SEMINAR ON SIMULATION

INTRODUCTORY COMMENTS

J. F. Nunamaker

The comment has been made that what we really need is a design theory for demand-actuated transportation systems (DATS). The following approach is suggested as one way of viewing the DATS design problem.

The DATS design process has a number of similarities to any physical design process, such as a production plant or a bridge. In each case there must be an initial recognition of a need. Next, preliminary studies are conducted in which major alternatives are considered, the technical feasibility is determined, and costs of alternatives are estimated. If a decision to proceed is made, the requirements must be stated in sufficient detail for designing the system. The design phase consists of preparing a set of specifications (blueprints) that are detailed enough for the construction phase.

The major functional activities are decision points in the design process of DATS are shown in Figure 1. After the requirements have been documented, the systems designers consider the equipment available (or equipment desired) and any constraints (such as the existing system) on the design activity. The design phase consists of producing the specifications for the 3 major parts of the system: transit equipment (characteristics and type of vehicles); transit facilities (loading and unloading facilities); and system scheduler.

The specifications must be detailed enough to verify feasibility and to evaluate the performance of the proposed system but only detailed enough to specify construction because producing specifications that are too detailed is costly; and they may have to be changed in any case. In addition, specifications that are too detailed tend to bind the design unnecessarily at too early a stage with negative payoff.

An important aspect of this formulation of the DATS design problem is that it should include an explicit statement of the performance criterion by which performance of the system is measured. A consequence of including performance measures is that the emphasis is focused on the overall performance of the system rather than on any one part. DATS design is defined as the process of producing design specifications necessary for the construction of the system from a problem statement and knowledge of the capability of the components of the system.

Simulation is one way of evaluating alternative designs. It should be kept in mind that simulation is the least desirable solution technique; however, it has the greatest applicability. The major classifications of solution techniques from most desirable to least desirable are as follows: analytical, iterative (mathematical programming), branch and bound, enumeration, heuristic, and simulation.

Simulation is, however, the most applicable solution technique available for handling large unstructured problems. Because we are concerned with simulation models, it is necessary that sensitivity analysis be discussed with respect to the various models. It is not uncommon for the systems designer as decision-maker to be presented with point estimates of the uncertain parameters to be used in his analysis. Sensitivity analysis

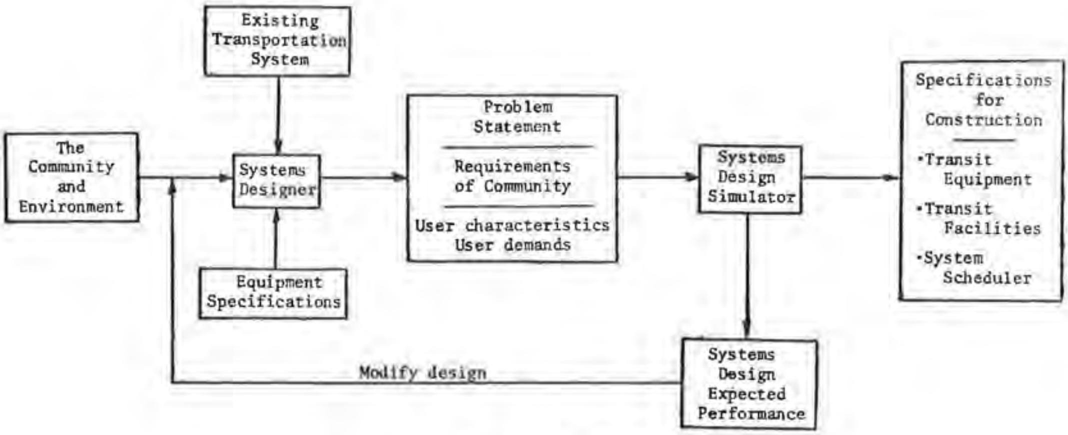

Figure 1. The DATS design process.

is one method of demonstrating the effects of unknown variances in these uncertain data and of identifying their most crucial elements.

The purpose of sensitivity analysis (or parametric analysis) is to show how the results of an analysis can change when either the data change or when the assumptions imposed on an analysis or model are altered. Treating the parameters in this way is a meaningful method of expressing uncertainty to the decision-maker.

Therefore, when reporting on simulation models, one should comment on the extent to which sensitivity analysis was used, the validity, and the testing of the models. In addition, one should indicate the decision-making structure of the models as described by (a) specification of the inputs and outputs, (b) specification of decision variables and determination of feasible alternatives, (c) selection of an objective function, (d) expression of objectives as a function of decision variables, (e) explicit statement of constraints that limit the value of the decision variable, and (f) solution, i. e., determination of the values of the decision variables.

There are other questions and topics that should be discussed. Was a data management language used for the analysis of survey data? Because of extensive questionnaire analysis, we are confronted with a large data base problem. To aid in the handling of large data bases, data management language has been developed during the past 8 years (1, 2, 3). A data management language is quite useful in situations where one is not sure what questions will be asked about the data. What computer configuration was used to run the simulation model? What language was used to program the model? Describe the scheduling algorithms. Discuss the implementation problems with respect to the scheduling algorithms. What thought has gone into the problem of selecting a computer for the implementation of DATS?

The selection of the "right" computer configuration is itself a very difficult problem in well-structured problems such as routine business data processing. There are the additional problems of selecting a computer that a community can afford for a real-time DATS. The DATS could be well designed for a community, but the entire system might fail because of the selection of the wrong computer. The first discussion is of a DATS simulation model developed by General Motors and described in a number of reports (4, 5, 6).

GENERAL MOTORS DATS SIMULATION MODEL

Larry Howson

The General Motors DATS simulation was displayed on an IBM Model 2550 Model 3 graphic terminal connected to an IBM 360/67 time-shared computer. A list of items such as number of buses in the system, number of seats on a bus, the velocity of buses in the various networks, and sizes of the zones in the simulation can be modified at the display terminal. In addition, the algorithms for routing and keeping track of buses can be modified at the display terminal. The number of demands per hour can also be displayed. A scale model of the area in which the bus is being routed is displayed on an upper corner of the scope. The route of one bus is on display continuously while the simulation is in progress, and the provisional tour is displayed and changed. A flashing or flickering occurs on the screen to indicate that a customer has been accommodated. The customer has been either picked up or delivered at his destination. Some additional numbers are also displayed to indicate the time and the number of vehicles in the system. The simulation can be interrupted at any point, and all of the tours can be observed. The tours of all the buses currently in the system can be displayed on a point-by-point basis. The points are labeled with P's and D's to indicate either a pickup point or a delivery point. The systems designer has the option of observing the tour of a single bus or the tours of all buses in the system. The system designer also can "dump" out data on each of the buses and analyze who is going where and why. In addition, he can check the constraints for each rider on each bus.

The display starts to get messy when approximately 10 or 11 buses are displayed in the system. However, you can get a mental picture of what a manual dispatcher might be looking at if he were trying to schedule all of the buses.

There are 5 IBM 2550 graphic terminals being supported on the IBM 360/67 plus a background partition. Therefore, the simulation does not run as fast as we would like, but it is not too slow for our purpose. The simulation with 100 demands to be satisfied required approximately 80 seconds of CPU time and required about 7 minutes of real time. The simulated time for this example was 1 hour of operation.

Karl Guenther

Do you need the ability to pull out the history of one vehicle and look at what it has done for the previous hour?

Larry Howson

We have not done that. We are preparing to punch a set of cards every time a demand is served so that at the end of the simulation we can evaluate the bus route. One of the ideas we have for demonstrating feasibility is to take one bus route from the simulation and drive a real bus along that route. In this way we can see whether the route is feasible and whether it indeed can do what we had proposed.

Karl Guenther

Did you present the same problem to a manual bus dispatcher and let him generate routes? Did you attempt to compare the computer-generated routes with manually generated routes?

Larry Howson

We have not done that at this time.

Eugene T. Canty

We analyzed one approach to manual dispatching for the simulation that was based on a particular case study involving 36 square miles. For a low demand, we had a few buses spread over a large area. We looked at the 3 closest buses to a customer, and that is the way we determined our cost based on somewhat inefficient scheduling. From this case study, it was determined that more buses were needed to meet the same requirements.

Larry Howson

There are essentially 2 ways that this scheduling problem could be structured. We could use the objective function technique that is similar in many respects to what has been done at M. I. T. The other approach is to use a technique similar to one that was developed at Northwestern University and that essentially looks at the closest bus. We said, "How might a manual system operate?" and we concluded that it might operate more like the Northwestern simulation rather than like the M. I. T. As a result it was proposed that we might be able to do somewhat a better job by looking at the first 2 or 3 closest buses and evaluating the first few positions in their tours. This was thought to be an enhancement of the approach to looking at the closest bus. This approach is very similar to the selection of routes by using strings and pins. The next step is to choose from the 2 or 3 closest buses 1 bus that looks best from the customers' constraints. In effect, we simulated a manual system in that respect. For certain dimensions of the problem, this approach was almost as good as a complete analysis of the objective function. Potentially it may be cheaper because fewer buses are scanned by the program.

J. F. Nunamaker

How long does it take to make the next decision on a route?

Larry Howson

That depends, of course, on the size of the system and how many buses we have to scan to make the decision. Some typical numbers are 56 to 60 milliseconds to scan a bus. If we have about 100 buses in the system, it would take several seconds to make a decision. If there are only 10 buses in the system, the time is considerably less. This seems to be a straight line function and seems to support the notion that there is no change of scale.

J. F. Nunamaker

What scheduling techniques are used in the simulation?

Larry Howson

The objective function of the scheduling model considers both waiting time and travel time. The scheduler evaluates the waiting time and travel time of customers who have requested service and who have not been assigned as well as the waiting time and travel time of those customers who have been assigned. In addition, the scheduler also considers the waiting time of those people who have been assigned but not picked up and customers who have been assigned and picked up but may be delayed. Weights are assigned to each of these parameters. The weights can be varied, and we can determine whether there is some combination of weights that gives a better system and a better solution for a set of input parameters.

The scheduling algorithms consist of the Northwestern system and the modified Northwestern system. The Northwestern system assigns the closest bus, and that is the only thing the scheduler considers. The modified Northwestern system evaluates the closest N buses, where N is less than the number of buses in the system, and then looks at the first j points in a tour of each of the buses, where j is less than the number of points on the tour of the bus. Part or all of the objective function can be used to determine which of that subset of buses would be the best bus to assign. That is the range of the assignment technique used in the General Motors simulation.

J. F. Nunamaker

What is the language for the DATS simulation model?

Larry Howson

The model is written in PL/1. This particular version of PL/1 is run under IPSS, an interim operating system for the IBM 360/67. PL/1 was not found to be as efficient as FORTRAN. The execution time for PL/1 was 5 times as long as the execution time for FORTRAN (with full optimization) for the minimum path algorithm. PL/1 is probably not the best language for our particular needs. However, we had little choice in the selection of a language because the systems we were working with only supported PL/1. This was not necessarily a bad situation because the graphics systems we used were only supported by PL/1. In some areas PL/1 was found to be very valuable. For example, PL/1 is very good in the area of address retrieval and data manipulation.

J. F. Nunamaker

Can you discuss components and structure of the objective function?

Larry Howson

The components considered in the objective function are the user-related attributes of the system. Weighting factors or values are assigned to the attributes, and the model is essentially a linear utility model. The model is formulated as a minimization problem. The model has the interpretation of minimizing the social and political costs of transportation of the customers in the system. The customers in the system include customers on board a bus and those called in or waiting for a bus. In that sense, there is a total correlation between the decision-making processes in the simulation model and the problem of minimizing the social cost of people involved in the transportation system.

Karl Guenther

In your objective function approach, did you find that specific coefficients used in the utility model vary over a extremely wide range? If not, then did a narrow range of coefficients make it possible for you to handle a large variety of different situations such as number of vehicles and demand rates?

Larry Howson

We have not done any sensitivity analysis on the parameters. The simulation program was completed in February 1970, and shortly thereafter the project was slowed down considerably. The inputs necessary for the economic analysis could not be obtained, and the simulation model was checked out with weights of unity. We have not had the opportunity to go back and to see what happens if we change the weights and to observe how sensitive the various weights are to changes in operating conditions. We would like to determine the full range of the parameters that can be used in the operation.

J. F. Nunamaker

What additional work is planned for the DATS simulation model?

Larry Howson

We intend to perform a sensitivity analysis and to investigate the efficiency of the program. We would like to reduce the cost of running the program.

J. F. Nunamaker

What are some of the reasons for the high cost?

Larry Howson

I am sure part of the problem is the language we are using and part is possibly due to inefficient programming.

Kenneth W. Heathington

Can the travel time constraint or any other service constraints be violated? Can you lose a customer in the system so that he ends up waiting 30 minutes for service when you guaranteed him a 15-minute pickup?

Larry Howson

It is impossible to miss a guarantee.

Nigel Wilson

How do you achieve that guarantee?

Larry Howson

If a guarantee is approached as customers are assigned a vehicle, each customer's projected arrival time is updated and the slack time is decreased and approaches zero. At this point the system goes critical for this one individual, and he is delivered. We have not looked at the possibility of something unexpected happening like a vehicle breakdown. It is assumed that the system operates perfectly at all times.

Nigel Wilson

Such an ambition, though laudable, is unattainable simply because the unexpected does happen. In addition, the GM simulation model assumes an unrealistically large supply of vehicles.

Kenneth W. Heathington

The problem of breakdown can be scheduled within a simulation; however, the only effect is to create a larger supply of vehicles with which to serve the demand. It is most important to recognize the fact that if one did not meet the service guarantees the system is in trouble. Some people on the DATS project felt that if we met the guarantees 95 percent of the time that would be sufficient. However, how do you keep that 5 percent that is not satisfied from shifting to another system? In addition, how do you keep from losing customers over a period of a couple of years? What does that do to the public's confidence in the system? It is not hard as far as scheduling algorithms go to meet all constraints with no one getting lost in the system and everyone being delivered on time. The question is, Can we afford to not meet the constraints that we have guaranteed?

Nigel Wilson

If on a very snowy day the demand is 3 times larger than the expected demand, how would you handle the situation?

Kenneth W. Heathington

We would propose that when a person calls in we would communicate with him and at that time give him a guaranteed time. This would be done only on days with adverse weather conditions. In both the Northwestern and General Motors studies, the time for pickup is fixed. The maximum waiting time is 15 minutes, and the delivery time is fixed. However, it is realized that in bad weather or as a result of breakdowns some problems will occur. It is not so detrimental to the overall concept of dial-a-bus to wait in adverse situations. However, on a bright sunny day in July it is upsetting to wait 45 minutes for a bus instead of the anticipated 5 minutes, and this becomes a very critical situation.

Nigel Wilson

There are 2 very different situations involved: whether a general system policy guarantee on waiting time and travel time may be violated (a) when adverse weather conditions are experienced or (b) when unexpected demand is experienced. As I understand the proposed system, at the time of the service request, the customer is notified of the service expectation and is given the option of saying, "I don't want to be picked up." This is a feasible way of operating, and it is probably the only feasible way of operating. The other approach is to have guarantees that one hopes the system never violates, and this is not a feasible approach.

Eugene T. Canty

We want to distinguish between the malfunctions in the system itself and the routing and scheduling. It is important to establish whether the system can pick up and deliver people on time if there are no acts of nature. The objective is to determine whether these assurances can be met with a computerized system and to determine whether the computerized system can do the job and a manual system cannot. The other aspect of the question is concerned with the extent to which failures and abnormal events can be accommodated, i. e., vehicle failure, traffic jams, and inclement weather.

Nigel Wilson

There is no great mystery about satisfying service guarantees. It is simply a question of how many vehicles are available to cover a specified demand in a given number of square miles. There is no way of satisfying guarantees if the system was designed with the wrong number of vehicles for a given number of demands.

Eugene T. Canty

The argument should be focused on whether the system is 100 percent responsive to the customer demands with everything else being fixed. In addition, the system can then be evaluated to determine if it is adaptive to the exogenous forces such as weather and bus breakdown. In the next set of simulations, it would be desirable to input these external forces as well as customer demands and to observe how well the system can adapt. Then it will be possible to calculate the frequency of occurrence in which the system did not meet its guarantees.

Joseph H. Stafford

None of us has done a crucial piece of work on this problem from the behavioral or customers' point of view. What is the trade-off? What are people willing to pay for increased reliability? Let me suggest an approach to the problem, if we had the time and energy to do it this way. What customers will do if the system is unreliable is to simply allow more time for the trip. If they are 95 percent confident of meeting their appointments at the doctors' offices, then they will begin to make the appropriate allowances for the trip. The customer will have enough experience with the system if he is using it at all to know or to have a "feel" for the trip time probability distribution. Then we can start asking the following question: What is a customer willing to pay to reduce the amount of safety time he has to allocate to a trip? To date, we have not done a complete analysis on that problem. It is necessary that we start thinking of the problem in terms of a confidence interval for travel time, i. e., the amount of time a customer must allow to be 90 to 95 percent sure of meeting appointments. This also suggests that travel time reliability may be rather different for a trip originating at home and for a trip ending at home. The problem of arriving home 15 minutes late is quite different from being 15 minutes late for work or 15 minutes late for an appointment.

Kenneth W. Heathington

One point we have not investigated concerns determining the exact size of the system required to hold to system guarantees. I am not convinced that it requires a large number of extra buses; maybe it does not require any, maybe one. I do not know of anyone that has investigated this problem.

M. I. T. SIMULATION MODEL

Nigel Wilson

At M. I. T., we have been involved in a project similar to the GM project. We have developed a series of algorithms for operation in the many-to-many problem or in the many-to-one problem (7, 8). These algorithms have been tested in a simulation model to determine how suitable they are for the operation of a DATS and to provide data for economic analysis of the system on which cost figures have been based. The simulation model was operational in March 1969, and since that time considerable testing has been done. Formulas have been developed for relating the number of vehicles to an area, demand level, and level of service (8). This formula is a result of a great deal of parametric work that was done to make it possible for comparisons to be made of the M. I. T. and General Motors results. It was realized about 15 months ago that, in order to get a "good" intuitive feel for how well the scheduling algorithm is doing compared with human logic, we really needed a graphic capability. As a result, a display system was implemented on the ARDS terminal. The ARDS is a much less expensive storage tube version of the IBM 2250 display scope. The display scope really increased our confidence in two respects: (a) that the algorithm had been programmed correctly and (b) that the logic of the algorithm was also correct. Human intuition did not indicate that we had a poor set of computer decision rules. The scope was found to be a very important aid in explaining to someone that was not familiar with the concept of dial-a-bus what it is all about. It is possible to demonstrate on the scope the type of service a customer might expect to receive. The ability to interrupt the simulation and to input new requests for service to the screen through a light pen was found to be very useful. With this type of device it is possible to specify an origin and destination on the screen and to let the computer algorithm specify alternative assignment of a customer to several vehicles. The customer is then given a first, second, and third choice, and then the customer (or analyst) can decide which choice is best and compare it with the computer's choice of which assignment is best. This was found to be a very useful approach to evaluating the scheduling algorithm.

In terms of the other questions raised in the introduction, we do specify constraints and an objective function that can be used in the way Larry Howson and Kenneth Heathington mentioned. One constraint is concerned with whether a request is acceptable or unacceptable. For instance, an option within the program permits the analyst to specify those constraints as firm constraints or soft constraints. Therefore, if there is no feasible way to service a new request and satisfy the service guarantee, the analyst can reject service for that customer. Whenever this situation occurs, the customer is notified of the best available service. This is clearly where the objective function comes in. The objective function is to maximize the satisfaction of all users or the social welfare of the system. The objective function considers everyone who is currently on the system, whether they have been picked up or whether they are still waiting. It also attempts to take into consideration potential customers who are not yet in the system and who will undoubtedly request service; i. e., it is necessary to keep a measure of slack to preserve service for those potential customers who will enter the system in a future time period.

The simulation was programmed in FORTRAN and is as a result operational on many computers. Most of the work has been done on the IBM 360/67, which has worked out very well with respect to the graphic capability. Some work has also been done on the Sigma 7.

J. F. Nunamaker

Can you comment on the details of the scheduling algorithms and on the practicality of the algorithms for an operational system?

Nigel Wilson

A result of the simulation work is the development of a primitive operational system. This primitive system is operational on the IBM 360/67, and in this version of the model the simulation environment has been removed from the software and replaced with external stimuli, i.e., an analyst seated at a console typing in a request from a person in terms of an origin (street address) and a destination. These street addresses are translated into coordinates that are put into the algorithm that exists in the same form as it exists in the simulation model. The instructions are printed on another console that represents the printout for a vehicle in an operational system. From that point of view, we have come a very long way toward having an operational system. Considerable testing has been done on the operational version, and we have reached the point where we have stopped finding elementary "bugs." There are no doubt a lot of other bugs in the system, and if it were put into a field environment we would find out the hard way. However, we would like to test the system in a field environment.

J. F. Nunamaker

Can you represent an existing street network, or do you use a rectilinear grid to approximate the street network?

Nigel Wilson

We can represent an actual street network. The operational system was programmed for Cambridge, Massachusetts. All the streets in a 4 square mile area are represented by a hash code and then translated into a grid map that is fed into the algorithm.

J. F. Nunamaker

Have you considered the selection of a computer for the implementation of an operational dial-a-bus system?

Nigel Wilson

That problem has been given considerable thought, and I would say we virtually have an operational system already, or we are very close to it. Naturally, there are some reservations about the system because it would be clearly a very experimental system. It would probably be prone to failures because it was designed from the simulation model, and as a result we patched various parts of the simulation model together to build the operational system. At present, it does not incorporate good systems design techniques (file design techniques) to make it a "good" operational system. Another drawback is that the system is tied to a fairly large computer because of the size of the program and the fact that data were stored redundantly to make it possible for someone not involved in the design of the program to modify it easily. This work was done in an academic environment, and the person who wrote the program for his master's thesis graduated. Since that time, everyone has been running it and modifying it, and it had to be very easy to modify and change. One of the initial design criteria was that the program must be written so that it could be easily modified.

J. F. Nunamaker

What are the minimum core requirements for the system as implemented on the IBM 360/67?

Nigel Wilson

The minimum core requirements exclusive of graphics is approximately 150,000 bytes. The model fits into one partition on the IBM 360/67 without any difficulty. The core requirement including graphics is approximately 200,000 bytes. We have considered using a mini-computer for implementing an operational system. This would necessitate developing the software from scratch, although basically the same scheduling algorithms could be used. The aim is to generate a system for a small computer very quickly. It is necessary to program it for speed of execution and to minimize the storage requirements. Whether we take this approach or not is undecided at this time.

J. F. Nunamaker

Do you have any particular mini-computer system that you are considering?

Nigel Wilson

There are several mini-computers on the market; possibly something like a VARIAN 620 is one that has been considered. There are many machines available; some can probably do the job, and others cannot do the job.

J. F. Nunamaker

The problem is that at the present time we know very little about the capability of these small computers.

Daniel Roos

That is true. One thing is certain: The system will have to be programmed in an assembly language rather than in a higher level language such as FORTRAN. This is necessary from the point of view of being able to fit the system into the memory available. Clearly that would have implications in terms of transferability, and one could imagine a proliferation of many dial-a-bus systems on small computers. The other approach would be to go to a medium- or large-scale computer where one is not so constrained as going in the FORTRAN route.

J. F. Nunamaker

Very few cities could afford to go the route of the medium- or large-sized computer for scheduling of a transportation system. Do you see a problem in having to customize the mini-computer dial-a-bus system for each city that would use the system?

Daniel Roos

There is a certain amount of customizing one has to do whether it be a small, medium, or large computer.

J. F. Nunamaker

But it is harder to make the changes in assembly language.

Daniel Ross

I agree that as one gets smaller and smaller machines one is faced with more and more customizing. However, for any system, the minimum amount of customizing involves modifying the street networks for each city.

Kenneth W. Heathington

Did you make an economic evaluation of using a large time-sharing system as opposed to using a dedicated smaller machine? Can you obtain the priority required for a DATS on a time-sharing system?

Daniel Roos

The one thing that bothered us about time-sharing was not so much the question of priority, because at least in the tests we ran the model needed a relatively small amount of the total computer, but the question of competition from the other users in terms of reliability. The fact is that when you are running with 23 other users and 1 of those 23 happens to "bomb" out the system you are down. That was our main reservation about time-sharing, and I do not think we ran into many problems actually getting into the system. At times during the day, when the load was very heavy, we did encounter some problems. In the environment we run under at M. I. T., one can set the number of users and set the priority, and problems can usually be worked out if one is willing to negotiate.

Kenneth W. Heathington

I realize that you can pay a premium and get top priority, but can you afford to do that 24 hours a day in order to meet all your schedules? However, the time-sharing costs seem to be decreasing all the time, and we now can have access to large capabilities at relatively low cost.

Nigel Wilson

To elaborate on a point made by Daniel Roos, I would say the amount of reprogramming that must be done for each new city will probably be quite small. The algorithms developed by M.I.T. and General Motors are very much independent of the demand distribution for an area. The grid network representation is also independent. Most of the software would not have to be reworked. However, the street address coordinate translation schemes would have to be reworked.

Kenneth W. Heathington

If, for example, people in Dayton, Ohio, want to conduct a feasibility study to estimate the number of vehicles required for a given demand, can they take the M.I.T. program and run it on a comparable machine in Dayton without needing the help of the people from M.I.T.?

Daniel Roos

We developed the simulation model at M.I.T. primarily for our own use. That was the extent of our initial grant application. As further work, we extended the system so that we hope we can use the system for that very purpose. We hope that a community could use the system to test out the feasibility of dial-a-bus analysis. How effective the model is for that purpose we will just have to wait and see.

Kenneth W. Heathington

Is M.I.T. attempting to market the system?

Daniel Roos

No, it is the property of the Urban Mass Transportation Administration.

Larry Howson

As far as General Motors is concerned, a community could try to use the system for a feasibility study. However, I have always experienced difficulty getting a program running that was developed elsewhere. This seems to be the case even with good documentation. Even if the installations are similar, there is usually some strange thing that is a bit different and that causes problems. The General Motors DATS simulation leans heavily on specialized software for the IBM 360/67. This implies that anyone wishing to use our model would have to do a lot of rewriting on the graphics part of the system.

Kenneth W. Heathington

Would you perform a feasibility analysis for a community if it provided the description of the existing street system.

J. F. Nunamaker

The consensus from representatives from both M.I.T. and GM is that they are quite willing to perform a feasibility study for a community or city. Did either M.I.T. or General Motors investigate the vehicle scheduling system (VSP) package available through IBM?

Nigel Wilson

We have looked at VSP and found it to have somewhat limited capabilities. The vehicle scheduling area has tremendous potential for many types of transportation systems.

Larry Howson

We have not really had an opportunity to evaluate the VSP package. The original intent was to use VSP as an adjunct to the dynamic scheduler. Vehicle scheduling looks like a great area for development work.

Nigel Wilson

In addition to the dynamic scheduling problem, we are doing some work on the prescheduling problem for work trips that are repeated day after day. There is a great potential for a more sophisticated vehicle scheduler that can be run off line and not in real time. This type of scheduler could be run once a week or once a month, whenever it is necessary. This approach gives tentative routes for vehicles that can be updated in real time with a less sophisticated algorithm to modify the schedule for day-to-day variations in demand. The basic system then might consist of a set of standing requests for service that could be overridden by a phone call on any particular day. Combining these 2 approaches to vehicle scheduling appears to be a good area for algorithm development.

Eugene T. Canty

I would like to comment on something Nigel Wilson said earlier because it has important implications. He was talking about the fact that when an operational computer program is developed in one community there is a minimal amount of change necessary to adapt it to another community. The key change centers around the characteristics of the local community, i.e., the street network. It is important that we have new transportation systems that have a steep learning curve. If the system is proved successful and operational for one community, then it is relatively easy to implement it into a second, third, and fourth community. It is important to have that learning curve both with regard to the hardware and the software. A dial-a-bus system is one that should have a very steep learning curve. It is different from a moving sidewalk where there is a very high level of architectural and engineering content and a very shallow learning curve. Almost as much architectural and engineering time is expended on the hundredth system as on the first system; however, the learning curve for the system is shallow. Dial-a-bus is different. With regard to hardware, software, communication system, and vehicle, it should be rather easy to standardize on those components. This is important to the community because if the federal government sponsors the development of components for the first community then the other communities around the country profit from the experience of the first community. This should also hold with respect to software. Intuition indicates to me that a computer routed and scheduled vehicle program will be more adaptable to a second city. The automated routing system has perhaps a steeper learning curve than a manually routed system. For that reason, I think it would have been better to go the manual scheduling route at first rather than start with computerized routing.

J. F. Nunamaker

It is obvious that we have a good start on the simulation of DATS. However, much work remains to be done with respect to sensitivity analysis, development of more powerful and sophisticated schedulers, and implementation problems in the community.

REFERENCES

1. CODASYL Systems Committee. A Survey of Generalized Data Base Management Systems. Assn. for Computing Machinery, May 1969.
2. Data Management Software. Diebold Research Program, Document T24, Sept. 1969.

3. Dixon, P. J., and Sable, J. DM-1: A Generalized Data Management System. Proc. Joint Computer Conf., May 1967.
4. Bauer, H. J. A Case Study of a Demand-Responsive Transportation System. General Motors Corporation, Res. Pub. GMR-1034, Sept. 10, 1970.
5. Howson, L. L., and Heathington, K. W. Algorithms for Routing and Scheduling in Demand-Responsive Transportation Systems. Highway Research Record 318, 1970, pp. 40-49.
6. Vitt, J. E., Bauer, H. J., Canty, E. T., Golob, T. F., and Heathington, K. W. Determining the Importance fo User-Related Attributes for a Demand-Responsive System. Highway Research Record 318, 1970, pp. 50-65.
7. Roos, D. Dial-A-Bus System Feasibility. Paper presented at Conf. on Demand-Actuated Transportation Systems and published in this Special Report.
8. Wilson, N. H. M., Sussman, J. M., Higonnet, B. T., and Goodman, L. A. Simulation of a Computer-Aided Routing System (CARS). Highway Research Record 318, 1970, pp. 66-76.