Comparative Evaluation of Bus Route Costing Procedures

DOUGLAS W. CARTER, SUBHASH R. MUNDLE, and BRIAN E. McCOLLOM

ABSTRACT

Because policies regarding transit funding are changing at all levels of government, transit planners will be required to more carefully monitor existing bus systems as well as intensively examine the net cost or savings of proposed service changes. In the past two decades, research has been focused on demand and, hence, revenue estimation. Current financial constraints suggest that in the coming years more effort will be focused on operating cost estimation and the underlying relationships that affect expenditures. Although a variety of cost estimation techniques has been developed and used by transit operators, no single technique has been generally recognized as more accurate than the others. The purpose of the project, of which this paper is a product, is to develop a technique (i.e., the proposed method) that is sensitive to the salient cost characteristics of route-level service changes and is relatively easy to understand and apply. A key feature of this project involved a comparative analysis of the performance of the proposed method, and several other prominent cost techniques, against a scheduling-based route costing model.

Operating in an era of diminished public funding, transit agencies are pressed to find ways to reduce, or constrain, expenditure of scarce operating dollars. Many systems, facing severe financial constraints, have already initiated substantial service changes to balance costs with available funds. This recent trend in the transit industry will place greater demands on transit planners to forecast, with reasonable accuracy, the financial impacts of planned service changes. Although several approaches have been suggested or used in the past, no single technique or approach has proven entirely satisfactory. Recognizing the need for a reliable and relatively simple incremental cost estimation technique, UMTA commissioned this research effort to develop and test a bus route costing procedure.

The study is comprised of several interrelated tasks. The initial task entailed review and evaluation of cost estimation techniques used in the industry and procedures suggested in the technical literature. Following this assessment, a proposed method was developed for identifying incremental cost implications of bus route service changes. Next, a techniques test was conducted for the proposed method and other prominent models using actual and hypothetical service changes at the Metropolitan Transit Commission (MTC) in Minneapolis-St. Paul. On the basis of the test results, the proposed method and several other models were revised to enhance applicability and accuracy. The concluding step of the study entailed documentation of the proposed costing technique and preparation of a step-by-step manual for calibrating and applying the model.

The results of the cost estimation techniques test are described and evaluated. The test entailed application of five incremental cost models to 12 actual and hypothetical service changes at the MTC. The model results are compared and evaluated in terms of accuracy, sensitivity, and level of effort. Although the test was limited in scope (i.e., only 12 service changes were involved), it was guided by scientific research principles to ensure objectivity. The results of the techniques test cannot be assumed to be statistically valid for all situations, but the test does provide important insights into model applicability and relative strengths and weaknesses.

TEST FRAMEWORK

Four key groups participated in the techniques test: the consultant, MTC staff, the review panel, and UMTA staff. Each group performed a different role in executing the test. The consultant directed all test activities and was responsible for orientation of the MTC staff, quality control in model application and evaluation of test results. The MTC staff calibrated and applied each of the models to the 12 service scenarios comprising the test. The review panel, comprised of industry costing experts, provided direction and critiqued findings and analysis at critical points in the test. UMTA staff members also provided project guidance and assisted in the orientation of MTC staff to each of the cost estimation techniques.

Five costing techniques were calibrated and applied during the test:

- Proposed method,
- Modified Adelaide model,
- Peak and base cost allocation model,
- Two-variable cost allocation model, and
- Scheduling-based cost model.

Proposed Method

The proposed method focuses on driver-related costs because driver wages and benefits comprise by far the largest portion of cost impacts resulting from a service change. Other, nondriver, incremental costs are estimated using a traditional two-variable (i.e., hours and miles) cost allocation approach. The proposed method, like all techniques involved in the test, is sensitive only to those costs that typically vary in response to changes in the scale or characteristics of fixed-route service (i.e., variable costs). Fixed costs are neither considered nor estimated by the technique.

In the proposed technique, detailed analysis of driver cost begins with the number of platform hours, stratified by time of day, for both before and after the service change. Run-type ratios calibrated from existing driver and service schedules at the division level are modified to reflect the unique characteristics of the route being changed. Next, driver assignments, spread premium hours, and overtime hours are estimated by applying the calibrated ratios to the new platform hours at the division level. Weekly driver requirements are then estimated from assignment projections and the daily driver availability ratios are developed in the calibration phase. Finally, wage and benefits costs are determined in terms of regular, overtime, spread premium, show-up wages, paid absences, and variable and fixed benefits. The incremental wage and benefits cost is estimated by taking the difference between cost projections for before and after the service change. Nondriver costs are estimated by applying calibrated hour and mile rates to the net change in each resource unit.

Modified Adelaide Model

Developed by the British firm of R. Travers Morgan and Partners, the Adelaide model incorporates several novel approaches to cost estimation. One attractive feature of this model is a simplified driver scheduling algorithm that transcribes buses-in-service, by time period, into driver work assignments. This model is applied both before and after the service change to estimate the incremental cost impact.

The original Adelaide model does not address trippers because this type of assignment is not used in Australia or Great Britain. Trippers, however, occur with great frequency in the United States. To make the model meaningful for this study, both the scheduling and the costing algorithms were modified to address tripper assignments. Scheduling ratios were calibrated by assessing the weighted average of worked to pay hours for split runs, tripper combinations, overtime trippers, and part-time trippers separately for the morning and evening peak periods. Average worked and penalty hour costs were determined at the division level and appropriate rates were applied to route changes.

Application of the Adelaide model, as modified for this test, requires four primary steps that are applied at the route level for before and after the change. First, vehicle requirements are estimated by time of day based on round-trip time and service headways on the subject route. Second, vehicles in service are transcribed into driver requirements using a simplified scheduling algorithm. Third, incremental worked and penalty hours are estimated on the basis of division-wide scheduling and pay practices. Finally, the total route cost is estimated by applying the net change in worked hours, penalty hours, vehicle-miles, platform hours, and peak vehicles to calibrated resource unit costs. The incremental cost is the difference between before and after estimates.

Peak and Base Model

This model represents an enhancement of the traditional cost allocation models in that it allows hourly costs to vary relative to the amount of peak and base service. Two vehicle-hour cost rates (i.e., cost per peak hour, cost per base hour) are obtained by adjusting the total unit cost regularly produced in cost allocation. This adjustment is performed by calculating two indices--one representing relative labor productivity (i.e., pay hours per platform hour by time of day) and the other representing the ratio of peak to base service.

Cost Allocation Model

The cost allocation model is probably the most commonly used method of estimating cost impacts of service changes. The model is predicated on the concept that cost is a function of one or more resource quantities, such as vehicle-hours, vehicle-miles, or peak vehicles. Unit cost rates per resource unit are found by assigning each individual expense in the system's financial statement to a resource, summing the expenses, and dividing by the respective resource quantity. To find the cost of a service change, these cost rates are simply multiplied by the net change in each respective resource and then summed. The method is easy to understand, calibrate, and apply, but it may underestimate or overestimate key cost elements such as driver wages. This error will increase in magnitude to the degree that the service change characteristics differ from the system average characteristics.

Scheduling-Based Cost Model

As part of the techniques test, the consultant, in conjunction with MTC staff, developed a scheduling-based cost model against which all other models were tested. The model uses complete run-cut information and system cost characteristics to develop cost estimates. Regular driver costs are based on the number of runs, overtime hours, and spread premium hours as scheduled by RUCUS software (MTC's normal scheduling procedure). Extra board costs are based on tripper pay requirements as determined by MTC's manual assignment (RUCUS does not explicitly address tripper pay requirements). Driver benefits (i.e., paid leave, FICA, pension, fixed benefits) are developed based on wages, past experience, and contractual provisions. Nondriver, variable operating costs are estimated with a cost allocation approach corresponding to that used in the proposed method. The model produces a cost estimate that MTC staff and the consultant believe best reflects true cost impacts.

The scheduling-based cost model was applied to the entire division for each route change because routes are not generally scheduled independently, and routes frequently contain foreign pieces (i.e., work from another route). The incremental cost impact is determined by taking the differences between the cost estimates for before and after the change. It should be noted that, unlike the modified Adelaide model, this model need only be applied once for the before condition--assuming that all route changes occur in the same division.

Service Scenarios

Route changes were the basic unit of analysis during the techniques test. A summary of the 12 service scenarios used in the test is given in Table 1. Route changes encompassed a variety of time periods, including weekday (peak only), weekday (midday only), weekday (all day), and weekend.

The service scenarios are comprised of three basic change types: change in running time, addition or deletion of an entire route, and addition or deletion of single trips. Changes in running time may result from any number of factors including extending or shortening a route, changes in load factors, and changes in traffic conditions or controls. All of these can contribute to a change in driver and vehicle use. Additions or deletions of entire routes, or of single trips, are situations faced by transit planners in tailoring service to match new fiscal or ridership conditions. All of these conditions can contribute to changes in driver and vehicle use, with corresponding cost implications.

It should be noted that the magnitude of the service changes was generally quite small. Daily

TABLE 1 Service Scenarios Used in Model Testing

Scenario Number	Type of Change	Description	Net Change in Daily Hours	Percent of Division Hours
1	Weekday-Peak Only	Extend an express route/MTC Express Rte. 35C	2.50	0.121
2	Weekday-Peak Only	Reduce service by half/MTC Express Rte. 35C	(10.05)	0.484
3	Weekday-Peak Only	Discontinue one A.M. and one P.M. Trip MTC Express Rte. 35LU	(2.88)	0.139
4	Weekday-Peak Only	Discontinue one A.M. and one P.M. trip MTC Local Rte. 47	(2.12)	0.102
5	Weekday-Midday Only	Discontinue midday service/MTC Local Rte. 9	(32.70)	1.576
6	Weekday-Midday Only	Reduce midday service by half/MTC Local Rte. 2	(6.67)	0.322
7	Weekday-All Day	Double midday service/MTC Local Rte. 2	22.27	1.074
8	Weekday-All Day	Discontinue weekday service/MTC Local Rte. 47	(63.35)	3.054
9	Weekday-All Day	Reduce service on express route MTC Express Rte. 52B	(15.68)	0.756
10	Weekend	Discontinue Saturday service/MTC Local Rte. 21	(116.22)	1.142
11	Weekend	Discontinue Sunday service/MTC Local Rte. 21	(79.42)	0.871
12	Weekend	Discontinue six Sunday trips/MTC Local Rte. 9	(4.03)	0.044

changes in vehicle-hours range from 2.12 to 116.22 hr. This translates to a range of less than 0.05 percent of division hours to about 3 percent of total hours. Specifically, eight of the scenarios represent a change of less than 1 percent of total division hours, and four scenarios represent changes of between 1 and 3 percent of service hours.

EVALUATION OF TEST RESULTS

Evaluation of the test results was an interpretive process based on simple statistical measures. An important consideration throughout the evaluation was the limited sample size. The test entailed application of five cost models to a total of 12 route-level service changes. All of the service changes occurred at one transit system and within a single operating division. The sample size suggests that the test results may not reflect actual model capabilities under all circumstances. The test is intended to provide an indication of how well these cost models perform against one another under field conditions, and does not preclude additional testing.

The incremental cost estimates produced in the techniques test are given in Table 2. Using the scheduling-based cost estimate as a reference, the annualized incremental cost impacts of the 12 service changes ranged from $4,893 (Scenario 12) to $780,897 (Scenario 8). These cost estimates formed the basis for the evaluation of model performance.

The evaluation of test results focused on three primary areas of concern: model accuracy, model sensitivity, and level of effort. Each of these issues is discussed further hereafter.

Model Accuracy

A significant limitation is placed on this element of the test because a true cost value does not exist; that is, it is practically impossible to determine the actual cost attributable to a change in service. For the purposes of this test, the scheduling-based cost model estimate serves as the yardstick against which the other models are measured. This comparison provides a reasonable indication of relative model accuracy.

Relative model accuracy was examined using the percentage of deviation from the scheduling-based

TABLE 2 Annualized Incremental Cost Estimates

	Incremental Cost Model				
Scenario	Scheduling-Based	Proposed	Adelaide	Peak/Base	Cost Allocation
1	$ 25,584	$ 15,692	$ 47,496	$ 16,870	$ 16,532
2	(145,652)	(93,116)	(104,096)	(96,541)	(95,193)
3	(29,902)	(30,239)	(17,019)	(29,537)	(29,092)
4	(22,019)	(22,265)	(36,809)	(20,571)	(20,271)
5	(229,310)	(267,770)	(239,678)	(271,132)	(275,828)
6	(61,597)	(71,799)	(58,023)	(51,321)	(52,285)
7	178,627	155,202	154,671	166,393	183,696
8	(780,897)	(701,155)	(785,099)	(622,610)	(661,758)
9	(209,973)	(175,949)	(154,298)	(173,694)	(173,738)
10	(233,830)	(225,595)	(234,569)	(199,980)	(203,405)
11	(178,704)	(175,187)	(200,180)	(151,820)	(154,430)
12	(4,893)	(8,983)	(4,305)	(7,438)	(7,570)

cost estimate as the primary measure. This method for evaluating relative model accuracy examines the percentage difference between each model's results and the best estimate. This measure can be examined from several different perspectives, including

- By type of service change and
- By size of service change.

Each of these measures, as applied in the techniques test, is further discussed.

Percent Deviation by Type of Service Change

This measure examines the percentage of deviation from the scheduling-based cost estimate by type of scenario for each model (Table 3). All models show substantial variability in their performance for peak period changes, with the modified Adelaide model experiencing the greatest overall deviation in the test. It should be noted that the proposed method comes within 1 percent of the scheduling-based cost estimate two out of four times. Each of the models appears more stable on midday only and all-day service changes. The weekend changes offer a mixed bag of performance--all models perform reasonably well on two of three changes. The models perform poorly on one scenario that entails a change of less than 1/20th of 1 percent of total division hours.

Percent Deviation by Magnitude of Change

Another way to view model accuracy is based on how well it performs in relationship to the magnitude of the service change. This relationship is shown in Figures 1-4 for the proposed method, modified Adelaide model, peak and base model, and cost allocation model, respectively. One preeminent trend is evident for each model--overall accuracy improves with increases in the magnitude of the service change. This trend is most pronounced in the proposed method and the modified Adelaide model. The trend also exists in the peak and base and cost allocation models, although to a smaller degree. These two models show lower variability, and they

TABLE 3 Percentage of Deviation from Best Cost Estimate

Scenario	Type of Change	Proposed Method	Adelaide Model	Peak/Base Model	Cost Allocation Model
1	Weekday-Peak	(39)	86	(36)	(35)
2	Weekday-Peak	(36)	(29)	(34)	(35)
3	Weekday-Peak	1	(43)	(1)	(3)
4	Weekday-Peak	1	67	(7)	(8)
5	Weekday-Midday	17	5	18	20
6	Weekday-Midday	17	(6)	(17)	(15)
7	Weekday-All Day	(13)	(13)	(7)	3
8	Weekday-All Day	(9)	1	(15)	(15)
9	Weekday-All Day	(16)	(27)	(17)	(17)
10	Weekend	(4)	1	(14)	(13)
11	Weekend	(2)	12	(15)	(14)
12	Weekend	84	(12)	52	55

( ) Denotes percent underestimated

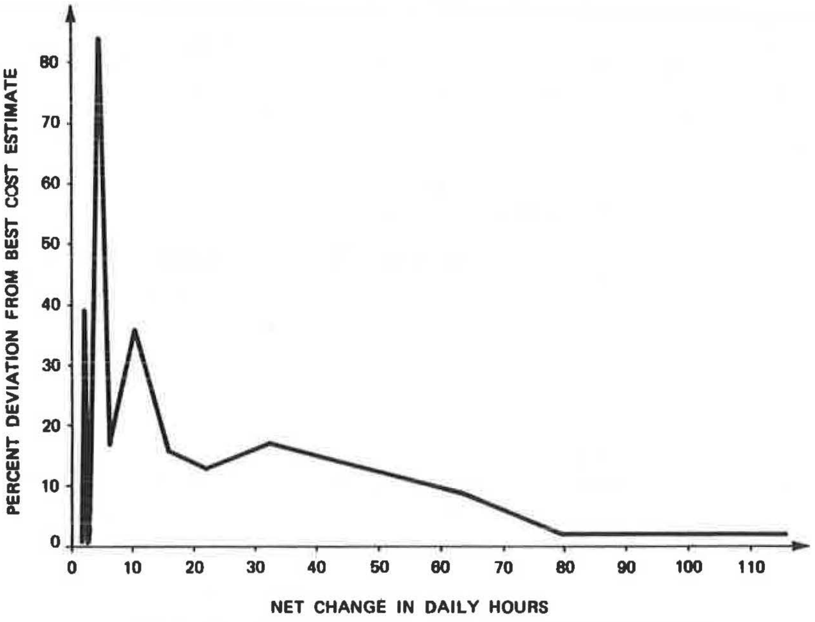

FIGURE 1 Relationship of model accuracy to magnitude of service change—proposed method.

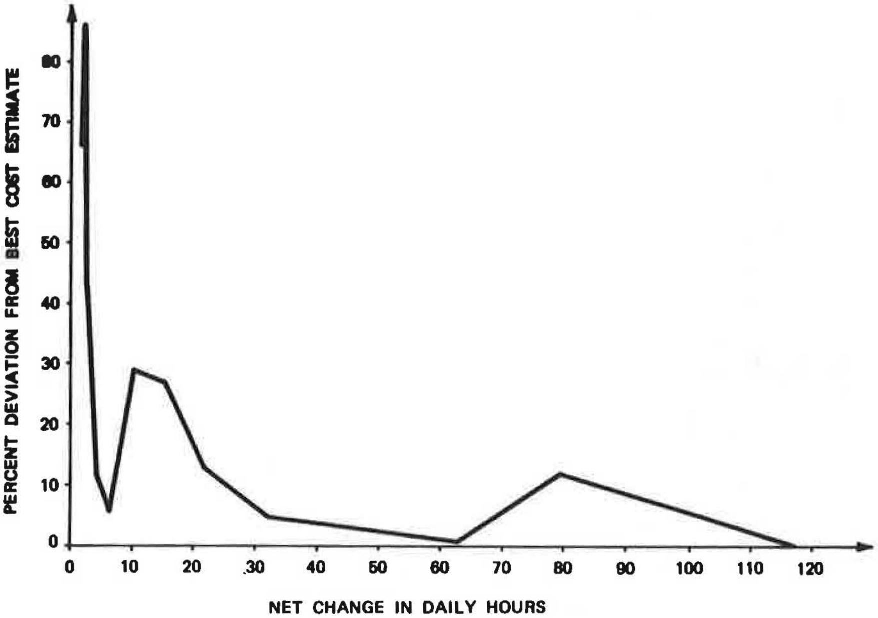

FIGURE 2 Relationship of model accuracy to magnitude of service change—modified Adelaide model.

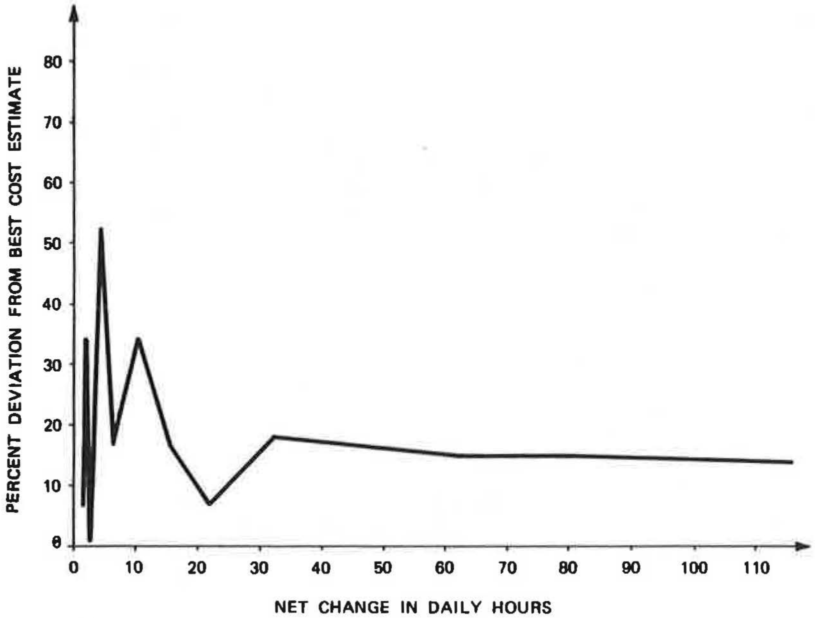

FIGURE 3 Relationship of model accuracy to magnitude of service change—peak and base model.

also exhibit greater deviation from the scheduling-based estimates overall.

Although each of these measures examines model accuracy in a different light, one common theme appears to prevail: All of the models tested exhibit high variability in their ability to replicate scheduling-based cost estimates for minute service changes (i.e., less than 1 percent of division hours). Overall performance and consistency improve with increases in the magnitude of the service change. Although the proposed method and the modified Adelaide model appear highly accurate on many individual scenarios, they are also quite inaccurate on several scenarios. In comparison, the peak and base and cost allocation models tend to be less accurate overall, but they also show less variability in their performance.

Model Sensitivity

The sensitivity of a model was measured by examining the amount of variation in its unit cost estimates of different service changes. It has been assumed that the unit cost (i.e., cost per hour) of a service change should be variable, reflecting the differential cost impacts of particular service characteristics (e.g., peak-only service versus weekend service, tripper versus regular runs, express versus local service). When the unit costs produced by a

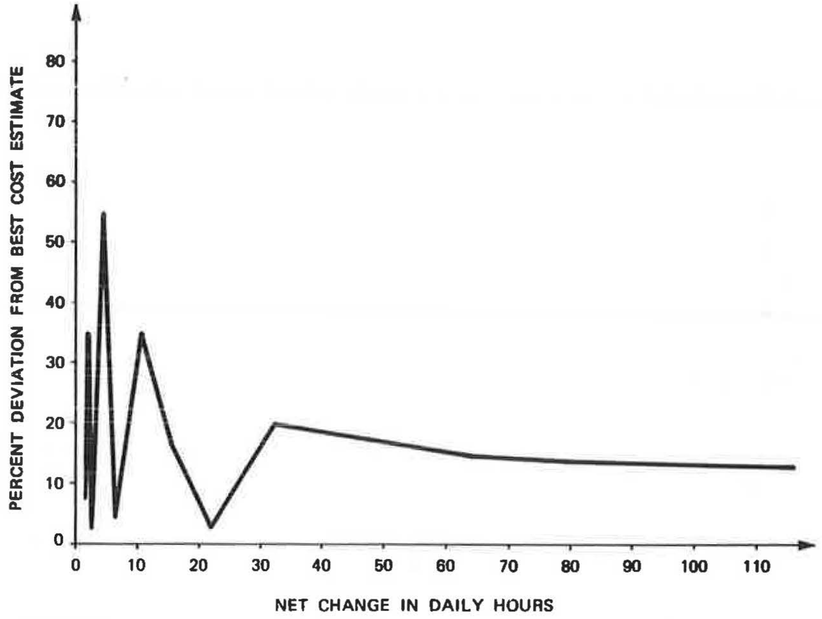

FIGURE 4 Relationship of model accuracy to magnitude of service change—cost allocation model.

model are relatively constant over a series of service changes, the model's sensitivity may be questioned.

Evaluation of model sensitivity is an interpretive process. The question of why one model shows greater variance than another in unit costs must be considered. This is essential because variation may be due to inaccurate cost estimates as well as to sensitivity to particular conditions. In the test, this analysis was conducted by contrasting sensitivity findings with accuracy results by model.

The primary measure of model sensitivity in this test is the coefficient of variation (CV). The CV is calculated as the ratio of a model's standard deviation to its mean unit cost (i.e., cost per hour) for all the test situations. A low CV value indicates that all unit costs are grouped in a narrow band around the mean. A high CV value indicates a high degree of variance from the mean unit cost.

The results of this test, given in Table 4, indicate that the proposed method comes closest to replicating both the mean hourly cost and the CV produced by the scheduling-based cost model. The modified Adelaide model produces the highest CV value, which exceeds the scheduling-based cost model's coefficient by more than 60 percent. This is partially attributable to the inaccurate cost estimates produced by the modified Adelaide model and demonstrated in the accuracy evaluation. The peak and base and cost allocation models produced coefficient of variance values only slightly lower than that of the proposed method. However, the mean unit cost estimate from these two models was significantly below the scheduling-based estimate.

TABLE 4 Model Sensitivity Measured by the Coefficient of Variance

	Aggregate for 12 Scenarios		
Model	Mean Cost* Per Hour	Standard Deviation	Coefficient of Variance
Scheduling-Based	$39.4	$10.1	0.26
Proposed Method	37.2	6.2	0.17
Modified Adelaide Model	41.0	17..2	0.42
Peak/Base Model	34.7	5.2	0.15
Cost Allocation Model	34.6	5.2	0.15

* Incremental Cost

Level of Effort

Another important consideration in model evaluation is the level of effort required to produce an incremental cost estimate. The level of effort needed to use a particular model falls into two categories, calibration and application. Each model must be calibrated before actual use, with the level of effort proportional to the number of steps required and the amount of time required. When applied, the level of effort is primarily a function of the data inputs required and the time needed to apply the algorithms.

Model Calibration

The purpose of model calibration is to prepare the costing method for application to route change scenarios. Calibration requires three basic activities: data collection, data processing, and calculation of unit costs and coefficients. A primary source of data for each of the models included in this test is the Section 15 accounting report. Several of the test methods required additional information about driver assignments, which was obtained from assignment and dispatching data.

Because each of the test models uses commonly available data for calibration, the level of effort required in this activity was examined in terms of relative time to calibrate each model. In the techniques test, the proposed method required the greatest amount of time for calibration--about 24 working hours. The modified Adelaide model also required a high level of effort, about 18 hours total. The peak and base model, as calibrated at the test site, required expenditure of 10 person-hours. Each of these models required analysis of scheduling practices, which accounted for a large portion of the time expended. The remaining two models, the scheduling-based cost and the cost allocation models, required 3 hr and 1 hr, respectively.

Although the level of effort required in model calibration varies significantly, it should be noted that calibration occurs only once a year. The effort expended in the calibration phase can be spread over the number of times the technique is applied, thus lowering its burden significantly. Because calibration data are generally valid for a year, the level of effort required for model application may be of greater concern to potential users.

Model Application

The level of effort required in applying costing techniques is primarily a functon of the data input needs and the time required to complete application algorithms. Each of these elements is discussed hereafter.

Data Requirements for Application

Each model's application can be classified according to the amount of data needed to apply the model. The two-variable cost allocation model can be applied when only the magnitude of the service change (i.e., net change in miles and hours) is known. The peak and base model and the proposed method require not only the scale of the change but the span as well (i.e., time periods in which change occurs). The modified Adelaide model estimates incremental cost based on changes in headways, round-trip time, platform hours, and vehicle-miles--thus application occurs later in the planning process. The scheduling-based cost model uses complete run-cut information to predict cost implications and, therefore, can only be applied after scheduling is completed.

Relative Time Requirements for Model Application

Experience in the techniques test indicated that the cost allocation and the peak and base models required the lowest level of effort, with complete application averaging about 5 to 10 min per scenario. The proposed method required between 35 and 50 min for application; the major time driver is whether the change occurs on a weekday or weekend schedule. Application of the modified Adelaide model averaged 30 to 40 min when the average headways and round-trip time had been determined for the five time periods. Establishment of average headways and round-trip times for before and after the change increases the expended time by 1 to 3 hr. Thus, total application time, relative to the other models, is generally 1.5 to 3.5 hr per scenario.

Application of the scheduling-based cost model averaged between 30 and 40 min per scenario after all required scheduling data were made available. To make a valid comparison with the other models, scheduling time must be included. At the test site, legal runs are scheduled at the division level using RUCUS software, and trippers are scheduled manually. The entire process took 14 to 16 hr per scenario. This brings the total application time up to 14.5 to 16.5 hr per service change--a level of effort not likely to be expended in the planning phase of service development.

Although the actual level of effort required to apply these models at other properties may vary somewhat with data processing capabilities, the overall relationship of the models should remain constant.

CONCLUSIONS

Although the techniques test was limited in scope and an actual value for incremental costs was not available, several interesting conclusions are suggested by the study results. First, none of the costing techniques appear to be consistently accurate for extremely small service changes (i.e., less than 1 percent of total division hours). Each of the cost models experiences high variability and substantial deviation from the best cost estimate for service changes of this magnitude. Second, the size of the estimated cost implications of minute service changes (i.e., less than 1 percent) is so small that transit properties may not wish to expend the resources necessary to estimate these costs. It may be more productive to focus service planning resources on more substantive, although still small, service changes.

The proposed method and the modified Adelaide model were the best overall performers in this test. The proposed method performed best in some measures of accuracy (i.e., aggregate rank, percentage deviation from scheduling-based estimates) and in measures of sensitivity (i.e., mean unit cost and coefficient of variation). The proposed method performs reasonably well in the level of effort category. The modified Adelaide model also performs well in measures of accuracy (i.e., magnitude of deviation and percentage deviation from the best estimates), but it is not as good a performer in sensitivity measures or level of effort. Although these two models frequently come closest to replicating the scheduling-based cost model, they also incur substantial variations from the best estimates in several instances (e.g., minute service changes). They are both deterministic models and are generally sensitive to cost differentials of a variety of service changes.

The peak and base model and the cost allocation model exhibit similar performance trends in all evaluation categories. Each of these models experiences less variability, even in minute changes, than do the proposed method and the modified Adelaide model. This is chiefly attributable to their average costing algorithms, which are less sensitive to the cost differentials of a variety of service changes. Some degree of sensitivity is surrendered, but simplicity in model application is maintained. On the larger of the small service changes (i.e., 1 to 3 percent of division service), these models are the poorest performers.

The techniques test provides some insights into the relative performance of each of the subject models. Both the proposed method and the modified Adelaide model represent more sophisticated attempts to simulate the complex factors driving incremental driver wage and benefits cost. The peak and base and cost allocation models, on the other hand, use a statistical approach in which systemwide average characteristics determine the extent to which incremental costs are affected. The accuracy, sensitivity, and ease of use evaluation measures suggest that no single model is preferred for all situations.

The use of a particular model would be a function of the extent of the service changes and the use of the cost estimates. For example, investigation of the cost consequences of a relatively minor service change would suggest the use of the peak and base or cost allocation model. The increased sensitivity and complexity of the other procedures do not appear to increase relative model accuracy for minute bus service modifications. This may be attributable to the numerous intermediate solutions (e.g., number of trippers and drivers) possible. For more substantial service changes, the proposed and modified Adelaide models may be preferred. In cases in which the cost impacts are expected to be relatively high, use of a more rigorous, and potentially more accurate, evaluation tool may be warranted.

Another issue related to the selection of an incremental costing procedure is the intended use of the resulting cost estimates. For a preliminary investigation of a wide range of bus service options, the simplistic techniques may be appropriate. In this case, the resources required to apply the technique would not unduly constrain the number of service changes that could be investigated. If a relatively limited number of changes were considered for implementation, a more accurate, but more time-consuming, model might be appropriate. Such an approach is consistent with other transportation analyses in which sketch-planning techniques are applied initially to screen a large field of options and then followed by more rigorous and detailed procedures for the most promising scenarios.

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

Estimating Bus Ridership

HERBERT S. LEVINSON and ORIKAYE BROWN-WEST

ABSTRACT

A route-based approach to estimating bus ridership is described. Bus riders on a typical Hartford, Connecticut, route were classed by walking distance and car ownership status, and compared to the number of dwelling units in each status. A series of "ridership penetration curves" shows how the number of bus riders per dwelling unit relates to car ownership and walking distance. For each level of car ownership, these curves decline with increasing walking distance. They show a drop of about 5 rides per 100 dwelling units for every 100-ft increase in walking distance, and a decrease of about 10 rides per 100 dwelling units when the first car is acquired and again when the second car is acquired. The data appear consistent with patterns derived from origin-destination surveys.

Public transportation planning and operation in today's urban environment increasingly concentrate on adjustments to existing services. They emphasize ways to increase transit service efficiency and to reduce operating deficits, instead of trying to assess impacts of large-scale investments. This involves adapting service to changing ridership patterns and cutting or restructuring service to bring costs and revenues into better balance. It calls for route-sensitive ridership estimation techniques that are keyed to fine-tuned service changes.

Much work, of course, has been done on estimating transit ridership. There is an extensive literature on network-based modal split models keyed to the relative disutilities of car and bus travel (i.e., logit modal split models). At the other end of the spectrum there is a growing body of elasticity factors that are keyed to service frequencies, fares, and travel times. Neither of these techniques properly addresses the question: If a new route is extended into a residential suburb, how many riders will it attract? Conversely, if a route is cut back, what will be the net loss in patronage?

Most of the current ridership estimation techniques are either too complex or too general to provide timely and meaningful responses to these fine-grained service changes. For these reasons, simplified and reliable estimating techniques that can be applied at the route level remain an important research need (1).

RESEARCH APPROACH

The results of a ridership research study conducted during 1981 and 1982 for the Connecticut Department of Transportation are discussed. The research objective was to develop a method for quickly estimating the ridership impacts of bus service changes in Connecticut cities (2).

Pilot surveys were conducted on six Hartford bus lines in June 1981 to identify parameters and refine the research approach. This was followed by a resurvey of riders on Line U-3 in June 1982. The ridership data for U-3 were compared with 1980 census data to obtain "penetration ratios" (ridership rates) by car ownership and walking distance strata. Finally, comparisons, applications, and extensions of the research were developed.

The research approach is shown in Figure 1. Detailed steps were as follows:

1. June 1981 on-board surveys obtained the travel patterns of 1,224 inbound riders out of a total of 21,130 weekday (two-direction) riders. The