MAINTENANCE MANAGEMENT FOR SOPHISTICATED GROUND TRANSPORTATION VEHICLES

C. W. Collins
Engineer, Systems Effectiveness Section
Vought Aeronautics Co., A Division of LTV Aerospace Corp.
Dallas, Texas

The next generation of ground transportation systems will use sophisticated vehicles and computer controls. Maintenance planning for these systems requires correspondingly advanced management techniques. Simulation has proved to be a valuable management tool for evaluating these plans.

This paper discusses a GPSS simulation used to investigate vehicle maintenance strategies for a proposed airport people-mover system. The identification of critical interfaces between operating philosophies and maintenance concepts is described. Model output giving managers insight into relative merits of candidate maintenance system designs is presented.

Introduction

Suppose you're project manager in charge of development for a radically new ground transportation system. You realize that even a technically successful design won't sell in today's market unless its price can be kept low. You're also aware that potential buyers are increasingly conscious of total system cost, and are looking beyond initial purchase price to consider the full spectrum of expenses lumped into the term "life cycle cost." Studies showing that operating and maintenance costs can exceed purchase price by factors of 10 to 100 have received sufficient publicity in recent years to convince you to make maintenance a major consideration in your design, and thus a major feature in your marketing strategy to demonstrate a successful and inexpensive system.[1]

Having made this decision you are faced with the task of answering questions such as:

1. What maintenance concepts will assure vehicle availability at least cost?
2. How can maintenance schedules and traffic demands be reconciled?
3. How many direct maintenance personnel will be required?
4. How many spare vehicles will be needed to insure that peak traffic demands are met?
5. What physical lay-out of the maintenance area(s) will support necessary vehicle maintenance activities?

New Concepts - New Problems

The situation described above recently occurred when a proposal for an airport "people-mover" system was developed. The technical concept was highly complex, involving a horizontal-elevator approach using sophisticated, guideway-suspended vehicles moving automatically around a transportation system serving five major airline terminals. Computer controls and guideway sensors would regulate traffic flow while allowing passengers to choose vehicle destinations by means of on-board pushbuttons. Other cars, carrying baggage, trash, or terminal supplies would be integrated into the traffic stream under direction of a central computer and dispatched to appropriate destinations for automatic loading/unloading.

The system was to serve a terminal complex five miles in length, with a small maintenance area connected to the main guideway at one end. What facilities to provide within the maintenance area was to be determined.

An initial operations research study, based on projected passenger and baggage flow rates, indicated the horizontal elevator concept was theoretically sound. The study also determined traffic demand schedules.

Model Design and Operations

The next step was to translate theory into system design. Actual detail design could take many forms, and it was soon evident that the vehicle was the most volatile design element. This realization resulted in a project to deal with vehicle maintenance through use of a GPSS model. The choice of language was partially influenced by the fact that Vought had previously produced GPSS simulations for investigating aircraft maintenance systems. The major determinant, however, was the ease with which GPSS models can be altered, and the ready understanding of its event-oriented operation by most engineers.[2] These qualities allowed both quick response and simplified conversations with persons providing inputs or using output information.

It was decided to deal only with the maintenance of major on-board subsystems since vehicle design details were in a state of flux. This freed the model builder from concerns about individual equipment items and allowed statistical distributions to represent systems. Daily contact with designers provided continuous distribution updating based on latest detail changes.

Vehicles were represented as GPSS transactions with 25 parameters. After initially generating the entire vehicle population, a predetermined number were positioned on the operating guideway, represented by a simple GPSS user chain. The other cars were placed on another user chain to represent vehicles waiting in a "ready area" off the main guideway.

The traffic demand schedule fed cars back and forth between the two user chains in response to airport traffic demands. As traffic increased, vehicles were called onto the guideway from the ready area. Then, when traffic decreased, excess cars were returned to the ready area.

The availability of vehicles to meet traffic demands was used as the prime measure of maintenance system effectiveness. The number of cars in the ready area immediately prior to peak traffic periods was therefore of extreme importance, and maintenance planning was directed toward insuring that sufficient vehicles would always be available to meet anticipated traffic demands. This objective was the primary determinant of the scheme developed to provide scheduled maintenance inspections to each car, since some way had to be found to do these inspections on a regular basis that wouldn't interfere with supplying cars for traffic demands.

Continuing study of evolving vehicle designs provided estimates for inspection frequencies and inspection duration distributions. With each new estimate a manual comparison was made between inspection and traffic schedules to see if any obvious conflicts were present. When there appeared to be none, the new schedule of maintenance was entered into the model to confirm its non-interference with traffic demands. When it appeared likely that conflicts would occur, revised maintenance schedules were entered and run. The necessary revisions were ordinarily suggested either by the manual schedule comparison or by the model output, which showed when, where, why, and for how long vehicles were being delayed by maintenance/operating interfaces. The model was thus used to find out how to do scheduled maintenance systematically, without degrading transportation operations. In most cases the model indicated scheduled maintenance should be done during the night-time hours when traffic loads were low.

But since maintenance personnel receive a premium wage for night work it wasn't apparent whether this strategy was the most economical. So a variety of other possible schedules were tried in the model.

Complicating the situation was the need to provide for repair of random breakdowns. Maintenance analyses could estimate distributions of failure occurrences and repair times, but couldn't indicate whether unscheduled maintenance should be done only during certain hours or should receive immediate attention whenever the need arose. If unscheduled repair jobs could be handled on a single shift, then perhaps facilities money could be saved by using the same maintenance bays for both scheduled and unscheduled maintenance. Success of this type of scheme would depend on the phasing of activities, since "scheduled" cars would have to vacate maintenance bays before "unscheduled" vehicles could be serviced and vice versa.

Investigations of this phasing, carried out with the model, showed it was desirable to provide a compromise mixture of facilities, some of which were dedicated to unscheduled maintenance activities while others were equipped for both scheduled and limited types of unscheduled maintenance operations. A series of GPSS facilities were used to represent maintenance bays, and received vehicles released from the "guideway" or "ready area" user chains. Vehicle parameter interrogations separated cars into three categories: those due scheduled maintenance, those needing unscheduled maintenance, and those requiring "light" unscheduled maintenance. The latter category applied to any vehicle that could be repaired in 30 minutes or less.

Experiments with different maintenance and operating philosophies involved changes to the "light" maintenance criterion, number and types of bays provided, personnel numbers and kinds, maintenance scheduling, and proposed physical arrangements and access provisions for the maintenance area. What finally evolved was the layout in figure 1. Analysis eventually determined that unscheduled activities could be handled during a single day shift. The anticipated unscheduled maintenance rate and the number of spare vehicles available were such that night-time vehicle failures would not interfere with traffic operations. A skeleton emergency crew would remove any disabled vehicles to the maintenance area to await repair during the following day shift. This was economically desirable, since the personnel involved in repair work would be the most skilled and highly paid, and night work would entail a higher premium wage rate than that for workers doing scheduled tasks.

Scheduled maintenance would begin each evening after the last traffic peak of the day. All vehicles were processed through janitorial services every night in the scheduled/light maintenance area, and this provided an ideal opportunity to check cars for inspection requirements that might be due.

Vehicles in the actual system would carry elapsed time meters to indicate when scheduled maintenance was due. Simulation of this feature was easily accomplished by recording clock time at the completion of each inspection and placing this value in a vehicle parameter. As the cars were processed through janitorial services this parameter was interrogated for use in a comparison with a constant and current clock time, representing time-between-inspections.

A variable was used to "initialize" the scheduled maintenance elapsed time parameter when vehicle transactions were originally generated. The objective was to get the simulation started with an approximately equal number of cars having requirements falling due each day. An adequately-designed maintenance plan should then result in a continuation of the sequence. For instance, if each car was due to receive a minor inspection every three days the initializing variable loaded parameters with dummy times such that 1/3 of the vehicles would fall due on the first simulated day, another 1/3 the next day, etc.

The variable took advantage of the fact that each vehicle transaction had a unique and sequential identification number assigned to it, in the same way an aircraft has a tail number or a ship a hull number. This I.D. number was multiplied by a value representing the theoretical time separation between vehicles, so that if 4 vehicles should ideally become due for scheduled maintenance during any given day, each vehicle's number would be multiplied by the simulation clock equivalent of six hours. A constant value was then added to compensate for the time elapsing between vehicle generation and the first

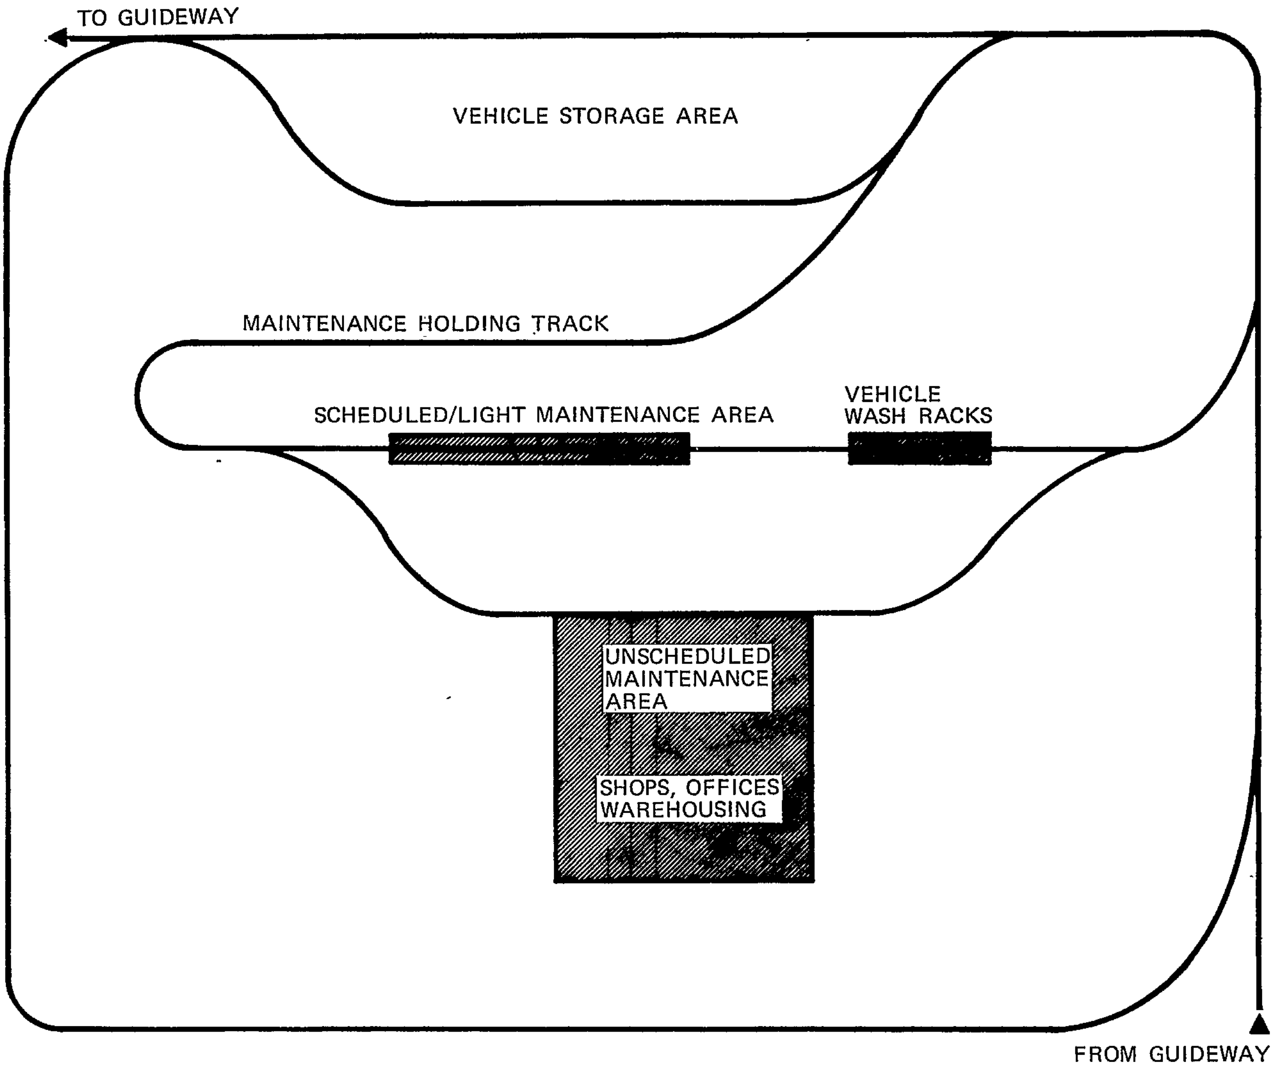

FIGURE 1.

check of scheduled maintenance requirements, resulting in the type of sequence shown in figure 2.

The validity of this philosophy is apparent when it's realized that a major maintenance system objective is to schedule smooth and even workloads. This helps simplify problems of efficient personnel and facilities utilization.

Unscheduled maintenance was introduced by using a GPSS function for the exponential distribution. A continuously-updated estimate of vehicle mean-time-between-failures (MTBF) was used in a generate block to modify the exponential function. Failure transactions were thus produced according to an exponential time distribution whose mean equalled the supplied MTBF.

These transactions caused vehicles to be released from either the guideway or ready area chains and sent to a model segment which determined, through probability distribution functions, which on-board system had failed and how

Example of parameter loading in which four vehicles will have elapsed scheduled maintenance times that are separated by six hours. (GPSS clock unit = two minutes)

Vehicle I.D. Number	X Simulation Clock Equivalent of Six Hours	+ Constant	= Elapsed Time Since Last Scheduled Maintenance
1	180	720	900
2	180	720	1080
3	180	720	1260
4	180	720	1440

FIGURE 2.

long it would take to repair once the necessary spare parts, tools, and maintenance personnel were assembled.

Failed cars entered a queue for space in the unscheduled maintenance facility, operated on a single daytime shift each day. After obtaining a bay, a queue was entered for personnel, represented in GPSS as storages.

After personnel were obtained, a probability distribution was sampled to see if necessary spares could be procured. Personnel were returned to the pool where they became available for use by other waiting vehicles if a delay for parts greater than 20 minutes was indicated. The rationale was that efficient personnel utilization dictated that no more time be consumed in a waiting-for-parts attitude than necessary, while at the same time realizing that new jobs are not ordinarily started by a maintenance man if the current parts delay is small.

More probability distributions were sampled to determine delays for tools, after which an advance block simulated the actual repair time contained in a vehicle parameter. Cars then vacated the unscheduled maintenance facility and underwent a simulated test of on-board systems before being returned to the ready area pool. A failure probability distribution governed whether or not the vehicle passed this test.

Unscheduled maintenance that could be done in 30 minutes or less was categorized as "light," and done in the same bays that were used at night for scheduled maintenance. The chronology of events for these cars was the same as for those in the unscheduled facility.

Additional unscheduled requirements were also generated during the course of regular scheduled inspections. It was, in fact, expected that this would be the primary source of repair actions and provided the main motivation for doing these inspections, since detection of incipient failures was deemed of prime importance. Obviously, minimizing the risk of a sudden and unexpected vehicle breakdown on the operating guideway was a major goal of scheduled maintenance operations.

Some Discoveries

Approximately 24 major experiments were conducted on the model over a two-week period. They ranged from varying the sizes of personnel pools to altering access arrangements of maintenance areas to changing maintenance schedules. Some results are shown in figure 3.

It soon became apparent that the physical arrangement of the maintenance area was going to exert a major impact on system performance. It was desirable to hold down costs by minimizing the length of track, number of switches, and sizes of buildings and other facilities in the maintenance area, and yet the model showed a minimum cost layout was not effective in providing service. Many of these layouts were simulated in a search for the best compromise solution.

An interesting discovery related to vehicle MTBF's. It was found that with certain sizes of maintenance organizations, and within certain ranges of MTBF, vehicle reliability could be allowed to fluctuate quite widely with virtually

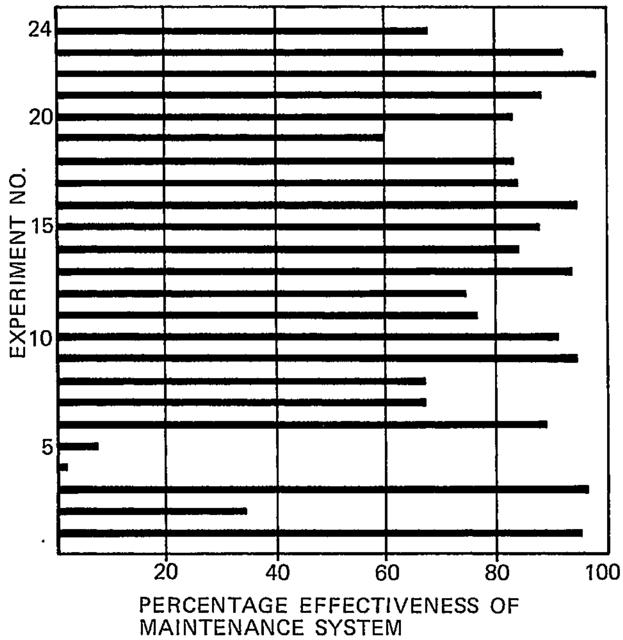

FIGURE 3.

no effect on operational availability. Examination of model detail outputs revealed that the elastic nature of maintenance organizations within particular operating boundaries allowed adjustment to absorb MTBF's varying by a factor of two or more. Excess maintenance system capacity, available for high MTBF's, was progressively employed as vehicle reliability was lowered. Eventually, of course, a point was reached where MTBF matched maintenance system capability, as shown in figure 4. Any further decrease in reliability began to overload the maintenance system and showed up as an immediate decrease in vehicle availability.

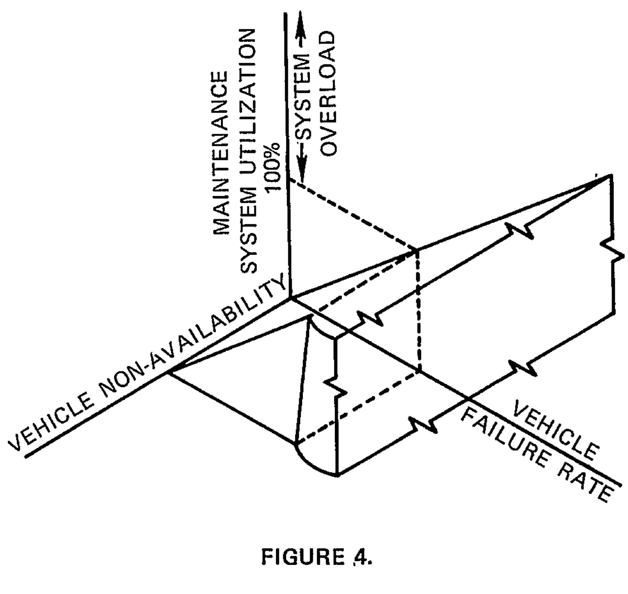

FIGURE 4.

In theory, the most effective plan for integrating vehicle reliability and maintenance system design would lie at the intersection of the vehicle failure line and 100% maintenance system utilization. Money spent to improve reliability beyond this value would not be justified since vehicle availability wouldn't be significantly improved, while decreasing reliability would reduce availability and exceed maintenance system capabilities. In practice, however, it's extremely difficult to reach a perfect conciliation between these elements. This somewhat idealized representation also does not recognize the costs associated with each vehicle failure, even though that failure doesn't adversely affect availability.

Input Data

It's difficult to get good data for a new system that has no counterpart. A prototype test track provided the foundation data, and was supplemented by maintenance information from several transportation systems operating in various parts of the country. An extensive store of aircraft maintenance data was useful in certain applications.

The evaluation of available data, like other aspects of simulation, often requires use of judgment and the employment of simplifying assumptions. Maintainability Engineering literature contains many studies relating elements of maintenance to statistical distributions, and knowledge of these relationships was used to simplify the input data task.[3] As an example, repair times have been shown to follow a log-normal distribution, which reduces the problem of gathering information to one of finding out what mean value applies, since the literature also provides standard deviations for given classes of equipment.[4]

The accuracy and reliability of available information ranged from questionable to good. Maintenance data from some sources results from formal collection procedures in operation for a number of years. Much of this has been subjected to numerous analyses and provides a large data base of known, and generally good, accuracy. Other sources, however, using move casual methods to record maintenance actions, were more valuable for the insights they provided on types and frequencies of various kinds of maintenance than for any quantitative descriptions they provided.

None of the data, however, was expected to fit the proposed maintenance system exactly. The objective was to obtain a rational point of departure for making estimates of parameters. Modifications to basic numbers could be made based on detail design decisions, once analysis of available information revealed approximate values for basic numbers.

Model Development and Use

A basic vehicle maintenance model was constructed during the pre-proposal stage based on very broad ideas concerning the probable form that the transportation system would take. This generalized model was written in GPSS by one man in one month and used approximately 53 minutes of development/checkout time on three different IBM 360-series computers. The fundamental increment of time was chosen to represent two minutes, since this was sufficient for the degree of detail desired and also approximated the time duration expected for the shortest maintenance task to be undertaken. Model size varied from about 100K to 150K, depending on the experiment, and took between three and five minutes of time on a machine equipped for GPSS/360 Version I or II.

The model was developed and operated by a single graduate engineer within the Vought logistics organization, which is the company agency responsible for maintenance planning and analysis. Prior experience as a maintenance technician and Maintainability Engineer was combined with formal training in statistics and computer programming to equip this individual with the necessary tools for the job. The same engineer had previously done aircraft simulation work and was thoroughly familiar with the problems and techniques involved.

The Vought logistics organization contains a number of engineering specialists in the fields of reliability, spares provisioning, manning levels, maintainability, safety, and human factors. Many of these specialists work closely with designers in a consultant capacity, and are cognizant of all design changes. The model operator's personal familiarity with individual analysts within the logistics group was of extreme importance, since they provided information on the actual design as the proposal effort got underway. Requests were received from these analysts to investigate design changes for their maintenance impact or to explore alternative maintenance policies for feasibility. Results of each major simulation were summarized on a standard output sheet as shown in figure 5. These

Maintenance Simulation Output Summary

Simulation I.D. number	____________
Operational availability (Ao)	____________
Vehicle utilization	____________
Maintenance man-hours per operating hour	____________
Mean time between maintenance actions	____________
Number of maintenance actions per day	____________
Average restore time per maintenance action	____________
Percentage time vehicles not operationally ready	____________
Scheduled maintenance rate	____________
Percentage time that traffic peaks were met	____________
Number vehicles short on guideway in "worst case"	____________

Percentage time spent in/on:

1. Guideway ____________
2. Ready area ____________
3. Storage ____________
4. Departure test ____________
5. Scheduled maintenance ____________
6. Not operationally ready status ____________

FIGURE 5.

summaries were distributed to the logistics manager and to the analysts mentioned above.

The logistics manager, in turn, used simulation results in discussions with senior proposal managers to insure that maintenance provisions were fully incorporated into overall system design. These discussions frequently resulted in suggestions for further configuration changes, which were then examined with the model.

Validation and Model Utility

Since the airport transportation system under consideration has yet to be built, and there are no comparable systems in existence, validation of the vehicle maintenance model remains incomplete. Certain indicators of validity are available, however. One of these is that a certain amount of overlapping output data from both the vehicle maintenance model and an operations research traffic model were similar. Although neither model could be verified with actual operational information, their agreement in common output areas was encouraging, the more so as they had been independently developed by persons with different objectives but working from common initial assumptions.

Another factor affecting validity was the process of subjecting the output from each model segment to expert scrutiny as that segment was completed in the initial modeling effort. The rationale was that accurate pieces would make for an accurate whole, assuming that the relationships between segments were adequately represented. Engineering personnel systematically reviewed outputs pertinent to their area of responsibility during model development and rendered judgments on the degree of reasonableness of results. The observation that an engineer can do a more accurate job in estimating some small detail of the logistics problem than he can in trying to make predictions for the complete system provided a logical basis for this kind of incremental proof process.

Realistically, however, an analysis of vehicle maintenance policies is not complete unless it includes all the major elements of maintenance for the entire transportation system.[5] The guideway, communications system, block control and surveillance systems, terminal platforms, and passenger information displays all require spare parts and personnel and form a part of the maintenance problem. Spare parts for these items require warehousing somewhere, and if stored in the maintenance building will increase the size and cost of that facility, at the possible expense of maintenance bays or capabilities. Personnel required to maintain these elements become part of the overall labor force available for the transportation system, and should be included in manning equations.

The model discussed in this paper was concerned with only the vehicle segment of the maintenance problem. It provides a logical basis for expansion into investigation of total system maintenance. Until a complete model has been developed the simulation analysis must be regarded as partial, and information obtained therefrom handled accordingly.

What Did Simulation Contribute?

Much of the model consisted of the subjective estimates and philosophies of many engineers concerned with different aspects of the same overall maintenance problem. The primary contribution of the simulation model was providing a systematic, logical, and orderly way to tie all this diverse mass of expertise together into a coherent and believable whole. The simulation framework gave managers a tool for "previewing" the operations of the vehicle maintenance system before that system was actually built, and for evaluating the overall system impact created by any given detail design change.

As an example of one way to approach the problem of evaluating whether or not a simulation model is worth the resources invested in it, consider the economics involved in answering one of the original questions addressed by the model under discussion: How many spare vehicles should be provided?

An answer to this question can be produced by either an outright guess, or a guess extrapolated from data from some other transportation system, or from a simulation model. Providing more spare vehicles than absolutely necessary raises the cost of the system, and in a highly competitive environment may be fatal to a proposal. On the other hand, insufficient spares may degrade system performance, reduce rider acceptance, and be harmful for future business if the proposal is accepted and the system is built.

The economics of the situation argue strongly for the logic of the simulation model approach. This can be appreciated when it is realized that a very small number of extra cars can easily add a million dollars or more to the price of a system, while a simulation model similar to the one described in this paper costs only a few thousand dollars complete, including development and use. For this expense the manager gets a spare car estimate based on about five simulated years of system operation, and which logically combines all the expert inputs from the whole spectrum of his engineering staff.

Although guesses or estimates from other systems may initially be cheaper to produce than those from simulation models, the degree of confidence reposed in these estimates must be weighed against the consequences of inaccuracy. Given a choice, an estimate based on a logical process must be preferred, especially for new and radical designs.

A Communication Approach

Even when managers understand the advantages conferred by simulation, they will have greater confidence in the technique if the model operator gives some consideration to the question of how best to communicate the results of his experiments. While the use of an output summary sheet to transmit simulation results to managers has much to recommend it, there are several communications problems better approached in other ways.

Information theory emphasizes the need for simple and easily-understood symbology for efficient data transmission.[6] Once it is accepted that print-out sheets, or even summary sheets, contain symbolic information that frequently must be interpreted before it has any meaning for the person who must use it, it becomes simple to choose an output format.

Although films made from simulation may well be the most effective communication tools yet devised they are too expensive for many simulation uses. But a relatively straightforward solution is available if the time and trouble are taken to provide graphic and tabular outputs in the manager's own terminology. Speaking the manager's language will go far toward solution of the symbology problem.

It's desirable to avoid insertion of the simulation analyst between the manager and the data, which is easily done if the required extra effort (and expense) is taken to provide directly the measures needed by managers. For instance, the value for operational availability appearing on the summary sheet was produced from a hand calculation made by the simulation operator using model outputs. With a little extra effort (and expense), this value could have been provided directly as a computer output, as could all of the other summary data. Managers could then receive information straight from the printout sheets, without feeling dependent on a third party for interpretations. In addition, use of GPSS graphs would make it easier for managers to compare the outputs from different runs, and provide more visible impact for the simulation results.

The decision to include cosmetic programming to dress up outputs is one that has to be made on the basis of the resources available for the simulation effort and the degree of familiarity that the managers have, both with simulation techniques and the subject being simulated. This latter becomes important when it is realized that maintainability analysis is a relatively new discipline with measures and terminology uniquely its own, and has not yet been fully accepted as an equal by older and established sciences.[5] Use of simple graphs and tables as direct outputs can be expected to help ease the communication problems when presenting and discussing maintenance-sensitive design parameters with supervision.

Future Uses

As managers gain experience using simulation to evaluate candidate maintenance system designs, they can be expected to become more intimately involved in the simulation effort. Simulations ability to assess alternative approaches is rapidly being recognized as a powerful tool in the search for optimum system configurations. It is likewise evident that a model of a new design has a high likelihood of providing accurate estimates when compared to other estimating techniques.

A model can have even greater significance for follow-on business if second-generation systems are developed using the original system for a basis. Validation or correction of the model, accomplished with data from operation of the actual system, provides a natural departure point for evaluating the derivative designs.

The simulation feature that proved of most value to managers transcended specific numerical outputs. This was the ability of the model to discover which aspects of the particular system being designed had significant impact on the maintenance problem, and which ones were minor. This characteristic may well be simulation modelings greatest contribution to the analysis of new systems. For the first time there is available a methodology that lets a manager see where his attention is most urgently needed in the design process.

Growing appreciation of the importance of utilizing managerial talents in the most effective manner is focusing attention on the use of simulation for this purpose. A technique that allows managers to concentrate on the specific areas of greatest impact while at the same time permitting evaluation of different options in acheiving economically-desirable design is in increasing demand. For example, the physical layout of the maintenance area was mentioned above as having significant impact on the performance of the maintenance system. Once in possession of this knowledge, the logistics manager was able to devote a greater portion of his attention to this item and relatively less to the problem of determining maintenance personnel requirements, which was easily solved. Another area needing close surveillance is the interface between maintenance schedules and traffic schedules. Design of the entire maintenance scheme starts here, since this is the place where maintenance policies will impact the ability of the transportation system to provide service, which is the only reason the maintenance activity exists. The logistics manager thus finds himself as concerned with traffic flow rates and times of occurrence as with vehicle failure rates and spares provisioning, once approximate values have been established for these latter items. The models ability to show dynamically the inter-actions of different aspects of the overall transportation system furthers the awareness of managers and others for the need to carefully integrate design activities.

It is these characteristics that have established simulation as a management tool of major importance.

BIBLIOGRAPHY

1. Maintainability Engineering, Department of the Army Pamphlet 705-1, Headquarters, Department of the Army, June 1966.

2. General Purpose Simulation System/360 OS Version 2 User's Manual, IBM Program Product SH20-0694-0, October 1969.

3. Goldman, A. S., and T. B. Slattery, Maintainability, A Major Element of System Effectiveness, John Wiley and Sons, Inc., New York, 1967.

4. Maintainability Engineering Handbook, NAVORD OD 39223, Commander, Naval Ordnance Systems Command, February, 1970.

5. Blanchard, B. S., and E. E. Lowery, Maintainability, Principles and Practices, McGraw-Hill Book Company, New York, 1969.

6. McCormick, Ernest J., Human Factors Engineering, McGraw-Hill Book Company, New York, 1964.