Interfaces
Vol. 40, No. 1, January–February 2010, pp. 5–16
ISSN 0092-2102 | EISSN 1526-551X | 10 | 4001 | 0005
DOI 10.1287/inte.1090.0465

THE FRANZ EDELMAN AWARD
Achievement in Operations Research

CSX Railway Uses OR to Cash In on Optimized Equipment Distribution

Michael F. Gorman
Department of Management Information Systems, Operations Management, and Decision Sciences, School of Business, University of Dayton, Dayton, Ohio 45419, michael.gorman@udayton.edu

Dharma Acharya, David Sellers
CSX Transportation, Jacksonville, Florida 32202
{dharma_acharya@CSX.com, david_sellers@CSX.com}

Each day, CSX Railway allocates hundreds of empty railcars among hundreds of customer car orders. In 1997, it implemented the US rail industry's first real-time, fully integrated equipment-distribution optimization system, the dynamic car-planning system (DCP). DCP seamlessly integrates operations research modeling into CSX's process that assigns empty cars to customer car orders. CSX estimates that the DCP system saves the company more than $51 million annually and has saved $561 million since its implementation. DCP has also provided $1.4 billion in capital-expenditure avoidance because of more efficient car allocation. Fewer railcars yield improved return on assets and reduced congestion on the CSX rail network. Customer satisfaction has also increased because of improved empty-car delivery. Public benefits include improved highway safety; reductions in congestion, pollution, and greenhouse gases; and reduced tax-supported road maintenance, thus saving an estimated $600 million.

Key words: freight; transportation; rail; equipment; rolling stock; railcar; CSX; distribution; allocation.

CSX Transportation, Inc., one of the major freight railroads in the United States, employs 35,000 people and earns $11 billion in annual revenue. It provides a crucial link to the transportation supply chain through its 21,000-mile rail network, which serves over two-thirds of the US population in 23 states east of the Mississippi River and in parts of Canada (Figure 1). CSX serves 70 ports along the Atlantic and Gulf Coasts, the Mississippi River, the Great Lakes, and the St. Lawrence Seaway, and thousands of production and distribution facilities through track connections to more than 230 short-line and regional railroads that provide service over all of North America.

To serve its general merchandise customers (e.g., those using boxcars, gondolas, hoppers, tank cars, etc.), CSX first delivers an empty car from its railcar fleet to the customer's required loading location. After loading, the car moves to the customer's load destination and is emptied. It then becomes available for a subsequent customer order, and the cycle begins again.

The implications of improved empty-car distribution decisions are significant for CSX. Repositioning empty railcars to customer order locations creates hundreds of thousands of empty-car miles each day. Empty-car miles generate wear and tear on the railcars and tracks, can require that additional trains be run, and create rail-yard congestion. Finally, empty cars represent lost revenue opportunities.

The Equipment-Distribution Problem

CSX and its predecessor, the B&O Railroad, have delivered empty cars to customers for more than

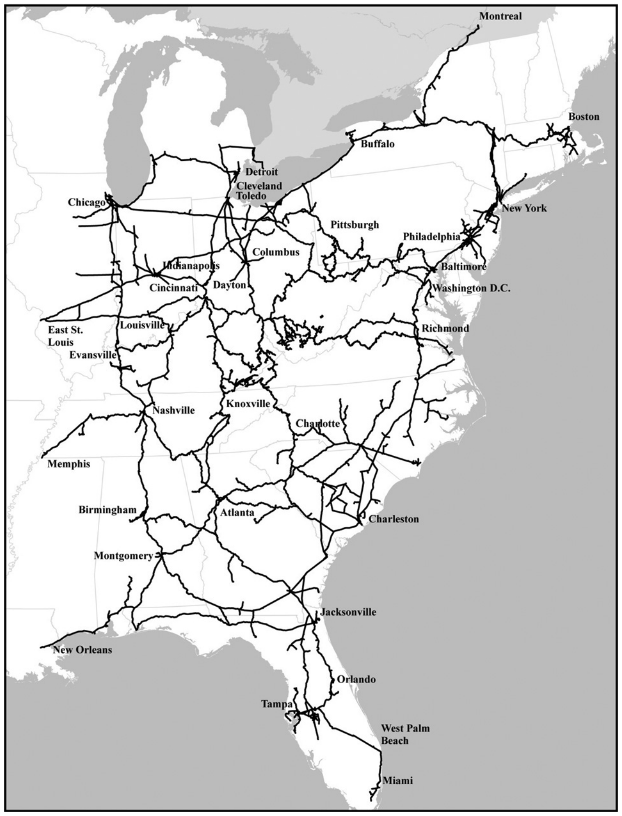

Figure 1: CSX's railway's 21,000-mile network spans most of the eastern United States and connects to other railroads to provide service over all of North America.

180 years. Historically, decisions on which empty cars to deliver to which customers have been made locally because timely and accurate centralized data did not exist. Local decision makers had fewer cars and car orders to manage; however, local control required many decision makers, led to inefficiencies (e.g., regional surpluses and shortages), and hindered the railroad's car-distribution visibility and its control over service levels. A long history of mergers has created today's large CSX rail network. Driven by competitive forces brought about by rail deregulation, and enabled by computer use, equipment-distribution decisions have been centralized, thus providing network efficiencies and improved service.

Empty-car distribution is complex. Each day, hundreds of CSX customers order hundreds of railcars. The source of the cars is a 90,000-car fleet (e.g., boxcars, gondolas, and hoppers) in a network with thousands of geographic locations. The car distributor tries to find the best empty car for a customer order, with "best" established by multiple cost and customer car preferences and service requirements. The problem has hundreds of thousands of possible solutions, and the constantly changing conditions exacerbate it. Throughout the day, equipment availability changes as customers return empty cars and send new and updated car orders, and CSX takes cars offline for cleaning or maintenance.

Confronted with the steady flow of information updates on customer car orders and empty-car availability, CSX had used several methods, including *shipper pools*, *single-car allocation* rules, and the *Sentinel weekly optimization model*, to assist the car distributor and reduce empty-car miles. We describe each method below.

Shipper Pools

A shipper pool is a set of railcars that is dedicated to a single customer and is based on the customer's average demand and shipment cycle time. It eliminates daily allocation decisions, gives customers the comfort of a guaranteed car supply, and simplifies fleet management. Unfortunately, dedicating equipment to one customer hurts car utilization, and any variability in order patterns and cycle times hampers timely deliveries. For example, Corona might ship beer from Mexico to New York in its dedicated pool, whereas International Paper ships paper products from the southeast to the southwest United States and Mexico in its dedicated pool. The transcontinental movements of empty cars in opposite directions are clearly inefficient.

Single-Car Allocation Systems

A single-car allocation system, an expert-system approach that is common in the North American rail industry, enables car distributors to use rules to automate the large volume of car assignments. For example, based on experience and repeated traffic patterns, a distributor might decide that "send boxcars in Birmingham to Smurfit Stone" or "gondolas released in Buffalo go to Newark" are good heuristic rules for decisions that occur regularly. As each empty car becomes available, the appropriate rule is applied and the car is allocated to an order. This approach codifies and standardizes the car distributors' knowledge of empty-car flows.

More importantly, it automates the execution of thousands of car assignment decisions and allows for cost-effective, continuous attention to assignments and immediate disposition of each empty car as it becomes available.

However, a single-car system has drawbacks. Effective rules are hard to design and must change as demand patterns shift, making them cumbersome to manage. Moreover, the approach provides suboptimal solutions because myopic rules apply to cars on a first-come, first-served basis. Each car gets the best assignment available, given the set of rules and the remaining unfilled car orders. The resulting heuristic assignments are influenced by the sequence of events. However, the timing of car orders and empty-car availability is unpredictable, making the performance of the rules difficult to predict.

Figure 2 illustrates the sequence-dependence problem. Four cars become available in sequences 1 through 4. In Figure 2(a), the sequence of car releases requires that the last car must move from Boston to New Orleans, because New Orleans is the only car order remaining when the fourth car becomes available. If the cars had arrived in the order shown in Figure 2(b), Boston would have received a better assignment and fewer miles would have been required to service all orders.

(a)

Empty available locations	Customer order locations	Miles
(1) Chicago	→ (1) Detroit	300
(2) Newark	→ (2) Philadelphia	100
(3) Atlanta	→ (3) Cincinnati	500
(4) Boston	→ (4) New Orleans	1,600
Total miles		2,500
Average empty travel distance: 625 miles/empty car		

(b)

Empty available locations	Customer order locations	Miles
(4) Boston	→ (1) Detroit	300
(1) Chicago	→ (2) Philadelphia	300
(2) Newark	→ (3) Cincinnati	850
(3) Atlanta	→ (4) New Orleans	650
Total miles		2,100
Average empty travel distance: 525 miles/empty car		

Figure 2: Sequential single-car heuristic assignments depend on the sequence of car-order and car-release events; the result is poor and unpredictable results.

When possible, car distributors handle exceptions to the rule-based assignments by overriding these assignments; however, distributors cannot evaluate each rule-based assignment because of the volume of assignments. Therefore, opportunities for improving empty-car distributions are lost.

The Sentinel System

In 1990, CSX implemented the Sentinel system, the first known optimization-based equipment-distribution system in US freight railroads (Turnquist 1986, 1994; Markowicz and Turnquist 1990) to remedy the single-car system problems. Sentinel performs a weekly, fixed-horizon optimization for forecasted orders and car supply to establish optimization-based allocation that the railroad uses to derive rules that are manually entered into a single-car system.

The weekly plan that Sentinel produces provides a more global network view for the car distributor; however, because it is a seven-day, fixed-horizon model, it becomes outdated soon after being generated because of rapidly changing supply and demand conditions. It relies heavily on week-ahead forecasts of expected car supply and demand; however, when forecasts are incorrect, the results are no longer useful because Sentinel cannot take advantage of unforecasted shifts in car supply and demand. It does not integrate with a single-car system, and implementing its solutions is labor intensive and costly. Moreover, its recommendations must still be executed via the sequence-dependent, single-car system.

The Dynamic Car-Planning System

CSX needed to combine the automation of the single-car system with Sentinel's holistic network view of supply and demand. In 1997, it implemented its state-of-the-art, real-time dynamic car-planning (DCP) system to more effectively distribute empty railcars to meet customer orders. The system cost $5 million and took two years to develop. DCP is written in C++ on a UNIX/RISC midtier computer linked to a mainframe computer. It uses a proprietary solver and a Sybase database for data management. In the following sections, we describe DCP's operations research (OR) modeling philosophy and the system, process, and organizational and commercial considerations crucial to its success.

OR Modeling Considerations

CSX could have deployed other modeling alternatives for DCP. Gorman et al. (2009), Cordeau et al. (1998), Powell et al. (1995), Crainic et al. (1993), and Dejax and Crainic (1987) provide detailed discussions of some alternatives.

Shipped cars follow a trip plan (Ahuja et al. 2007, Ireland et al. 2004) that determines a series of trains and intermediate connection yards from origin to destination and the estimated time in transit. The CSX network includes both cars and orders of varying types at various locations, at various dates, and connected by trip plans; thus, modeling the problem as a muticommodity, flow-time-space network of sources, sinks, and connecting trains or trip plans might seem intuitive, as Joborn (1995) and Holmberg et al. (1998) describe for Swedish National Rail. Such a specification, however, has substantial data requirements and long solution times, and it is subject to the uncertainties of train travel times and supply and demand.

To address these uncertainties, Topaloglu and Powell (2006) and Powell and Topaloglu (2005) employ stochastic, multicommodity flow modeling to empty-car distribution at Norfolk Southern Railroad. When CSX was developing DCP, neither railroad had fully researched or deployed these projects. Moreover, CSX did not feel the need to introduce such modeling complexity given the near-term focus of most allocation decisions and the sphere of influence of the car distributors within CSX's organizational structure. It adopted a simpler minimum cost flow formulation. This modeling approach enables quick formulation and solution and a modicum of information to be communicated from DCP to operations. The BNSF Railway (Gorman et al. 2009) and Union Pacific Railway (Narisetty et al. 2008) also subsequently implemented similar basic designs.

System Design and Integration

Figure 3 shows a diagram of DCP's input sources, modules, and integration into production systems. Table 1 shows its key input data. It has five input sources: three external and two internal to the equipment-distribution organization. Key external inputs to the system are the customer car orders, the available cars that might meet these needs, and the transit-time standards. A customer order specifies the car type, any required and preferred features (e.g., capacity, door height), and the required date and location. Each car has a set of attributes, a location, and a date of availability. The illustration in Table 1 shows a 50-foot rigid boxcar available on day 1; in some cases, empty cars are anticipated as available on future dates. Transit-time standards and yard handlings are established weekly for all origin-destination pairs based on current train schedules and car-trip plans for each origin and destination on the CSX network.

The equipment-distribution organization manages two key input files: customer profiles that specify customer priority, car preferences, and allowable substitute cars for each customer and an acceptable earliness and lateness window for car deliveries, and DCP cost parameters that define the hard and soft costs that DCP uses to establish cost parameters.

The feasibility engine (Figure 3) checks the attributes of each empty car and the customer-preference profile to assure the permissibility of each match. Car-order pairs that the feasibility engine deems infeasible for either service or car-preference reasons are screened from the DCP optimization engine. Based on transit-time standards and customer-profile specifications, the feasibility engine identifies permissible cars that can meet a customer-specified time window and excludes those that cannot. For example, the last two rows of Table 1 show examples of infeasible car assignments. The car is not an acceptable type to customer 6 and, by CSX transit standards, cannot be delivered to customer 7 within the lateness tolerance constraint. The infeasibility is shown as a high cost; the optimization model excludes these pairings from consideration because they are deemed infeasible. Simply put, is this the car the customer wants, and can it get to the customer on time?

The costing engine (Figure 3) includes numerous complex hard and soft cost components of each assignment. As Table 1 illustrates, the primary hard-dollar costs in empty-car movements are *car travel distance* (fuel, depreciation), *car handling costs* (the cost associated with car movements at yards between two trains), and *car travel time* (opportunity cost and car hire costs). CSX also includes soft-service costs to assure quality service. It applies a *customer priority*

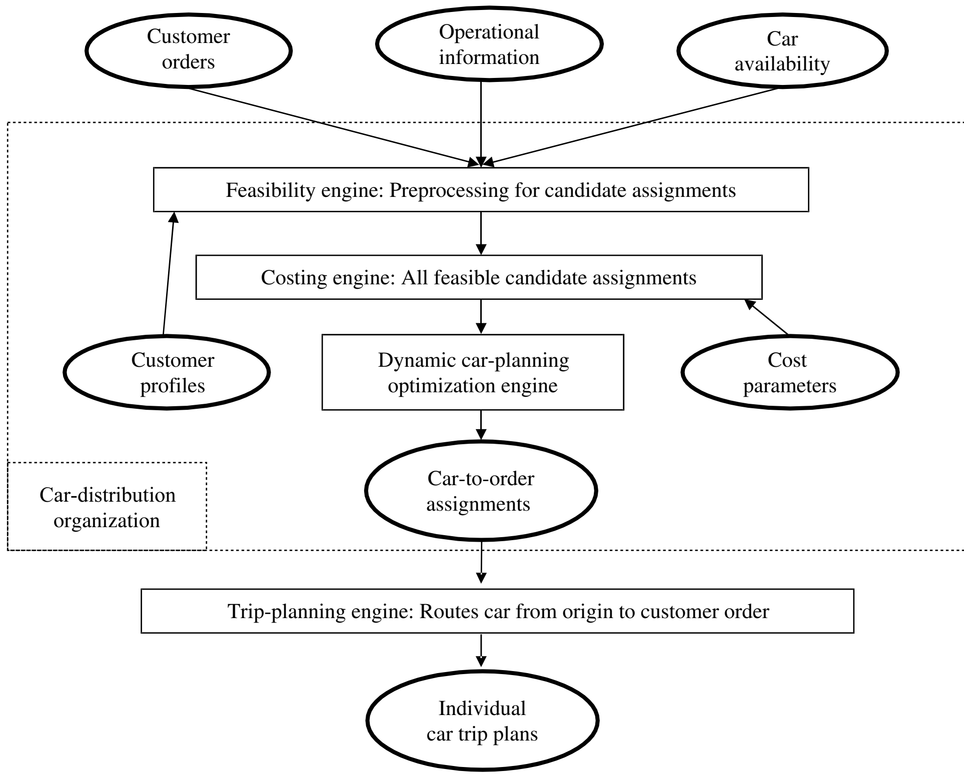

Figure 3: DCP is tightly integrated with operational and marketing data sources as well as production systems. It checks the customer preference and operational feasibility of each assignment, assigns a cost to it via the costing engine, optimizes the network, and communicates car-order assignments to operations on an as-needed basis. Operational systems constantly update DCP on an empty car's next available location for possible replanning.

value to customer orders when making equipment assignment decisions during periods of short supply and applies an *early and late penalty* to encourage on-time delivery. The cost of being late is lost customer goodwill. The cost of being early is based on shipper dock congestion and equipment idle time. The *car preference mismatch* cost is applied when equipment attributes do not match a customer's first choice but can suffice; for example, the car is not a perfect match because of an incorrect door height, length, or cubic capacity but is still usable. In Table 1, we see that customers with a desire for a 50-foot rigid box car have no mismatch penalty, but customers desiring other boxcar types are assessed a penalty to discourage the DCP assignment of a 50-foot rigid boxcar.

The costing engine can handle any level of complexity (including convexity of lateness, demand priorities, fixed costs of car handling, customer priorities, and preferences). It converts the nonlinearities and discontinuities in the cost function to a single-cost coefficient of assigning a specific car to a specific order, as the far right column in Table 1 shows.

DCP Model Formulation

DCP's optimization engine finds the best set of car-to-order assignments and provides a car and a destination to operations to create a trip plan for each car. The optimization engine is based on a straightforward minimum cost flow problem, made possible by preprocessing the feasibility and costing engines. We summarize the model below; the appendix shows the mathematical formulation.

$$
\begin{aligned}
&\text{Minimize total hard car costs and soft penalty costs}\\
&\quad = (\text{car mile cost} + \text{car time cost} + \text{car handling cost})\\
&\qquad + (\text{early penalty} + \text{late penalty} + \text{priority penalty}\\
&\qquad\quad + \text{car mismatch penalty}),
\end{aligned}
$$

Car to allocate: 50-foot rigid boxcar; available on day 1 at location 6

Customer information							Operational information			Feas. eng	Costing engine cost parameters						Optimization engine coefficient
Customer car order				Customer profile							Hard costs			Soft costs			
Customer	Location	Desired car type	Want day	Customer priority	Max early/late (days)	Car preference rank	Distance (miles)	Trip plan transit standard (days)	Trip plan yard handling	Expected delivery day	Car travel distance ($)	Car handling ($)	Car travel time ($)	Customer priority penalty ($)	Early/late penalty ($)	Car preference mismatch ($)	Total cost ($)
Cust1	Loc1	Plate F Box	1	HIGH	2	LOW	100	2	1	3	83	50	46	100	178	228	686
Cust2	Loc2	50FtRigidBox	3	HIGH	2	HIGH	300	4	2	5	250	100	93	100	178	—	721
Cust3	Loc3	Plate F Box	3	MED	1	LOW	50	1	2	2	42	100	23	200	39	228	632
Cust4	Loc4	50FtRigidBox	4	MED	1	HIGH	200	2	3	3	167	150	46	200	39	—	602
Cust5	Loc5	HighCubeBox	5	LOW	2	MEDIUM	80	2	2	3	67	100	46	300	78	152	743
Cust6	**Loc6**	**Gondola**	**2**	**HIGH**	**2**	**N/A**	**0**	**1**	**0**	**2**	—	—	**23**	**300**	—	**999,999**	**999,999**
Cust7	**Loc7**	**50FtRigidBox**	**2**	**HIGH**	**2**	**HIGH**	**240**	**4**	**2**	**5**	**200**	**100**	**93**	**100**	**999,999**	—	**999,999**

Sample cost parameters:

Car mile cost ($)	Car time cost/day ($)	Handling cost ($)	Late penalty ($)	Early penalty ($)	Customer priority ($)	Car preference ($)
1	23	50	89	39	100	76

Table 1: DCP applies a number of hard and soft costs to each possible car-to-order assignment. The costing engine calculates these costs; the DCP optimization engine applies the total of these costs when making car allocation decisions. The feasibility engine applies constraints to cars that do not meet the car type or timely service needs of the customer as indicated by each customer's profile (shown in bold).

Notes. Total cost of car assignment = car mile cost + car handling cost + car time cost + customer priority penalty + late penalty + early penalty + car preference penalty. The total cost of car assignment acts as the optimization engine coefficient in the DCP optimization problem.

subject to

- All car orders are met,
- Supply of each car type is not exceeded,
- Service constraint: Car is delivered within the allowable time window,
- Car-type constraint: Car delivered meets the allowable car-type constraint. This formulation fits with the organizational, operational, and commercial considerations, and CSX's rapidly changing conditions, as described below.

Organizational Considerations

The car distributor is concerned with the feasibility and cost of the car-order pairing but cannot affect the operational details of the car's trip plan, which are the responsibility of CSX's operations department. The model formulation focuses solely on the car assignment, using the information of the underlying trip plan for determining feasibility but treating it as exogenous. The time-space network concept is preserved through differing supply and demand nodes for each car-availability date and customer order and is taken into account by the feasibility and costing engines. Thus, the time-space modeling that Swedish National Railway uses (Joborn 1995, Holmberg et al. 1998) is less appropriate and unnecessary given CSX's organizational structure, utilization of the trip-planning engine, and decision sequence.

Operational Considerations

To handle the inherent future uncertainty of the problem, CSX solves the deterministic DCP model frequently (i.e., every 15 minutes throughout the day) to capture changing supply-and-demand conditions. The model takes 1 minute to load and 10 seconds to solve. The result is a best-deterministic solution that responds quickly to changing conditions, such as supply-and-demand shifts.

The OR modeling in DCP is simple. The approach enables CSX to reoptimize quickly and adjust plans as new information becomes available. From its Sentinel experience, CSX realized that solving the model frequently and continually refreshing solutions based on the current best information produced the most useful, effective, and implementable solutions. Unlike Sentinel, DCP bases its solutions on a rolling horizon; at any given time, DCP looks two weeks ahead to avoid the short-sightedness of single-car systems. Although DCP might generate multiple plans for a car over time, CSX defers decisions until they are required, i.e., the time at which the operations department needs disposition on the cars. Thus, it avoids thrashing (i.e., frequent reversals of previous decisions). This is a fundamental difference between DCP and the single-car system, in which a car is given a single and final assignment immediately after the event that makes the car available (Figure 3). Other modeling approaches that endogenize uncertainty (e.g., Topaloglu and Powell 2006) are more data intensive, take longer to run, and are thus run less often. CSX has been successful with this simpler, intuitive, and real-time approach that focuses on immediate executable decisions while keeping a placeholder for future supply-and-demand parings.

DCP goes beyond deferring a decision until the trip plan starts. It can revisit and revise the decision even after an empty car has begun its trip plan, and the operations group can seamlessly execute according to the change in plan. Figure 4 shows the integration of DCP with operations. The DCP car assignment is available when a trip begins. The empty car typically has multiple trains and intermediate handling yards in its trip plan. Each handling yard removes the car from the inbound train and switches it to the outbound train. The yard can easily handle a reassignment if it receives notice prior to the car's arrival at the yard. After an empty car has started its trip, it appears to DCP as supply at the car's future intermediate handling yard and date of availability. If in subsequent

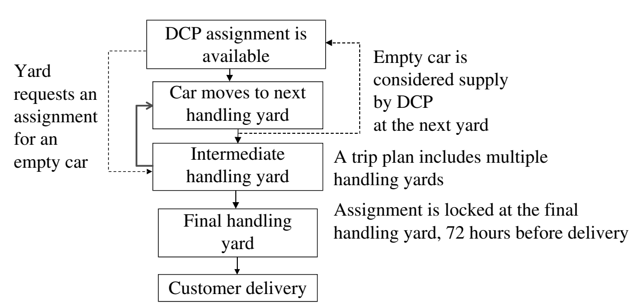

Figure 4: DCP integrates tightly with the CSX operational process. A car might receive many assignments before it begins moving and is a candidate for assignment at each intermediate handling yard. The car assignment is held by DCP until needed by operations, effectively deferring final decisions until necessary.

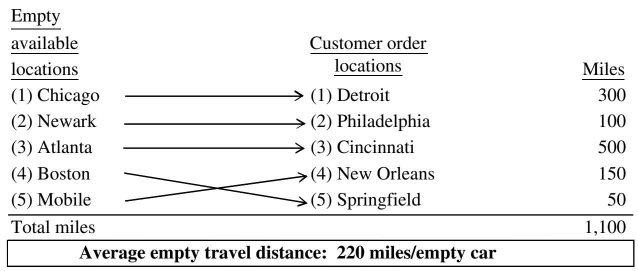

Figure 5: OR allows DCP to take a network view, making the sequence of events immaterial. Deferring and revisiting decisions can save many miles. In this example, empty miles are reduced by more than 50 percent, and an additional customer is served.

DCP optimization runs the original assignment is still optimal, DCP will "reassign" the car to the same customer; however, DCP might also revise the final destination at that time because new supply or demand information has changed the car's optimal assignment. In either case, the car assignment executed by the intermediate handling yard is finalized when operations pulls the current car assignment from DCP before the car arrives at the handling yard. This capability gives DCP opportunities to improve the assignment of each empty car after its trip has begun.

Figure 5 shows the effect of one more car and order on the original empty-car distribution network (Figure 2). The last order and available car might become available after the other cars have been assigned and have begun their trips. By dynamically reassigning previously assigned cars, total empty-car miles decrease even as the number of customer orders served increases.

Commercial Considerations

Some CSX customers were accustomed to pools of cars dedicated to their individual use and were not comfortable with relinquishing that control. Others were familiar with the single-car allocation systems and could track "their" empty car as it traversed the CSX network toward their facility. This is akin to requiring Hertz to provide a customer with the license plate number of a car reserved for the following week. To improve operational efficiency and customer delivery, CSX had to remove the artificial constraints, and its customers had to trust that CSX would deliver on time a car that matched their specifications. CSX "locks" the final assignment at the last intermediate handling yard, 72 hours prior to delivery. DCP is not likely to change the car's destination, and providing this insurance relieves any customer concerns, thus improving customer service.

DCP relies on the accuracy of customer orders and forecasts. CSX requires a one-week lead time on all car orders and requests its customers to provide a six-week forecast of car orders. Although not binding, this forecast is vastly superior to any statistical forecast that CSX might develop. CSX and its customers both benefit from this collaborative relationship.

Challenges and Best Practices

Challenges

CSX was the first railroad to implement real-time decision support systems for empty-car distribution. Understandably, its customers had concerns about how the system would affect their service. Moreover, CSX field-operating personnel were used to having more autonomy in making car allocation decisions; some did not understand or did not want to follow DCP recommendations. CSX had to provide considerable training and communications to ensure that these constituencies would embrace the dramatic change. The communications efforts were made more difficult because CSX chose a simultaneous rollout of DCP on the entire 90,000 DCP-managed car fleet, an approach that accelerates benefits but increases cost and risk.

Because of the difficulties of trading off hard and soft costs, model cost-parameter tuning is critical to obtaining good results. Car distributors and fleet managers work with the DCP system, adjusting parameters by fleet, geographic region, and customer. The car distributors' job description changed with the inception of DCP; they became costing technicians rather than car allocators. Car distributors manage DCP exceptions rather than each car assignment. CSX was putting this innovative system in the hands of savvy car distributors; however, they did not know OR. Training on how the model works and why it differed from previous methods was essential for user acceptance. Illustrations such as Figures 1 and 4 helped users to understand the new system.

Best Practices

CSX found that having internal senior-level "champions" helped garner resources for DCP development and implementation. It identified key car-distribution experts who supported the more sophisticated approach; they were instrumental to overcoming problems during implementation.

CSX's experience with Sentinel paved the way for the more advanced DCP system. First, it simplified the financial justification of the new system; Sentinel showed the value of optimization and the need for real-time data feeds and system integration. Second, it led to quicker user acceptance by the fleet managers. CSX's evolution from single-car heuristics, to weekly fixed-horizon optimization, to DCP's dynamic and integrated optimization provided insights into how to improve empty-car distribution.

DCP Benefits

DCP has provided substantial benefits to all stakeholders, including CSX, its customers, the rail industry, and the US public, as we describe below.

Financial Benefits to CSX

Since its DCP implementation, CSX has needed 10 fewer car distributors, thus saving $1 million per year. More importantly, empty cars now travel far fewer miles to their next load destination. Each saved empty-car mile reduces fuel, crew, and equipment depreciation costs. We show our calculation methodology below.

We developed a savings method based on the improvement in empty miles for CSX's biggest fleets, boxcars, and gondolas. Table 2 shows the empty miles and loaded miles in 2006 and estimated savings relative to the 1:1 empty-load ratio that CSX attained prior to implementing DCP. We projected the empty miles for 2006 (without DCP) based on the historical ratio prior to the DCP implementation and the actual loaded miles.

We used multiple methods to estimate and validate our savings before and after implementing DCP. We looked at additional years, and compared DCP with non-DCP in the same year and CSX with other railroads in the same fleet. First, we looked at the years 2002 through 2008 for consistency of results within the boxcar and gondola fleet over time. Second, we compared the empty-load ratio for DCP-managed cars with "foreign" (owned by other railroads) and "private" (owned by shippers) railcar ratios in the same year. Third, we compared the CSX empty-load ratio for the cars in the North American boxcar fleet (a shared railroad pool of cars) with the empty-load ratio achieved by other railroads. In each case, we found that these comparisons supported our savings estimate.

Empty movement cost per mile	$0.80
Boxcar savings	
Total estimated empty miles without DCP	26,938,800
Total empty miles with DCP	17,959,200
Mileage savings from DCP	8,979,600
Annual DCP boxcar savings (at $0.80 per mile)	$7,183,680
52-foot mill gondola savings	
Total estimated empty miles without DCP	47,566,478
Total empty miles with DCP	37,050,167
Mileage savings from DCP	10,516,311
Annual DCP gondola savings (at $0.80 per mile)	$8,413,049
Total savings: Gondolas and boxcars	$15,596,729
Estimated annual total savings for entire fleet	$50,312,028

Table 2: The gondolas and boxcars in these examples are approximately 30 percent of the DCP-managed fleet and are representative of the savings achieved in other fleets. Applying these savings across the remaining DCP-managed fleets results in $51 million in annual savings and $561 million savings in the years between 1997 and 2007.

Capital-Avoidance Savings

When cars spend less time empty, they can spend more time loaded. One could evaluate DCP's benefits by estimating volume increase and revenue enhancement from the railcar fleet's more efficient use. More conservatively, we quantify DCP's benefits using the estimated reduction in the size of railcar fleets required to serve *existing* business. CSX has avoided capital expenditures by making better use of its existing fleets rather than buying more cars. If CSX had to support its current business with its 1997 empty miles per load, the fleet would require an additional 18,000 cars at a replacement cost of $75,000. The railcar capital investment avoided is approximately $1.4 billion, with a resulting increase in CSX's return on assets because of using each car more efficiently.

By having fewer cars on its network, CSX reduces yard congestion; by reducing the number of trains that it must run, it reduces track congestion, a major rail-industry problem (Gorman 2009). Highly utilized yard, line, and train capacity can result in delayed empty deliveries, or worse, lost orders and revenue.

By reducing car fleets and empty miles, DCP alleviates these problems, reducing cost and improving service to customers. Although we do not quantify them in our estimate, these savings are significant.

Customer Satisfaction

Using DCP, the car manager can more easily substitute equivalent cars when a shortage exists and can more actively track fulfillment status, improving customer satisfaction by making it more likely that the customer will get the right car on the right day. Largely because of the improved order-fill rates and timely deliveries made possible by DCP, CSX's customer "car-order" score was the highest of all customer-satisfaction scores that CSX achieved in its 2008 J.D. Power and Associates survey (Figure 6). Thus, CSX has been able to retain and grow business that might have otherwise gone to a competing railroad or to a truck.

Rail-Industry Benefits

CSX has led the North American rail industry in the development of optimized empty-car distribution systems. Ireland et al. (2004) report that in 2002, Canadian Pacific (CP) used the methodology that Turnquist developed (Turnquist 1986, 1994; Markowicz and Turnquist 1990), and CSX had implemented as the Sentinel system. CP subsequently purchased DCP from CSX. Other railroads have benchmarked DCP, implemented similar systems, and attained improvements. BNSF implemented a system of similar design in 2000 (Gorman et al. 2009); Union Pacific Railway followed suit in 2003 (Narisetty et al. 2008). Norfolk Southern Railroad is also benchmarking DCP.

The use of empty-car distribution systems by other railroads indirectly benefits CSX because these systems enable the railcar fleet in the US rail network to be used more efficiently. Notably, the major railroads have developed a shared boxcar fleet, effectively breaking down the railroad-specific ownership constraint in car-fleet management. Systems such as DCP enable the management of such a fleet.

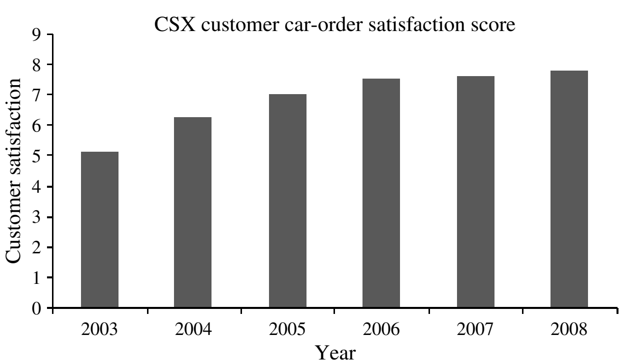

Figure 6: CSX's customer car-order satisfaction rating has increased steadily since the DCP implementation. CSX has gradually increased its J.D. Power car satisfaction ratings since 2003, the first year of tracking.
Source of data. J.D. Power and Associates.

US Public Benefits

The public benefit derived from DCP results from rail's inherent societal advantages over trucks (Gorman 2008) and CSX's expanded ability to serve rail customers. If CSX had not invested in either DCP or 18,000 new cars, carrying today's rail load would have required 383,250 trucks, given typical cycle times, a 3.5:1 railcar-to-truck weight ratio, and a 50 percent rate of conversion to truck (Table 3).

Rail is a green technology; Gorman (2008) estimates that rail freight incurs only 20 percent of

Societal savings calculations

DCP cars saved (CSX estimate)	18,000
Car cycle time (days)	30
Cycles per year (loads/car)	12.17
Railcars per year (cars ∗ loads/car)	219,000
Percent diverted from truck (CSX estimate)	50%
Tons/truck (CSX estimate)	20
Tons/car (CSX estimate)	70
Trucks/car (Ratio: tons per railcar/tons/truck)	4
Annual reduction in truckloads (diversion ∗ railcars ∗ truck/car)	383,250
Avg miles/load (CSX estimate)	614
Tons/year (cars ∗ cycles ∗ tons/car)	15,330,000
Ton miles/year (miles/load*tons/year)	9,412,620,000
Truck social costs per ton mile (Gorman 2008)	$0.0144
Int. rail social cost per ton mile (Gorman 2008)	$0.0028
Annual social cost truck (cost/ton mile truck ∗ ton miles)	$67,770,864
Annual social cost rail (cost/ton mile rail ∗ ton miles)	$13,177,668
Annual social cost savings maximum (truck-rail)	$54,593,196
Estimated 11-year social savings	$600,525,156

Table 3: Given a CSX average of 614 miles per load and typical weight of a railcar, we calculate 15 million tons moved, or 9.4 billion ton-miles. Using published values for relative public costs of truck and rail, we calculate truck cost at $68 million and rail at $13 million, thus calculating a $55 million annual incremental truck cost (without DCP). Over 11 years, this represents approximately $600 million in public costs avoided.

the public costs of comparable truck freight, resulting in improved road safety and reduced pollution, greenhouse gases, congestion, and tax burdens from unfunded road maintenance.

Given the number of shipments and their typical weight and distance, DCP has contributed approximately \$600 million in public benefits since its inception. Of course, because of the vagaries of developing such an estimate, we intend it only as an indicator of the order of magnitude of the public benefit. However, we note that this benefit estimate is based on conservative savings values and from CSX's experience, not including road-to-rail diversions made possible by similar railroad car-distribution systems. DCP's annualized public and private benefit is approximately \$250 million per year, not including benefits to and from other railroads.

DCP Improvements and Future Use

Over the last decade, CSX has continued to invest in and enhance DCP. It has improved costing parameters, accuracy and completeness of customer profiles, and real-time reporting information on train operations and car status to make DCP more effective. Web-based order-management and visibility tools, known as ShipCSX, have made DCP more customer friendly and have improved car order-forecast accuracy. DCP's use over the previous decade indicates both the system's stability and its fit with CSX's long-term equipment-distribution needs. Simply stated, CSX could not revert to previous empty-car distribution methods and continue to operate with its existing fleet and network capacity.

Summary

CSX required a real-time empty-car distribution system that was nimble and robust enough to (1) adjust to rapidly changing conditions throughout the day by constantly refreshing the current best solution, (2) defer decisions until they must be made, (3) revisit previously made decisions when needed, and (4) seamlessly integrate with the car-distributor decision process and communicate decisions to both customers and operations. DCP provides all these functionalities.

Without tight real-time systems and process integration of both inputs and outputs, DCP could not (1) have the required current information for quality model results, (2) make the information available to the field on an as-needed basis, and (3) change the routing of a car as new information becomes available and only as operationally feasible. These features, coupled with fast OR model-solution times, drive DCP's success.

DCP enables CSX's continued quality customer service and growth. Its sustained use over the last decade gives evidence to its stability and potential future use. DCP is responsible for almost \$2 billion in capital and operating savings to CSX and an additional \$600 million in public benefits, or \$250 million per year in annualized public and private benefits, not including benefits to other railroads. The subsequent adoption of such systems by other railroads is further evidence of the widespread transferable benefits of OR-based empty-car distribution systems. DCP exemplifies the importance of operations research in managing the ever-increasing complexity of the US rail network and is perhaps the most successful OR application in the rail industry.

Appendix

The optimization model (Equations (1)–(6)) is based on a minimum cost flow formulation. Let a be a vector of car attributes such as car type, date and location of availability, etc. Let b be the set of customer preferences on desired and substitute car types, location, date, and delivery window. A is the set of all attribute vectors on current car supply, and B is the set of all attribute vectors on current orders. S_a, $a \in A$, is the number of cars with particular attribute vector a, and D_b, $b \in B$ is the number of orders with particular attribute vector b. Φ is the set of allowable car to order pairings (a, b), $a \in A$, $b \in B$.

The objective function, Equation (1), charges a cost c_{ab} for allowable assignments Φ of supply a to car orders b. Flows from demand locations to the sink node k earn a bonus (a contra-cost) b_d based on customer priority for providing service. The total costs of assignments are minimized through optimal flows x_{ab}, which is an integer variable (6). The number of supply units R at the source node r and the volume of demand K at the sink node k are set equal to the sum of all supply over the horizon, $R = K = \sum_{a \in A} S_a$. The flows from the source node r are equal the supply at each supply node S_a to force equipment into the network at the appropriate supply nodes (2). All cars are accounted for (3), either in an assignment to demand or no assignment (storage). Flow to a demand location to not to exceed its demand (4), and the flow into a demand location equals the flow out (5).

$$\text{Min} \quad \sum_{ab\in\Phi} c_{ab}x_{ab} + \sum_{a\in A} C_k x_{ak} - \sum_{b\in B} b_d x_{bk} \tag{1}$$

$$\text{subject to} \quad x_{ra} = S_a \quad \forall a \in A, \tag{2}$$

$$\sum_{b\in B} x_{ab} + x_{ak} = S_a \quad \forall a \in A, (a, b) \in \Phi, \tag{3}$$

$$\sum_{a\in A} x_{ab} \le D_b \quad \forall b \in B, \tag{4}$$

$$\sum_{a\in A} x_{ab} = x_{bk} \quad \forall b \in B, \tag{5}$$

$$x_{ab} \ge 0, \quad \text{and integer.} \tag{6}$$

Acknowledgements

We would like to thank the Edelman Committee for its consideration of this work as an Edelman Competition finalist, and our competition coaches, Anthony Brigandi and Sudhansu Baksi, for their help in developing the paper and presentation. We would like to thank the faculty, staff, and administration of the University of Dayton for their support during the competition. Finally, we would like to thank the scores of dedicated CSX employees who contributed to the original and ongoing success of DCP; in particular, we thank Dave Bell, Marty Blue, Alan Blumenfeld, Kathleen Brandt, Cheryl Crow, Ellen Dear, Gene Hartley, Ken Hinson, Mary Hollin, Joe McMillan, Bob Muenz, Tim Poineau, and Mickie Wittig.

References

Ahuja, R. K., C. K. Jha, J. Liu. 2007. Solving real-life railroad blocking problems. *Interfaces* **37**(5) 404–419.

Cordeau, J.-F., P. Toth, D. Vigo. 1998. A survey of optimization models for train routing and scheduling. *Transportation Sci.* **32**(4) 380–404.

Crainic, T. G., M. Gendreau, P. Dejax. 1993. Dynamic and stochastic models for the allocation of empty containers. *Oper. Res.* **41**(1) 102–126.

Dejax, P. J., T. G. Crainic. 1987. A review of empty flows and fleet management models in freight transportation. *Transportation Sci.* **21**(4) 227–248.

Gorman, M. F. 2008. Evaluating the public investment mix in U.S. freight transportation infrastructure. *Transportation Res. Part A* **42**(1) 1–14.

Gorman, M. F. 2009. Statistical estimation of line congestion delay in U.S. freight rail. *Transportation Res. Part E* **45**(3) 446–456.

Gorman, M. F., K. Crook, D. Sellers. 2009. North American freight rail industry real-time optimized equipment distribution systems: State of the practice. Unpublished manuscript, University of Dayton, Dayton, OH.

Holmberg, K., M. Joborn, J. T. Lundgren. 1998. Improved empty freight car distribution. *Transportation Sci.* **32**(2) 163–173.

Ireland, P., R. Case, J. Fallis, C. Van Dyke, J. Kuehn, M. Meketon. 2004. The Canadian Pacific Railway transforms operations by using models to develop its operating plans. *Interfaces* **34**(1) 5–14.

Joborn, M. 1995. Empty freight car distribution at Swedish Railways—Analysis and optimization modeling. Ph.D. thesis, Department of Mathematics, Linkoping University, Sweden.

Markowicz, B. P., M. A. Turnquist. 1990. Applying the LP solution to the daily distribution of freight cars. Presentation, TIMS/ORSA national meeting, Cornell University, Ithaca, NY.

Narisetty, A. K., J.-P. Richard, D. Ramcharan, D. Murphy, G. Minks, J. Fuller. 2008. An optimization model for empty freight car assignment at Union Pacific. *Interfaces* **38**(2) 89–102.

Powell, W. B., H. Topaloglu. 2005. Fleet management. S. Wallace, W. Ziemba, eds. *Applications of Stochastic Programming*. Math Programming Society—SIAM Series in Optimization, Philadelphia, 185–215.

Powell, W. B., P. Jaillet, A. Odoni. 1995. Stochastic and dynamic networks and routing. M. O. Ball, T. L. Magnanti, C. L. Monma, G. L. Nemhauser, eds. *Handbooks in Operations Research and Management Science, Network Routing*, Vol. 8. North-Holland, Amsterdam, 141–295.

Topaloglu, H., W. B. Powell. 2006. Dynamic programming approximations for stochastic, time-staged integer multicommodity-flow problems. *INFORMS J. Comput.* **18**(1) 31–42.

Turnquist, M. A. 1986. MOVE-EM: A network optimization model for empty freight car distribution. Report, School of Civil and Environmental Engineering, Cornell University, Ithaca, NY.

Turnquist, M. A. 1994. Economies of scale and network optimization for empty car distribution. Technical report, Cornell University, Ithaca, NY.