



INFORMS Journal on Applied Analytics

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

Dispatch Optimization in Bulk Tanker Transport Operations

Ted Gifford, Tracy Opicka, Ashesh Sinha, Daniel Vanden Brink, Andy Gifford, Robert Randall

To cite this article:

Ted Gifford, Tracy Opicka, Ashesh Sinha, Daniel Vanden Brink, Andy Gifford, Robert Randall (2018) Dispatch Optimization in Bulk Tanker Transport Operations. INFORMS Journal on Applied Analytics 48(5):403-421. <https://doi.org/10.1287/inte.2018.0956>

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 © 2018, 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>

Dispatch Optimization in Bulk Tanker Transport Operations

Ted Gifford,^a Tracy Opicka,^a Ashesh Sinha,^a Daniel Vanden Brink,^a Andy Gifford,^a Robert Randall^b

^a Engineering and Research, Schneider National, Inc., Green Bay, Wisconsin 54306; ^b Princeton Consultants, Inc., Princeton, New Jersey 08540

Contact: giffordt@schneider.com,  <http://orcid.org/0000-0002-7861-354X> (TG); opickat@schneider.com (TO); sinhaa@ksu.edu (AS); daniel.vandenbrink@thyssenkrupp.com (DVB); gifforda@schneider.com (AG); rrandall@princeton.com (RR)

<https://doi.org/10.1287/inte.2018.0956>

Copyright: © 2018 INFORMS

Abstract. All modes of freight transportation are subject to flow imbalances that affect the efficiency of asset utilization. The use of mathematical programming optimization models has a rich history of application to this problem. We use variations and extensions of the classical assignment problem to find globally optimal or near-optimal solutions to the problem of assigning a large number of individual resources (transportation assets) to tasks (freight movements). We address a particularly difficult variant of this problem that occurs in the bulk transport (fuels and chemicals) division of Schneider National Inc. This group accepts 350 customer orders per day, involving 10,000 distinct commodities, with pickup and delivery locations across the continental United States. The corresponding freight movements are executed using a fleet comprising a 1,000 drivers and 1,600 tanker trailers. Chemical interaction properties of these commodities impose complex product-sequencing constraints, interorder tanker wash and preparation processes, and the selection of specific trailer configurations. Schneider National must consider these complexities in addition to those encountered in more common fleet dispatch problems. To address this problem, the engineering group at Schneider National designed and implemented a multiphase, multidimensional matching algorithm and developed new business processes that enable business planners to leverage optimized solution recommendations. We documented over \$4 million in annualized operational and capital cost savings, as well as significantly improved productivity and customer service, which this new system has been delivering since its implementation.

Keywords: vehicle routing • driver scheduling • multiphase optimization • column generation • mixed-integer programming • resource assignment • freight transportation networks

Introduction

The efficient, timely transport of commercial freight across public road networks depends on many complex operational decisions. A significant percentage of freight in North America is transported by large-scale for-hire full-truckload carriers. In operations of this type, carriers are tendered freight orders that originate and are delivered across a wide geographic range of customer locations. Because of service and cost considerations, full-truckload carriers operate random one-way networks. This contrasts with less-than-truckload (LTL) and small-package operators (e.g., UPS, FedEx) who operate structured networks consisting of predefined hubs and lanes.

In random one-way operations, resources (i.e., drivers with tractors and trailers) are typically assigned to specific freight orders with a lead time of several hours to several days prior to pickup activity. Within a specific lead-time horizon, several thousand drivers may become available (after completing earlier orders or returning from time off) and ready to be assigned to a similar number of customer orders that the carrier has accepted. Most large fleets maintain trailer pools that allow for loading and unloading to occur separately from

pickup and delivery activity, which enables more efficient utilization of assets. In this scenario, tractor and trailer assignments are considered separately and solution techniques based on a classical two-dimensional assignment problem are no longer applicable.

In this paper, we address a particularly challenging variant of this problem that arises in the transport of liquid-chemical products and fuels with bulk tanker trailers. The nature of this freight requires the consideration and modeling of two additional complicating factors: (1) prevention of hazardous interactions between different commodities and (2) the washing and cleaning of tanker compartments between orders.

Problem Description

Schneider National's Bulk Transport division operates a national fleet comprising roughly 1,000 tractors and 1,600 tanker trailers, across which it dispatches 350 new orders per day. Over the course of a year, 10,000 distinct commodities may be transported. Chemical-interaction properties of these commodities impose complex product-sequencing constraints, interorder tanker wash and preparation processes, and the selection of specific trailer configurations. In almost all cases, tanker trailers

must be routed to independent facilities where they are washed and prepared for their next use.

The dispatch problem that we address can be viewed as a multidimensional matching problem in which several types of constrained resources are matched to a collection of complex tasks. These tasks, customer requests to transport liquefied chemicals, are characterized by a set of attributes consisting of an origin-destination location pair, pickup and delivery time windows, a product-commodity specification, and possibly special equipment or handling instructions. The execution of each task requires several interdependent sourcing decisions that determine an appropriate tanker trailer and one or more drivers to complete execution of the order. The selection of the tanker trailer, in turn, requires several subdecisions related to tanker attributes, wash activity, and the product-sequence constraints we allude to above.

The tank-choice decision comprises several steps: (1) determining a collection of suitably cleaned, prepped, and configured tanker trailers that are compatible with order requirements, (2) further reducing this trailer set by checking whether previous contents of the tanker meet compatibility rules relative to the prospective contents, (3) selecting a tank wash location where such tankers will be available, and (4) selecting a second tank wash location at which the selected trailer will be washed and prepped for a subsequent order.

The bold lines in Figure 1 illustrate a single episode (i.e., coverage of a single order) in the life cycle of a tanker trailer. In this case, a customer order specifies loading and pickup at customer location P1 and delivery at D1. A suitable tanker is identified at wash location W1 and moved *clean* to P1 for loading. The *loaded* tanker is then transported to consignee location D1.

After unloading at D1, the *dirty* tanker is then repositioned to wash location W4, where it will be cleaned and prepped for its next order. The nonbold lines in Figure 1 depict an alternative (nonselected) assignment choice for the selected trailer to a different order (W1-P2-D2-W4) or use of a different tank for this order (W2-P1-D1-W4). The task of the tanker route optimization model (Phase 1) is to determine optimal (and near-optimal) solutions from the feasible combinations that are identified. Prewash location alternatives (e.g., W2-P1) are considered in the context of both costs (e.g., distance) and constraints (e.g., inventory balancing). Postwash alternatives (e.g., D1-W3) attempt to account for relative opportunity future value based on tanker type and attributes. The “Planned Improvements” section will describe plans to integrate this work with our order forecast system to improve the accuracy of these future-value estimates.

At most customer locations, loading or unloading may take several hours to a day or longer. Consequently, most orders require two or more distinct driver assignments so that drivers do not incur undue idle time. Driver-assignment options are evaluated for cost and feasibility against a subset of trailer routes that were determined to be either optimal or near optimal in the tanker route optimization phase. Generally, the direct cost difference between alternative drivers executing tanker moves is minimal. Consequently, the focus of the second phase optimization is on minimizing unproductive driver activity, while ensuring that driver work rules and other considerations are met.

The maps in Figure 2 indicate the number and geographic distribution of the locations that compose the network over which bulk dispatch optimization occurs. Maps (a) and (b) depict pickup and delivery

Figure 1. (Color online) Tankered Trailers Revisit Wash Facilities After Servicing Customer Orders





Figure 2. (Color online) The Geographic Distribution of Customer and Tank Wash Facility Locations



locations, respectively. Although many locations enter the problem in both roles (e.g., manufacturing facilities in which raw material commodities are inbound and finished product commodities are outbound), many sites are exclusively pickup or exclusively delivery points. Schneider National currently has approximately 9,000 distinct customer and consignee locations.

Map (c) indicates wash locations, which are operated by external vendors and shared across the bulk transport industry. These facilities vary in both handling-capacity size and wash capability (i.e., commodities

washed and prep procedures offered), and the dispatch optimization system must account for this. Of the roughly 150 active wash facilities, approximately half are used on a regular basis. Additional complexity is introduced by customer-specific requirements that often specify different procedures, although the commodities may be chemically equivalent.

Customer-specific requirements also drive prior-product (i.e., previous commodities that have been in the tank) compatibility rules that together with trailer physical attributes (e.g., lining material, heaters, pumps) determine the subset of tanker trailers that are feasible for a given order. One of the early challenges of this project was to gain control over this process and establish comprehensive, consistent, and reliable data stores to enable automated, high-speed checking and validation of prior-product conditions. Prior to implementation of new systems at Schneider National, customer requirements and compatibility rules were delivered using various formats (e.g., PDF, XLS, DOC) with no standardization and no process in place to convert and maintain this information in digital format. Some requirements apply across generic chemical-compound names; however, many specifically pertain to unique customer products. We developed a new data model and associated maintenance and rules-engine logic tools to facilitate automation. We will describe aspects of these in the subsequent sections.

Related Work

The dispatch decision process for a network of tanker trucks transporting chemicals or liquid bulk lies in the intersection of heterogeneous-fleet vehicle routing, driver scheduling, and integrated fleet-assignment problems. Gifford (2011) discusses the general problems of optimizing ongoing operations of for-hire fleets of tractors and trailers in a large-scale network.

One of the classes of rich vehicle routing problems deals with a system of heterogeneous fleets of vehicles. The heterogeneous vehicle routing problem, which was first introduced by Golden et al. (1984), addresses operations with fleets of vehicles that differ in terms of vehicle type, capacity, and costs. Later, Taillard (1999) introduced the heterogeneous fixed fleet vehicle routing problem (HFVRP) using a set of predefined vehicles. Unlike classical vehicle routing problems with multiple identical vehicles (Dantzig and Ramser 1959), to the best of our knowledge, the HFVRP has not yet been solved by an exact approach (Koç et al. 2016).

Studies on heterogeneous vehicle routing focus mainly on developing heuristic approaches and determining good lower and upper bounds on the optimal solution (Desrochers and Verhoog 1991, Yaman 2006, Li et al. 2007, Euchi and Chabchoub 2010). Choi and Tcha (2007) develop a combination of column-generation and dynamic programming-based schemes to generate tight

bounds on the optimal solution for a heterogeneous vehicle routing problem. They first formulate the problem as a set-covering model, then use dynamic programming to efficiently generate feasible columns and, finally, solve the linear programming (LP) relaxation. Subramanian et al. (2012) analyze a single-depot vehicle routing problem with heterogeneous fleets using a hybrid algorithm that utilizes local-search-based heuristics to generate routes and follow with a set-partitioning formulation. Later, Penna et al. (2013) improve the hybrid approach by integrating the local-search heuristic with a variable neighborhood descent procedure. Nazi-Azmi and Salari (2013) develop an integer programming-based heuristic for a heterogeneous vehicle routing problem that destroys and repairs the initial solution to solve the model to optimality.

Another class of vehicle routing problems addresses integrated fleet assignment and driver scheduling (Savelsbergh and Sol 1998, Sherali et al. 2013, Goel and Vidal 2014). From that perspective, Xu et al. (2003) used column generation and dynamic programming to analyze a multiperiod vehicle routing problem with multiple time windows for pickup and delivery, and Department of Transportation (DOT) driver hours-of-service rules. Goel and Irnich (2016) develop an exact branch-and-price-based algorithm to schedule drivers for a vehicle routing problem that also considers DOT rules. Cacchiani and Salazar-Gonzalez (2017) develop a multiphase approach that utilizes column-generation and dynamic programming procedures to solve an integrated fleet-assignment and crew-scheduling problem.

In contrast to these studies, our work on bulk dispatch involves four levels of decisions—order selection, tanker trailer, driver scheduling, and wash locations—in addition to practical constraints (e.g., chemical

interactions and wash and tanker-type requirements) that create additional complexity. Except for approaches to relatively small problems and those with special structures or constraints that limit the space of feasible solutions, approaches to complex problems must rely partly on heuristic strategies. In addition to elements represented by the problems described in related literature, the practical aspects of our problem tend to achieve short processing times, while discounting the value of the exact solutions. We believe our approach exhibits a theoretically appropriate and operationally pragmatic balance between heuristic procedures and explicit optimization steps. For small test problems, our approach becomes provably optimal because the generation allows us to uncover all reasonable candidate routes and schedules.


Business Processes

In this section, we describe the business processes associated with the life cycle of a customer order and explain the interactions between the “human in the loop” and the various software components that support dispatch decisions.

The process begins with order tender offers from a customer to Schneider National via electronic data interchange (EDI), telephone, or email attachments. In almost all cases, pricing and other customer specifics have been agreed to in advance and are automatically supplied to processes and models as needed. Figure 3 depicts the main decision steps and their typical timing relative to the beginning of actual execution.

Order information includes specific-product commodity, origin-pickup location, destination-delivery location, and associated time constraints. Commodity, customer and location details, tanker trailer characteristics, wash

Figure 3. (Color online) Sequence of Business Steps Associated with an Order



instructions, and required driver certifications are retrieved and made available for subsequent selection steps. Prospective prewash and postwash locations are designated using static tables that are periodically reviewed by business analysts. These preliminary selections serve as an aid to regional planners and are often overwritten by the optimization solver as current information regarding trailer inventories at wash facilities and driver availability by geography becomes available. As we note in Figure 3, these preliminary wash-location assignments are normally made four days before the pickup date.

Two days before the scheduled pickup, results from the trailer route optimization solver (TROS) provide preliminary trailer-assignment recommendations. These are provisional assignments that may be revised either in subsequent TROS runs or when the combined trailer driver-optimizer (TDOS) is run.

An integral component of selecting a feasible tanker trailer is ensuring that the previous contents of the container are chemically compatible with a prospective order. Restrictions are such that even with consideration of intervening washes, the commodities associated with the three prior loads may need to be considered. Products that

may appear to be chemically similar may also be subject to customer-specific rules. Prior to the implementation of the new system, this process was manual, inefficient, and subject to errors and rework. With a conservative perspective, trailers were often unnecessarily held idle to compensate for inaccurate or incomplete information. In the “Business Benefits and Challenges” section, we describe savings associated with reducing the number of customer rejections of incompatible trailers.

As a prerequisite to developing the optimization models, we had to develop an integrated tool to manage and enforce prior-product constraints. The supporting data model can catalog restrictions by customer, location, and product-commodity specifics.

Figure 4 gives a view of the user interface for this tool. The tool supports compatibility and incompatibility specifications at various levels of aggregation. A hierarchy is maintained in order to minimize the number of rules that must be examined. The strictest level supports rules that pertain to customer-specific proprietary compounds. Lower levels allow the grouping of products with a similar chemical composition (e.g., dioxides) or for industrial purposes (e.g., dyes and inks). The lowest

Figure 4. (Color online) Prior-Product Relationships Are Maintained for All Products

Prior Product Requirement Database

Shipper & Ship Unit Info

Liabe Party: CHEMICAL CO - Location: CHEMICAL CHANNA... Manufacture Name: ISOPENTANOIC ACID, MIXED IS... # of Prior Check: 1

Add Delete Copy profile from Search Profile Save

Prohibited Prior Products Approved Prior Products Contact Info

By Manufacture Name

- #1 CASTOR OIL
- 100 OIL
- 110RLV (CHEVRON NEUTRAL OIL)
- 2-ETHYLHEXYL THIOGLYCOLATE
- 240 NEUTRAL BASE OILS
- 40-100 AND LIGHT PALE OIL
- 4700 OIL
- 500 OIL
- 680 BLOWN SOYBEAN OIL
- 680 BLOWN SOYBEAN OIL (TK T SBO BLOWN OIL)

578 records Edit

By Product Family

- BIOCIDES
- CHLORINATED HYDROCARBON
- DYES OR INKS
- FATTY ACID
- FERTILIZER
- HERBICIDE
- LATEX_SYNTHETIC RESIN
- LINSEED OILS
- LUBRICANTS_NON OIL
- OIL ADDITIVES

24 records Edit

By Product Attribute

- WATER_SOLUBILITY_NO

1 records Edit

level can refer to general properties (e.g., water solubility, odor, or pH). Rules are structured to flexibly allow for the combining of inclusion and exclusion logic.

As we mention in the previous section, most orders require multiple drivers. The driver changes are referred to as relays and cause an order to be split into *relay legs* or *shipments*. Many relays are planned in advance to improve driver productivity or accommodate circumstances such as driver-certification requirements or international-border crossings. A number of unplanned relays are, however, predicated by conditions not known at the time of the original assignments.

Examples of such conditions include breakdowns, traffic delays, and driver work-schedule issues. Approximately 60% of orders are relayed at least once and over 50% are relayed two or more times. In the “Planned Improvements” section, we discuss ongoing work to provide additional automation and optimization capabilities to the relay determination process.

Finally, drivers are assigned to individual shipment legs. Assignment recommendations are provided by the combined TDOS, which often changes trailer solutions when a lower overall cost can be achieved. The full optimization cycle runs approximately every 10 minutes to be responsive to ongoing changes in order, driver, or trailer availability information. It is important to note that the dispatch optimization system is a decision support tool, not a decision-making tool. Despite best efforts, it is impossible for such a system to accurately capture and adjust to all relevant real-world information in such a way that all decision recommendations remain durable through to execution. The stated objective of the system is


to produce recommendations that are accepted with little or no human review at an 80%–90% rate. As such, we devoted considerable effort to providing an integrated user interface or cockpit control panel that allows the planner and dispatcher to interact with the optimization functions and perform various queries that facilitate intelligent overrides and responses to conditions outside the scope of the optimizer. Figure 5 provides a screenshot of this user interface.

The main screen of this application is organized into several key sections: list of shipments, trailer list, driver list, and assignment evaluation information. Planners apply various filters to focus on their areas of responsibility. Optimization model output includes reduced-cost estimates (LP relaxation dual values) to provide planners with information to evaluate alternatives if they believe that the optimal solution recommendations are not practicable. As we will describe in the “Route and Assignment Optimization” section, we implemented several trade-offs and related judgment parameters as soft costs or constraint thresholds. We also provided additional user-interface elements to allow network managers to adjust these parameters and other business-rule conditions. Examples include costs incurred for nonrevenue miles, unproductive driver hours, maximum allowable deadhead distance, and penalties to be assessed if orders are uncovered (i.e., not delivered) or wash-facility capacity limits are violated.

Route and Assignment Optimization

In this section, we describe key components of the software and mathematical models that support the

Figure 5. (Color online) Dispatchers Use a Multipanel Screen to Manage Dispatch Activity




dispatch decision process. As we described in the “Business Processes” section, this process comprises the following sequence of assignment decisions: (1) a specific tanker trailer and the prewash facility from which it will be sourced, (2) a potentially different driver and (or) trailer to execute each move, and (3) the postwash facility to which the trailer will be repositioned after delivery of the product to the customer. Figure 6 gives an overall view of the sequence of resource assignments required to execute an order. In this diagram, the path of arcs represents covering order O3 with a trailer at wash facility W1 and assigning drivers D2, D3, and D1, respectively, to the prewash, loaded, and postwash shipment legs, and finally positioning the trailer at a postwash site. The groups of connected dots within the driver nodes (i.e., D1, D2, and D3), each of which represents an individual driver, depict different driver work-assignment sequences, only one of which can be active for a given driver. The other arc paths suggest how the remaining portions of *activated* driver work-assignment sequences might participate in the execution of other orders.

During a single solution cycle, this assignment process will typically attempt to cover several thousand individual orders extending over a multiple-day horizon. Approximately 1,000 potentially available tanker trailers will be spread across 100 prospective wash locations.

The number of drivers whose next available time falls within the horizon is usually around 500. Consequently, the number of feasible order-execution scenarios is in the billions.

To address this size complexity, we designed a multi-phase process that breaks the problem into a sequence of interleaved candidate-generation and optimization steps


Figure 6. (Color online) Assignment Decisions Required to Execute an Order



(Figure 7). The overall process is heuristic rather than true optimization because the candidate-generation steps implement search trees that prune (locally) unpromising candidates, but do not guarantee that such a candidate would not be in an optimal solution. In Appendix A.4, we provide empirical analysis that strongly supports our confidence that solutions obtained are practically optimal.

Given the dynamic nature of the transport operations, the recommendations provided by the optimizer must accurately reflect changes to the underlying data in near real time. As such, the solver processes run continuously, completing each cycle in under 10 minutes. Only a modest percentage of the recommendations associated with a given solver run (i.e., either trailer routes or work-assignment sequences) actually may be dispatched before the next cycle completes. However, most recommendations

Figure 7. (Color online) The Solution Technique Comprises Interleaved Generation and Optimization Processes



remain durable over a long sequence of consecutive solver runs. Between successive solution cycles, 10% of order information is typically changed because of new additions, cancellations, or other information updates. Similarly, approximately 10% of tanker or driver availability information changes. These changes typically affect about 20% of trailer route recommendations and about 10% of driver work-assignment sequences.

Tanker Trailer Feasibility

The first step of the solution process is to determine the set of feasible, available tanker trailers, which depends on the chemical composition of products, customer requirements, and intervening wash cycles. This is the preprocessing step for trailer route generation, which we will describe in the “Trailer Route Generation” subsection. As mentioned in the previous section, this process is enabled by a business rules engine and corresponding database that capture prior-product compatibility and incompatibility relationships and support rapid identification of allowable tankers for prospective order of tanker new-order feasibility. Figure 8 depicts that logic flow to generate a set of feasible tanker trailers.

This process comprises a series of filtering checks to narrow the set of tank trailers to a subset that meets the required conditions. The sequence of these steps has been optimized to minimize the average evaluation time. In Phase 1, we check for manufacturer-specific constraints on a tanker profile and follow by matching product specifications and restrictions against the previous contents of the remaining trailers. To facilitate the rapid elimination of unacceptable trailers, this matching algorithm leverages both inclusion and exclusion lists and hierarchically descends from broad product families (similar chemical properties) to attribute classes and then to rules defined for specific compounds. The initial project plan assumed use of commercial rules-engine software; however, it resulted in unacceptable response times. With a redesign using a purpose-specific data structure and rules processing, complete prior-product feasibility filtering for a single order completes in about 0.05 ms. This represents an improvement of two orders of magnitude over the commercially available alternative.

Trailer Route Generation

After a set of product-compatible tanker trailers has been identified, we entered Phase 1 of the optimization process to generate and cost out possible trailer assignments for each order (Box 1A in Figure 7). Because the number of possible distinct tanker trailer route combinations required to cover even a modest number of orders is extremely large, the process makes extensive use of parallel processing to generate and determine costs for hundreds of thousands of trailer-assignment


options. To limit the number of routes, we consider prewash location options, but postwash choices are deferred unless they are internal to a multiple-order assignment. Each trailer route consists of the sequence of moves required to execute one or more orders and associated wash activities (excluding the final postwash). Multiple-order assignments are generated only when their total elapsed time falls within the planning horizon. In addition to standard costs reflecting trailer repositioning and loaded moves, additional cost factors contribute to the total route cost. These include penalties for late arrival (also subject to hard-constraint bounds), tank wash and prep costs (which may vary according to prior products), and bonus or penalty values reflecting the associated inventory impact (of trailer choice) at specific wash locations.

The penalty (and bonus) values used in costing phases fall into two categories: those that can be directly inferred from financial impacts; (e.g., empty miles, driver wait time, and wash cost differences) and those that are intended to influence solution characteristics, which are more difficult to quantify (e.g., late arrivals and inventory drawdowns). These values are based on a combination of business judgment and interactive tuning, wherein representative model problems are rerun with different values until a related outcome metric is achieved. A good example of this is the relationship between the late-arrival penalty and the achievement of an on time service level that meets customer expectations. Another example is the value associated with selecting the wash location to which a tanker will be repositioned after unloading. In the “Planned Improvements” section, we will discuss a new automated forecast-aware approach that will be used to determine this factor.

Trailer Route Optimization

These routes are then provided as columns to an LP relaxation of the initial trailer route set-covering optimization model (i.e., TROS) to find a reasonably sized set with promising routes to which final postwash options will be added for further evaluation; Box 1B in Figure 7 and Appendix A.1 show a detailed model. The optimization model minimizes the total cost of the route while ensuring the following constraints: (1) each order and trailer can be assigned to at most one route, (2) capacity at the tank washes should not be exceeded on a given day or week, and (3) number of trailers of a given type meet the minimum and maximum requirements for each region and area. To limit the number of trailer routes provided to the next step, we generate postwash options for the best several hundred thousand feasible options based on reduced costs from this LP relaxation (Boxes 1C and 1D in Figure 7). These final trailer routes are then entered into the TROS integer program (Box 1E in Figure 7). Although a set of optimal trailer assignments represents a lowest-cost deployment of tankers from

Figure 8. (Color online) The Logic Used to Validate Prior-Product Compatibility



wash location to order completions, the locations, capabilities, and available times of drivers may suggest that nonoptimal tanker assignments (in this narrow sense) may lead to better overall solutions

when drivers are considered. For this reason, many nonoptimal assignments (columns) with relatively low reduced cost (from the LP relaxation) are identified and included in the final phase.

Driver Work-Assignment Schedule Generation

Phase 2 in the optimization process includes a column-generation process that identifies feasible drivers and corresponding costs to execute the candidate trailer-order combinations (Box 2A in Figure 7). This is a many-to-one scheme in which several drivers with different work profiles (i.e., local, regional, nationwide) may complete different legs of the order (e.g., prewash to pickup, pickup to delivery, delivery to postwash). The process begins by decomposing orders, covered by the trailer routes, to collections of work assignments. Next, the set of available drivers is evaluated against each work assignment to find feasible driver work-assignment combinations, called driver work-assignment schedules. Several feasibility considerations, including hazmat requirements, power-unit accessory needs, and customer-specific certifications, may prevent a driver from being matched to a given work assignment. As feasible work-assignment schedules are generated, corresponding costs are accumulated. These costs consider empty miles, driver idle time, and penalties for undesirable (but sometimes unavoidable) consequences, such as late arrivals or driver-type mismatches (e.g., long-haul driver on regional order). Estimates of unused hours and arrival times are determined by simulating the driver execution of a work assignment, considering location open hours and including estimated breaks necessary for hours-of-service compliance. As subsequent assignments are added, the number of work-assignment sequences could grow into the millions. To limit the number considered to several hundred thousand, we empirically determined threshold ceilings that we use to discard options unlikely to be of further interest. These are discussed in Appendix A.4.

Combined Trailer Driver Optimization

Finally, a candidate set of driver work assignments and trailer routes are entered into the combined TDOS, which recommends a feasible combination of driver and trailer assignments to minimize overall operating costs subject to a variety of structural and business generated constraints; Box 2B in Figure 7 and Appendix A.2 provide the detailed model. The model retains all costs from TROS (trailer only) in addition to other driver-related costs, such as unused hours cost, and driver-based bonuses and penalties. The model requires that each driver be associated with at most one work-assignment sequence and links constraints between driver work assignments and the relevant trailer. Although this model provides a recommendation for the postwash location, planners sometimes need to consider other options. To facilitate this, additional postwash options are added to the model (Box 2C in Figure 7) and an LP version (Box 2D in Figure 7) of this expanded problem is solved to provide reduced costs associated with alternative wash locations. Planners then

use these values as partial guidance when choosing among alternatives.

Business Benefits and Challenges

This new system was developed and implemented over a one-year period from May 2016 to May 2017. The data depicted in Figure 9 outline the milestone dates for the major components. In Appendix A.3, we give implementation details.

The principal users of the system are area planning managers (APMs) who have the responsibility to plan and assign the resources required to execute customer orders. We will describe their roles in the “Planned Improvements” discussion.

As with almost all large-scale software development projects, we had numerous challenges to address and overcome. The primary technical challenge was to balance the need for fast response time and solution quality, which would stand up to user scrutiny, with the dynamic nature of the data inputs and the large problem size resulting from the combinatorial explosion of feasible options as wash location, trailer choice, and driver selections are considered. The business challenges were driven by the usual scope, schedule, and cost considerations. User acceptance was identified as a concern early on, due to the perception by APMs that their work was heavily dependent on extensive “tribal” knowledge, intuition, and human judgement. One of the outstanding successes of the project has been to move these users to an appreciation of the strengths of computer-based decision support and the realization that their work has become more enjoyable as they were able to focus more effort areas that require expert human judgement.

The new system substantially increases automation and uses mathematical optimization techniques to select optimal assignment choices. In the first few months of operation, this system generated significant gains in productivity for assets, drivers, office staff, and management personnel, as well as significant direct cost savings in fuel and other direct expenses.

Schneider’s engineering and business teams jointly performed an initial project assessment prior to

Figure 9. Start and End Dates of the Project Phases

Project Phase	Begin	End
Prior Products & Eq. Optimization	9-May-16	12-Aug-16
Equipment Optimization	8-Apr-16	12-Aug-16
Tank Wash Optimization	6-Jun-16	7-Oct-16
Driver Optimization	31-Oct-16	16-Dec-16
Floor Wide Rollout	17-Mar-17	17-Feb-17
APM Role Consolidation	17-Mar-17	7-Apr-17
Program Completion		6-May-17
Warranty Period	9-May-17	20-May-17
Maintenance Status	21-May-17	on going

authorizing the financial investment in this project. This analysis determined three main benefit areas. These included (1) \$2 million in cost avoidance, due primarily to a reduction in the number of nonrevenue miles driven, (2) \$2.1 million in additional revenue opportunity, driven primarily by increased driver and tractor productivity, and (3) a 28% improvement in the productivity of office staff engaged in planning and dispatch functions. In addition, several other benefit areas that are either difficult to quantify or not easily interpretable in monetary terms were identified. The benefits described below have been reaffirmed and are now based on performance for the three-month period from June to August 2017 and are net of year-over-year (YOY) differences that are unrelated to the project (e.g., changes in business climate). In the subsections below, we will describe these benefits in more detail.

Our development cycle leveraged an agile methodology in which we deployed a succession of minimum viable products to the business (in production) at one-to-two-month intervals with incremental releases every two weeks. This approach allowed us to benchmark benefit metrics before and during the project and facilitated course corrections that kept the realization of benefits in focus. This agile approach was key to the team's ability to exceed the ROI targets set at project inception.

Nonrevenue Miles Reduction

Almost all bulk orders incur nonbillable miles both prior to and following the tank wash, pretank wash and posttank wash, respectively. The pretank wash miles occur prior to pick up as a clean and prepped trailer is brought from a tank wash to the shipper. The posttank wash miles occur after delivery when the empty trailer is repositioned to another tank wash to be washed and prepped for its next order. As the new system was implemented, we saw an average 16-mile reduction prior to the tank wash side and an 8-mile reduction following the tank wash, a reduction of over 1 million nonbillable miles annually. In our south region, we saw a reduction of 86 miles per order YOY. This will provide an annualized reduction of over 1 million empty miles. We have also seen a significant reduction in the volatility of nonbillable miles, from a range of 290–460 in 2016 to 260–310 in 2017. This decrease is having a pronounced positive impact on planning functions.

Driver Productivity Improvement

In addition to selecting pre- and posttank wash trips, the dispatch system optimizes the selection and scheduling of drivers to service both these and loaded moves (i.e., commodity carrying from customer to consignee). This optimization has led to a significant decrease in the number of miles that need to be driven

to position drivers for order-based activity. Although some of this positioning supports tank wash moves, which are needed anyway, a significant portion consists of bobtail trips (i.e., trips without a trailer). Our analyses indicate that each 2.5-mile reduction in nonbillable driver miles enables one additional revenue-producing mile over the same asset base. We are now forecasting \$2.5 million in additional annual revenue.

Unused Hours

For the three-month period ending August 31, 2017, we realized a reduction of unused driver time of 2.5 hours per shipment without an appreciable change in shipment characteristics. At current average monthly volumes, this represents roughly 202,500 additional available hours per year over a constant driver population. Using a conservative estimate of the percentage of these hours that can be applied to additional freight and the resultant margin, this represents \$1.8 million in additional annual earnings before interest and taxes (EBIT).

Planner Productivity Improvement

Prior to full implementation of the new system, APM roles were separated; four planners covered the selection of tanker trailer and wash locations, nine were responsible for driver-tractor assignments, and three were qualified to do either task as needed. We now have 11 planners in the combined role. In addition, two full-time planners provide flex support. Consequently, in a YOY comparison, we are now able to plan an additional 900 orders with three fewer associates. This represents a 28% productivity improvement, considering both reduced headcount and increased throughput. We expect additional productivity gains when associates have gained expertise with the new processes and additional enhancements to the optimization models are made. While this represents a modest cost benefit, the greater advantage is in allowing associates more time to focus on problems and complex issues that need attention and are beyond the scope that automated systems can address. As we stated above, our goal is to achieve 80%–90% compliance with model recommendations. Currently, these compliance values are 60% for driver selection and 75% for trailers. However, as model tuning and other enhancements are implemented, we fully expect to reach our target.

Better Equipment Optimization

The system has enabled a structure change on the network planning team. Managing both driver and trailer assignment has been combined into a single planner role; this was not possible prior to implementing the new system. The APM now has complete visibility to all dispatch activity in a market.

Consequently, we have been able to reduce idle days by one day per trailer-month. This equates to approximately 12,000 additional days of trailer capacity availability per year, effectively adding 60 trailers to the fleet at no additional cost.

Prior-Product Validation

Prior to implementation of the subsystem to manage and enforce prior-product compatibilities, feasible trailer identification and verification would often take 30 minutes. The process now completes in 3–5 minutes and often in much less time. This translates to several thousand person-hours per year.

Improved Customer Experience

Because of errors in prior-product validation, drivers would sometimes arrive at a customer location with an infeasible trailer. A resultant rejection by the customer is expensive with respect to unbillable miles and erosion of customer good will. Trailer rejections for incompatible prior products have been reduced substantially and customer feedback has been strongly positive; as we provide better information to customers, we are helping them improve their own processes. Additionally, when trailers are rejected erroneously, the new system supports quick recognition of the customer error and enables timely billing for costs incurred. Conversely, when customers correctly reject trailers, we are now able to update rules and data to prevent repeat occurrences. We are still working on efforts to record and quantify this impact.

Order-Acceptance Response Time

The new system has enabled quicker and more accurate visibility to available and projected capability. This has led to a reduction in the time required to make an acceptance decision. Assessing the financial impact of this benefit is a task still to be completed.

We note that these are conservative estimates based on actual performance during the first several months of operation. As the system matures and learnings are incorporated into improvements and enhancements, we are confident that these benefit measures will increase.

Planned Improvements

During the development of this system, several important features were deferred in accordance with our Agile development methodology and the time-to-market benefits of a minimum viable product. With the system described in this paper fully operational, we are completing the design and beginning development of additional capabilities. We mention these briefly in this section.

Relays

A considerable number of orders are subject to a relay between drivers, which may occur at wash facilities, customer sites, or some specified location along a longer load route. Safety and security requirements dictate that these relay locations may have certain physical characteristics that will necessarily limit their number. We will be implementing a network-design solution to determine a limited and fixed set of locations that will best service network freight flows, as well as a real-time rerouting tool that will consider out-of-route (extra) miles, commodity and driver limitations, and relay location characteristics to provide updated solutions to meet driver needs and to help alleviate driver capacity imbalances across space and time.

Local Driver Optimization

Bulk tractor drivers are subdivided into three general work configurations: (1) over-the-road (OTR) drivers are assigned primarily to long-distance legs of several hundred miles or more; (2) regional drivers are assigned work that limit the number of consecutive nights away from home; and (3) local drivers are assigned work that allow them to return home every night. Currently, driver-assignment optimization uses various bonus and penalty adjustments to direct drivers to appropriate shipments. We believe that we can achieve significant additional cost reductions and productivity improvements, particularly for local drivers, by introducing additional algorithmic enhancements. We recently began work in this area.

Integration with Demand Forecasts

When considering wash location options, particularly for postwash, an understanding of near-term geographic differences in demand for various tanker configurations can lead to positioning choices that reduce the nonrevenue miles associated with subsequent prewash to customer moves. The current system makes limited use of historical patterns to inform these choices. Planned improvements will more actively integrate location-selection evaluations with demand forecasting and order-acceptance processes.

Dynamic Wash Rebates

The costs associated with tank washes, which are performed by third-party vendors, can vary significantly and are sometimes subject to volume discounts. The current costing tools are driven by transactions and are unable to model the step-function nature of volume discounts. Standard volume-discount techniques have limited value because distance-to-wash-facility trade-offs must be considered. We are investigating a new scenario analysis approach that will

better incorporate demand forecasts to project when volume discounts will outweigh additional-distance considerations.

Conclusion

Although Schneider National has been developing and implementing optimization-based decision support for over 25 years, this project has offered important insights and new ideas into the process of developing data-driven decision support models to govern and respond to real-time changes in a dynamic and complex operational environment. In particular, we have gained an enhanced appreciation of the importance of building feedback mechanisms directly into the process rather than as an afterthought. Providing timely and actionable feedback to frontline planners and dispatchers and their direct managers has been invaluable in driving the adoption of a new system that was met with considerable initial skepticism by business experts who were insistent that their work was too complex to allow even modest automation. The true power of the system is that it has freed these business experts from tasks that are well-suited for mathematical sophistication and allowed them to focus on the 10% of dispatch planning work that requires human judgment and resolution of issues that are affected by factors that the automated processes cannot address.

Feedback mechanisms also provide a systematic way for the engineer and developer to monitor and improve model performance. The close connection of this optimization system to real-world conditions and exigencies requires the model to balance sometimes-competing objectives. Because these are generally handled with soft constraints and associated bonus or penalty factors, we are able to efficiently link the tuning of model parameters to the solution characteristics desired by the business.

From a software development perspective, this project has provided an opportunity to develop a methodology, platform, templates, and reusable components to facilitate development and delivery of new dispatch optimization systems for our intermodal dray and standard-trailer truckload businesses.

From a technical and algorithmic perspective, our decomposition approach, which reduces problem complexity by initially separately considering the decision components (i.e., wash locations, tanker trailers, and drivers) and then combining them in a multiphase optimization framework, represents an innovative way to address and successfully solve a practical problem that is both structurally complex and of significant size. We have been able to achieve this with processing times that meet the requirements of a real-time dispatch system and still provide solution quality that

exceeds business expectations and is theoretically near optimum.

Appendix

A.1. Trailer-Only Optimization Model

In this section, we describe the “trailer-only” optimization model that determines an optimal set of trailer routes to cover a set of orders and their associated wash operations.

Let $\mathcal{E}$ be set of equipment (i.e., tanker trailers), $\mathcal{N}$ be set of network nodes that include wash locations, $\mathcal{W} \subseteq \mathcal{N}$, order pickup locations, $\mathcal{P} \subseteq \mathcal{N}$, and final delivery locations, $\mathcal{F} \subseteq \mathcal{N}$, for customer orders.

We let $t, t \in \mathcal{T}$, index the sequence of periods (days) under consideration; $e, e \in \mathcal{E}$, denote the individual equipment items, and $w, w \in \mathcal{W}$, the wash facility locations. For each customer order $o, o \in \mathcal{O}$, we are provided the pickup and delivery locations (p_o, f_o) , $p_o \in \mathcal{P}$, $f_o \in \mathcal{F}$, and will generate a number of route options for equipment e as “walks,” represented by $r_e = (w_1, p_o, f_o, w_2)$, $w_1, w_2 \in \mathcal{W}$, $r_e \in \mathcal{R}$, where w_1 is the prewash location and w_2 is the postwash location. Each equipment route r_e incurs a cost C_{r_e} comprising empty-mile costs, late-arrival costs, and appropriate bonuses and penalties. We also classify the

Table A.1. Sets of Equipment Routes

Tours	Description
$\mathcal{R}^e$	Set of routes for equipment $e \in \mathcal{E}$, $\mathcal{R}^e \subseteq \mathcal{R}$
$\mathcal{R}^o$	Set of routes that cover order $o \in \mathcal{O}$, $\mathcal{R}^o \subseteq \mathcal{R}$
$\mathcal{R}^w$	Set of routes that use wash $w \in \mathcal{W}$, $\mathcal{R}^w \subseteq \mathcal{R}$
$\mathcal{R}^g$	Union of routes for equipment contained in group $g \in \mathcal{G}$, $\mathcal{R}^g \subseteq \mathcal{R}$

Table A.2. Three Additional Sets for Driver Optimization

Sets	Description
Ω^d	Collection of candidate work-assignment sequences for driver $d \in \mathcal{D}$, $\Omega^d \subseteq \Omega$
Ω^s	Collection of work-assignment sequences that contain work assignment $s \in \mathcal{S}$, $\Omega^s \subseteq \Omega$
Ω^i	Collection of work-assignment sequences that intersect equipment route $i \in \mathcal{R}$

Figure A.1. Typical Problem Sizes and Run Times for the Generation Processes

Process	Method	# Generated	Time
Generate Trailer Routes w/o Post Wash	Tree Search with Pruning	20,000	< 5 secs
Trailer Routes w/ Post Wash	Tree Search Extension	500,000	< 5 secs
Select Routes for Driver Phase	Heuristic Pruning	20,000	< 3 secs
Generate Driver Work Sequences	Tree search with Simulation	1,000,000	< 1 min
Generate additional Wash Options	Tree Search Extension	500,000	< 2 secs

Figure A.2. Typical Problem Sizes and Run Times for the Optimization Steps

Model	Method	# Variables	# Constraints	Time
Preliminary Trailer Optimization	Linear Program	70,000	3,000	< 10 secs
Final Trailer Only Optimization	Mixed Integer	225,000	3,000	< 1 min
Main Driver-Trailer Optimization	Mixed Integer	300,000	9,000	< 2 mins
Final Optimization (adjust Post Wash)	Linear Program	50,000	3,000	< 10 secs

Table A.3. Typical Problem Sizes and Run Times for the Generation Processes


Process	Method	No. generated	Time
Generate trailer routes without postwash	Tree search with pruning	20,000	<5 s
Trailer routes with postwash	Tree search extension	500,000	<5 s
Select routes for driver phase	Heuristic pruning	20,000	<3 s
Generate driver work sequences	Tree search with simulation	1,000,000	<1 min
Generate additional wash options	Tree search extension	500,000	<2 s

Table A.4. Typical Problem Sizes and Run Times for the Optimization Steps

Model	Method	No. of variables	No. constraints	Time
Preliminary trailer optimization	Linear program	70,000	3,000	<10 s
Final trailer-only optimization	Mixed integer	225,000	3,000	<1 min
Main driver-trailer optimization	Mixed integer	500,000	9,000	<2 min
Final optimization (adjust postwash)	Linear program	50,000	3,000	<10 s

equipment into nonoverlapping groups with similar physical characteristics, such as the interior lining material, the number of compartments, and the location of valves. Let $g, g \in \mathcal{G}$ index these groups. We define additional sets of routes in Table A.1.

Let $x_i, x_i \in X$ be a binary variable that equals 1 if equipment route i is selected, and 0 otherwise. For a given set of feasible routes $\mathcal{R}_{\text{feasible}}$, we develop a set-partitioning formulation with side constraints, called *TO*, as follows:

Figure A.3. Cost Values for the Route Options Generated for Several Different Trailers

$$\begin{aligned}
 TO : \min \quad z_1 = & \sum_{i \in \mathcal{R}} C_i x_i + \sum_{o \in \mathcal{O}} B_o \xi_o + \sum_{\substack{w \in \mathcal{W} \\ t \in \mathcal{T}}} B_{w,t} \xi_{w,t} \\
 & + \sum_{g \in \mathcal{G}, w \in \mathcal{W}} (B_{g,w}^L \xi_{g,w}^L + B_{g,w}^U \xi_{g,w}^U). \quad (A.1)
 \end{aligned}$$

Subject to

$$\sum_{r_e \in \mathcal{R}^e} x_{r_e} + \xi_e = 1, \forall e \in \mathcal{E} \quad (A.2)$$

$$\sum_{r_o \in \mathcal{R}^o} x_{r_o} + \xi_o = 1, \forall o \in \mathcal{O} \quad (A.3)$$

$$\sum_{r_w \in \mathcal{R}^w} x_{r_w} - \xi_{w,t} \leq K_{w,t} \quad \forall w \in \mathcal{W}, \forall t \in \mathcal{T} \quad (A.4)$$

$$\sum_{r_{g,w} \in \mathcal{R}^g \cap \mathcal{R}^w} x_{r_{g,w}} + \xi_{g,w}^L \geq \theta_{g,w}^L, \forall g \in \mathcal{G}, \forall w \in \mathcal{W} \quad (A.5)$$

$$\sum_{r_{g,w} \in \mathcal{R}^g \cap \mathcal{R}^w} x_{r_{g,w}} + \xi_{g,w}^U \leq \theta_{g,w}^U, \forall g \in \mathcal{G}, \forall w \in \mathcal{W}. \quad (A.6)$$

In the above formulation TO , the objective defined by Equation (A.1) minimizes the total costs that include the following costs terms: (a) total cost of the route, $\sum_{i \in \mathcal{R}} C_i x_i$, (b) total penalty cost for not covering an order by any equipment route, $\sum_{o \in \mathcal{O}} B_o \xi_o$, where B_o and ξ_o are the corresponding penalty cost and slack variable associated with not covering order o with any equipment route, and (c) total penalty cost for exceeding the

tank-wash capacity, $\sum_{w \in \mathcal{W}, t \in \mathcal{T}} B_{w,t} \xi_{w,t}$, where $B_{w,t}$ and $\xi_{w,t}$ are the corresponding penalty cost and slack variable associated with exceeding the tank-wash capacity $K_{w,t}$ at tank wash w at period t , respectively, and (d) the total penalty cost for not meeting the lower and upper thresholds of equipment group at a tank wash, $\sum_{g \in \mathcal{G}, w \in \mathcal{W}} (B_{g,w}^L \xi_{g,w}^L + B_{g,w}^U \xi_{g,w}^U)$, where $B_{g,w}^L, B_{g,w}^U$ are the penalty costs associated with lower threshold $\theta_{g,w}^L$ and upper threshold $\theta_{g,w}^U$, respectively, and $\xi_{g,w}^L, \xi_{g,w}^U$ are the corresponding slack variables.

Next, we describe the constraints for the model. Equation (A.2) ensures that for each piece of equipment e , we can only have one route scheduled in the optimal solution. Equation (A.3) ensures that each order o is either assigned to a route or incurs a penalty cost for being unscheduled. Equation (A.4) enforces tank-wash capacity $K_{w,t}$ for tank wash w for each period t . Finally, Equations (A.5) and (A.6) ensure that for each equipment group (e.g., type, valve location) g , a minimum threshold $\theta_{g,w}^L$ and maximum threshold $\theta_{g,w}^U$ of equipment must arrive at the tank wash w at the end of time t .

A.2. Driver Scheduling with Trailer Optimization

In this section, we present the combined driver-trailer wash optimization model that schedules drivers to perform work assignments against an order and may select a different trailer than the one that was provisionally selected in the

Figure A.4. Cost Values for the Work-Assignment Schedules for Several Different Drivers





Figure A.5. Cost Values for the Options Generated to Cover Several Different Orders

trailer-only phase. We let $\mathcal{S}$ represent the union of work assignments (shipments) across all orders. Each order $o \in \mathcal{O}$ is thus decomposed into a sequence of work assignments, $s_o, m, m = 1, \dots, l_o$ that include wash-to-pickup, pickup-delivery, delivery-to-wash legs, and possibly additional stops or relays. We use $\mathcal{D}$ to represent the set of drivers.

If we designate the pickup and delivery location pair as $(p_s, f_s), p_s \in \mathcal{P}, f_s \in \mathcal{F}$ for each work assignment $s \in \mathcal{S}$, then a work-assignment sequence for driver d can be considered as a walk represented by $\omega_d = (w_1, p_{s_1}, f_{s_1}, \dots), w_1 \in \mathcal{W}, \omega_d \in \Omega$. Note that a typical driver work-assignment sequence will include shipment legs from several different orders. Each candidate driver work-assignment sequence $\omega_d \in \Omega$ incurs a cost C_{ω_d} comprising unused hour costs and various bonuses and penalties. We define some additional sets related to work-assignment sequences in Table A.2.

Let $y_j \in Y$ be a binary variable that equals 1 if driver work-assignment sequence j is selected, and 0 otherwise. For a given set of feasible work-assignment sequences $\Omega_{feasible}$, we develop a set-partitioning formulation with side constraints, called *DO* as follows:

$$DO : \min \quad z_2 = z_1 + \sum_{j \in \Omega} C_j y_j + \sum_{s_1 \in \mathcal{S}} B_s \xi_s + \sum_{d \in \mathcal{D}} B_d \xi_d \quad (A.7)$$

Subject to

$$\sum_{r_e \in \mathcal{R}^e} x_{r_e} + \xi_e = 1, \forall e \in \mathcal{E} \quad (A.8)$$

$$\sum_{r_o \in \mathcal{R}^o} x_{r_o} + \xi_o = 1, \forall o \in \mathcal{O} \quad (A.9)$$

$$\sum_{r_w \in \mathcal{R}^w} x_{r_w} - \xi_{w,t} \leq K_{w,t} \quad \forall w \in \mathcal{W}, \forall t \in \mathcal{T} \quad (A.10)$$

$$\sum_{r_{g,w} \in \mathcal{R}^g \cap \mathcal{R}^w} x_{r_{g,w}} + \xi_{g,w}^L \geq \theta_{g,w}^L, \forall g \in \mathcal{G}, \forall w \in \mathcal{W} \quad (A.11)$$


$$\sum_{r_{g,w} \in \mathcal{R}^g \cap \mathcal{R}^w} x_{r_{g,w}} + \xi_{g,w}^U \leq \theta_{g,w}^U, \forall g \in \mathcal{G}, \forall w \in \mathcal{W} \quad (A.12)$$

$$\sum_{\omega_d \in \Omega^d} y_{\omega_d} + \xi_d = 1, \forall d \in \mathcal{D} \quad (A.13)$$

$$\sum_{\omega_s \in \Omega^s} y_{\omega_s} + \xi_s = 1, \forall s \in \mathcal{S} \quad (A.14)$$

$$\sum_{\omega_i \in \Omega^i} y_{\omega_i} \leq m x_i, \forall i \in \mathcal{R}, m \text{ sufficiently large} \quad (A.15)$$

Figure A.6. Percentage Rank of the Selected Optimum and Candidate-Set Size for the Collection of Orders in a Representative Solution Run



In the above formulation *DO*, the objective defined by Equation (A.7) minimizes an overall cost comprising the following costs terms: (a) the total cost associated with selected trailer routes TO , z_1 , (b) the total cost of selected driver work-assignment sequences, $\sum_{s_1 \in \mathcal{S}} B_s \xi_s$ including penalty and (or) bonus costs related to pickup time violations and work-assignment priorities and a penalty for not covering work assignment s with any driver, and (c) a penalty for leaving an available driver idle (i.e., with no work-assignment sequence within the time scope of the optimization run).

Next, we describe the constraints used in model *DO*. Equations (A.8)–(A.12) are similar to the Equations (A.2)–(A.6) used in the trailer-only optimization model *TO*. Equation (A.13) ensures that a driver d can only have one driver work-assignment sequence scheduled in the optimal solution. Equation (A.14) ensures that each work assignment s can be assigned to a route or it can be unscheduled. Equation (A.15) matches the driver work assignments to the corresponding equipment routes.

A.3. Implementation Details

The system is primarily written using Java (v1.8). The code that implements user interactions was developed using Java Swing class libraries; the persistence model and data storage are built on an Oracle database; optimization models are implemented using IBM Java Concert APIs; and the linear programming/mixed-integer programming (LP/MIP) solver is IBM CPLEX (v12.7). With the current CPLEX settings, the realized MIP gap ranges from 0.05% to 0.17%. Typical objective function values (less penalty terms) are in the range of \$1.6 million, so the achieved solutions are within \$300 to \$1,000 of the optimal solution.

The system interoperates with several commercial software systems also developed by Oracle. These include Seibel order management and Oracle Transportation Management (OTM) execution management. Order and customer data are transmitted from the order management system to the

dispatch system. Driver and trailer information is maintained by the execution management system and assignment recommendations are passed back to that system. The dispatch system runs on a virtual server running Oracle Enterprise Linux (v6.8) with 10 cores and 48 GB of memory. The underlying physical hardware is an active-active failover cluster of 10 Lenovo blade servers, each with 40 CPUs at 2.26 Ghz.

Figures A.1 and A.2 give representative instance sizes and processing times for the solution components referenced in Figure 7. For the optimization models, the number of variables is equivalent to the number of potential solutions.

Tables A.3 and A.4 give representative instance sizes and processing times for the solution components referenced in Figure 7. For the optimization models the number of variables is equivalent to the number of potential solutions.

A.4. Empirical Analysis of Heuristic Steps

As we mentioned in the main paper, the heuristic aspect of our methodology is that candidate solutions for trailer routes and for driver work-assignment schedules are not generated exhaustively. While the generation methodology ensures the new candidates for routes and schedules are discovered in increasing-cost order, it is possible that we may terminate the search for candidate solutions before the true *optimal* choice is uncovered. To keep processing time under the business-imposed limit, we currently have cutoff thresholds set as the maximum cost of a driver schedule = \$1,000; the maximum number of drivers per work assignment = 50; the maximum number of work assignments per driver = 100; and the threshold for adding additional assignments to a work schedule = 36 hours. For a random sample of specific actual problem instances, we have increased these limits and have not achieved overall solution improvements. It should also be noted that these cutoffs are well above what a planner would evaluate when *manually* considering assignment options.

We have also looked at how the costs associated with an *optimization-selected* option compare with those of all other candidates generated. Figures A.3–A.5 give this information for the candidate sets of a few trailers, drivers, and orders, respectively. Here, the intuition is that it is highly unlikely that ungenerated options (i.e., options with yet higher costs) would have been selected.

In the case of order coverage, where the candidate-set size can be in the thousands, we performed an additional analysis. We observed that in rare cases where the cost of the optimal-selection cost is near the maximum of the candidate set, the size of the candidate set is relatively small. This is a strong indication that in these cases there are no ungenerated higher-cost options. Figure A.6 shows that orders covered by selections near the maximum have relatively small candidate-set sizes.

References

- Cacchiani V, Salazar-Gonzalez JJ (2017) Optimal solutions to a real-world integrated airline scheduling problem. *Transportation Sci.* 51(1):250–268.
- Choi E, Tcha DW (2007) A column generation approach to the heterogeneous fleet vehicle routing problem. *Comput. Oper. Res.* 34(7):2080–2095.
- Dantzig G, Ramser J (1959) The truck dispatching problem. *Management Sci.* 6(1):80–91.
- Desrochers M, Verhoog T (1991) A new heuristic for the fleet size and mix vehicle routing problem. *Comput. Oper. Res.* 18(3):263–274.
- Euchi J, Chabchoub H (2010) A hybrid tabu search to solve the heterogeneous fixed fleet vehicle routing problem. *Logist. Res.* 2(1):3–11.
- Gifford T (2011) Operations research applications in truckload freight transportation. Cochran JJ, Cox LA Jr, Keskinocak P, Kharoufeh JP, Smith JC, eds. *Wiley Encyclopedia of Operations Research and Management Science* (John Wiley & Sons, Hoboken, NJ).
- Goel A, Irnich S (2016) An exact method for vehicle routing and truck driver scheduling problems. *Transportation Sci.* 51(2):737–754.
- Goel A, Vidal T (2014) Hours of service regulations in road freight transport: An optimization-based international assessment. *Transportation Sci.* 48(3):391–412.
- Golden B, Assad A, Gheysens F (1984) Fleet size and mix vehicle routing problem. *Comput. Oper. Res.* 11(1):49–66.
- Koç Ç, Bektas T, Jabali O, Laporte G (2016) Thirty years of heterogeneous vehicle routing. *Eur. J. Oper. Res.* 249(1):1–21.
- Li F, Golden B, Wasil E (2007) A record-to-record travel algorithm for solving the heterogeneous fleet vehicle routing problem. *Comput. Oper. Res.* 34(9):2734–2742.
- Nazi-Azmi Z, Salari M (2013) A complementary tool to enhance the effectiveness of existing methods for heterogeneous fixed fleet vehicle routing problem. *Appl. Math. Model.* 37(6):4316–4324.
- Penna P, Subramanian A, Ochi L (2013) An iterated local search heuristic for the heterogeneous fleet vehicle routing problem. *J. Heuristics* 19(2):201–235.
- Savelsbergh M, Sol M (1998) DRIVE: Dynamic routing of independent vehicles. *Oper. Res.* 46(4):474–490.
- Sherali H, Bae K, Haouari M (2013) An integrated approach for airline flight selection and timing, fleet assignment, and aircraft routing. *Transportation Sci.* 47(4):455–476.
- Subramanian A, Penna P, Uchoa E, Ochi L (2012) A hybrid algorithm for the heterogeneous fleet vehicle routing problem. *Eur. J. Oper. Res.* 221(2):285–295.
- Taillard E (1999) A heuristic column generation method for the heterogeneous fleet VRP. *RAIRO (Oper. Res.)* 33(1):1–14.
- Xu H, Chen Z, Rajagopal S, Arunapuram S (2003) Solving a practical pickup and delivery problem. *Transportation Sci.* 37(3):347–364.
- Yaman H (2006) Formulations and valid inequalities for the heterogeneous vehicle routing problem. *Math. Programming* 106(2):365–390.

Verification Letter

John Bozec, Senior Vice President/General Manager of Schneider Bulk Division, writes:

“Please consider this letter as *written verification of success in practice* for this application for the 2017 Daniel H. Wagner Prize for Excellence in Operations Research Practice.

The Schneider Bulk Transport Dispatch Optimizer System is a suite of software components and associated business processes that support and recommend real-time decisions allocating truck drivers, tractors, and bulk tanker trailers to orders from customers to transport liquid chemicals. As described in the accompanying abstract, this assignment problem is rendered particularly difficult by operational complexities introduced by the chemical properties of the commodities transported, interorder tank wash requirements, tanker trailer configuration options, and driver work rules and conditions.

Prior to implementation of this system, the process was largely manual (and in our estimation remains so throughout the industry) and overly reliant on human judgment, tribal knowledge, and disparate, difficult-to-maintain data sources. The new system substantially increases automation and uses mathematical optimization techniques to select optimal assignment choices. As a consequence, it has, in its first few months of operations, already generated significant gains in productivity for assets, drivers, office staff, and management; as well as significant direct cost savings in fuel and other direct expenses. I will quantify some of these savings below.

Project Financial Justification

An initial project assessment was performed prior to authorizing the financial spend. Three main benefit areas were identified: cost avoidance, revenue generation, and associate productivity. Since the inception of the project, annualized benefits have been

\$2.3 million cost avoidance (empty mile improvement):
19% more than projected

\$2.5 million additional revenue (improved driver productivity): 20% more than projected

28% planner productivity: 18% more than projected

Empty Mile Reduction

Every bulk order incurs both pretank wash (TW) and post-TW empty miles. The pre-TW miles are prior to pick up when a clean and prepared trailer is brought from a TW to the shipper. The post-TW miles occur after delivery and the empty trailer needs to be brought to a TW for cleaning. Combined, these are termed “unbilled miles per order” (UBMPO). The project has led to a 24-mile reduction (16 pre, 8 post) in UBMPO. This provides an annualized reduction of over 1 million empty miles. With complete rollout in the initial market (South), we are seeing a total reduction UBMPO has been reduced year over year by 86 miles per order (bobtail + deadhead) in March. It is also less volatile with the daily UBMPO ranging from 260 to 310

in 2017. The same date range in 2016 saw a fluctuation 290 to 460 miles per order.

Driver Productivity Improvement

Avoiding unbilled empty miles has resulted in drivers able to run more revenue generating miles. Approximately 40% of avoided empty miles directly result in a billed mile. This generates \$2.5 million additional revenue.

Planner Productivity Improvement

Year over year, an additional 900 orders have been planned by four fewer associates. This is a 28% productivity improvement. Additional productivity gains are expected when associates have gained expertise with the new processes and additional enhancements to the optimization models are made.

Additional Project Benefits

There are additional process improvement, asset utilization, and customer satisfaction benefits as a result of the project. Some benefits have objective metrics and others are anecdotal.

Better Equipment Optimization. The system enabled a structure change on the network planning team. Managing both driver and trailer assignment has been combined into a single planner role (this was not possible prior to the new system). The Area Planning Manager now has complete visibility to all dispatch activity in a market. Consequently, we have been able to reduce idle days by 1 day per trailer-month. This equates to approximately 12,000 additional days of trailer capacity availability per year, effectively adding 60 trailers to the fleet at no additional cost.

Associate Productivity: Prior Product. Prior to implementation of the business rules and associated databases that maintain and enforce prior product (previous commodity contents of tank) compatibilities, feasible trailer identification and verification would often take 30 minutes. The process now completes in 3–5 minutes and often much less. This translates to several thousand man-hours per year.

Improved Customer Experience. Prior products were previously searched for manually. Because of the technical nature of the process, errors are common and drivers arrive with trailers that don't meet requirements. This results in rejected trailers which is both costly in unbilled miles and frustrating for customers. Trailer rejections for bad prior products have already been reduced substantially and customer feedback has been very positive (as better information from us helps them improve their own processes. Additionally, when trailers are wrongly rejected, the new systems and processes support quick recognition of this and allow timely billing of customer for costs incurred. Conversely, when customers correctly reject trailers we are now able to update data to prevent repeat occurrences.

Unused Hours. The percentage of driver assignments with excess unused hours (waste) has already significantly decreased in the South in cases in which model recommendations are

followed. This has improved driver retention and increased driver and asset utilization.

Order Acceptance Response Time. The new system has enabled quicker and more accurate visibility to available and projected capability. This has led to a reduction in the time required to make an acceptance decision. It remains to assess the financial impact of this benefit.

This conservative estimate is based on actual performance during the first several months of operation. As the systems mature and learnings are incorporated into improvements and enhancements, this number will increase."

Ted Gifford is a distinguished engineer at Schneider National Inc., where he previously served as Vice President of Engineering & Research. Prior to joining Schneider, Gifford was a member of mathematical sciences faculty and an associate dean at the University of Alaska. His various technical leadership and research roles include President of Computer Consultants of Alaska, Director of Quantitative Research at McKinley Capital Management, and Senior Engineering Manager at Symantec Corporation.

Tracy Opicka is a senior optimization engineer in the Engineering and Advanced Analytics group at Schneider National Inc. She earned a bachelor's degree in mathematics/computer science from St. Norbert College and a master's degree in industrial engineering from Purdue University.

Ashesh Sinha is an assistant professor at the Department of Industrial and Manufacturing Systems Engineering at Kansas State University. He received a bachelor's degree at the Indian Institute of Technology in Kharagpur, India. He earned a master's degree in manufacturing systems engineering and a doctorate in industrial engineering at the University of Wisconsin-Madison. Before coming to Kansas State, he worked as an optimization engineer at Schneider National Inc. from 2016 to 2017.

Daniel Vanden Brink is head of Global Analytics for ThyssenKrupp Aerospace. In his previous role, he was Vice President of Engineering and Analytics at Schneider National Inc. Prior to this, he led IBM's Worldwide Optimization and Supply Chain Team. He has a master's degree in Operations from IIT and a bachelor's degree in Industrial Engineering from Iowa State University.

Andy Gifford is a software developer at Schneider National Inc. He received a bachelor's degree in computer science from the University of Wisconsin Oshkosh in 2010. His primary focus has been on the development of transportation planning and dispatch applications.

Robert Randall is a senior optimization specialist at Princeton Consultants, where he develops custom production optimization models and decision support systems in areas such as fractional airline scheduling, less-than-truckload (LTL) flow plan optimization, rail car allocation, and bulk chemical trailer assignment. He earned a PhD in industrial engineering from Clemson University, with a focus on metaheuristic design and development.