Formerly ISSN 1046-1469

JOURNAL OF THE TRANSPORTATION RESEARCH FORUM

Volume 39, Number 1 Winter 2000

TRANSPORTATION RESEARCH FORUM

Revenue Management in Railroad Applications

This article explores common characteristics and critical differences between a variety of railroad revenue management problems. Most railroad problems tend to focus primarily on origin destination traffic management rather than overbooking or price discrimination. These problems tend to be highly network oriented, have booking arrival patterns independent of fare class value, and except for long distance passenger service, have very short booking lead times. Any of these features of railroad problems can cause difficulties for traditional airline-style, leg-based "EMSR" approaches. A bid price methodology effectively addresses many of the common core requirements of these railroad applications.

by Edwin R. Kraft, Bellur N. Srikar, and Robert L. Phillips

This article explores common characteristics of and critical differences between a variety of railroad industry revenue management problems, comparing these problems to their trucking and airline counterparts. Passenger and freight railroad revenue management problems share a number of common characteristics since the services are produced using similar technology, but each has its own peculiarities due to the differing nature of markets served. Even within freight or passenger applications, distinctive market segments exist having different characteristics from a revenue management perspective. Current revenue management systems for passenger railroads can be improved by moving to a "bid price" approach; this method can be applied to freight railroad problems as well.

The next section describes several key components of any revenue management process. The remainder of the article discusses applicability of these approaches to railroad problems in particular, identifying both common elements of and important differences between freight and passenger problems. The current implementation status of revenue management at passenger railroads is reviewed, but since no applications are yet known in freight railroading, a survey of current academic literature is offered in this area. Finally, future research opportunities and needs will be noted.

Key Components of Revenue Management

Revenue management increases revenues through application of three techniques: overbooking, discount allocation and origin-destination traffic manage-

Journal of the Transportation Research Forum, Vol. 39, No. 1, Winter 2000 (157-176)

ment.

Overbooking: The purpose of overbooking is to compensate for reservations made but not used. The need for overbooking depends on the statistical "no show" rate, as well as the carrier's willingness to risk an oversold situation.

Discount Allocation: The basic principle of discount allocation is to protect seats for high valued future demand by limiting current availability of low priced fares, so that the risk between revenue dilution and inventory spoilage is optimally balanced. Discount allocation implements a form of price discrimination, exploiting the fact that different customers may be willing to pay different prices to receive essentially the same service. For example, airline business travelers usually pay a higher fare to receive essentially the same service as leisure travelers (although business customers gain more flexibility to change their plans and book travel at the last minute.) Low fares offered in price sensitive market segments can stimulate demand to fill capacity which might otherwise go unused. Fare restrictions, such as requirements for advance purchase and Saturday night overstays, as well as limitations on the total amount of low-priced capacity offered, attempt to prevent price-insensitive segments from being able to utilize the discounted fares.

Traffic Management: Traffic management addresses the situation where different customers receive different services, requiring the use of different combinations of resources in the process. Typically, a long distance passenger pays a higher fare than a short distance rider, but shorter trips may produce higher revenue per passenger mile. The solution to a traffic mix optimization or *traffic management* problem determines the optimal selection of short- and long-distance traffic, or short- and long-term car or hotel room rentals, to most effectively utilize capacity and maximize total revenues. Solving a traffic management problem typically requires network modeling and cannot be approached by individual leg or time period.

Historically, the development of revenue management literature has been strongly influenced by the capabilities and limitations of airline reservations systems. As reported by Williamson (1992), most airline reservations systems were designed 20-30 years ago when the market environment was much simpler. These systems only allowed for physical control of seat inventories at the fare class and flight leg level, rather than by origin-destination.

As airlines evolved hub-and-spoke network structures, they began to recognize the need for origin destination traffic management. Considerable effort has been expended trying to devise ways to "work around" limitations of old reservations systems: see, for example, "virtual nesting" by Smith, Leimkuhler, and Darrow (1992), which groups origin-destination fares into buckets by value based on clustering algorithms, rather than directly determining fare class hierarchy based on the fare class code. More recently, underlying airline reservation system limitations have started to be directly addressed based on "Seamless Availability" (Phillips, 1994a, see Appendix), making "bid price" approaches to Revenue Management feasible, even for airlines.

The "bid price" method works by calculating "opportunity cost" for units sold. Utilizing this method, one should never sell a unit of capacity for less than its opportunity cost, even though the direct revenue impact may be positive. Following Williamson (1992, p. 90-92):

> The idea behind the bid price approach is to establish a "cut-off" value for each flight leg which can be used to make decisions whether to accept or reject differ-

ent ODF (origin-destination-fare class) requests. The difference in the methodology of the bid price approach, when compared to other conventional seat inventory approaches, is that ODF inventories are either open to bookings or closed; there are no explicit booking limits for different ODF's. For a single leg itinerary, a fare class is open for bookings if the corresponding fare is greater than the bid price, or shadow price, for the leg. For a multi-leg itinerary, the total fare must be greater than the sum of the bid prices from the respective flight legs it traverses.

One advantage of the bid price approach is that it is a very simple method of managing seat inventories. Hence, it would be very easy to implement in a reservations system when compared to OD and fare class approaches. The disadvantage of the bid price approach, however, is its open/closed control philosophy. If a given ODF passes the bid price criteria, that ODF remains open to bookings until the bid prices are revised. Thus, in order for the network bid price approach to be an effective seat inventory control approach, frequent revisions would be necessary, requiring both reoptimization and reforecasting. For a truly optimal system, revisions would be necessary on a real-time basis.

Several recently-implemented real world airline, hotel and rental car revenue management systems utilize bid prices as their main control methodology. Many of these systems continuously update the bid prices in real time and in many cases, they update the demand forecasts as well. Following Phillips (1994a, p. 8):

> Supporting the on-line functions of a dynamic system such as the ODRMS is a number of "off-line operations" . . . These off-line operations include optimization (calculation of bid prices), demand forecasting and fare forecasting. Typically, the optimization and demand forecasting functions are run often—possibly as often as every booking and cancellation while the fare forecasting function is run more infrequently to reflect fare changes. The underlying philosophy is that the data in the system at any time should reflect the latest information available.

In the "pure" bid-price approach, bid-prices alone are used to control availabilities. While appealingly simple, the "pure approach" has two drawbacks:

(1) It is inherently incremental. The bid price only indicates whether or not a carrier should accept the next incremental booking. It does not tell the carrier what to do about multiple booking requests (groups) or non-incremental bookings (large freight shipments with substantial volume and weight.)

(2) The "pure" bid-price approach requires that the bid-price be updated regularly, especially following bookings and cancellations. While this might be technically feasible, it is a very heavy requirement to place on an information system, particularly in light of the possibility of occasional hardware failures, communication links going down, etc.

For these reasons, the bid-price is generally supplemented by other control mechanisms that deal with the two is-

sues above. One is a simple "safety net" limitation that prevents bookings from exceeding capacity (possibly adjusted for overbooking). Another approach is to define "gradients" to allow the reservation system to adjust the bid prices itself after each booking without requiring full reoptimization. A third approach is to define increments at which the bid price should be recalculated ("triggers", see Phillips, 1994b). Thus, we might calculate the bid price as $100 and also specify triggers of (-5, +8), meaning that the bid-price model should be recalculated if we get more than 8 net bookings or more than 5 net cancellations. Summarizing, the bid price method is attractive for railroad revenue management applications for the following reasons:

(1) The bid price method easily handles network traffic management problems using a simple control mechanism: based on legs traversed, if the sum of the bid prices is less than the revenue for the whole trip, the booking request should be accepted, otherwise it should be rejected.

(2) Bid-price based revenue management systems must be designed to update the bid prices frequently, in real time if possible. These operational characteristics also make a bid price system well suited to cope with the short booking curves that are characteristic of many railroad passenger and freight problems.

(3) Updating bid prices frequently also eliminates the need for any fare class nesting assumptions. Belobaba (1989) found, as the frequency of updates to bid prices or allocations approaches real time, the impact of nesting diminishes and eventually disappears entirely. Thus, as long as the model is solved frequently enough, it is not necessary to include nesting in the mathematical formulation of the optimization model. A bid-price system can accommodate any booking arrival pattern, as opposed to an EMSR-based system (Belobaba, 1987), which tends to "overprotect" space for the higher value fare classes.

(4) A bid price-based management approach could readily combine elements of Powell's (1987) "Regional Impact Model" with traditional leg-based revenue management, to jointly optimize equipment allocation with line haul capacity utilization.

(5) A "safety net" function using "triggers" (Phillips, 1994b) or other similar control mechanism can be used to detect when a multiple group booking request should be flagged for manual intervention or other special handling in the system.

Revenue Management as Practiced by Railroads

Most railroad revenue management problems tend to focus primarily on traffic management rather than overbooking or price discrimination. Amtrak's long distance passenger trains carry a very small proportion of business trips—nearly all the ridership consists of leisure class travelers. In the northeast, premium fare Amtrak Metroliner service is primarily targeted to time sensitive business travelers—most leisure class customers are accommodated on lower fare Northeast Direct trains. The relative homogeneity of markets served by each train service does not allow much of an opportunity to improve revenue through price discrimination. Standees are un-

acceptable on long distance trains or premium fare Metroliners, so overbooking must be practiced very conservatively. However, there is still a tremendous opportunity for traffic mix optimization in both freight and passenger rail applications.

In general, railroad revenue management problems tend to be *highly network oriented*, have *booking arrival patterns independent of fare class value*, and except for long distance passenger service, have *very short booking lead times*. Any of these features of railroad problems can cause difficulties for traditional airline-style, "EMSR" leg-based revenue management approaches.

Railroad Problems Tend to be 'Network Oriented'

Nearly all passenger railroad problems are heavily network oriented, due to the large number of intermediate station stops made by the typical train. Each pair of adjacent station stops defines a "leg" for which the opportunity cost and/or capacity allocations by fare class must be determined.

Figure 1 illustrates the markets, which can potentially be served by a single train. Since this WAS to NYP train makes 5 intermediate station stops, it has 6 legs and can serve 21 possible origin-destination city pairs, or markets. Having 6 fare classes per OD pair results in 126 possible origin destination market classes, which can be served by this single train departure for each of which demand forecasts must be individually developed.

Considering only full fares for the sake of simplicity and without loss of generality, a request for NYP-WAS full fare must be evaluated by the traffic optimization model against all other possible requests, such as NYP-WIL and WIL-WAS. If one considers only NYP-WIL full fare versus NYP-WAS full fare, then the decision to accept NYP-WAS full fare

Figure 1: Traffic Displacement Example

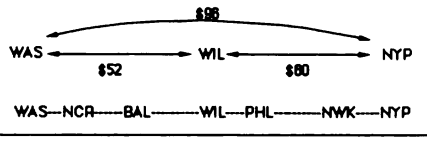

is trivial. However, if WIL-WAS can be sold also, then the choice becomes more complicated. Now the value of NYP-WAS is no longer equivalent to the $96 full fare but has to be reduced to compensate for the probability of down line displacement of expected WIL-WAS revenue.

In Amtrak's system, these relative values are estimated by solving a deterministic linear programming model (Williamson, 1992, pp. 68-69). The shadow prices and the reduced costs are then used in a heuristic model (Powell, 1989) to determine the relative revenue value of each market after scaling for displacement costs, called the Cumulative Relative Revenue (CER).

On low demand trains, since there is a low probability of closing any leg, it will always be beneficial to accept a long-haul request such as NYP-WAS because there is no short-haul revenue displacement. However, if a train has peak demand legs, then market allocations among NYP-WAS, NYP-WIL, and WIL-WAS have to be based not only on fare values but also on the probability of selling short hauls versus long hauls. The same principle can be extended to multiple market fare classes to determine the optimal market class allocations to maximize revenue from the traffic mix accepted.

Table 1 gives the total number of legs and the average number of legs traversed by the average passenger for some typical Amtrak trains in July 1997. Normally the average passenger does not travel the entire length of the train's route but still traverses a large number of legs. A no-

table exception (not shown) is Train 52/53, the Auto Train, which operates as a shuttle between Lorton, VA and Sanford, FL with no intermediate stops. Clearly this problem is quite difficult as compared to a typical airline "hub and spoke" problem, which may have perhaps two or three flight legs per average passenger.

Table 1. Number of Legs Per Amtrak Train and Average Legs per Passenger

Train	# of Legs	Avg Legs Psgr
Sunset Limited ORL-LAX	40	14
Cardinal WAS-CHI	27	12
California Zephyr CHI-OAK	38	16
Silver Meteor NYP-MIA	33	14
Metroliner #101 NYP-WAS	9	3
Northeast Direct #95 BOS-NPN	24	8

Freight railroad problems are comparable in level of difficulty to passenger rail problems, with an added twist: the rail freight network must explicitly represent allowable origin to destination connections, or "blocks," so that not only train capacity utilization, but also shipment routing can be determined by the optimization code. A "block" is a set of cars temporarily joined for the duration of a trip between a common origin and destination (Campbell, 1996). Grouping cars in this manner is required for both convenient assembly of trains in yards and also to facilitate efficient pickup or setoffs of groups of cars at intermediate locations.

Legs or route segments, to which train capacity constraints apply, can be derived from block pickup and setoff locations. For example, Figure 2 shows a train which handles two blocks of cars: Block "A," picked up at node 1 and set off at node 3; and Block "B," picked up at node 2 and set off at node 4. This train's route would be divided into 3 legs or segments; break points occur whenever pick up or set off activity occurs.

Figure 2: Representation of Blocks and Legs

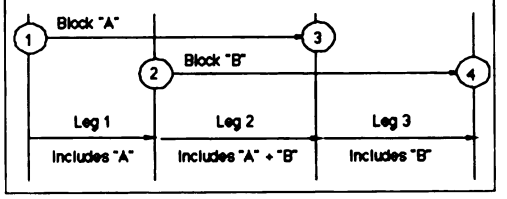

A single freight shipment typically rides on several trains, passing through several intermediate terminals, before it reaches its final destination. Figure 3 shows the number of trains used between origin and destination, for a small 12 yard, 16-train MIT test problem (see Kwon, 1994 and Kraft, 1998.) By comparison, Figure 4 shows that shipments often traverse many more "legs" than the number of trains. Clearly real world railroad freight problems would be much larger and more complex than the test data presented here. But even this small test problem exhibits a high level of interdependency across trains and time periods.

The 'Highest Fare Class Books Last' Assumption is Seldom Satisfied

"EMSR" (and related approaches) explicitly assume that bookings occur in reverse fare order—with the highest fare booking last. If this assumption is violated, it is well known that "EMSR" tends to "overprotect" allocations, or set aside too much capacity for only the highest fares. The cost of this overprotection depends on how frequently the allocations are updated.

In industries other than airlines, the basic paradigm of "highest fare books last" is just not a good assumption. For example, in the rental car industry, leisure customers tend to pay some of the highest fares and they tend to book earlier than lower-fare business customers

Figure 3: Trains Traversed in MIT Shipment Routing Problem

Figure 4: Legs Traversed in MIT Shipment Routing Problem

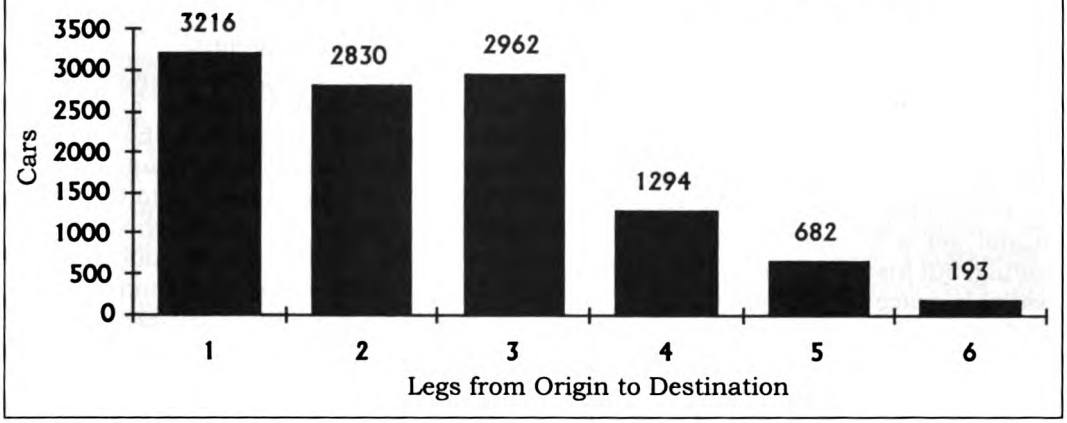

Table 2. Booking Lead Times for Selected Amtrak Trains

Trains	Avg. Lead Days	Avg. Lead exc. Cancel	% Depart Day Res.	% Depart Day Cancel
Sunset Limited ORL-LAX	115	97	5.9%	24.2%
Cardinal WAS-CHI	100	78	4.0%	16.4%
California Zephyr CHI-OAK	146	141	8.0%	16.3%
Silver Meteor NYP-MIA	93	53	10.5%	26.9%
Metroliner #101 NYP-WAS (AM)	5.6	3.7	19.5%	11.7%
Metroliner #119 NYP-WAS (PM)	6.6	4.5	43.4%	25.7%
Northeast Direct #95 BOX-NPN	41	26	20.6%	19.6%

(who are often eligible for substantial corporate discounts.) Similarly, it is certainly not the case that lower paying freight always books earlier than higher-paying freight—for example, some high-value, high-paying freight may book very early in order to be assured of space available. For most of these industries, a more realistic assumption is that bookings over time reflect a mixture of high-value and low-value customers, with the balance possibly shifting over time. As airlines exercise more fare flexibility, this assumption is tending to become truer for them as well.

Availability of last minute discounts such as "standby" pricing may encourage some customers to wait to the last minute to get the cheapest possible price, gambling that space will still be available, or to cancel existing reservations and rebook the space at a lower price. Freight and rental car applications have natural price "fences" that discourage this kind of customer gaming behavior. For example, in the rental car example, pricing depends on having a corporate discount, and in freight, prices normally are contracted in advance, a customer cannot get a cheaper price simply by waiting, but just risks the possibility that available space might sell out.

For any given flight leg, network revenue management also works to invalidate the "highest fare books last" assumption. It seems to be the case that customers who are booking on a longer or more complex itinerary (and thus tend to be higher value) tend to book earlier. American Airlines' virtual nesting approach (Smith, Leimkuhler, and Darrow, 1992) suffers the same problem: If a short distance business fare is virtually nested in the same "bucket" (grouping of fares having approximately equivalent value) alongside a long distance discount fare, the timing of these demands would be spread across time, not concentrated at the end of the booking period. This would lead to overprotection of the higher value fare classes under a "virtual nesting" approach.

Short Booking Lead Times are Typical

For passenger services, as shown in Table 2, long distance trains usually are reserved with a long lead time, but high speed corridor services experience short booking curves. This is consistent with the market orientation of these services, whereby Amtrak's long distance and Northeast Direct trains mostly serve a price sensitive personal or leisure travel market, but the New York to Washington Metroliner service is oriented mainly towards business travelers. For selected July 1997 Amtrak trains, Table 2 gives the average booking lead time, average lead time excluding cancelled reservations, percent of same day of departure bookings and percent of departure day cancellations.

In general, high frequency corridor services such as Metroliner and Northeast Direct experience both shorter booking lead times and higher cancellation/rebooking rates, since customers have more traveling options to choose from. Evening departures have more cancellations and late bookings than morning departures since many travelers change their reservations if planned activities end later or earlier than planned. Daily updates to allocation levels are just not frequent enough when, as in the case of Metroliner #119, over 40% of demand does not materialize until the day of departure. Because of this system limitation and other marketing considerations, Amtrak's Metroliners have been removed from discount allocation and traffic mix control optimization. They are overbooked to a very limited extent.

Freight shippers, in general, are not accustomed to having to reserve space far in advance, if at all, although some notable exceptions do exist—particularly in international container shipping.

Campbell (1996, p. 54) reports that intermodal shippers generally provide no more than a 24-hour advance notification of demand. Some customers who ship when goods are ordered, or in response to spot-market commodity prices, may simply be unable to predict their requirements far into the future. Others manufacturing goods according to preestablished production schedules, or those who are shipping to warehouses for inventory replenishment might be able to better predict their needs, given the right marketplace incentives to share the information.

In the intermodal distribution chain, the presence of third party freight forwarders prevents railroads from sharing information directly with the ultimate customer. Suppose a shipper agent knows a particular load is not immediately needed by the consignee. The agent also knows this information is valuable to the railroad, that the railroad can derive an operational benefit from knowing it. The shipper agent can force the rail carrier to share some of the benefits of revenue management by controlling information. Under this circumstance, the railroad may need to establish a differentiated pricing structure—possibly including accompanying advance purchase and nonrefundability restrictions—to induce shipper agents to reserve space earlier, to show up at the origin loading point as arranged, and to advise the carrier when the shipment is really needed at the destination.

By contrast, rail carload freight has a long history of strong price and service differentiation based on commodity. Railroads normally market their service directly to shippers and consignees located along their rail lines. The promise of more reliable service may provide sufficient incentive for railcar shippers to make reservations earlier. Shippers of low-priced commodities will find they can get the best service commitments if they are willing to make reservations early, or are willing to commit to purchase capacity on a "take or pay" basis. By comparison, leisure airline travelers generally get their best deals by booking early—and a nonrefundable fare is the equivalent of the "take or pay" shipping contract. For the most part, freight railroads already possess the information they need to differentiate rail carload service without having to offer additional price incentives.

Critical Differences Between Freight and Passenger Problems

The following sections explore critical differences between railroad passenger and freight revenue management problems. As compared to passenger implementations, freight revenue management has an "operational" characteristic, which might create organizational barriers to its successful implementation. Other differences include the ability to "bump" low priority freight at intermediate terminals; the focus of freight revenue management on developing achievable service commitments and on improving service reliability, as opposed to price discrimination; and differing service network structures and marketplace environments of the respective businesses.

The Operational Characteristic of Freight Problems

Freight applications differ from passenger applications in a very fundamental way. While passengers can board, deboard, and make connections in terminals on their own, freight has no inherent mobility—it must always be handled operationally. Therefore, any freight revenue management application takes on operational process control characteristics, focusing on management of terminal operations to ensure that the proper connections are made—as opposed to strictly maintaining a sales and

marketing focus as passenger applications do. Thus, implementation of a freight revenue management process tends to be more intrusive in daily operating practices, compared to a passenger implementation.

In rail carload freight, as shown in Figure 5, although booking lead times may be short, the traffic retention period is quite long, several days at least. Thus, particularly at intermediate terminals with little originating traffic, a railroad's ability to plan terminal operations is better than might at first appear, in spite of short booking lead times.

Figure 5: Future Workload Projection Uncertainty Traffic Retention Period

Booking Lead Time	Short		Long
Short	Short Distance Rail Passengers Trucking Airline Business Travelers	Rail Intermodal Freight	Rail Carload Freight
Long	Airline Leisure Travelers		Long Distance Rail Passengers

Still, at terminals with a lot of originating traffic, a decision whether or not to classify a shipment onto a particular outbound train must often be made with imperfect information, and once a shipment has been classified, the decision becomes difficult and costly to change. Therefore, it is important to get the decision "right the first time" as often as possible to avoid adverse service impacts, or the expense of reworking incorrect decisions later.

Fortunately, Revenue Management provides a rational means to approach this kind of decision-making, in spite of uncertainty in future demand. The only requirement is that demand must be understood at least well enough to calibrate a probability distribution. Such a distribution can be estimated based on historical experience as well as leading indicators, such as empty car orders or advance information received from connecting railroads. An exact "point estimate" of demand is not required. The fundamental reasoning underlying the revenue management approach is well explained by Elkins (1991, pg 7-8):

> If we reserve a unit of capacity (an airline seat or a hotel room or 30 seconds of television advertising time) for the exclusive use of a potential customer who has a 70% probability of wanting it and is in a market segment with a price of $100 per unit, then the expected revenue for that unit is $70. Faced with this situation 10 times, we would expect that 7 times the customer would appear and pay us $100 and 3 times he would fail to materialize and we would get nothing. We would collect a total of $700 for the 10 units of capacity or an average of $70 per unit.
>
> Suppose another customer appeared and offered us $60 for the unit, in cash, on the spot. Should we accept his offer? No; because as long as we are able to keep a long-term perspective, we know that a 100% probability of getting $60 gives us an expected revenue of only $60. Over 10 occurrences we would only get $600 following the "bird in the hand" strategy.
>
> We should never sell a unit of capacity for less than we expect to receive for it from another customer, but if we can get more for it, the extra revenue goes right to the bottom line.

This implies that management must be willing to take some calculated, short-term risks in order to maximize long term

gains. The inevitability of an occasional poor decision must be accepted. Operating performance must be evaluated on an appropriate long-term basis, not post-audited on a 20/20 hindsight basis. A fear of this kind of critical post-auditing can lead to a highly reactive, rather than proactive management style. In many instances, even current trip planning systems may be able to confidently foresee a future opportunity or problem, but it may still be difficult to get railroad operating managers to proactively act on this information.

As well, external factors such as the operating budget can inappropriately influence management, often leading to poor decisions. As one railway manager put it, "It is easy to justify running extra trains to clear out a congested terminal. It is not so easy to justify spending the money from the extra train budget ahead of time to prevent the terminal from going down."

This hesitancy to act on less-than-perfect information may stand as a significant barrier to successful railroad implementation of Revenue Management systems. The best solution is not to wait for perfect information, which is an unattainable goal; but instead to make the best use of information already available to support rational decision-making, based on the objective of maximizing expected profits. This strategy must be understood and clearly supported by top management—otherwise, system recommendations will not be followed, and implementation will fail.

Shipment Priorities, 'Bumping' and Service Reliability

Another difference between freight and passenger problems is the ability in freight to displace a lower priority shipment, even at an intermediate terminal, in favor of a newly arrived load. Airlines do not normally "bump" passengers based on the fare class of the ticket they hold nor do they routinely schedule connections in intermediate terminals on anything other than a "next flight out" basis. In railroads, this ability to "bump" shipments at intermediate terminals is very important, since a railcar or trailer may require several days to reach its destination, which provides ample opportunity for a higher priority shipment to materialize in the meantime. This problem is compounded by the typically short booking curves, which exist in the freight rail industry, as previously discussed.

In railroad carload freight, the rationale for establishing shipment priority typically has been first-come-first-served rather than any criteria based on customer needs or due date information. Even in intermodal shipping, according to Jay Hirst of Alliance Shippers (*Railway Age*, 1993, p. 60), the ability to prioritize traffic to match customer expectation is still far from ideal:

> Terminals have a bad habit of not being able to prioritize. Terminal operations just seem to allow for one set format to process trailers and containers, rather than being able to do what the customer requires . . .

The ability to establish shipment priorities is important to maintaining service reliability, because it is not always possible or cost effective to move all traffic on the first available train. In the case of traffic overflowing capacity, it is essential to make certain that cars having no remaining slack in their commitment delivery times have first access to available space. Kraft (1995) and Kwon (1994) directly link traffic volume variability to railroad freight service reliability. Following Kraft (1995, p. 28):

> Service failures can result if there is a mismatch between demand and the amount of capac-

ity provided. If demand is higher than the capacity of the train, excess cars spill over to the next day's train, unless an extra train is operated, which may not always be physically possible or economic. If demand is much lower than planned, some trains may be annulled, consolidated, or held for tonnage, leading once again to unreliable transit times. There is a direct link between the variability of customer demands and the reliability of transit time produced by a railroad freight transportation system.

Management has some latitude to allocate capacity among different customers and traffic lanes, but in the short term only within fixed limits determined by the number of locomotives in the system, train crew availability, and requirements to reposition both crews and locomotives to handle future demands. Passing siding lengths and train handling considerations determine the maximum train size, which can be operated over any route. But, there may not be sufficient locomotives to power all the trains at this maximum size.

The traditional freight railroad response to volume variability has been to annul or consolidate trains if volume is too low, or to operate extra trains or second sections (resources and budget permitting) if volume is too high. Operation of extra trains is generally not harmful unless unplanned departures create line capacity problems with excessive train delays, or throw operating resources such as locomotives or crews out of balance. Some carriers plan for extra schedule "slots" to allow for extra train operations. However, the strategy of annulling or consolidating trains has an extreme adverse effect on service reliability (Kwon, 1994). Recently, some freight rail carriers have started to emphasize "running to plan" or the operation of a fixed set of scheduled trains every day. However, even if all trains run on time every day, reliability problems may still be caused by overflowing available train capacity (Kraft, 1995).

Campbell (1996) and Kraft (1998) both propose revenue management formulations which can lead to establishment of shipment delivery "due dates" and associated penalty costs for missing these delivery targets. Once due dates and penalty costs are established, these can be used to determine shipment priorities in real time to determine which shipments should actually be loaded onto a train.

Market Structure of Railroad Freight Problems

Another difference between passenger and freight applications relates to the fundamental market structure. While airlines and passenger rail applications have a "mass market" orientation, freight railroads provide service to a relatively small number of industrial customers. Particularly in carload freight, the railroad should know its customers individually; marketing may negotiate transportation contracts, specifying unique price and service characteristics required by each customer. These contracts establish a framework for a long-term business relationship, whereas most airline and passenger rail revenue management models view the customer relationship only in terms of the current transaction.

While air travelers may be able to choose from several airlines, and freight shippers might choose from a long list of trucking companies, most likely a rail customer is directly served by only one or two railroads. For certain commodities, it might not be economical to pro-

duce a product in a certain place or ship to certain markets except by rail. Still, when a contract is signed, the customer expects that transportation will be provided at the price agreed to, and that service levels will generally fall within a contractually agreed upon range of transit time and reliability.

Although trucking companies routinely reject offered loads that don't meet their revenue management criteria (see for example, Powell et al., 1988), outright rejection of loads is more difficult for railroads than if the customer had a large number of competitive rail options. The irony is that high value loads having many competitive alternatives (where trucking is an option) are ones the rail carrier is least likely to want to reject. Still, there could be room for a railroad to negotiate delivery times for individual shipments, within the overall parameters of the governing transportation contract.

In railcar freight, since prices are contracted in advance, when the customer calls to offer a shipment, the discussion should focus on the question of when service can be provided, not at what price. The railroad does not reject any offered loads. However, the *customer* can always choose to reject service offers and ship by another mode.

Survey of Railroad Revenue Management Applications

The application of revenue management is well established in passenger railroads. Amtrak was first in the railroad industry to recognize the need for controlling seating availability by fare class in the markets they served. A milestone was reached in July 1991, when Amtrak implemented the world's first automated railroad revenue management system. On average, Amtrak realizes an additional 3% to 5% in incremental revenues from the current revenue management practices.

The primary focus of revenue management at Amtrak has been to ensure that short distance, low revenue riders do not block capacity across peak load segments, which could be sold to longer distance, higher revenue passengers. Traffic management considers the tradeoff between the revenue value of accepting a booking request, versus the expected opportunity cost of alternative future requests that may be displaced.

Amtrak's ARROW reservation system supports serial nesting of fare classes, where a higher value fare can always be sold if a lower value fare class is still open, but not "virtual" nesting based on origin-destination clustering algorithms (Smith, Leimkuhler, and Darrow, 1992). Because ARROW is a leg-based reservations system, when a fare class is closed, all markets in that class using that leg will be restricted for sale. To prevent high-valued long haul markets being entirely shut out by closing discount fare classes, traffic control is affected by specific origin destination market sales limits. This allows short distance low revenue origin destination pairs to be restricted or closed entirely, while still allowing sales for long-haul higher revenue markets in the same fare class. Hence, the Amtrak revenue management optimization must not only generate leg class authorizations, but also origin destination market class authorizations. AROW can override leg class restrictions if the market class still has availability based on market class limits.

At Amtrak, the overbooking model currently is used very conservatively and causes negligible standees. More often, standees have resulted due to other causes such as passengers presenting invalid tickets for the train they are on, or last minute equipment changes. Amtrak's current discount allocation model is based on Belobaba's (1987) "EMSR" approach. The "EMSR rule" is a heuristic rule to allocate capacity on a single flight leg by equalizing the expected

marginal revenue of each fare class.

Sabre Decision Technologies (SDT) implemented an integrated decision support system for the French National Railways (SNCF). When SNCF's long term "Railplus" and short term "Railcap" (Ben-Khedher et al., 1998) schedule development and capacity optimization systems were developed, the revenue management system was integrated with it to support intermediate stage planning for both marketing and operations. This produced a fully integrated rail decision support system for planning of train schedules, equipment allocation, pricing and revenue management of SNCF's high speed TGV train service. The Eurostar revenue management system is basically the same as the SNCF system with very similar models and also implementing virtual nesting controls. The main difference in the Eurostar system is that certain additional variable costs are included in the models.

VIA Rail Canada is controlling reservations by origin and destination in order to maximize revenues (Berwick and Therrien, 1997). VIA's demand forecasting system is integrated with both their capacity allocation and their revenue management system. The capacity allocation system considers the marginal cost and the revenue generated from additional cars, and demand characteristics that may vary over different legs of the train cycles. (Cordeau, Desaulniers, Lingaya, Soumis and Desrosiers, 1998).

For freight applications, the literature on shipment routing and scheduling is well developed. The Less-than-Truckload (LTL) trucking network design problem, and certain air cargo problems are very closely related to the railroad shipment-scheduling problem. Powell implemented interactive optimization systems for Ryder/PIE (Powell and Sheffi, 1989) and later for Yellow Freight (Braklow, Graham, Hassler, Peck, and Powell, 1992) to find opportunities to bypass break bulk facilities, and performed interactive "what if" analysis in real time on various shipment routing strategies. Barnhart and Sheffi (1993), Farvolden, Powell and Lustig (1993), Jones, Lustig, Farvolden and Powell (1993) all solve similar LTL shipment routing problems. Kwon (1994) outlines a railroad freight car scheduling problem formulated in path variables, and solved it using a standard column generation approach. Kraft (1998) proposed a "Dynamic Car Scheduling" process using a customized dual adjustment heuristic to solve an integer multicommodity network flow problem.

However, all these routing and scheduling models have *cost minimizing* formulations—they do not take revenues, or in some cases even delivery time commitments into account. A true revenue management application should be based on a *profit maximizing* formulation, and should also address "load selection" in some manner, not just optimize empty equipment repositioning. Examples of research meeting these criteria include Powell's (1988) work on optimal load selection for full truckload carriers, Ph.D. dissertations by Nozick (1992), Campbell (1996) and Kraft (1998) on intermodal and railcar revenue management, and a survey paper by Kaslingam (1996) on air cargo revenue management.

Although revenue management has become commonplace in passenger railroading, no freight railroads are yet known to have implemented it. However, railroads' direct competition—the trucking industry—has done so, which may account for some of trucking's recent competitive success against railroads. This is one reason why understanding Powell's (1988) work should be important to the railroad industry.

Powell's (1988) work is also relevant because the need to account for empty equipment repositioning makes freight revenue management more complicated

than its passenger counterpart. On an airline between the same city pairs, full fare business travelers will *always* be preferred over leisure travelers paying a discounted fare. The airline would like to fill the airplane exclusively with business travelers, if it could. This leads to the airline concept of "nested" fare classes, where a full fare ticket can always be sold to a business-class traveler even if the predetermined business class allocation happens to be sold out.

This strict hierarchical relationship may not always hold true in the case of freight. The first few vehicles into a zone have a high probability of finding backhauls, representing highly valuable business. Additional loads become less profitable, such that some other business may take priority at a certain point. Powell's (1987) system develops and maintains opportunity cost information for equipment supply, in a continuous real-time mode. Although Powell (1987) never labeled it as such, his model clearly implements a special case of bid-price revenue management, as applied to a freight transportation problem in the trucking industry.

A comprehensive revenue management system for rail freight should address both equipment supply and also the allocation of space on trains. Powell's (1988) model addresses equipment supply but does not address train capacity utilization. Campbell (1996) and Kraft (1998) address train capacity utilization but not equipment supply. Nozick's (1992) model includes both equipment supply and train capacity constraints, but assumes deterministic demand. Nozick's (1992) model has been used to address the efficiency and organization of intermodal drayage, and the effects of a traffic priority system on fleet sizing and intermodal car and trailer fleet management.

Campbell (1996) researched the application of revenue management techniques in railroad intermodal applications, focusing on the allocation of railcar capacity to origin-destination shipping lanes. His research extends techniques originally developed for fixed capacity networks, adapting Belobaba's (1987) "EMSR heuristic" to apply to flexible capacity networks as well. Campbell's dissertation uses an origin-destination rather than leg-based definition of "EMSR." Campbell's proposed booking control mechanism is primarily based on origin destination market sales limits, very similar to Amtrak's system.

Kraft (1998) focuses on the process of developing appropriate and achievable delivery time appointments using a "bid price" revenue management approach, and then managing the operational service delivery process to conform to these commitments. The concept for scheduling delivery appointment times is modeled after current motor carrier industry practice, where a delivery appointment time is established for each shipment at the time the initial order is placed.

Kaslingam (1996, p. 43) proposed a chance-constrained formulation for air cargo discount allocation. Interestingly, some early airline overbooking and seat allocation models also proposed a "chance constrained" approach (see Charnes and Cooper, 1963) . Kraft, Oum, and Tretheway (1986) suggest that seat allocations be set "by choosing a probability level for seating all full fare passengers. For example, the airline might choose to allocate seats to full fare passengers such that 95% of the time all full fare passengers will be accommodated . . . "

In passenger applications, overbooking conditions are usually resolved on a *voluntary* basis, in which case the expected cost of overbooking constraint violations can be precisely and accurately quantified. Only rarely must an overbooking condition be resolved on an involuntary basis, which leads to a

typically small-expected penalty for lost customer goodwill. Thus a penalty cost can be assigned for overbooking violations and the optimization program is allowed to choose the overall best level, by trading off the penalty cost versus expected revenue gain.

However in freight applications, capacity constraint violations must generally be resolved by involuntarily "bumping" excess shipments off the train, truck or airplane. The decision which freight to forward versus which freight to hold back is typically made by the carrier, seldom in consultation with the customer. The consignee typically receives notification of a "bumping" decision or missed connection after-the-fact, if at all. Even if a freight carrier must pay a financial penalty for late deliveries, the cost of a service failure will still be dominated by such "soft" considerations as lost customer goodwill. Since the cost of constraint violation is hard to quantify, then a chance-constrained approach may be both a more direct and honest treatment for freight applications.

Future Opportunities and Research Needs

Since the primary focus of most railroad revenue management problems is on origin destination traffic management and not on price discrimination, the traditional airline leg-based "EMSR" approach really does not address well the central issue of railroad revenue management. A bid price methodology is a very attractive basis for both passenger and freight railroad revenue management, because it effectively addresses many of the common core requirements of these applications.

(1) For long distance passenger trains, a bid price approach is highly attractive because of the large number of legs traversed by the average intercity passenger, as well as the high number of interconnecting trips between trains.

(2) For short distance, high speed services, real time, frequent updates provided by a bid price system allows the application to cope with extremely short booking curves, and high cancellation and rebooking rates.

(3) For freight applications, the bid price approach provides an intuitive means for developing achievable delivery due dates and clear movement priorities for each shipment, using a modified shortest path algorithm, as in Kraft (1998).

The beauty of the bid price approach lies in its simplicity. A bid-price system can readily accommodate any kind of booking arrival pattern, as opposed to an "EMSR"-based system that tends to "overprotect" space for the higher value fare classes. Updating the bid prices frequently eliminates the need for any nesting assumptions. A "safety net" function using "triggers" (Phillips, 1994b) or other similar control mechanism can be used to detect when a multiple group booking request should be flagged for manual intervention or other special handling in the system—including consideration of whether additional train capacity should be provided.

Talluri and Van Ryzin (1998) discuss strengths and weaknesses of the bid price control method in general, and present several conditions under which bid price controls might lead to suboptimal decision-making. However, they conclude that a bid price control scheme is "close to being globally" optimal, and are continuing to research improved methods for developing more accurate bid prices.

To succeed, freight revenue management implementations should be inte-

grated within a higher-level capacity management framework to ensure that operating resources are properly positioned to meet anticipated demands—as SNCF and Sabre Decision Technologies have already done for passenger rail applications. Such a system would be demand forecast-driven and would address global issues of resource management, such as locomotive and crew management, availability of rail line and terminal schedule slots, and empty equipment repositioning. Implicitly assumed is the ability to run trains on time, and that connections can be made in terminals as scheduled. For freight carriers, this would be facilitated by moving towards a preplanned, scheduled train operation.

Most intercity passenger rail carriers have already implemented some form of revenue management. It would appear to be relatively straightforward to extend revenue management to intermodal freight services, where trains operate on strict schedules and terminal operations are more flexible than for railcars. Although rail carload service potentially stands to gain the most, fundamental improvement in classification yard and train-operating discipline is necessary before revenue management can succeed there.

From a research perspective, an integrated framework is needed to incorporate both empty equipment allocation and revenue management of train capacity in a single model. Nozick's (1992) dissertation probably comes closest right now to addressing this need, but her approach needs to be generalized to handle stochastic, rather than deterministic demand.

In theory, rail carriers should provide capacity so they can sell space on trains to generate revenue and make a profit. In the past, however, decisions to provide train capacity have often been based on cost minimization, not profit maximization. Integrating capacity management into a revenue management framework provides, for the first time, an ability to understand the revenue, as well as the cost implications of a decision to provide capacity, and the ability to incorporate that information into real time decision-making.

Seamless Availability

Historically, airline carriers have controlled reservation system availability by sending "batch updates" of booking class/leg departure availability levels at various intervals. More recently, some reservation systems have provided the capability of message-switching booking requests in "real time" to the carriers. In some sense, this capability has been around since the early days of the reservation systems. However, initially the booking requests were relayed to the carriers "one leg at a time" with no easy way for the carrier to determine whether or not the requests came from the same passenger or from two different passengers. For example, a booking request for:

UA 178 JFK - ORD C-Class, connecting to UA 150 ORD - SFO C-Class

would be received by the carrier as two separate requests, one for the JFK-ORD leg and one for the ORD-SFO leg with no way to determine that this was a single request rather than two requests. This effectively made Origin-Destination based revenue management impossible.

With a Seamless Availability capability in the reservations system, the entire booking request, including all connecting legs, is transmitted in a single message to the carrier. This enables the carrier to determine the full requested itinerary and manage accordingly. This capability has been available in some form from most of the reservation systems for several years, and all the major systems have or are now implementing Seamless Availability. The implementa-

tion of Seamless Availability has been a major incentive for airlines to develop true origin-destination revenue management capabilities.

Reservation system limitations have often been relevant outside the airline industry—particularly in the hotel and rental car industries. For example, most early hotel/rental car reservation systems did not allow availability controls to be established by length-of-stay/length-of-rental. Since control by length-of-stay/length-of-rental is an extremely important aspect of revenue management in these industries, this was not a trivial limitation. Fortunately, these reservation systems have been somewhat easier to change than the airline systems and most major systems have been brought up to date.

References

Barnhart, C. and Sheffi, Y. (May 1993) A Network-Based Primal-Dual Heuristic for the Solution of Multicommodity Network Flow Problems, *Transportation Science*, 27(2), 102-117.

Belobaba, P.P. (1987) Airline Yield Management: An Overview of Seat Inventory Control, *Transportation Science* 21, 63-73.

Belobaba, P.P. (1989) Application of a Probabilistic Decision Model to Airline Seat Inventory Control, *Operations Research* 37(2), 183-197.

Ben-Khedher, N., Kintanar, J., Queille, C. and Stripling, W. (January-February 1998) Schedule Optimization at SNCF: From Conception to Day-of-Departure, *Interfaces* 28(1).

Berwick, J. and Therrien, C. (1997) Revenue Management at VIA Rail Canada, *Transportation: Emerging Realities Forum* presentation.

Braklow, J.W., Graham, W.W., Hassler, S.M. Peck, K.E. and Powell, W.B. (January-February 1992) Interactive Optimization Improves Service and Performance for Yellow Freight System, *Interfaces* 22(1), 147-172.

Campbell, K.C. (1996) *Booking and Revenue Management for Rail Intermodal Services*, Ph.D. Dissertation, Department of Systems Engineering, University of Pennsylvania, Philadelphia, PA.

Charnes, A. and Cooper, W.W. (1963) Deterministic Equivalents for Optimizing and Satisficing Under Chance Constraints, *Operations Research* 11, 18-39.

Cordeau, J.F., Desaulniers, G., Lingaya, N., Soumis, F. and Derosiers, J., (1998) *Simultaneous Locomotive and Car Assignment at VIA Rail Canada*, Technical Report, Les Cahiers du GERAD G-98-61, Ecole des Hautes Etudes Commerciales, Montreal, Canada.

Elkins, S. (September 1991) *The Basics of Yield Management by IDeas*, Integrated Decisions and Systems, Inc., 3500 Yankee Drive, Suite 350, Eagan, MN 55121.

Farvolden, J.M., Powell, W.B. and Lustig, I.J. (July-August 1993) A Primal Partitioning Solution for the Arc-Chain Formulation of a Multicommodity Network Flow Problem, *Operations Research* 41(4), 669-693.

Jones, K.L., Lustig, I.J., Farvolden, J.M. and Powell, W.B. (1993) Multicommodity networtk flows: The Impact of formulation on decomposition, *Mathematical Programming* 62, 95-117.

Kaslingam, R.G. (1996) Air Cargo Revenue Management: Characteristics and Complexities, *European Journal of Operational Research* 96, 36-44.

Kraft, D.J., Oum, T.H., and Tretheway, M.W. (1986) Airline Seat Management, *Proceedings of the Transportation Research Forum*, 27(1), 340-348.

Kraft, E.R. (1995) The Link Between Demand Variability and Railroad Service Reliability, *Journal of the Transportation Research Forum*, 34(2), 27-43.

Kraft, E.R. (1998) *A Reservations-Based Railroad Network Operations Management System*, Ph.D. Dissertation, Department of Systems Engineering, University of Pennsylvania, Philadelphia, PA.

Kwon, O.K. (1994) *Managing Heterogeneous Traffic on Rail Freight Network Incorporating the Logistics Needs of Market Segments*, Ph.D. Dissertation, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA.

Nozick, L.K. (1992) *A Model of Intermodal Rail-Truck Service for Operations Management, Investment Planning and Costing*, Ph.D. Dissertation, Department of Systems Engineering, University of Pennsylvania, Philadelphia, PA.

Phillips, R.L. (September 1994a) *Using Itinerary and Point-of-Sale Informaton to Improve Profitability*, Decision Focus Inc. working paper.

Phillips, R.L. (October 1994b) *State-Contingent Yield Management*, Decision Focus Inc. presentation.

Powell, W.B. (1987) An Operational Planning Model for the Dynamic Vehicle Allocation Problem with Uncertain Demands, *Transportation Research* B 21(3), 217-232.

Powell, W.B., Sheffi, Y., Nickerson, K.S., Butterbaugh, K. and Atherton, S. (January-February 1988) Maximizing Profits for North Amiercan Van Lines Truckload Division: A New Framework for Pricing and Operations, *Interfaces* 18(1), 21-40.

Powell, W.B. (November 1989) A Review of Sensitivity Results for Linear Networks and a New Approximation to Reduce the Effects of Degeneracy, *Transportation Science* 23(4), 231-243.

Powell, W.B., Sheffi, Y. (January-February 1989) Design and Implementation of an Interactive Optimization System for Network Design in the Motor Carrier Industry, *Operations Research* 37(1), 12-27.

Railway Age (October 1993) What's Needed to Keep Intermodal Growing, *Railway Age* 194 (10), 60.

Smith, B.C., Leimkuhler, J.F. Darrow, R.M. (January-February 1992) Yield Management at American Airlines, *Interfaces* 22(1), 8-31.

Talluri, K. and Van Ryzin, G. (November 1998) An Analysis of Bid-Price Controls for Network Revenue Management, *Management Science* 44 (11 - part 1 of 2), 1577-1593.

Williamson, E.L. (1992) *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*, Ph.D. Dissertation, Department of Aeronautics and Astronautics, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, MA.

JTRF

Edwin R. Kraft is responsible for costing systems development and new computer systems for Mail and Express operations at Amtrak. He holds a Ph.D. in Systems Engineering from the University of Pennsylvania.

Bellur N. Srikar is the project manager at i2 Technologies, Inc., managing future implementation of supply chain optimization for a major multinational pharmaceutical corporation. He received a Ph.D. in Operations Management and an M.B.A. in Management Science from the University of Texas at Austin, and a B.S. from Indian Institute of Technology in Bombay, India.

Robert L. Phillips is Chief Technology Officer with Talus Solutions in Mountain View, CA. He is responsible for creating new business opportunities by identifying and developing products and services that deliver high customer value in all areas of pricing and revenue management. Phillips received a Ph.D. in Engineering-Economic Systems from Stanford University and holds undergraduate degrees in Mathematics and Economics from Washington State University.