Exploring the Multiple Factor Concept for Bus Maintenance Using Simulation

GEORGE LIST, MYSORE SATISH, and MARK LOWEN

ABSTRACT

The transit industry has clearly shifted to an emphasis on fleet maintenance, with operators trying to improve their control of this activity by using tools such as maintenance management information systems. One advantage of these systems is their ability to put within reach a wide range of scheduling rules or algorithms. Explored in this paper are the benefits of a scheduling rule that relies on more than one independent variable or factor. It is based on a premise that the failure distributions of vehicle components are functions of different factors. Alternatively, the components are sensitive to different measures of use. The benefits are clear. For systems where buses accumulate use at widely varying rates, one factor to another, or where the services are in a state of flux, multiple factor control provides much lower in-service failure rates than does single factor control. Moreover, sensitivity analyses indicate that the extent of these benefits is dependent on whether on-condition or planned replacement is employed and whether the component failure distributions are normal or exponential.

The transit industry has clearly shifted to an emphasis on fleet maintenance as a result of the recent federal austerity and state and local government belt tightening. Moreover, it appears that many operators are striving to improve their maintenance practices as well, as evidenced by the popularity of recent bus maintenance workshops.

Operators are searching for better maintenance procedures, up-to-date training aids, solutions to specific problems, and better ways to manage the overall maintenance process--especially ways that take advantage of computerized tools such as maintenance management information systems (MMISs). One advantage of an MMIS is its ability to put within reach a great number of maintenance activity scheduling rules or algorithms.

Although maintenance managers have previously had to rely typically on just one factor for practicality, an MMIS allows them to specify more sophisticated algorithms based on several factors, such as oil analysis results, hours, and stops, in addition to miles. But this raises the question as to whether such sophistication has significant value and, if so, when. This question is examined in this paper by analyzing the value of multiple factor control in situations where it is likely to prove useful, such as systems whose routes are different from one another (e.g., in terms of average speed or stopping frequency) or systems whose routes are in a state of flux (e.g., expanding or contracting).

THE MULTIPLE FACTOR CONCEPT

The multiple factor concept states that the failure distributions of the vehicle's components may be functions of different independent variables or factors. Alternatively, its components are sensitive to different measures of use. For example, lights and other electrical equipment may be sensitive to hours of use, air conditioners may be sensitive to equivalent full-load hours, and brakes may be sensitive to the number of stops made. Hence, the vehicle's overall reliability is a function of a vector of factors, not just one. The concept also states that these factors may themselves be functions of other factors (e.g., engine wear may be a function of both miles and hours).

Under these conditions, a maintenance program based strictly on one factor will have significant shortcomings compared to one that uses a vector of factors unless the buses "age" at proportional rates for all factors in the vector. If, for example, the buses accumulate mileage at rates that vary widely from one bus to another, even though they all operate the same number of hours per day, then it will be important to include bus-hours along with bus-miles in the vector of factors.

EVIDENCE OF THE USE OF MULTIPLE FACTORS

The concept of using multiple factors was, at one time, quite popular and it is under consideration again today (1). Evidence of attempts to use the multiple factors concept can be found in the recent Bus Maintenance Workshop Proceedings (2):

- In Syracuse, New York, an inspection program is used that combines mileage and hourly factors;
- In Los Angeles, California, it is predicted that on-board electronics will necessitate better monitoring of bus hours;
- In San Antonio, Texas, it is preferable to schedule engine maintenance based on hours although there is a lack of confidence in hour meters; consequently, mileage is used;
- In Cleveland, Ohio, hours are used (instead of miles) to schedule the city's preventive maintenance.

Moreover, a recent TRB-sponsored study found that some bus property authorities think fuel consumption

Department of Civil Engineering, Rensselaer Polytechnic Institute, Troy, N.Y. 12180.

could be useful as a basis for specifying engine, and perhaps other, component maintenance (3).

ANALYSIS OPTIONS

Clearly, empirical data should be used for analyses whenever possible because actual situations tend to have features that model builders cannot, or fail to, account for, a criticism that can be lodged against the work presented here as well. But researchers have discovered that cross-sectional or time-series maintenance data are difficult to obtain. It seems the industry simply has not computerized its maintenance data and is only now making progress in that direction, which is due, in part, to the microcomputer (4). That bus maintenance data can be used to investigate specific issues has been illustrated by Maze, Dutta, and Kutsal, who sought to determine whether a technological fix to a transmission problem produced any quantifiable improvement in reliability (5).

Section 15 data seem to be of some use, but Fielding, Babitsky, and Brenner clearly showed that many maintenance-related data items have either missing or ambiguous entries (6). In their study, road calls had to be dropped from the analysis because the variable's definition led to inconsistent entries, active vehicle count-related entries had to be deleted because they were ambiguous, and fuel had to be dropped because there was no obvious way to combine the data for the four different types of fuel in use. In a separate Section 15-based study, Foerster, Miller, Kosinski, and Rueda could not obtain a coefficient of determination (R^2) greater than 0.04 in their maintenance-oriented regression analyses (7).

Under these conditions, simulations can often be used to generate synthetic data. For example, Dutta (8) developed a simulation model, including resource allocation suboptimization routines, that allows for experimentation with radically different bus maintenance strategies. Maze, Dutta, and Kutsal (9) illustrated the potential problem of maintenance demand peaking that can occur when all new buses are purchased. Muthukumaran, Miller, and Foerster (10) used MASSTRAM (11) to study optimal maintenance planning, and Sinha and Guenther (12) combined a maintenance planning model with an operations model using a dependability factor to study the impacts of maintenance strategies on service reliability.

For purposes of this analysis, however, although each of these studies approaches the maintenance planning problem with a different methodological framework, they have one significant aspect in common: they use just one variable, mileage, to determine when a vehicle is going to fail and when it should be scheduled for maintenance. Hence, it was necessary to develop a simulation model that incorporated this feature.

THE MULTIFACTOR MODEL

The multifactor bus maintenance model (13) provides a simple representation of a transit system's operation, moving buses from one stage to another in a four-stage system as shown in Figure 1. The stages are as follows:

• In storage, which is either overnight or as a service spare;
• In service, which is differentiated by type (e.g., urban, suburban, or express);
• Awaiting maintenance, which can be repair (high priority) or inspection (low priority);
• In the shop (in the repair facility).

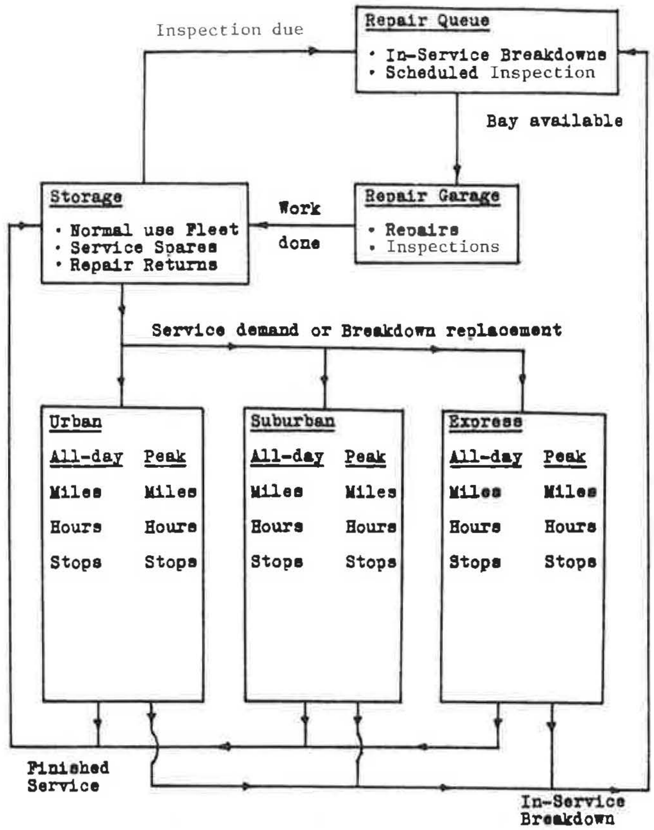

FIGURE 1 Structure of the simulation model.

INITIALIZATION

At the beginning of simulation, each bus is placed In Storage. Its components are assigned "times" to failure (e.g., miles, hours, or stops) by drawing values from failure distributions that use the appropriate causal factors as the independent variable. For example, if component X is sensitive to miles, it samples a distribution relating the probability of failure to the number of miles traveled.

MOVEMENT BETWEEN STAGES

On each day of simulation, buses In Storage are placed In Service by filling route assignments on a random basis. After this process is complete, any buses still In Storage are left In Storage as spares for future assignment.

Routes can be of various types (e.g., urban, suburban, or express) and can have different durations (e.g., peak or all day). Moreover, each one (e.g., an all-day urban route) has a set of probability density functions for the causal factors (e.g., distributions for bus miles, hours, and stops per day). Hence, when a bus is assigned to a given route, these distributions are sampled to obtain a "use" vector for the day (e.g., the day's incremental miles, hours, and stops).

If the day's use vector will push one of the components past its point of failure, the bus will have a breakdown while In Service. The bus accumulates a percentage of the use vector proportional to the component's point of failure, leaves the assigned route, records the failure as being peak or off-peak depending on when it occurred, and goes to the Awaiting Maintenance stage.

Replacement buses are dispatched from In Storage to fill vacancies created by the in-service failures. Each one accumulates the remainder of the failed bus's use vector plus an increment to reflect travel (e.g., miles, hours, and stops) to the point where its service starts. If no replacement bus is available, the model records the lost hours of service, peak or off-peak.

At the end of each day, buses In Service return to storage except when they are due for inspection, in which case, they go to the Awaiting Maintenance stage. Buses Awaiting Maintenance sit in queue until it is their turn to occupy a bay in the repair facility. A bus needing repair has priority over one scheduled for inspection, and within each of these categories, buses are sequenced according to the time when they joined the queue.

Once a bus is In The Shop, it is either repaired or inspected as appropriate. If it is to be fixed, a repair time distribution is sampled for the component being replaced to determine how long it will be In The Shop. Once this time has elapsed, the bus leaves the shop, releases the facility capacity it had employed, and returns to storage, to await its next service assignment.

If the bus is In The Shop for a component inspection, a test is performed to see whether the component is still serviceable or needs replacement. Ideally, this test would be based on the probability that the component shows significant wear given its present age plus a conditional probability that the inspecting mechanic will decide to replace the component given this information. As a simple approximation, the model assumes that at a given point in time, expressed as a percent of the component's time to failure (e.g., 85 percent of its life, measured on the basis of the factor that dictates failure), it will be obvious that the component needs to be replaced. Hence, if the component's percent of time to failure is beyond this point (e.g., less than 15 percent of its life remaining), the bus will be shopped (i.e., put in the Awaiting Maintenance queue) so that the component can be replaced; otherwise, it will be returned to the In Storage stage to await its next service assignment.

THE BENEFITS OF USING MULTIPLE FACTORS

To investigate the benefits of multiple factor control, a hypothetical transit system was developed that was assumed to have the following characteristics.

Its buses have three components, the first of which has a failure distribution dependent on miles; the second, hours; and the third, stops. Corresponding mean times to failure are 50,000 mi, 3,000 hr, and 200,000 stops, and the failure distributions are normal with a standard deviation equal to 20 percent of the mean (see later text for sensitivity analyses regarding these assumptions). The numbers are intended to represent engine-transmission combinations, air conditioners, and brake systems, but there is no claim that the numbers are representative of any specific system. [Note that Foerster et al. (7) did develop such statistics for several components based on miles, and the statistics used here are loosely related to these.]

It is also assumed that three types of routes are being operated: urban, suburban, and express. The urban routes have 10 stops per mile and an average speed (V_{avg}) of 10 mph; the suburban, 1 stop per mile and V_{avg} = 20 mph; and the express, one stop every 10 mi and V_{avg} = 30 mph. Fifty percent of the routes are urban, 30 percent are suburban, and 20 percent are express. For all three types of routes, one-third of the assignments are all-day (16 hr) and the remaining two-thirds are peak (3 hr in the morning and 3 hr in the evening).

The maintenance schedule is based on planned replacement, with an assumed MMIS being used to schedule buses for component change-outs at mileages predicated on the last change-out. For example, the change-out interval for the hours-sensitive component might be set to 40,000 mi. Every time the component fails or is changed-out, the mileage counter is reset, so that the next change-out will be scheduled for precisely 40,000 mi after the preceding one. Figure 2 shows that this minimizes the number of component replacements required while still meeting a given in-service failure rate goal.

The main question is whether a multifactor strategy would offer significant advantages over the present strategy. Consider the situation where service cutbacks are planned in the near future because of fiscal constraints. Assume one of two scenarios is most probable. Either the suburban and urban services will be retained (Scenario A) or only the urban service will be kept (Scenario B), as shown in Table 1.

The maintenance problem under both scenarios is to keep the in-service failure rate under control (e.g., below 20 percent) in spite of the drastic changes in service. This goal is difficult to achieve because in both scenarios buses will be

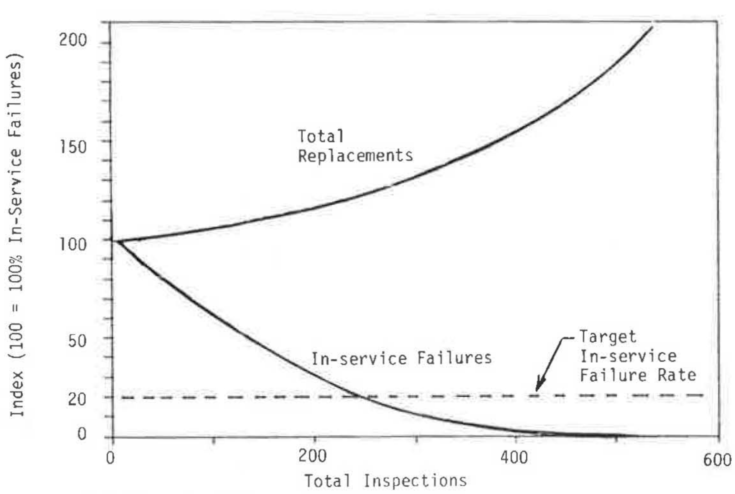

FIGURE 2 In-service failures and total replacements versus total inspections.

TABLE 1 Operating Environment by Scenario

Characteristic	Base Case	Scenario A	Scenario B
Average speed (mph)	17	13	10
Distribution of routes (%)			
Urban[a,b]	50	70	100
Suburban[c,b]	30	30	-
Express[d,b]	20	-	-
Hours of service			
Peak periods	6	6	6
All day	16	16	16

[a]10 mph average speed, 10 stops per mile.
[b]Two-thirds peak hour buses, one third all day.
[c]20 mph average speed, one stop per mile.
[d]30 mph average speed, one stop every 10 miles.

accumulating hours and stops at faster rates per bus-mile than they are presently. Figure 3 shows that while the in-service failure rate in the base case is 20 percent for the hours-sensitive component (using a 40,000-mi change-out interval), it is 70 percent in Scenario A and 90 percent in Scenario B.

One potential solution is to identify a new change-out interval for each scenario. To stay at 20 percent in-service failures, Figure 3 shows that the interval should be set to 32,000 mi for Scenario A and 25,000 mi for Scenario B. But the problem is that this means a different change-out interval for each scenario, and new change-out intervals if other scenarios unfold.

However, multifactor control produces much better results. As Figure 4 shows, a change-out interval of 2,400 hr yields failure rates under 20 percent for all three scenarios, meaning the mix of services can change constantly and yet the in-service failure rate will remain under control.

SENSITIVITY ANALYSES

A number of key questions can be asked about how sensitive the findings are to the underlying assumptions. Most important, the questions deal with the failure distribution (e.g., type mean and variance) and the relative merits of planned change-outs versus on-condition replacements [see Etschmaier (14) for a discussion of the relative merits of these two strategies]. The critical thing to focus on is the relationship between the maintenance interval

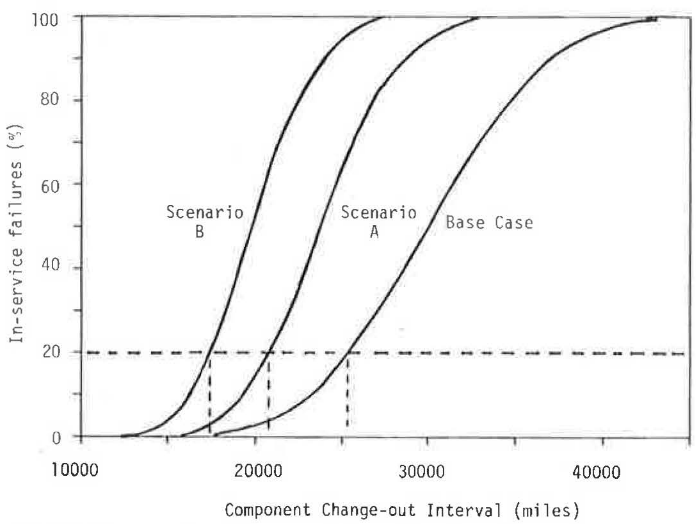

FIGURE 3 In-service failure trends for mileage-based component change-outs and time-dependent, normally distributed failure intervals.

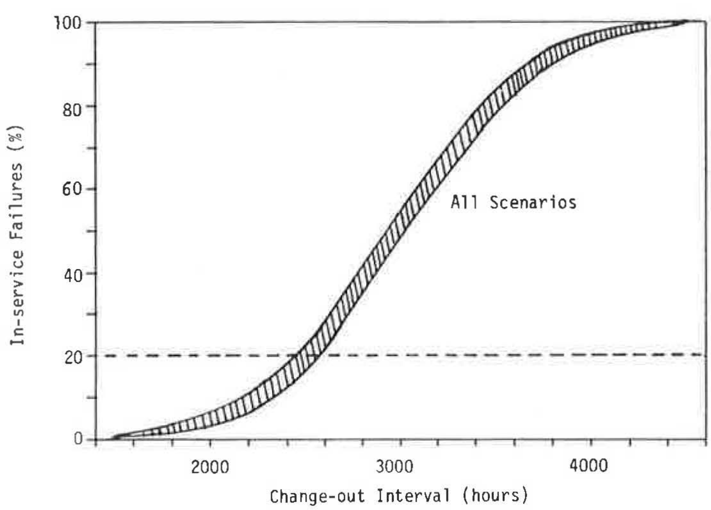

FIGURE 4 In-service failure trends for hours-based component change-outs and time-dependent, normally distributed failure intervals.

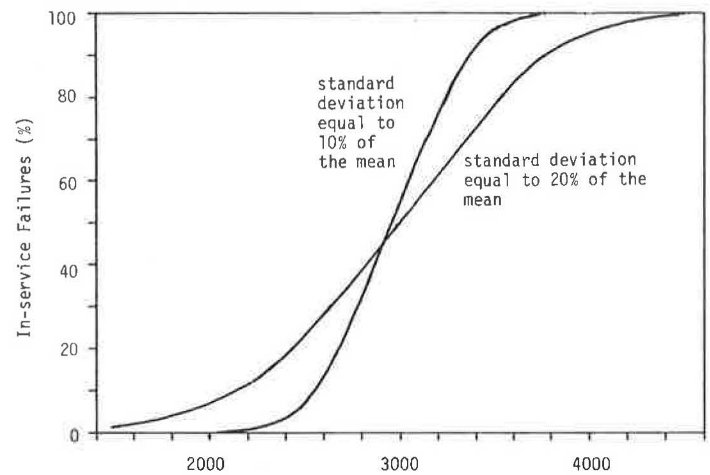

FIGURE 5 Changes in the in-service failure trends due to a change in the standard deviation of the failure interval distribution—time-dependent, normal distribution.

(planned replacement or inspection) and the in-service failure rate. As Figures 3 and 4 show, the key attributes are (a) the shape of the relationship and (b) the range of maintenance intervals over which the in-service failure rate undergoes significant change. Using Figure 4 as an example, the in-service failure rate increases monotonically as the maintenance interval widens, and the failure rate undergoes its significant change as the maintenance interval rises from 1,000 to 4,000 hr.

SENSITIVITY ANALYSES USING A NORMAL-BASED FAILURE DISTRIBUTION FUNCTION

When the times to failure follow a normal distribution, the effects of changes in mean and variance are clear. If the mean increases, the midpoint of the effective range of maintenance intervals increases but the range remains constant. For example, if, in Figure 4, the mean shifts to 4,000 hr, the curve shifts to center on 4,000 hr, but the range of effective intervals remains plus or minus 1,500 hr. If the variance increases, as shown in Figure 5, the midpoint of the range remains fixed, but the width of the range increases, proportional to the change in the standard deviation.

Understanding the effects of a shift to on-condition replacement is more complex. Remember that the model assumes there is a small window of time before failure (the near-failure window) when the component indicates replacement is required (e.g., within 15 percent of the end of its life). For the on-condition replacement strategy to be effective, the component must be inspected during this near-failure window.

Figure 6 shows that the shift in strategy yields a complex relationship between the inspection interval and the in-service failure rate. Most important, the timing of the inspections is critical. When the inspection interval is short, there is a high probability that an inspection will occur during the near-failure window and a low in-service failure rate results. As the interval widens, however, the in-service failure rate rises sharply because the last inspection before failure increasingly comes too early to be useful. In fact, at slightly below two inspections per expected lifetime, the failure rate reaches a local maximum because the synchronization between inspections and

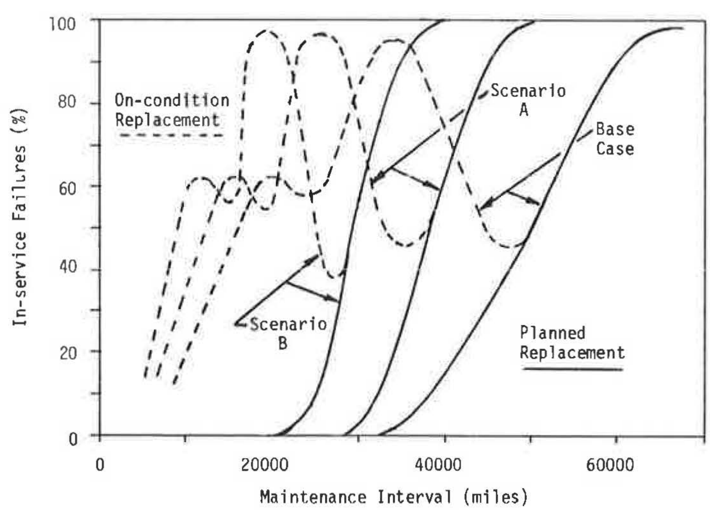

FIGURE 6 On-condition versus planned replacement (change-outs), time-dependent, normally distributed failure intervals.

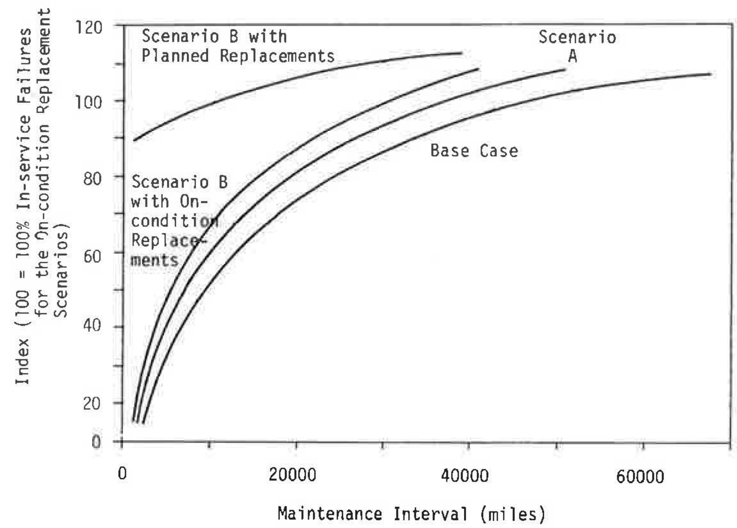

FIGURE 7 In-service trends for on-condition and planned replacement; time-dependent, exponentially distributed failure intervals.

failures is poor. After decreasing slightly at one-half the expected lifetime, the in-service failure rate rises sharply again reaching a rate as high as that encountered when no inspections are conducted because the timing problem is most acute. At first, the problem is severe because there is a low probability that any inspection will occur during the near-failure window. However, as the inspection interval widens still further, approaching the length of the expected lifetime, the in-service failure rate drops markedly because an increasing percentage of the inspections are occurring during the near-failure window. In fact, at intervals slightly smaller than the expected lifetime, there is a local minimum because the number of inspections during the near-term failure window reaches a local maximum. Once past this maximum, the length of the expected lifetime, on-condition replacement appears to be the same as planned change-out with a monotonically increasing in-service failure rate.

SENSITIVITY ANALYSES USING AN EXPONENTIAL-BASED FAILURE DISTRIBUTION FUNCTION

When the times to failure follow an exponential distribution, the effects of planned change-out and on-condition replacement are reversed. A planned change-out strategy keeps the in-service failure rate high no matter what change-out interval is selected (with similarly large total replacements) while on-condition replacement produces small in-service failure rates, provided the inspection interval is kept short relative to the expected lifetime.

As Figure 7 shows, dropping the planned change-out interval from 4,000 hr down to 125 produces only an 8 percent drop in the in-service failure rate. Moreover, although not shown in the figure, the total replacements increase almost ninefold! Switching to an on-condition replacement strategy over the same range drops the in-service failure rate from 100 percent to 13 percent. Moreover, although it is not shown in the figure, total replacements do not increase at all. The figure does show, however, that under these conditions shifts in the service characteristics of the bus system are not as critical because a short inspection interval must be used to keep the in-service failures under control in any event.

IMPLICATIONS FOR FUTURE RESEARCH

There are many implications from this research, but four seem most important. First, the industry should try to determine whether, and to what extent, factors other than mileage are critical in the failure distributions of various components. Second, when gathering historical maintenance data, analysts should strive to measure such things as bus-hours, fuel consumption, and stops, in addition to bus-miles so that these causal relationships can be identified. Third, analysts should also attempt to determine the precise nature of the failure distributions because this paper indicates that they are critical to the selection of an appropriate maintenance strategy. Finally, there is a need to explore further the issue of on-condition versus planned replacement using models such as the one that has been presented here.

REFERENCES

1. Report of the Bus Maintenance Workshop, July 22-24, 1985, Buffalo, New York. American Public Transit Association, Washington, D.C., forthcoming.
2. Report of the Bus Maintenance Workshop, March 5-7, 1985, Houston, Texas. American Public Transit Association, Washington, D.C., 1985, pp. 7-8.
3. Bus Inspection Guidelines. Draft Final Report on NCTRP Project 60-1, Topic TS-13. Giuliani and Associates, Maple Glen, Pa., undated.
4. Urban Mass Transportation, Microcomputers in Transportation: Software and Source Book. Report UMTA-URT-41-85-1. UMTA, U.S. Department of Transportation, Feb. 1985, pp. 7-24.
5. T.H. Maze, U. Dutta, and W. Swanger. Estimating Bus Component Failure Distributions from Censored Samples. Transportation Research, Vol. 18B, No. 3, June 1984, pp. 201-208.
6. G.J. Fielding, T.T. Babitsky, and M.E. Brenner. Performance Evaluation for Bus Transit. Transportation Research, Vol. 19A, No. 1, Feb., 1985, pp. 73-82.
7. J. Foerster, F.G. Miller, M. Kosinski, and A. Rueda. Management Tools for Bus Maintenance. Report DOT-I-84-42. UMTA, U.S. Department of Transportation, May 1983.

8. U. Dutta. Development of a Bus Maintenance Planning Model. Ph.D dissertation. The University of Oklahoma, Norman, May 1985.
9. T.H. Maze, U. Dutta, and M.D. Kutsal. Role of Quantitative Analysis in Bus Maintenance Planning. In Transportation Research Record 915, TRB, National Research Council, Washington, D.C., 1983, pp. 39-48.
10. N. Muthukumaran, F.G. Miller, and J.F. Foerster. A Methodology for Cost Effective Maintenance Scheduling of Transit Buses. Terotechnica, Vol. 2, 1981, pp. 289-300.
11. J.D. Herniter, S.R. Rosenthal, and V.P. Wellan. The Development of a Computer System for the Cost Effective Maintenance of Rail Equipment in Urban Mass Transit Systems. UMTA, U.S. Department of Transportation, June 1977.
12. R.P. Guenther and K.C. Sinha. Maintenance, Schedule Reliability and Transit System Performance. Transportation Research, Vol. 17A, No. 5, Sept. 1983, pp. 355-362.
13. M. Satish. Simulation of Bus Maintenance. M.S. thesis. Rensselaer Polytechnic Institute, Troy, N.Y., June 1985.
14. M.M. Etschmaier. Transit Bus Maintenance Management in the United States. Transportation Research, Vol. 19A, No. 1, Feb. 1985, pp. 1-14.

Effectiveness of Improved Repair Scheduling in the Performance of Bus Transit Maintenance

UTPAL DUTTA, T. H. MAZE, and ALLEN R. COOK

ABSTRACT

Described in this paper is a computer simulation model that is used to investigate the efficiency improvements that are possible through the scheduling of bus maintenance repairs through a maintenance shop. The scheduling rules that are investigated rank repair jobs in priority order according to the length of time the bus has been waiting for repair and the length of time the job will take. It is found that scheduling, as opposed to not scheduling, can make dramatic improvements in the maintenance system's efficiency. Further, once scheduling policies are identified that result in superior performance, it is found that these same policies are superior under a variety of system conditions. The conditions varied include the number of spare buses carried, the fleet size, the failure distribution parameters, mechanic labor availability, and the maximum length of time a bus can wait for a repair.

The general financial dilemma faced by transit operators is well documented in the literature (1-3). This condition is a result of escalating operating costs and efforts by the federal government to reduce federal operating subsidies. This financial pinch is placing pressure on members of the transit industry to strive to operate as economically as possible. Many have argued that cost efficiency gains are possible if transit agencies institute more effective fleet management principles (4-6).

The purpose of this paper is to present computer simulation experiments used to determine the potential for efficiency gains from improved fleet management policies. The policies investigated deal with the effective use of maintenance activity scheduling. The scheduling rules rank in priority order the making of corrective repairs. For example, one simple rule would be to schedule for repair first those jobs that require the fewest mechanic-hours to complete. Improved repair scheduling rules have been shown to result in better system performance for a fixed level of resources (labor, spare units, and repair facility resources) in other industries (7).

U. Dutta, Goodell-Grivas, Incorporated, 17320 West Eight Mile Road, Southfield, Mich. 48075. T.H. Maze, Oklahoma Highway and Transportation Engineering Center, The University of Oklahoma, Norman, Okla. 73019. A.R. Cook, School of Civil Engineering and Environmental Science, The University of Oklahoma, Norman, Okla. 73019.

EXPERIMENTAL APPROACH

To determine if similar efficiency gains are possible in transit bus maintenance as a result of improvements to repair scheduling, simulation experiments are conducted. Simulation allows the analyst to build a symbolic model of a system on the com-