




Interfaces

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

An Integrated Load-Planning Algorithm for Outbound Logistics at Webb Wheel

Burcu B. Keskin, İbrahim Çapar, Charles R. Sox, Nickolas K. Freeman

To cite this article:

Burcu B. Keskin, İbrahim Çapar, Charles R. Sox, Nickolas K. Freeman (2014) An Integrated Load-Planning Algorithm for Outbound Logistics at Webb Wheel. Interfaces 44(5):480-497. <https://doi.org/10.1287/inte.2014.0760>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2014, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

An Integrated Load-Planning Algorithm for Outbound Logistics at Webb Wheel

Burcu B. Keskin, İbrahim Çapar, Charles R. Sox, Nickolas K. Freeman

Department of Information Systems, Statistics, and Management Science,
The University of Alabama, Tuscaloosa, Alabama 35487
{bkeskin@cba.ua.edu, icapar@cba.ua.edu, csox@cba.ua.edu, nkfreeman@cba.ua.edu}

We present an integrated model for simultaneous optimization of the loading and routing decisions associated with an automotive supplier's outbound supply chain. The supplier, Webb Wheel (WW), is a manufacturer of brake drums, rotors, hubs, and spoke wheels. WW accepts new orders from customers each day. Given sufficient inventory, it combines these orders into loads, releases them based on various dispatch criteria (e.g., truck-utilization, route-utilization, or penalty-based dispatch policies) and due-date considerations, and ships them in truckloads, less-than-truckloads, and containers. Dynamically changing demand information, inventory rationing, inventory interactions among orders, and lead-time considerations are some of the challenging aspects of the problem. Our optimization model is based on the decomposition of the problem into assignment and routing subproblems. The assignment subproblem determines the transportation mode and carrier choices, while considering total transportation costs. These costs depend on a variety of factors, including destination, number of drop locations on the route, and needs of customers on the route. Given the customer clusters and transportation modes from the assignment subproblem, the routing subproblem determines the sequence of drops and the true cost of the shipment using a modified traveling salesman problem. A scalable database with a graphical user interface supports the optimization model. We test our algorithm using four months of WW data and compare these data to the company's practice. Our results demonstrate the impact of transportation mode-specific capacities, customer locations, inventory availabilities, and due-date restrictions on outbound logistics costs. Since implementing our load-planning algorithm, WW has achieved cost savings of 4.4 percent over its previous load-planning process.

Keywords: coordinated logistics; vehicle routing; multimodal transportation; dispatch policies.

Webb Wheel (WW), headquartered in Cullman, Alabama, is a manufacturer of brake drums, rotors, hubs, and spoke wheels for trucks, trailers, and buses. With manufacturing plants located in Cullman, Alabama and Siloam Springs, Arkansas, WW services the commercial vehicle aftermarket. The Cullman plant serves customers east of the Mississippi River and the Siloam Springs plant serves customers west of the river. WW serves 4,642 customers with 626 unique product offerings. It strives to improve operational performance in all aspects of its business, especially in its outbound transportation operations, which have a significant impact on operating costs. We collaborated with WW in building an integrated load-planning model for simultaneous optimization of the loading and routing decisions associated with its outbound supply chain.

In particular, our primary goal is to develop an operational decision support tool to determine shipment

quantities for customer orders at both plants. The tool considers production output and available inventory, enabling the company to satisfy dynamically changing demand in a cost-effective and timely manner. Each day, WW receives new orders that trigger the master production plan and stimulate inventory buildup. Given sufficient inventory, WW builds loads daily and releases them based on various dispatch criteria, including truck-utilization, route-efficiency, and due-date considerations. For outbound shipments, WW relies on several common carriers that provide truckload (TL), intermodal (IM) (i.e., container), and less-than-truckload (LTL) delivery options.

The main challenge associated with this problem is the dynamically changing and incomplete demand information. WW accepts orders any time of any given day; however, a lack of information about the size and location of the next customer order imposes challenges

for rationing the inventory, building the routes, and dispatching the loads in a cost-effective manner. Limited inventory presents an additional complication. In cases in which sufficient inventory is not available to cover all the open orders, WW needs rationing rules to decide how to allocate inventory among all open orders. In addition, an inventory-rationing decision made on a given day can impact future periods because of inventory interactions among orders and consolidation opportunities. These issues force WW to look for load-dispatch policies that eliminate waste and take advantage of consolidation opportunities. Therefore, one of our secondary goals is to characterize and evaluate load-dispatch policies. Finally, the products offered by WW are very large and heavy; for example, a typical brake drum and rotor measures about 18×8 inches and weighs between 60 and 120 pounds. Therefore, splitting the customer orders is common practice (Archetti and Speranza 2012, Bolduc et al. 2010, Lei et al. 2012). Order splitting also helps achieve desired truck-utilization levels.

Given the aforementioned characteristics, the problem under consideration is a multiple-item, multimode, heterogeneous vehicle routing problem (VRP) with split delivery and incomplete information. To solve this problem, we develop a solution that decomposes the problem into two subproblems in which we assign customer orders into clusters, and route these orders, while simultaneously considering customer- and carrier-specific limitations. We know of no other published work that addresses this problem (i.e., a multiple-item, multimode, integrated load-planning and routing problem with incomplete information). Hence, both the model formulation and solution development are novel contributions to the literature.

Running our algorithm for a given day allows WW to assign a route to each open order for which it has sufficient inventory to fill. However, dispatching all the routes that the algorithm generates may not be in WW's best interest because of potential consolidation opportunities (i.e., opportunities to consolidate shipments) that may arise later. To determine which routes to dispatch, we evaluate three policies: (1) weight-based truck utilization (TU), (2) route-based utilization (RU), and (3) penalty-based dispatch (PP). TU is the percentage of the total weight in a truck to its capacity. RU is the percentage of the total weighted distance to

the total ton-miles available in the route, whereas PP combines RU with the cost of the route to help dispatch the cheaper routes. RU and PP are novel dispatching policies for evaluating the quality of a route, and are especially beneficial for routes with multiple drops.

We compare the performance of the algorithm against WW's actual practice using four months of shipment data for both the Cullman and Siloam Springs plants. The results demonstrate the importance of transportation mode-specific capacities and the impact of customer locations and their geographical spread, availability of inventories, and due-date restrictions on outbound distribution costs; the results also provide evidence that an efficient load-planning process with consolidation prospects presents an opportunity for cost savings over the costs in the current system.

We organized the remainder of the paper as follows. In the next section, we include a brief review of the related literature. Before presenting the specifics of our models, we discuss cost structure and restrictions of transportation modes in *Transportation Mode Characteristics*. In *Modeling Consideration*, we provide existing practice of WW and the model formulations for assignment and routing problems. We then explain the details of the dispatch policies. In *Computational Results* we present how each dispatch policy performs and compare the overall cost performance of our algorithm against WW's actual performance. Finally, we conclude the paper with a summary of contributions and offer some future research directions in *Conclusions and Future Research Directions*.

Literature Review

WW's outbound supply chain problem is a form of the VRP. VRPs have been the subject of intensive research for more than 50 years, because they are difficult combinatorial optimization problems, which are relevant in many application fields, including transportation, logistics, communications, manufacturing, military, and relief systems. The broad range of applications with routing issues lead to the definition of many VRP variants with additional characteristics and constraints. Vidal et al. (2013) provide a review of multiattribute VRPs. Other important references on VRPs include Bräysy and Gendreau (2005a, b), Golden et al. (2008), and Toth and Vigo (2002).

The problem at hand is classified as multiple-item, multimode, heterogeneous VRP with split delivery, incomplete information, and other practical constraints. Dror and Trudeau (1989, 1990) originally introduced the split-delivery VRP (SDVRP). Its main characteristic is that a customer's demand can be split among a number of delivery vehicles. Some researchers (Chu 2005, Bolduc et al. 2007, Côté and Potvin 2009, Bolduc et al. 2010) consider SDVRPs with private fleet and common carrier options. Chu (2005) develops a heuristic algorithm to route the private trucks and to make a selection of LTL carriers by minimizing a total cost function. Considering a similar problem, Bolduc et al. (2007) show that the solution in Chu (2005) has errors, and they develop a two-phase heuristic with λ -exchange based on an improvement heuristic. Côté and Potvin (2009) describe a tabu-search heuristic for a SDVRP in which the owner of a private fleet can either visit a customer using one of the fleet vehicles or assign the customer to a common carrier. Expanding this line of work, Bolduc et al. (2010) also consider a tabu-search-based algorithm for a dynamic, multiperiod SDVRP in which each customer may be served by private fleet vehicles, common carriers, or any combination of these.

Our work departs significantly from this existing line of work in two ways. First, WW does not own its fleet; it relies on carriers that ship using TL, LTL, and IM policies. We consider a SDVRP with multiple transportation modes, each with specific constraints and mode-dependent cost functions. Second, we approach the problem using a cluster-first, route-second type of solution approach as opposed to local- or tabu-search-based metaheuristics. Many VRP problems are solved using a cluster-first, route-second approach (Bramel and Simchi-Levi 1995, Fisher and Jaikumar 1981). These problems range from theoretical applications of VRP variants (Stacey et al. 2007, Natarajathinam et al. 2012) to practical applications, as Ambrosino and Sciomachen (2007) present for a food-distribution problem in an Italian company. We customize this solution approach in a unique way to handle multimode-related restrictions and costs. We discuss these specifics for the various transportation modes in the *Transportation Mode Characteristics* section.

We do not explicitly consider inventory decisions at the WW plants or by its customers, as they are considered in the well-known inventory-routing problem

(Bertazzi et al. 2008, Campbell et al. 2002, Gaur and Fisher 2004). However, we acknowledge the possibility that WW may not have enough inventory to fulfill all its open orders. With limited inventory, WW gives priority to critical orders whose promise-to-ship date has already passed or is within a few days. In the inventory literature, this practice is known as inventory rationing, and is used to differentiate among multiple customer classes, typically to serve higher-priority customers (Cattani and Souza 2002, Deshpande et al. 2003, Gayon et al. 2009, Hung and Hsiao 2013). Our work is the first paper in the VRP literature to incorporate inventory rationing based on due dates.

Transportation Mode Characteristics

As we mention previously, WW relies on common carriers that use TL, IM, and LTL options to fulfill customer orders. Both plants can accommodate all three transportation modes. Each transportation mode is subject to different capacity and transportation mode-specific constraints: coverage limitation, cost structure, and order drop sequence. If WW decides to split an order, it might ship the split order via different modes. Next, we discuss these specifics for each transportation mode.

Truckload (TL) Characteristics

WW procures full truckloads from private carriers and fulfills most customer orders using the TL delivery method. The cost of a TL route is a function of the distance traveled, the mileage rate, and the number of drops on the route. The location of the final customer on a route determines the mileage rate for that route. Using location-based rates is a common practice in the trucking industry to penalize routes ending in locations at which backhaul opportunities are low. A mileage surcharge based on the weekly average national fuel price is also added. Depending on the total route mileage and mileage rate, total cost is determined in one of two ways:

(1) Minimum charge met: If the total mileage cost, calculated using the total route distance times the mileage rate, is greater than or equal to the minimum charge of the TL carrier, then the total TL cost is given as the sum of the total drop charge and the total mileage cost, amended by the surcharge rate.

(2) Minimum charge not met: If the total mileage cost is less than the TL carrier's minimum charge, then the TL carrier charges a minimum charge and

calculates the total cost as the sum of the total drop charge, total route distance times surcharge rate, and minimum charge.

TL mode also requires two additional considerations. First, the capacity of the trucks is limited. To avoid high penalties for exceeding the weight capacity, WW uses a capacity of 43,500 pounds in making load decisions. The second consideration is the final drop location-based mileage rate. The TL carrier can dictate the final drop location for some routes to avoid incurring empty miles. Specifically, if a route includes a drop to the New York Islands, such as Manhattan Island or Long Island, this drop should be the last stop on the route. This restriction is motivated mainly by the high toll and bridge costs around this area. The trucking companies force this restriction on their customers because of excessive exploitation of the lower rates in the neighboring states, such as New Jersey. In the trucking industry, this restriction is known as the “elimination of the hook”; see Figure 1.




Figure 1: TL carriers reject any route that contains “the hook.”

Intermodal (IM) Characteristics

The use of intermodal (IM) shipments is fairly recent by WW. Because of the potential transportation cost savings IM shipments can provide, the company desires to increase its utilization of this transportation mode. When using IM shipments, customer orders are first loaded onto a container and then transported to a ramp location via train. At the ramp location, the container is loaded onto another truck for final delivery. IM delivery has two cost components: a ramp location-based fixed cost and a drop charge per order.

Intermodal constraints complicate the formation of routes and the use of ramps. In particular, the total service time from a ramp location is limited. Hence, the customers assigned to an IM route must be serviced within a total service time limit. That is, the final delivery trucks (i.e., the trucks that make the deliveries to the customers) must return to the original ramp location after unloading the shipments within a fixed time. This period includes the total traveling time spent on the road and the fixed time for unloading each customer order. One challenge associated with this restriction is that because the load is determined before the actual route is known and without the actual routing information, we cannot determine if the load will violate the total service time limit constraint. To overcome this limitation, we approximate the total travel time of the loads within a container using the star-distance method (Bramel and Simchi-Levi 1995) (see Figure 2). Although this approximation overestimates the total distance traveled, it ensures the feasibility of the IM container routes.

In addition to this service time limit constraint, as in the TL mode, the containers are subject to a capacity constraint. Excluding the product packaging, this limit is set at 41,000 pounds.

Less-Than-Truckload (LTL) Characteristics

LTL shipments are generally used to ensure that due dates are met. However, the LTL delivery method may be economically attractive if a delivery location is isolated. LTL delivery cost is a complicated piecewise function of the customer order location and the total weight of the order. It is calculated per order and the fixed cost associated with a LTL shipment is zero.




Figure 2: We use star-distance approximation to calculate the total distance from the ramps.

Modeling Considerations

WW's Current Practice

Two planners at WW manually generate routes each day; each planner is responsible for a specific plant. For each plant, the corresponding planner typically determines the route as follows. (1) The planner selects as an anchor point the open order that has a delivery location farthest from the plant, and (2) manually adds the additional open orders to the route one by one, trying to preserve a straight-line route from the plant to the anchor point until the capacity limit is reached. If the route is fairly straight and has a high TU rate, the planner dispatches the route. In contrast, our algorithm determines the routes from the same plant through an optimization-based decomposition approach. In making a dispatch decision, we consider both the TU rate of the route and its RU or PP. As an example, we present a case for the Cullman plant (see the black square in Figures 3(a), 3(b), 3(c), and 3(d)). On day 1, the plant has two open orders, indicated by black circles in Figure 3(a). The orders originate from Clarksville, Tennessee and Cincinnati, Ohio, and each requires half a truckload. The planner, following the described protocol, sets the order from Cincinnati as an anchor point and then adds the order from Clarksville to the route. Because the TU rate of the route is 100 percent, the planner dispatches the route that we refer to as WW Route 1. On day 2, two new orders, indicated by black triangles in Figure 3(b), arrive from Nashville, Tennessee and Columbus, Ohio, each requiring half a

truckload. To handle the new orders, the WW planner dispatches a second route, WW Route 2, to Nashville and Columbus by following the same logic. Figure 3(c) shows WW's Route 1 (solid line) dispatched on day 1 and WW's Route 2 (dashed line) dispatched on day 2. The total costs of WW Route 1 and WW Route 2 are \$1,120.35 and \$1,205.82, respectively. Table 1 provides the total costs and the corresponding RU rates of 63 and 64 percent for WW Route 1 and WW Route 2, respectively.

Our algorithm would generate the same route as the WW planner on day 1; however, it would not dispatch the route because the associated route utilization is lower than the threshold value of 75 percent. On day 2, our algorithm will dispatch the two routes shown in Figure 3(d), the University of Alabama (UA) Route 1 and UA Route 2. Both routes have RU values much higher than the acceptable preset value. Moreover, the routes that our algorithm determines save \$751.72 and 292 miles. Next, we discuss our optimization-based approach that generates these savings.

Optimization-Based Algorithm

Our solution considers all open orders that are available at run time and for which sufficient inventory is available. Each order is assigned to a route(s) via a specific transportation mode regardless of its order date. Therefore, we do not consider a time index for shipments or routes; however, we consider due dates while rationing limited inventory over open orders and later deciding which loads to dispatch. To determine the order clusters, ration inventory, select transportation modes, and route the orders, we divide the overall problem into two subproblems: a generalized assignment problem and a routing problem.

The first subproblem, the generalized assignment problem, is based on an extension of the capacitated-concentrator location problem. This assignment problem determines the clusters of orders for each transportation mode, while considering transportation mode-specific cost structures and restrictions, as we explain in the *Transportation Mode Characteristics* section. Additionally, the assignment problem ensures the demand from critical (i.e., soon-to-be-due) orders is satisfied first.

Given the clustering information and transportation mode selection for each cluster from the assignment problem, the routing subproblem determines the




Figure 3: The four graphics illustrate differences in route formation as a result of Webb Wheel (WW) and University of Alabama (UA) dispatching policies.

sequence of drops and the true cost of the shipment using a modified traveling salesman problem (TSP).

The Assignment Problem

The capacitated-concentrator location problem (Bramel and Simchi-Levi 1995) defines unique setup costs for seeds of concentrators and a specific connection cost from each terminal to each seed. The objective is to minimize the sum of setup and connection costs

without violating any of the location-related constraints. In a similar fashion, we define a seed as a virtual anchor point for a cluster of orders that can be served by only one vehicle (truckload, container, or one LTL shipment) without violating any constraint of the routing problem. More specifically, we define seed sets for each transportation mode. Our model includes TL, IM, and LTL seeds. Additionally, we define customer order sets and product sets. For each customer order,

Routes	Drop order	Mileage	TU (%)	RU (%)	Mileage rate (\$)	Cost (\$)
WW 1	AL-Clarksville-Cincinnati	485	100	63	2.31	1,120.35
WW 2	AL-Nashville-Columbus	522	100	64	2.31	1,205.82
	WW total mileage	1,007			WW total cost (\$)	2,326.17
UA 1	AL-Cincinnati-Columbus	522	100	90	2.31	1,205.82
UA 2	AL-Nashville-Clarksville	193	100	87	1.91	368.63
	UA total mileage	715			UA total cost (\$)	1,574.45
	Mileage savings	292			Cost savings (\$)	751.72

Table 1: The table shows routes and costs for Webb Wheel (WW) and University of Alabama (UA) dispatching policies.

we calculate the assignment cost of assigning that order to each member of the seed sets.

Decision Variables: In our solution approach, we introduce three sets of decision variables. The first represents the seed location setup and checks whether a seed location is selected; the second represents assigning orders to seed locations; and the third represents the amount of a product shipped for a customer order through a specific seed location.

Constraints: To handle the transportation mode-specific capacities, we define a new capacity parameter. Depending on the transportation mode, the capacity parameter could be set to the truck capacity or to the container capacity. In the appendix, we present the mathematical form of all the constraints. Constraints (1) state that the capacity of a seed, regardless of its type, is not violated. Note that the capacity limitation is weight-based because the products offered are heavy.

To define the necessary amount of shipment for open orders, we need to ensure that the number of units shipped is less than the order quantity and the inventory, whichever is smaller. Constraints (2) and (3) state this restriction: the amount of shipment via all the transportation modes for a particular product of a particular order should be less than the order demand and available inventory for the product, respectively.

One complicating constraint in our model relates to allocation of available inventory to customer orders. WW sometimes has inventory shortages. In such cases, it allocates inventory to critical orders that will be due soon. The criticality is set as a control parameter. For orders that are due within three days, the criticality parameter is set as 1; otherwise, it is set as 0. Constraints (4) ensure that the demand for critical orders is satisfied first from the available inventory.

Constraints (5) state that the demand of the remaining orders is satisfied as much as the inventory permits. These two constraints allow us to ration the limited inventory across critical items.

To determine the number of truck-related seed locations, we calculate the number of trucks required to fulfill orders in each unique five-digit zip code location. Then, we create one truck-seed location for each truck required. Truck-seed location cost and the assignment cost for an order to a truck-seed location are calculated based on truck cost structure. In the appendix, we provide detailed calculations for both costs.

As we mention previously, the TL carrier can dictate the final destination of a route. To handle this restriction, we modify the assignment cost to be a very large number, if the order location is around New York Islands and the seed location is not. Otherwise, the assignment cost stays as calculated.

As we discuss in the *Transportation Mode Characteristics* section, to ensure the feasibility of the IM routes, we use the star-distance approximation. For this purpose, we add constraints (6) to ensure that the sum of the drop times and the round-trip travel time from a container-seed location to the customer order location is less than or equal to the total service time.

To determine the number of container-related seed locations, we create one seed location for each customer-ramp pair if the customer is within the total service time from the ramp. If an order's total weight is more than the container's capacity, the number of container-seed locations is calculated as the rounded-up ratio of the total weight of order to the container capacity. We note, however, that although we determine the number of container-related seed locations by looking at the customer-ramp pairs, we do not assign any

customer orders to any container-seed locations; that is, any order can be assigned to any IM seed as long as constraints (6) are satisfied.

Unlike the other seed-location types, an LTL seed location is tied to customer orders. To implement this restriction for LTL seed locations, we use an indicator parameter. Specifically, we add constraints (7) to ensure that a customer order can only be assigned to a LTL seed location if the LTL seed location is created for that customer order. LTL delivery cost is a function of the customer order location and the total weight of the order. In our model, we set the LTL seed cost to the LTL delivery cost and the assignment cost to zero. One limitation of this approach is that even if an order is partially shipped via LTL, the model still uses the LTL cost calculated for the whole order.

The appendix provides a summary of the notation and the overall formulation for the assignment problem.

The Routing Problem

The solution of the assignment problem provides the cluster of orders that is to be shipped via specific transportation mode. However, the assignment problem does not provide the drop sequence. Determining the drop sequence is important only for TL-based shipments. To find the optimal drop sequence in a TL route, we use N-MTSP, a modified TSP formulation proposed by Dantzig (1963), which includes N different stopping points in a modified TSP formulation. The appendix shows the overall formulation of the routing problem.

The integer solution of the N-MSTP formulation clearly forms a tour (Flood 1956). Although many TSP formulations exist, we choose N-MSTP for two reasons. First, we can easily dictate the start and end locations of the tour. Second, we can calculate the rate per mile of the tour. One drawback of this formulation is that the objective of N-MSTP is a quadratic function. Although most commercial solvers can handle quadratic functions, we eliminate the nonlinearity by parameterizing the final customer. With this linearization, the advantage of this approach is twofold: (1) solving the linear N-MSTP N times is much faster than solving one quadratic N-MSTP, especially for our problem, because (N —the number of drops in a route) is typically small, and (2) we can incorporate the TL carrier restriction of orders to New York Islands. To linearize N-MTSP, we

remove constraints (15) and (16), fix the last customer order in a route, and set the rate of the route to the rate of the last customer location. To find the optimal route, we use each eligible customer in a route as the last customer in the route. The customer order that provides the lowest objective function is the last customer in the optimal route.

To calculate a route cost of TL and container shipments, we use the costs discussed in the *Transportation Mode Characteristics* section.

As we discuss previously, the cost of LTL routes may be higher than the actual value, because LTL seed cost is evaluated at the full order weight; however, the possibility exists that only part of the order is fulfilled with LTL. To eliminate this overvaluing, we recalculate the cost of LTL routes at the end with updated order weights.

Dispatch Policies

The routing decisions for a particular day affect all future periods because of inventory interactions among orders and consolidation opportunities. Once the algorithm is run for a given day with inventory and order information, all open orders are assigned to a route. However, dispatching all the routes that the algorithm generates may not be in WW's best interest; consolidation opportunities may arise later. Thus, a dispatching policy is needed to determine which routes to release and which to hold back. In this section, we discuss the details of three dispatching policies: truck-utilization, route-utilization, and penalty-based dispatch policies.

Truck-Utilization (TU) Dispatch Policy

The TU policy is commonly employed by shippers and common carrier companies. TU is calculated as the ratio of total weight on the truck (or container) to its capacity at the time it leaves the plant.

The basic concept is that if a company utilizes the total volume of a truck (i.e., cubes out the truck) for high-volume and lightweight items, or if it utilizes the total weight limit of the truck for heavy items (e.g., as in WW's case), the company would need fewer trucks. It would eliminate the extra route(s) that would be required if the trucks were not full, and thus save money. Because all parties in the trucking industry understand this policy, it is implemented widely. WW is also concerned about sending trucks as full as possible

in terms of weight; otherwise, it would send more trucks than necessary and generate additional costs.

Route-Utilization (RU) Dispatch Policy

To facilitate order consolidation from similar locations, WW uses a RU policy to dispatch loads. This policy seeks to attain high truck utilization throughout a route, rather than only measuring utilization as a truck or container departs the plant, as in the TU policy. We present exact formulations of TU, RU, and PP policies in the appendix.

Compared to TU, RU is more useful in capturing the full potential of the route, as we explain in the example in *WW's Current Practice* section and Figure 3. The RU policy helps in evaluating the utilization of the truck throughout the route and helps dispatch the correct routes, given the uncertainty in the size and location of future orders.

Penalty-Based (PP) Dispatch Policy

WW's overall goal is to reduce transportation costs. Obviously, better route planning and order consolidation help reduce the costs. As an additional alternative, we consider a route-dispatching policy that is based on the cost of the inefficiency of a route. For example, if a route is cheap (i.e., the rate, the total mileage, or both, are much lower than the average), inefficient dispatching will have little effect on overall transportation costs. However, if a route is expensive (i.e., the distance and rate are higher than average), ensuring that the truck capacity is used more efficiently is important.

This route penalty represents the opportunity cost of using a less efficient route. Obviously, if the route is cheaper, even if the route utilization is low, the opportunity cost will be lower than a more expensive route with the same route efficiency. In contrast, for expensive routes, the opportunity cost will be more sensitive to the route utilization.

To assess the quality of a route based on this penalty, we determine an acceptable penalty, or a cutoff level, using historical route and cost information from WW. For this purpose, we calculate route penalties for all the WW routes for a four-month period and sort them in descending order. As in the Pareto principle (i.e., 80–20 rule), we notice that we may attribute approximately 80 percent of the total route penalties to roughly 20 percent of the routes. That is, most of the

routes are built efficiently, but 20 percent of them have significant route penalties, ultimately affecting WW's profit. To avoid expensive routes and transportation waste, we set the cutoff value as the minimum route penalty observed from the 20 percent of high-penalty routes. Moreover, we do not dispatch a route if its calculated penalty is higher than this lower bound. For each plant, we calculate this cutoff value separately using routes initiated from the Cullman and Siloam Springs plants, respectively.

Once we determine the cutoff level, we dispatch a route if its PP is less than the respective cutoff value; otherwise, we hold it for further consolidation opportunities. We handle the orders that are late or due soon in a different way, as we specify in the *Comparison of Dispatch Policies* section.

Computational Results

We compare the performance of our algorithm with WW's practice. For this purpose, WW provided four months of data, including customer orders, inventory, and actual routes. During this period, WW received more than 2,300 unique orders for more than 10,000 line items. On average, each order had approximately five different products. The data were generated cumulatively at the end of each workday. That is, each day's data also included data starting from the first day. Because the orders for each plant are handled separately, we run our algorithm separately for each day and each plant.

We implement our solution approach in C++ programming language, Concert Technology, and IBM ILOG CPLEX 12.4. We perform the runs on a machine with a 3.6 GHz Intel Xeon processor and 64 GB RAM. We use Microsoft Visual Basic for Applications 7.0 to communicate between the database and optimization algorithms. We observe that our solution approach for the optimization algorithm results in very reasonable run times using WW data. Specifically, the solution for each plant and for one day is obtained within one hour.

Before discussing the results, we present the details of the graphical user interface (GUI) for the database and performance of TU and RU as dispatch policies. Then, we present a detailed cost breakdown for both plants. To ensure data confidentiality, we scale all dollar values with arbitrary coefficients, but keep the percentage values the same.

Graphical User Interface

A scalable database supports our optimization model. The database has a GUI that represents the various quantities of interest when modeling the outbound supply chain. The database collects and aggregates data, checks for data integrity, performs any necessary cleanup, and generates interdistances between WW customers and plants. The database also acts as an interface for calling the optimization models and generating reports based on model output. The key components of the database include application program interfaces for Microsoft MapPoint, optimization algorithms, and report generation tools. The database also allows users to change common carrier providers, capacities, costs, and demand data. Figure 4 shows a snapshot of the GUI.

Comparison of Dispatch Policies

We demonstrate the savings opportunities of the RU dispatch policy by comparing it to that of the TU. In particular, we use the first month of data and dispatch trucks at five different TU rates. Specifically, we consider dispatching trucks at TU rates of 75, 80, 85, 90, and 95 percent. With each dispatch rate, we calculate the total cost over the period and total weight shipped.



Figure 4: The graphical user interface (GUI) for the optimization model facilitates access to the data, optimization models, and results.




Figure 5: A RU dispatch policy provides better savings opportunities in comparison to that of a TU policy.

In Figure 5, we compare each rate using a ratio of total cost to the total weight shipped, that is, cost per pound (CPP). The reason for using this ratio, as opposed to a pure cost comparison, is to eliminate the effect of excessive consolidation and not shipping, and hence, not incurring any costs. As TU increases, the CPP over this pilot period decreases and levels off at roughly 90 percent. Similarly, we dispatch trucks with RUs of 50, 60, 70, 80, and 90 percent. Even with the lowest RU rate of 50 percent, the CPP returned for the pilot time (approximately 4.7 cents) is lower than the best CPP returned by TU (5.0 cents) alone. As we increase the RU rates, the CPP continues to decrease.

After observing the importance of RU in cost savings and TU in practice, we decide to combine them into a dispatch rule. To dispatch a route, we use the following rule: (1) If both TU and RU of a route are larger than the preset TU and RU rate, respectively, dispatch the route; (2) else, if the sum of the TU and RU of a route is larger than a preset level, dispatch the route; (3) else, do not dispatch.

Item (2) identifies the routes that are high quality but slightly miss the cutoff criteria for both TU and RU. A route may have slightly lower utilization (TU or RU) in one dispatch criterion, but it may still be a high-quality route.

A problem arises when a route containing a late order fails to meet the described criteria. WW does not allow orders with available inventory to be late. To address this situation, we develop two late-order dispatching methods: (1) direct-cost comparison, and (2) CPP comparison. In the former, if the algorithm fails

to dispatch a route containing one or more late orders, we first add up the LTL shipment costs for all late orders in the route. We then compare this sum with the route cost. If the sum of LTL shipment costs for the late orders is lower than the route cost, we ship the late orders via LTL; otherwise, we dispatch the route even if it does not meet the efficiency dispatch criteria. In the CPP method, after summing the LTL shipment cost, we divide the sum by the total weight of the late orders. Similarly, we divide the route cost by the total weight of the route. If the LTL shipment cost-to-weight ratio is lower than the route cost-to-weight ratio, we ship the late orders via LTL; otherwise, we dispatch the route.

To determine the critical efficiency cutoff level and late-order dispatching method to use, we ran our algorithm on a sample data set. We observed that the CPP method always outperforms the direct-cost comparison method in both plant locations. Thus, we set the utilization rate to 90 percent and use the CPP method as a late-order dispatching method for both

plants. A RU of 85 percent gives the best result for Cullman; however, a RU of 75 percent performs better for Siloam Springs. Therefore, in our final test, we use these values for Cullman and Siloam Springs, respectively. Sometimes, a high-quality route may have slightly lower utilization (TU or RU) in one dispatch criterion. To avoid not dispatching the high-quality routes, we determine if the sum of TU and RU values is higher than 175 percent (85 percent + 90 percent) for Cullman or 165 percent (90 percent + 75 percent) for Siloam Springs. If so, we dispatch the route.

A Detailed Analysis of Cost Breakdown and Insights

In Table 2, we present the cost analysis for Cullman, including total shipments over four months in pounds (lbs.), total cost, CPP, and percentage savings of our algorithm, compared to WW's practice. We refer to our algorithm as integrated loading-routing optimization (ILR). During the test period, WW shipped more than 19 million pounds of products at a cost of

	Total shipment (lbs.)	Total cost (\$)	No. of shipments	CPP (\$)	% save
WW	19,347,372.82	1,000,000.00	539	0.05169	
ILR (TU:90/RU:85) TC: 43,500 lbs. DD: 0 day	19,316,138.89	1,004,029.07	488	0.05198	−0.57
Relaxed ILR TC: 43,650 lbs. DD: 3 days	19,019,864.92	971,866.92	457	0.05110	1.14
DD relaxed ILR TC: 43,500 lbs. DD: 3 days	19,012,620.95	977,038.93	464	0.05139	0.58
TC relaxed ILR TC: 43,650 lbs. DD: 0 day	19,358,273.29	1,002,738.14	474	0.05180	−0.22
ILR (PP) TC: 43,500 lbs. DD: 0 day	19,351,543.69	1,009,496.01	487	0.05217	−0.93
Relaxed ILR (PP) TC: 43,650 lbs. DD: 3 days	19,056,236.26	972,441.06	462	0.05103	1.27
DD relaxed ILR (PP) TC: 43,500 lbs. DD: 3 days	19,027,553.37	973,820.02	471	0.05118	0.98
TC relaxed ILR (PP) TC: 43,650 lbs. DD: 0 day	19,358,273.29	1,008,809.11	480	0.05211	−0.82

Table 2: The table shows cost and weight breakdowns of shipments from the Cullman plant. To ensure data confidentiality, we scale all dollar values with arbitrary coefficients, but keep the percentage values the same.

\$1,000,000 in 539 shipments, including TL, IM, and LTL transportation modes. For Cullman, we run ILR with two dispatch policies. The first policy uses a TU of 90 percent and RU of 85 percent; the second policy is a PP policy in which we set the penalty level at \$460. Per WW's request, ILR uses a truck capacity of 43,500 pounds with no due-date extensions. Surprisingly, we observe a lack of savings with both policies (−0.57 percent with TU-RU and −0.93 percent with PP) when we compare ILR to WW's performance, although ILR ships a total amount that is comparable to WW, with fewer shipments. A careful inspection of the results reveals that WW occasionally violates the constraints imposed on ILR claiming that exceptions occur. These exceptions are generally capacity and due-date violations. Specifically, we find 20 truck-capacity violations, sometimes by as much as 1,000 pounds, and due-date violations of up to two weeks. However, achieving a fair comparison is not easy. When we modify ILR to match WW's observed violations, our results show that most routes are overloaded and (or) contain late orders.

As a compromise, we relax the truck capacity by only 150 pounds and extend the due date by three days. With this relaxation, we save about 1.14 and 1.27 percent with TU-RU and PP policies, respectively, over WW's practice. WW welcomed the savings for two reasons. First, it manufactures and distributes products with relatively low profit margins; therefore, any operational improvement is significant. Second, the test data correspond to a period of low-demand volume for the company. Since WW implemented our optimization-based load-planning algorithm, it has achieved annualized savings of 4.4 percent in its outbound transportation costs.

To understand the sources of these savings, we devise two additional runs for each dispatch policy: (1) due-date extension (DD relaxed ILR), and (2) truck-capacity relaxation (TC relaxed ILR). We find that due-date relaxation contributes the most to the savings with both dispatch policies, principally from the savings from LTL shipments. This is evident from Table 3 in which we present the transportation mode details of the shipments.

In Table 3, we show the details of TL, IM, and LTL shipments for each version of ILR. We again present the total amount shipped with each mode, total cost, number of shipments, and CPP. Although the majority

of shipments use the TL mode, most savings come from efficiently handling LTL shipments. Specifically, WW ships 77,110 pounds of product in 86 LTL shipments. In our runs, the relaxed ILR ships 160,851 pounds in 19 shipments with TU-RU. Additional runs suggest these results are consistent.

Next, we compare the performances of the TU-RU and PP dispatch policies. For the restricted version of ILR, TU-RU performs better (although some negative savings are present) than the PP policy. However, PP performs better for the relaxed ILR and DD relaxed ILR, because of the PP policy's strength in evaluating the use of truck capacity and due dates. In both cases, ILR with PP policy ships more products in a less costly way.

The results for the Siloam Springs plant (see Tables 4 and 5) are very similar to the results for the Cullman plant. With Siloam Springs, however, we observe savings of 2.02 percent over WW's practice, even with the restricted ILR using the TU-RU policy. Using the relaxed ILR, truck and container capacities increase by 200 and 750 pounds, respectively. This is consistent with the average shipment weights of WW from the Siloam Springs plant. Furthermore, due date is extended by one day. With these relaxations, the savings increase to 4.93 percent with TU-RU and 3.05 percent using PP. The savings result from both LTL (an improvement of as much as 92 percent) and TL shipments (an improvement of as much as four percent).

ILR outperforms WW's practice for two primary reasons, even in the restricted case, at the Siloam Springs plant. First, the Siloam Springs plant serves a larger geographic area. Hence, a more sophisticated approach, such as ILR, is necessary to generate efficient routes. Because Cullman serves a much smaller and more densely populated region, identifying consolidation opportunities is easier for WW. Second, the number and size of outstanding customer orders from the Cullman plant is larger, also simplifying the process of identifying good routes. This is not the case for Siloam Springs; consequently, this plant has a greater need for a refined tool to identify quality routes.

In terms of transportation modes, a major difference between the plants is the extensive use of IM shipments that comprise approximately 10 percent of shipments. Therefore, the ability of ILR to check IM constraints and generate feasible, cost-effective routes is important.

	Total shipment (lbs.)	Total cost (\$)	No. of shipments	CPP (\$)	% save
WW	19,347,373	1,000,000.00	539	0.0517	
TL	19,232,073	982,764.67	452	0.0511	
IM	38,190	3,063.79	1	0.0802	
LTL	77,110	14,171.54	86	0.1838	
ILR (TU:90/RU:85)	19,316,139	1,004,029.07	488	0.0520	−0.57
TL	18,886,774	976,636.30	440	0.0517	−1.19
IM	38,190	3,063.79	1	0.0802	0.00
LTL	391,175	24,328.98	47	0.0622	66.16
Relaxed ILR	19,019,865	971,866.92	457	0.0511	1.14
TL	18,820,824	958,323.59	437	0.0509	0.36
IM	38,190	3,063.79	1	0.0802	0.00
LTL	160,851	10,479.54	19	0.0652	64.55
DD ILR relaxed	19,012,621	977,038.93	464	0.0514	0.58
TL	18,799,709	961,095.17	437	0.0511	−0.04
IM	38,190	3,063.79	1	0.0802	0.00
LTL	174,722	12,879.97	26	0.0737	59.89
TC ILR relaxed	19,358,273	1,002,738.14	474	0.0518	−0.22
TL	19,069,962	981,914.13	442	0.0515	−0.76
IM	38,190	3,063.79	1	0.0802	0.00
LTL	250,121	17,760.21	31	0.0710	61.36
ILR (PP)	19,351,544	1,009,496.01	487	0.0522	−0.93
TL	19,108,816	991,213.95	456	0.0519	−1.51
IM	38,190	3,063.79	1	0.0802	0.00
LTL	204,538	15,218.28	30	0.0744	59.52
Relaxed ILR (PP)	19,056,236	972,441.06	462	0.0510	1.27
TL	18,977,267	965,616.99	448	0.0509	0.43
IM	38,190	3,063.79	1	0.0802	0.00
LTL	40,779	3,760.27	13	0.0922	49.83
DD relaxed ILR (PP)	19,027,553	973,820.02	471	0.0512	0.98
TL	18,848,636	959,406.99	447	0.0509	0.39
IM	38,190	3,063.79	1	0.0802	0.00
LTL	140,727	11,349.23	23	0.0806	56.12
TC relaxed ILR (PP)	19,358,273	1,008,809.11	480	0.0521	−0.82
TL	19,139,993	991,054.90	451	0.0518	−1.33
IM	38,190	3,063.79	1	0.0802	0.00
LTL	180,090	14,690.42	28	0.0816	55.61

Table 3: The table gives the transportation mode breakdown from the Cullman plant.

Table 4 shows that the TU-RU policy outperforms the PP dispatch policy in Siloam Springs. We set the penalty for Siloam Springs to \$750 based on historical data. The higher penalty and geographical differences cause this performance difference. Additionally, we observe that DD relaxed ILR with PP policy returns slightly higher savings than the relaxed ILR with PP policy. Understanding this counterintuitive result requires two observations. First, the performance of the ILR is sensitive to the due-date extension. Second, the increased truck and container capacities of the relaxed ILR policy create opportunities for consolidation, but also increase

the penalty for routes. Thus, a route that is acceptable with DD relaxation of ILR may not be acceptable with the relaxed ILR. Furthermore, depending on the route cost, this route may not be eligible to dispatch via truck or container. When we analyze the detailed cost breakdown in Table 5, comparing relaxed ILR (PP) and DD relaxed ILR (PP), we see that although IM CPP is less, both TL and LTL costs per pound increase. Moreover, we observe that the amount of product shipped via LTL in the relaxed ILR doubled compared to DD relaxed ILR. This indicates that fewer routes are eligible for shipment by truck or container because of

	Total shipment (lbs.)	Total cost (\$)	No. of shipments	CPP (\$)	% save
WW	11,059,357	750,000.00	327	0.06782	
ILR (TU:90/RU:75)	10,996,688	730,699.62	265	0.06645	2.02
TC: 43,500 lbs.					
CC: 41,000 lbs.					
DD: 0 day					
Relaxed ILR	10,908,633	703,344.06	270	0.06448	4.93
TC: 43,700 lbs.					
CC: 41,750 lbs.					
DD: 1 day					
DD relaxed ILR	10,951,853	716,064.31	269	0.06538	3.59
TC: 43,500 lbs.					
CC: 41,000 lbs.					
DD: 1 day					
TC-CC relaxed ILR	11,058,497	731,034.26	282	0.06611	2.52
TC: 43,700 lbs.					
CC: 41,750 lbs.					
DD: 0 day					
ILR (PP)	11,065,620	752,950.89	303	0.06804	−0.34
TC: 43,500 lbs.					
CC: 41,000 lbs.					
DD: 0 day					
Relaxed ILR (PP)	10,957,624	720,411.89	290	0.06575	3.05
TC: 43,700 lbs.					
CC: 41,750 lbs.					
DD: 1 day					
DD relaxed ILR (PP)	10,922,687	717,422.20	284	0.06568	3.15
TC: 43,500 lbs.					
CC: 41,000 lbs.					
DD: 1 day					
TC-CC relaxed ILR (PP)	11,058,497	737,724.24	299	0.06671	1.63
TC: 43,700 lbs.					
CC: 41,750 lbs.					
DD: 0 day					

Table 4: Cost performance of shipments from the Siloam Springs plant is better than from the Cullman plant; however, the total weight of the shipments from Siloam Springs is approximately half the weight of shipments from Cullman.

higher penalty costs. Overall, these results highlight the sophisticated inventory interactions and consolidation opportunities inherent in the problem.

Conclusions and Future Research Directions

In this research, we investigate WW's outbound supply chain. WW uses common carrier companies with TL, IM, and LTL options. Each transportation mode has different advantages and unique complicating constraints. Additionally, the dynamic nature of demand imposes planning challenges. We present a cluster-first, route-next type of heuristic to minimize the total outbound supply chain cost, while addressing all the system-wide restrictions. A generalized assignment problem creates

the clusters and loading plans, and a modified TSP solution finds optimal route and unloading sequences for customer orders. We propose and evaluate three dispatch policies: TU, RU, and PP. A combination of TU and RU provides the best results.

Incorporating our tool into WW's load-planning process has resulted in cost savings of 4.4 percent for the company over its previous process. Today, WW continues to use and benefit from this tool. In addition to the savings in outbound supply chain operations, this load-planning tool has also saved a significant amount of WW staff time.

WW is interested in further improvements to the work we discuss in this paper. One potential extension of this research is that instead of preassigning

	Total shipment (lbs.)	Total cost (\$)	No. of shipments	CPP (\$)	% save
WW	11,059,357	750,000.00	327	0.068	
TL	9,742,097	629,633.25	230	0.065	
IM	1,291,736	104,600.12	31	0.081	
LTL	25,524	15,766.63	66	0.618	
ILR (TU:90/RU:75)	10,996,688	730,699.62	265	0.066	2.02
TL	9,698,116	625,706.50	229	0.065	0.17
IM	1,177,306	99,176.22	30	0.084	−4.03
LTL	121,266	5,816.90	6	0.048	92.23
Relaxed ILR	10,908,633	703,344.06	270	0.064	4.93
TL	9,643,547	603,116.82	226	0.063	3.23
IM	1,107,027	89,006.67	27	0.080	0.71
LTL	158,059	11,220.58	17	0.071	88.51
DD ILR relaxed	10,951,853	716,064.31	269	0.065	3.59
TL	9,600,501	600,356.89	224	0.063	3.24
IM	1,244,443	105,843.04	31	0.085	−5.03
LTL	106,909	9,864.39	14	0.092	85.06
TC ILR relaxed	11,058,497	731,034.26	282	0.066	2.52
TL	9,637,194	608,148.16	226	0.063	2.36
IM	1,221,099	101,313.47	30	0.083	−2.46
LTL	200,204	21,572.63	26	0.108	82.56
ILR (PP)	11,065,620	752,950.89	303	0.068	−0.34
TL	9,409,037	601,910.69	243	0.064	1.02
IM	1,503,920	131,540.38	39	0.087	−8.01
LTL	152,664	19,499.82	21	0.128	79.32
Relaxed ILR (PP)	10,957,624	720,411.89	290	0.066	3.05
TL	9,408,471	586,600.67	239	0.062	3.53
IM	1,444,092	121,472.37	36	0.084	−3.88
LTL	105,061	12,338.85	15	0.117	80.99
DD relaxed ILR (PP)	10,922,687	717,422.20	284	0.066	3.15
TL	9,342,359	579,561.63	238	0.062	4.01
IM	1,519,648	131,320.03	39	0.086	−6.72
LTL	60,680	6,540.54	7	0.108	82.55
TC relaxed ILR (PP)	11,058,497	737,724.24	299	0.067	1.63
TL	9,391,750	587,608.93	241	0.063	3.19
IM	1,538,218	132,817.89	39	0.086	−6.63
LTL	128,530	17,297.43	19	0.135	78.21

Table 5: The table gives the transportation mode breakdown from the Siloam Springs plant.

customer orders to the Cullman or Siloam Springs plant, we develop an allocation module and let the optimization model assign customer orders to plants. Another extension would address late orders. An analytical approach for determining the degree of due-date relaxations would give WW a better tool to serve its customers for a lower outbound cost. Additionally, potential postprocessing may improve the handling of some late or critical orders, eliminating possible human intervention.

Although the problem and the solution approach we describe in this paper are tailored for WW, the model can be extended to address similar problems that shippers in other industries face. In particular, the RU and PP dispatch policies would be useful for other

shippers who experience similar dynamic demand challenges.

Acknowledgments

The authors thank Randall Robinson and Uday Rao for their detailed comments and three reviewers for their useful suggestions. This paper was sponsored by Webb Wheel, Inc.

Appendix

Sets and Indices

i : index of seed locations, $i \in \mathcal{I}$.

j : index of orders, $j \in \mathcal{J}$.

p : index of products, $p \in \mathcal{P}$.

Problem Parameters

D_{jp} : demand of product $p \in \mathcal{P}$ in order $j \in \mathcal{J}$.

a_{ij} : cost of adding order $j \in \mathcal{J}$ to seed location $i \in \mathcal{I}$.

c_i : cost of using seed location $i \in \mathcal{J}$.
 Q_i : capacity of seed location $i \in \mathcal{J}$.
 w_p : weight of product $p \in \mathcal{P}$.
 I_p : available inventory of product $p \in \mathcal{P}$.
 k_j : 1, if order $j \in \mathcal{J}$ is critical, 0, otherwise.
 HL : total available driving time for containers.
 DT : drop time of an order for container-seed location.
 v : average speed.
 d_{ij} : distance from container-seed location $i \in \mathcal{J}''$ to order location $j \in \mathcal{J}$.
 b_{ij} : 1 if the seed location $j \in \mathcal{J}$ is created for order $i \in \mathcal{J}'''$; 0 otherwise.

Decision Variables

Y_i : 1 if seed location $i \in \mathcal{J}$ is used; 0 otherwise.
 Z_{ij} : 1 if order $j \in \mathcal{J}$ assigned to seed location i ; 0 otherwise.
 X_{ijp} : amount of product $p \in \mathcal{P}$ shipped for order $j \in \mathcal{J}$ via seed location $i \in \mathcal{J}$.

Calculation Details

Truck-Seed Location Cost

$$c_i = \begin{cases} d_{0i}sr + mc, & \text{if } d_{0i}r_i \leq mc; \\ (sr + r_i)d_{0i}, & \text{otherwise,} \end{cases}$$

where d_{0i} is the distance from the plant to the seed location i , r_i is the rate per mile for the seed location i , sr is the surcharge rate, and mc is the minimum charge requested by the trucking company.

Assignment Cost to a Truck Seed

$$a_{ij} = \begin{cases} d_{0j} + (d_{0j} + d_{ji})sr + mc - c_i, & \text{if } (d_{0i} + d_{ji})r_i \leq mc; \\ d_{0j} + (d_{0j} + d_{ji})(sr + r_i) - c_i, & \text{otherwise,} \end{cases}$$

where d_{0j} is the drop charge, d_{0j} is the distance from the plant to the order location j , and d_{ji} is the distance from the order location j to the seed location i .

Weight-Based Truck Utilization (TU)

$$TU_i = \frac{\sum_{j \in \mathcal{J}} \sum_{p \in \mathcal{P}} w_p X_{ijp}}{Q_i} \times 100.$$

Route-Based Utilization (RU)

$$RU_i = \frac{\sum_{j \in \mathcal{J}} \sum_{p \in \mathcal{P}} d_{0j}(w_p X_{ijp})}{L_i Q_i} \times 100,$$

where d_{0j} is the distance from the plant to order location j , L_i is the length of route i , and Q_i is the capacity of the vehicle.

Penalty-Based Policy (PP)

$$PP = (1 - RU) \cdot \text{Route Cost}.$$

Overall Formulation of the Assignment Problem

$$\min \left\{ \sum_{i \in \mathcal{J}} c_i Y_i + \sum_{i \in \mathcal{J}} \sum_{j \in \mathcal{J}} a_{ij} Z_{ij} \right\} \quad (\text{AP})$$

subject to

$$\sum_{j \in \mathcal{J}} \sum_{p \in \mathcal{P}} w_p X_{ijp} \leq Q_i, \quad \forall i \in \mathcal{J}' \cup \mathcal{J}'', \quad (1)$$

$$\sum_{i \in \mathcal{J}} X_{ijp} \leq D_{jp}, \quad \forall j \in \mathcal{J} \text{ and } p \in \mathcal{P}, \quad (2)$$

$$\sum_{i \in \mathcal{J}} \sum_{j \in \mathcal{J}} X_{ijp} \leq I_p, \quad \forall p \in \mathcal{P}, \quad (3)$$

$$\sum_{i \in \mathcal{J}} \sum_{j \in \mathcal{J}} X_{ijp} k_j \geq \min \left(\sum_j D_{jp} k_j, I_p \right), \quad \forall p \in \mathcal{P}, \quad (4)$$

$$\sum_{i \in \mathcal{J}} \sum_{j \in \mathcal{J}} X_{ijp} \geq \min \left(\sum_j D_{jp}, I_p \right), \quad \forall p \in \mathcal{P}, \quad (5)$$

$$\sum_{j \in \mathcal{J}} \left(DT + 2 \frac{d_{ij}}{v} \right) Z_{ij} \leq HL, \quad \forall i \in \mathcal{J}'', \quad (6)$$

$$Z_{ij} \leq b_{ij}, \quad \forall i \in \mathcal{J}''' \text{ and } j \in \mathcal{J}, \quad (7)$$

$$\sum_{p \in \mathcal{P}} X_{ijp} \leq \sum_{p \in \mathcal{P}} D_{jp} Z_{ij}, \quad \forall i \in \mathcal{J} \text{ and } j \in \mathcal{J}, \quad (8)$$

$$Z_{ij} \leq Y_i, \quad \forall i \in \mathcal{J} \text{ and } j \in \mathcal{J}, \quad (9)$$

$$X_{ijp} \in \mathbb{Z}^+, Y_i \text{ and } Z_{ij} \in \{0, 1\}, \quad \forall i \in \mathcal{J}, j \in \mathcal{J} \text{ and } p \in \mathcal{P}. \quad (10)$$

In the objective function (AP), the first term is the total seed-location usage cost. The second term is the total assignment cost of orders to seed locations. Constraints (1) dictate that the capacity of the seed location is not violated. Depending on the type of seed (i.e., truck, container, or LTL), the capacity can take on different values. Constraints (2) and (3) ensure that the number of units shipped is less than the order quantity and the inventory, respectively, whichever is smaller. Constraints (4) and (5) ensure that the demand for critical and regular orders is fulfilled until either demand is satisfied or no inventory remains for all products $p \in \mathcal{P}$. Constraints (6) limit the travel-time service requirements from a container-seed location. Constraints (7) ensure that an order can only be assigned to a LTL seed location, if the LTL seed location is opened for that order. Constraints (8) and (9) establish the relationship among binary variables Z_{ij} and Y_i for all $i \in \mathcal{J}$ and $j \in \mathcal{J}$. Finally, constraints (10) specify the bounds for the decision variables.

Notation for the Routing Problem

N : the number of orders/drops in a route.
 e, f, g : set of locations in a route.
 t : route taken from e to f at step t , $t \in \{0, \dots, N\}$.
 r_f : rate per mile of order location f .
 R : rate per mile of the tour.
 d_{ef} : the distance from the order location e to the order location f .

d_{0f} : the distance from the plant to the order location f .
 d_{e0} : the distance from the order location e to the plant,
 where $d_{e0} = 0$.

Decision Variable

W_{eff} : 1 if we drive from the order location e to the order location f at step t .

Overall Formulation of the Routing Problem

The N-MTSP formulation is as follows:

$$\min \left\{ \sum_{e=0}^N \sum_{f=0}^N \sum_{t=0}^N d_{ef} W_{eff} R \right\} \quad (\text{N-MTSP})$$

subject to

$$\sum_{e=0}^N W_{e,f,t} - \sum_{g=0}^N W_{f,g,t+1} = 0, \quad \forall f \in \{0, \dots, N\}, t \in \{0, \dots, N-1\}, \quad (11)$$

$$\sum_{e=0}^N W_{e,f,N} - \sum_{g=0}^N W_{f,g,0} = 0, \quad \forall f \in \{0, \dots, N\}, \quad (12)$$

$$\sum_{f=0}^N \sum_{t=0}^N W_{e,f,t} = 1, \quad \forall e \in \{0, \dots, N\}, \quad (13)$$

$$\sum_{f=0}^N W_{0,f,0} = 1, \quad (14)$$

$$\sum_{f=0}^N W_{f,0,N} = 1, \quad (15)$$

$$R = \sum_{f=0}^N r_f W_{f,0,N}, \quad (16)$$

$$W_{eff} \in \{0, 1\}, \quad \forall e, f, t \in \{0, \dots, N\}. \quad (17)$$

The objective of N-MTSP is to minimize the total cost of the tour. Because d_{e0} equals zero for all $e \in 0, \dots, N$, the tour is also the route for our calculation. Constraints (11) and (12) are flow constraints that balance incoming and outgoing flows to each location. Constraints (13) ensure that each location is visited once. Constraints (14) and (15) dictate that the tour starts and ends at the plant. Constraint (16) sets the rate per mile of the tour to the rate per mile of the last order location.

References

- Ambrosino D, Sciomachen A (2007) A food distribution network problem: A case study. *IMA J. Management Math.* 18(1):33–53.
- Archetti C, Speranza MG (2012) The split delivery vehicle routing problem: A survey. *Internat. Trans. Oper. Res.* 19(1–2):3–22.
- Bertazzi L, Savelsbergh M, Speranza MG (2008) Inventory routing. Golden B, Raghavan S, Wasil E, eds. *The Vehicle Routing Problem: Latest Advances and New Challenges* (Springer, New York), 49–72.
- Bolduc M-C, Renaud J, Boctor F (2007) A heuristic for the routing and carrier selection problem. *Eur. J. Oper. Res.* 183(2):926–932.
- Bolduc M-C, Laporte G, Renaud J, Boctor F (2010) A tabu search heuristic for the split delivery vehicle routing problem with production and demand calendars. *Eur. J. Oper. Res.* 202(1):122–130.
- Bramel J, Simchi-Levi D (1995) A location based heuristic for general routing problems. *Oper. Res.* 43(4):649–660.
- Bräysy O, Gendreau M (2005a) Vehicle routing problem with time windows, Part I: Route construction and local search algorithms. *Transportation Sci.* 39(1):104–118.
- Bräysy O, Gendreau M (2005b) Vehicle routing problem with time windows, Part II: Metaheuristics. *Transportation Sci.* 39(1):119–139.
- Campbell AM, Clarke LW, Savelsbergh MWP (2002) Inventory routing in practice. Toth P, Vigo D, eds. *The Vehicle Routing Problem* (SIAM, Philadelphia), 309–326.
- Cattani KD, Souza GC (2002) Inventory rationing and shipment flexibility alternatives for direct market firms. *Production Oper. Management* 11(4):441–457.
- Chu CW (2005) A heuristic algorithm for the truckload and less-than-truckload problem. *Eur. J. Oper. Res.* 165(3):657–667.
- Côté JF, Potvin J-Y (2009) A tabu search heuristic for the vehicle routing problem with private fleet and common carrier. *Eur. J. Oper. Res.* 198(2):464–469.
- Dantzig GB (1963) *Linear Programming and Extensions* (Princeton University Press, Princeton, NJ).
- Deshpande V, Cohen MA, Donohue K (2003) A threshold inventory rationing policy for service-differentiated demand classes. *Management Sci.* 49(6):683–703.
- Dror M, Trudeau P (1989) Savings by split delivery routing. *Transportation Sci.* 23(2):141–145.
- Dror M, Trudeau P (1990) Split delivery routing. *Naval Res. Logist.* 37(3):383–402.
- Fisher ML, Jaikumar R (1981) A generalized assignment heuristic for vehicle routing. *Networks* 11(2):109–124.
- Flood MM (1956) The traveling-salesman problem. *Oper. Res.* 4(1):61–75.
- Gaur V, Fisher ML (2004) A periodic inventory routing problem at a supermarket chain. *Oper. Res.* 52(6):813–822.
- Gayon J-P, Benjaafar S, de Véricourt F (2009) Using imperfect advance demand information in production-inventory systems with multiple customer classes. *Manufacturing Service Oper. Management* 11(1):128–143.
- Golden BL, Raghavan S, Wasil EA (2008) *The Vehicle Routing Problem: Latest Advances and New Challenges* (Springer, New York).
- Hung Y-F, Hsiao J-Y (2013) Inventory rationing decision models during replenishment lead time. *Internat. J. Production Econom.* 144(1):290–300.
- Lei H, Laporte G, Guo B (2012) The vehicle routing problem with stochastic demands and split deliveries. *INFOR: Inform. Systems Oper. Res.* 50(2):59–71.
- Natarajathinam M, Stacey J, Sox C (2012) Near-optimal heuristics and managerial insights for the storage constrained, inbound inventory routing problem. *Internat. J. Physical Distribution Logist. Management* 42(2):158–173.
- Stacey J, Natarajathinam M, Sox C (2007) The storage constrained, inbound inventory routing problem. *Internat. J. Physical Distribution Logist. Management* 37(6):484–500.
- Toth P, Vigo D (2002) *The Vehicle Routing Problem* (SIAM, Philadelphia).
- Vidal T, Crainic TG, Gendreau M, Prins C (2013) Heuristics for multi-attribute vehicle routing problems: A survey and synthesis. *Eur. J. Oper. Res.* 231(1):1–21.

Burcu B. Keskin is an associate professor of operations management at The University of Alabama. She earned her PhD in industrial engineering from Texas A&M University in 2007. She has conducted research in different areas of supply chain management including inventory theory, facility location, network design, and vehicle routing. She has obtained funding for her research through the University Transportation Center for Alabama grants and several companies. Her papers have appeared in *Naval Research Logistics*, *IIE Transactions*, *Computers and Operations Research*, *European Journal of Operational Research*, *Journal of Operational Society*, and *Interfaces* among others. Dr. Keskin serves on the editorial boards of *Handbooks in OR/MS* and *International Journal of Operations Research and Information Systems*. She is the recipient of 2014 Culverhouse Board of Visitors Teaching Award. In addition to her scholarly activities, Dr. Keskin is an active member of INFORMS, DSI, and POMS. She was a recipient of the INFORMS Judith Liebman Moving Spirit Award in 2007 and INFORMS Fora Moving Spirit Award in 2013.

İbrahim Çapar is a PhD student in the operations management program at the University of Alabama. He earned an MBA degree from the University of Alabama in 2011 with a dual concentration in supply chain operations management and business analytics. His research interests include online optimization and its applications related to vehicle routing, patrol routing, emergency response, and logistics.

Charles R. Sox holds the University Chair of Manufacturing Management at The University of Alabama where he also

serves as the department head for information systems, statistics, and operations management in the Culverhouse College of Commerce and Business Administration. He completed his M.S. degree in 1991 and his PhD degree in 1992 at Cornell University in the School of Operations Research and Industrial Engineering. His current areas of research and teaching are in supply chain management, planning and scheduling, inventory control, and logistics. His papers have appeared in *Management Science*, *IIE Transactions*, *Transportation Research*, *European Journal of Operational Research*, *International Journal of Inventory Research*, *Operations Research Letters*, *International Journal of Physical Distribution and Logistics Management*, and the *International Journal of Production Economics*. His research has been supported by grants from the National Science Foundation and the Department of Transportation. He was selected as the Outstanding Alabama MBA Faculty Member in 2009 and 2010. He was selected as the outstanding Auburn ISE Professor in 1997 and 1999. He has worked as a consultant for several companies including Mercedes-Benz U.S. International, Southern Company, Hyundai Motor Manufacturing of Alabama, Engelhard Corporation, and Milliken.

Nickolas K. Freeman earned his PhD degree in operations management from the University of Alabama in 2014. He is an assistant professor of supply chain management at the University of Houston. Modeling and optimization in the areas of scheduling, supply chain risk management, and logistics are among his research interests. He has published in the *International Journal of Production Economics* and *IIE Transactions*.