The difference between neighborhood pulse and full-scale pulse systems is the size of the system in which the pulsing routes are found. With neighborhood pulse, a set of local routes pulse together to facilitate travel within a neighborhood. Because this may occur in areas outside of the congested CBD, neighborhood pulse can be found in very large cities or on any property that has non-CBD subcenters that are logical transfer points. The actions required to do this are quite similar to the actions associated with pulse scheduling.

CONSEQUENCES OF TIMED TRANSFERS

The use of timed transfers does not inevitably lead to any particular set of consequences. Simple timed transfers, pulse scheduling, lineups, and neighborhood pulse clearly all require different levels of effort and generate impacts of different magnitudes. Even within properties that use pulse scheduling, impacts vary greatly depending on the required operator actions. This wide divergence of possible impacts follows directly from the multiplicity of actions that make up timed transfers (described earlier). For the purpose of detailing consequences, these operator actions will be divided into the five categories addressed above: service frequency, routing, schedule adherence, provision of space for buses, and provision of user information. The analysis of each type of consequence--cost, user satisfaction, ridership, and revenue--will focus on those categories of operator actions that have the greatest impact.

Costs

The greatest potential influence of timed transfers on cost arises from changes in bus hours and bus miles that may have to be made to match headways on different routes. In practice, however, it is not clear whether this is an important effect. Frequency changes for simple timed transfers and lineups seem to be small, especially since headways in the evening are often fixed by policy. Frequency changes for pulse scheduling and neighborhood pulse are potentially more significant, but it is impossible to tell in general whether frequencies will be raised, lowered, or both on any particular property. In practice, pulse properties appear to have somewhat longer peak headways and somewhat shorter base headways than comparable nonpulse properties. The operator may feel that, because of reduced transfer time, the peak-period headways can be raised without reducing the overall level of service. Alternatively, the operator may decide to maintain the peak headway to accommodate work riders.

In general, properties that use pulse scheduling do not attribute major cost consequences to frequency changes mandated by the use of timed transfers. Because of site-specific factors, however, it is not possible to anticipate the direction or magnitude of the changes in service frequency needed to implement pulse scheduling in cities that currently do not have it. These impacts must be assessed on the basis of the policies selected by the operators and the preexisting schedule.

The systematic dollar cost differences that do exist between pulse and nonpulse properties stem mainly from extra layover time built into the schedule to ensure schedule reliability. Because timed transfers are based on guaranteeing that buses will meet, more system resources are devoted to this end. As extra layovers are built into the system, two distinct effects can occur. With a greater fraction of vehicle time spent idle, cost as estimated on a per-mile basis will increase because of the decrease in vehicle miles of travel (VMT). Actual total operating costs may decrease because of savings in bus running costs (if no more buses are added). The conflict between these indicators and the small expected size of the impact are compatible with the indecision of many operators concerning the overall cost impacts of pulse scheduling.

Another cost of pulse scheduling that can be significant is the cost for the street space used by the pulsing buses. This cost is not normally a direct financial burden on the operator in the usual sense. However, consumption of street space by the buses can cause an increase in traffic congestion and a reduction in parking-meter revenues as well as aesthetic problems. These costs are not borne by the transit operator but may have to be taken into account in deciding whether to implement pulse scheduling.

User Satisfaction

User satisfaction among transferring passengers almost always increases significantly when any type of timed transfer is used. However, there are several factors that appear to influence the degree of change in user satisfaction, including service reliability, comprehensibility, frequency, and, to a lesser degree, coverage.

Reliability is the key element in determining whether user satisfaction increases sharply with timed transfers. If riders are assured of a very high probability of making their connection, both the mean and the variance of transfer wait time will go down. The variance is especially important because one bad experience can counteract the effects of a large number of good ones. Therefore, operator actions to ensure a high degree of reliability in making connections are essential for a large gain in user satisfaction.

The comprehensibility of the system is a second important determinant of the changes in user satisfaction that accompany timed transfers. If an operator makes riders aware of the timed transfers, then the system is easier to understand and use. Riders who want to transfer need not worry about when the connecting bus will arrive at the transfer point. Schedules are thereby simplified and made less confusing.

Service frequency also affects user satisfaction with timed transfers, although its effect is essentially inverse. Since high frequencies lead to low average transfer wait times, even without timed transfers, the implementation of timed transfers would have a reduced positive effect on user satisfaction. On the other hand, low frequencies mean that timed transfers can have a large positive impact on average transfer wait time and hence on user satisfaction. Simple timed transfers and lineups, which are typically used at times of low bus frequency, are thus more likely to greatly increase user satisfaction. It should be emphasized that this relation to frequency focuses on the change in user satisfaction induced by timed transfers. The overall level of user satisfaction typically would be higher with high service frequencies.

The final factor, coverage, is generally less significant than the first three in determining user satisfaction. It is true that operators often adjust routes to accommodate pulse scheduling or lineups. These changes can affect overall coverage on the outlying portions of routes, the streets used to reach the downtown terminal point in the allocated running time, or the location of the terminal point itself. In practice, pulse scheduling and lineups have had little effect on coverage of outlying areas. However, in at least two cases, changes in the terminal point have affected the level of service available to both transferring and nontransferring passengers. In Brockton, the off-street transfer facility was located several blocks away from the previously used pulse point, which had been closer to the center of town. This led to a net increase in the distances that many people had to walk to gain access to transit. In Lafayette, when the dual-pulse-point system was instituted, people who had formerly traveled on one bus from the west side of town to the CBD were compelled to transfer, which reduced their level of service. In general, though, coverage seems to have been affected in only a minor way.

It is important to consider how changes in user satisfaction affect different groups. Geographically, the four categories of timed transfers inherently have different consequences for different groups of riders. Simple timed transfers only increase user satisfaction for individuals transferring between two particular routes, whereas the effects of neighborhood pulse are restricted to riders in a particular area. Lineups, which are typically used on a systemwide basis, have consequences only for people traveling during off-peak hours. Pulse scheduling will affect the satisfaction of almost all riders; the elderly, the young, and people who transfer frequently will experience the highest gain. On the other hand, riders making peak-hour work trips may have much less of a gain from pulse scheduling because of the heavily radial nature of their trips. For such riders, transfer policy options such as through routing, which eliminates transfers altogether, may be much more beneficial.

Ridership

Simple timed transfers or lineups clearly do not produce large gains in ridership, since the typically long headways on the originating leg remain an important determinant of ridership. On the other hand, some properties have experienced substantial increases in ridership because of the use of pulse scheduling, although many of these properties instituted other service improvements simultaneously with the pulse scheduling. In Brockton, for example, ridership increased sixfold at a time when VMT was increased fourfold to fivefold. Since only 25 percent of passengers now transfer and the reliability of service has drastically improved, a reasonable estimate of the increase in ridership directly attributable to pulse scheduling may be on the order of 10 percent of current ridership. This estimate is substantiated by the experience in Superior, Wisconsin, where ridership rose 10-12 percent with the advent of pulse scheduling and there were no other important changes in service. Furthermore, several pulse operators (including those in Everett, Washington, and Lewiston, Maine) see no definite link between pulse scheduling and ridership.

It is possible to estimate the ridership changes caused by implementation of pulse scheduling. For example, consider a small property where all routes meet at one point, all have unsynchronized headways of 30 min, the overall transfer rate is 20 percent, and the transfer charge is zero. Before pulse scheduling, the average out-of-vehicle time for transferring passengers will be 30 min (15 min transfer time plus an assumed 15 min of initial walk and wait time and final walk time). With pulse scheduling, transfer time will drop to 5 min for a total average out-of-vehicle time of 20 min.

Under this scenario, the increase in ridership attributable to pulse scheduling can be calculated in two different ways. The first way uses the -0.7 elasticity of demand with respect to out-of-vehicle time presented by Domencich (2). Since, with pulse scheduling, out-of-vehicle time for transferees decreases by 33 percent, the number of transferring passengers will increase by 23 percent (0.33 x 0.7). If the initial transfer rate was 20 percent, the overall ridership will increase by 4.6 percent (0.23 x 0.20). This figure does not take into account the change in overall user perception of the system as conducive to reliable transferring. There is a belief shared by several pulse operators that timed transfers at the downtown terminal promote a comprehensible, easily "imaged", and popularly supported system that leads to more riding than simple reductions in waiting time between two connecting lines would suggest.

A second elasticity-based method for predicting the ridership consequences of pulse scheduling uses the pre-Bay Area Rapid Transit aggregate demand elasticity of San Francisco ridership with respect to transfer time (only) of -0.26 calculated by McFadden (3). For the above example, this yields an overall ridership increase of approximately 17 percent (0.67 reduction in transfer time x -0.26). It should be noted that this increase may be equated in the above example to an elasticity of -2.5 for all out-of-vehicle time alone (17/20 percent x 0.33). There is support for bus service elasticities this high under conditions of infrequent service (e.g., comparable to long waits at transfer points--prepulse) and relatively high fares (4,5).

Overall, 5-17 percent appears to be a reasonable range for the ridership effects of pulse scheduling. The higher increases would be more likely for systems that increased service reliability at the same time and/or had the potential for significant riding to nondowntown terminal locations because of the presence, at dispersed destinations, of major attractors of discretionary trips or trips by the elderly.

Revenue

The revenue consequences of timed transfers follow directly from ridership consequences as long as the distinctions between groups paying different fares are observed. The key question is whether the revenue gained from increased ridership covers the cost of setting up a reliable pulse-schedule system.

This question can be addressed in a hypothetical case by using the above example. Consider again the small transit property with 30-min headways and without pulse scheduling. To implement pulse scheduling and attract ridership, reliability may have to be increased by adding layover time. Assume this added layover time to be 5 min added to the previous running-plus-layover time of 55 min (two buses on each route).

In this example, assume pulse scheduling is to be implemented at no additional operating cost. Therefore, VMT must be decreased proportionally--that is, by 9 (or 5/55) percent. (In fact, the decrease may be slightly less because layovers decrease mileage-related costs.) Using the typical bus VMT service demand elasticity of -0.7 yields a VMT-related decrease in ridership of 6.3 percent. This decrease in ridership would probably be less because the VMT changes take place at the ends of the routes that are likely to be in low-density areas. In any case, this decrease in ridership from the added layover time is at the low end of the range of the above estimated pulse-schedule-induced ridership increase. Hence, a no-cost implementation of pulse scheduling under this scenario may still attract additional ridership and may be a productive option in this situation. Unfortunately, the actual cost and ridership effects of timed transfers in any real application depend heavily on policies undertaken by the operator to equalize roadways, provide adequate space for all buses to meet, etc. The site-specific nature of all of these factors makes it impossible to generalize results except to say that many operators believe that timed transfers of some sort are the most efficient means available by which to provide improved levels of service under many circumstances.

APPLICABILITY OF TIMED TRANSFERS

Property size is the principal criterion for assessing the applicability of the four different types of timed transfers. Transit properties whose service areas have fewer than 400 000 people are generally able to use pulse scheduling at their main transfer point. On the other hand, larger properties often have lineups but not pulse scheduling. Simple timed transfers can generally be used on any system, although they are more likely to be found on medium-sized properties. This is because small properties usually do not have significant outlying transfer points, whereas large properties have more complex systems for which the scheduling of simple timed transfers at numerous outlying transfer points may not seem worth the effort. Finally, neighborhood pulse is applicable to any system that has subcenters that serve as logical pulse points.

Several other factors, some of which may be related to property size, also affect the general applicability of timed transfers. The first is schedule reliability, which is very important for increasing user satisfaction. A disintegrating-pulse situation, where people cannot be assured of meeting their buses, eliminates the rationale behind a timed transfer system. Hence, cities in which the transit property has problems adhering to schedules would have difficulty using timed transfers in general and pulse scheduling in particular. In addition, on large properties that have severe schedule-adherence problems, increasing user satisfaction by means of timed transfers would tend to be prohibitively expensive. This is one reason why large properties tend not to use pulse scheduling during the day and instead concentrate on times when schedules are more reliable.

Service frequency also influences the applicability of timed transfers. At high enough frequencies (e.g., 15 min), the drop in average transfer time attributable to timed transfers is not significant enough to substantially increase user satisfaction and begin to offset the added costs of timed transfers in city centers (e.g., street space, congestion, etc.). Since larger properties tend to have high service frequencies even during the day, this constrains the applicability of pulse scheduling and possibly neighborhood pulse as well.

Space limitations on the number of buses that can meet at a point also have an important influence on the applicability of timed transfers. The difficulty involved in finding a place in a congested area where all buses can meet explains in part why large properties avoid daytime pulses, resorting instead, in some cases, to lineups in the evening when the CBD is less congested. Moreover, even if there is a place to meet, the distance between buses will have a very significant effect on the transfer time.

Given these size-related reasons why pulse scheduling is implemented only by small properties and lineups are implemented only by large properties, it is appropriate to outline the reasons why the use of pulse scheduling of buses varies among small cities. Clearly, widely dispersed origins and significant numbers of non-CBD destinations indicate that the city is a candidate for pulse scheduling. In addition, geographic layout--the CBD being in the center of the service area, for instance--can make scheduling easier. However, the most influential factor seems to be a political climate in which transit innovation can occur. If political factors determine the level of service allocated to different areas, pulse scheduling may not be feasible. This type of constraint must be addressed on a case-by-case basis. However, if the political climate is conducive to a major change and revamping of service, pulse scheduling has the potential of being a cost-effective way of increasing service and ridership without necessarily increasing operating costs.

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 of the Urban Mass Transportation Administration. We would like to thank Robert Casey of TSC for his many helpful comments and observations. We also gratefully acknowledge the contributions of the transit professionals who participated in this study. However, the opinions and conclusions expressed or implied in this paper are ours, and we are solely responsible for any errors of content or omissions that remain.

REFERENCES

1. Charles River Associates, Inc. Operator Guidelines for Transfer Policy Design. U.S. Department of Transportation, June 1980.
2. T.A. Domencich, G. Kraft, and J.-P. Valette. Estimation of Urban Passenger Travel Behavior: An Economic Demand Model. HRB, Highway Research Record 238, 1968, pp. 64-78.
3. D. McFadden. The Measurement of Urban Travel Demand. Journal of Public Economics, Vol. 3, 1974.
4. R.L. Carstens and L.H. Csanyi. A Model for Estimating Transit Usage in Cities in Iowa. HRB, Highway Research Record 213, 1968, pp. 42-49.
5. Peat, Marwick, Mitchell, and Company. Public Transportation Fare Policy. U.S. Department of Transportation, May 1977, pp. III.40-III.44.

Publication of this paper sponsored by Committee on Public Transportation Planning and Development.