The PR 2 extension had the highest $\Delta R/\Delta C$. This extension was made to a commercial area in order to serve reverse-commutation trips from Rochester. This extension of an express park-and-ride line was the only 1 of the 10 extensions to serve a commercial area. It has been suggested that high-quality transit service at a high price is most likely to be self-supporting.

The PR 1 and PR 2 extension into Kodak Park also proved to be profitable despite the fact that it was the longest of the 10 extensions. The previous comment concerning high-quality service is also applicable here. Extensions to areas of significant employee concentration appear to be most promising in terms of R/C ratio.

The RIT extension to a residential area was the shortest of the 10 extensions. This demonstrates the importance of the length of the route extensions.

The route 89 extension brought service within easy reach of public-housing residents, many of whom are captive transit riders. Local factors also contributed to the positive ridership response to this extension.

In conclusion, size of population, type of land use, quality of service, and length of extension are four major factors in the determination of the success of route extensions. Areas that have a significant concentration of employees seem most likely to support profitable extensions. Special local conditions can also influence ridership changes connected with route extensions. A general R/C model can be used to evaluate route extensions, and the criteria used to judge extensions can be left to the discretion of local operators. The problem of increased headways associated with route extensions resulting in a decline in service on the original portion of the route must be taken into account when it arises. Finally, conventional units of data collection (such as census tracts or TAZs) are too large for the purpose of evaluating route extensions.

Directions for further research in the area of route extensions are clear. Collection of data on a small scale commensurate with the area actually served by an extension and explicit correlation of these data with changes in ridership are the immediate next steps to be taken. The development and testing of attraction functions for different types of land use follow these steps. A general predictive model of the effects on transit ridership of route extensions can then be constructed. This paper has suggested the basics for such a model and has provided preliminary findings concerning the most salient factors in determining the outcome of a proposed route extension.

REFERENCES

1. J. P. Attanucci, L. Jaeger, and J. Becker. Bus Service Evaluation Procedure: A Review. Urban Mass Transportation Administration, Interim Rept., UMTA-MA-09-7001-79-1, March 1979.
2. D. L. Weiss. The Effects on Automobile Energy Consumption of Transit Fare and Service Changes. New York State Department of Transportation, Albany, memo, Feb. 16, 1978.
3. Public Transportation Operating Assistance Programs in New York State: 1978 Annual Report, Technical Supplement. Planning Division, New York State Department of Transportation, Albany, March 1979.

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

Hierarchical Procedures for Determining Vehicle and Crew Requirements for Mass Transit Systems

Lawrence D. Bodin and Robert B. Dial

This paper presents procedures for determining vehicle and crew requirements for mass transit systems. Some of these procedures are very fast computationally but only give lower bounds, upper bounds, or estimates of resource requirements. Other procedures are slower computationally but give actual crew and vehicle schedules. Depending on the type of analysis being performed (long-range planning, short-range planning, or operational planning), all of these procedures play a useful role in the design and analysis of proposed mass transit systems. The paper has two sections: (a) the first discusses techniques for determining vehicle requirements and (b) the second discusses techniques for determining crew requirements. Within each section are a set of procedures that range from the very simple to the complex, along with comments on their usefulness and shortcomings.

The design of mass transit systems occurs in various planning scenarios: long-range planning (5-20 years in the future), short-range planning (1-5 years in the future), and operational planning (less than 1 year in the future). The long-range planning analyst does not need (and cannot afford) the same information on crew and vehicle requirements as the operational planner. Whereas the operational planner needs actual feasible crew and vehicle schedules, the long-term planner may only need an estimate or lower and upper bounds on total crew and vehicle requirements for the analysis. Thus, the long-range transit planner should use fast crude estimation procedures to help evaluate a proposed transit system, since he or she may consider scores of alternative transit systems in attempting to find the optimal system.

In this paper, hierarchical procedures for determining crew and vehicle requirements are given. Some procedures require only manual calculations and furnish inexpensive (albeit crude) estimates. Others consume a significant amount of computer time and give more accuracy and detail. As will be seen, if the planner requires a more exact or more detailed vehicle or crew schedule, a higher cost must be absorbed in terms of computer time and human effort.

Table 1. Timetable 1.

	Time		Location			Time		Location	
Trip	Start	End	Start	End	Trip	Start	End	Start	End
1	7:03	8:13	l_1	l_2	8	8:48	9:48	l_1	l_2
2	7:18	8:28	l_1	l_2	9	9:04	10:14	l_1	l_2
3	7:35	8:45	l_1	l_2	10	9:18	10:28	l_1	l_2
4	7:48	8:58	l_1	l_2	11	9:35	10:45	l_1	l_2
5	8:04	9:14	l_1	l_2	12	9:48	10:58	l_1	l_2
6	8:18	9:28	l_1	l_2	13	10:03	11:13	l_1	l_2
7	8:35	9:45	l_1	l_2					

The results in this paper evolved out of the design and implementation of program UCOST (1) for the Urban Mass Transportation Administration (UMTA). Many of the procedures discussed here have been implemented or will be implemented within the various procedures contained in the Urban Transportation Planning System (UTPS) (2) or within future computer-based transportation planning systems to be implemented and distributed by UMTA. A more detailed description of these procedures can be found in Bodin and Dial (3).

All of the procedures described for the determination of vehicle requirements and line-by-line analysis for estimating crew requirements have been used in Dade County, Florida, for the design of the bus system that is to feed the proposed urban rail system. Those procedures allowed for the myriad of possible feeder bus systems to be reduced to a few by finding reliable capital and operating cost estimates. Some of the procedures have not been used in the field as yet: The histogram approach is included as part of program UCOST and the interactive procedures are currently under development. The RUCUS system (4) has been modified by several organizations and has been used with varying degrees of success in several cities.

BASIC STRUCTURE OF A TRANSIT SYSTEM

A transit system can be depicted by a set of transit lines that presents data for each line in one of two ways. The first way gives a timetable (headway sheet) for the system that shows, for each trip in the timetable, its line number, start time, end time (including layover), start location, and end location. This is the kind of data RUCUS (4, 5) requires as input. Preparation of the data in order to depict the transit system in this detailed manner is expensive.

The second way gives the length of time to cover any trip on the line, the time between runs on the line (called the headway for the line), and the start and end locations for each trip on the line. To take into account variable traffic patterns and demands for service, both the time to cover a trip and the headway for the line can be a function of time of day. The second way costs less to prepare but does not specify a timetable directly.

In long- and short-range planning, an actual timetable may not be necessary in order to perform the desired analysis. Moreover, Bodin and Rosenfield (1) showed that the determination of a well-designed timetable (in terms of passenger transfer times) from the line data specified in the second way is a challenging computational exercise. However, a daily timetable (which may not be well designed) can be quickly generated in the following manner. The first run of each line in a time period can be assumed to begin at the start time of the time period. Then, the other runs for the line in the time period are found by increasing the start and end times of the previous run for the line by the headway. The timetable generated in this manner may be unsatisfactory for operational planning since the lines are not synchronized, but this timetable may be adequate for long-range planning and for some short-range planning exercises.

The two ways of depicting a transit line can be illustrated as follows. A timetable for a line in a period is given in Table 1 (this will be referred to as timetable 1 in the remainder of this paper). The headway between adjacent trips in a timetable need not be the same; therefore, the start and end times for each trip in the timetable must be specified. Since many transit systems have several thousand trips, the preparation of the data (unless the headway for a line is constant) can be a significant undertaking.

In this paper we attempt to demonstrate what a planner can discern about crew and vehicle requirements when only headway information for each line is available. Furthermore, we attempt to show what additional information can be determined about crew and vehicle requirements when an actual timetable of trips is available. Finally, we assume that the layover time is a requirement of the system and is included in either the start or end time of the trip if a timetable is given or as part of the time to cover the trip if a timetable is not specified.

DETERMINATION OF VEHICLE REQUIREMENTS

Procedures for determining vehicle requirements range from a simple procedure that can be done manually to a complex optimization procedure that requires a computer. Although these procedures are not the only approaches available for determining vehicle requirements for transit systems, they illustrate a hierarchical approach to this problem and demonstrate the additional information gained by using a more complete (i.e., costly) approach.

Maximum Number of Vehicles: Line-by-Line Approach

The vehicle requirements for each line in a proposed transit system are estimated as follows:

$$\text{Vehicles for line } i \text{ in time period} = \lceil \text{Time to cover a trip} \div \text{headway of line } i \text{ in time period} \rceil \qquad (1)$$

where $\lceil x \rceil$ is the smallest integer greater than or equal to x. The number of vehicles to service the entire transit system in a time period is the sum over all the lines in the system of the number of vehicles needed to service each line as found in Equation 1. Thus, if the time to cover a trip is 70 min and the headway is 15 min, then an estimate of the number of vehicles needed to service the line in the time period is $\lceil 70/15 \rceil = 5$.

This quick procedure is useful for quick determination of a maximum number of vehicles (i.e., a capital requirements analysis). As such, this analysis need only be performed over the peak time periods. The vehicles required are the maximum of the vehicle requirements needed in each of these time periods.

If a vehicle is to service trips in both directions of a two-way line (or a trip on one line followed by a trip on the second line), then the time to cover a trip on the line is equal to the time to service a trip in each direction plus the time that the vehicle needs for turning around at each end of the line. For one-way lines, the time to cover a trip on the line is equal to the time to traverse the line in one direction plus the time to deadhead back to the beginning of the line plus the turnaround times.

Computation of the deadheading time for all pairs on terminal points can be an expensive enterprise. Therefore, an estimate of deadheading time (as a linear function of distance) may be appropriate for this procedure.

This procedure can determine feasible vehicle schedules if a vehicle is restricted to servicing only one line and the input data satisfy the requirements listed above. If the requirements are satisfied but the vehicle is allowed to deadhead between ends of the lines (i.e., service more than one line), then the above procedure is an upper bound on the estimation of vehicle requirements. If the deadheading and turnaround times are not known, then the procedure gives an estimate of vehicle requirements, but not necessarily an upper bound. In this situation, the procedure may underestimate vehicle requirements.

If the assumptions above are satisfied but deadheading between lines is allowed, the resulting upper-bound estimate of vehicle requirements may be a considerable overestimate of actual vehicle requirements. This procedure can be performed manually.

Example

Let line 1 have a headway of 15 min and a duration of 70 min and let line 2 have a headway of 15 min and a duration of 50 min. Assume that the end location of line 1 is the start location of line 2 and vice versa. Furthermore, assume that it takes 22 min to deadhead from the end locations of each line to its beginning location and assume that the turnaround time at the ends of the line is 4 min. The following estimate of vehicle requirements can be made.

Vehicles for lines 1 and 2 together $= \lceil(70 + 4 + 50 + 4) \div 15\rceil = 9$

Note that, in the case of a trip that has the same start and end locations, the deadhead time does not enter the computations.

Lower Bound on Vehicle Requirements: Histogram Approach

In the histogram approach, a timetable must be specified. In this timetable, the start and end times for each trip are known. Let a 1440-strata (= 60 min x 24 h) histogram be specified where stratum i corresponds to the ith minute of the day. If trip j starts at time k and ends at time e, then a vehicle is required for the k, k + 1, k + 2, ..., e - 1 minutes of the day. In this case, 1 is added to the values of strata k, k + 1, k + 2, ..., e - 1 in the histogram. The above procedure is repeated for all trips in the timetable. Let m_i be the number of vehicles required in the ith stratum (i.e., the ith minute of the day) and let $M = \max_i (m_i)$. Then M is a lower-bound estimate of the number of vehicles required.

M is a lower-bound estimate because M denotes the maximum number of vehicles required by the timetable but fails to consider any deadheading or dead time that may require additional vehicles in an operational schedule. When actually scheduling vehicles, it may be necessary to deadhead a vehicle over the stratum that designates the peak number of vehicles. Hence, this procedure gives a lower bound.

If the planner only wants to estimate the vehicle requirements, then this analysis need only be performed over the peak periods.

We have found that, in many cases, M is a surprisingly accurate estimate of actual vehicle requirements.

If all lines operate over the entire time period (i.e., no special trips needed over a small portion of the time period), then the results of this procedure are (for the most part) independent of the timetable used. In this case, a timetable may not be required to use this approach and only line data are used as specification of the transit system.

Example

Part of the histogram for timetable 1 is given below. The histogram oscillates between 4 and 5 until 10:12, when it begins to damp out. The peak number of vehicles estimated is 5.

Time Interval	No. of Vehicles
7:03-7:17	1
7:18-7:34	2
7:35-7:47	3
7:48-8:03	4
8:04-8:12	5
8:13-8:17	4
8:18-8:27	5
8:28-8:34	4

The line-by-line analysis, which generally gives an upper bound on vehicle requirements, may not be accurate, but the results can be found without having to use a computer program. The histogram approach generally gives a lower bound and is accurate, but it needs a simple computer program to derive the desired estimate. For capital cost estimation, both procedures can be used in a sketch-planning mode, and the histogram approach can be used in a short-range planning mode. The concurrent scheduler (described next) should be used in a short-range planning mode if both a capital cost analysis and operating cost estimate are needed.

Feasible Vehicle Schedules: Concurrent Scheduler

The concurrent scheduler is a straightforward heuristic that creates a feasible vehicle schedule (set of blocks) for a given timetable. The trips for all lines are merged together and are sorted by starting time from a specified beginning time of day. Although which beginning time of day to select for a 24-h timetable is not obvious, we have found that the results for the concurrent scheduler and the Dilworth chain decomposition procedure are not greatly affected by this beginning time of day, as long as the beginning time of the day is in an off-peak time period.

The concurrent scheduler operates as follows:

1. Orders the trips in the timetable by time of day; call this list of trips the sorted list;
2. Assigns trip 1 in the sorted list to vehicle 1 (i.e., block 1);
3. Assumes that the first k trips in the sorted list have formed m partial vehicle schedules (blocks). Then, it is possible to assign trip k + 1 in the sorted list to partial vehicle schedule n, n = 1, 2, ..., m if (a) E(n) = start time for trip k + 1 - end time for partial vehicle schedule n ≥ some minimum time as specified by the planner and (b) E(n) + safety factor ≥ time to deadhead from the end location of partial vehicle schedule n to the start location of trip k + 1;
4. If trip k + 1 can be assigned to more than one partial vehicle schedule, then the scheduler assigns the trip to the partial vehicle schedule that minimizes E(n), n = 1, 2, ..., m or to the first partial vehicle schedule found that satisfies the above conditions;
5. If trip k + 1 cannot be assigned to any partial vehicle schedule, then it creates a new partial vehicle

schedule m + 1 beginning with trip k + 1; and

6. Repeats steps 2-5 for all trips in the sorted list.

The above procedure gives vehicle schedules that are feasible but not necessarily optimal. An example is presented in Bodin and Dial (3) that illustrates this point. The Dilworth procedure (6) discussed in the next section determines a minimum number of vehicle schedules for a given timetable.

The procedure is very fast computationally because it has to pass only once through the sorted list of trips.

If only capital requirements are needed, then the concurrent scheduler need only be applied to the trips in each peak period and the maximum selected as the peak requirements.

Example

Let timetable 1 be specified in Table 1 (where l_1 is the start location for each trip in timetable 1 and l_2 is the end location). Furthermore, let the timetable for line 2 (called timetable 2) be specified in Table 2. The deadhead times are as follows: $d(l_1, l_2) = d(l_2, l_1) = 22$, $d(l_1, l_1) = d(l_2, l_2) = 0$. The turnaround time at the end of each trip is 4 min. The sorted timetable and the vehicle assignments of each trip, by using the concurrent scheduler, are given in Table 3. The number of vehicles required by the solution to the concurrent scheduler is 10. The estimated number of vehicles as found in the line-to-line analysis is 9. This number can be attained by the concurrent scheduler if the start and end time for each trip in timetable 2 is increased by 4 min. Therefore, trip 1 for timetable 2 would be the following:

Time		Location	
Start	End	Start	End
7:04	7:54	l_2	l_1

Note that the estimated number of vehicles that use the line-to-line analysis was made independent of the timetable used, whereas the results from the concurrent scheduler were based on a timetable.

Table 2. Timetable 2.

Trip	Time Start	Time End	Location Start	Location End	Trip	Time Start	Time End	Location Start	Location End
1	7:00	7:50	l_2	l_1	8	8:45	9:35	l_2	l_1
2	7:15	8:05	l_2	l_1	9	9:00	9:50	l_2	l_1
3	7:30	8:20	l_2	l_1	10	9:15	10:05	l_2	l_1
4	7:45	8:35	l_2	l_1	11	9:30	10:20	l_2	l_1
5	8:00	8:50	l_2	l_1	12	9:45	10:35	l_2	l_1
6	8:15	9:05	l_2	l_1	13	10:00	10:50	l_2	l_1
7	8:30	9:20	l_2	l_1					

Optimal Number of Vehicles

To derive feasible vehicle schedules (blocks) that minimize the number of vehicles needed, a procedure such as the Dilworth chain decomposition (6) must be used. The Dilworth chain decomposition finds the minimum number of chains needed to cover all the nodes of an acyclical directed network. Each chain corresponds to a vehicle schedule. The nodes in this network are the trips from the timetable, and the arc from node i to node j implies that it is feasible (vis-à-vis the conditions in step 3 of the concurrent scheduler) to service trip i and then trip j on a vehicle schedule. A description of the implementation of the Dilworth procedure for transit scheduling can be found in Bodin and Rosenfield (1).

The Dilworth procedure does not minimize deadhead requirements, and the solution from the concurrent scheduler (for the entire day) is a good starting solution to the Dilworth procedure. To our knowledge, there is no procedure available that can simultaneously minimize both vehicle requirements and deadhead distance. The vehicle scheduling procedure in the RUCUS computer system minimizes deadhead distance but not vehicle requirements, is much slower computationally, and only handles much smaller problems. The network in the RUCUS procedure is the same as the network in the Dilworth procedure except for the costs on the arcs of the network.

The minimization of vehicle requirements and then deadhead distance (given vehicle requirements) requires a two-step procedure. The first step performs the Dilworth procedure to minimize vehicle requirements. The second step uses the RUCUS vehicle scheduling procedure while fixing the number of vehicles to be allowed (as found in the Dilworth procedure). This is accomplished by fixing the lower and upper bounds on flow on the branch from the supersink to the supersource equal to the Dilworth solution and using the Dilworth solution as the starting solution from this minimum cost-flow problem.

To derive a solution that trades off between number of vehicles used and total deadheading requires that the RUCUS BLOCKS model be modified as follows. A relative weight is chosen to be associated with the number of vehicles; this weight reflects the value of a vehicle with respect to a deadheading unit (i.e., distance or time). This weight is used as the cost on the arc from the supersink to the supersource. The solution to the minimum cost network problem would then be the one that trades off vehicles with savings in deadheading. The cost of using this model would be essentially equal to that of using the present RUCUS model.

The three models described above give different answers to the same problem based on the objective the planner wishes to use. The planner must decide whether it is worth the investment in computer time to run either of the latter two models (the two-step model or the com-

Table 3. Sorted timetable.

Time Start	Time End	Location Start	Location End	Vehicle Assignment	Time Start	Time End	Location Start	Location End	Vehicle Assignment
7:00	7:50	l_2	l_1	1	8:35	9:45	l_1	l_2	5
7:03	8:13	l_1	l_2	2	8:45	9:35	l_2	l_1	4
7:15	8:05	l_2	l_1	3	8:48	9:58	l_1	l_2	7
7:18	8:28	l_1	l_2	4	9:00	9:50	l_2	l_1	6
7:30	8:20	l_2	l_1	5	9:04	10:14	l_1	l_2	9
7:35	8:45	l_1	l_2	6	9:15	10:05	l_2	l_1	8
7:45	8:35	l_2	l_1	7	9:18	10:28	l_1	l_2	10
7:48	8:58	l_1	l_2	8	9:30	10:20	l_2	l_1	1
8:00	8:50	l_2	l_1	9	9:35	10:45	l_1	l_2	2
8:04	9:14	l_1	l_2	1	9:45	10:30	l_2	l_1	3
8:15	9:05	l_2	l_1	10	9:48	10:58	l_1	l_2	4
8:18	9:28	l_1	l_2	3	10:00	10:50	l_2	l_1	5
8:30	9:20	l_2	l_1	2	10:03	11:14	l_1	l_2	6

bined RUCUS model) rather than the Dilworth procedure in order to determine the trade-off between vehicle requirements and deadhead distance.

DETERMINATION OF CREW REQUIREMENTS

In this section, we present procedures for determining crew requirements for a proposed transit system. The crew requirements problem is more complex than the vehicle requirements problem because a vehicle can operate the entire day without a break, but a crew has specific work rules that restrict the total amount of work that can be done during the day. Furthermore, the cost of a crew depends on the type of shift worked, the length of the shift, the time of day worked, overtime, and so forth.

In a simplified model, there are three basic crew workdays: full-time shifts, split shifts, and tripper shifts. A full-time shift is a complete workday for a crew with one embedded short break for lunch. A split shift is a complete workday for a crew with an embedded longer break of several hours that splits up the workday. A tripper shift is a part of a workday and has no scheduled breaks. Since transit systems generally have a morning and evening peak surrounded by lesser requirements during the off-peak hours, tripper shifts and split shifts usually exist to service the peak periods, and full-time shifts are scheduled to handle the nonpeak demands in both the peak and off-peak periods. The cost of a crew is a function of the type of shift, the time of day that the shift works (generally associated with the starting time of the shift), and the length of the shift (i.e., overtime).

The crew scheduling problem can be thought of as a very large set-covering or set-partitioning problem. A description of the set-covering and set-partitioning formulations of this problem can be found in Bodin and Dial (3).

The RUCUS implementation and the set-partitioning and set-covering formulations of the crew scheduling problem represent one-shot batch-optimization procedures for solving this problem. In a batch procedure, all parameters that guide the solution process are set prior to the computer run itself. The computer program then finds a solution to the problem based on the parameters set and the data. We believe that a batch-optimization algorithmic procedure for solving the crew scheduling problem is computationally prohibitive in most cases. Therefore, the development of heuristic procedures or man-machine interactive procedures for solving this problem appears necessary. The heuristic procedures were discussed in detail at a meeting on operator scheduling (Workshop on Automated Techniques for Scheduling of Vehicle Operators for Urban Public Transportation Services, April 27-29, 1975). Many heuristics exist for solving the crew scheduling problem, including a particularly effective one that adapts the RUCUS system (4, 5). Because of the diverse nature of these heuristics, they will not be discussed in any detail in this paper.

It is possible, however, to develop procedures for simply estimating or bounding total crew requirements that do not depend on costly crew scheduling heuristics. Such estimates are invaluable for cost-estimation purposes and provide targets at which schedulers who use run cutting can aim. Procedures that can play a central role in the planning and operation of transit systems are discussed below.

Estimation of Crew Requirements Without a Timetable: Line-by-Line Analysis

In time period p, let L(i,p) be the duration of a trip on line i, including layover in minutes, and let n(i,p) be the number of trips on line i in the period. Then CR(i), which is the estimated number of crews who work a full shift on trip i, is found as follows:

$$Cr(i) = \left\lceil \sum_p L(i,p)\, n(i,p)/E \right\rceil \quad (2)$$

where E is the number of minutes in an effective workday for a full shift. E is discussed below. The number of crews who work a full shift (DR) is found as follows:

$$DR = \sum_i CR(i) \quad (3)$$

This approach can be performed manually.

If a line is to operate over the entire duration of period p, the duration of period p is D(p) minutes, and the headway of line i is H(i) minutes, then

$$n(i,p) = \lceil D(p)/H(i) \rceil \quad (4)$$

where $\lceil x \rceil$ is the smallest integer greater than x. Thus, if x is an integer $\lceil x \rceil = x + 1$.

E, the number of minutes in an effective workday, needs some clarification. Let T be the duration of the workshift for a full-time crew. Let each crew spend, on the average, t minutes in nonrevenue activities such as deadheading to and from the garage, lunch break, or time between runs. Then E = T - t is the number of minutes in a day that a crew spends on revenue activities (i.e., actually serving passengers). The revenue activities are the runs specified in the line schedule or the timetable.

It is difficult to estimate E without having actual crew schedules (i.e., it is often difficult to discover the time to and from the garages or the time between trips). To find E requires the use of previous experience with the transit system. As a rule of thumb, we have found that an estimate of E between 6 and 6.5 h provides reasonable estimates of crew size for a traditional bus operation.

Given a more reliable estimate of E and L(i,p), we can better estimate CR. Thus, any preliminary scheduling that can be performed is useful. For example, let line 1 go from node A to node B and let line 2 go from node B to node A and let both lines have the same headway. If the following crew schedule is to be run: line 1, line 2, line 1, and so on, then L(i,p) can be redefined as the time to complete the round trip and start out on the next available trip on line 1. In this case, dead-time information is embedded within the computation of required work time. Thus, in the expression E = T - t, t equals the time to deadhead to and from the garage to the specified start and end points of the lines plus the time for lunch. Since this definition of t gets rid of much of the variability attributable to crew scheduling, experience has shown that this estimate of E gives a more reliable estimate of n(i,p) than the estimate of E described previously.

In many cases, transit planning is performed one time period at a time. To discover an estimate of operating cost for a time period, we need the estimated equivalent number of crews who work a full shift in time period p, which we call CRP(p). CRP(p) is found as follows:

$$CRP(p) = \lceil \sum_i L(i, p)\, n(i, p)/E \rceil \quad (5)$$

Since these time periods can be short in duration, no attempt is made to break down CRP(p) into a line-by-line analysis.

Let G be the cost per day of a full-time crew. Then GTOT [or GTOT(p)], the estimated operating cost attributable to the crews (or the estimated operating cost attributable to the crews in time period p), is given by

$$GTOT = G{*}CR \text{ [and } GTOT(p) = G{*}CRP(p)] \quad (6)$$

GTOT is a simple estimation procedure; it disregards pay differentials as a function of shift type and time of day, but it can give reasonable answers to crew size requirements with a minimal investment in data preparation and computer implementation.

Example

Suppose that timetables 1 and 2 are to operate from 7:00 a.m. to 7:00 p.m. The headways on both lines are 15 min, each trip for timetable 1 has a duration of 70 min, and each trip for timetable 2 has a duration of 50 min. Assume that it takes 30 min to get from the garage to either l_1 or l_2 and each driver is to get a 45-min lunch hour. If both lines can be serviced by the same crews, then the length of a trip = 70 + 4 + 50 + 4 = 128 min. Therefore, $CR(1 + 2, 1) = \lceil(128)(49)/375\rceil = 15$. If the cost of a crew is \$50/day, then the crew cost estimate in this example is (\$50) [CR (1 + 2, 1)].

The line-by-line analysis gives an estimate of crew requirements (assuming that each crew is full time) that may not be accurate, but the results can be found without a computer program. As such, it should be used in a long-range planning environment. The histogram procedure described in the next section gives a more accurate estimate of crew requirements, uses to some extent differing shift types in building its model, but requires a computer.

Histogram Procedure

The procedure in the previous section gives an estimate of crew size requirements assuming that crews only work a full-time shift. The procedure in this section gives the following:

1. Estimate of the number of crews needed by shift type (full-time shift, split shift, and tripper);
2. Estimate of the number of crews needed by time of day; and
3. Estimate of the total crew costs, taking into account pay differentials.

This procedure does not give actual crew schedules.

Input to this procedure is the set of aggregated trips or vehicle schedules. Each aggregated trip represents a collection of trips or blocks that must be serviced by a crew and vehicle. The concurrent scheduler or the Dilworth procedure can be used to create the set of aggregated trips. The aggregated trips are essential to avoid double counting the required number of crews and overestimating the number of crews required to service the transit system.

The first step in this procedure is to form a histogram (called the demand histogram) of the number of crews required to cover all the trips of the day. Input is the set of aggregated trips or blocks. To construct this histogram, the time of day is broken down into time intervals, where it is assumed that any blocks that fall into any part of these small time intervals require a crew to service them for the entire time interval. For example, a block that starts at 7:09 generates a crew requirement from 7:00 to 7:10 if a 10-min time interval is used in defining the histogram. Experimental evidence indicates that a 10-min time interval derives accurate estimates of crew size.

The crew estimation is based on allowable shift segments that the planner specifies. A shift segment consists of a consecutive number of work hours that a crew is to work. A shift segment specification is designated by the crew cost (including pay differential and overtime), first permissible time of day when crews can report to work on this shift segment, and last time of day when crews can report to work on this shift segment. The shift segment specifications form the alternatives on which the crew estimation is to be based. Split shifts are combinations of two shorter shift segments. Thus, if a crew is to work a split shift from 9:00 a.m. to 1:00 p.m. and from 5:00 p.m. to 9:00 p.m., then the crew estimation component assumes that the crew works two shift segments, each of 4-h length. One shift segment is from 9:00 a.m. to 1:00 p.m., and the second shift is from 5:00 p.m. to 9:00 p.m.

From the demand histogram and the shift segment specifications, a network is created. The out-of-kilter algorithm (6) is then employed to determine the number of crews of each type of shift segment that are needed to cover the demand histogram. A detailed description of this procedure along with an example that illustrates the procedure can be found in Bodin and Rosenfield (1).

This procedure gives an estimate of the number of crews required by time of day and an estimate of the number of crews to be assigned by time of day. Hence, the difference is how many crews are present but not assigned to a particular activity. The times of day when this difference is positive may be times of day when runs can be added to the timetable without requiring an additional crew. Hence, these runs can be serviced at no additional crew cost. This characteristic of the solution to this problem is useful in attempting to design a transit system.

The solution of the minimum cost-flow problem requires no more computer time than the generation of the network. Therefore, this procedure can derive an estimate of crew requirements in the same amount of computer time [about 10 s of central processing unit (CPU) time on an IBM 370/168] independent of the line schedule used.

A slightly more difficult problem is to allow the planner to place bounds on the number or percentage of crews of various types. This problem cannot be solved with a network-flow algorithm but can be solved with a moderate-sized linear problem (150 rows, 1500 columns). This latter model is a planned improvement to the UCOST software.

Optimal Crew Scheduling

We do not feel, given the current state of computing machinery, that the optimal crew scheduling problem, in general, can be solved by use of a batch algorithm. We do feel, however, that the problem would be solvable by using a man-machine interactive procedure. Such a procedure would begin with a feasible schedule composed of two subschedules. One subschedule would consist of those fixed crew schedules that cannot be altered and the other would consist of free schedules that could be changed. Also, certain partial schedules could be fixed as if they were a run and then joined as a block on a full-day schedule. The procedure would then improve the given schedule by manipulating the set of free schedules.

The procedure would then write the results of this analysis to a structured data base. The planner could then query the data base to find instances of crew schedules that were unacceptable. The planner would then alter the fixed and free subschedules to reflect complaints and reexecute the algorithm.

Mathematically this problem can be formulated analagous to the Dilworth chain decomposition procedure except that the length of each chain is restricted to be within certain bounds (to reflect length of shift requirements). The fixing of schedules corresponds to the forcing of a flow of one over certain branches in the network. The prohibition of the joining of two runs on a crew schedule corresponds to the fixing of a flow of zero on the appropriate branch. The Dilworth chain decomposition with length of chain restriction can then be solved over this smaller network.

Output from this analysis would then be a set of tripper shifts or half-day schedules. A heuristic or exact 1-match procedure (7) can then be employed to join these half shifts into full-day work shifts. The 1-match procedure can be designed to take into account secondary considerations, such as the allowable percentages of tripper shifts, fixed number of full-time shifts, and allowable overtime. If the solution of the 1-match procedure is not acceptable to the planner, then the half-day schedules or tripper-shift schedules can be changed to reflect the planner's complaints and the 1-match solution used. We are currently carrying out a project to design and implement the above procedure.

ADDITIONAL COMMENTS

We have presented a variety of manual and computer-based procedures for estimating, bounding, or determining exact vehicle and crew requirements for transit systems. We have also attempted to illustrate the strengths and weaknesses of these procedures, their assumptions, and their possible utility. A problem that planners encounter in analyzing their problems is (a) they expect too much from some models or (b) they overbuild and complicate their model in attempting to get their desired results. In the first case, their results are superficial and incomplete; in the second case, their results are extremely costly to derive. We have demonstrated that, if the goals of the planner are modest, simple procedures will derive the required answers. However, if a more detailed result is needed, a much more complex and costly model needs to be constructed. Therefore, the planner has to decide whether the more detailed result is needed and whether he or she is willing to pay the price (in terms of data collection, computer programming, and computer time) to find the results. Finally, we hope that the results in this paper demonstrate that estimates (and not necessarily bounds) of the solution are of use in a planning environment.

ACKNOWLEDGMENT

Research into the results presented in this paper was sponsored by an UMTA contract with the University of Maryland.

REFERENCES

1. L. Bodin and D. Rosenfield. Estimation of the Operating Cost of Mass Transit Systems. College of Urban and Policy Sciences, State Univ. of New York, Stony Brook, Rept. WAHCUPS-UMTA-1-76, Sept. 1976.
2. UTPS Reference Manual. Urban Mass Transportation Administration, 1975.
3. L. Bodin and R. Dial. Determining Crew and Vehicle Requirements for Mass Transit Systems. College of Business and Management, Univ. of Maryland, College Park, 1979.
4. Vehicle Scheduling and Driver Run Cutting (RUCUS) Package User Documentation. Mitre Corporation, McLean, VA, MTR-6678, May 1974.
5. MTA Run Cutting Requirements Analysis. De Leuw, Cather, and Company, Washington, DC, RUCUS Technical Rept., No. 1, Nov. 1976.
6. L. R. Ford, Jr., and D. R. Fulkerson. Flows in Networks. Princeton University Press, Princeton, NJ, 1962.
7. N. Christofides. Graph Theory: An Algorithmic Approach. Academic Press, New York, 1975.

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