REDUCING THE TRAVEL REQUIREMENTS OF THE AMERICAN CITY: AN INVESTIGATION OF ALTERNATIVE URBAN SPATIAL STRUCTURES

Jerry B. Schneider and Joseph R. Beck,
Departments of Urban Planning and Civil Engineering, University of Washington

Little attention has been given to investigating the potential for gradually restructuring the city to reduce its built-in requirements for transportation. This paper investigates this potential by analyzing the transportation requirements of some alternative urban spatial structures. A computer-directed search procedure is developed and tested by using two simple urban structures. These experiments form the basis for an investigation of a larger urban structure based on the 1970 urban pattern of western King County, Washington. The results of these investigations indicate that, by moving about a third to a half of the people and about a third of the jobs to other locations, substantial reductions (50 percent or more) in travel requirements could be obtained while also improving overall access levels. Although structural changes of this magnitude are not feasible in the near term, many cities may be growing by this much in the next 20 to 30 years. The potential of nontransportation solutions to transportation problems appears to be a significant but neglected area for policy-oriented research investigations.

•THE need to reduce the transportation requirements of the American city is becoming more urgent. Part of the impetus behind this need comes from a growing realization that we must find many ways to conserve energy in the future. Other problems such as poor air quality and the citizen revolt against urban freeways have also had an important role in stimulating renewed interest in searching for ways to reduce the need for transportation in cities. Proposed solutions to these problems often take the form of plans to build new transit facilities such as rail mass transit or automated personalized rapid transit. To date, little attention has been given to investigating the potential of gradually restructuring the city so as to reduce its built-in requirements for transportation. Few people have asked, "Can significant reductions in the transportation requirements of a city be achieved by changing its urban spatial structure in certain ways?" It is the purpose of this paper to investigate this question by analyzing the transportation requirements of some alternative urban spatial structures. The approach used is experimental in nature and involves the use of a computer-directed search for urban spatial structures that have minimal transportation requirements. The effort here is to deal with what we perceive to be the basic causes of the urban transportation requirements instead of examining only the symptoms of these problems.

Another way of conceptualizing the approach is to ask, "For a given transportation system, how much could travel be reduced by shifting or rearranging the location of people and jobs?" If, for example, we could show that by shifting the location of 10 percent of the people and jobs in a city, one could expect to see total journey-to-work travel in that city decline by 15 percent, it would seem logical to closely examine ways

---

Publication of this paper sponsored by Committee on Passenger and Freight Transportation Characteristics.

in which such a restructuring of the city might actually be accomplished. Few of the recent and very expensive proposals to construct transit facilities in major cities can demonstrate that they will reduce the travel requirements of the population they will serve. Almost all of them will allow and encourage more people to travel more. While this may be viewed as being "good" in terms of helping to satisfy our virtually insatiable desire for more and more individual mobility, it is inconsistent with our important needs to conserve energy, improve air quality, and minimize disruption to existing parts of the urban fabric. At some point in our history, we will have to face up to the question, "How much individual mobility is enough?" There are important trade-offs between mobility and the various aspects of environmental quality that are only dimly perceived by most people at this time. Our view is that, in general, more mobility means less environmental quality, unless very large sums of money are invested for environmental protection in the transportation arena. If this is true and if environmental quality is getting to be as highly valued as mobility, then it makes sense to think much harder about ways to preserve and enhance the environmental quality of our cities. If we can refrain from building new and expensive transportation facilities while preserving current mobility levels in the future, we can expect to free the resources needed to make our cities much more livable than they are now. Our quest, then, is to search for nontransportation solutions to the urban transportation problem where such solutions can be expected to (a) result in equal or better individual mobility and (b) produce a better level of environmental quality in the city. A nontransportation solution is one that involves no new transportation facility or service but instead involves a rearrangement of a particular urban spatial structure such that some part of the present transportation requirement of the city is reduced significantly.

Some additional background and perspective for this rather radical approach to the urban transportation problem are presented later in this paper. An automated search procedure designed to "discover" high-performance (i.e., low-travel-demand) urban spatial structures is presented and is then applied to two simple networks to test its utility. Then the search procedure is applied to an abstraction of the urban form of the western part of King County, Washington, an area which included more than 1 million people in 1970. Finally, some conclusions from this study and some suggestions for further research are given.

A PERSPECTIVE ON THE LAND USE-TRANSPORTATION RELATIONSHIP

The relationship between land use and transportation can be conceptualized as a circular chain as shown in Figure 1. Within this circular chain, land development determines the pattern of trips in the urban area simply because the land use pattern represents the locations of all trip origins and destinations. These trips, when aggregated into an urban travel pattern, define the transportation requirements for an area. As problems arise (or are forecast to arise), additional transportation facilities are constructed, providing increased accessibility to certain parts of the area. This new accessibility causes changes in the value of the parcels of land served by the new facilities. The land is then put to more intense uses, and the circular chain of relationships continues until no further development of the land is possible.

The nature of this circular chain raises some questions about the decision-making process that has traditionally been used to plan transportation facilities. Specifically, if new facilities are built to serve only those needs determined by present trip patterns and projections of present patterns, then it is clear that the land use implications of such decisions are not being given sufficient consideration. All too often, this emphasis on one part of the circular relationship results in a serious lack of attention to the remaining interacting elements.

Even if the circular nature of the relationship between transportation and land use in urban areas becomes widely recognized, another question arises concerning the extent to which this knowledge will be used to guide the future development of the area. Some type of answer to this question should be formulated by the people in each metropolitan area. While this has been rarely done in the past, at least one metropolitan

Figure 1. Circular relationship between land use and transportation activities.

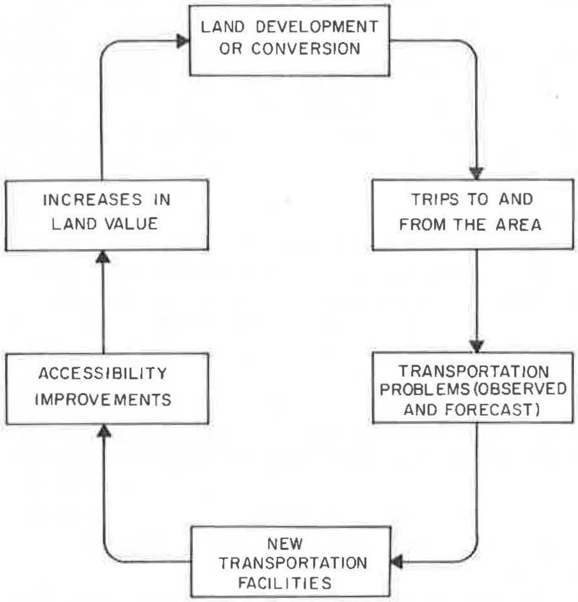

area (Minneapolis-St. Paul) has recently made the decision that capital investments in transportation facilities will be used as a tool to guide development as well as to serve it (2).

At the national level, the National Academy of Engineering has recently completed a study for the U.S. Department of Transportation that recommends areas for future research and development in the field of urban transportation (6). One area recommended for study was the effect of city design on urban transportation:

> The increasing focus on the quality of urban life clearly calls for a better understanding of the interactions and relationships between urban transportation systems and the functions of metropolitan areas. This, in turn, requires an enhanced program of analysis and real world experimentation. . . .Satisfactory urban transportation solutions depend to an important degree on the origin and arrangement of the city, on the creation of an attractive environment, and on a desirable growth policy. More satisfying urban communities depend to an important degree on the design of transportation systems, including the effective use of transportation infrastructure as an aid to good urban design and environment.

Because there is a possibility that more and more metropolitan areas will seek to use transportation investments as tools to guide their development in the future, some further investigation into the land use-transportation relationship is urgently needed. These investigations should make use of methods that adequately represent the circular nature of that relationship, and they should be capable of answering at least three questions:

1. What are the transportation criteria with which we can evaluate alternative land use patterns?
2. How can we use these criteria to discover those land use patterns that are most desirable?
3. After desirable patterns of land use have been identified, how can we determine which transportation policies can best assist the achievement of such patterns?

The first question is not intended to suggest that only transportation criteria should be used to determine favorable urban development patterns. Any choice of this type must depend on a combination of criteria and goals covering all the varied activities in the urban area. However, it will certainly be helpful to identify those transportation criteria that are important and to discover those land use configurations that exhibit the best performance based on the chosen criteria.

Further, other factors that influence land use in a manner similar to transportation may also be used to guide urban development toward desirable patterns, once those patterns are identified. Therefore any procedure that can answer the first question posed above will have application to capital-investment programs for utility planning (such as power, water, and sewage systems) and open-space planning as well as transportation-system planning.

Providing the answers to the second question will require the development of automated search procedures designed to discover high-performance land use patterns that best satisfy the criteria used. The next section of this paper describes our initial effort to develop such a procedure.

Providing the answers to the third question will require development of a model designed to simulate all of the interrelated elements shown in Figure 1. This model must include a feedback structure that will allow the simulation of the effects of alternative transportation and land use policies, so that one can identify policies that will move the metropolitan area toward a spatial structure that has been identified as most desirable. This task is well beyond the scope of this paper.

Basic to an understanding of any approach to the problem of identifying an optimal urban form is the concept of a combinatorial space. This term simply represents the set of all the possible urban configurations for any particular geographic setting. For example, if we were dealing with a situation where there were 100 different feasible and logical ways to distribute people in an urban area, 50 different ways to distribute jobs, and 5 different ways to design the transportation network, we would have to deal with 25,000 alternative urban situations ($100 \times 50 \times 5 = 25,000$). In most practical situations, the number of urban structure combinations that arise (even when the urban area is characterized in very abstract terms) is astronomical in size. One cannot hope to be able to examine all of these alternatives in any practical situation of this type. One must either eliminate most of the alternatives by (a) making a series of judgments, or (b) use a computer to search for a satisfactory solution within this set of all possible combinations (i.e., the combinatorial space), or (c) use some type of analytical procedure such as combinatorial programming to find a solution with the desired performance level (8). We wish to do something better than relying totally on intuitive judgment but have found our problem to be unsuited to the more elegant combinatorial programming approach. Thus, our attention has been directed to the development of ways to use the computer to search the combinatorial space formed by alternative urban configurations.

Five previous studies of this topic have influenced our approach to this problem. Hemmens (5) investigated the transportation requirements of a simplified urban form consisting of 37 zones arranged in a grid pattern. Thirty alternative urban structures were studied where residential, commercial, and other employment land uses were located in various locational patterns. Hemmens used a linear programming algorithm to assign home-to-work and home-to-shopping trips to shortest path routes. The programming model does not attempt to simulate the behavior of tripmakers, but rather makes assignments such that for each configuration the total travel, in man-minutes, is a minimum. The only criterion used to judge the desirability of one spatial structure over another was the aggregate time required to complete all the work and shopping trips in the city.

The 30 selected configurations evaluated by Hemmens included various location patterns for the two named land uses and also included variations in the transportation system connecting the zones. However, several restrictions were placed on the configurations that severely limited the number of possible cases. The 30 cases studied represent only a small fraction of the total possible cases, a number that is astronomical in size. The land use-transportation combinations that were chosen for examination were selected on an intuitive basis and may or may not represent the most significant subset of combinations. Hemmens did not find much variation in performance (i.e., total travel) among the alternatives he examined, and this may be the result of a scope that was too restrictive.

Harris (4) mentions the application of combinatorial methods to the regional planning process and discusses two reasons why such an application would be difficult. One application is the evaluation of alternative transportation system plans. Harris suggests that it might be possible to cast the set of all possible alternative transportation systems as a combinatorial space and evaluate all of them by means of high-speed digital computers, retaining only a few of the best combinations for output. Another suggested application is that of finding the best sequence of development for a particular urban region. This sequencing problem could be defined as a combinatorial space representing all possible development sequences, and then the set of sequences could be evaluated automatically, using a set of cost or other criteria representing a time dimension to the combinatorial space.

The disadvantages of this approach cited by Harris are

1. The space containing all possible combinations for any realistic problem is so large that the evaluation of all combinations would be prohibitively expensive and time-consuming.
2. The combinatorial space could not be expected to be smooth or continuous, and any search procedure that attempts to find "best" combinations on the basis of gradients would likely be ineffective due to the problems of local optimums.

This study is designed primarily to examine the second of these assertions within an experimental framework.

A recent study of the Detroit metropolitan area by Doxiadis (1) tends to confirm at least the first difficulty mentioned by Harris. The Detroit study used a combinatorial approach and quickly identified about 49,000,000 alternative future development patterns for the area. Most of these alternatives were eliminated by using intuitive judgment, and only about 40 were examined in detail by the study staff. This study required 5 years and cost approximately 3 million dollars, but in fact only a very small proportion of all of the possible alternatives were closely examined. The selection of the set for detailed study was based largely on human judgment and was not the result of any automated search or evaluation of the full set of alternatives.

A characteristic common to both the Hemmens and Doxiadis studies is that, while the combinatorial concept was used to establish a framework for the study, the actual set of combinations selected for detailed study was defined using only a series of human judgments. Each such judgment has the effect of eliminating thousands or tens of thousands of alternatives. In such situations, one is never quite sure what good alternatives might be lost in this type of elimination process.

In an effort to overcome the difficulties presented by a very large combinatorial space while at the same time making use of human judgment, Goldman (3) developed an interactive graphic computer program in which the human analyst chooses an urban configuration and submits the configuration to the computer for analysis. The results of the automated analysis are then presented to the analyst in the form of a series of graphic displays. The analyst then makes judgments based on these displays, modifies the configuration, and resubmits the modified configuration to the computer for analysis. Because the cycle of modification and evaluation can be performed very quickly, the analyst can construct and evaluate several alternatives quickly and can usually derive a series of successive configurations, each of which has a higher level of performance.

However, it has been pointed out by Rapp (7) that, in any situation in which the human analyst must deal with a very large combinatorial space, it may be very difficult to decide how to modify the present configuration in order to move it toward the established performance objectives. That is, the analyst will, more often than not, be overwhelmed by the large number of possible choices. In such cases, he will either give up or make a series of guesses which, more often than not, will not lead to the discovery of a satisfactory solution.

A possible solution to this problem is the automation of the process of modifying configurations in addition to the analysis of each chosen configuration. With the combinational framework established (i.e., when each possible configuration can be defined by a specific combination of the variables used), the choice of a modified configuration (i.e., the next configuration) for analysis requires only that the choice be based on some numerical measure of performance. In this way, the process of modification and evaluation can be completely automated. Human judgment is needed only to determine (a) the original specification of objectives, in numerical terms, as derived from the desired level of performance of the configuration and (b) the starting configuration or position where the search is to begin. A completely automated search process makes maximum use of the speed available in digital computers to search large combinatorial spaces. The results of this search will be the most desirable configurations that can then be subjected to further examination. In contrast to the approaches of both Hemmens and Doxiadis, an automated search process does not require one to intuitively select various configurations for further detailed study. Rather, it finds or discovers a set of configurations that satisfy the given objectives without any further input from the human analyst.

This paper examines the concept of automated search algorithms that operate within combinatorial spaces as useful tools in examining the transportation requirements of alternative urban structures. Our purpose is to identify those urban structures that are most satisfactory according to the transportation criteria selected.

Procedures will be presented that allow various possible urban structures to be represented as points in a combinatorial space. A method is adopted for evaluating each candidate configuration according to a set of transportation criteria chosen by the user. A prototype search algorithm will be presented and evaluated by application to two simple and small combinatorial spaces representing hypothetical urban situations. The performance of the algorithm in these test cases is discussed and provides a basis for the investigation of a large problem representing an actual urban configuration.

The criteria for desirable urban structures used in this paper are based exclusively on the internal transportation requirements of each structure, and therefore no claim is made that the structures identified are desirable in any other sense.

DEVELOPMENT OF THE AUTOMATED SEARCH ALGORITHM

This section presents a concise statement, in mathematical terms, of the problem of finding desirable urban structures (cast in combinatorial form) and gives a description of the search algorithm used for this purpose.

Problem Statement

The problem to be examined here is that of arranging a specified number of people and jobs on a fixed transportation network (a set of nodes, links, and travel times) that represents the transportation facilities available in an urban area. Measures of a societal utility of alternative arrangements of people and jobs will be based on the internal travel requirements of each arrangement. The search algorithm then is to find an arrangement that satisfies the objectives specified.

Each arrangement or configuration of people and jobs on the nodes of the network is defined by a set of configuration variables

$$X = [x_1 \ldots x_n] \tag{1}$$

These variables are the number of people (tripmakers) and the number of jobs at each node in the network. Since we consider only journey-to-work travel, these variables correspond to the spatial distribution of residences and work places on the network.

Constraints are specified that provide upper and lower limits for the number of people and jobs at each location (node) on the network. Another constraint specifies that the total number of trips made is a constant, since the objective is to find a more desirable arrangement for a constant number of trips. Each possible arrangement of people and jobs constitutes one configuration, and any configuration that does not violate any constraint is called a feasible configuration.

The second set of variables is called the impact set:

$$T = [t_1 \ldots t_m] \tag{2}$$

For our problem, the impact variables are the link flows and accessibilities generated by the configuration of people and jobs on the network.

A set of constants is specified that defines the network of locations (nodes) and their connecting links:

$$C = [c_1 \ldots c_p] \tag{3}$$

The constants describe the physical layout of the network and the travel times over each link of the network in the urban area.

The impact variables are related to the configuration variables and constants by the system equations:

$$\begin{aligned} t_1 &= g_1(x_1 \ldots x_n,\ c_1 \ldots c_p) \\ &\vdots \\ t_m &= g_m(x_1 \ldots x_n,\ c_1 \ldots c_p) \end{aligned} \tag{4}$$

A gravity model is used as a basis for these system equations in this study.

The relative desirability of each configuration is determined by a set of performance measures:

$$\begin{aligned} PM_1 &= h_1(t_1 \ldots t_m) \\ &\vdots \\ PM_r &= h_r(t_1 \ldots t_m) \end{aligned} \tag{5}$$

These measures are derived from the impact variables; that is, they are functions of the accessibilities and patterns of trips generated by any particular configuration of people and jobs on the network.

Performance Measures

Performance measures are one way of summarizing the characteristics of each configuration. These characteristics must be summarized and expressed numerically because they form the basis for the decision rules used by the search algorithm. The performance measures used in this study are discussed in the following.

Total Travel—Total travel is the sum of all travel, measured in person-minutes, required to complete the set of all work trips from all the origins to all the destinations in the network. It is not only a measure of the collective time required to satisfy all trip demands, but it is also an indicator of the magnitude of secondary effects associated with travel such as the consumption of energy and the level of exhaust emissions from vehicles:

$$PM_1 = \sum T_{ij} t_{ij} \tag{6}$$

where

PM_1 = total travel;
T_{ij} = number of trips between origin i and destination j; and
t_{ij} = time required to travel from origin i to destination j by the shortest path.

The search algorithm will seek those configurations that require the least total travel to satisfy their trip requirements (i.e., moving everyone from their home to their job).

Total Weighted Accessibility—Total weighted accessibility is a measure of aggregate nearness of each residential location to all employment locations in the urban area:

$$PM_2 = \sum_{i} P_i \sum_{j} \frac{A_j}{t_{ij}^{b}} \qquad (7)$$

where

PM_2 = total weighted accessibility;
P_i = trips produced at origin i (residential location);
A_j = trips attracted to destination j (job location);
t_{ij} = time required to travel from origin i to destination j by the shortest path; and
b = exponent reflecting the friction of space or average difficulty of overcoming spatial separation in an urban area.

The value of the travel time exponent used in this measure is 2.0, an average of those values commonly used in urban transportation studies. Total weighted accessibility is not a measure of travel but rather a measure of overall spatial relationship between home and work locations in the urban area. Configurations that have a higher total weighted accessibility are interpreted as having more social utility than those that have lesser weighted accessibilities.

Average Link Load—Average link load is an indication of the average level of use (i.e., average loading) of the transportation network. Because only internodal trips load the links, average link load relates only to internodal travel, as distinguished from the intranodal trips, which both originate and end within a single location (zone):

$$PM_3 = \frac{\sum_{k} LO_k}{NL} \qquad (8)$$

where

PM_3 = average link load;
LO_k = number of trips on link k; and
NL = total number of links in the network.

Maximum Link Load—Maximum link load is the largest of the loads on any link in the network:

$$PM_4 = \max_{k} \left[LO_k\right] \qquad (9)$$

where LO_k = number of trips on link k. Since all links are defined to be one-way links, the maximum link load is also associated with a direction. Maximum link load is a function of the concentration of internodal trips on a single link. This concentration may be considered undesirable under some conditions but may be desirable under some other conditions (e.g., the user may wish to find a configuration with a highly concentrated travel pattern in order to make best use of a high-capacity fixed-route transportation technology). However, since one of our objectives is to find urban configurations that do not require high-capacity transportation facilities, we will interpret high maximum link loads as being undesirable.

All of these measures are derived from an "all or nothing" assignment procedure. This means that all trips are assumed to use the shortest time path between each origin-destination pair of nodes. This procedure is a crude approximation to the behavior of actual tripmakers but is assumed to be sufficiently realistic for the purposes of this study.

Objective Function

In order to assess the overall utility of any particular configuration, some method is needed to combine all the performance measures into a single numerical score. However, because we wish to examine the performance of the search algorithm itself, we will use individual performance measures in separate searches and will not compute and use an overall score. This approach will permit evaluation of the search algorithm with respect to the individual performance measures and will also identify the configurations that are optimal for each performance measure. This simpler approach is viewed as a necessary step in the development of a search algorithm that can deal with multiple objectives simultaneously.

System Equations: Gravity Model

This paper used the gravity model approach described by Goldman (3) rather than the linear programming method of Hemmens (5). The gravity model is used because it better represents the behavior of actual tripmakers, as opposed to the linear programming method, which does not distribute trips in a realistic manner.

The gravity model equation is shown by Eq. 10. According to the gravity model concept, travel generated by persons in one location and jobs in another location is directly proportional to the number of persons and the number of jobs in both locations and inversely proportional to some power of the time or distance between the two locations:

$$T_{ij} = \frac{\dfrac{P_i A_j}{t_{ij}^b}}{\displaystyle\sum_j \dfrac{A_j}{t_{ij}^b}} \qquad (10)$$

where

T_{ij} = number of trips from origin i to destination j;
P_i = trips produced at origin i (residential location);
A_j = trips attracted to destination j (job location);
t_{ij} = time distance from node i to node j by the shortest path; and
b = exponent expressing the friction of travel.

The search algorithm does not necessarily require that the gravity model be used as the basis for trip distribution. Any other trip distribution technique could be substituted for it without changing the search procedure.

Specification of the Search Algorithm

The purpose of the search algorithm is to generate successively more desirable urban configurations of people and jobs as determined by the performance measures described earlier. The search algorithm is based on the concept that any distribution of people and jobs over the network of nodes can be represented as a combination of variables. Any such combination may be thought of as a point in a combinatorial space. The function of the search algorithm is to locate successively better combinations by moving about within the combinatorial space.

A flow diagram of our search algorithm is shown in Figure 2. Beginning with a starting configuration, the algorithm generates a new configuration by moving a specified number of people from the first node to the second. The travel requirements of the new configuration are calculated and compared to the values of the starting configuration. If no improvement has been found, the second configuration is eliminated and a third configuration is generated from the starting configuration by moving a block of people from the first node to the third. If this third configuration produces a score higher than the original score, it replaces the starting configuration in the memory of the computer. The process is continued until all possible node pair trip production shifts have been examined. The search is then repeated for job location shifts in an identical manner.

This is an extremely simple method of searching the combinatorial space, but it has the advantage of being very fast, and the high speed of the algorithm makes it possible to examine a vast number of configurations in a very short time.

The search algorithm can be made even faster by placing upper and lower bounds on the numbers of people and jobs located at each node. These upper and lower limits can be set by the user to restrict the search to a set of configurations considered reasonable. After each configuration is generated, it can be checked to determine if any constraints have been violated. If they have, the configuration is eliminated without being analyzed. These bounds can significantly reduce the size of the space that has to be searched.

A further increase in speed can be obtained by repeating those shifts that have generated an improved configuration. A successful shift may be repeated between the same node pair until it no longer generates a better performing configuration. In the next section we will test the performance of this algorithm in the context of two experiments.

TWO EXPERIMENTS DESIGNED TO TEST THE PERFORMANCE OF THE SEARCH ALGORITHM

The purpose here is to present a description of the two experiments that were carried out to examine the performance of the search algorithm. The results of the experiments are presented in graphic and tabular form and show both the performance of the algorithm and the characteristics of the urban patterns found by the algorithm.

Summary of Experimental Design

Two experiments were performed to evaluate the effectiveness of the search algorithm. These two experiments involved application of the search algorithm to a 3-node and a 5-node network. The combinatorial space associated with these examples was evaluated completely by examining all possible configurations before the search algorithm was applied. Thus it was possible to determine whether the algorithm was actually able to find the best configuration for each performance measure. Searches were conducted using each of the four performance measures as the objective, and several different starting configurations were used for each performance measure search as well.

Experiment I: A 3-Node Network

The first experiment consisted of an application of the search algorithm to a network of 3 nodes and 6 one-way links that formed a right triangle, as shown in Figure 3. Numbers beside the links indicate travel times along the links. These travel times are used as the measures of distance between the nodes and are represented by the symbol t_{ij} in the gravity model (Eq. 10). The distance decay exponent b in Eq. 10 has a value of 2.0 in all experiments. The intranodal time, or average time between people and jobs located at the same node, was set at 1.0.

The 3-node network is the smallest and simplest network that will yield useful information. The 3 pairs of one-way links form unequal legs of a triangle, and therefore the network is not symmetrical. The lack of symmetry means that even in this simple network one node is the most central node and one is the most remote node. In this case, node 2 is the most central and node 3 is the most remote.

The number of people and jobs to be located on this network was arbitrarily limited to a total of 300 people and 300 jobs. An upper limit of 200 people or jobs at any single node was used, and the minimum limit was set at 50. These limitations therefore allow any people-job combinations that sum to 300 people and 300 jobs, including concentrations of as many as 200 people and jobs at any node, or an even distribution of 100 people and jobs at each node.

Complete Enumeration of the 3-Node Network and Experimental Results

The process of evaluating all possible configurations within the limitations listed was carried out using a step size of 50 people or jobs. This means that each configuration differed from the previous configuration by the removal of 50 people or jobs from one node and the addition of 50 to some other node. Under these conditions, there are 100 possible configurations.

Figure 2. Flow diagram of search algorithms used to find improved urban spatial structures.

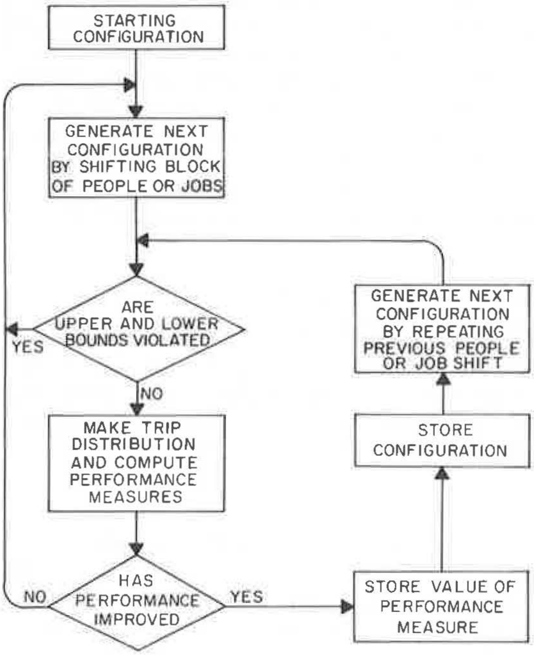

Figure 3. The 3-node network.

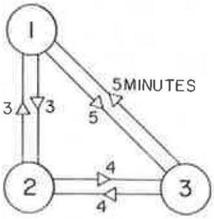

Figure 4. The best and worst people/job configurations for the 3-node network.

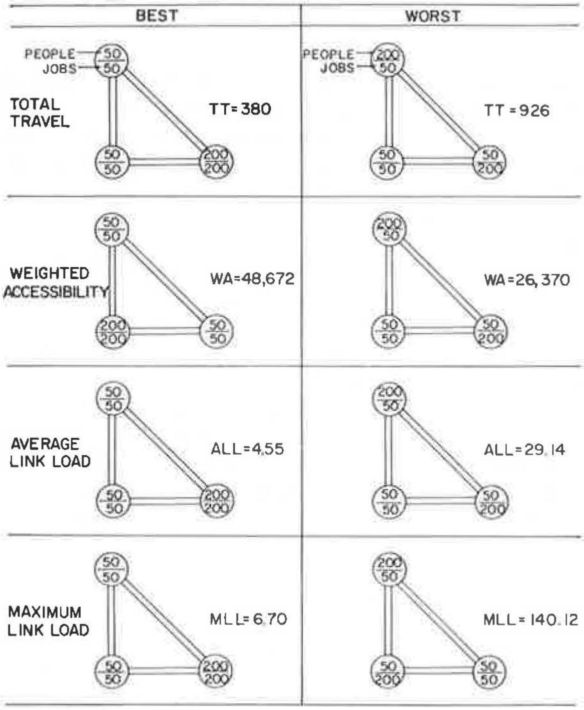

Each of the 100 configurations was generated and evaluated with the four performance measures. The best and worst configurations for each of the four performance measures are shown in Figure 4. As was expected, three of the four "best" configurations were identical, because a maximum concentration of people and jobs at the least central node will produce a minimum of total travel, a minimum average link load, and a minimal maximum link load. On the other hand, maximum weighted accessibility occurs when there is a maximum concentration of people and jobs at the most central node. In a similar vein, three of the worst configurations are identical in that a minimum number of people and a maximum number of jobs are located at the least central node, producing a maximum of total travel, a minimum weighted accessibility, and a maximum average link load. The largest maximum link load occurs when the two closest nodes are as unbalanced as possible (i.e., many people but few jobs).

With these results in hand, the next logical question was, "How often can the search algorithm find these best configurations?" Answers to this question are given in Table 1. Four different starting configurations were selected at random and the search algorithm was run 16 times, 4 times (from different starting conditions) for each of the 4 performance measures. As shown in Table 1, it was completely successful (i.e., it found the best configuration) only 2 of 16 times. However, when it did not find the best configuration, it did find one that was very nearly as good as the best, except in the case of average link load. Table 1 also shows how the average value of the 4 runs compares with the best value in each case. The differences are very small, with the exception of the average link load measure. These results were judged to be sufficiently encouraging to warrant a further round of testing with a 5-node network.

Experiment II: A 5-Node Network

The second series of tests of the search algorithm is similar to the first in that the algorithm was applied to a small network that had been previously completely enumerated. The difference is that the network configuration is slightly larger and more complex. Also, it is symmetric, as shown in Figure 5, and has a node that is clearly central, surrounded symmetrically by four others, which are equally least central. This experiment was designed to begin to approximate a symmetric urban configuration. As in the previous experiment, a pair of opposing one-way links connects each node pair and the intranodal distance, or travel time, was one unit.

A total of 400 people and jobs was distributed on this network, with a minimum of 50 people and jobs and a maximum of 200 people and jobs at any single node. Note that these limitations allow concentrations to be formed at any node but do not allow a uniform distribution over all the nodes.

Enumeration of the 5-Node Network and Experimental Results

An enumeration of all the possible configurations for the 5-node network was done with a step size of 50 for both people and jobs, producing a total of 1,225 configurations. In contrast to the 3-node network, enumeration of the larger 5-node network resulted in many sets of configurations that all had the same performance level. This is a direct result of the symmetry of the network. The best and worst of the 5-node network configurations are shown in Figure 6. Two of these best configurations are identical in that a maximum concentration at any one of the four least central nodes produces a minimum of total travel and a minimal average link load. Concentration of a maximum number of people and jobs at the most central node produces a maximum of weighted accessibility. The smallest possible maximum link load occurs when the people and jobs are dispersed as possible.

Three of the four worst configurations are identical. An unbalanced distribution of people and jobs (i.e., maximum people and minimum jobs) at two of the least central nodes produces a maximum of total travel, a minimum of weighted accessibility, and a maximum average link load. As before, when the two closest nodes are assigned a highly unbalanced people-job mix, the maximum link load occurs on the link that joins them.

How did the search algorithm perform on the 5-node network? Table 2 gives these

Table 1. Results of application of the search algorithm to the 3-node network from four different starting configurations.

Performance Measure	Success Ratio	Value of Best Performance	Average Value of Results of Four Searches	Average/Best
Total travel	2:4	380.00	380.5	1.00
Weighted accessibility	0:4	47,605.00	46,317.00	0.97
Average link load	0:4	6.70	10.45	1.56
Maximum link load	0:4	4.55	5.45	1.20

Figure 5. The 5-node network.

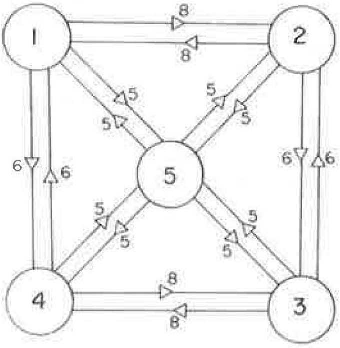

Figure 6. The best and worst configurations for the 5-node network.

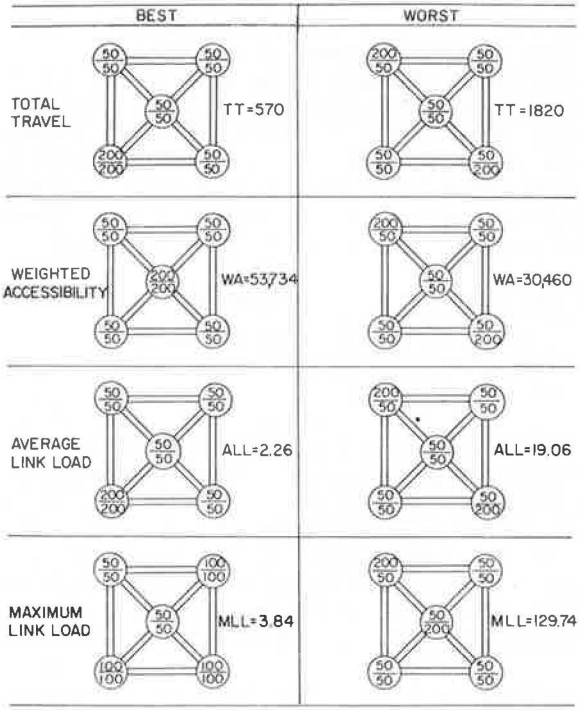

Table 2. Results of application of the search algorithm to the 5-node network from four different starting configurations.

Performance Measure	Success Ratio	Value of Best Performance	Average Value of Results of Four Searches	Average/Best
Total travel	4:4	569.70	569.70	1.00
Weighted accessibility	0:4	52,735.00	52,740.00	1.00
Average link load	0:4	5.44	5.66	1.04
Maximum link load	4:4	2.26	2.26	1.00

results. While the search algorithm found the best configuration only 8 of 16 times, it did find near-optimal configurations on all 16 tries. This result is shown in Table 2 by the very close correspondence between the average value of the 4 runs and the best value for that run in all 4 categories. These results were judged to be satisfactory enough to continue the test of the search algorithm using a 12-node network.

APPLICATION OF THE SEARCH ALGORITHM TO A SIMPLIFIED 12-NODE NETWORK REPRESENTING THE WESTERN PART OF KING COUNTY, WASHINGTON

The purpose of the third experiment is to use the search algorithm on a reasonably realistic network to investigate the effect on travel requirements that could be obtained by altering an existing urban spatial structure. To study this problem, a simplified network representation of the western part of King County was developed. This 12-node, 38-link network is shown in Figure 7. Several of these nodes represent the city of Seattle while the others represent surrounding suburban communities. Population and employment data for 1970 were developed for each node by aggregating 1970 census tract data for King County. The total population allocated among the 12 nodes is 941,000 people and the total employment is 478,000 jobs. The question to be addressed is, "How might we rearrange these people and jobs among the 12 nodes so as to substantially reduce the travel requirements of the system while maintaining a high level of accessibility?" Our approach to this problem is as follows:

1. Define a best and worst urban spatial structure for each of the four performance measures. Use these configurations to establish upper and lower bounds (i.e., a scale) for the investigation.
2. Calculate the travel requirements of the 1970 spatial structure and a uniform spatial structure (i.e., equal numbers of people and jobs at each node) to compare with the results of the search algorithm.
3. Use the search algorithm to find a good spatial structure for each of the four performance measures.
4. Determine the travel requirement reduction associated with each of the four spatial structures found by the search algorithm in relation to the 1970 base.

The results of each of these steps will be briefly discussed in turn.

Estimation of a Best and Worst Urban Spatial Structure for Each Performance Measure

Table 3 shows that the best and worst results for the 3-node and 5-node cases have quite distinct characteristics. These same characteristics were used to estimate the best and worst configurations for the 12-node case. For example, the best 3- and 5-node configurations for the weighted accessibility performance measure were found by assigning a maximum, balanced, people/job level to the most central node while placing a minimum number of people and jobs at all other nodes. The same logic was assumed to hold in the 12-node case. Other best and worst configurations were similarly determined by following the logical rules of Table 3. The range of performance between these best and worst configurations provides a scale that can be used to compare various configurations. This scale is shown in Figure 8 and is discussed in the following section.

Results of the Calculation of the Performance of the 1970 Spatial Structure and a Uniform Spatial Structure

The comparative performance of the 1970 spatial structure is shown in Figure 8. As can be seen, the 1970 system is within 25 percent of the best possible performance in all categories except weighted accessibility. These results suggest that our present urban configuration might not be as inefficient as the various critics of the American city would have us believe. These measures also suggest that the largest potential for improvement is in the total travel and weighted accessibility categories. For comparison purposes, a uniform spatial structure was constructed, and its performance is also plotted in Figure 8. The uniform spatial structure has an equal number of people

Figure 7. The 12-node network representation of western King County, Washington.

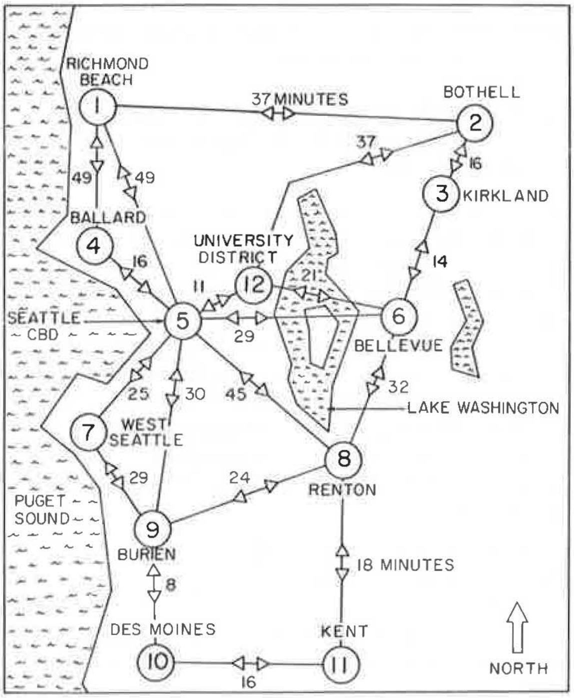

Table 3. Definition of best and worst configurations for each of the four performance measures.

	Configurations	
Performance Measures	Best[a]	Worst[b]
Total travel	Maximum people and jobs at least central node, minimum people and jobs at all other nodes, remainder at second least central node	Maximum people, minimum jobs at least central node, minimum people and jobs at all other nodes, remainder at second least central node or node as far removed from least central node as possible
Average link load	Same as total travel	Same as total travel
Maximum link load	Same as total travel	Maximum people and minimum jobs at two closest nodes
Weighted accessibility	Maximum people and jobs at most central node, minimum people and jobs at all other nodes, remainder at second most central node	Same as total travel

[a] Minimum total travel, average link load, maximum link load, and maximum weighted accessibility.
[b] Maximum total travel, average link load, maximum link load, and minimum weighted accessibility.

Figure 8. Comparative performance of alternate urban spatial structures.

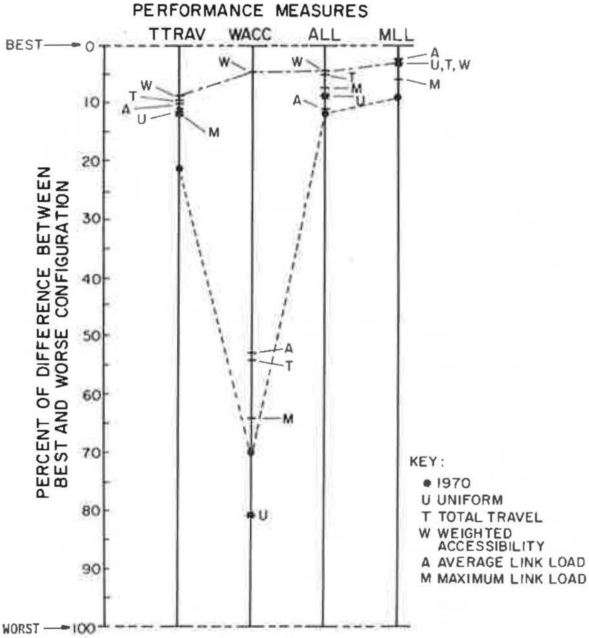

Table 4. Comparison of alternative urban configurations with the 1970 base configuration.

Configuration	Percent of People Moved	Percent of Jobs Moved	Percent Reduction/Increase of 1970 Performance				Sum of Nodal People/Job Ratio Deviations From 1.97
			Total Travel	Weighted Accessi-bility	Average Link Load	Maximum Link Load	
1970 base	0.0	0.0	n.a.	n.a.	n.a.	n.a.	52.5
Best configuration (total travel, average link load, maximum link load)	84.6	94.0	-66	n.a	-91	-87	0.03
Best configuration (weighted accessibility)	75.2	55.8	n.a	+243	n.a.	n.a.	0.03
Uniform	22.2	55.2	-46	-35	-27	-76	0.0
Search result I (total travel)	36.0	36.7	-54	+51	-51	-57	12.5
Search result II (weighted accessibility)	56.5	33.5	-61	+222	-78	-54	16.6
Search result III (average link load)	40.1	37.7	-53	+55	-36	-25	12.8
Search result IV (maxi-mum link load)	23.8	25.6	-43	+20	-34	-63	5.6

and an equal number of jobs located at each node. As can be noted, a uniform spatial structure is better than the 1970 situation in all categories except weighted accessibility, where it performs quite poorly.

Utilization of the Search Procedure to Find Good Configurations for Each Performance Measure

For the search process, the upper and lower limits on people located at any single node were set at 600,000 and 2,000. The bounds on jobs at any node were set at 300,000 and 1,000. In both the people-shifting and job-shifting phases of the search process, a step size of 1,000 was used. The 1970 configuration was used as the starting condition for all four searches. The results of these four searches are shown in Figure 8. As can be seen, the search procedure found configurations that were substantially better than the 1970 situation in all cases. Configuration W, which was found by using weighted accessibility as the single objective, is the best of these four configurations in three of the four performance categories and is the best spatial structure found by the search procedure in relation to all four performance measures. In no case did the search procedure find one of the four best configurations.

Travel Requirement Reductions Associated With the Alternative Spatial Structures in Relation to the 1970 Base

Table 4 gives the proportion of people and jobs that would have to be moved to other locations for the best, uniform, and each of the four search configurations in relation to the 1970 base. It also shows the percentage reduction of each of the three travel performance measures and the percentage improvement of the accessibility performance measure for each alternative configuration in relation to 1970 as a base. Table 4 shows that very large changes in the current urban structure would be required to achieve the best configuration. Although changes of this magnitude are clearly infeasible, even in 20 to 30 years' time, they do provide a type of upper limit in terms of what might be ultimately possible insofar as travel requirement reductions and maximization of accessibility are concerned. The uniform configuration requires far fewer people and job location shifts and is associated with substantial improvements in all performance measures except accessibility, which is worse than the 1970 level. Most interesting are the results of the four search runs. Very generally, these results indicate that by moving about a third to a half of the people and about a third of the jobs, reductions in total travel, average link load, and maximum link load of about 50 percent or more (below 1970 levels) could be expected. Increases in accessibility of from 20 percent to 220 percent could also be expected. These are very substantial reductions but would also involve quite large structural changes in the current urban pattern.

Changes in an existing urban structure of this magnitude are certainly not feasible in the near future (5 to 10 years), yet most large cities will probably grow by this much during the next 20 to 30 years. If this growth could be guided into appropriate locations, then we might experience a concurrent growth in transportation requirements that would be far less than might occur in a laissez faire situation where present trends were continued unchanged into the future.

Another way of describing the difference between the 1970 spatial structure and the five alternative spatial structures is to examine the balance between people and jobs at each node in the network for each case. The people/job ratio for the study area is 1.97. In the uniform distribution, the people/job ratio in each node is therefore 1.97. If we sum the differences between the people/job ratio at each node and 1.97, we obtain a rough index of how "balanced" a particular configuration is. The closer this index is to 1.97, the greater is the balance between people and jobs at each node in the network. These data are shown in the right-hand column of Table 4. They show that the land use balance index of each of the configurations found by the search algorithm is much closer to 1.97 than is the 1970 base. This means that the balance of people and jobs at each location in the system is a key factor in restructuring urban areas so as to reduce their transportation requirements.

CONCLUSIONS AND RECOMMENDATIONS

Our conclusions will address two questions: "How useful is the analytical approach used in the study?" and "What are some possible policy implications that can be derived from our results?"

Utility of Analytical Approach

This study has demonstrated that a simple search algorithm can be a useful tool for finding spatial structures that have desired characteristics. More powerful and reliable algorithms are needed because our simple algorithm performed only reasonably well on problems with known optimal solutions. The complexity of the search algorithm will undoubtedly have to be increased as the size and complexity of the problem increases. Most useful would be a search algorithm that will look for configurations that are better with regard to some combination of performance measures rather than for only one performance measure at a time.

The alternative to using a search algorithm is the fabrication of alternative spatial structures in one's mind. It is certainly possible that one could construct an adequate sample of all possible configurations judgmentally or by following systematically some logical decision rules. One would need to conduct a series of experiments along these lines before any definite conclusions on this issue could be reached. Until it can be shown that judgmental searching is more cost-effective than computer-directed searching, it seems reasonable to continue the development and testing of search algorithms for urban systems design problems.

Some Policy Implications of the Experiments

It has been shown that some dramatic reductions in travel requirements could be achieved by altering urban spatial structure. By logical extension, it has been argued that by guiding the growth of a city it should be possible to substantially reduce its needs for travel and transportation facilities and services relative to those that would result if present trends in growth patterns continued unchanged into the future. It has also been shown that substantial improvements in accessibility can be achieved by altering an existing spatial structure. However, these results have been derived using a very simplified representation of a real-world urban system and by using a very simple predictive model to generate estimates of the travel requirements of various urban configurations. This means that our results must be interpreted with caution and represent only a rough idea of the potential of altering urban spatial structure to reduce the need for transportation in our cities.

What appears to be most needed at this time is a way of identifying those particular locations where it would be most beneficial to encourage new people or jobs to locate. If such locations could be identified, public programs and policy could then be oriented to encouraging growth to occur in locations where the associated transportation requirements would be minimal. Other complementary programs oriented to the encouragement of particular changes in the existing urban structure could also be formulated with the aid of such a technique. Such a program would typically specify several locations where increases or decreases of residences and/or jobs would do the most good in terms of reducing future travel requirements as well as maintaining a high level of accessibility. The development of such a technique is high on our list of priority research tasks and will be the subject of a future research report.

ACKNOWLEDGMENTS

This paper is based on an unpublished Master of Science in Civil Engineering thesis by Joseph R. Beck. This thesis, completed in 1972, was edited and extended during the early part of 1973 by J. B. Schneider and James Clark.

This report was produced as part of a Research and Training Program in Urban Transportation sponsored by the Urban Mass Transportation Administration of the U.S. Department of Transportation. The results and views expressed are the independent products of university research and are not necessarily concurred in by the Urban Mass Transportation Administration.

REFERENCES

1. Doxiadis, A. D. Emergence and Growth of an Urban Region: The Developing Urban Detroit Area. Detroit Edison Co., 1966.
2. Einsweiler, R. C. A Case Study: Twin Cities—Mass Transit as a Factor in Metropolitan Planning. In Urban Mass Transit Planning (Homburger, W. S., ed.), Institute of Transportation and Traffic Engineering, Univ. of California, Berkeley, 1967.
3. Goldman, M. A Man-Computer Graphics System for the Evaluation of Transportation Impacts of Alternative Urban Spatial Structures. Unpublished Master of Science thesis, Dept. of Civil Engineering, Univ. of Washington, 1972.
4. Harris, B. The City of the Future: The Problem of Optimal Design. In Papers of the Regional Science Association (Thomas, M. D., ed.), International Academic Printing Co., Tokyo, Vol. 19, 1967.
5. Hemmens, G. C. Experiments in Urban Form and Structure. Highway Research Record 207, 1967, pp. 32-41.
6. Urban Transportation Research and Development. National Academy of Engineering, Committee on Transportation.
7. Rapp, M. H. Planning Demand-Adaptive Urban Public Transportation Systems: The Man-Computer Interactive Graphic Approach. Urban Transportation Program, Univ. of Washington, Research Rept. 71-4, 1972, 203 pp.
8. Scott, A. J. Combinatorial Programming, Spatial Analysis, and Planning. Methune and Co., London, 1971.

DISCUSSION

Fred L. Hall, Department of Civil Engineering and Department of Geography, McMaster University

Schneider and Beck suggest that it is feasible to attempt to reduce travel requirements in an urban area by a long-term restructuring of the city's spatial patterns. Certainly to judge by their conclusions there is great potential in such an approach to the urban transportation problem. The purpose of this discussion is to suggest that further studies of this topic must (a) give closer scrutiny to the performance measures used to judge urban spatial configurations, (b) find better search algorithms and means for evaluating them, and (c) employ clearer techniques for presenting results and assessing the potential of urban restructuring.

PERFORMANCE MEASURES

The authors state that investigations of land use-transportation relationships should be capable of answering the question, "What are the transportation criteria with which we can evaluate alternative land use patterns?" They do not return to this question, despite the fact that their own findings provide sufficient data to answer it for at least two of the four performance measures applied in the study.

With respect to its centrality to the study, the more important of these two measures is the weighted accessibility measure. It appears to be used as a surrogate for mobility, but the results suggest it is not a good one, and other considerations support that contention. Early in the paper Schneider and Beck state that they wish to find urban configurations that "(a) result in equal or better individual mobility and (b) produce a better level of environmental quality in the city." In the application to the Seattle area, this is translated as "substantially reduce the travel requirements of the system while maintaining a high level of accessibility." However, their results (Fig. 8) indicate that the 1970 Seattle spatial structure did not have a particularly high level of accessibility. In fact, its score is closer to the worst possible level than to the best. Yet they surely imply that mobility in the area is good.

Figure 8 also shows that the 1970 spatial pattern in Seattle performs well on the three other measures. One inference might be that these three are reasonable measures of the way people choose to behave but that the measure of weighted accessibility does not correspond to any form of present locational behavior. Closer examination of this measure supports at least the latter part of the inference. Their accessibility measure is a function of separation between all homes and all employment in the area. More likely, individuals locate with respect to only one job, but also with respect to various cultural, recreational, or environmental amenities that are not present in this particular measure. Hence to rely on it as a justification for restructuring an urban area seems dubious at best.

The other performance measure that can be assessed is that pertaining to the maximum link load. The authors suggest that large maximum loads may be either desirable (permitting high-capacity fixed route systems) or undesirable. They then assert that they wish "to find urban configurations that do not require high-capacity transportation facilities", so that the presence of a high maximum link load is taken to be a negative feature. This seems rather arbitrary, considering the ambiguity of the measure. The findings suggest it adds nothing to the identification of good configurations. (Figure 8 indicates that spatial configurations that perform well on other measures also tend to have low maximum link loads.) It therefore seems reasonable to omit this measure.

A replacement for it might be a measure relating to the amount of construction of new facilities required for any particular urban spatial configuration. It makes little sense to decrease total travel in a region if the bulk of this reduced travel is anticipated to take place along routes that are at present of very low capacity. This appears to be the kind of minimum travel solution to which several of the performance measures would lead, with the bulk of the people and jobs in the least accessible location. Would the existing transportation facilities be adequate to deal with such redirected traffic? Should this not be one of the primary criteria for a revised urban structure?

SEARCH ALGORITHM

The authors state that their study is primarily intended to examine the assertion that the combinatorial space representing urban structures is not smooth and that "any search procedure that attempts to find 'best' combinations on the basis of gradients would likely be ineffective." Their conclusions state that their simple search algorithm is useful but that better algorithms are needed. They do not refer to the assertion they claim to be examining, but the implication is that gradient-based searches can be effective.

However, data from their applications of the algorithm suggest that the algorithm is not particularly effective and hence that the assertion is reasonable. Figure 8 provides the clearest demonstration of the algorithm's failings. The urban structure with the highest weighted accessibility also gave the best algorithm-determined values for two other measures. That is, when attempting to minimize total travel explicitly, the algorithm did not arrive at as good a total travel figure as it did while minimizing a different measure. This failing is even more apparent for the maximum link load measure. Here every other application of the algorithm (i.e., to optimize each of the remaining three measures) produced a better value for the maximum link load than did the run in which it was directly optimized.

Clearly there exists a better logic for a search algorithm. The authors use heuristics derived from their first two tests to construct "best" and "worst" configurations. As this procedure produces a solution roughly 5 to 10 percent better than anything the search algorithm found, it would seem reasonable to reject the algorithm and devise a new one based on the heuristics described in Table 3.

RESULTS

Schneider and Beck conclude that "dramatic reductions" in travel could be accomplished by restructuring urban spatial patterns. However, they present their findings about the potential for improvement in a rather curious way. They state that reductions "of about 50 percent or more (below 1970 levels)" can be expected from moving roughly a third of the people and jobs in the Seattle area. But this reduction is not in terms of actual 1970 travel levels; rather, it is in terms of differences between the 1970 value and the best possible. To clarify what is meant, let X represent the best possible value and Δ the difference between the best and the worst. Then, reading approximate values from Figure 8,

$$\text{Total travel (1970)} = X + 0.22\,\Delta$$

and

$$\text{Total travel (T structure)} = X + 0.10\,\Delta$$

For this value of T-structure total travel to represent a 54 percent reduction from the 1970 total travel, Δ must be roughly 450 times the size of X. That is, the worst possible value of total travel must be more than 450 times as great as the best value. Although they do not supply actual numbers, it seems unlikely that this is the case: In the 3- and 5-node problems, the ratios of worst to best were 2.4 and 3.2 respectively. Using a similar magnitude for an example, a worst-to-best ratio of 10 in the Seattle problem would imply a reduction in total travel of about 36 percent of the 1970 level. While this is not as impressive as their 54 percent figure, it still represents a sizable amount of travel. Expressing potential reductions in terms of present conditions would provide a number that is simpler to understand and forms a more reasonable basis for decisions. Further, it would be a much stronger indicator of the importance of any future studies of urban restructuring.

AUTHORS' CLOSURE

Professor Hall's discussion suggests that further studies of this topic are needed. We fully agree and are presently engaged in such work. Beyond this, however, we find that his comments are either misdirected or are based on a mistaken interpretation of our results. Initially, he discusses the performance measures used to gauge the travel requirements of alternative urban spatial structures. His suggestion that the accessibility measure used could be improved by considering access to other than workplaces is good and we concur. This was done by Hemmens in a study we referred to conducted in 1966. Our reason for using only workplace access is that all available empirical evidence to date suggests that access to workplace is a far more important determinant of residential locational choice than is access to other nonwork activities. Our reason for interpreting high link loads as being undesirable was not arbitrary. Instead, it is based on the general philosophy of the paper, which is that we were seeking nontransportation solutions to current transportation problems. We stated our preference for the "no-build alternative" early in the paper, but this was apparently overlooked by Hall. His suggestion that an additional indicator related to "new construction required" is needed also fails to recognize the theme of our investigation.

The difficulties we experienced with the search algorithms were clearly stated by us and are only reemphasized by Hall. We have made no claims that the search algorithm we used is highly effective and we note in our conclusions that a better search algorithm is needed. Since this paper was written, we have developed an algorithm that uses a gradient search procedure, and it is now operational. Initial tests have shown that it is much more effective than the one used in the investigation being discussed here. Our purpose in this paper was to take a quick and rough cut at the problem, and this meant that extensive work on refining the algorithm could not be justified. The algorithm worked sufficiently well for us to produce results that we feel are encouraging enough to warrant a second cut at the problem. This investigation will be more detailed, rigorous, and elegant.

Hall's claim that our results are computed improperly is based on his mistaken interpretation of them. Our general conclusion that substantial reductions in travel could be accomplished by restructuring the city is not derived from the data presented in Figure 8, as Hall asserts. They were calculated using the 1970 situation as a base, and this is clearly indicated in Table 4. We feel that our conclusions, while derived from a crude, macro-scale analysis, are sufficiently encouraging to warrant further investigation of this topic, and there is nothing that Hall has included in his comments that gives us any reason to think we are not pursuing a proper course of action.