A Knowledge-Based System for Transit Bus Maintenance

PETER WOOD

ABSTRACT

A knowledge-based system for bus maintenance has the potential for improving the efficiency and effectiveness of bus maintenance by making the knowledge of the most highly skilled maintenance personnel available throughout the transit industry. A knowledge-based system is a practical application of the research that has been performed on artificial intelligence. These systems have been developed for a wide range of applications, including the diagnosis of diesel-electric locomotive problems. The four basic elements of a knowledge-based system are described, and an example is provided of the application of such a system to bus maintenance. From a review of the impact on performance of two other techniques for bus diagnosis, spectrochemical oil analysis and the New York City Transit Authority's Automated Bus Diagnostic System, it is concluded that both a reduction in the time required for diagnosis and an increase in the quality of maintenance, measured in terms of a reduction in road calls, would be achieved. It is recommended that a prototype system be developed so that both the costs of implementation and the savings that would result can be determined.

Effective and efficient maintenance is a key element in the operation of a transit system. Poor maintenance will have an adverse impact on both costs and revenues, because the public has demonstrated quite clearly that reliable service is essential if ridership is to be sustained or increased. Although there has been a considerable improvement over the last decade, the quality of bus maintenance still continues to be a problem. One concern that faces the industry now is one that has faced it for almost two decades--the shortage of adequately trained labor. A combination of unattractive working conditions, uncompetitive salaries, and frequently the requirement that all new employees, irrespective of experience, start at an entry-level position has made it difficult to recruit and employ the highly skilled staff necessary for an efficient maintenance organization.

Many steps have been taken to alleviate this problem. Several transit systems have established training programs to raise the level of performance of their staff. New and improved maintenance manuals have been introduced by both industry suppliers and individual transit systems. Special programs dealing with particular systems such as engines, transmissions, and air conditioning have been sponsored by the American Public Transit Association (APTA). UMTA's Technical Assistance Program has made a wealth of material available to the transit industry. TRB also has sponsored meetings on maintenance research and development and effective tools and techniques (1-2).

One common issue that pervades all the discussions on maintenance needs has been that of better information exchange. In spite of the efforts that have been made to date, there is a general feeling that better information exchange would improve the effectiveness of maintenance. In this paper an approach to information exchange is discussed that combines the latest in computer techniques with the most expert knowledge of maintenance procedures that exists within the industry.

Expert, or knowledge-based, systems are a branch of artificial intelligence. They are now being used in many applications to facilitate the transfer and use of expertise that has been built up over a period of years by one or more highly qualified practitioners. Only recently, the use of an expert system to assist in the diagnosis of diesel locomotive problems was announced. The question that needs to be answered is whether there is any role for this approach in the area of transit bus maintenance.

COST OF BUS MAINTENANCE

The costs of maintaining a bus are difficult to establish, because virtually every transit system has a different basis for allocating costs. APTA has indicated that vehicle maintenance expense for all modes of transit is in excess of $1.1 billion a year, or approximately 19 percent of the total cost of operations. Of this sum, more than one-half is attributed to wages (3). In a more detailed analysis by Etschmaier (4) devoted to bus operations, vehicle maintenance was found to cost $0.526 per vehicle mile. This is close to 20 percent of the average cost per mile for bus operations and amounts to $16,400 per vehicle per year. Etschmaier compared this with a maintenance cost of between $2,100 and $3,000 per year for trucks. The comparative annual mileages were 31,200 for buses and 44,000 for trucks. Although the differences in the operating profiles for the two types of vehicles limit the value of any further comparison between their maintenance costs, this comparison suggests that there is room for improvement in bus maintenance efficiency.

One measure of performance that provides an indication of the effectiveness with which bus maintenance is being performed is the number of road calls that occurs. Again, precise comparisons between transit systems are difficult because of the different definitions used by each transit agency. However, data prepared by APTA (5), shown in Table 1, clearly reveal the great differences that occur between agencies in the number of road calls experienced. The lowest mileage reported between road calls for mechanical failures is 672 miles--about once a week. The median values range between about 2,500 and 3,700 miles (around once a month), whereas the highest mileage reported is about 38,000 miles--less than once a year. Although the 60:1 range of

TABLE 1 Miles Between Road Calls

Population of City	Lowest	Median	Highest
Less than 200,000	1,522	3,116	25,707
200,000-500,000	750	3,701	10,257
500,000-1,000,000	1,294	2,488	28,920
More than 1,000,000	672	2,665	7,652

mileages is undoubtedly due to the differences used for defining a road call, it is obvious that there is a big difference between the lowest and the highest mileages that are being achieved at this time.

Road calls are a particularly appropriate measure of the quality of maintenance because each road call represents an unscheduled activity and frequently an interruption of service. On the one hand, there is an unnecessary expense, on the other, a failure to provide a reliable service. Both have an adverse effect on the transit system. Although a dollar value can be expressed for the cost of a road call (even though it is difficult to find a transit system that can report what the cost is), it is more difficult to express the adverse impact on transit revenues of unreliable operation. Because of this, frequently a comparison is only made between the costs of road calls versus the additional maintenance costs that could be incurred in an attempt to reduce them. Obviously some balance must be maintained. The costs of maintenance of a system that prided itself on having no road calls would probably be so high as to make the size of subsidy necessary to cover the costs completely unacceptable. On the other hand, a system that did the absolute minimum of maintenance and was prepared to accept a large number of road calls would offer such a poor level of service that it would be completely unacceptable to the public.

What would be desirable would be an approach that would result in a reduction in road calls without an increase in maintenance costs. This is what the transit industry is doing in its current efforts to improve the quality of its workforce through improved training. A knowledge-based system is an attempt to improve the quality of work performed by the existing labor force by making the knowledge of the most experienced maintenance staff immediately available. All skills are, by and large, the result of the transfer of knowledge. A knowledge-based system uses modern computer technology to facilitate this process.

FUNDAMENTALS OF A KNOWLEDGE-BASED SYSTEM

A knowledge-based system is a practical application of the research that has been performed over the last 20 or more years on artificial intelligence. Although knowledge-based systems use sophisticated computer and programming techniques, it should be made clear that the power of a such a system lies in the knowledge that it contains rather than in the hardware and software that make that knowledge available.

The history of knowledge-based systems dates back more than 20 years, but it is only in the last decade that they have been considered for general use. Systems are now being used in a wide range of applications, varying from the diagnosis of internal diseases to the configuration of computer systems. Other applications include air traffic control, simplification of symbolic mathematical expressions, and detection of geological deposits. More relevant to our immediate concern, General Electric Company is using a knowledge-based system to diagnose faults in diesel electric locomotives (6). Recently, General Motors invested in one of the leading companies engaged in the development of knowledge-based systems, so it is likely that in the future this approach will be applied to automotive diagnostics.

At one time it was assumed that the ability of a computer to calculate at a high speed would allow problems to be solved by a process of evaluating all the possible alternative solutions. It soon became apparent that the combinatorial explosion of possibilities that arose with even the simplest of problems made this approach completely impractical. At the same time, these problems were being solved on a daily basis. The game of chess is computed to have 10^{120} possible solutions, and yet chess players of all levels play and win each day. The answer to this apparent paradox is the introduction of heuristic rules, or rules of thumb, that human experts have developed to solve such problems. A simple chess problem such as knight and king against rook and king has roughly 20 million possible configurations. Michie has shown that only 30 heuristic rules are necessary to achieve expert performance in solving such a problem (7).

The use of heuristic rules is the basis of a knowledge-based system. As Feigenbaum, one of the pioneers in knowledge based systems, states (8):

> An expert system is an intelligent computer program that uses knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution....
>
> The knowledge of an expert system consists of facts and heuristics. The "facts" constitute a body of information that is widely shared, publicly available, and generally agreed on by experts in a field. The "heuristics" are mostly private, little discussed rules of good judgement that characterize expert level decision making in the field.

A knowledge-based system manipulates symbols rather than numerical values. Because of this, programs are generally written in LISt Processing language (LISP). A number of higher-level languages have been written that facilitate building a knowledge-based system, including EMYCIN, produced by Stanford University, and ROSIE, by the Rand Corporation. Many other programming systems have been developed by private industry.

Figure 1 shows the basic structure of a knowledge-based or "expert" system. It is made up of four major elements:

- A knowledge base--the "facts" referred to by Feigenbaum,
- An inference engine--the "heuristics,"
- A control section for handling the overall operation, and
- An input-output section for communication with the user--generally in regular English.

The knowledge base contains the basic data regarding the topic that is the subject of the expert system. For example, if the expert system was concerned with diagnosis of faulty engines, the knowledge base would contain such data as the number of cylinders for a particular class of engine, nominal pressure measurements, idling speed, injector settings, and so forth.

The inference engine contains a set of rules that it has been determined are used by the expert when dealing with his area of expertise. These are ar-

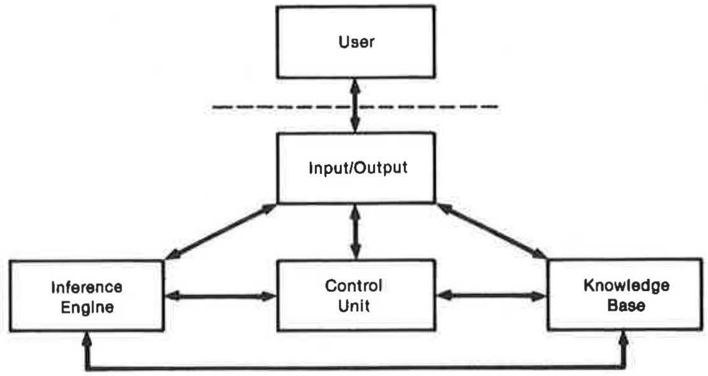

FIGURE 1 Knowledge-based system.

ranged to be encoded in a form IF [conditions] THEN [actions]. A simple example for a gasoline engine would be

> IF [the engine does not start] AND [the starter turns the engine] AND [the fuel gage does not show EMPTY] AND [the accelerator pedal has been depressed several times], THEN the engine has been flooded. The accelerator must be depressed fully and held down while the engine is started. Release the accelerator once the engine fires.

The control section is responsible for controlling the strategy to be used by the system.

Two basic strategies of problem solving are used by knowledge-based systems--data driven (or forward chaining) and goal driven (or backward chaining). The first approach starts with the initial set of data and examines the most likely alternatives. Where additional data are required to resolve an ambiguity, the user is requested to provide the necessary information. In the ideal case, a unique solution is obtained for a given set of data.

The goal-driven approach starts by assuming a certain solution and then works back to determine whether all the data are consistent with this solution. When an item of data is detected that renders the initial hypothesis invalid, an alternative solution is investigated.

In practice many complex systems rely on both approaches. This is particularly helpful when the data are inadequate. A combination of data- and goal-driven search is used, and the approach that most nearly provides a direct and unambiguous connection between data and solution is considered to be the most probable.

Knowledge-based systems have a number of characteristics that make them particularly suitable for applications such as maintenance diagnostics. One of these is the ability to operate successfully in the presence of data that are incomplete or imprecise or both. There does not have to be a complete and unambiguous path from the data to the solution.

Knowledge-based systems also have the capability to explain the reasoning that led to the preferred solution and indicate the probability that it is the correct one. It is accepted that a system will not always provide the correct answer (its capabilities are only as good as those of the experts whose knowledge it contains), and allowing the user to examine the approach that the system followed in obtaining a solution will frequently indicate what other steps might be introduced in the system to improve its performance.

This in itself is an important feature of a knowledge-based system. The structure is such that reprogramming is not necessary in order to modify and improve its performance. Most knowledge-based systems are continually expanding both the knowledge base itself and the heuristics in the inference engine in order to refine the performance of the overall system.

The design of the input-output section of a knowledge-based system is critical to its success. The majority of users are not usually experienced in the use of computers and yet are relatively expert in their own field. It is highly desirable that any information that must be supplied by the user be requested in a form to which the user can relate, and the output must similarly be provided in the specialized language of the user. Up to this time, a conventional computer keyboard has been the most common input mechanism for data that the user must supply; with the advent of techniques such as the touch screen and the "mouse," it is probable that in future systems most input requirements will be handled without requiring any typing skills on the part of the user.

APPLYING A KNOWLEDGE-BASED SYSTEM TO BUS MAINTENANCE

A knowledge-based system for bus maintenance will differ somewhat from other systems in that both the knowledge and the heuristics will be derived from the composite expertise of a number of sources. In the initial stages of development of a knowledge-based system, however, it is usual to concentrate on one subsystem that can be used in the development of a prototype. After the potential of the system has been shown through application of the prototype to real-world problems, the knowledge base can be extended to cover the balance of the system. At least in the early stages, therefore, it will probably be possible to develop an effective prototype based on the experience of one or two experts working on one particular subsystem--preferably one that is generally considered to present significant difficulty in the attempt to diagnose the cause of specific failures.

The development of a prototype system should concentrate on the diagnosis of difficult faults. Examples would be those that take a considerable amount of time to diagnose under normal conditions or cases where a fault is reported frequently and repetitively, even though the correct repair action is supposed to have been taken.

Development of a knowledge-based system is an interactive process between the software developer (or knowledge engineer) and the expert. Initially an attempt is made to determine the logical processes that the expert follows when he makes a diagnosis and the information that he requires. Frequently diagnosis is made through a process of elimination, and this must be included in the heuristic. This is a most important feature, because much of the power of a knowledge-based system is the ability to prune the large number of possibilities that exist into a small and well-defined set that can then be investigated in more detail. Frequently it is not possible to come up with a unique solution, and it is then necessary to develop a weighting procedure that will allow the alternatives to be presented in order of probability. At all times the procedure that led to a decision must be available, so that the expert can review the process to confirm that it reflects his normal diagnostic process.

After an initial set of heuristics has been established, it must be refined by practical application in a real-world situation. When a correct diagnosis has been obtained, the procedure can be annotated to show that, for the given set of conditions, the procedure was correct. Where an incorrect diagnosis occurred, or the knowledge-based system was unable to reach a conclusion, the decision process must be analyzed and additional heuristic rules of thumb or data introduced to provide the correct solution. Finally, after an acceptable level of accuracy has been achieved, the system can be put into general use. Even then the refinement process continues, because cases of incorrect diagnosis will continue to occur, and the expert must again be consulted to improve the system.

As an indication of the accuracy that can be achieved, a diagnostic system for pulmonary diseases has achieved an accuracy in initial diagnosis of approximately 80 percent (9). Although this might not appear high, it is slightly better than the accuracy of a specialist in the disease and almost 20 percent higher in accuracy of initial diagnosis than that of a general practitioner.

Fortunately, the high degree of commonality in the subsystems that are in common use in the transit industry (even with buses from different manufacturers) would make a knowledge-based system generally applicable. Also, in contrast to most other computer-based systems, the value of a knowledge-based system would probably increase as the size of the transit system decreased, because the smaller transit systems are less likely to have the in-house experts that are available at larger agencies. For this reason, it may be desirable to incorporate diagnosis of less obscure problems in a system so that it would be of more use to small transit agencies.

Although modern computer systems are both smaller and less sensitive to their environment than their predecessors, they are still generally not ideally suited to an industrial or workshop environment. (Computers capable of working under adverse environmental conditions are generally available, but usually at a considerably higher price.) This does not appear to be too much of a problem for most transit applications, because it is likely that any computer system would be installed in a relatively benign environment such as a supervisor's office rather than on the shop floor. Alternatively, the computer itself could be installed in a separate room, with ruggedized terminals available for use by the maintenance personnel.

Experience in other applications has shown that although there may be some initial resistance to such a system, the benefits soon become obvious and the users become the strongest proponents. Probably the only exception is in the medical profession, where fears have been publicly expressed that the introduction of an expert system for diagnosis would make the necessary skills too widely available, adversely affecting the specialists in the field.

USER INTERACTION WITH THE SYSTEM

An expert system normally uses a dialog with the user in order to determine the information on which it can base its decision. In order to minimize the skills required of the user, a multiple-choice menu is provided so that the user is only required to answer yes or no or key in a number. Even this is not required with the most recent installations; the user is only called on to place a cursor over the appropriate answer by using a joystick or a mouse.

With the computer display indicated by () and the response by [], a typical sequence of operations would be as follows:

(What is the bus number?)

[User inputs bus number.]

(You have put in bus number xxxx. Is this correct?)

[User types in "y" for "yes" or "n" for "no." If the answer is no, the computer repeats the first question.]

(What is the mileage?)

[User keys in the mileage.]

(You have entered xxx,xxx miles. Is this correct?)

[User types in "y" for "yes" or "n" for "no." If the answer is no, the computer repeats the question.]

(What is the problem?)

At this point the computer would show a list of the most likely problems, such as:

1. Brakes,
2. Air conditioning,
3. Low power,
4. Starting.

Because it is unlikely that all the possible problems could be displayed on a single list, the final item would be "Display next list." If this was selected, a further list of problems would be displayed.

When the user selected a problem, another list of questions would be asked appropriate to the particular problem selected. It is feasible, in many cases, to answer "don't know." After a certain number of questions had been asked and answered, the computer would provide a message saying:

(Please wait. I am working on the problem.)

It is quite possible that the computer will find that it does not have all the information it needs to make a decision, in which case it would ask for the user to provide it. Eventually, however, a probable solution would be obtained, and the computer would display a message, such as

(Since you say the freon pressure is correct, the most likely cause is a slipping clutch. Do you know how to check that the clutch is working correctly?)

If the user inputs "y," the computer could display:

(If the clutch is operating correctly when you check it, we will have to look for the next most likely cause. When you log on, type "c" to continue this diagnosis.)

It is also possible for the computer to provide additional information that could be helpful. For example, one output could be

(The problem is most likely fuel dilution due to transmission fluid entering the sump through a faulty bearing seal. Bill Smith of XYZ had this problem with the same model of engine and transmission in 1984. You can call him at xxx.yyy.zzz for additional information.)

This is typical of the type of information that a knowledge-based system can hold in its data base to assist the user.

IMPACT OF KNOWLEDGE-BASED SYSTEMS ON TRANSIT BUS MAINTENANCE

Although no knowledge-based system has been developed for transit bus maintenance, General Electric has developed a system, CATS-1, for diagnosing diesel engine malfunctions. It can display locomotive components on a graphics display and demonstrate repair procedures on a video monitor. CATS-1 uses both forward and backward chaining, and is based on a knowledge base of diagnostic rules and facts derived from their top locomotive repair expert (6). Unfortunately, no information on cost or performance is available at this time, and because there are no other systems directly comparable with that required for transit bus maintenance, it is difficult to obtain a quantitative assessment of the impact of such a system on transit maintenance. However, at least two other diagnostic techniques have been used in the transit industry. These are spectrochemical oil analysis (10), used by a number of transit systems, and the Automated Bus Diagnostic System that has been demonstrated by the New York City Transit Authority (11). Experience with these techniques provides an indication of the benefits that might be expected from knowledge-based systems.

Spectrochemical oil analysis of engine defects is based on analyzing the chemical properties of a sample of oil drawn from the sump of a diesel engine (or transmission when transmission defects are being analyzed). An expert can, using his experience, predict with a high degree of accuracy the probability and type of failures that are likely to occur on the basis of the trace chemicals that are detected in the analysis. This allows the appropriate corrective action to be taken before a catastrophic failure, leading to a breakdown, occurs. One of the earliest studies on the application of spectrochemical oil analysis to the transit industry clearly showed the benefits of such a system (10).

Spectrochemical oil analysis was started as a normal preventive maintenance procedure at the Autoridad Metropolitana de Autobuses (AMA) in San Juan, Puerto Rico, in December 1970. Each coach was scheduled for a minimum of one sample per month. Figures 2 and 3 show the trends in the analyses of the oil. A "red" report indicates a critical abnormality; a "white" report, no abnormality. It can be seen that the number of buses that showed critical abnormalities initially increased and then subsequently fell, whereas the number of buses with no abnormalities steadily increased. Although there is no direct evidence that shows a correlation between the improved diagnosis of potential failure conditions and a reduction in road calls (unfortunately, road call data for this period were unavailable), the reduction in the number of buses with critical abnormalities and the increase in buses with no abnormalities suggests that an improvement in the reliability of operation also occurred.

The Automated Bus Diagnostic System, demonstrated in New York in 1982, is even more closely related to a knowledge-based system in that the data obtained from an instrumented bus are analyzed to diagnose the cause of problems that have been detected in an earlier check performed at the service island. The diagnosis is purely automatic, but is based upon that of experts in bus maintenance (11).

Unfortunately, deficiencies in the data collection process make it difficult to draw conclusions that would be considered statistically valid. Nevertheless, some general conclusions can be made based on

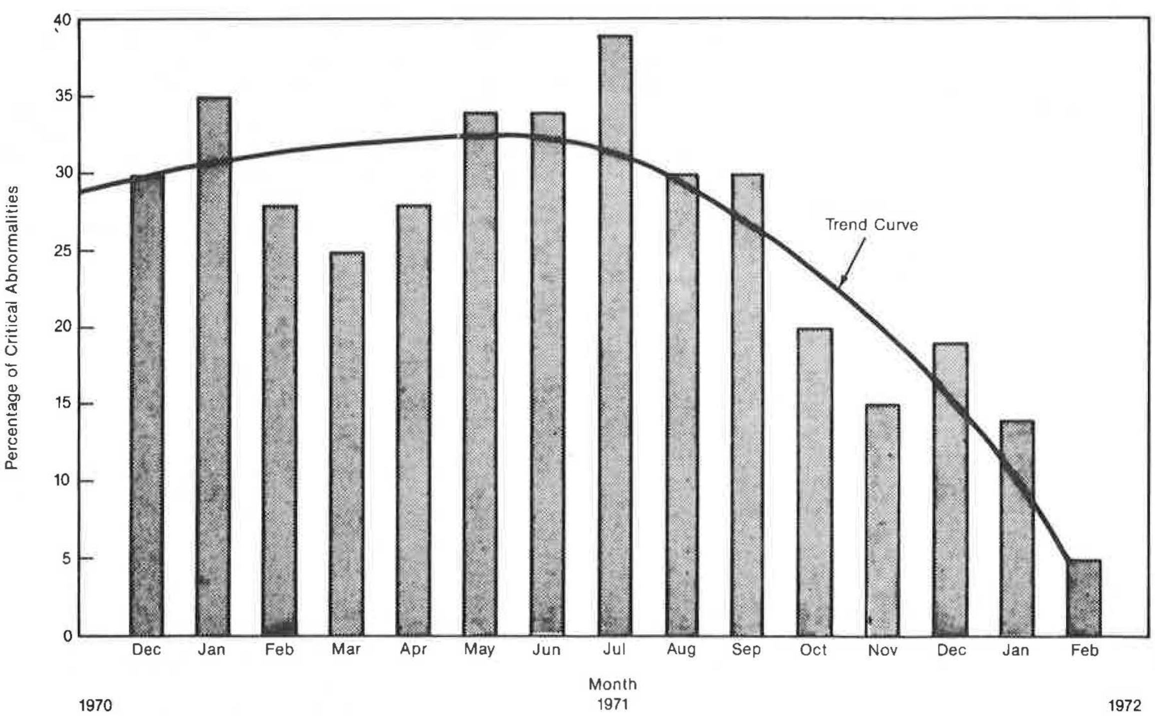

FIGURE 2 Summary of red reports for AMA.

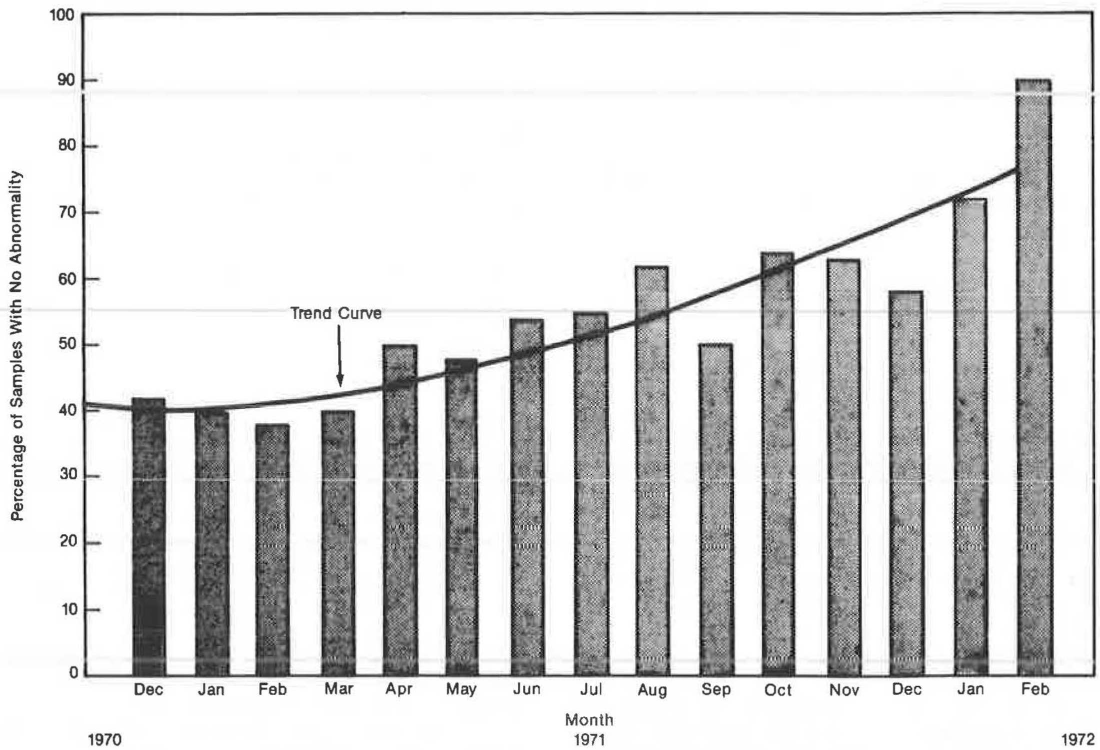

FIGURE 3 Summary of white reports for AMA.

the data that were collected. First, as expected, there were fewer road calls for the buses that were being evaluated on a daily basis by the automated diagnostic equipment than for the control group that were receiving the normal manual inspection and repair (Table 2). Second, the buses in the test group were out of service for more days for repair than the control group. This suggests that defects were being detected by the diagnostic equipment and repaired, which led to the lower number of road calls. Significantly, though, at the end of the 6-month test period the test group had fewer days out of service for repair than the control group, which suggests that all the incipient failures had been detected.

TABLE 2 Number of Road Calls and Out-of-Service Days

	Number of Road Calls		Out-of-Service Days	
Date	Test Group	Control Group	Test Group	Control Group
February 1982 (part)	8	11	86	14
March	6	18	250	42
April	16	21	181	74
May	15	22	84	86
June	12	33	86	134
July	21	16	51	157
Total	78	121	738	507

No measurement was made of the time necessary to diagnose a fault with the automated equipment versus that with conventional expertise. However, individuals were asked to estimate the time that an expert would take to diagnose a fault. This was compared with the time taken using the automated diagnostic equipment. In every case there was a significant time saving. On the basis of the experience in New York, an improved version of the Automated Bus Diagnostic System will be tested at four other transit systems, including Flint, Michigan; Syracuse, New York; and Nashville, Tennessee (12).

It can be seen that there is a strong indication that the introduction of a knowledge-based system would benefit transit bus maintenance, both in terms of fewer road calls and shorter times for diagnosis. Without more directly related experience, however, it is not possible to determine whether the reduction in costs of road calls and diagnosis would offset the additional repair costs that would apparently result, and also whether the costs of operating and maintaining the system could be offset by savings in the overall maintenance operation. Because of the potential benefits, however, there is good reason to develop and test a prototype system so that both maintenance cost savings and improvement in the quality of performance can be determined as well as the costs of operation and maintenance.

REFERENCES

1. Bus Maintenance Improvement. Special Report 198. TRB, National Research Council, Washington, D.C., 1983.
2. Proceedings of the 1984 Bus Maintenance Productivity Workshop. Transportation Research Circular 289. TRB, National Research Council, Washington, D.C., 1985.
3. Transit Fact Book. American Public Transit Association, Washington, D.C., 1981.
4. M.N. Etschmaier. Transit Bus Maintenance Management: Assessment and Needs. Staff Study SS-62-U.3-01. Transportation Systems Center, U.S. Department of Transportation, Cambridge, Mass., Oct. 1982.
5. Transit Operating Statistics for Calendar/Fiscal Year 1982, Vol. 2. American Public Transit Association, Washington, D.C., Oct. 1982.

6. Computers That Think Like Experts. High Technology, Jan. 1984, p. 30.
7. D. Michie. Knowledge-Based Systems. Report 80-1001. University of Illinois at Urbana-Champaign, Jan. 1980.
8. E.A. Feigenbaum. Knowledge Engineering for the 1980's. Computer Science Department, Stanford University, Palo Alto, Calif., 1982.
9. F. Hayes-Roth, D.A. Waterman, and D.B. Lenat (eds.). Building Expert Systems. Addison-Wesley, Reading, Mass., 1983.
10. D.F. O'Sullivan. Spectrochemical Oil Analysis. Report DOT-UT-10005-OS1. Mitre Corporation, McLean, Va., Jan. 1973.
11. R. Casey. The Automated Bus Diagnostic System Demonstration in New York City. Report UMTA-IT-06-0102-83-1. Research and Special Projects Administration, Transportation Systems Center, U.S. Department of Transportation, Cambridge, Mass., Dec. 1983.
12. Automated Bus Diagnostic System. UMTA Technical Assistance Briefs, Vol. 3, No. 2, Spring 1984.

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

Information Systems in Bus Fleet Management

ELLIOT I. GITTEN, T. H. MAZE, ALLEN R. COOK, and UTPAL DUTTA

ABSTRACT

The rate of use of computerized fleet management information systems in the transit industry is investigated, and it is concluded that transit fleet management underutilizes automated systems. Next the evolution of the state of the art of fleet management systems is examined. Directions are hypothesized for the future development of computerized systems. Last, elements that should be considered in the planning of a system are discussed.

Computerized management information systems were first introduced to industry in the late 1950s (1), typically for accounting and other applications-oriented activities (such as inventory tracking and labor and payroll reports). By the mid-1960s many industries recognized the need for and value of information in a broader sense. At that time computer technology had evolved to the point where large, integrated systems were feasible. Since then, information systems have had more than 20 years to develop and they have permeated most aspects of business and government.

One would expect that bus fleet management would be no different than similar areas of industry and that record keeping and information preparation would be largely computerized by now. As an indication of the rate of computerization of bus fleet information in the transit industry, Kliem and Goeddel summarized the results of a 1980 American Public Transit Association (APTA) survey of computer applications at transit agencies (2,3). They found that "of the 54 transit properties identified (representing approximately 65 percent of the total industry's vehicle fleet), 28 reported the use of automated information systems for vehicle fleet maintenance." Slightly more than half does not indicate an overwhelming rate of computerization, but by 1984, certainly more have become computerized. Even the 1980 rate (52 percent of those surveyed) suggests that computers have a strong foothold in bus fleet management.

After a closer look at the list of systems claiming computerized maintenance information systems, one with which the authors were familiar was spotted that appeared odd. This transit system was a department of the city's government. A clerk on the bus maintenance staff retyped the information from work orders into text files on the city's mainframe computer. The text files were used to produce paper copies of work histories, but they were never machine processed for summary information. Therefore, the computer was acting largely like an electronic file cabinet. Further, because the records were not machine analyzed, there was no need to be totally accurate in data entry and repair cases were often lost. Technically this system kept maintenance records by using a computer, but this could hardly be termed an information system. Unfortunately, the APTA survey used by Kliem and Goeddel is insensitive to the degree of computerization of record keeping. Therefore, the rate of computerization in fleet management is probably better measured by whether the system uses computers as well as the degree of sophistication of the use of computerized systems.

A better indication of the rate of computerization is given in the results of a 1983 survey (more than