Research issues in rail transport operations

by

CHARLES E. TAYLOR
Association of American Railroads

INTRODUCTION

This paper will examine some current research issues in North American rail transport operations and some on-going and anticipated research activities related to those issues. The specific research activities examined were selected to illustrate some issues associated with research priorities, approach, data requirements and analytic methodology.

A convenient framework for structuring a review of some of these issues and research activities is the so-called car cycle. Car cycle, as used here, refers to the sequence of communication, decision and physical activities which is repeated each time freight is transported by a rail vehicle. A typical load-to-load car cycle begins with a shipper's order for an empty car, includes such intermediate cycle elements as the assignment of an empty car to fill the order, placement of the car at the shipper's rail siding, release of the loaded car by the shipper, followed by the initial, intermediate and final terminal and line haul handling of the car over one or more railroads and the placement of the car at the receiver's siding. The cycle is completed when the car is unloaded, released back to the rail carrier and, finally, reassigned to a new shipper car order (see Exhibit 1).

This car cycle framework is particularly appropriate in light of the growing recognition of the need for more attention, and research, directed to the disposition of the individual freight car as it moves through the rail network, both loaded and empty. Too frequently, past attention was directed almost exclusively to optimization of train operations, to the neglect of the physical and economic utilization of individual freight cars and the closely related dock-to-dock service provided to rail customers.

Exhibit 1 – Sequence of a typical load-to-load freight car cycle

1. Shipper orders empty car for loading
2. Railroad assigns empty car to shipper's order
3. Railroad delivers empty to shipper's siding
4. Shipper loads car
5. Shipper notifies railroad car is loaded and ready to move
6. Railroad pulls loaded car from shipper's siding, moves it to (origin) terminal
7. Car is processed for outbound movement from origin terminal
8. Car is hauled to intermediate yard by road train
9. Car is processed for outbound movement from intermediate terminal
10. Car is hauled to next terminal where it is interchanged with connecting railroad
11. Connecting railroad hauls car to destination terminal
12. Car is delivered to receiver's siding
13. Receiver unloads car
14. Receiver notifies delivering railroad car is unloaded and ready to move
15. Railroad assigns destination to empty car
16. Railroad pulls empty car from receiver's siding
17. Railroad delivers empty to new shipper for loading

Dramatic increases in the costs of freight car ownership and operation, combined with increased loss of business to competing modes, have resulted in a significant shift in research emphasis to the physical and economic utilization of the freight car as it cycles from load to load throughout its lifespan. An important example of this new emphasis was the establishment, in 1975, of the Freight Car Utilization Program. This Program, which is jointly supported by North American railroads, shippers, labor, and government, is conducting research in such areas as the definition and measurement of utilization, the design and development of a system to collect and analyze car cycle data, improved freight car information and control systems, the effects of car service loading restrictions on utilization, and the relationship between car utilization and service reliability. [1] [2]

CAR CYCLE ANALYSIS

An initial of the Freight Car Utilization Program is the development of a system to continuously capture and analyze data on a sample of freight cars as they move throughout the rail network. A primary objective of this system is to provide the identification and quantification of transport operation problems associated with the movement cycles of selected homogeneous car groups. The system will thereby provide a quantitative basis for the identification of research needs and priorities. The system will also support analyses of specific research experiments designed to correct or minimize such problems, by providing a measurement system to determine the effects of experimental changes on car cycles.

Movement data on some 8,500 cars have been reported to this system since March, 1976. This system will collect, edit and logically sequence car movement data to create complete load-to-load car cycles for each car in the sample fleet. For a typical load-to-load cycle record for a car in the sample, the location data would be recorded for each of the cycle events described in Exhibit 1. Thus, the complete load-to-load cycle could be examined to determine the amount of time the car spent under shipper control, how long the car was in the origin terminal, in each intermediate terminal, how long the consignee took to unload the car, and how long the empty car then took to move to its next loading point.

The sample of some 8,500 cars was selected to provide statistically significant cycle data for specified subsets of the freight car fleet. The freight car characteristics used to partition the freight car fleet for subsample studies were car type, car age, and type of service. Type of service was defined by both commodity carried and by the degree of carrier and shipper control over car movement by partitioning the fleet into three groups: free-running, assigned and private cars. Thus, for selected combinations of car type, car age, car assignment and commodity, analysis of the distributions of overall load-to-load cycle times or specific cycle event times can provide valuable insight for the identification of car movement and utilization problems. Equally important, the same system can be used to evaluate the effects of changes introduced into the transportation system which were designed to improve cycle times.

The attention and resources that are being devoted to this car cycle analysis system are in response to a growing appreciation of the complexity of the rail transport system and the attendant need for a more comprehensive and quantitative approach to the analysis of the operations of that system. The judgement of those capable and experienced in rail operations and analysis is an essential ingredient to this research, however, the complexity of the rail operating network demands that the researchers also have available data systems designed to identify the source and magnitude of network operating problems, and to measure the total system effects of changes designed to reduce those problems.

Some preliminary analyses of such car cycle data have helped document the need for further research in such areas as empty car distribution, and rail yard and terminal operations. The remainder of this paper willl discuss some of the on-going research and perceived research needs in these two major areas on the premise that, collectively, they contain a fairly representative mixture of some of the more important research issues related to rail transportation operations.

EMPTY CAR DISTRIBUTION

Car cycle data frequently reveal inefficient handling during the empty portion of the cycle. A major source of this inefficiency is lack of adequate information and techniques for the efficient matching of empty cars to shipper car orders. Examination of this process has led to the identification of research needs in three areas:

1. Forecasting of shipper demand for empty cars,
2. Forecasting of the supply of empty cars, and
3. Efficient decision rules and assignment algorithms to match supply and demand.

Research to date has concentrated primarily on the first and third needs.

Demand Forecasting

A current project, supported jointly by the Federal Railroad Administration (FRA) of the U.S. Department of Transportation, and the Association of American Railroads, has provided valuable insights into the problems associated with shipper demand information and forecasting, and has resulted in the development of a car demand forecasting model which has significantly advanced the state of the art. [3]

The analysis of the data collected for this research, combined with discussions with rail shippers and railroad personnel, led to the following conclusions:

1. Most railroad forecasts of shipper demand for empty cars are based on historical car loading data. Such forecasts typically understate demand, since the percentage of demand not filled is ignored. Such data is also misleading since it ignores substitution of equipment, i.e., it describes equipment loaded rather than equipment requested. These observations led to a recommendation to record historical shipper car order data in addition to historical car loading data.

2. Analysis of shipper car order data indicated that shipper demand is highly variable and lead time on orders extremely short. Analysis of sample data collected on two major railroads for this project revealed fluctuations in demand ranging from plus or minus 50% of the average, and the average lead time for the order data was only 1.5 days. This combination of widely fluctuating demand and short lead time makes it extremely difficult for rail carriers to efficiently distribute equipment.

3. Since most areas unload either more or less cars than they load, railroads must maintain large inventories of empty cars at loading points and continously move empty cars from surplus to deficit areas.

4. The size and impact of the time lag to recognize and react to changes in demand patterns can be reduced if potential surplus and deficit areas can be forecasted and car flows adjusted to reduce imbalances before they occur.

As a result of this research, a forecasting system was developed which was designed to be used as part of the freight car allocation and distribution process of major railroads. The forecasting system objectives included accuracy sufficient to aid empty car distribution, information requirements consistent with the capabilities of rail information systems, and simplicity and robustness sufficient to require only infrequent maintenance and updating of the estimation and forecasting algorithms imbedded in the information system.

The forecasting model developed provides one, two, and three-week ahead forecasts of empty car demand, by car type. The model utilizes three sources of information:

1. Historical car loadings,
2. Historical car orders, and
3. A subjective order forecast by the local railroad agent.

More specifically, the demand forecast is produced by a linear composite model using as input variables, (1) a loading forecast based on previous loadings, (2) an order forecast based on previous orders, and (3) the local railroad agent's subjective forecast of orders. The first two forecasts are developed using discrete linear stochastic models, often referred to as Box-Jenkins Models. The model is of the form:

$$O_t = W_0 + W_1 \hat{L}_t + W_2 \hat{O}_t + W_3 A\hat{O}_t$$

where:

O_t = the number of cars of a given type ordered during week t

$\hat{L}_t$ = the forecast of loadings during week t given by a loadings forecasting model using actual loading data

$\hat{O}_t$ = the forecast of car orders during week t given by an orders forecasting model using actual order data

$A\hat{O}_t$ = the subjective agent order forecast for week t

W_0, W_1, W_2, W_3 are coefficients estimated by regression analysis.

This composite model has the logical advantage of using predictor variables which are subject to somewhat offsetting biases. Loadings tend to be biased downward as they neglect unfilled demand, whereas orders tend to be biased upward as shippers overorder to protect their requirements. Additionally, the local railroad agent is in a unique postion to assess the probable effects of irregular occurrences such as strikes and severe weather conditons. (The mean absolute forecast error for the study sample of 27 region-car type combinations was 11.77% of the mean weekly car orders.)

Interestingly, for the 12 car types operating in six geographic regions examined during this research study, an autoregressive model of order 1 [AR(1)] provided the best all-around forecast model for both the loadings forecast ($\hat{L}_t$) and the orders forecast ($\hat{O}_t$).

Plans are currently under development to install a linear composite forecasting model of the type described herein on one or more North American systems for demonstration and further development.

A critical limitation of existing rail information systems to the implementation of any such forecasting models is the lack of current and historical car order data. For the model described, data on both loadings and car orders are required for the previous 26-week period. Such data are usually recorded for car loadings. Car orders typically enter the rail system at the local level. However, these orders are usually transmitted to the central data system only if the order cannot be filled by an empty car available in the local area. Even the orders reported to the central system are purged once they have been acted upon. Historically, car orders have not been considered accurate measures of demand, and this attitude has been carried into the design of most rail information systems.

Supply Forecasting

There is a recognized need for improved techniques to forecast the supply of empty rail cars for more efficient distribution and to minimize empty car inventory requirements at terminal loading points. Empty cars are moved to loading points to fill anticipated demand whenever local car availability is estimated to be inadequate. Too frequently, such moves are inefficient because, subsequent to their initiation, unpredicted changes occur in supply or demand. As an example, a railroad may assign and begin transporting more empty cars to a terminal than will ultimately be required because predictions of cars received in interchange from connecting railroads, or cars released empty by local rail customers were too low. Since techniques currently employed to forecast both supply and demand are imprecise, most railroads tends to protect their customers by maintaining unnecessarily large empty car inventories at terminal loading points. With the average purchase price of freight cars now in excess of $25,000, there is strong economic incentive to develop more efficient methods to forecast car supply as well as demand, in order to improve the utilization of such equipment.

A car supply forecasting system would likely be a composite model with input from such sources as:

1. real-time terminal inventory of empties on hand,
2. cars estimated to be released from local cleaning and repair facilities,
3. estimates of empty cars to be released by local customers,
4. estimated arrivals of empties enroute to a terminal from points within a railroad, and
5. estimated arrivals of empty cars from connecting railroads.

A task to develop an empty car supply forecasting procedure which can be widely used with existing or planned railroad data is currently under consideration for the next phase of the Freight Car Utilization Research Program.

Empty Car Assignment

The distribution of empty cars on most North American railroads is based on variations of the "flow-rule" concept. Flow rules are essentially a set of instructions which specify the disposition of empty cars for every terminal point on a rail system. These rules are now based on best available estimates of the location and magnitude of empty car surpluses and empty car deficits, i.e., terminal-specific supply and demand estimates or forecasts. These rules also attempt to minimize total empty car distribution costs using techniques which range from fairly fixed sets of empirically derived distribution rules which designate preferred supply points for each demand point, to frequent updates of the rules using linear programming techniques. Some railroads have operating control systems which permit flow rule adjustments one or more times per day, whereas other roads make such adjustments on a weekly or monthly basis.

At present, most of these flow rules are based on experience and rough estimates of average terminal area supply and demand by car type.

There has been considerable research and development devoted to improve empty car decision rules and assignment algorithms. [4] The major constraints to the further development and implementation of such systems appear to be more related to the availability of data necessary for supply and demand forecasting than to the availability of the necessary analytic techniques or operating information and control systems technology.

The functional requirements and specifications of a viable empty car assignment algorithm must take proper account of such factors as:

1. current and anticipated accuracy of supply and demand information and forecasting systems,
2. response time of the rail system to assignments and assignment changes,
3. comparison of costs to maintain local inventories of empty cars with costs to move cars from supply points elsewhere in the rail system, and
4. efficiency requirements of assignment algorithms (computer operating costs).

Proposed research to examine relationships between current and anticipated supply and demand forecasts and allocation logic requirements is currently under consideration by AAR and FRA. Preliminary investigations may include rail network simulations designed to examine the efficiency of various empty car assignment algorithms beginning with assumptions of perfect supply and demand information and progressively degrading forecast quality. Such research would help determine what distribution logic is best suited for various assumptions of uncertainty. A research objective would be the development of heuristics for desensitizing assignment models to avoid unnecessary empty car moves intended to satisfy spurious demand caused by forecast errors.

The assignment algorithms currently available are batch processors. Car orders and location and status data on empty cars may be accumulated for a 24-hour period. The algorithm is then run to generate car assignments. This batch procedure often leads to inefficiencies due to the continuous nature of this process. More frequent iterations of such algorithms would avoid inefficient assignments caused by out-dated empty car supply and demand information. Obviously, the optimum frequency of iterations for assignment algorithms is closely related to achievable levels of forecast accuracy. Thus, requirements for frequent iterations of the empty car assignment process provides incentive for the development of efficient algorithms which require minimum computer processing time.

YARD AND TERMINAL OPERATIONS

Analyses of car cycle data consistently point to yards and terminals as major sources of delay to both loaded and empty car movements. These delays have been identified as the primary cause of transit time unreliability, a major source of complaints by rail transport users. Preliminary examinations have revealed that two of the sources for such delay and unreliability are (1) the management of operations within a classification yard, and (2) the coordination and control of car movements over two or more railroads within a multi-road interchange terminal.

Classification Yards

Considerable research and development have been directed at the operation of classification yards. The emphasis of most of this research and development has been in such areas as improved information and process-control systems. Modern classification yard systems, for example, now have the capability to automatically generate switch lists from verified train consists, to automatically control locomotives engaged in classification switching, and to automatically line switches and control car velocity to insure cars are switched into the correct classification tracks and that they couple to the standing cars in those tracks at the proper velocity. In addition, these systems can monitor cars pulled from the classification track for train makeup and automatically maintain a real-time inventory of the location of every car within the yard. While these systems have contributed significantly to the efficiency of yard operations, substantial opportunities for productivity improvements remain in even the most modern of these automated yards. A major source of this improvement potential derives from the fact that most of the work planning and scheduling, and much of the real-time monitoring of work progress is still performed manually in even the most modern yards. Recognition of this potential has led to plans for research in such areas as computer-based decision systems for the advanced planning and scheduling of yard work.

The productivity of classification yards is extremely sensitive to the human ability of individual supervisors to plan, schedule and monitor yard tasks. These tasks typically include switching cars from receiving tracks to classification tracks, pulling cars from classification tracks to assemble outbound trains, and inspection and repairs of cars and locomotives. Great variability has been observed between individual yard supervisors, with differences of over 30% in the number of cars handled per eight-hour shift not uncommon. These observations confirm the need for research which would examine the types of decisions these supervisors are required to make, the information available upon which to base those decisions, and comparisons of decisions made by the more productive supervisors with those of the less productive. This significant relationship between the "human factor" and productivity suggests opportunities for improved planning and decision aids. The more productive supervisor may be the one who has greater natural ability to mentally visualize, organize and recall the key elements of his yard operation. Once these ability requirements and decision processes are documented and analyzed, it should be possible to develop information and analysis systems to support the management process. Such a management support system could contain current information on yard configuration, locomotive and car inventory, arriving and departing trains, and available work resources such as yard engines and mechanical inspection and repair crews. The system could evaluate work scheduling plans proposed by yard supervisors by simulating them in advance of execution and developing estimates of the results. A more advanced system could actually develop an optimum solution to the yard work planning problems and recommend a sequence of work tasks for consideration by the yard supervisor.

A management support system for yard operations could provide supervisors with continuous monitoring of all important work events, and prompt them whenever new decisions or updated work plans are required. Such a system could also be used for training yard supervisors and as a planning and research tool to evaluate proposed yard system changes.

Terminals

Car and train movements between yards, to and from local rail customers, and between two or more railroads within multi-road interchange terminals are receiving increasing attention by rail researchers. The justification for this growing attention also derives from analysis of car cycles, which consistently point to operations within these large terminal complexes as major sources of delay. A major industry response to these terminal problems has been developed by the Labor-Management Task Force on Rail Transportation, comprised of representatives from railroad labor and management and supported jointly by labor, management and government. This Task Force has established project teams in several large multi-road terminals. These terminal project teams are identifying the sources of delays to car and train movements within these terminals.

Once these sources have been identified, the project teams, working with local rail management and labor, identify the underlying causes of the delays. Specific experiments designed to correct these problems are then developed and proposed. If approved by all affected parties, the experiments are conducted under closely controlled conditions and the effects are carefully measured and analyzed.

This research approach has proven to be very successful. Several terminal project experiments have already resulted in permanent changes in labor agreements and operations which have resulted in substantial reductions in delays and improved service to rail customers. As an example, the average time a car spent in the St. Louis Terminal on the Missouri Pacific Railroad decreased by nearly 25% in the six-month period from March to October, 1975 as a result of this research approach. [5] The St. Louis project on the Missouri Pacific Railroad was the first of these terminal projects. Further testimony to the success of this approach to rail transport research is the recent installation of project teams in the terminals in Chicago, Illinois and Houston, Texas.

The initial requirement for each of these terminal projects is a data analysis system for the evaluation of car movements within the terminal. The experience of the first project team on the Missouri Pacific Railroad in St. Louis established the need for such a data system as essential to both the identification of specific car movement problems and the evaluation of experiments. The system design is oriented to our old friend the car movement cycle. The car cycle within the terminal can begin either when the car arrives at the terminal on a road train or is released by a local customer within the terminal. Intermediate terminal cycle events can include arrivals to and departures from yards within the terminal, and delivery of the car to other rail carriers that also operate within the terminal. The terminal cycle can close with either delivery of the car to a customer within the terminal or its departure from the terminal on a road train.

The AAR currently has under development a data collection and analysis system designed to support the Labor-Management teams in Chicago, St. Louis and Houston. The system will be designed to support project teams in additional terminals as they come on line. The railroads that operate in these terminals will create a weekly magnetic tape record of all car movements on their rails within those terminals from their master car movement data files. These tapes will then be sent to the AAR where they will be loaded into a program that will chronologically sequence the individual cycle elements of each car as it moves through the terminal. Thus, the terminal cycle for a car that moved over more than one railroad within a terminal will be created by merging the cycle event data for that car from each carrier that participated in the movement, in chronological sequence. This terminal cycle data base can then be used to generate reports designed to identify specific car movement problems. The system will provide weekly reports on the movements of car groups which should move through the terminal in the same manner, receiving identical handling. These weekly reports can provide statistics on the distributions of time for the overall terminal cycle and for the individual cycle elements for each homogeneous group of cars. The system will also be designed to generate special reports and support special analyses. Thus, by inputting specific parameter values, the terminal project teams can call statistical reports on any one or combination of terminal car movements.

This data collection and analysis system will not only provide for the identification of specific terminal car movement problems, but will also provide insights and quantitative assessments essential to the establishment of research priorities. Once the preliminary problems and priorities have been established, the next step is for the terminal project team, working closely with terminal labor and operating management, to identify the specific characteristics and underlying causes of the high priority research problems that have been identified. Specific experiments to correct or improve these problems are then designed. If all parties that would be affected by the experimental change are willing to proceed with, and participate in, the experiment, and the Labor-Management Task Force on Rail Transportation authorizes the experiment, it is conducted. The effects of the experiment on the total terminal rail network will be measured using the data collection and analysis system. Some experiments may require supplemental data collection and analysis using either the data systems of the participating railroads or manual data collection.

Experiments to date have included both temporary changes in operating procedures, and temporary relaxation of labor agreements. Many of the experiments involved yard-to-yard movements of cars. Cars handled by more than one yard within a terminal frequently experience excessive delay and are not moved consistently. For example, temporary relaxation of labor agreements, which permit more flexible use of terminal switching crews and road train crews for the transfer of cars between yards have proven very successful in reducing the time cars spend in yards. Additional experiments either underway or under consideration involve changes in schedules for intra-terminal movements, revisions in procedures for processing arriving road trains, and the addition of switch engine assignments. All of these experiments are in response to problems initially identified through analysis of the terminal car cycles of specific car groups.

At the completion of each experiment, the data are analyzed to assess the impact of the change on car and train movements. A final report is prepared for review by the Task Force and all directly affected parties. It is then up to local railroad management and labor to review the findings and determine whether permanent changes to operating practices and labor agreements are in everyones' best interest. As mentioned previously, this research approach has already resulted in permanent changes with significant improvements in operations and service to rail customers.

CONCLUSIONS

The research activities and associated problems discussed in this paper were selected to illustrate some of the more important research issues in North American rail transport operations. Among these is the need for more systematic approaches to the identification of research needs and the quantification of those needs to enable the rational establishment of priorities. Too frequently research has been based on the judgement of individuals whose knowledge encompassed only a subset of an overall problem. At best, such projects consume scarce resources that could have earned much higher returns if invested in research more carefully conceived. At worst, they reduce or correct subset problems at the expense of the rest of the system. The complexity of rail operating networks requires a system for the collection and analysis of operations data to monitor the performance of transport operations, identify undesirable performance, and thereby direct the intelligent allocation of research resources.

A related issue is the need to measure the total system effects of changes introduced into the rail operating network. Complex network interactions often make it impossible for even the most knowledgeable operations and research personnel to predict where and how the system will respond to change. The introduction of operational changes often results in secondary system responses which are difficult or impossible to anticipate. Thus, data collection and analysis systems are necessary to identify the location and magnitude of the response of the overall system to experimental changes in operations.

It was not until transportation researchers began examining individual car movements that some of the really significant rail transport problems and associated research needs were identified. Until recently, however, most computer information and control systems for rail transport operations in North America were not designed to provide the data necessary for such examinations. Fortunately, there is increasing recognition of the need for such evaluation and analysis, and many railroads are now providing for this capability in their information and control systems.

The development of this data and analytic capability by individual railroads must be complemented by a similar industry-level capability to provide data for the evaluation and analysis of multi-road car movements. Recognition of this need formed the basis for the development of the car cycle analysis and terminal car movement systems previously described.

Thus, the real challenges facing researchers in rail operations derive from the need for better information with which to describe those operations. The analytic methodologies required for such research are, for the most part, within the present state of the art. The challenge is to develop systems and procedures to capture data to monitor rail transport operations, identify problems and support research analysis.

REFERENCES

[1] Shaffer, Frank E.; Taking Car Utilization To Task; **Modern Railroads**; July, 1975, p 56.

[2] Car Utilization Final Task Force Report, Recommended Research Program; **Association of American Railroads**; July, 1974.

[3] Freight Car Demand Information and Forecasting Research Project, Phase I Final Report; **Federal Railroad Administration, U.S. Department of Transportation**; March, 1975.

[4] Development and Evaluation of a Computer-Based Model For Optimal Railroad Freight Car Distribution: Final Report; **Federal Railroad Administration, U. S. Department of Transportation**; July, 1975.

[5] A Cooperative Program of Experiments Involving Changes in Railroad Operations: St. Louis Project 1975 Progress Report; **Association of American Railroads**, April, 1976.