Value of High-Quality Service: How Should the ARES-Equipped Railroad Operate?

MICHAEL E. SMITH AND RANDOLPH R. RESOR

The North American railroad industry is beginning to implement new, advanced train control technologies that will significantly change railroad operations. Collectively, these technologies are referred to as the Advanced Train Control Systems (ATCS). Burlington Northern Railroad's specific version of ATCS is called the Advanced Railroad Electronics System (ARES). In an extensive operations analysis, Burlington Northern found that it could use ARES to greatly reduce costs or improve service, or a little of both. To determine the optimal course of action, the railroad conducted a market study to determine the value of better service. The study results indicated that customers were willing to pay much more for small improvements in service. Burlington Northern's market managers disputed this, believing that better service would not significantly increase prices or market share. Meanwhile, ARES operations analysis indicated that operating improvements could be targeted very precisely using the new technology. Travel time improvements could be allocated at will among various classes of trains. Given these results, Burlington Northern should concentrate its initial implementation of ARES functionality on reducing cycle times for bulk commodity trains. This will result in the need for fewer coal sets to move a given amount of coal. Then, the railroad will receive the certain payoff of reduced assets as opposed to the uncertain payoff of increased revenue. As implementation proceeds, Burlington Northern should use ARES capabilities to test the value of improving service. Then implementation strategies can be adjusted to improve the outcome.

The introduction of new technology into an existing operation often provides the opportunity to improve the processes making up the operation. Advanced Train Control Systems (ATCS) have the potential to improve railroad operations. Through investments in this new technology, railroads have the opportunity to lower costs and improve service simultaneously.

This happy state of affairs does not frequently present itself. Usually, providing an improved level of service requires that more resources (more cost) be used in the operation. An existing process can be used to translate input resources into desired outputs. More output, or higher-quality output, requires more input resources.

Now, suppose the process is changed. A new set of assets, better equipment, is used to produce the output. Now, less input is required to produce the same output. An example is a tailor using a sewing machine instead of hand stitching to make clothes. The tailor can now produce a garment in an hour when it used to take a day. This is a cost-reducing approach to using the new equipment. There is, however, a quality improvement approach. With a sewing machine, the tailor may be able to spend all day on the garment and produce one of higher quality.

Which approach is best? Should the machine be used to produce the output using fewer input resources (and reduce cost)? Or should it be used to produce the same quantity, but higher-quality, output (and improve quality)? These same questions are equally valid for ATCS.

The Burlington Northern Railroad is struggling with the decision as to whether it should install its version of ATCS, the Advanced Railroad Electronic System (ARES). In the process of making that decision, the railroad analyzed both the cost reduction and quality improvement potential of ARES. The way in which ARES is used and designed depends on the relative values of improving quality and reducing cost.

Two alternatives present themselves for the use of ARES on Burlington Northern—one that improves quality and reduces cost, and one that only reduces cost. The different methods of looking at the value of higher quality need to be weighed along with the impact on design and operations that the alternative ways of using ARES will impose.

DESCRIPTION OF ARES

In the process of developing ARES, Burlington Northern has spent considerable effort in determining how the system affects existing processes from the top down. This discussion presents a top-down flow of the control and process envisioned by ARES and then describes the individual components used in that process.

For a top view of ARES, consider Figure 1. This shows a hierarchical planning and command structure for a railroad operation. First, a planning and scheduling group establishes a schedule for operating the system. Depending on the level of sophistication, this could include train schedules, terminal schedules, equipment rotation plans, and maintenance schedules.

The first component of ARES is shown just under the dotted line in Figure 1. This top-level component is called the strategic traffic planner (STP). The function of the STP is twofold. First, it will translate the train schedule into time goals and priorities for trains handled by each dispatcher. Second, if schedules cannot be met, the STP will assist system

M. E. Smith, ARES Traffic Planning Systems, Burlington Northern Railroad, 9401 Indian Creek Parkway, Overland Park, Kans. 66201. R. R. Resor, Costing and Pricing, Zeta-Tech Associates, P.O. Box 8407, Cherry Hill, N.J. 08002.

FIGURE 1 Hierarchy of ARES command and control functions.

control personnel in selecting alternatives that do the least damage to the network.

The second component of ARES is shown as the next level down in the hierarchy. This component is called the tactical traffic planner (TTP). The TTP assists the dispatcher in finding the most efficient meet-and-pass plan for each train in the dispatcher's territory. Further description of how the STP and TTP are intended to work can be found in Smith and Resor (*1*).

Once the TTP has generated a meet-and-pass plan, it passes that information to a computer-aided dispatching tool that automatically generates authorities to vehicles in the field. Using global positioning system (GPS) location technology and digital data radio, field vehicles are continuously tracked to ensure that they are operating according to plan.

If a train is not operating according to plan, the engineer will be prompted to speed up or slow down, as appropriate, by the display on board. If the train is falling behind schedule, the engineer will be prompted to give a reason for the delay. If the on-board computer determines that the train will not arrive at the next significant event location (e.g., meet location or terminal) within the time required, a signal will be sent to the TTP requesting a new plan. This provides a closed-loop process that ensures the system responds to external disturbances in an optimal way.

The response of the system depends greatly on how "optimal" is defined. The TTP can provide meet-and-pass plans that minimize travel time, minimize deviation from schedule, minimize either of those with different weights for each train, minimize fuel, or any combination of the preceding. Which should be done?

When Burlington Northern evaluated ARES, it was assumed that the system would be used to minimize travel time, weighted by the value of each train. The results showed that this approach could reduce the cycle time of bulk trains and increase the reliability of carload freight trains. That left the task of determining how much each of these things was worth.

The worth of decreased cycle time for bulk trains was easy to understand. It simply amounted to the savings generated by requiring less equipment sets to move the same amount of commodity. The worth of increased reliability for carload freight was more difficult to determine. It depended on how much better service was worth to the customer and on how much of that value the railroad may be able to extract through increased marketshare or price.

To address this difficulty, Burlington Northern evaluated ARES with two different objective functions for the TTP. In the first objective function, the evaluation minimized train travel time weighted by the cost of train delay for each train

type. This train-delay cost was based on both the value of the railroad equipment and the value of delay to the lading (*2*). This approach assumes implicitly that reliability to the customer is important and valuable. In the second objective function, the evaluation minimized travel time on bulk trains with no increase in the delays experienced by other trains. This approach implicitly assumes that additional reliability to the customer has no value to the railroad.

Which approach is right? How should ARES be used when it is put into service? That depends on how much reliability can be improved by reducing train delay and how much that is worth to the customer. Further analysis pointed toward the answers. Simulations showed how the operation would improve with the two different objective functions and the value of improved service to the customer was estimated.

THE IMPACT OF ARES ON SERVICE

ARES, as designed by Burlington Northern and Rockwell International, will incorporate a sophisticated optimization program that will provide dispatchers with a mathematically "best" dispatching plan. However, users must specify the inputs—such as the cost of train delay—that are used to determine the optimum. Different train-delay costs and even different delay functions (e.g., quadratic versus linear) can produce radically different dispatching plans.

The meet/pass planning algorithm within ARES is intended to devise an optimal meet/pass plan for each "planning line" (segment of railroad). This plan will minimize a weighted combination of fuel consumption and running time, taking into account the differing priorities of different trains. It is important to note that ARES will use weighted priorities rather than "cardinal" train priorities. That is, rather than being ranked by category, with trains in the highest category always receiving priority over lower-ranked trains, trains will be assigned varying delay costs. Thus, an intermodal train may sometimes be delayed for a lower-priority train if by doing so the total delay cost is minimized.

Computer-aided meet/pass planning is expected to yield large benefits for Burlington Northern when used in conjunction with the real-time location information provided by ARES. Areas of benefit include an increase in line capacity due to less time lost in making meets, improved reliability, and fuel savings through the avoidance of unnecessary stops and through "pacing" of trains so that they arrive at meet points exactly on schedule, rather than proceeding at maximum speed, arriving early, and having to sit and wait for an opposing train. However, these benefits cannot all be realized simultaneously; that is, a railroad cannot achieve increased capacity, improved reliability, and reduced fuel consumption together.

To evaluate the effect of this optimization program on Burlington Northern operations, a series of simulations was carried out using actual train movement data. Although the optimization model used (the SChedule ANalyzer or SCAN, developed for the University of Pennsylvania) is not identical to the model developed for use in ARES, it was thought that SCAN would provide a good approximation of the benefits to be expected from the use of dispatching optimization within ARES.

Data used in this study were originally gathered in spring 1988 from dispatchers' sheets and centralized traffic control (CTC) on switch (O/S) reports for 12 "lanes" covering a total of about 3,500 route-miles. They include a mix of CTC, automatic blocking system (ABS), and dark territory. On each lane, train movement data were gathered for a period of about 24 hours; on lanes with light volume, the period was longer, while on lanes with very heavy volume the period was as short as 16 hours. The number of trains on each lane varied from 7 (Madill to Irving) to 45 (Alliance to Edgemont). A total of 846 trains were included in the analysis.

To use the SCAN model, several data elements are required:

- Scheduled, or desired, running time for each train;
- Location and duration of all delays to trains;
- A delay cost for each train type;
- Route topology, including speed limits, grades, and siding locations;
- An estimated unconstrained (minimum feasible) running time for each train; and
- A cost per gallon of diesel fuel.

Train-delay costs were developed with the help of Burlington Northern accounting and marketing personnel and the Association of American Railroads. Burlington Northern calculates an hourly ownership cost for each piece of railroad-owned equipment; to this cost was added an estimated value of the lading carried. To estimate lading value, four broad categories of trains were created:

- Loaded bulk (coal, grain, ore),
- Empty bulk (empty return movements),
- Mixed freight (carload traffic), and
- Intermodal (trailers and containers).

For each category, an "average" train was defined. Table 1 shows the cost per hour of equipment, of lading, and the total cost used in the simulations.

Fuel cost was assumed to be 50 cents per gallon in all cases. (The analysis predated the recent upheaval in the Middle East.) Energy consumption for each train was calculated from gross weight, horsepower, and the route topology, and was then converted to fuel consumption by the use of appropriate factors.

Delays, along with train consist information, were obtained from dispatching records. Route characteristics were taken from track charts and timetables. Minimum feasible running times were estimated by use of Burlington Northern's train performance simulator.

The most difficult problem was the development of schedules for all trains. Many trains on Burlington Northern do not have schedules, and many that do will often run hours ahead of or behind scheduled times. For the purposes of this analysis, then, the desired schedule was assumed to be the same as the actual time each train operated. Actual running time was further assumed to be an upper bound; in determining the benefits of computer-aided dispatching, the first analysis tried to better the actual running times of all trains in the lane. Later, in a sensitivity analysis, running times of certain train types were held constant. These results will be discussed later.

Table 2 shows the results of the analysis by train type. Two kinds of mixed freights have been defined here; "priority" freights carry more time-sensitive commodities and have more locomotive horsepower assigned per ton than "secondary" freights.

These results clearly show that the greatest improvement in both mean running time and in variability of running times occurs for low-priority trains. This finding seems intuitively correct; a normal human response in situations where many different items must be considered simultaneously is to handle the most important ones first. Human dispatchers are simply letting the low-priority trains sit until there is time for them.

Table 3 summarizes the percentage changes in mean travel times by train type that are expected to occur with the installation of ARES.

When the results of this optimization were considered, a question arose about the different ways in which optimization might be used. The initial focus had been on minimizing total travel time for all trains. But suppose for a moment that the high-priority trains were already running as fast as market conditions required. To put it another way, let us say that there is no additional revenue or marketshare to be gained from shorter transit times for intermodal and mixed freight trains. Can running times of bulk commodity trains be reduced even further?

Most coal on Burlington Northern moves under contract in fixed annual volumes. The cycle time (round-trip time plus loading and unloading time) is known approximately for each coal movement. Thus, although these trains do not have schedules as such, equipment needs are determined by the number of train sets (cars plus locomotives) required to move the contractual coal volume. The longer the cycle time, the more equipment is required.

Burlington Northern moves a great deal of coal, and Burlington Northern and shippers have dedicated a large fleet of cars and locomotives to this service. Significant reductions in the running times of coal trains could reduce equipment requirements very substantially. Potential savings are very large.

TABLE 1   AVERAGE DELAY COST BY TRAIN TYPE

TRAIN TYPE	LADING DELAY COST ($/train hour)	EQUIPMENT DELAY COST ($/train hour)	TOTAL DELAY COST
Loaded Bulk	10.37	172.00	182.37
Empty Bulk	0.00	172.00	172.00
Mixed Freight	35.98	127.00	162.98
Intermodal	136.47	130.00	266.47

TABLE 2   COMPARISON, MEAN, AND STANDARD DEVIATION OF
RUNNING TIME FOR ALL TRAINS BY CLASS (IN MINUTES PER TRAIN)

TRAIN TYPE	BASE CASE		OPTIMIZED CASE	
	MEAN	STANDARD DEVIATION	MEAN	STANDARD DEVIATION
Bulk	253.96	125.39	190.95	102.37
Intermodal	196.09	114.60	175.61	105.77
Priority Freight	203.71	149.96	163.31	120.57
Secondary Freight	245.90	132.29	179.13	94.08
All Trains	227.52	130.76	179.37	105.08

TABLE 3   SUMMARY OF CHANGES IN
MEAN TRAVEL TIMES—BASE CASE

TRAIN TYPE	TRAVEL TIME CHANGE
Intermodal	(10.4%)
Priority Freight	(19.8%)
Secondary Freight	(27.1%)
Bulk	(24.8%)
All Trains	(21.1%)

SCAN allows the user to specify delay-cost functions for each class of trains. In the original analysis, a large cost penalty was assigned to lateness (running time exceeding the actual time each train took to traverse a lane), while a linear savings was assigned to earliness. Thus, savings could be achieved by running trains ahead of schedule.

A second and more limited analysis was carried out in which a large penalty was assigned to intermodal and mixed freight trains for both lateness and earliness. Thus, the SCAN model tended to try to run these trains as close to schedule as possible. For bulk trains, by contract, a linear benefit was assigned to earliness, just as in the original analysis. By doing this, it was hoped that all the benefits of optimization on bulk trains could be conferred.

The results of the analysis, shown below, were not surprising.

Train Type	Travel Time Change (%)
Bulk	32.9
All other	0.9

Although the total benefits were less in this constrained case (as might be expected), the net reduction in travel time for the bulk trains was nearly 33 percent—without any penalty to the mixed freights or intermodal trains.

A one-third reduction in the running time of every coal train on Burlington Northern will produce a very substantial benefit in terms of reduced equipment purchase requirements. Clearly, dispatching optimization within an ATCS can produce very substantial benefits. How the benefits are allocated among traffics will be determined by the perceptions of railroad managers regarding the value of service. Faster and/or more reliable schedules can be operated for high-priority trains, or for all trains, or alternatively the benefits can be taken entirely in the form of reduced cycle times (and therefore reduced equipment requirements) for the low-priority trains. There is a tendency to regard service quality improvements as a "soft" benefit, while savings in equipment requirements are a "hard" benefit. Historically, railroads have favored the hard benefits—and cost minimization—over the soft benefits of service quality improvements.

VALUE OF IMPROVED SERVICE

It became obvious from the operations analysis that the benefits of ARES, as well as the best way to use it, would depend greatly on the value of good service to the customer. Therefore, the marketing department on Burlington Northern was consulted in estimating this value. Based on the recommendation of the marketing department, the John Morton Company (JMC) was retained to perform the study.

JMC used a method called conjoint analysis to interview customers and map their preferences. This analysis method is not described here; there is sufficient description in Johnson (3) and in Johnson and Squeo (4). Briefly, conjoint analysis is a very sophisticated, computer-based interview technique that ensures that the respondent's preferences have been accurately mapped. The one weakness of the technique is that it maps what customers say they will do, not what they actually do (more on this later).

With the assistance of the marketing department, five commodities were selected for study: pet foods, aluminum, plastics, paper, and tires. The marketing department believed that these commodities were representative of the carload freight market.

The marketing department was then asked to assist in defining service variables that should be measured in the survey of the customers. Although there were nine service variables defined, only the three most important are presented here:

1. Reliability of cargo delivery—the percentage of time that a loaded car arrives at the customer's dock within the time window desired by the customer;
2. Reliability of empty equipment delivery—the percentage of time that a customer's request for an empty car is satisfied with an acceptable car within the time window desired by the customer.
3. Dock-to-dock transit time—the time required for the shipment to move from the shipper's dock to the consignee's dock.

In their survey of the customers, JMC intended to develop estimates of elasticity. That is, for each service dimension, how much marketshare gain could be expected from a 1 percent improvement in that dimension? Also, JMC estimated price elasticity. That is, for each 1 percent increase in price, how much marketshare would be lost? From these two numbers, service-price cross elasticities could be estimated. That is, for each 1 percent improvement in service, how much can price be increased without losing marketshare?

The results of the JMC study are shown in Table 4. These results were quite startling. If the customers have revealed what they really would do, there is tremendous potential for increased revenue from even small increases in service levels. For example, if the reliability of cargo delivery is improved by just 1 percent, these surveyed customers say they would be willing to absorb an average price increase of 4 percent. This would imply that ARES should be used primarily for improvements in reliability, not reductions in operating cost.

The marketing department at Burlington Northern was skeptical that these elasticities were truly representative of customer behavior. Sure, the customers may say they will do that, but will they really? And even if they were willing to pay extra for improved service, would they be sure enough that the service had actually improved? And, were we perceptive enough to capture all that a customer would be willing to pay when it came time to negotiate the price?

These nagging doubts led to a search for studies that had been done based on data representing actual customer choices as opposed to customers' stated preferences.

Kansas State University (KSU) had performed some analyses in the 1970s that used regression techniques to correlate rail and truck services to actual customer choices. After some discussion, we decided that insufficient data existed to support KSU's efforts to prepare similar models specific to Burlington Northern. However, KSU did provide their estimates of service elasticities from their previous studies (see Table 5).

But KSU did not provide price elasticities or cross elasticities. These price and cross elasticities were estimated based on an assumed range for price elasticity that will be discussed later. (Assumed quantities are shown in the table as italics.) The KSU numbers are based on a different definition of service. They defined service as car-miles per car-year. The presumption is that if the railroad is providing better service, cars will turn faster. KSU used this surrogate for service because there was no better measure from publicly available data.

Because KSU was unable to develop a Burlington Northern–specific analysis, and because their existing analysis used a questionable definition for service, Burlington Northern's market managers were interviewed and asked to forecast the market gains that could be achieved from improved service. The results of that survey are shown in Table 6.

TABLE 4   RESULTS OF JOHN MORTON COMPANY SURVEY

COMMODITY	SERVICE ELASTICITIES			PRICE ELASTICITY	CROSS ELASTICITIES		
	CARGO RELIAB.	EMPTY RELIAB.	TRANSIT TIME		CARGO RELIAB.	EMPTY RELIAB.	TRANSIT TIME
Paper	6.0	3.3	-1.1	-1.1	5.5	3.0	-1.0
Pet Food	6.9	2.8	-1.4	-1.5	4.6	1.9	-0.9
Aluminum	4.3	1.9	-1.3	-1.3	3.3	1.5	-1.0
Plastics	4.7	2.1	-0.9	-1.6	2.9	1.3	-0.6
Tires	6.2	2.3	-1.6	-0.9	6.9	2.6	-1.8
AVERAGE	5.3	2.5	-1.2	-1.3	4.1	1.9	-0.9

TABLE 5   RESULTS OF KANSAS STATE UNIVERSITY RESEARCH

COMMODITY	SERVICE ELASTICITY	PRICE ELASTICITY		CROSS ELASTICITY	
		LOW	HIGH	LOW	HIGH
Food Products	0.7	1.3	3.0	0.2	0.5
Tobacco Products	1.4	1.3	3.0	0.5	1.1
Textile Products	1.3	1.3	3.0	0.4	1.0
Lumber & Wood Products	0.6	1.3	3.0	0.2	0.5
Furniture	1.1	1.3	3.0	0.4	1.2
Paper Products	0.6	1.3	3.0	0.2	0.5
Chemicals	1.0	1.3	3.0	0.3	0.8
Stone, Glass & Clay	1.5	1.3	3.0	0.2	1.2
Primary Metal Products	1.8	1.3	3.0	0.6	1.4
Fabricated Metal Products	3.0	1.3	3.0	1.0	2.3
Non-Electrical Machinery	4.3	1.3	3.0	1.4	3.3
Electrical Machinery	1.7	1.3	3.0	0.6	1.3
TOFC/COFC	1.2	1.3	3.0	0.4	0.9

TABLE 6 RESULTS OF MARKET MANAGER SURVEY

BUSINESS UNIT	SERVICE ELASTICITY	PRICE ELASTICITY	CROSS ELASTICITY
Industrial Products	0.1	3.0	0.03
Forest Products	0.2	7.0	0.03
Food & Consumer	0.5	1.2	0.4
Automotive	0.04	infinite	0
Agricultural	0.1	3.0	0.03

TABLE 7 ELASTICITY ESTIMATES

SOURCE	SERVICE ELASTICITY		PRICE ELASTICITY		SERVICE-PRICE CROSS ELASTICITY	
	LOW	HIGH	LOW	HIGH	LOW	HIGH
John Morton	4.3	6.9	0.9	1.6	2.9	6.9
Kansas State	0.6	4.3	1.3	3.0	0.2*	3.3*
Market Mgrs.	0.01	0.5	1.2*	infinite*	0	0.4

The market managers were asked to provide service elasticities and cross elasticities. The price elasticities were derived by dividing the service elasticity by the cross elasticity. The definition of service used for this exercise was equivalent to the one used by JMC, that is, the percentage of shipments that arrive in the customer's desired time window. The high estimate on price elasticity is infinite. This implies that Burlington Northern could capture the entire market for that commodity by a very small cut in its price. We felt that this might be unrealistic.

The most common estimate of price elasticity from the market managers was 3.0. This value was then used as the high estimate for price elasticity to derive implied figures in the KSU work. The low value used for that purpose was 1.3, equal to the mean price elasticity reported by JMC.

The results of all three studies are reported in Table 7. There is not much guidance from this table. The range of numbers here is so broad that it is impossible to determine with any degree of reliability whether an investment in good service will pay off.

IMPLICATIONS FOR ARES IMPLEMENTATION

The results of these analyses have significant implications for the implementation of ARES. First, what value of service should be used to determine whether ARES should be pursued? Second, how should the railroad implement ARES? The second question has two parts: (a) What form should the objective functions take? and (b) What implementation strategy should be used for phasing in the functions of ARES over different parts of the railroad?

Because there was a very wide range of estimates on the value of service, we believe that ARES implementation should follow three principles:

1. The initial fielding of ARES capability should concentrate on its capabilities for reducing cost.

2. The ARES design should allow for flexible objectives. That is, the same system should be usable to meet a wide variety of business goals.

3. The implementation process should include a plan for testing the value of improving service.

In following these principles, we have recommended that ARES be installed first on routes where traffic is predominately unit coal trains. If ARES succeeds as projected, substantial reductions in coal train cycle times should be achieved. This will allow Burlington Northern to reduce the number of train sets in service while hauling the same amount of coal. Alternatively, the railroad will be able to haul more coal without putting additional train sets in service.

Once this has occurred, Burlington Northern will then be well assured that ARES can deliver its promised benefits. As the railroad continues to spread ARES capability across the system, it will use ARES' inherently flexible methods for establishing objectives to test the value of service. When ARES is installed in areas where carload freight trains predominate, the railroad can set objectives for the system that allow for more reliable service. Then Burlington Northern will be able to test that more reliable service in the marketplace and see if it pays off.

REFERENCES

1. M. E. Smith and R. R. Resor. Keeping Trains on Schedule: On-Line Planning Systems for the Advanced Railroad Electronics System (ARES). *Journal of the Transportation Research Forum*, Vol. 31, No. 1, 1990, pp. 17–24.
2. M. E. Smith and R. R. Resor. The Influence of Train Delay Costs on Optimal Train Dispatching Plans. Presented at 5th World Conference on Transport Research, Yokohama, Japan, 1989.
3. R. M. Johnson. Trade-Off Analysis of Consumer Values. *Journal of Marketing Research*, Vol. 11, May 1974, pp. 121–127.
4. R. M. Johnson and D. G. Squeo. A Proven Technique for Product Development. *Bank Marketing*, April 1982, pp. 20–24.