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 computer. Once constructed, the model can be used to experiment with system changes without disrupting the real operational system. Besides not disrupting the actual system with an experiment, the simulation model has two other important advantages. First, the results are obtained quickly, perhaps within a few minutes. The same experiment with the actual system might take years before the result would be known. Second, because all of the system variables in the model are controlled, the analyst knows that the results from the experimentation were produced by the variable(s) that were manipulated. In other words, results obtained from an experiment with a real system may be affected by uncontrollable variables that change during the course of the experiment, such as the weather or a new union contract. These factors can be held constant in the computer model. Thus, a computer-based simulation model can be less disruptive, faster, and more accurate than a real-life experiment in the analysis of a complex system.

Despite simulation's many positive attributes, the user of a simulation experiment's results must recognize that most complex systems include a larger number of variables than what can be practically considered in one simulation model. Therefore, to make it economical to conduct a simulation, the analyst must limit the number of parameters used and the variables included to just those that are considered important or representative of the entire system, or both. For example, in a study of maintenance practices at the Chicago Transit Authority, Haenisch and Miller estimated that bus mechanics regularly perform 1,800 different jobs (8). If an analyst were attempting to simulate this maintenance system, it would clearly be uneconomical to model the distribution of each and every event and enter the distribution parameters into a computer simulation. However, simulation studies that use only a fraction of the system's elements in the analysis are more than sufficient for policy studies where the primary emphasis is to determine the existence of relationships and to gain inferences of their strength.

Reducing the complexity of systems down to a manageable problem leaves the results of the simulation analysis vulnerable to those who question the model's relevance because of its lack of specific details. However, the model's results should be judged with respect to whether any of the missing details would affect the validity of the relationships discovered. If the missing details do not impact the validity of the relationships, then their inclusion is not necessary at the policy analysis stage.

The first step in the experimentation is to prove that systematic scheduling of repairs, as opposed to nonsystematic repair scheduling (random scheduling), can improve the productivity of the maintenance system. The experimentations show that the efficiency gains that result from systematically scheduling repairs are quite striking. Once scheduling is proven as a robust means for improving the efficiency of the maintenance system, the next step is to search for the most effective repair policies. Eight scheduling rules are developed and tested to determine which is the most efficient on the basis of a series of performance measures. The last step in the experimentation is to investigate whether the same policies remain superior under a variety of conditions. This is done by measuring the sensitivity of the system's performance to changes in fleet size, component failure distribution patterns, number of spare buses (spare factor), and the amount of labor resources available for conducting repairs.

In this paper, only a brief description of the computer simulation model is provided. The interested reader will find a thorough description of the model elsewhere (9).

MAINTENANCE SYSTEM CHARACTERIZATION

The simulation model is structured to represent a 2-tiered maintenance system. The two-tiered system is one in which there are two levels of maintenance performed (10). Light maintenance (e.g., preventive inspection, brake overhauls, and tire maintenance) is performed at storage garages. Heavy maintenance (major corrective component overhauls) is performed at a central maintenance facility. Further, the model is restricted to experimentation with only the work flow at the central maintenance facility.

From the perspective of a storage facility (the first tier), a bus's operating status may be classified into one of several categories. For example, if a bus is due shortly for a preventive inspection, the manager can wait for a convenient time to perform the inspection without taking the bus out of service by assigning the bus to single-trip, peak-period commuting runs (tripper runs) while the maintenance manager waits for an opportunity to schedule the bus for an inspection between tripper runs. Alternatively, the bus could be taken out of service and held while it waits for an inspection, or the inspection could be deferred while the bus is scheduled for regular service. There are other possible categories of status, thus making the classification of a bus's status (from the perspective of the storage garage) a complex problem to model.

From the perspective of the central maintenance facility, categorizing status is less difficult. Because buses are generally only brought to the central facility when they require a major unit overhaul, buses within the system may be classified into one of only three categories: (a) active buses that are operative and scheduled for service, (b) spare buses that are operative but not in service, and (c) failed buses that are out of service and inoperative because of a mechanical failure. Over time, each bus will cycle among the three categories of status.

For purposes of the simulation and in relation to the central facility, the day-to-day events occurring to buses are assumed to be limited to the following scenarios:

1. An "active" bus is assigned to daily service.
2. If a bus fails while in service, it is replaced by a spare bus, if one is available.
3. A failed bus is inspected to determine the cause of the failure and, if the cause is a failure of a major component or part, then the bus is driven or towed to the central maintenance facility.
4. At the end of the day, the central maintenance shop schedules repair work for the next day on the basis of the number of failed buses, mechanic labor, and parts required as well as the availability of parts and labor.
5. The buses that are not scheduled to be repaired the next day wait in the bad order parking lot of the central maintenance facility until they can be scheduled for repair.
6. After being repaired, the bus joins either the pool of active buses or the pool of spare buses depending on the number of buses required to meet scheduled service and the number of operable buses.

Repair Scheduling Policies

The purposes of the simulation experiments are to determine: (a) whether systematic scheduling im-

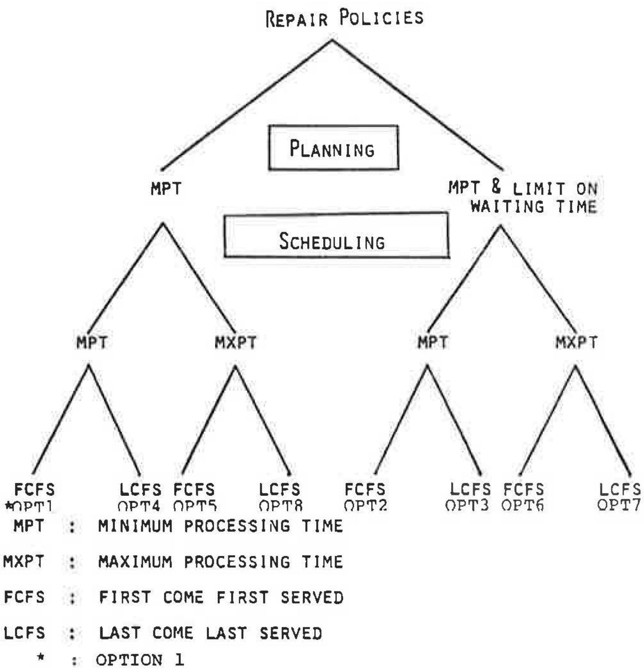

FIGURE 1 Repair tree.

proves the performance of the maintenance system, (b) which repair scheduling policies are the most effective if performance is improved, and (c) how the superiority of scheduling policies is affected by changes in the system's condition. The first step in conducting these experiments is to create scheduling rules and policies. Later, these policies will be modified to represent systems without systematic repair scheduling. A repair policy tree is shown in Figure 1. There are two steps in the repair process. These are

1. Planning. Selection of the number of repairs to be made by component type (e.g., remove and replace transmission or remove and replace air conditioning compressor) is made during the planning step. The selection process is conducted by using an optimization technique. The optimization seeks to maximize the number of repairs made with the available resources (labor and facilities). In planning, it is assumed that the length of time required to make repairs is deterministic (constant).
2. Scheduling. This step determines the execution of the planned repair work. The time required to fix a component is considered to be stochastic (variable). In other words, the time required to conduct each repair is a random variable that follows some typical distribution. Depending on the difference between the stochastic times (assumed in the scheduling step) and the deterministic times (assumed in the planning step), all planned repair activities may or may not be scheduled for repair on a particular day. If repair resources are exhausted before the completion of planned repair work, then the remaining planned repair work is cancelled. If repair resources are available after the completion of planned repair work, then additional repair work is scheduled.

Planning is the first step of the repair process and it follows one of the following two rules. These rules are identified by the upper two branches of the repair tree in Figure 1. The rules are

I. Optimization techniques are used to select the number of repairs to be conducted by repair type. The objective is to maximize the number of repairs by effectively utilizing available resources.

II. Repair those failed buses that have been waiting for a repair more than a certain number of days and utilize Rule I to allocate the remaining resources.

Once planning is completed, the next task is scheduling. The first step in the scheduling process is to determine which type of job waiting for repair (e.g., the buses waiting to have their transmissions removed and replaced) is to be scheduled for repair first. The selection of which waiting line (failed bus queue) to process first is based on either the minimum or maximum time required to complete each type of repair (processing time). In Figure 1, this is represented by the four branches of the repair tree at the scheduling level. The second and final step of scheduling is the selection of the specific bus to be repaired from the selected failed bus queue. Selection of the bus to be repaired from the queue is either first-come-first-served (FCFS) or last-come-first-served (LCFS). This is shown in the last eight branches of the repair policy tree, which defines the final eight repair policies. The repair policies are labeled as Options 1 through 8. For example, if the leftmost branches are followed through the repair tree, planning is based on minimum processing time (MPT), scheduling is based on MPT, and buses are selected from the failed bus queue based on FCFS. This combination of branches is Option 1 (OPT1).

The measures of system performance selected for determining the effectiveness of the scheduling policies are:

1. Average time spent by each bus waiting to be repaired plus the time required for the repair (time in the system = TSTS).
2. Average daily number of vehicles failed and tied-up in maintenance (total number of failed buses = TQUEUE).
3. Average number of buses in all the repair queues (WQUEUE).
4. Average mechanic overtime required per day (OTIME).

BASE CASE STUDY

The development of a simulation model requires that the model be constructed such that it depicts the characteristics of an actual system. This requires that certain assumptions be made regarding system operational procedures and parameters developed that identify the relationships between the various elements of the simulated system. Further, there may be too many possible events in real systems to economically simulate all possibilities. However, it is generally possible to include only the major events in the simulation and assume that the entire system of all possibilities would perform similarly under the same circumstances.

In the simulation's characterization, only 16 types of component or part failures are considered. These components were selected by staff members of the Detroit Department of Transportation (DDOT) as those that are the most common repairs made at their heavy-repair facility. Other assumptions made were that

- Maintenance workers are interchangeable and can perform all repairs made at the central maintenance facility.
- Repair times and miles until failure are stochastic.

- All repairs are corrective.
- Maintenance equipment and tools are always available.
- All buses are the same model.

Model Parameters

The parameters for the base case are

1. Total active fleet. In this study, the active fleet consists of 500 vehicles. This means that at the beginning of the simulation run, 500 entities are created to represent the number of buses in service. Five hundred is also a large enough number so that the sample size is great enough for any statistical test.
2. Spare factor. This is the ratio of spare buses to active buses. For the base case, the factor is assumed to be 10 percent, which is a figure reported as a level that the industry desires to achieve (11,12).
3. Available labor hours. This is the total labor hours available for daily repair. The quantity of labor hours required per day is time-dependent. In other words, the number of labor hours required to repair enough buses to meet service requirements will depend on the number and types of corrective maintenance activities required by the buses in the failed queue, which varies with the age of the buses. Early in the life of the buses, most components will be relatively reliable and, as they age, components will become less reliable and more prone to failure. In the maintenance shop, based on the composition of the failed queue, the amount of labor resources should be varied. When relatively stable (long-term) increases in failure rates occur, labor resources will be increased; similarly, they will be reduced when failure rates are low. It is found in the simulation experiments that as buses age, higher levels of failure occurrence take place after an initial break-in period [see Maze et al., for illustrations of this phenomenon (13)]. Therefore, in practice, adjustments in labor needs would not necessitate abrupt changes in the number of mechanics in the labor pool. Similarly, gradual changes in the labor pool could be obtained in an actual maintenance system through normal mechanic attrition and new hires. In this study, a simple rule is established to specify the available labor resources. When the failure rate is high, it is assumed that the available resources (in man-hours per day) is equal to a factor multiplied by the number of active buses. For example, for the base case, a factor of 0.40 is used and, because there are 500 active vehicles, 200 man-hours are available per day. During periods when failure rates are uniform, the total resource available is assumed to be 75 percent of the peak. These rules may not replicate normal staffing requirements for an actual system; however, the simulation only considers a fraction of the actual activities conducted by a maintenance facility.
4. Overtime. When the number of failed vehicles is so great that the system's ability to meet service demands is jeopardized, then overtime labor resources are used to repair failed vehicles. The use of overtime is also limited by the two following rules: (a) if the total number of failed buses exceeds the number of spare vehicles then, and only then, overtime is permitted; and (b) once overtime is permitted, it is limited to 30 percent of the regular hours if the number of failed buses is more than 15 percent of the total fleet (critical conditions); otherwise, it is limited to 25 percent of the regular hours.
5. Failure patterns. The failure patterns of 16 different bus components are identified from maintenance records of several transit agencies, including the Detroit Department of Transportation (DDOT), the Central Oklahoma Transportation and Parking Authority (COTPA), the Dallas Transit System (DTS), and the Austin Transit System (ATS).
6. Repair time distributions. Repair time distributions of the components considered are determined using repair times recorded by DDOT.

Tables 1, 2, and 3 give the model parameters for the base case study. Table 1 presents the specification of total active fleet, spare factor, repair labor resources, and overtime for the base case. Parameters of the failure distributions of the 16 components considered are given in Table 2. The failure distribution of the components follows two distinct patterns: (a) the Weibull distribution, and (b) the exponential distribution.

Table 3 gives the repair time distribution param-

TABLE 1 Model Parameters

Parameter	Value
Total active fleet	500 buses
Spare factor	10 percent
Repair resource	200 hr (peak) 150 hr (off-peak)
Overtime	30 percent

TABLE 2 Failure Distribution Parameter

Component	Distribution	Minimum Life (mi)	Parameter 1[a]	Parameter 2[a]
Gear train	Weibull	3,051.9	2.751	113,504.0
Control arm	Weibull	8,634.6	1.364	98,489.0
Blower motor	Weibull	22,323.6	1.431	85,776.0
King pin	Weibull	11,056.5	1.507	84,100.0
Bell crank	Weibull	16,263.0	1.397	83,602.0
Fan torous	Weibull	8,649.9	1.165	76,250.0
Destination sign	Weibull	20,439.9	2.049	82,994.0
Power steering	Weibull	3,448.8	1.263	83,823.0
Condenser core	Weibull	6,507.9	1.446	58,183.0
Engine	Weibull	80,302.5	2.173	167,373.0
Dome light	Weibull	14,223.0	2.930	32,726.0
Transmission	Weibull	3,487.0	1.518	55,107.0
A.C. compressor	Weibull	19,983.5	2.107	123,592.0
Starter motor	Exponential	10,300.0	-	27,666.0
Door engine	Exponential	264.0	-	42,187.0
12-V charger	Exponential	127.0	-	27,497.0

[a]For the Weibull distribution, Parameter 1 is the shape parameter and Parameter 2 is the scale parameter. For the exponential distribution, Parameter 2 is the mean mileage.

TABLE 3 Repair Time Distribution Parameter

Component	Mean (hr)	Standard Deviations (hr)
Gear train	65.00	5.000
Control arm	9.75	2.024
Blower motor	4.27	2.036
King pin	14.00	0.250
Bell crank	2.33	1.780
Fan torous	15.29	1.090
Destination sign	1.31	.637
Power steering	5.00	1.000
Condenser core	4.66	1.895
Engine	80.00	5.000
Dome light	1.21	1.110
Transmission	37.97	5.500
A. C. compressor	10.07	.902
Starter motor	2.68	1.880
Door engine	6.00	0.250
12-V charger	1.73	.680

Note: The distribution for all the components is normal.

eters for all 16 components. In this study, repair times are assumed to be normally distributed. The normal distribution is assumed because (a) of limited data, which makes it difficult to ascertain the validity of other distributions, and (b) several distributions have been used to represent repair time distributions. Sinha and Bhandari (14) used the Gamma distribution, Kelly and Ho (15) found that repair times followed the log-normal distribution, and Conway et al. (16) identified repair times to be normally distributed. Because there does not appear to be a consensus, the normal distribution was selected because of its ease of use and familiarity with its properties.

Now that the model parameters have been presented, the next step is to present the results of the simulation experiments. The results are presented in three steps: (a) the running of experiments that schedule repairs randomly (without systematically ordering the priority of repairs) followed by a comparison of these results with the results of the simulation model when comparable systematic repair rules and policies are used; (b) the running of experiments with the systematic scheduling policies and the selection of the superior systematic scheduling policy; and (c) the determination of the sensitivity of the superiority of scheduling policies to changes in system parameters.

RANDOM SCHEDULING

It has been observed that at several transit systems, buses are not scheduled for repair using specific scheduling rules that take into account the expected work content (processing time) involved in repairing the vehicle. Examples of scheduling without regard to the work content would include ordering bus repairs according to the order in which they arrived at the maintenance facility or even with regard to the preferences of mechanics to conduct certain types of repairs. To model a system that does not schedule repairs with regard to job processing time, the experiments assume that repairs are scheduled randomly using the following procedures:

1. Option 1. In this option, the job that arrives in the failed queue at the earliest date will be selected for repair first (FCFS).
2. Option 2. In this option, the failed vehicle queues are separated by type of failed component or part into separate failed vehicle queues. Then, a failed vehicle queue is randomly selected and a bus is scheduled for repair from the queue based on FCFS, unless a bus in the selected queue has been waiting longer than 2 days, and then is it repaired first.
3. Option 3. In this option, if any job in the randomly selected failed vehicle queue has waited longer than 2 days, then selection is made among these jobs using LCFS. If no failed buses have been waiting longer than 2 days, failed buses are selected for repair using LCFS.
4. Option 4. In this option, all jobs are selected according to LCFS without a waiting-time limit.

All four of the random scheduling repair policies are similar to the systematic repair policy with the options numbered identically (see Figure 1 for systematic repair scheduling policies). The difference in each case is that the random policies do not schedule jobs according to minimum processing time.

Four runs of the simulation model are made using the four random scheduling policies. In all cases, the seed value for the random number generator is kept constant. By keeping the seed value constant, the same stream of random numbers is used in all runs. Hence, the same sequence of random samples will be generated for each run of the model. The system performance indicators of these runs are presented in Table 4 along with the performance indicators for the comparable systematic repair policies.

TABLE 4 System Performance of Two Repair Processes

Repair Policy Options	TSYS	OTIME	TQUEUE	WQUEUE
1				
Random	4.279	29.48	68.73	53.19
Systematic	3.384	26.74	58.71	42.73
2				
Random	3.931	22.81	62.84	47.06
Systematic	2.975	18.41	49.48	33.68
3				
Random	4.639	27.55	76.06	60.10
Systematic	2.345	23.79	54.31	38.28
4				
Random	3.345	29.02	71.91	56.35
Systematic	1.721	24.43	54.96	39.10

Comparison of the Two Repair Processes

In Table 4, TSYS for randomly scheduled repair options varies from 2.541 days (Option 4) to 3.424 days (Option 1). For the systematic repair options, TSYS (the average time spent by buses in the maintenance system) varies from 1.721 days (Option 4) to 3.384 days (Option 1). From this experimentation, it is observed that TSYS for each systematically scheduled repair option is always lower than that of the comparable randomly scheduled repair option. Other performance indicators also prove the superiority of systematically scheduling repairs. A t-test is conducted to compare the performance indicators of the two repair processes for similar options and, for all options, they are statistically different at the 95 percent confidence level.

This comparison demonstrates that the system performance for systematically scheduled repairs is superior to that of randomly scheduled repairs. For example, while using systematic scheduling rules, the time that buses are tied up in the maintenance system (TSYS) under the best conditions (Option 4) for both processes (random and systematic) is roughly one half the time required under random scheduling. In the next section, systematically scheduled repair options are compared.

SYSTEMATICALLY SCHEDULED REPAIR POLICIES

Table 5 gives the performance indicators for all eight systematically scheduled repair options. In Table 5, it should be noted that the Option 4 (minimum processing time and LCFS) performance for TSYS is significantly better than the other options. More specifically, the application of Option 4 results in buses being tied up for maintenance a shorter average time than any other repair scheduling policy.

In Option 4, failed vehicles are scheduled for repair based on processing and arrival times. The vehicle that joins the queue at the last moment and needs the minimum time to be repaired is given the highest priority. This causes the repaired vehicles to spend the minimum average time in the maintenance system. It is important to note that other performance indicators are not at their least value for

TABLE 5 System Performance (base case) Systematically Scheduled Repair

Repair Policy Options	TSYS	OTIME	TQUEUE	WQUEUE
1	3.384	26.74	58.71	42.73
2	2.975	18.41	49.48	33.68
3	2.345	23.79	54.31	38.28
4	1.721	24.43	54.96	39.10
5	3.208	15.50	49.31	33.13
6	3.015	14.94	47.22	30.91
7	2.565	16.65	49.71	33.47
8	2.412	14.99	48.49	32.37

Option 4. This trait has also been observed by researchers who have studied scheduling in other industries (16-18). According to Conway et al., under the minimum processing time rule, the mean time spent in the system is small but some individuals' jobs (those requiring long processing time) will be intolerably delayed (19). Thus, although some jobs will take short times to flow through the system, a few will require inordinate lengths of time to be processed through the system. Because of the variability in the time spent in the system, other performance indicators are not at their minimum for Option 4.

The performance indicators, OTIME, TQUEUE, and WQUEUE, are at their lowest values for Option 6. In Option 6, failed buses are scheduled for repair by using maximum processing time (MXPT) and a maximum waiting time constraint. The Option 6 waiting time constraint places vehicles that have waited longer than 2 days first in line for repairs on the next day. Later, waiting-time limits will be explored to determine if 2 days is the most efficient limit and, if not, how many days the limit should be.

The values of OTIME, TQUEUE, and WQUEUE are close for Options 2 and 6. The only difference between Options 2 and 6 is that in Option 2, the repair work is scheduled using minimum processing time (MPT) and FCFS rules. The waiting time constraint is common to both options. A t-test is conducted to compare the performance indicators of these two options. It is found that, at the 95 percent confidence level, there is no statistically significant difference between the performance indicators of Options 2 and 6.

In all the policies tested, it is observed that the system is operating at capacity almost all of the time. In other words, the utilization of available resources is approximately the same under all policies. The main objective of scheduling is to maximize the number of repairs using available resources. Therefore, all the policies utilize repair resources equally.

Another important observation is that the total number of failed vehicles waiting for repairs (performance indicator TQUEUE) attributed to Options 1, 4, 5, and 8 is significantly greater than that of Options 2, 3, 6, and 7, respectively. However, the only difference between the two sets of repair options is that Options 2, 3, 6, and 7 have waiting time constraints. This difference permits the measurement of the waiting time constraint's impact on system performance.

SENSITIVITY ANALYSIS

This phase of the experimentation is designed to determine the extent to which the performance of the simulated system is affected by changes in model parameters. The model parameters considered in the sensitivity analyses are as follows:

1. The failure distribution parameters of components,
2. The spare bus factor,
3. The fleet size,
4. The man-hours (repair resources) available, and
5. The maximum waiting time limits for Options 2, 3, 6, and 7.

The impacts on the superiority of the various options as the parameters are changed are examined in the following sections.

Failure Distribution Parameters

The distribution of component failures with respect to wear varies with environment, duty cycle, terrain, and so forth. In this part of the sensitivity analysis, the failure distribution parameters of bus components are modified and two different sets of simulation runs are made. The outputs of the two sets of runs are compared with the base case. The two sets of runs have two distinct features:

1. In Case I, the Weibull distribution has three parameters: (a) shape, (b) scale, and (c) minimum life. The components whose failure distribution shape parameters are close to 2 and above are considered to have age-dependent and predictable failure rates (20). Those with shape parameters close to 1 or lower are considered to have random failure patterns. Those components with failure distribution shape parameters close to 1 have their shape parameter changed to 2. By doing this, the failure distributions are all age-dependent. One run is made for each of the eight scheduling options using the age-dependent parameters and the results are given in Table 6.
2. In Case II, the component failure distribution parameters are changed from age-dependent to

TABLE 6 Distribution Parameter and System Performance

Repair Policy Options	TSYS	OTIME	TQUEUE	WQUEUE
1				
Case I	6.056	36.29	107.03	91.48
Case II	6.110	36.24	108.20	91.96
Base case	3.384	26.74	58.71	42.73
2				
Case I	4.834	33.92	83.14	65.87
Case II	4.914	35.35	85.10	67.77
Base case	2.975	18.41	49.48	33.68
3				
Case I	5.490	35.30	95.03	78.21
Case II	5.830	35.60	101.40	84.73
Base case	2.345	23.79	54.31	38.28
4				
Case I	3.149	36.29	108.30	92.00
Case II	2.906	36.29	113.90	97.56
Base case	1.721	24.43	54.96	39.10
5				
Case I	5.161	36.19	87.14	70.23
Case II	5.393	35.99	90.27	73.42
Base case	3.208	15.50	49.31	33.13
6				
Case I	4.841	33.23	82.73	65.54
Case II	4.749	33.34	81.41	64.12
Base case	3.015	14.94	47.22	30.91
7				
Case I	5.181	33.19	86.17	69.20
Case II	5.330	33.19	89.40	72.52
Base case	2.565	16.65	49.71	33.47
8				
Case I	3.176	36.19	85.52	68.44
Case II	3.593	36.20	92.56	75.64
Base case	2.412	14.99	48.49	32.37

random. Similar to Case I, eight runs of the model are made and the results are also given in Table 6.

The system performance indicators for the base case and Cases I and II are given in Table 6. The average time spent by buses being repaired reaches a minimum under Option 4 for all three cases. For both Cases I and II, the repair policy, which resulted in the minimum value of OTIME, TQUEUE, and WQUEUE, is Option 6. It should be noted that there is no statistically significant difference between the performance indicators of Options 2 and 6.

For Cases I and II, as well as the base case, the same repair policy is superior. From this observation, it can be concluded that the superiority of scheduling policies is not sensitive to the values of the failure distribution parameters. This means that if one policy is superior in one environment, it will be the superior policy in another.

Spare Factor

Sinha and Bhandari found that the number of spare buses has a significant influence on the reliability of transit service (14). To analyze the impact of the spare factor on the simulated system's performance and the superiority of scheduling policies, the spare factor (i.e., spare buses/active buses) is varied from the base case value.

The base case spare factor is 10 percent. The modified spare factors chosen are 8 and 12 percent. Sixteen different runs are made using the eight repair scheduling options and the two new spare factors. The observed performance indicators are given in Table 7.

Performance indicators for the base case and modified spare factors are tabulated in Table 7. In all cases, the minimum value for time in the system (TSYS) is observed for Option 4. The values for OTIME, TQUEUE, and WQUEUE are at their minimum in Option 6. The superior repair policy remains unchanged under all spare factors.

TABLE 7 Spare Factor and System Performance

Repair Policy Options	TSYS	OTIME	TQUEUE	WQUEUE
1				
Spare factor (8%)	2.942	29.54	52.11	35.80
Base case (10%)	3.384	26.74	58.71	42.73
Spare factor (12%)	3.546	22.44	61.29	45.49
2				
Spare factor (8%)	2.599	21.52	44.27	27.60
Base case (10%)	2.975	18.41	49.48	33.68
Spare factor (12%)	3.492	15.52	57.72	41.80
3				
Spare factor (8%)	2.849	27.03	52.84	36.44
Base case (10%)	2.345	23.79	54.31	38.28
Spare factor (12%)	2.005	23.88	64.58	48.53
4				
Spare factor (8%)	2.086	29.94	55.11	38.78
Base case (10%)	1.721	24.43	54.96	39.10
Spare factor (12%)	1.678	23.49	63.97	48.10
5				
Spare factor (8%)	2.897	21.77	45.08	28.47
Base case (10%)	3.208	15.50	49.31	33.13
Spare factor (12%)	3.707	13.89	57.16	41.10
6				
Spare factor (8%)	2.661	18.27	42.77	25.97
Base case (10%)	3.015	14.94	47.22	30.91
Spare factor (12%)	3.527	13.48	55.49	39.31
7				
Spare factor (8%)	2.826	20.59	44.44	27.52
Base case (10%)	2.565	16.65	49.71	33.47
Spare factor (12%)	2.355	15.03	58.10	41.90
8				
Spare factor (8%) 1	2.547	20.88	44.18	27.18
Base case (10%)	2.412	14.99	48.49	32.37
Spare factor (12%)	2.487	14.89	58.86	42.75

In Table 7, note that all system performance indicators except OTIME are higher for the spare factor of 0.12 relative to 0.08. While modeling, it is assumed that if the number of failed vehicles exceeds the number of spare vehicles, then, and only then, will overtime be permitted. Through time and by-random-chance failures will occur in surges. How well the system can absorb these surges depends on the number of spares that is available to replace the failed vehicle. Therefore, the simulation experiments demonstrate that there is a relationship between the spare factor and the labor hours required (both overtime and regular time), which indicates the relationship between transit system operating costs and capital costs. In other words, there is a definite trade-off between the capital costs invested in spare vehicles and the operating expenditures on mechanic labor.

This finding has serious transit industry policy implications. The Urban Mass Transportation Administration (UMTA) is currently evaluating its policy on permissible spare ratios (12). Presumably the emphasis in UMTA's spare ratio policy will be to place a reasonable cap on the number of spare buses that a transit system may carry. If spare ratios are reduced, it will come at the cost of additional operating costs. Because the portion of operating costs of U.S. public transit systems subsidized by the federal government is less than the portion of capital costs that is federally subsidized, capping spare ratios will have the impact of pushing more of the total costs of transit service back on the transit systems that currently have spare ratios that are higher than the cap. However, the trade-off between maintenance labor hours and spare buses has not been quantified and, without this information, policy makers placing a cap on spares (in the name of cost savings) may select an inefficient cap. An inefficient limit may ultimately increase the total cost of transit service (operating plus capital cost) for those systems that are forced to reduce the number of spares they carry.

Fleet Size

The fleet size varies with the transit system and depends on the quantity and quality of transit services provided. In this experiment, the fleet size is changed from 500 to 600 vehicles. Although the fleet size is changed, the spare factor is kept at 10 percent. The system performance indicators for all eight options and fleet sizes of 500 and 600 buses are given in Table 8.

When the fleet size is 600 buses, the minimum value of TSYS is 1.971 days for Option 4. It is 1.721 days for a fleet of 500 buses. When the fleet size is 600, more buses are put in service resulting in more failures than with 500 buses. This creates a higher level of competition among the failed entities to be selected for repair. Because maintenance resources are held constant for both fleet sizes, the failed entities spend more time in the maintenance system waiting to be scheduled for repair. This causes a higher value of TSYS for the 600-bus fleet. The other performance indicators also have higher values for the 600-bus fleet.

With the exception of TSYS, the performance indicators are at their minimum for either Option 2 or Option 6 for the 600-bus fleet. Further, no statistically significant difference is found between the performance indicators for Options 2 and 6. Because the same result occurs when the fleet size is 500 buses, it can be concluded that scheduling policy superiority is insensitive to fleet size.

TABLE 8 System Performance for Various Fleet Size

Repair Policy Options	TSYS	OTIME	TQUEUE	WQUEUE
1				
600-bus fleet	6.890	36.16	128.51	114.50
500-bus fleet	3.384	26.74	58.71	42.73
2				
600-bus fleet	5.027	33.32	92.44	74.19
500-bus fleet	2.975	18.41	49.98	33.68
3				
600-bus fleet	3.101	36.16	133.50	116.20
500-bus fleet	2.345	23.79	54.31	38.28
4				
600-bus fleet	1.971	36.16	136.50	119.40
500-bus fleet	1.721	24.43	54.31	38.28
5				
600-bus fleet	5.833	36.16	103.90	86.09
500-bus fleet	3.208	15.15	49.31	33.13
6				
600-bus fleet	5.122	33.82	93.31	75.06
500-bus fleet	3.015	14.94	47.22	30.91
7				
600-bus fleet	3.852	35.82	97.60	79.53
500-bus fleet	2.565	16.65	49.71	33.47
8				
600-bus fleet	2.752	36.00	111.80	94.06
500-bus fleet	2.412	14.99	48.49	32.37

Labor Availability

Labor availability is the most important element of maintenance activities. It controls the number of failed vehicles not scheduled for repair on a particular day. In this part of the analysis, the sensitivity of repair scheduling policy superiority to labor availability is tested. For the purpose of the simulation, the following equation is used to specify the level of available labor hours:

$$\text{Labor hours available per day} = (\text{Number of buses}) \times (\text{a Factor}) \qquad (1)$$

Then, the factor is given a variety of values, including 0.20, 0.35, 0.40, 0.425, 0.45, and 0.50. In the base case, the labor available per day was 200 hr [(Number of buses) x 0.40)]. Forty-eight runs of the simulation model are made using all combinations of the varied number of labor hours available and the eight scheduling policy options. Plots of the values of TSYS, OTIME, TQUEUE, and WQUEUE for all the combinations are shown in Figures 2-5, respectively.

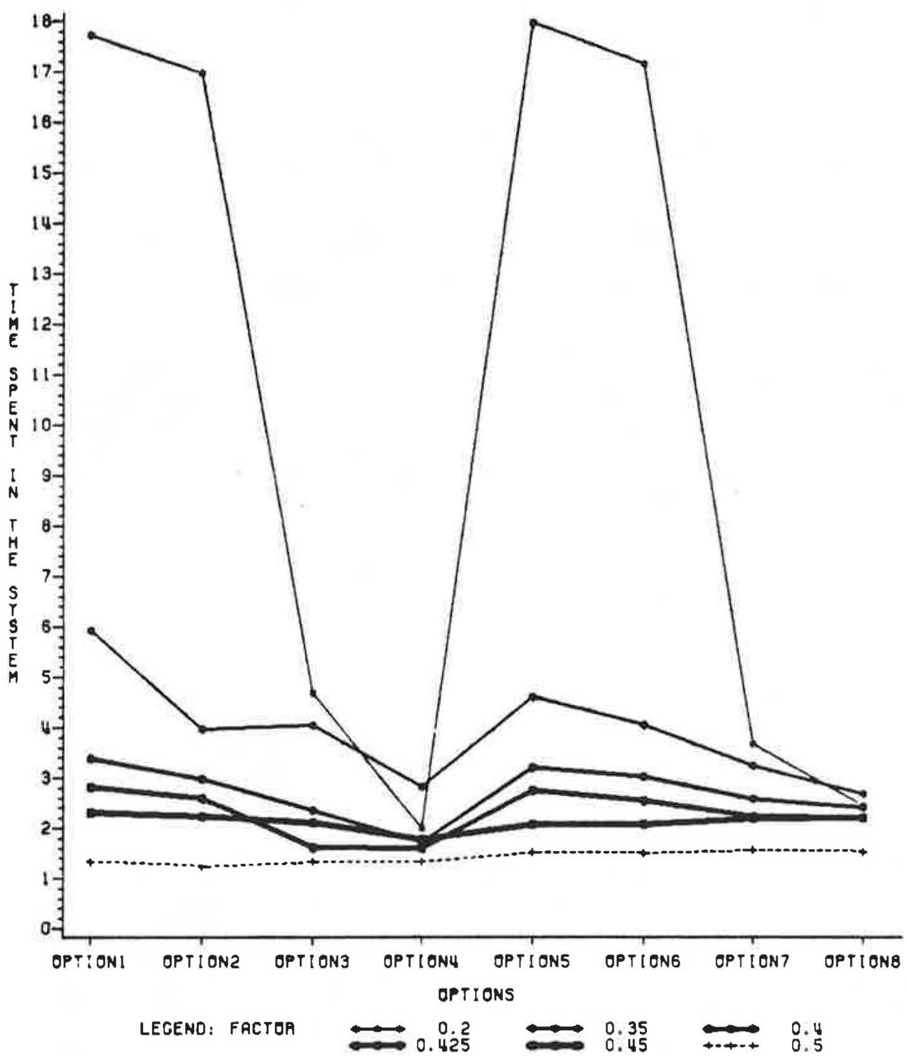

FIGURE 2 System performance (TSYS) of different repair policies.

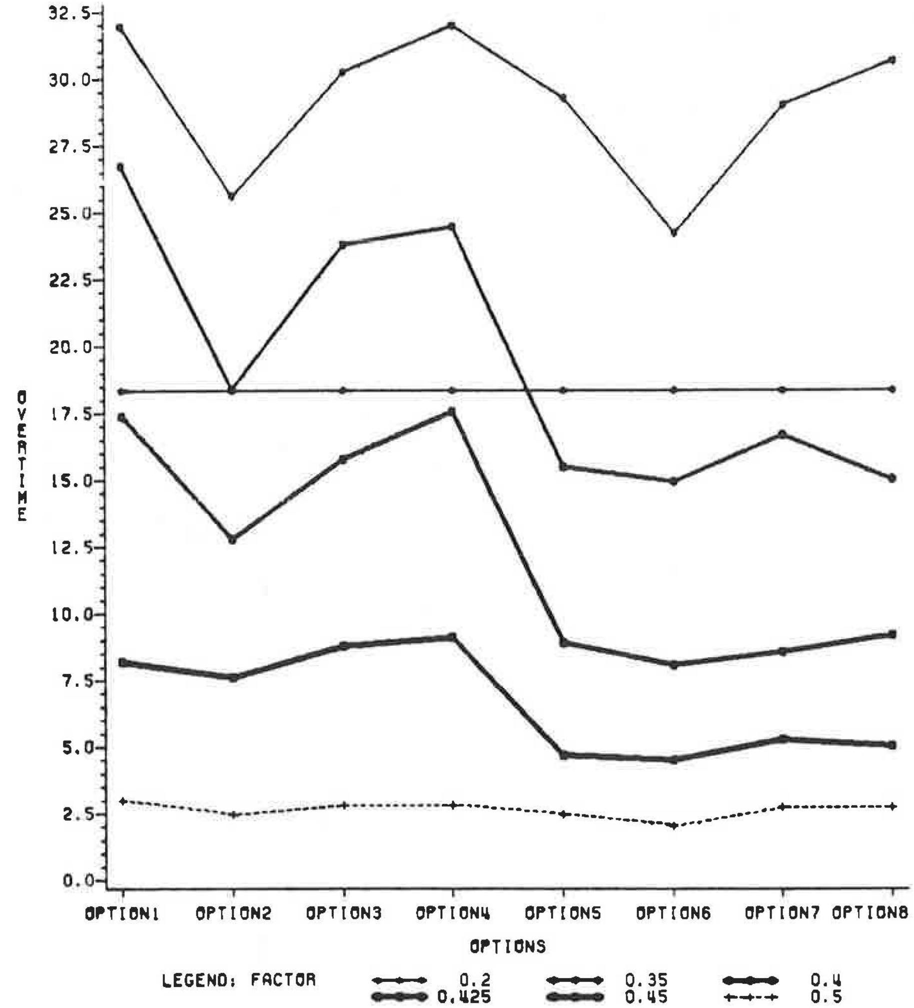

FIGURE 3 System performance (OTIME) of different repair policies.

As shown in Figure 2, it is evident that for values of the factor up to 0.425, the same repair policy is superior (i.e., Option 4). For other indicators, Option 6 is superior up to a factor of 0.425. After 0.425, repair resources move toward saturation. This means that when the value of the factor is more than 0.425, there is no competition among the failed entities for repair resources because there are more than enough available. As a result, when the system is saturated with available labor, the system performance for all options becomes approximately the same because efficient scheduling no longer matters. This means that when labor availability is excessive, there is no need for scheduling. Spinner, while researching the importance of scheduling, found that the same is true in other industrial applications of scheduling (18).

Waiting Time

For the simulation runs made with the base case parameters, Options 2 and 6 provided nearly the same level of performance. The important feature of both options is the limit on the maximum number of days a bus could wait before being scheduled for repair work. In this experiment, the sensitivity of system performance to the length of the waiting time constraint is analyzed.

The analysis is performed using Option 2 and varying the waiting time limit. The maximum waiting time limits considered in the experiments include 2, 3, 4, 5, and 10 days. The observed performance indicators for the various waiting times are given in Table 9. It is observed that the average time spent in the system (TSYS) increases with increased waiting time constraints. On the other hand, based on the value of TQUEUE and WQUEUE, the 4-day waiting time constraint seems to be better in comparison to the other waiting time limits. This result points out the importance of a proper waiting time limit on the system performance.

CONCLUSIONS

Presented in this paper were the results of a series of simulation experiments. The experiments were conducted to determine superior repair policies and test the sensitivity of repair policies with varied system conditions. From this study, it is concluded that

1. The performance of transit maintenance can be dramatically improved with the use of systematic repair scheduling rules.
2. The performance of transit maintenance varies widely with different repair scheduling policies.
3. Specific repair policies for scheduling are almost always superior regardless of the values of the system parameters.
4. The importance of efficient scheduling is increased when labor resources are constrained.

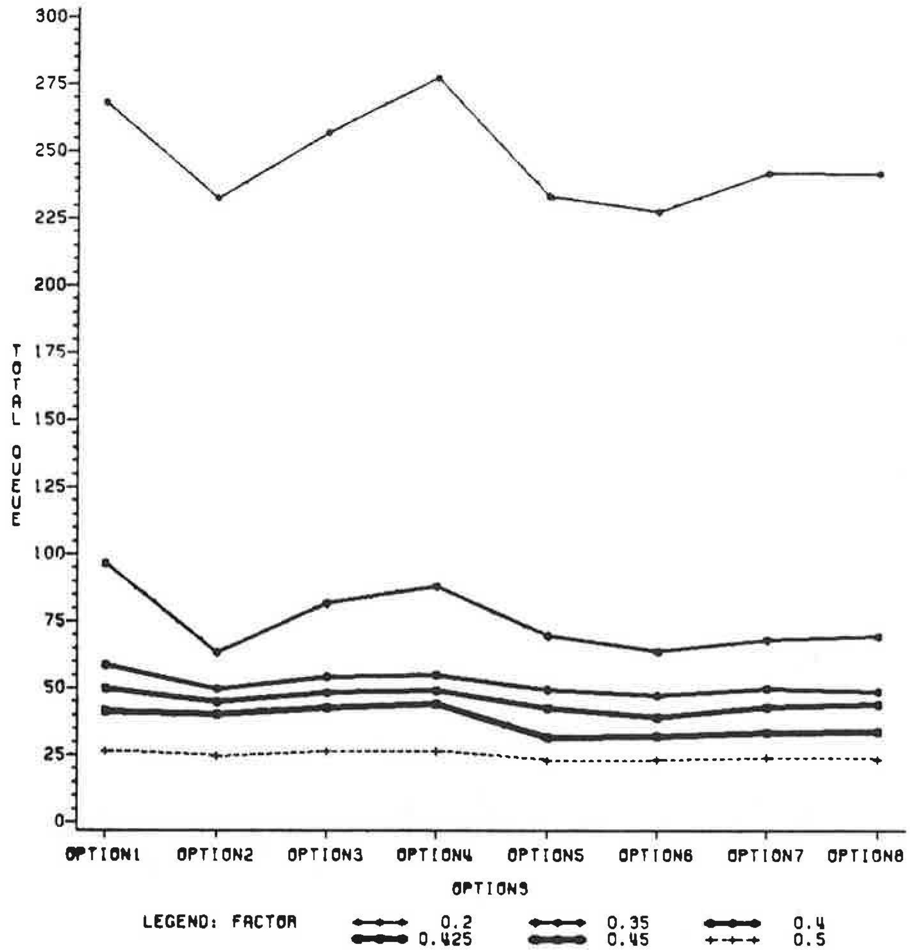

FIGURE 4 System performance (TQUEUE) of different repair policies.

5. Capital cost savings through reduction in spare buses can be accomplished at the expense of increased maintenance labor costs.

6. By assigning a higher priority to those failed buses that have waited for repairs more than the maximum waiting time, the system performance can be significantly improved.

RECOMMENDATIONS

Recommendations were derived for two subjects. The first involves transit industry policy designed to regulate the management of bus fleets (e.g., spare factor limits, maintenance standards, and age requirements for vehicle replacement). The second level deals with the future use of simulation analysis to study bus fleet management issues.

Policy Recommendations

In 1981 UMTA attempted to develop standard policy guidelines for transit maintenance (21). However, this effort was finally abandoned because of a lack of agreement on universally acceptable standards. The important point that UMTA's experience illustrates is the inability to prescribe specific sets of blanket minimum standards that are applicable and acceptable under all circumstances.

Experiments conducted in this study found that the superiority of specific scheduling policies is universal to all conditions. This finding is another demonstration that tested management methods (e.g., scheduling policies and other techniques) are universally applicable. This suggests that, if some assurance of proper maintenance is required, transit agencies should be advised by UMTA to institute proven management methods (e.g., repair scheduling techniques and other management techniques) instead of adopting blanket standards (e.g., minimum spare ratio requirements). Through the use of proven fleet management methods, the transit system's management has the flexibility to efficiently adjust their maintenance procedures to fit their own circumstances (e.g., fleet age, vehicle mix, duty cycle, labor wage rates, and terrain). On the other hand, blanket minimum requirements leave no room for flexibility.

Methodology Issues

As outlined earlier in the paper, the simulation methodology utilized in this study has the drawback of including only a limited number of events. This is an inherent problem in any simulation model that utilizes probability distributions to generate event occurrences. There are simply too many events to be able to economically derive probability distributions for each one. This necessitates using a limited subset of the possible events in the model. Simulation studies that use only a fraction of the system's activities in the analysis are generally appropriate for policy studies, but such simulations are of limited value to the study of operational is-

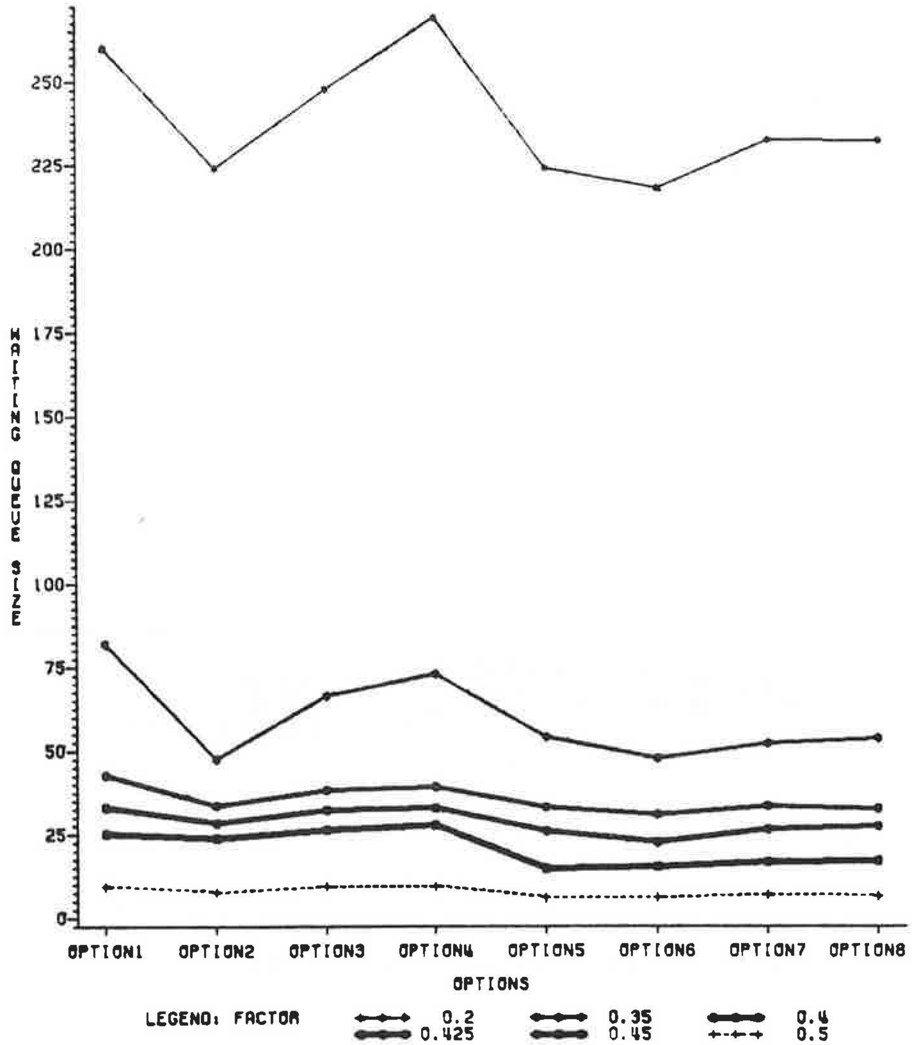

FIGURE 5 System performance (WQUEUE) of different repair policies.

TABLE 9 System Performance for Waiting Time Limit

Waiting Time (days)	TSYS	OTIME	TQUEUE	WQUEUE
2	16.96	18.32	232.2	223.8
3	17.29	18.32	226.7	217.3
4	17.27	18.32	224.7	215.1
5	19.91	18.32	232.2	222.9
10	18.45	18.32	229.8	220.3

Note: The total resource for this run during peak was 100 hr per day and the total resource for this run during offpeak was 75 hr per day.

sues. Operational issues require that the analysis provide information on the strengths of relationships with a high degree of confidence in the results.

A possible alternative to the use of a probability distribution-driven simulation model is the use of a trace-driven simulation model (22). A trace-driven model does not generate a stream of events from distributions. It uses a stream of historical events to drive the simulation. In other words, a simulated bus fleet assumes events in the same order that they were experienced by an operational fleet of buses. Therefore, all events that occurred in the period during which the data were collected are included in the simulated stream of events. Through the use of a trace-driven simulation, detailed analysis could be conducted of specific operational issues. However, whether future researchers use trace-driven simulation or some other approach, and before any detailed analysis can be conducted, richer and more complete data sets than those currently in existence must be made available to researchers.

ACKNOWLEDGMENTS

The authors would like to thank William Swanger, Richard Golembiewski, and James Gregory Mitchell, all formerly with the Detroit Department of Transportation, for their thoughts, advice, and help in obtaining data. Others who provided assistance in procuring data include Sam Cox, Dallas Transit System; Bill Loftgren, Austin (Texas) Transit System; Timothy Lett, formerly with the City Transit Service of Fort Worth, Texas; and John Shuggart, Central Oklahoma Transportation and Parking Authority (Oklahoma City). Further, the senior author would like to thank L.B. West, J.R. Stock, C.E. Barb, and W. Fisher for their assistance in the initial research that led to this paper. Finally, this paper and the research described were, in part, made possible by grants to the University of Oklahoma from the UMTA University Research and Training Program. The authors are grateful to have had the opportunity to conduct research through this program.

REFERENCES

1. M.M. Smith. Mass Transit: A Perspective for the Future. Special Report 198. TRB, National Research Council, Washington, D.C., 1983.
2. B.W. Kliem and D.L. Goeddel. Management Tools for Improving Maintenance Performance. Special Report 198. TRB, National Research Council, Washington, D.C., 1983.
3. T.H. Maze, G.C. Jackson, and U. Dutta. Bus Maintenance Planning with Computer Simulation. Journal of the Transportation Engineering Division, ASCE, Nov. 1982, pp. 550-569.
4. E.T. Etschmair and G. Anagnostopoulus. Systems Approach to Transit Bus Maintenance. In Transportation Research Record 994, TRB, National Research Council, Washington, D.C., 1984, pp. 65-71.
5. J. Foerster, F.G. Miller, M. Kosinski, and A. Rueda. Management Tools for Bus Maintenance--Current Practices and New Methods. Urban Mass Transportation Administration, U.S. Department of Transportation, 1983.
6. T.H. Maze, A.R. Cook, and U. Dutta. Bus Fleet Management Techniques Guide. Urban Mass Transportation Administration, U.S. Department of Transportation, 1985.
7. W.H. Hausman and G.D. Scudder. Priority Scheduling Rules for Repairable Inventory Systems. Technical Report 80-3R. Industrial Engineering and Department of Management, Stanford University, Stanford, Calif., 1981.
8. G.C. Haenisch and F.G. Miller. Increasing Productivity in Bus Maintenance Functions. Proc., American Institute of Industrial Engineers Spring Conference, Norcross, Ga., 1976.
9. U. Dutta. Development of a Bus Maintenance Planning Model. Ph.D. dissertation. School of Civil Engineering and Environmental Science, University of Oklahoma, Norman, 1985.
10. Bus Maintenance Facility Planning. Committee 6A-30. The Institute of Transportation Engineers, Washington, D.C., 1985.
11. J. Barber. A Conceptual Framework for the Applicability of Technology to the Transit Bus Industry: Staff Paper. UMTA, U.S. Department of Transportation, 1983.
12. Section 9 Formula Grants Application Instructions. UMTA, U.S. Department of Transportation, 1984.
13. 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.
14. K.C. Sinha and A.S. Bhandari. A Comprehensive Analysis of Urban Bus Transit Efficiency and Productivity, Part III. UMTA, U.S. Department of Transportation, 1978.
15. A. Kelly and C. Ho. Topics in Terotechnology: An Instructional Series. Maintenance Management International, No. 3, 1982, pp. 49-72.
16. R.W. Conway, B.M. Johnson, and W.L. Maxwell. An Experimental Investigation of Priority Dispatching. Journal of Industrial Engineering, Vol. 11, June, 1960, pp. 221-229.
17. J.E. Day and M.P. Hottenstein. Review of Sequencing Research. Naval Research Logistic Quarterly, Vol. 15, No. 1, 1970, pp. 319-330.
18. A.H. Spinner. Sequencing Theory--Development to Date. Naval Research Logistic Quarterly, Vol. 15, No. 1, 1968, pp. 319-330.
19. R.W. Conway, W.L. Maxwell, and L.W. Miller. Theory of Scheduling. Addison Wesley Publishing Company, New York, 1974.
20. A. Kelly. Maintenance Planning and Control. Butterworths, London, England, 1984.
21. DOT Needs Better Assurance that Transit Systems are Maintaining Buses. Report GAO/RCED 83-57. U.S. Government Accounting Office, 1983.
22. A.A.B. Pritsker and C.D. Pegden. Introduction to Simulation and SLAM. John Wiley and Sons, Inc., New York, 1979.

Discussion

Peter Wood*

Early last year, as in prior years, I had the opportunity to review two papers that were submitted for presentation at the Transportation Research Board's Annual Meeting. Both related to simulations of maintenance strategies. As usual, both papers were well written, scientifically sound, and included extensive bibliographies.

Some quotes from these papers follow. In the paper "Effectiveness of Improved Repair Scheduling in the Performance of Bus Transit Maintenance," the authors wrote "Maintenance workers are interchangeable and can perform all the repairs made at the central facility . . . all buses are the same model . . . maintenance equipment and tools are always available." In the paper "Exploring the Multiple Factor Concept for Bus Maintenance Using Simulation" (elsewhere in this Record), the authors wrote "The fleet is brand new . . . the maintenance manager must promulgate different PM interval guidelines for each of the garages."

Simplistic assumptions and unrealistic procedures such as these characterize virtually all the papers on this subject that I have reviewed over the past few years. This is unfortunate, because many of them contain useful ideas that, if implemented, could lead to some improvements in efficiency. However, when a paper based on artificial restraints, hypothetical data, and broad assumptions states that: "From this study it is concluded that the performance of transit maintenance can be dramatically improved . . ." it is not surprising that the transit industry dismisses it as yet another paper produced by an academic with no knowledge of the real world.

What can be done to make this work more useful, more usable, and, most important, more acceptable? First, let me state some assumptions of my own:

1. Data are, and always will be, inaccurate, incomplete, and out of date,
2. We should concentrate more on decision support tools and less on optimization under steady-state conditions, and
3. Any program, however good, that increases the workload of the maintenance manager, is likely to be ignored.

I will examine each of the assumptions in turn. First, I will consider data. Typical of the comments that appear in papers are: "Data are not presently available . . . not viewed as particularly meaningful . . . if reliable data could be collected the

*The MITRE Corporation, 1820 Dolley Madison Boulevard, McLean, Virginia 22101.

concept would have significant merit . . . the results were misleading because the data were inconsistent." Similar statements have been made.

And yet, one model calls for "failure distributions by component, preventive replacement times by component, emergency replacement times by component, probability of bus-accident upon in-service failure of component, costs of and times for replacement, average costs of replacement, average cost of an accident, and bus preparation costs." Even if these data were provided, would anyone be prepared to guarantee their accuracy? In the remarks of British economist Sir Josiah Stamp (1880-1941), "The government is very keen on amassing statistics. They collect them, add them, raise them to the nth power, take the cube root and prepare wonderful diagrams. But you must never forget that every one of these figures comes in the first instance from the village watchman, who puts down what he damn well pleases."

It should not be believed that a 10 or even a 1 percent improvement in efficiency can be achieved if this is dependent on the generous availability of accurate data. Models should be designed using industry averages, modified where appropriate by local estimates, and refined whenever possible by validated data. These are the first steps toward utility.

Now to the second point. Most work on maintenance modeling today is based on maximizing efficiency in a steady-state environment. We have a given number of buses, a certain number of miles operated, component failures occurring at statistically established intervals, preventative maintenance performed at specified times, and so forth. A common objective is to minimize the maintenance cost per vehicle mile. If a more sophisticated model is being dealt with, an element will be included that relates to road calls (it is undesirable to run buses until they break down) and spares ratio (it is undesirable to concentrate on simple repairs only).

But how should the situation be handled where, for example, an attempt is being made to service and repair all the air conditioning equipment before the start of the summer? Or where a new fleet of buses is being introduced and several of the key mechanics are placed at the manufacturer's plant? What about the staff to handle the inevitable high level of initial failures? Because these are warranty repairs, they do not affect costs, but they certainly affect labor availability. The situation is even worse when a new bus design is introduced. And yet, these are real-world problems that a maintenance manager has to face. They are precisely the kind of problems that could usefully be handled through a simulation.

There is a class of software systems now being introduced under the general heading of "decision support systems." These are not intended to replace the manager, but to provide him with information on which he can make informed decisions by providing a range of acceptable alternatives, together with the advantages and disadvantages of each. In contrast to management information systems, which report the results of previous actions, decision support systems attempt to predict the results of future ones. They allow the manager to say "What if . . .?" and look at the results.

A simulation that, for example, aids in resource allocation (such as labor) is based on the data that are regularly available, takes into considera- tion the constraints that exist within a specific system, and allows the user to choose from a range of alternatives, is precisely the type of mainte- nance simulation that would be useful.

Such a system would satisfy my third assumption: It should be configured to minimize the demands on the user--for example, through the use of graphics and menus, and by requiring inputs to be of the simple "yes/no" variety--and to provide an explanation facility so that the user understands why a particular answer has been given.

Such a simulation could provide answers to questions such as "Which buses should be worked on so that the maximum number will be available for a special event?" "How should the work be scheduled during the period when two of my key employees will be at the manufacturer's plant inspecting the new bus order?" and "The level of service is being reduced by 10 percent; by how much can my maintenance costs be reduced?" Note that all of these are dynamic conditions, not the static conditions that have been assumed for most simulations.

How can such a simulation be worked toward? By concentrating on researching how a maintenance department is managed, rather than on how it is operated. An essential first step would be to establish the decisions that are being made by the maintenance manager in his day-to-day operational role. What are the steps that he takes in reaching these decisions? What information does he need? What information would make the decision-making process easier? Based on this information, the requirements of the simulation that would answer the maintenance manager's needs could then be examined. I have no doubt that such a simulation would be both useful and accepted, for at least four reasons:

1. Most decision making could be improved if more time were available to analyze the alternatives. An effective simulation would present a greater range of alternatives to the manager, together with an analysis of the impact of each.
2. The simulation capability would provide for improved decision making at abnormal times (e.g., if a type of bus developed a defect that required that all buses of that type be removed from service).
3. The manager could spend less time planning and more time managing, and
4. The simulation would be a valuable training tool, allowing a new or potential manager to assess the impact of various decisions "off line."

I have tried to provide some suggestions about how the many valuable ideas that these papers (on bus PM systems) contain can actually be "reduced to practice." I believe that this can be achieved easily by concentrating less on scientific abstractions, and dealing more with practical realities.

Authors' Closure

Although it is apparent that Wood's comments are directed at a number of papers and not just the authors' paper (Effectiveness of Improved Repair Scheduling in the Performance of Bus Transit Maintenance), it is perhaps fitting that the authors should respond to Wood's comments. In past years, the authors have written many of the papers to which Wood is referring.

From Wood's comments, two responses come to mind. First, Wood has articulately outlined responses the authors have received from many practitioners regarding their work. Indeed, practitioners have tended to view the authors' work as "yet another paper produced by an ácademic with no knowledge of the real world." However, the authors believe that Wood and other practitioners should not summarily dismiss academic studies solely because they are constrained by simplifying assumptions that fail to entirely duplicate real-world situations. Academics may be in an "ivory tower," but, from this perspective, they can perhaps "see the forest" while practitioners get distracted by the "trees" of assumptions.

Second, Wood's recommendation regarding the development of dynamic computer modeling tools that can be directly applied to day-to-day maintenance problems is sound. In fact, in two previous papers, the authors reached the same conclusion and suggested approaches for the development and use of such systems (1,2). However, there are many reasons why such models have not been developed and, on closer inspection of the state of the practice of bus maintenance management, the authors believe that such models may not even be warranted.

MAINTENANCE MANAGEMENT DECISION SUPPORT SIMULATIONS

Wood is correct in asserting that the lack of quality data is a common scapegoat for the lack of computerized simulations of maintenance management decisions. However, lack of data is not the only problem. It has been the authors' impression that transit maintenance managers do not value these tools or recognize the need for the research needed to develop them. All too often, this is because transit maintenance managers attained their positions because of their experience and knowledge of maintenance, not because of their formal (or informal) training in management.

Transit maintenance managers, like all other managers, should be managers first. Only when maintenance management is raised to the same level of professionalism as other transportation system managers (e.g., transportation engineers, planners, and accountants) will the need for better management support systems be recognized. Because there is no perceived need, there is no pressure for the development of more sophisticated tools. Without such pressure, there will be little funding for the development of maintenance management decision support tools. Without dramatically increased levels of funding, it is unlikely that useful decision support systems will be developed.

It seems realistic, however, to believe that the modest funding that may be available could support the research required to develop static management principles to direct decision making under a number of significant "real-world" situations. To illustrate the value of applying sound management principles to maintenance management, consider the San Juan Metropolitan Bus Authority, which, 10 years ago, was troubled by having too many of its buses tied up in the maintenance shop. Even though they had a spare ratio of almost 50 percent, some runs were missed because of the unavailability of buses (3). Management asked an academic industrial engineer, who was not a bus maintenance expert, how to increase the vehicle flow through the maintenance shop. He drew on scheduling-sequencing theory, which has proved that the flow through a simple system is maximized when the backlogged jobs that require the shortest time are done first (4). Therefore, he advised that when the shop supervisor assigns a job to a mechanic from the maintenance backlog, the job that appears to require the least time to repair should always be selected. The shop supervisors of the San Juan Metropolitan Bus Authority followed this simple management principle and within 3 months, bus unavailability was decreased by nearly 50 percent.

Given that conditions in a bus maintenance system are subject to dynamic change, maintenance management is probably better equipped if they have simple management principles for instant application to day-to-day decisions rather than a cumbersome, data-intensive computer simulation model.

THE "IVORY TOWER" SYNDROME

At the 1986 Annual Meeting of the Transportation Research Board, two bus maintenance simulation papers were presented that were apparently reviewed by Wood. One paper, by List, Satish, and Lowen (elsewhere in this Record), sought to show that it is important to use other variables besides mileage (e.g., hours of bus use and the duty cycle) to trigger the need to conduct preventive maintenance. The other paper was the authors', which sought to show that is is important to rationally sequence the order of processing maintenance work through a maintenance facility. Both papers resulted in findings that seem obvious: (a) maintenance managers should consider service attributes other than the mileage traveled when deciding on preventive maintenance intervals, and (b) they can improve the flow of bus repairs through the maintenance facility if they sequence repairs with regard to the length of time required to make a repair.

The significance of these papers is that they confirmed their findings through the use of computer simulations that are dramatically less expensive and time consuming than experiments with an actual maintenance system. During the simulation experiments, the researchers made simplifying assumptions to allow the work to fit within the meager resources allotted to them (both studies originated through the modeling work of a graduate student completing a thesis). Both made assumptions which, as Wood noted, do not reflect actual bus maintenance operations. However, both papers had useful findings that can be converted to sound maintenance management principles. It would be unfortunate if practitioners ignored these principles only because they were based on work that made simplifying assumptions. They would then fail to see the forest because of overconcern with suspect details (the "trees").

CONCLUSIONS

Wood's and the authors' arguments may be moot, however. Given the current austere conditions for funding of academic research on transit maintenance management issues, it is likely that there will be little research to create any sort of simulation of bus maintenance systems. However, the authors believe that if there is any funding available in the area of bus maintenance research, it would be more fruitful to examine the relationship between management actions and system performance. From such examinations, the researchers could recommend management principles that appear to improve performance.

REFERENCES

1. T.H. Maze, U. Dutta, and M. 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.
2. E.I. Gitten, T.H. Maze, A.R. Cook, and U. Dutta. Information Systems in Bus Fleet Management. In Transportation Research Record 1019, TRB, National Research Council, Washington, D.C., 1985, pp. 91-98.
3. L.A. Martin-Vega. SPT, Data Analysis, and a Bit of Common Sense in Bus Maintenance Operations: A Case Study. American Institute of Industrial Engineers Transactions, Vol. 13, No. 4, 1981, pp. 349-356.
4. K.R. Baker. Introduction to Sequencing and Scheduling. John Wiley and Sons, Inc., New York, 1974.