Railroad Car Distribution Performance: Conceptual Framework and Underlying Mathematical Relationships

CARL D. MARTLAND AND MARK McCORD

Car distribution performance, which can be measured in terms of car days, car miles, or other measures, depends on many factors. The mathematical relationships underlying car distribution performance are presented, including equations for analyzing situations in which cars are in surplus or shortage situations. Improving performance requires a coordinated effort involving many organizations and railroad departments, not simply changes in the way that car distribution decisions are made. A framework for structuring this coordination is given.

The railroad freight car distribution process involves moving empty cars from an unloading point to the next loading point. Car distributors assign specific empty cars to specific customers and issue orders for the operating department to move the cars to their designated customers. Car distribution performance is normally measured in terms of resource consumption, productivity, and service levels, as outlined below:

1. Resource consumption
   a. Empty-car miles by type of equipment, region, or time period
   b. Empty-car days by type of equipment, region, or time period
   c. Cost of empty movements and storage
2. Productivity or efficiency of fleet of cars
   a. Ratio of empty-car miles to loaded or total car miles
   b. Empty-car miles per load originated
   c. Empty-car miles per empty-car day
   d. Empty-car days per load originated or per load handled (including loads handled by a railroad that originated on another railroad)
3. Service provided to group of customers by fleet of cars
   a. Unfilled cars orders (i.e., the number of requests for empty cars that could not be filled by customers)
   b. Number of cars rejected by customers as unsuitable for loading

Financial measures can be developed by attributing costs to these measures. A central thesis of this paper is that car distribution performance cannot be solely attributed to the decisions of car distributors. Indeed, car distribution performance is intimately related to other car management functions (especially fleet sizing) as well as to marketing and operating practices and to institutional relationships among railroads.

Car distribution is an extremely complex activity. Any measures of car distribution performance, therefore, must be regarded only as indicators of the general performance of a complex system. To improve car distribution performance, it is necessary to understand not just the performance measures but also the relationships among fleet size, the demand for freight cars, and the car distribution process. It is especially important to understand that the size of the fleet can be a dominant factor in car distribution performance.

The size of the desirable fleet can be determined for any traffic projection by assuming efficient utilization of the cars in the fleet. The desirable fleet balances the costs of overutilization and underutilization. When the fleet is overutilized, profitability suffers because some loads are not handled in the ideal car, and others cannot be handled at all. Adding cars to the fleet would reduce this problem and increase profitability. When the fleet is underutilized, either many cars are idle or improvements in use of the existing fleet would offer a cheaper means of expanding capacity than would the purchase of new cars.

In this paper, a framework is proposed that places the components of railroad car distribution in a unified environment. The need for such a framework became evident during an investigation and evaluation of a plethora of proposals for improving car distribution performance. [More than 60 alternatives for improving car distribution performance were identified and evaluated (1). A more complete discussion of car distribution, including case studies of practices on three railroads, may be found elsewhere (2).] The framework was also useful for assessing the increasing number of modeling efforts that attempt to optimize specific (and sometimes not explicitly defined) portions of car distribution performance. By describing the total environment, it is possible to describe proposed changes in common terms and to determine which levers the increasingly sophisticated analytical tools are pulling.

The unifying framework is provided by portraying the car distribution problem, in its most general sense, as a system control problem. Then, by presenting the underlying mathematical identities that shape the system, one can see, without further complicating assumptions, the most basic limits on the problem--limits that any action must acknowledge. The framework is intentionally general. No black-box model is developed that estimates, for a set of input parameters, a set of performance measures. Rather, it is shown how to link specific proposals for improving or modeling performance to the overall environment of freight car management.

CAR DISTRIBUTION AS A CONTROL PROBLEM

Car distribution as used here involves moving empty cars from one region and status to another region and status. The possible statuses of empty cars include moving to an assigned distribution point, awaiting distribution, awaiting repairs, and being stored. The control-problem (CP) model is shown in Figure 1 and can be stated as follows:

$$CP = [T, S, (A,D), Y, f(a,p), g]$$

where

- T = number of days in period and t_k is kth day;
- S_k = vector describing location and status of all cars at time k;
- A_k = particular decision applied to system on day k (which cars should be assigned to what status at what location at time k + 1);

Figure 1. Car distribution as a control problem.

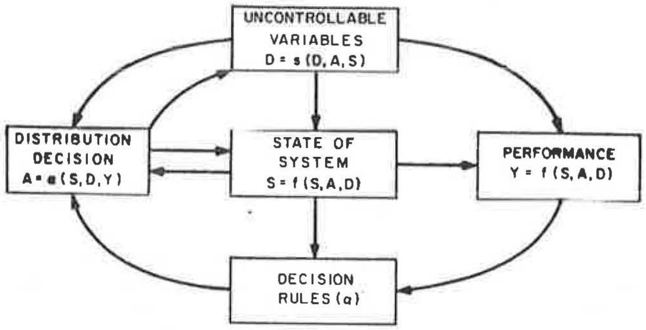

D_k = input to system at time k concerning demand for cars, institutional constraints, weather, bad orders, and other variables not under direct control of person making decision A_k;

Y_k = vector of performance measures for day k that could include empty-car miles, empty-car days, unfilled orders, and other measures;

f_k = function that determines state of system at k + 1 given state at k, distribution decision A_k, and uncontrollable input D_k;

a_k = rule or policy on day k that determines decision A_k as function of S_k and D_k (this may be formally or informally stated by railroad);

p = function that determines D as function of state of system S and decisions A (for example, how will demand for cars vary with state of system and decisions that allocate empties to shippers); and

g = function that relates performance to state of system S, decisions A, and uncontrollable input D.

The objective of the car distribution problem is to optimize performance Y over the time period T. This problem is complicated by the large numbers of freight cars and statuses of freight cars (i.e., the complexity of S); the numerous options available to decision makers; the uncertainty inherent in f, a, p, and g; and the intricacy of the market for freight transportation and the rules governing car distribution, both of which are included in D. We clarify this model by looking at its various components.

As the basis for any control system, one must be able to identify the state of the system at any time (S_k). In car distribution, the state of the system is given by the location of each car and its status. This is the same information that is kept by most railroads as part of their computer information system. To simplify matters, the state of the system need not consider every car individually but can refer to the number of cars in each status in each region.

The next major part of the control process is the statement of objectives that describe the desired state of the system. The goal of the car distribution function is to improve car distribution performance (Y), which can be a single performance measure or a vector of measures. Some possible measures include total empty-car miles, empty-car days, and various ratios relating empty-car miles and days to loads originated or terminated. Other measures relate to the availability of cars when desired by shippers and the quality of cars. Each of these performance measures can be obtained by analysis of computerized car movement records or, in principle, from network models.

Control variables are the means by which an organization seeks to change the state of the system to be more in line with the objectives. In car distribution, these control variables can be grouped into a number of categories:

1. Car distribution decisions by which empty cars are routed from unloading points to loading points,
2. Operating decisions by which cars actually move along these routes,
3. Marketing decisions by which a particular mix of traffic is solicited,
4. Fleet acquisition decisions by which railroads and other organizations expand and replace their car fleets, and
5. Investment decisions by which railroads and customers build, replace, or downgrade fixed facilities.

Note that these variables are controlled by different groups within the railroads and in some instances by other organizations. Only the first group of these variables is controllable by car distributors. Therefore, if one wishes to model the problem from their perspective, these decisions are included in A, and the rest must be modeled in the environment as uncontrollable variables. From the perspective of car management (as opposed to car distribution), however, all of the potential control variables must be considered.

The environment consists of factors that, for one reason or another, are not controllable. The cyclical patterns of business activity, the location of raw materials and markets, the maximum speed of rail freight trains, and the weather are certainly uncontrollable within the context of car distribution. The rules and regulations governing the movement of cars, the operating practices of railroads and their customers, and the physical facilities of the rail system are relatively uncontrollable except over a period of several years or longer.

It should be emphasized that this is but a conceptual model of the car distribution problem. It would be difficult to solve this problem with any degree of generality. There have been recent attempts to optimize parts of this car distribution problem, however.

Turnquist and Jordan (3) looked at a limited system over a short time period (T). They acknowledge uncertainty in the functions that generate future supplies and demands, part of D_k, but assume that the mean and variance of these inputs are known for each day. They also allow for uncertainty in the state transition function due to uncertain travel times between yards. A subset of the possible performance measures is selected and put into a single dimension by combining revenue from filling orders with costs attributed to holding unused cars, failing to fill orders, and repositioning empty cars. The uncertainty is factored out by considering the expected value of this financial performance distribution. The output is a set of distribution decisions (A_k) that maximizes the expected performance measure.

Turnquist and Jordan have made headway in showing the effect, measured by their specific definition of Y, of uncertainty in the state transition function (f) and some of the input (D). They did not test alternative decision rules of the type used by a railroad. To address this problem, Mendiratta (4) and Mendiratta and Turnquist (5) separated the system-level decisions concerning empty-car movements from the terminal-level decisions concerning empty-car inventories. They conclude that their model can be used as a policy evaluation tool by railroad central management and an operational tool for the daily distribution of empty cars by terminal personnel.

It should not be forgotten that these models, useful in highlighting the effects of certain parameters of the car distribution environment, assume that most of the parameters are held fixed. When car distribution performance is viewed from the broad perspective of Figure 1, it is evident that the above studies investigate only portions of the first of the five major approaches to improving car distribution performance:

1. Try to improve the car distribution function itself by establishing better policies (a) and monitoring car distribution decisions (A) to make sure they are consistent with these policies.
2. Improve the information systems so that decision makers have better data on the state (S) and performance (Y) of the system and can learn more about the transformation functions (f and g).
3. Modify the institutional framework, traffic patterns, and other factors represented by D and p. Although the environment may be outside the control of car distributors, other railroad officials and other organizations can make changes.
4. Change the composition or ownership of the fleet, which will influence S and therefore the decisions made and the resulting performance.
5. Improve the technology or operating policy of the rail system, which would change the transformation functions (f and g).

Controlling car distribution, therefore, is a complicated problem involving coordination among various groups within each railroad as well as among railroads, shippers, and other organizations. Underlying all of the above approaches and overriding any type of coordination, however, are basic relationships among fleet size, traffic volume, traffic mix, and physical utilization. The relationships, which are presented next, are important because they are based on identities concerning the car cycle that both determine and limit the possible performance effects of any change in the car distribution environment.

OVERVIEW OF CAR CYCLE

The car cycle is the total time between successive loadings of a particular car, i.e., a period of time that begins when the car is placed for loading and ends when it is next placed for loading. If we define a set of mutually exclusive statuses that cover all possible situations, the car cycle equals the sum of the time spent in each status between two successive loadings. [The first comprehensive, published analysis of freight car cycles was by Reebie Associates (6).]

One can divide this car cycle into components at different levels of detail, but it is instructive to begin at a fairly coarse level and to subdivide only those components that call for more detailed analysis. The first level (Figure 2) divides the car cycle time into six major components:

$$\begin{aligned}\text{Car cycle time} = {} & \text{loading time} + \text{loaded storage time}\\ & + \text{loaded transit time} + \text{unloading time}\\ & + \text{empty-transit time} + \text{empty-storage time}\end{aligned} \quad (1)$$

All of these components can be divided further in a manner that depends on what is to be analyzed. Because the average of a sum is equal to the sum of the averages, we can obtain the next component:

Figure 2. Basic components of car cycle.

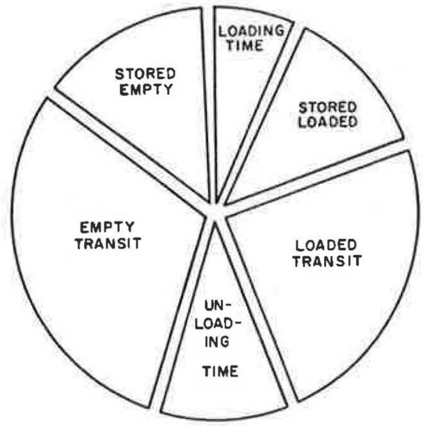

$$\begin{aligned}\text{Average car cycle time} = {} & \text{average time in loading}\\ & + \text{average time in loaded storage}\\ & + \text{average time in loaded transit}\\ & + \text{average time in unloading}\\ & + \text{average time in empty transit}\\ & + \text{average time in empty storage}\end{aligned} \quad (2)$$

But for any period of time, the average car cycle time can also be calculated as follows:

$$\begin{aligned}\text{Average car cycle time} = {} & (\text{car days available in period})/\\ & (\text{number of cycles in period})\end{aligned} \quad (3)$$

By equating Equations 2 and 3 and writing the average times in Equation 2 as the number of car days in each component divided by the number of cycles in the period (equal to the number of loads handled), we obtain

$$\begin{aligned}\text{Car days available} = {} & \text{car days in loading} + \text{car days in loaded storage}\\ & + \text{car days in loaded transit} + \text{car days in unloading}\\ & + \text{car days in empty storage}\\ & + \text{car days in empty transit}\end{aligned} \quad (4)$$

Equation 4 suggests an additional approach to studying car distribution. The car days available during any period are determined by the fleet size and the length of the time period:

$$\text{Car days available} = (\text{fleet size}) \times (\text{length of time}) \quad (5)$$

For this reason, the car days available is a useful accounting measure.

The basic unit of this accounting framework is the elapsed time spent by a single car in a particular status. Because a car is always in one and only one status, the summation of these basic units will be the total car days available for the fleet under consideration (Equation 5). We can aggregate these units in many ways, e.g., by car type, by status (as in Equation 4), by time period, or by cycle. However we choose to do this, the result is still determined by Equation 5. The interrelated measures for car days available, cars handled, and average cycle times will be shown to be useful in analyzing the car distribution problem.

STANDARDS FOR CYCLE-TIME COMPONENTS

If standard times for each component can be developed, they can be combined to find a standard time for the car cycle. Such a standard would be directly relevant for fleet management because, in combination with demand projections, it would provide an estimate of car days required in the future. In this section it is shown how standards can be derived from both theory and empirical evidence. The intent is not to show how to define particular standards but to emphasize that standards for each cycle component can be combined to obtain a standard for the entire cycle. In addition, the discussion assumes a basic familiarity with the use of standards in railroad management control systems.

In all of the following functional relations, only the main variables are identified. The parts of the cycle needed to carry a load are the loading, loaded transit, and unloading components. The loading and unloading portions depend basically on the number and types of loads and the loading and unloading procedures used. This can be expressed as follows:

$$\text{Average loading time} = f\,(\text{number of loads, type of loads, loading procedure}) \tag{6}$$

For example, if the average loading time remains constant, the total time will be as follows:

$$\text{Loading time} = (\text{average loading time}) \times (\text{number of loads}) \tag{7}$$

A similar equation can be developed for unloading time. [The Boston and Maine Corporation uses such a standard in its weekly operating and service plan performance report (7, p. D-19).] The average loaded transit time would depend on the network, the traffic mix, and the operating plan. Models that have been developed to determine standard transit times include, for example, the Massachusetts Institute of Technology (MIT) Service Planning Model (8). Basically, the total number of car days in loaded transit for a period of analysis could be considered a weighted average of the standard trip times thus calculated.

Even though the empty and storage components of the cycle are unnecessary to carry the load, some such time will normally be required. Once the car has been unloaded, it must either await loading (empty storage) at the same location or be moved (empty transit) to another point for reloading. Also, variations in demand or in fleet size will cause periodic surpluses of equipment, which leads to empty time. Finally, customers will not always be able (or desire) to unload a car precisely when it arrives, which leads to loaded storage.

Although empty time is not necessary to carry a load, it is nonetheless inevitable. How much empty time is reasonable is a difficult question because of the many alternatives for moving empty cars to reloading points and because of the variability in demand for freight cars.

By identifying the various causes of the empty time and estimating how much empty time each cause implies in the car cycle, the levers that must be adjusted to reduce the empty portion of the cycle can be identified. It is proposed here that the reasons for empty time can be classified into four broad categories. Despite some overlap among the categories, they are distinct enough to present an interesting classification.

One of the basic causes of empty time is that the fleet is at times simply too large for the traffic. Because freight cars last many years, the fleet cannot be quickly reduced if demand slackens. When total car days available is nearly constant, the average cycle time increases as the number of loads declines (see Equation 3). If the average time for the customer and loaded components remains constant, empty time must increase. This effect of fleet size is evident in times of slack activity and will be discussed in the section on surpluses below. In this case it is not the quantity of empty time but its distribution between transit and storage that is important. Because demand is low, more cars will spend more time empty, but they do not necessarily have to be moving (9).

In times of heavy business (see section on shortages below), the quantity of empty time is more important because it can be the cause of car shortages and lost business (10). Empty time persists despite the demand for cars because of spatial and temporal imbalances in loading and unloading patterns, institutional restrictions on the use of empty cars, and operating policies on an individual railroad.

Although the total number of cars loaded equals the total number of cars unloaded, the demand for loads is not uniform in time or in space. From a spatial perspective, one region may be a net originator of loads, whereas another is a net terminator of loads. In this case of spatial imbalances, empty cars will have to be sent from the terminating region to the originating region. From a temporal perspective, loads may originate only some time after the car to carry the load has been unloaded or arrived empty. In this case of temporal imbalance, empty cars will be stored in a yard. Various techniques have been proposed for routing empty cars to minimize the costs of spatial and temporal imbalance. Philip identified 16 car distribution planning models, 7 of which were actually implemented and used. Nevertheless, he concluded (2, p. 34): "Models to support car distribution decisionmaking have not found widespread application in the industry, despite the numerous attempts to describe how such models might be useful, and several attempts to do so."

Over and above the inevitable empty time caused by the spatial and temporal distribution of the demand is the empty time caused by institutional regulations (11). Regulations may prohibit or restrict the loading of cars in certain areas and encourage empty cross-haul, which can be defined as the simultaneous movement of empties in both directions between two regions.

Even in the absence of institutional regulations, there would be more empty time than absolutely necessary due to spatial and temporal imbalances because of the car distribution practices of individual railroads. [The Freight Car Utilization Program (FCUP), administered by the Association of American Railroads (AAR) and supported by FRA, has published numerous reports addressing the management of empty-car distribution. For an annotated bibliography, see AAR Report R-453 (12).] These practices may increase empty time for a variety of reasons. Insufficient information systems would force decisions to be made without valuable data. Poor management structure could preclude efficient operations. A lack of training of distribution personnel could be another reason. Another cause related to distribution practices is that of competing objectives. Empty time must be balanced with the costs of managing cars, repositioning cars, and failing to provide cars when desired by shippers.

In summary, the causes for empty time can be grouped into four broad categories:

1. Excessive fleet size,
2. Spatial and temporal imbalances in demand,
3. Institutional restrictions, and
4. Operating policies.

Given these causes for empty time, a standard for empty time can, at least in theory, be estimated:

$$\text{Empty transit standard} = f\,(\text{traffic volume, fleet size, spatial imbalance, temporal imbalances, institutional regulations, operating policy}) \quad (8)$$

In short, standards can be developed for each component of the car cycle. By summing the average standards for each component, one can define a standard for the average cycle time:

$$C^* = \text{standard cycle time} = \Sigma_{\text{comp}}\,(\text{standard time, component i}) \quad (9)$$

where comp stands for components of car cycle. The standard cycle time can be used in various ways. The next section shows how C* can be used to study fleet sizing issues with a number of equations developed for understanding car shortages and surpluses. C* also provides a link to the control theory and network models described above. Finally, C* provides a link among the analysis of car distribution decisions, operating plans, traffic flows, and fleet sizing, all of which affect empty time and empty mileage.

FLEET SIZING, SURPLUSES, AND SHORTAGES

The object of this section is to sketch the relations among the standard cycle time C*, the actual cycle C, investment decisions, the time period of analysis, and marketing practices. Investment decisions affect the fleet size. The time period of the analysis influences the measured imbalances in traffic flow by smoothing out or accentuating the random and cyclical variations in demand. Marketing practices and the general business environment affect the number of loads carried.

For simplicity, the following notation will be used:

- F = fleet size, assumed to be constant over period of analysis;
- T = number of days in period of analysis;
- C = actual average car cycle time (i.e., the cycle time realized during the period of analysis);
- C* = standard average car cycle time (i.e., the cycle time calculated from Equation 9);
- L_o = number of loads carried during the period; and
- L_d = number of loads demanded during the period, which may exceed L_o.

From the definition of car days available in a period, we have

$$\text{Car days available} = (F)(T) \quad (10)$$

Based on the standards, the number of days needed to fill all of the orders in the period would simply be the product of the standard average car cycle and the number of loads demanded in the period:

$$\text{Car days required} = (C^*)(L_d) \quad (11)$$

The differences between the car days available and the car days required will be defined to be the surplus car days for the period. The extent of such a surplus is determined by the relationships defined above:

$$\begin{aligned}\text{Surplus car days} &= (\text{car days available}) - (\text{car days required})\\ &= [(F)(T)] - [(C^*)(L_d)]\end{aligned} \quad (12)$$

If this difference is greater than zero, the period will be called one of surplus. If the difference is negative, a period of shortage ensues. These two cases will be investigated separately.

A surplus represents idle capacity. Some cars, which require investment and maintenance and therefore accrue costs, cannot be used to move shipments that generate revenue. As a result, the actual cycle time C rises above the standard cycle time C*:

$$\begin{aligned}C - C^* &= [(F)(T)]/[(L_d)] - C^*\\ &= \text{surplus car days}/L_d\end{aligned} \quad (13)$$

This increase is required by the definitions of C, C*, and the assumption of a fixed fleet, which is reasonable for periods of analysis much shorter than the average life of a freight car. During surplus conditions, therefore, the average cycle time will remain higher than the standard until the fleet (F) can be reduced to a more desirable size or until traffic (L_d) rises to the expected level. From the perspective of the entire fleet, the increase in the cycle time is, in short, inevitable when the fleet is too large for the traffic.

It is useful to consider where the extra time shows up in the car cycle. One or more of the six major components of the car cycle (Equation 1) must increase. If customer practices and traffic mix are not affected by the surplus, we would expect the standards for loading time, unloading time, loaded transit, and loaded storage time to remain the same. The extra time would therefore be absorbed either in the empty storage or the empty transit components of the cycle. If the fleet were controlled by a single organization, the need to store a portion of the fleet would be evident. The oldest or least-reliable cars would be obvious candidates for storage (13), but it might also be desirable to slow down the repair of cars or to undertake a rebuilding program.

The number of cars to store or to take out of service in other ways (F_s) can be determined from Equations 5 and 12:

$$(\text{Average number of cars stored})\,(T) = (F_s)(T) = (\text{surplus car days})$$

$$\begin{aligned}F_s &= [(F)(T)] - \{[(C^*)(L_d)]/(T)\}\\ &= F - (C^*L_d/T)\end{aligned} \quad (14)$$

$$F_s = F - F^* \quad (15)$$

where F* is the fleet size that would make the surplus equal zero. Because the number of cars stored could vary during the period because of variations in demand, Equations 14 and 15 deal with the average number of cars stored.

If cars are not stored during a surplus, other components of the cycle must increase. To reduce expenditures, the railroads might reduce the number of trains operated as well as the number of yard-crew assignments. This could increase the standards for both the loaded and empty transit portions of the cycle, which reflects the fact that the opportunity cost of the marginal car day has dropped substantially to zero. Railroads might also offer customers greater leeway in the use of cars for temporary storage. If the surplus continues even after these measures have been taken, surpluses of empty cars will become increasingly evident at numerous locations. To what extent these cars sit in yards and to what extent they shuttle back and forth among yards depends greatly on the industry rules and procedures governing the use of foreign cars (i.e., cars belonging to a railroad other than the one using them) and the management objectives and practices of the railroad where the car is unloaded. A single-fleet manager would reduce unnecessary empty movements to avoid the associated costs for fuel, crews, and maintenance. A number of interrelated railroads, however, might well attempt to shift the burden of the surplus to other carriers; such suboptimal behavior could easily increase the cost of car distribution despite the inevitability of the empty time that the railroads individually seek to avoid.

Shortages can also be investigated by using the basic equality of Equation 12. When there are fewer car days available than required to handle the loads demanded, a shortage of capacity exists, and all of the loads demanded cannot be carried in the period of analysis. The amount of shortage is shown by the following:

$$\text{Shortage of car days} = [(C^*)(L_d)] - [(F)(T)] \quad (16)$$

As with surpluses, the extent of the shortage is defined in terms of the standard rather than the average cycle. Equation 16 defines the relative impact that innovations affecting the variables can have on the extent of the shortage. Whereas the costs of surpluses are those of idle capacity, the costs of shortages are unfilled orders and delays to shippers.

When there is a shortage, railroads either change their operating practices to reduce components of cycle time below their standards or are unable to provide cars for loading by shippers. Because the car days available is assumed constant (FT) for the period, the number of loads originated (L_o) will drop below the number of loads demanded (L_d). For the period, our identities give us

$$F(T) = C(L_o) \quad (17)$$

and therefore

$$\text{Shortage of car days} = C^*(L_d) - C(L_o) \quad (18)$$

If the actual cycle equals the standard cycle, we can easily calculate the number of unfilled orders (U):

$$U = (L_d - L_o) = \text{shortage}/C^*$$
$$= L_d - [F(T)/C^*] \quad (19)$$

By the end of the period, if all orders are eventually to be filled (i.e., no loads are lost because of delays), the average delay in waiting for a car can be found by relating the unfilled orders to the average daily demand, which is L_d/T:

$$\text{Average delay} = U/(L_d/T) = UT/(L_d) \quad (20)$$

The mathematics becomes more complicated if we attempt to consider the possibility that shippers unable to obtain cars will either decide not to ship or use another mode. Because computer models have been developed to handle such situations, there is no need to pursue such issues in this paper (14).

What happens where there is a shortage? From Equation 17, we see that the actual cycle time multiplied by the actual loads handled must equal the available car days. The longer the cycle time, therefore, the fewer the cars that are originated and the greater the delays to shippers. Clearly, during such periods, reductions in the cycle time can reduce the delays in placing cars for loading, which may provide an immediate benefit by keeping shippers from diverting traffic to other modes and a long-term benefit by keeping shippers happy.

There is evidence that the car service rules used by the U.S. rail industry promote increases rather than decreases in the car cycle (15). When certain types of cars are in short supply, owners may restrict the ability of railroads terminating these cars to reload them. This causes additional empty mileage and increases the empty-car days required to reposition the cars. Hence the capacity of the fleet is reduced precisely when capacity is most in demand. The justification of this system of car service rules is that the owners deserve first priority in loading cars when shortages occur because they incur the risk of having surplus cars when demand is low. Alternative means of distributing empty cars, however, may achieve the same protection for owners with much lower requirements for empty movements. [For example, the clearinghouse railroads have pooled their general-purpose boxcars and use a linear program to determine the required movements of empty cars from one clearinghouse road to another. FCUP has recommended a new approach to freight car management that would extend the clearinghouse concepts (16).]

SUMMARY AND CONCLUSIONS

Car distribution involves moving empty cars from an unloading point to an appropriate reloading point. Car distribution performance can be measured in terms of empty time, empty mileage, the cost of car distribution, and the quality of service (availability of suitable cars when desired by shippers). Car distribution performance is a function of

1. Fleet size relative to the average demand,
2. Spatial imbalances in demand,
3. Temporal imbalances in demand,
4. Institutional restrictions on the use of cars, and
5. Operating policies.

The car days required for a particular traffic volume can be estimated as the product of the number of loads and the standard car cycle, which consists of the following standard amounts of time:

1. To load car,
2. To move car to unloading point (loaded transit),
3. To unload car,
4. To move car to loading point (empty transit),
5. For loaded storage, and
6. For empty storage (including repairs, cleaning, and idle time).

The car days required varies substantially because of the variations in demand, operating policy, and weather. The total number of car days available in any time period is the product of the average fleet size and the number of days in the period. Because cars have long lives, this number is fairly constant for the fleet as a whole. The variability in the number of car days required plus the invariability in car days available combine to complicate management of car utilization in general and car distribution in particular. When requirements are substantially different from the available car days, the result is a noticeable surplus of equipment or delays in placing cars for loading by shippers.

During surpluses, a single-fleet manager would seek to minimize distribution costs, possibly by storing the oldest or least-reliable equipment. During shortages, a single-fleet manager would seek to reduce the time required for each component of the cycle time, including the time required for car distribution, in order to reduce delays in placing cars for loading.

When car distribution and freight car management involve many railroads and other organizations, however, the industrywide response to surpluses and shortages may be less than optimal. Instead of trying to share the surpluses and shortages in an equitable manner, railroads have an incentive to use the applicable rules so as to shift the burden to other railroads or organizations. Hence, studies of these rules and management practices may identify ways to improve car distribution performance. Many such studies have been conducted by AAR through its various committees and FCUP.

In the long run, car distribution performance can be affected by many activities not commonly perceived as car distribution activities (17):

1. Marketing policies (pricing, sales efforts, etc.) that affect the traffic volume, the traffic mix, and the imbalances in traffic flows;
2. Investments (and disinvestments) by railroads and shippers that affect the structure of the rail network and the location of shippers and receivers on that network;
3. Ownership of the fleet, because different sets of rules and objectives apply to different owners; and
4. Degree of standardization of the fleet, because the larger the group of cars being managed, the less important the random variations in traffic volume and the various types of imbalances.

The equations presented in this paper provide the analytic framework necessary to categorize and evaluate the many alternatives available for improving freight car distribution performance. Because performance varies with the actions of many groups and organizations, a coordinated approach will be necessary to achieve substantial improvements. The mathematical relationships provide a logical basis for this coordination.

ACKNOWLEDGMENT

This paper was prepared with support from the University Grant Program of the U.S. Department of Transportation. It does not, however, necessarily represent the views of DOT or any other government agency.

REFERENCES

1. C.D. Martland, M.D. Meyer, J. Sloss, P. Bray, and M. McCord. Improving Freight Car Distribution Performance: Overcoming Organizational and Institutional Barriers to Change. Massachusetts Institute of Technology, Cambridge, MA, M.I.T. Studies in Railroad Operations and Economics, Vol. 36, 1982.
2. C.E. Philip. Improving Freight Car Distribution Organization Support Systems: A Planned Change Approach. Massachusetts Institute of Technology, Cambridge, MA, M.I.T. Studies in Railroad Operations and Economics, Vol. 34, 1980.
3. M.A. Turnquist and W.C. Jordan. Management of Railroad Car Distribution. Office of University Research, U.S. Department of Transportation, Final Rept., March 1982.
4. V.B. Mendiratta. A Dynamic Optimization Model of the Railroad Empty Freight Car Distribution Process. Department of Civil Engineering, Northwestern Univ., Evanston, IL, Ph.D. dissertation, 1981.
5. V.B. Mendiratta and M.A. Turnquist. Model for Management of Empty Freight Cars. TRB, Transportation Research Record 838, 1982, pp. 50-54.
6. Toward an Effective Demurrage System. Reebie Associates, Greenwich, CT, Rept. FRA-OE-73-1, July 1972.
7. C.D. Martland, M.G. Messner, and V. Nowicki. Implementing and Operating/Service Plan: A Case Study of the Boston and Maine. Massachusetts Institute of Technology, Cambridge, MA, M.I.T. Studies in Railroad Operations and Economics, Vol. 29, 1980.
8. J.R. McCarren and C.D. Martland. The M.I.T. Service Planning Model. Massachusetts Institute of Technology, Cambridge, MA, M.I.T. Studies in Railroad Operations and Economics, Vol. 31, 1980.
9. P. French. Managing Freight Cars in a Period of Surplus: Does Utilization Matter? Railway Age, Feb. 9, 1981, pp. 34-38.
10. J.R. Felton. The Economics of Freight Car Supply. Univ. of Nebraska Press, Lincoln, 1978.
11. Task Force II-4. Analysis of Car Service Rules, Orders, and Directives--The Impact of Car Service Rules on Car Utilization. FCUP, Association of American Railroads, Washington, DC, AAR Rept. R-369, May 1979.
12. Catalog of Projects and Publications, 2nd ed. Association of American Railroads, Washington, DC, AAR Rept. R-453, 1980.
13. Task Force II-4. Car Cycle Analysis System. FCUP, Association of American Railroads, Washington, DC, AAR Rept. R-442, Sept. 1980.
14. J.A. Sussman, C.D. Martland, R.D. Juster, R.A. Gray, and J.A. Krueger. The U.S. Railroads: Some Alternatives for Improvement. Multisystems, Inc., Cambridge, MD., 1977. NTIS: PB-295-793.
15. C.D. Martland. Overcoming Fundamental Problems in the Car Hire/Car Service System. Proc., 23rd Annual Meeting of the Transportation Research Forum, R.B. Cross Co., Oxford, IN, 1982, pp. 549-560.
16. Task Force II-4. Proposal for Change in the Railroad Industry-Wide Car Management System. Association of American Railroads, Washington, DC, AAR Rept. R-379, June 19, 1979.
17. C.D. Martland, M.D. Meyer, J. Sloss, R.A. Gray, and M. McCord. Improving Freight Car Distribution Performance: Overcoming Organizational and Institutional Barriers to Change. Massachusetts Institute of Technology, Cambridge, MA, M.I.T. Studies in Railroad Operations and Economics, Vol. 36, 1982.

Publication of this paper sponsored by Committee on Railroad Operations Management.