Chapter 7
Empty Railcar Distribution

Michael F. Gorman

7.1 Introduction

Each year in North America, approximately 30 million carloads are shipped via rail in "general merchandise" or carload service (AAR 2012). In each case, the railroad must deliver a rail-owned empty railcar (such as a box car, gondola, or hopper depending on the commodity) to the origin of the shipper to begin loading. (This process does not apply to private fleets owned and managed by the shipper, as is common for some car types such as tank cars.) After the loaded railcar is delivered to the shipper's destination and emptied, the rail car is released back to the railroads' custody and the cycle begins again. The challenge of repositioning a multitude of rail-owned railcars to various origins is known as the empty railcar distribution problem.

The empty railcar distribution problem is complicated by a number of considerations, including the specificity of the wants and needs of the customer (such as capacity and door height), rail ownership and rent paid to other railroads for use of their cars (known as "foreign car hire"), and the distance and time the empty must move over. Further, orders are received and cars released unpredictably, so the problem is constantly changing throughout the day. Finally, there are a number of "soft" trade-offs such as the desire for timely delivery (not too early and not too late) and customer car preference.

Effective solution of the problem is extremely valuable to the rail industry. First, customer service can be improved. Second, cars spend less time empty and more time loaded. Third, with effective empty railcar distribution fewer cars are needed, and those that are in service travel fewer empty miles, producing lower wear and tear on cars per load handled. Fourth, the train space required for moving empties is reduced, effectively expanding capacity for loaded movements, and saving locomotive fuel.

M.F. Gorman (✉)
Department of Operations Management, University of Dayton School of Business, OH, USA
e-mail: mgorman1@udayton.edu

B.W. Patty (ed.), *Handbook of Operations Research Applications at Railroads*,
International Series in Operations Research & Management Science 222,
DOI 10.1007/978-1-4899-7571-3_7

Railroads have reported saving tens of millions of dollars per year and billions of capital avoidance from the implementation of railcar distribution systems (see Gorman et al. 2010, 2011; Narisetty et al. 2008).

These tactical railcar systems have been extensively applied to general merchandise traffic, which accounts for about 20 % of all rail traffic. The systems are described in detail for the most of the remainder of this chapter. Automotive railcar distribution follows similar rules, but is managed differently in U.S. rail, as described at the end of this chapter. Coal and grain typically move in "unit" trains, cycling between origin and destination collectively as a train set. Intermodal railcar distribution has different requirements because shippers do not order or load the railcar, but rather the container. (This problem is discussed in the intermodal chapter of this book.)

7.2 Background on Empty Railcar Distribution

7.2.1 Local Distribution and Shipper Pools

Before centralized information systems became commonplace in the rail industry, railcar supply was managed locally. A pool of cars was managed locally and allocated among local shippers. Decentralized control led to inefficiencies such as hoarding behavior and regional shortages. Rail mergers have created larger and more complicated rail networks, creating the need for more sophisticated methods. Improved information systems have allowed for centralized car tracking information. "Shipper pools" (sets of cars with similar characteristics that were dedicated to a shipper) were still used to manage car supply after centralization, but such constraints on car usage vastly reduced flexibility and did not allow for more efficient assignments.

7.2.2 Rules-Based Transaction Processing Systems

Early equipment distribution systems were rules-based transaction processing systems. As a car was released by a consignee, or an order for equipment placed by a shipper, various criteria (such as car type, dimensions, capacity, and ownership) were checked and a car was assigned to an order. In the case of a car becoming available and no orders for that type of car present in the system, a generic "flow order" was used to get the car moving in the general direction of the demand for such cars. This expert-system style of rules codified the knowledge and heuristics used by car distributors to manage car supply and fill orders. Importantly, it automated a labor-intensive task.

However, these systems were lacking in a number of ways. First, the copious rules had to be managed and changed as very seasonal shipping patterns changed. Effective rules are hard to create, and harder to keep up to date when shipping patterns shift. Second, a heuristic system had rules that worked in general, but often failed to manage the fleet well in specific instances because of the sequence dependency of the execution of such rules. For example, a car might receive a rule-based assignment to an order 500 miles away from the shipper, and subsequently another car could become available only 50 miles away from the shipper, but the first car would remain on the order. This assignment would have been reversed if the near-by car was released first, demonstrating another problem; the execution of the rules was highly sequence dependent. Car distributors could manually override such poor assignments, but because the volume of transactions was high, often such inefficiencies would go unnoticed.

7.2.3 Nonintegrated Optimization Systems

Early attempts at optimization of railcar distribution were not integrated with the transactional systems. (Published examples include Jordan and Turnquist 1983; Turnquist 1986; Turnquist and Markowicz 1990.) Typically, a week's worth of actual and forecasted orders were optimally allocated according to an objective such as minimizing total miles of empty car movements, subject to customer service constraints. The problem was formulated as a transportation problem in which supplies of empty cars are assigned to customer orders, minimizing the distance the cars travel, among other considerations. Such a system showed potential improvements over rules-based systems because of their more global view of the problem.

However, these optimization programs generally did not achieve anticipated benefits. The weekly forecast was, of course, subject to error. Often, the optimal results were out of date long before they could be used, and worse, recommendations could be wrong because of errant forecasts and execution failures. Finally, model results would be implemented in the transactional system, causing a large number of assignments to be manually entered. As a result, nonintegrated optimization was not a successful attempt at introducing optimization to empty car distribution.

7.3 Current Day Integrated Real-Time Optimization Systems

In the late 1990s and early 2000s, railroads began investing in integrated, near-real time optimization systems. CSX Railroad implemented its "Dynamic Car Planning System" (DCP) in 1997; BNSF developed its "Equipment Distribution Optimization" (EDO) system in 2000 (Gorman et al. 2010). The Union Pacific developed its system in 2003 (Narisetty et al. 2008).

7.3.1 Model Inputs

The systems have remarkably similar characteristics; below we describe the common components found in such systems: Car supply, shipper orders, marginal shipping cost factors, and customer preferences.

7.3.1.1 Car Supply: Actual and Predicted

The primary source of car supply is the location, date and time of release of an empty car from a consignee. In some cases, equipment is moved from the consignee's location to storage locations in anticipation of future orders.

Forecasted supply is also often used for predicting future anticipated supply. In some cases, empty equipment is in transit, and its "location" is the next location, date and time the car is planned to be available when the train is at the next yard where the car is switched. In each of these cases, the actual car and its full set of attributes (dimensions, etc.) are used in matching the car to customer orders. Often, cars are interchanged between railroads, and future supply is predicted to be delivered from the other railroad where they meet. In lieu of information shared from the "foreign" railroad, interchange volumes of general equipment types are forecasted based on historical patterns and general equipment attributes.

7.3.1.2 Car Orders: Actual and Predicted

Car orders are placed by shippers at the time they plan a loaded shipment. Railroads request (but do not always receive) sufficient lead time to plan empty economical and on-time deliveries, usually 1 or 2 weeks in advance. Some railroads request more advance time on orders for longer term planning, others supplement actual car orders with statistical forecasts based on historical patterns. Such forecasts are often simply moving averages with some day of week and time of year seasonality. Orders are notoriously hard to predict, so actual car orders are vastly preferred. Often, such orders are aggregated into large geographic regions, and storage yards are used as the center of aggregation. Thus, a car that is assigned to a forecasted load (planned for a storage yard) is superseded by an actual car order.

7.3.1.3 Shipper Preferences

Shipper preferences such as maximum allowable early and late delivery, specific physical car requirements (capacity, door heights, and other attributes), and allowable substitutes are kept to balance customer car needs with the need for some flexibility in meeting orders. Hard requirements act as constraints on allowable equipment assignments to orders; preferences are included as a component of the cost of car-to-order mismatches.

7.3.1.4 Cost Parameters

Railroads consider a number of hard dollar costs for empty car distribution, including empty car mileage, travel time costs (including car rents for use of foreign railroad cars), and car handling costs for switching between trains at yards. Shipper preferences are also used in costing for capturing the soft service costs of empty car assignment. Within allowable car assignments, slightly early, late and mismatched cars are assigned a service cost. All of these costs are applied to the feasible railcar-to-order pairings and combined into a single cost coefficient in a costing module for use in the optimization model.

7.3.1.5 Operational Information

Service times from empty supply locations to shipper origins based on train service helps to identify the feasibility of assignments of empties to orders from a timing perspective.

7.3.2 Model Framework

7.3.2.1 Model Preprocessing

The complexity of train operations is simplified through preprocessing. Allowable matches are found by matching car attributes to customer requirements, and checking service feasibility based on empty availability date, customer order date, and the service time between the two locations. In this way, the complexity of rail movements and operations is reduced to core information needed by the model: where is the car, when will it be available, and how long does it take to get to candidate destinations. Each car is considered for assignment to orders for which it meets the customer service criteria.

7.3.2.2 Model Formulation

Problem preprocessing allows the empty railcar assignment problem to be solved via a transportation problem or transshipment problem formulation. Because the two formulations are similar, only the transportation problem is shown below. A detailed comparison of the two formulations can be found in Gorman et al. (2011).

We define a as the vector of permanent car attributes such as car type and ephemeral attributes such as next available date and available location. We define b as the vector of attributes on a customer order, including specific requirements on car type as above, and other shipper attributes such as location, priority, date equipment required, customer preferences on acceptable substitute equipment, acceptable early or lateness, forecasted or actual order, etc. We let A be the set of cars in the planning period and let B be the set of orders in the planning period. We define S_a, $a \in A$ to be the number of planning period cars with particular attribute vector a, and D_b, $b \in B$ be the number of planning period orders with particular attribute vector b. As described earlier, the set of attributes a and b includes not only physical car attributes and customer requirements (e.g., car type, dimensions, etc.), but also the date and location of each car and order.

We define Φ as the set of allowable pairings (a,b), with $a \in A$ and $b \in B$ to assign a car with a vector of attributes a to an order with a vector of attributes b established in preprocessing. Φ limits the number of decision variables, x_{ab}, considered by the model by eliminating parings of cars to orders that are not acceptable. Φ is not only based on the customer's car acceptance profile which defines allowable assignment of a car of attributes a to a customer car order with requirements b, but also the feasibility of the railroad to deliver the car from its location and available date within an acceptable time window of the customer's desired date to the customer at a given location.

As discussed in the previous section, the hard and soft costs of any allowable assignment in Φ of supply to demand (whether actual cars and orders, or forecasted groups of car types and order types) are established via preprocessing, and are included in the single cost coefficient, c_{ab}.

In order to assure feasibility of any model run regardless of the data, a phantom supply source (r) and super sink (l) are created prior to optimization. The source and sink capacity are calculated prior to model formulation. The number of supply units at r, $R = \sum_{b \in B} D_b$, and the total demand at k, $K = \sum_{a \in A} S_a$. Thus, source and sink volumes meet all customer and car attribute requirements (the source node is connected to all demands and the sink node is connected to all supply) so that all supply and demand constraints are met with equality: $R + \sum_{a \in A} S_a = K + \sum_{b \in B} D_b$. By definition, every model run is feasible because if necessary all orders can be met by node r and all cars can be sent to node k.

The transportation problem formulation is given by the optimization model in Eqs. (7.1)–(7.6).

$$\text{Min} \sum_{ab \in \Phi} c_{ab} x_{ab} + \sum_{a \in A} C_k x_{ak} + \sum_{b \in B} C_r x_{rb} \tag{7.1}$$

Subject to:

$$\sum_{b \in B} x_{ab} + x_{ak} = \mathrm{S}_a \forall a \in A,\ (a,b) \in \Phi \tag{7.2}$$

$$\sum_{a \in A} x_{ab} + x_{rb} = D_b \forall b \in B,\ (a,b) \in \Phi \tag{7.3}$$

$$x_{rk} + \sum_{b \in B} x_{rb} = R \tag{7.4}$$

$$x_{rk} + \sum_{a \in A} x_{ak} = K \tag{7.5}$$

$$x_{ab} \geq 0, \text{ and integer } \forall a \in A, \forall b \in B \tag{7.6}$$

The vectors a and b on empty equipment and customer orders contribute to the cost of each assignment, c_{ab}. To the extent that a customer might accept a car that is not a perfect match or not delivered on the exact want date, the cost coefficient c_{ab} is increased accordingly. The total costs of assignments are minimized through optimal assignments x_{ab}, which is a nonnegative integer variable (7.6).

The flow of cars from each supply node to demand locations or the super sink must equal the supply at each node (7.2), and all customer orders of each type must be met from allowable supply or the super source (7.3). Sizable penalties of using phantom cars or car storage (C_k and C_r) are used to discourage flows directly from source to sink. The cost parameter C_r explicitly captures the cost of not meeting a customer order with the decision x_{rb} to supply the order from a phantom car source, r. Similarly, C_k captures the cost of the decision x_{ak} to not use car and moving it to a super sink location, k. Constraints (7.4) assure that all units of supply at the super source flow to demand nodes or to the sink, and constraints (7.5) assure excess supply and super source cars flow to the sink. In the case of a balanced network, $x_{rk} = R = K$.

7.3.3 Model Output Post Processing

The result of the model run is a set of car to car order assignments. The highest priority assignments are actual cars and orders, but the car order-specific assignments are supplemented by forecasted cars and orders.

In the case of oversupply, the model sends cars to a super sink. Car distributors "flow" cars into regions where they are needed or to storage facilities when the optimization model does not have a use for them. Such a flow generally is a function of a forecast or experience. Cars are flowed to storage yards. If no order is received, the cars remain in storage as supply. In the case of a deficit, car distributors must prioritize orders and ration cars between orders. This delicate balance is based on customer priorities and equitable treatment.

No model is perfect; specific operational complexities not known to the model (such as yard configurations which affect desirability of pulling cars in various locations), or information discrepancies (missing reportings), and the like cause distributors to make revisions to model outputs. Where the car distributors formerly worked with rules to allocate all cars, they now focus only on problematic exceptions. The exceptions are managed in the form of car assignment instruction overrides in an environment similar to the rules-based system described earlier.

7.3.4 Systems Integration

A critical component to the success of equipment distribution systems is a deep integration with operational systems. To overcome challenges of early optimization-based methods, the equipment distribution optimization engines must be deeply integrated with other production systems. Recall early attempts' solutions became "stale" as unexpected events occurred, and manual translation of model solutions into car movement instructions was laborious. Updated information with automated translation make integrated equipment distribution systems both more efficient and effective. See Gorman et al. (2010) for more details.

7.3.4.1 Optimization Engine: Customer Car Order System

The optimization engine receives live car orders (and cancelations) from customers in near real-time to ensure the engine is considering the most up-to-date demand information. This includes both individual car orders, as well as customer order preferences (car types, acceptable earliness, and lateness) described above.

7.3.4.2 Optimization Engine-Transactional Equipment Distribution System

Model results are communicated to the field via movement instructions through tight integration of the rules-based system. Model assignments are translated into car movement instructions consistent with the previously described rules, and the optimization model simply provides assignments through the rules-based system. In fact, in many cases, the rules-based system is still in use as a safety net for unallocated cars as they become available. It is worthy of note that, while the optimization engine may have a network-wide plan for current and future empty cars, only the empty cars that require a decision are acted upon operationally. Thus, the transactional rules-based system provides a means to implement the network solution one car at a time.

7.3.4.3 Transactional Equipment Distribution System: Car Movement Management and Tracking System

Car movement management and tracking systems allow railroads to monitor key operational events on the network that spur management action. Two examples are when a customer releases an empty car after it is unloaded (a "release empty" event) and another is when another railroad sends an empty car back to its owning railroad (an "interchange" event). Both of these events constitute new supply for equipment distribution to assign to customer orders; these are "trigger events" for the transactional equipment distribution system to provide disposition for an empty car. These events are automatically transferred to the equipment distribution system so that the equipment distribution system has current information on car supply. But, this is just one example of the deep integration of the car management and equipment distribution systems.

Once the origin–destination pair of the empty car is established by the model, it is translated to an empty car movement instruction. This origin–destination pair is then transferred to the car movement system for the automated creation of a "trip plan" for the empty. A "trip plan (see Ireland et al. 2004; Ahuja et al. 2007) is generated for the car. The chapter on Car Scheduling in this book provides more specifics on the development of trip plans. Similar to an itinerary in air travel, the trip plan maps the sequence of trains from origin to destination to get the car to destination with appropriate cost and service. This trip plan is a live version of the more static "operational information" (described above) that is used on input to the model for the original model preprocessing to determine the timing feasibility of empty car assignments.

As cars move across the network and, critical "events" are tracked (such as "In yard," "On train," etc.) so the progress of the move can be monitored. Generally, if cars move according to their trip plan, they are on time and will meet the timing for the customer's request. Cars that are in jeopardy of being late can, at a minimum, be managed by exception by equipment managers, or at a minimum, status updates given to the customer. But, more importantly, such event information can be integrated directly into the optimization model to optimize the network as critical events occur.

7.3.4.4 Optimization Model: Operational Systems: Decision Making Process Integration

A critical insight into the trip plan helps drive empty assignment flexibility, improved dynamic decision making, and reduce costs: At yards where cars are sorted between trains and reblocked, there is little or no incremental costs of changing to which block a car goes (Gorman et al. 2011). The car can easily be reassigned at any yard where it is reblocked. As such, the empty car may be considered "available supply" just like a release empty or interchange empty event. Thus, the empty car optimization can consider empty cars on assignment for reassignment simply by treating those empty cars as available supply, and their orders as open orders in need of a car. Any changes that have taken place since the last optimization run (i.e., new car releases, cars break down, orders are made and canceled) can be taken into account, and the entire network reoptimized. This capability allows the static optimization model to incorporate dynamic information.

The optimization model can be solved frequently because of the simple and efficient formulation. In fact, the optimization itself is a small fraction of the time to resolve the network because of the data retrieval and transfer times between systems. Railroads typically reconfigure the network ever 10–30 min so that the optimal results are "fresh." The reoptimized network also considers assignments that car distributors have "locked" in place (for example, to ensure a delivery of a particular car to a particular customer), and these will not be changed; they are treated as a hard constraint. Through near continuous resolving the network problem, a current best solution is always available. Though future events might modify the optimal solution, the solution and resulting empty car distribution instructions are automatically calculated and generated; obviating two key problems of prior systems by generating a network optimal solution and automating the communication of that plan to operations.

7.3.5 Reported Benefits

These systems are among the most success examples of the application of operations research. Railroads have claimed dramatic benefits of such systems, based on a reduction of empty car miles (7–15 %), improved customer order fulfillment and customer satisfaction, and very high return on investment.

For examples, CSX railroad reports approximately a $50 million per year benefit from their system, and BNSF has estimated $13 million; their systems cost approximately $3 to $5 million. The Union Pacific reports a 35 % return on investment, but does not report dollar amounts.

CSX also estimates that based on higher utilization of its rail car fleet, it has avoided purchasing additional $1.4 billion dollars in railcars to support its base of business.

The U.S. public also benefits from reduced truck traffic from reduced pollution, road congestion and the like; based on the CSX diversion of road traffic to rail, that benefit is approximately $50 million per year.

7.3.6 Other Implementation Considerations

7.3.6.1 User Acceptance

Car distributors must go through a big change in the way their work is conducted when such systems are implemented. Railroads report a number of strategies to improve user acceptance and model adherence. First, railroads spend copious time setting model cost and constraint parameters to improve model solutions, though balancing soft and hard costs and constraints is an ongoing challenge. Finding key modeling advocates who also know the problem domain helps build acceptance. Finally, rail customers can be uncomfortable with switching cars on their orders; only through extensive communications, changes in policy, and improved performance can customers be persuaded.

7.3.6.2 Model Thrashing

One concern facing repeated model optimization runs is possible "thrashing" from model run to model run. Thrashing occurs if model recommendations change regularly between runs, resulting in operational confusion or lack of trust in model results. Railroads limit thrashing in a number of ways. First, they leverage automatic locking or freezing of assignments as the empty car is near the customer (for example, 72 h from delivery). Second, model releases assignment information on a "need to know" basis; that is, though the model may have a number of different possibilities for an empty car, a decision is only communicated when disposition is required at interchange, empty release, or at an intermediate yard. As a result, the flexibility afforded by repeated optimization results in relatively infrequent (5–10 %) decision thrashing. Yet, the changes that are made are of great economic value.

7.3.7 Other Modeling Considerations

7.3.7.1 Endogenizing Stochasticity

One approach to addressing the inherent stochasticity facing this problem is to endogenize it within the optimization methodology. As reported above, railroads have solved a deterministic problem repeatedly as the input data change. An alternative might be to endogenize the stochasticity and solve a stochastic model, as is done in Topaloglu and Powell (2006) using an approximate dynamic programming approach. That approach is reportedly under development at Norfolk Southern. While endogenizing stochasticity has potential, specifying the form of that stochasticity can be problematic, and the complexity of modeling and implementation grows.

7.3.7.2 Including Blocking Costs in Empty Car Assignment

When cars move in collections (known as blocks), the handling cost of each car falls. To the extent that empty car assignments can consider such handling considerations, car sorting and handling can be reduced. As noted above, US freight rail organization and processes separate the assignment and routing decisions organizationally, separating these decisions; thus, such a modeling paradigm would not be appropriate. However, Joborn (1995), Holmberg et al. (1998), and Joborn et al. (2004) explore methodologies to exploit economies of scale for repositioning multiple empty cars in the same group in large blocks, effectively combining the assignment and routing decision. This line of work strives for improved equipment distribution methods for the Swedish National Railway using a deterministic capacitated multicommodity time–space network. They address uncertainty of delivery time by explicitly modeling empty capacity of train routes; resulting in better empty delivery reliability.

7.3.8 Other Areas of Application in Rail

To this point, this chapter has focused on the distribution of traditional mixed merchandise rail freight cars (e.g., box cars, condoles, etc.). Some modeling efforts have been made in intermodal and automotive as well.

In intermodal, Powell and Carvalho (1998a, b) approach intermodal flat car distribution using approximate dynamic programming. The problem is different in intermodal, as the network has fewer nodes, and individual cars are not ordered by customers (rather, they carry customer's trailers and containers).

The automotive industry has much higher concentration of shippers, therefore, shifts in shipping patterns can be a large to railroads' operations and car management. In automotive railcar management, in an attempt to reduce the empty miles traveled by empty automotive railcars between destinations of loaded shipments and origins of subsequent shipments, railroads have created a common pool of automotive railcars that are shared among railroads, and are managed by a jointly owned subsidiary, TTX Corporation. While this arrangement greatly increases the options for railcar assignment to loads (and with the increased options comes lower empty miles), it creates a challenge for the fleet sizing and management of the railcars amongst competing organizations with disparate objectives. Because of the limited size of the network (relatively few origin and destination nodes), the distribution problem is simpler. Thus, Sherali and Tuncbilek (1997) and Sherali and Maguire (2000) discuss the modeling challenges of developing fleet size strategies and help equitably distribute cars and allocate empty repositioning costs amongst the shippers. In this case, annual forecasts are developed, along with estimated monthly fluctuations. Car allocations are a function of relative demand; empty car costs are programmatically distributed amongst the participants based on the results of a time–space network, a series of cost allocation rules, and negotiated agreement amongst the carriers.

References

Ahuja RK, Jha CK, Liu J (2007) Solving real-life railroad blocking problems. Interfaces 37(5):404–419

Cordeau J-F, Toth P, Vigo D (1998) A survey American Association of Railroads (AAR). www.aar.org, 2012

Gorman MF, Acharya D, Sellers D (2010) CSX railway cashes in on optimization of empty equipment distribution. Interfaces 40(1):5–16

Gorman MF, Crook K, Acharya D (2011) North American freight rail industry real-time optimized equipment distribution systems: state of the practice. Transport Res C 19:103–114

Holmberg K, Joborn M, Lundgren JT (1998) Improved empty freight car distribution. Transport Sci 32(2):163–173

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

Joborn M, Crainic T, Gendreau M, Holmberg K, Lundgren J (2004) Economies of scale in empty freight car distribution in scheduled railways. Transport Sci 38(2):121–134

Joborn M (1995) Empty freight car distribution at Swedish Railways – analysis and optimization modeling. Ph.D. Thesis, Department of Mathematics, Linkoping University, Sweden

Jordan WC, Turnquist MA (1983) A stochastic, dynamic network model for railroad car distribution. Transport Sci 17(2):123–145

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

Powell WB, Carvalho T (1998a) Dynamic control of logistics queueing networks for large-scale fleet management. Transport Sci 32(2):90–109

Powell WB, Carvalho T (1998b) Real-time optimization of containers and flatcars for intermodal operations. Transport Sci 32(2):110–126

Sherali, Hanif D., and Cihan H. Tuncbilek. “New reformulation linearization/convexification relaxations for univariate and multivariate polynomial programming problems.” Operations Research Letters 21.1 (1997):1–9.

Topaloglu H, Powell WB (2006) Dynamic programming approximations for stochastic, time-staged integer multicommodity flow problems. Informs J Comput 18(1):31–42

Turnquist MA (1986) MOVE-EM: a network optimization model for empty freight car distribution. School of Civil and Environmental Engineering, Cornell University

Turnquist MA, Markowicz BP (1990) TIMS/ORSA conference presentation, Vancouver, BC