to one another. Urban representations need not be complete circles. Lakes or harbors can be represented by assigning zero population and employment. Discontinuities can be introduced into corridors to reflect rivers or other geographical barriers. Once the corridor structure is established, residential and major activity center modules are identified and located on the circular structure. The product of the representation work is an urban structure that can be entirely or partially analyzed by using the SMART model.

REFERENCES

1. D. E. Ward. Algorithms for Efficient Transit Systems Design. Transportation Systems Center, U.S. Department of Transportation, Cambridge, MA, 1976.
2. F. J. Mason and J. R. Mumford. Computer Models for Designing Dial-a-Ride Systems. SAE Automotive Engineering Congress, Detroit, 1972.
3. L. S. Englisher and K. L. Sobel. Methodology for the Analysis of Local Paratransit Options. TRB, Transportation Research Record 650, 1977, pp. 18-24.
4. R. J. Smeed. Traffic Studies and Urban Congestion. Journal of Transport Economics and Policy, Jan. 1968.
5. L. Bronitsky, M. Costello, C. Hoaland, and S. Schiff. Urban Data Book. U.S. Department of Transportation, Rept. DOT-TSC-OST-75-45.I, 1975.

Discrete Optimization in Transportation Networks

Paulo A. R. Lago*, PROMON Engenharia, S. A., São Paulo, Brazil

In most cases, planning capital investment in transportation networks is an unwieldy job because the number of investment options grows so rapidly. The real situation faced by the transportation planner is, in general, when, where, and by how much to allocate available resources. The transportation investment problem can be characterized as the location and timing decisions to be made by the planner. A branch-and-backtrack algorithm is presented that tackles both location and timing aspects of the capital investment problem in small and medium transportation networks. The results presented are encouraging for future research in which the technique can be applied to larger, actual transportation networks.

The problem addressed in this paper is a common one in transportation planning. Given an existing network of M links, a set of future supplies at each origin in the network, and a set of future demands at each final destination, when, where, and by how much should additional investment be dedicated to each link?

For the purposes of this paper, a link can be either a physical connection between two geographically separated points, such as a rail line or a big highway, or it can be a transshipment facility such as a port. It is assumed that demand at each destination and supply at each origin are inelastic—that is, independent of transportation cost. Under this assumption, minimization of present-value social costs aggregated over the network is consistent with maximum national income, and this is the objective function used throughout. The algorithm presented can be extended to price-sensitive supply and demand without computational difficulty by using Devanney's method (1) and replacing cost minimization with maximization of the present value of the sum of the consumer's and producer's surplus. It is also assumed that, whatever the investment, all links are priced at their marginal social cost. In the transportation planner's vernacular, "system-optimized" rather than "user-optimized" operation is assumed. This is in part a reflection of my interest in freight rather than in passenger transportation and in part a reflection of my philosophy that, wherever the results of user-optimized operation differ greatly from those of system-optimized operation, the costs of administering a marginal-cost toll system on nonurban networks can and should be borne.

Finally, it is assumed that future growth of demand and supply is known with certainty. Before we can tackle uncertainty, we must have an efficient algorithm for handling investment with certainty (2). Indeed, one of the goals here is an algorithm that is efficient enough to be routinely run over a range of hypotheses for growth of demand and supply.

Even given all the assumptions above, the magnitude of the problem can be appreciated by considering a network with M links, T possible investment periods, and N possible levels of investment on each link in each period. Then the number of possible investment strategies is N^{MT}. Consider a very small network with four links, three levels of investment on each link, and five investment periods. The number of potential investment strategies is 3.5×10^9. To solve such a problem, for each such investment strategy examined one must

1. Compute the set of equilibrium flows in the network for each period and present value and the resulting link flow costs and
2. Combine the present-value flow costs with the present value of the fixed costs associated with this investment strategy.

In short, each investment strategy examined requires the solution of T network flow problems. Even with the very efficient available algorithms for network flow, direct enumeration of all investment strategies is clearly out of the question for even a very small network.

The major reason that our problem is so different is that we have combined allocation in space with allocation in time. Most work on investment across links has assumed only a single possible investment point in time. Either an investment in a link is made at that time, or it is never made. In reality, given the generally continuous growth in transportation demand, investment timing is as important as investment location. Yet most work on investment scheduling has assumed a network that consists of a single link. This clearly is of little use in determining which of a number of competing links should receive the planner's attention. In short, in many real-world cases, both the location and the timing of investment are crucial. Furthermore, these two dimensions are so closely coupled that, unless they are handled simultaneously, seriously misleading results can be obtained.

The algorithm discussed in this paper is derived from the branch-and-bound technique. Because of its specific branching procedure, it is called in the literature the branch-and-backtrack method. Computational results of tests on several small networks are presented. Although the sample of test results is small and computation time can be very sensitive to the specifics of a particular problem, it appears that the algorithm can efficiently handle the dimension of timing of capital investment. In addition, even if the algorithm must be cut off before optimality, it results in a feasible solution that can then be compared with the best that has been obtained by other means, including the intuition of the network designer.

DESCRIPTION OF THE PROBLEM

The problem can be briefly described as follows:

1. At each of T time periods indexed by t and for each of I origins indexed by i and each of J destinations indexed by j, there are a given demand $D_i(t)$ and supply $S_j(t)$.
2. In addition, there are M different links (actual or potential) in the network that connect the I origins to the J final destinations. On each such link, there are N possible levels of investment. Associated with each possible level of investment on each link is a fixed investment cost $F_m(n)$. This fixed investment cost should include not only the initial cost of the improvements to the link but also any future maintenance costs that are independent of the level of flow on that link present-valued back to time of investment. In addition, there are the flow-dependent costs $V_m(x_t, n_t)$ in each period t, where x_t is the level of throughput on link m and n is the level of investment already in place in that period.
3. Finally, we are considering an overall time horizon of T periods indexed by t. Investment on any link may occur at the beginning of each of the periods. Whatever the level of investment is in each link at each period, it is assumed that the network is operated so as to minimize the total flow-dependent cost of satisfying the given demands from the given supplies in that period. This is also the short-run equilibrium set of flows under textbook competition, such as that observed in the world tanker network. Our problem is to compute the investment strategy—$n_m(t)$, the level of investment in each link in each time period—that minimizes the sum of the present-valued fixed cost of the network and the present-valued flow-dependent costs of operating the network under the chosen pattern of investment.

The first step in tackling such a problem is to separate the fixed investment decisions from the resulting short-run flow patterns. For any trial investment pattern $n_m(t)$, the problem of determining the resulting set of short-run flows for each time period t is a simple network flow problem for which a number of extremely efficient algorithms exist. One of the most efficient algorithms—some would argue the most efficient—is the University of Texas primal algorithm PNET (3), which is used in this study. For a given investment pattern in both space and time, it is a relatively straightforward problem to apply PNET to the network T times and determine the minimum flow-dependent costs in each period and the present value. This process can then be repeated for other trial investment patterns. The overall scheme then can be shown as it is in Figure 1.

This basic decomposition makes a great deal of computational sense in that it separates the overall problem into two parts, one of which can be solved very easily. In addition, it represents a natural separation from an economic point of view, dividing the problem as it does into its short- and long-run components. Such an explicit treatment of the short run allows us to model demand growth and, among other things, generates the optimal tolls on each link in each period for each investment pattern studied. These tolls will appear as the duals associated with the corresponding link variable cost function. If one is unwilling to assume one has direct control over the network operations or the ability to levy congestion tolls, then the actual short-run flows will be user optimized. The basic decomposition can still be used to generate "second-best" investment patterns. In so doing, one would use one of the network flow algorithms that generate user-optimized flow patterns to simulate the network in each period. From this flow pattern, one can compute the corresponding flow-dependent social costs and then proceed as before.

Transportation investment by its nature tends to come in large, discrete chunks. It simply does not pay to add half a lane to a highway or half to a port. Given such indivisibilities and a continuous growth in transportation demand, even under an optimal investment pattern, individual links will almost never be operating at design capacity. At any given point in time, some links will be operating below design capacity and some above. Hence, analysis that assumes such an unattainable long-run equilibrium will not only yield misleading results with respect to investment but will also yield no results as to how the network should be operated and priced through time. A more complete discussion of the coupling between long-run investment and short-run pricing, given indivisibilities in capital investment, is presented by Devanney (1).

The real problem lies in the left-hand side of the diagram in Figure 1—that is, in the method for choosing the investment patterns to be costed out. We have already seen that direct enumeration of all possible investment patterns is clearly infeasible because of the

Figure 1. Overall decomposition into fixed investment decisions and short-run operating decisions.

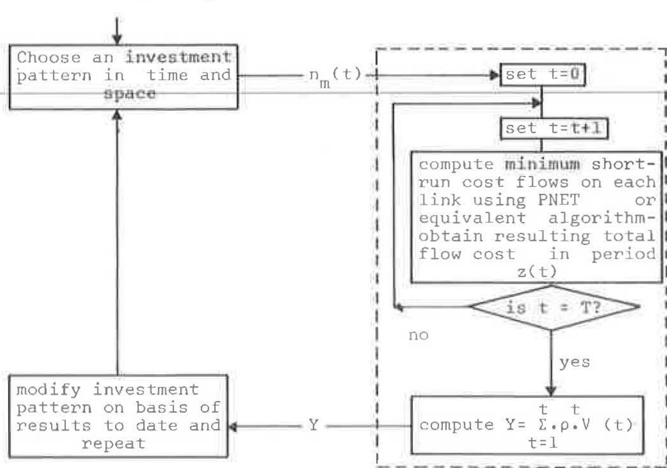

number of such alternatives. Two basic approaches that use the decomposition of Figure 1 have been suggested: (a) Bender's decomposition, which at each iteration of the left-hand portion of Figure 1 generates an integer problem, and (b) the branch-and-bound technique. The method examined in this paper is a variant of the branch-and-bound technique. A comparative effort at the Massachusetts Institute of Technology (MIT) is studying Bender's decomposition.

Several authors have applied branch-and-bound to capital investment in transportation networks. Perhaps the most significant work for our purposes is that of Ochoa-Rosso (4), who in 1968 presented formulations for four different transportation problems. All of those problems, however, dealt with the static problem and used a single target year. Ochoa-Rosso did not deal with the problem of multistage improvement—the scheduling dimension—but only cited it as a potential field needing further study. He used the bound criteria previously used by Ridley in 1965. In essence, the formulations presented by Ochoa-Rosso and later by Tan (5) assume that the planner has the option to perform a single investment now or reject it completely.

Ochoa-Rosso and Silva subsequently applied one of Ochoa-Rosso's formulations in a case study on the Puerto Rico system of seven nodes and four links (6). They assumed one target year and applied two methodologies—branch-and-bound and branch-and-backtrack (a variation of branch-and-bound with a different branching sequence)—to select the optimum choice out of 16 possible alternatives. There do not appear to have been any significant improvements in the application of branch-and-bound to transportation network investment since Ochoa-Rosso's work.

BRANCH-AND-BACKTRACK METHOD

Key Assumptions

To solve the network investment problem outlined above by branch-and-bound, it is necessary to make two basic assumptions about the form of the flow-dependent cost functions on each link.

The first assumption is that the variable flow-dependent cost on each link does not increase with the amount of investment committed to each such facility. In other words, as the investment level in a transportation link increases, the variable cost associated with handling a given amount of traffic decreases. Still more concisely, it is assumed that the partial of all link flow-dependent cost functions with respect to link investment is nonpositive: $[\partial VC(x, I)/\partial I] \leq 0$, where VC (x, I) is the variable cost associated with an investment level I and a flow of x. This hypothesis is not particularly limiting. In general, this is the situation for the bulk of transportation facilities. If the planner of a roadway invests to provide two lanes, the variable cost associated with a certain level of traffic will be higher than the variable cost associated with the same level of traffic if the road had four lanes instead. This will certainly be true for all but very low levels of throughput.

Another typical example would be the level of investment in a road in terms of the construction material. As long as the quality of material to be applied on the surface of the road is increased, the variable cost associated with a specific level of traffic will be the same as or less than the cost that would have resulted from investing in a cheaper, lower-quality pavement (see Figure 2, where F_i = fixed cost associated with investment i, q = traffic flow, and h = variable cost).

For traffic flow q′, the variable cost is the same, independent of the fixed costs. For q_1, the variable costs for investments I_2 and I_3 are the same but the variable costs for I_1 are higher. The same thing happens for q_2 and q_3.

The second basic assumption is even less controversial. Assume that the partial of the flow-dependent link cost functions with throughput is nondecreasing: $[\partial VC(x, I)/\partial x] \geq 0$. The partial is generally called the marginal social cost. For most transportation technologies, this marginal unit is constant or nearly so at low levels of throughput and increases rapidly as throughput approaches and passes the design capacity of the task (see Figure 3, where MC = marginal cost). Independent of the level of investment, as traffic flow increases from q_1 to q_3, the marginal associated cost will first be constant and then, as the flow approaches capacity, its value will abruptly increase.

In any event, these two rather weak and generally realized assumptions are the only requirements in the functional form of network costs that must be improved so that the branch-and-bound algorithm can be used. All sorts of functions are possible within this general framework, including economies and diseconomies of scale with respect to investment.

Example

Consider a single-period problem that involves three links, for each of which low, medium, and high levels of investment are possible. Assume further the sets of costs given in Table 1 by level of investment and link. Note that the link flow costs decrease with increasing investment, as required by the algorithm.

The problem can be represented by a decision tree such as that shown in Figure 4. The left-hand three-way branch represents the decisions for link 1. Whatever is decided for link 1, we are then faced with the decision for link 2. This three-way choice for each possible link 1 decision is represented by the second set of branches on the tree. Finally, the link 3 choices are represented by the right-hand set of branches. There are 27 terminal nodes to the tree, representing the 3^3 possible patterns of investment. The number in parenthesis at the tip of each right branch in Figure 4 represents the order of evaluation as the algorithm proceeds. More detail about the algorithm can be found elsewhere (7).

APPLICATION OF TECHNIQUE TO TRANSPORTATION NETWORKS IN VENEZUELA

Caracas Transportation Links

Problem

The branch-and-backtrack algorithm was tested on a very simplified version of a transportation investment problem that currently faces Caracas, the capital and largest city of Venezuela. Caracas attracts commodities from the inland and from abroad as well. As the most developed city in the country, Caracas also provides the rest of the country with products manufactured by its industries, some of which are also exported.

Since Caracas is located 12 km (7.5 miles) from the seashore, all exports and imports currently travel over a single highway that links Caracas to the nearest port, La Guaira. The road was built in 1952 as a two-lane, two-way highway and was later widened to a four-lane, two-way highway and was later widened to a four-lane, two-way highway. Another important characteristic of the geography of Caracas is that the city is about 939 m (3100 ft) above sea level. The mountainous route between Caracas and La Guaira has two two-way tunnels and two two-way bridges, each of which is about 970 m (3200 ft) long.

Traffic on the system is currently congested, particularly at peak hours. Aside from the import and export traffic, local goods destined for various parts of the Venezuelan coast (cabotage) as well as traffic to and from Venezuela's only international airport at La Guaira contribute to the demands on the road.

The problem faced by planners in Caracas is how to improve the transportation system so as to avoid future congestion. Their options include investment in the road to increase its capacity, investment in the La Guaira port to improve the service capacity provided there, and the construction of a new rail system to link Caracas with Puerto Cabello, a port about 121 km (75 miles) away. In connection with the construction of this new rail line, the Caracas planners must also decide on the amount of investment that is required to make Puerto Cabello a feasible alternate port to La Guaira (see Figure 5).

In essence, this represents a typical investment problem in which there are four possible transportation links: the Caracas-La Guaira road; the port of La Guaira; the Caracas-Puerto Cabello rail line; and the port of Puerto Cabello. Although it is simple, such a transportation network can be used to test the applicability of the previously developed branch-and-backtrack technique in investment decisions. More details about the cost model used in this example can be found elsewhere (7).

Figure 2. Total investment costs.

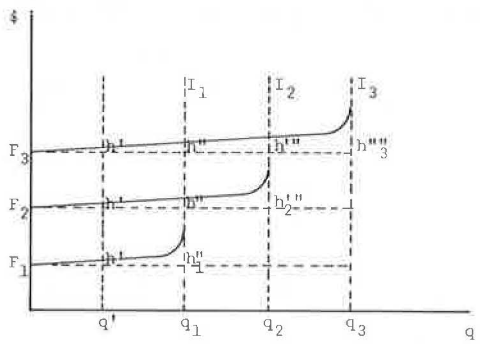

Figure 3. Total investment and marginal costs.

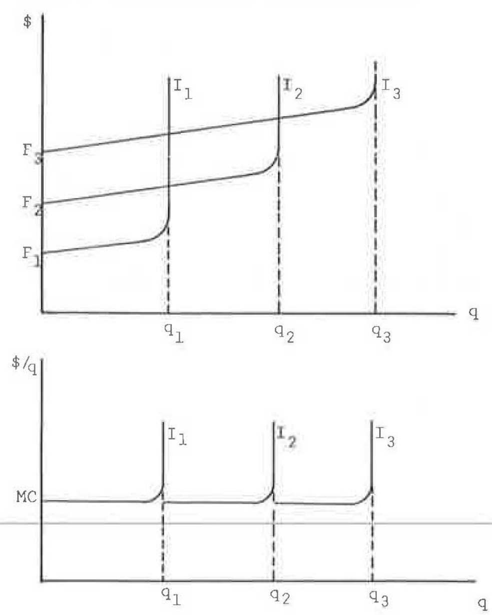

Table 1. Investment and flow-dependent costs for three links and three levels of investment.

Link	Fixed Costs Low	Fixed Costs Medium	Fixed Costs High	Present-Value Flow Costs Low	Present-Value Flow Costs Medium	Present-Value Flow Costs High
F_1	6	8	20	20	15	12
F_2	9	15	24	17	10	8
F_3	2	10	18	7	5	3

Results

It was assumed that the Venezuela planners would determine the optimal investment policy in 1959 and then, assuming no disinvestment for the subsequent years, would at discrete points in time (say, years) search for new investment plans from that point on to improve the previously determined option. Thus, in the initial tests we did not attempt a true optimization over the time dimension but a series of suboptimizations in

Figure 4. Example of branch-and-backtrack method.

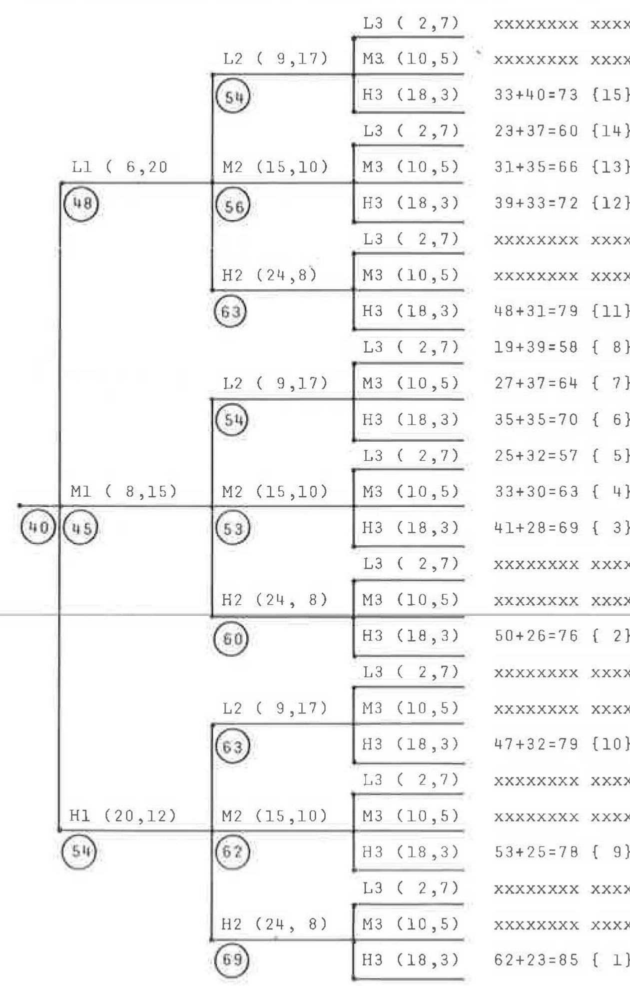

which, at any point in time, the transportation planner could make a particular investment "now or never". Thus, the total number of possible investment patterns was 625 × 10 rather than the true 625^{10}.

Five different problems of this sort were run with randomly varied coefficients, and the results are given in Table 2. For the five cases given, the common information shown in Figure 4.4 of the report by Lago (7) was assumed. Furthermore, the link characteristics associated with each investment set, such as handling rate (for each port), road capacity and operating cost, and rail capacity and operating cost, were varied.

Two comments should be made about this set of suboptimizations. In 1959, Venezeula planners are assumed to be faced with only the 625 immediate investment options. After the given number of iterations, they obtain their answer for 1959. In 1960, the same planners, under the restricted ground rules of this test, are faced with the same problem but given the investment that has already been made. This procedure is then repeated for each of the remaining years. The last five computations, given in Table 3, show the cost involved.

The algorithm was developed and run on an IBM 370-168 computer at the MIT Information Processing Center. Figures given include central processing unit, memory, and input-output (I-O) cost but not setup and print charges.

Unfortunately, we did not keep a record of the exact modifications established for the first, second, and fourth cases. It appears that sometimes only the demand schedule was changed and at other times the link characteristics were changed. In the last computation, the system was "forced" to evaluate all possible investment alternatives through its lifetime—i.e., 625 evaluations at each decision or 6250 evaluations. The resulting execution cost gives an idea of the economy that can be achieved by using the branch-and-backtrack algorithm.

On the basis of these encouraging results, it was decided to try to solve the investment scheduling problem in a truly optimal fashion. Another trial was therefore conducted to verify the practicability of the method (see Table 4). At each decision point, four facilities and two levels of investment were assumed so that there were 16 alternatives open to the planner and, in the 10-year life, 16^{10} or 1.1×10^{12} options. In these three cases, the basic information given in the Lago report (7) was assumed. The magnitudes of investment were changed in the first two cases in such a way that the first values were half of the second values. It is worth considering the great sensitivity in terms of the number of performed evaluations that is apparent whenever the investment level is increased. For the third case, although the optimum solution has not been reached, the best feasible point up to that printing limit is obtained. Thus, a good result is obtainable by use of the algorithm even under computer-time or budget constraints.

Given that the proposed algorithm evaluates the vector of the highest investments as its first point in the tree, if the associated fixed costs are not significant in comparison with the total cost, then this point is closer to the optimum one. In such a case, solutions are obtained in few iterations (cases 1 and 2 in Table 4). On the other hand, if fixed costs are very large in comparison with total costs, then the actual optimum will be much farther from the initial solution and it will take the algorithm a great deal longer to come up with the optimum (case 3). Once again, we see the importance of a good initial solution and the role of the transportation planner's judgment.

Figure 5. Caracas rail and road links.

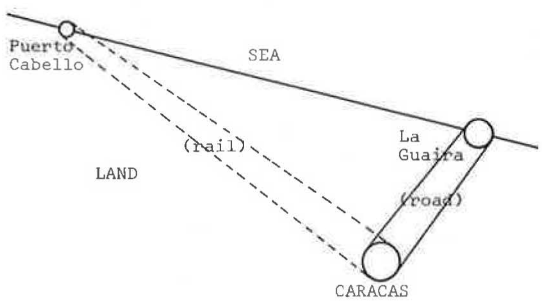

Table 2. Initial suboptimization runs for Caracas transportation links.

Case	Number of Links	Number of Investments on Each Link	Total Number of Investment Patterns	Number of Iterations	Number of Evaluations Performed
1	4	5	6250	29	140
2	4	5	6250	18	173
3	4	5	6250	133	214
4	4	5	6250	76	140
5	4	5	6250	117	164

Figure 3. Computations for Caracas problem showing cost of execution.

Case	Number of Facilities	Number of Investments	Total Number of Iterations	Number of Evaluations Performed	Cost of Execution ($)
1	4	5	6250	158	0.76
2	4	5	6250	150	0.71
3	4	5	6250	95	0.61
4	4	5	6250	109	0.79
5	4	5	6250	6250	12.27

Table 4. Final computer runs for Caracas problem.

Case	Number of Periods	Number of Possible Investments at Each Period	Total Number of Investment Patterns	Number of Evaluations Performed	Cost of Execution ($)
1	10	16	1×10^{12}	16	0.47
2	10	16	1×10^{12}	385	1.38
3	10	16	1×10^{12}	7196	26.55

Venezuela Transportation Network

Problem

On the basis of the preliminary but rather encouraging results discussed above, it was decided to apply the proposed branch-and-backtrack algorithm to a larger, more realistically sized transportation network. To exploit the information already obtained in the Caracas case, we decided to analyze a simplified transportation network in Venezuela.

The first problem was to characterize the Venezuela transportation network in 1959, the year for which cost information was available. This proved to be impossible. A scenario was therefore hypothesized for 1959. This assumption did not defeat the purpose, which was to test the validity of the proposed algorithm for a dynamic situation in a medium network.

The problem is presented schematically in Figure 6. Twenty-two cities were selected, and 48 transportation facilities were defined. Each arc in the figure represents one facility and each node the city of transshipment. It was assumed that in 1959 the Venezuela transportation network had basically two modes of trans-

Figure 6. Representation of Venezuela transportation network for analysis of investment problem.

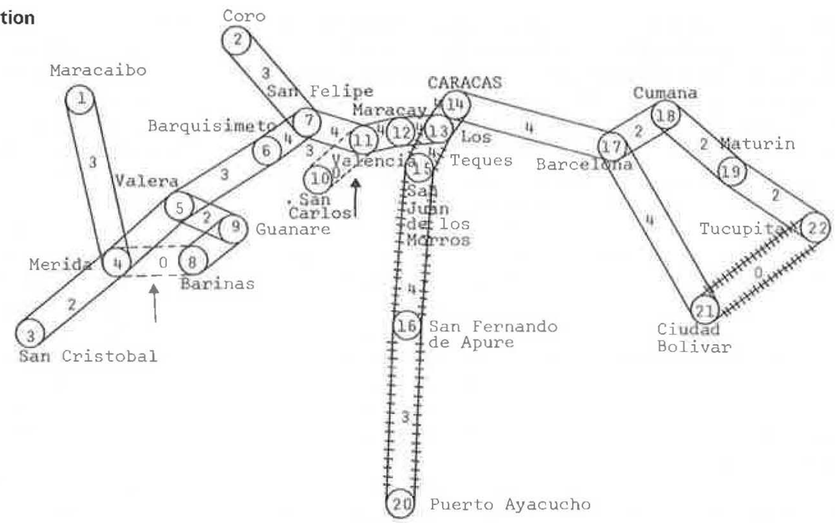

Table 5. Computer analysis of investment problem for simplified Venezuela transportation network.

Case	Number of Facilities	Number of Investments	Number of Iterations	Number of Evaluations Performed	Cost of Execution ($)
1	5	27	1.4×10^7	27	3.00
2	5	27	1.4×10^7	40	3.38
3	5	27	1.4×10^7	51	3.55
4	5	27	1.4×10^7	385	17.65

portation—road and rail. Although the country has navigable rivers, no data on water transportation were available, and so it was eliminated from consideration. To simplify the transportation investment problem, it was also assumed that only internal movements within Venezuela were being studied. Except for the broad view of costs, the problem has all the same characteristics, in terms of investment, costs, etc., as the previously described Caracas case.

Two railroads were considered (Figure 6): the one from Los Teques (13) to Puerto Ayacucho (20) and the one from Barcelona (17) to Ciudad Bolivar (21). All other facilities were assumed to be roads. For each facility, one level of capacity was arbitrarily assigned (see the numbers inside the arcs in the figure). Thus, whenever potential investments were analyzed, there was a basis for evaluating possible capacity improvements.

Three potential locations for the construction of new transportation facilities were hypothesized (a) a road from Merida (4) to Barinas (8), (b) a road from San Carlos (10) to Valencia (11), and (c) a railroad from Tucupita (22) to Ciudad Bolivar (21). Three levels of investment (low, medium, and high) were also hypothesized for each of the above facilities. The period of analysis was assumed to be five years. Thus, at each decision point in time there were 3^3 or 27 possible combinations and, through all five years, 27^5 or 1.4×10^7 investment strategies to be studied.

For this problem, we selected the mathematical programming system for solving network flow problems, PNET. This system was easily incorporated in the algorithm as a subroutine of calculation. In each evaluation, this equilibration procedure was used and it was assumed that the demand-supply pattern was inelastic (if desired, this assumption can be relaxed).

Results

The investment options used (7) represent the investment values faced by the planner at each period of time. In addition, a unit cost of transportation, besides the upper and the lower bounds, was associated with each link. The results obtained are given in Table 5 (the last case was interrupted by the limit on computer time).

The same level of investment was assumed at each decision point. Associated with these fixed costs were various link characteristics such as capacity and operating costs. Furthermore, each link on the network had a fixed unit cost and upper and lower bounds. For supply and demand schedules, it was assumed that at each period one origin point supplies one particular value and one destination point demands another fixed value.

CONCLUSIONS AND RECOMMENDATIONS

1. The results obtained for the Caracas problem were encouraging. The capital investment aspects of location and timing decisions were tested, and feasible solutions were found at reasonable expense. The major finding is that in most cases the aspect of investment timing can be jointly analyzed with location and that this can be done within a reasonable range of work. Both small and medium transportation networks were examined under multistage investment decisions, and feasible solutions were obtained.
2. Both of the problems examined show that, by using the proposed algorithm, the planner should come up with a very good solution within assumed budget constraints.
3. The planner's initial feelings should be incorporated in the algorithm to save many extra computations before convergence toward the optimum solution. This is verified in the example shown in Figure 4: Depending on each investment assumed for link 1, different costs are found in obtaining the optimum result.
4. The Caracas problem illustrates the feasibility of running the branch-and-backtrack algorithm under different conditions. Let us say that five demand schedules that cover a reasonable range are hypothesized. This would give a good idea of possible strategies to be selected by the planner.
5. As a result of observations made during the running of the Caracas problem, it was concluded that, whenever more constraints are presented in the sys-

tem, fewer iterations will be performed before the optimum solution is reached. It would be advisable to incorporate, let us say, budget constraints to be faced by the planner at each decision point.

6. Both problems analyzed present completely different structures. The first concerns a small network, and the user-optimized rule is used for network equilibration. The second concerns what can be regarded as a medium network, and the equilibration procedure used is the system-optimized rule. The proposed algorithm could be applied in both cases. This illustrates its versatility.

7. As Ochoa-Rosso (4) points out, more research should be devoted to the study of the trade-off between the branch-and-bound and branch-and-backtrack methods. Although the first requires greater computer memory, the second is more time consuming. This needs to be verified.

8. The technique proposed here for capital investment problems should be compared with another optimization procedure such as the Bender method.

ACKNOWLEDGMENT

I would like to acknowledge the support provided in this research by the Research Support Foundation of the State of São Paulo, Brazil, and the U.S. Institute of International Education. I would also like to extend my appreciation to John William Devanney III for his supervision of this work and his interest at various points.

REFERENCES

1. J. W. Devanney III. Marine Economics. Massachusetts Institute of Technology, Cambridge, lecture notes, 1977
2. J. W. Devanney III. Marine Decisions Under Uncertainty. Cornell Maritime Press, 1971.
3. F. Glover, D. Karney, and D. Klingman. Implementation and Computational Comparisons of Primal, Dual, and Primal-Dual Computer Codes for Flow Problems. Center for Cybernetic Studies, Univ. of Texas, 1973.
4. F. Ochoa-Rosso. Applications of Discrete Optimization Techniques to Capital Investment and Network Synthesis Problems. Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, 1968.
5. P. A. Steenbrink. Optimization of Transport Networks. Wiley, New York, 1974.
6. F. Ochoa-Rosso and A. Silva. Optimum Project Addition in Urban Transportation Networks Via Descriptive Traffic Assignment Models. Department of Civil Engineering, Massachusetts Institute of Technology, Cambridge, 1968.
7. P. A. R. Lago. Capital Investment in Transportation Networks Using the Branch-and-Backtrack Method. Department of Ocean Engineering, Massachusetts Institute of Technology, Cambridge, 1977.

**P. A. R. Lago was a graduate student at the Massachusetts Institute of Technology when this research was performed.*

Residential Area Location Preference Surfaces

W. Young and A. J. Richardson, Department of Civil Engineering, Monash University, Clayton, Australia

Although an understanding of the interaction between land use and transportation is essential to a rational evaluation of urban and regional policy, it is frequently complicated by the introduction of sophisticated mathematical techniques. In an effort to make this interaction more visible to the decision maker, two of the more advanced techniques—multinomial logit analysis and mental maps—are placed in a common framework of analysis and presentation. The strength of a rigorous theoretical background is thus combined with the simplicity of a visual presentation. The theory and development of the technique are outlined, and its use in a case study of the residential location preferences of residents of the inner suburbs of Melbourne, Australia, is described.

An understanding of the ways in which transportation investment, activity placement, and residential location interact is essential to a rational evaluation of urban or regional policy alternatives. Frequently, however, the methods used by planners to examine these interactions are complicated by the introduction of sophisticated mathematical models. Although such models may improve the explanatory power of the planning method, such an improvement is frequently made at the expense of the layman's understanding of the method.

If one wishes to make the interactions clearly visible to the decision maker, who frequently is not aware of the mathematical complexities involved in the modeling process, a clear, concise method for the presentation of results and implications must be devised. This paper attempts to provide such a method and at the same time to use two of the more advanced mathematical techniques in the analysis of location decision: multinomial logit choice modeling and the concept of mental, or cognitive, maps.

The approach, which is shown schematically in Figure 1, has essentially three stages. In this paper, the model is developed in the context of urban residential location. However, the basic model structure, as outlined in Figure 1, could well be applied to problems that involve regional development policies, decentralization, or alternatives of facility location.