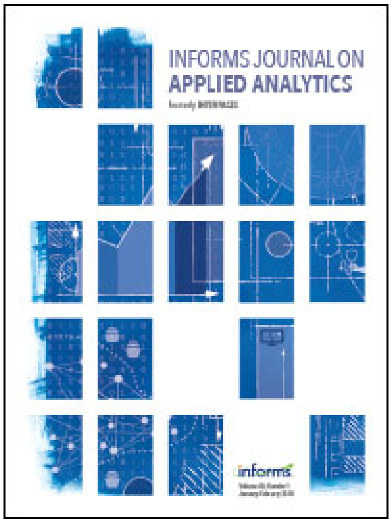

INFORMS Journal on Applied Analytics

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

Network Mode Optimization for the DHL Supply Chain

Yibo Dang , Manjeet Singh , Theodore T. Allen

INFORMS JOURNAL ON APPLIED ANALYTICS
Vol. 51, No. 3, May–June 2021, pp. 179–199
ISSN 0092-2102 (print), ISSN 1526-551X (online)
http://pubsonline.informs.org/journal/inte

Network Mode Optimization for the DHL Supply Chain

Yibo Dang,[a,b] **Manjeet Singh,**[a] **Theodore T. Allen**[b]

[a] DHL Supply Chain, Westerville, Ohio 43082; [b] The Ohio State University, Columbus, Ohio 43210

Contact: yibo.dang@dhl.com, https://orcid.org/0000-0002-2474-3234 (YD); manjeet.singh3@dhl.com, https://orcid.org/0000-0001-9558-4640 (MS); allen.515@osu.edu, https://orcid.org/0000-0002-9522-3252 (TTA)

Received: September 12, 2019
Revised: February 4, 2020; May 27, 2020
Accepted: June 7, 2020
Published Online in Articles in Advance: February 5, 2021

https://doi.org/10.1287/inte.2020.1046

Abstract. DHL Supply Chain North America moves more than 20 million packages each year. DHL transportation planners perform routing and cost-deduction tasks for many business projects. We refer to the associated planning problem as the *Vehicle Routing Problem with Time Regulations and Common Carriers* (*VRPTRCC*). Unlike ordinary vehicle routing problems, which use only a single type of transportation mode, our VRPTRCC applications include make–buy decisions because some of the package deliveries are ultimately subcontracted to organizations other than DHL. *Time regulation* means that the problem considers not only delivery-time windows, but also layover and driving-time restrictions. Our developed Network Mode Optimization Tool (NMOT) is an ant-colony optimization (ACO)-based program that aids DHL Supply Chain transportation analysts in identifying cost savings in the ground logistic network. By using the NMOT, DHL and its customers have saved millions of dollars annually. Also, the NMOT is helping DHL to win new customers against bidding competitors and reducing estimation times from multiple weeks to hours. The results show an actual increase in profits compared with the previous process by more than 15% through a combination of new projects enabled and reduced current operational costs. The NMOT is implemented and evaluated by using data from ongoing projects.

History: This paper was refereed.

Keywords: transportation • vehicle routing problem with time regulations and common carriers • network mode optimization • make–buy decisions • ant colony optimization

Introduction

DHL Supply Chain North America operates one of the world's largest logistics networks, delivering more than 1 billion packages annually for corporate customers. Owing to limited labor and truck resources, many companies utilize more than a single transportation mode to reduce their shipping cost. The two transportation modes considered in this paper and used most by DHL are (1) using the company's own dedicated fleet of trucks and (2) subcontracting the delivery to one or more third-party common carriers.

Transportation costs for shipments depend on factors such as total number of miles on the route, number of stops, and load weight and volume. Subcontracting to third parties is preferred when the total customer demand exceeds the capacity of DHL's own fleet, if the fleet cannot guarantee the delivery within the customer's time window, or if it is more economical to do so.

We recently developed the Network Mode Optimization Tool (NMOT) to help DHL improve their transportation-mode decisions. The NMOT uses future demand information to determine which packages should be delivered by DHL and which packages should be outsourced to a third-party carrier. DHL estimates that the NMOT has already contributed more than $5 million in savings for its corporate customers.

DHL runs the NMOT for each corporate customer separately. Most customers have similar needs over time; therefore, the solutions give a highly accurate estimate of real transportation cost. The NMOT is being utilized (a) for helping existing customers and (b) during the bidding process for new customer contracts. For existing customers, NMOT solutions can be used to switch the shipment mode of planned deliveries, from dedicated fleet to third-party carriers or vice versa, to reduce shipping cost. Alternatively, during the bidding process, the NMOT can help in determining a practically viable least-cost solution that maximizes the chance of winning profitable bids and avoiding bids that are not profitable.

Background

DHL is uniquely positioned in the logistics world with a comprehensive range of international express, freight transportation, e-commerce, and supply chain management services. The group employs approximately 550,000 employees in more than 220 countries and territories worldwide. As the world's leader in contract logistics, DHL Supply Chain offers standardized warehousing, transport, and value-added services that can be combined to form customized supply chain solutions.

Because of the complexities associated with running a logistics network, many companies employ professional third-party logistic companies to manage and operate their supply chain business. Large third-party logistics companies move millions of packages every year. These companies are constrained by limited capacity and growing demands, which is a bottleneck for the logistics industry. In this paper, we discuss DHL Supply Chain's advanced analytical methodology called the NMOT that addresses this problem. DHL uses the NMOT to prepare low-cost practical proposals during the new-customer contract-bidding process.

To prepare competitive bidding proposals, analysts must determine the appropriate division of deliveries between the two transportation modes. During this process, the analysts construct detailed routes with the associated costs for both the modes. This facilitates the make–buy decisions—that is, which deliveries should be done by DHL's dedicated fleet, in what sequence, and which others should be outsourced to a third-party freight provider (or referred to as a common carrier), to minimize the total transportation costs.

The make–buy assessment process has two steps: (a) routing and (b) cost estimation. (a) First, routing is performed by using a combination of commercial routing software and the analyst's experience or tribal knowledge. (b) Second, based on the routes from step (a), overall costs of all deliveries for both transportation modes are manually estimated and used in make–buy decisions. It is difficult to ensure that all practical routing and customer constraints are satisfied while achieving profitability.

Before the start of this project, we talked to DHL employees in several regions from the operations teams, the solution-design teams, and the sales teams, all of whom agreed that an optimization tool for the transportation modes is crucial for marketing their business and guaranteeing operational effectiveness. In addition, virtually all the analysts mentioned that the previous process usually takes several weeks and requires extensive geographical and market knowledge. Because of the complexity and the size of the associated delivery problem, current off-the-shelf optimization tools are not a viable solution.

Objectives

The NMOT helps DHL analysts in preparing low-cost proposals for bidding on new customer contracts. To generate feasible solutions, the NMOT considers various practical aspects within DHL's road-transportation network. For example, the scheduled delivery locations can be stores, warehouses, a consolidation hub, or a mix of these locations. Additionally, the NMOT considers routing constraints and the associated costs based on geography and truck types, product type being shipped, driver's layovers, and time windows.

The inputs include unrouted shipments, common-carrier costs, geographical information of distribution centers and demand points, and product classes. The outputs are the expected costs, routes, schedules, and designated transportation modes (DHL's dedicated fleet or a common carrier). With NMOT solutions, the bidders can approach customers with operational transportation-cost estimates from DHL. Then, customers usually compare these estimates with costs from the previous processes or quotations from other competitors. NMOT solutions have helped in winning more than 10 transportation operations bids with significant associated gains in profitability.

The primary goal of this project was to replace a semimanual and iterative planning activity with a reliable advanced analytical model (NMOT) to support the bidding process and assess the shipment mode for current corporate customers. Another objective of the project was to shorten the planning time, so that fewer analysis hours would be needed for bid generation. The question is how to quickly and optimally route the huge less-than-truckload network. Although we developed the NMOT to accommodate a wide variety of real situations of DHL Supply Chain transportation networks, it is also capable of solving a standard problem. The most basic and standard version of our problem is the well-known Vehicle Routing Problem (VRP). In general, the VRP is concerned with the optimal design of routes in a transportation network. Specifically, the objective of the NMOT is to identify opportunities to design or modify the transportation networks for DHL's corporate customers to reduce costs. The timely and high-quality solutions generated by the tool will provide DHL with a competitive and feasible estimate of the operational costs that can be quoted in the bidding process. In fact, costing was previously the most time-consuming part in the make–buy assessment process.

The NMOT deals with a complex transportation network that requires sophisticated attentions to the optimization constraints. The dedicated fleet must visit the assigned customers within allowed delivery-time windows. And the trucks cannot be overloaded or overused. In other words, the truck-capacity, route-distance, and time-limits constraints should be satisfied. In addition, after a driving and unloading period of 14 hours, the driver shall take an uninterrupted daily layover period of 10 hours. Modeling challenges include (1) a combination of operational constraints; (2) a make–buy decision; (3) robustness to various transportation networks, areas, and classes of shipments; and (4) solving a large-scale, real-life problem.

This paper does not consider travel and work time stochasticity; therefore, we approximate the model parameters with historical averages. Additionally, the objective is to generate efficient routes to get operational cost estimates for the bidding process. Therefore, we did not consider detailed routing parameters like unusual traffic impacts or route resilience factors.

Literature Review

Since the first VRP problem proposed by Dantzig and Ramser (1959), tremendous research has been focused on solving this combinatorial problem. Our study of the literature on various VRP problems shows the limitations of effective approaches to handle large-scale instances. The VRP has many variants, depending on the parameters and constraints, which are generally NP-hard (Savelsbergh and Sol 1995). The basic VRP is the Capacitated Vehicle Routing Problem (CVRP), which assumes a fixed fleet with uniform demands supplied by a central depot (Dondo and Cerdá 2007). This paper solves an extension of the CVRP, for DHL, which we call the *Vehicle Routing Problem with Time Regulations and Common Carriers* (VRPTRCC). This section focuses on the following three aspects: (1) previous research that addresses the constraints introduced in VRPTRCC, (2) existing solution methods used for solving variants of the VRP, and (3) reasons for choosing the solution technique used in the NMOT.

The VRP with subcontracting or time regulations has been researched in few studies. Moon et al. (2012) were the first to consider outsourcing options in the VRP. Yet, these authors only considered the time-window constraints, whereas this paper considers layover time and other detailed operational constraints. Vidal et al. (2016) considered common carriers to maximize the vehicle-routing profits, but did not consider time regulations studied in this paper. Gahm et al. (2017) introduced a new variant of the VRP with common carriers considering a heterogeneous dedicated fleet and multiple cost options for the last-mile decisions. Yet, those authors only considered the last-mile routing, whereas this paper addresses the mode decisions for the entire route. Kok et al. (2010) and Goel (2018) studied the impact of the drivers' time regulations on VRPs. Note that none of the aforementioned works merged the outsourcing decisions and distance constraints with the time regulations, as studied in this paper. More recently, Alcaraz et al. (2019) combined heterogeneous fleet types with last-mile make–buy decisions and drivers' hour regulations (layovers). Alcaraz et al. (2019) compared their results with simple intuitive heuristics on small samples. Although they considered comprehensive constraints similar to our work, their problem considered outsourcing decisions for last-mile deliveries only. In its place, their work focuses on other aspects, such as incompatibilities of goods and split deliveries.

Heuristics and metaheuristics have been widely studied in solving complicated vehicle routing problems. Wu et al. (2017) adjusted the local search heuristic for a VRPTW with different less-than-truckload carriers' selections. However, their experimental results only supported up to 10 customers. Another work related to our research was conducted by Cassettari et al. (2018), in which they developed a multistage clustering-based heuristic to solve the CVRP with time-window and distance constraints. Vidal et al. (2013) studied more than 64 metaheuristics on 15 classic variants of VRPs, suggesting the relevance of ant-colony optimization (ACO) and other hybrid methods. Probably the most convincing evidence to use swarm intelligence (e.g., ACO) to solve Rich VRPs was inspired by the work from Pellegrini et al. (2007). The authors presented a case study on a hybrid VRPTW with two versions of ACOs. The results, compared with tabu search, nearest-neighbor search, and simulated annealing, significantly favored ant-colony systems as the best solution heuristic. Another ACO implementation, proposed by Rizzoli et al. (2007), has been applied to real contexts addressing separately a heterogeneous fleet, time windows, pickup and delivery constraints, and time-dependent deliveries. The authors tested four ACO algorithms using data from real distribution companies with 15–600 customers. Admittedly, none of the heuristic structures can guarantee optimality, but such well-designed local search procedures can greatly improve the solution performance.

Apart from the heuristic methods, we researched large-scale exact methods in the literature for solving VRP variants. These are branch-and-cut, Lagrangian relaxation, and column generation. Bard et al. (2002) discussed an exact branch-and-cut method in solving the VRPTW. Kallehauge et al. (2006) developed a new Lagrangian duality algorithm for VRPTW. Their experimental results showed that the method can solve the problem to optimality for sizes up to 200 customers. Bettinelli et al. (2011) considered a complex variant of the VRP, which involves heterogeneous fleet sizes, multidepot, and time windows. The authors presented a column-generation framework with multiple exact and heuristic pricing and cutting techniques. Moreover, Wen et al. (2011) considered a similarly constrained problem, which contains multiperiod horizon, time windows for delivery, heterogeneous vehicles, drivers working regulations, and other constraints. The author proposed a mixed-integer linear program (MILP) embedded in a multilevel local search algorithm. Good-quality solutions dealing with as many as 500 customers are generated by using real case information. In addition, Desaulniers et al. (2017) proposed exact algorithms for Electric VRPTW. The results were tested with a maximum of 100 nodes.

Typically, exact methods usually work for hundreds of customers, whereas metaheuristics or heuristic-based relaxation algorithms can provide optimal or near-optimal solutions for larger data sets. Furthermore, there is limited work discussing the implementation of algorithms on real-world, large-scale VRP problems. Finally, it can be concluded that none of the previous research combines all the comprehensive constraints and detailed problem space as solved in this paper. Our research fills this gap by solving VRPCCTR on large-scale, real-life problems at DHL.

The contribution of our research can be summarized in three points: (1) It studies an integrated vehicle routing problem, which schedules complicated route allocation and transportation-mode selection; (2) it provides a high-efficiency metaheuristic that routinely solves large-sized business problems in real situations; and (3) the developed tool and solutions are successfully applied to the network-mode optimization at DHL Supply Chain North America, a leading logistic company with complex operational protocols and solution needs.

The data used in this study are provided by the transportation team at DHL Supply Chain North America, which includes networks in the United States, Mexico, and Canada. The remainder of this paper is organized in the following sequence of sections: Problem Descriptions, Modeling and Solution Methods, NMOT Implementation, Practical Applications, Conclusions, and Appendix (Mathematical Models and Results Comparison).

Problem Descriptions

DHL Supply Chain provides transportation network operation, optimization, and consulting services to their corporate customers. One of the key elements for the success of DHL's logistic services is optimization of transportation costs. Data sent from DHL's existing customers or potential customers are analyzed, and tentative transportation-mode decisions are made among dedicated fleets and other carrier providers. The network assignments are then created and optimized. A high-quality assignment solution will not only increase the possibility of success in bidding for new contracts, but also offer useful guidance for positioning fleet resources on the network. DHL's NMOT addresses a novel problem that has received relatively scarce attention in the previous studies, perhaps because few organizations operate at our large scale.

• Make–buy decisions: Make–buy decisions are the subcontracting decisions as to whether DHL operates certain shipments with their dedicated fleets or outsources them to a common carrier.

• Shipments: Shipments are the items to be delivered between locations. In the model, they represent the weights of cargo that requires transportation from a DHL warehouse to a destination site in a specific time window (i.e., the time between the moment the cargo becomes available at the origin warehouse and the time the cargo must be delivered at the destination location, subject to the destination's work shifts). Each shipment likely represents a specific cargo class that has a specific rating scheme or service standard (e.g., machine classes may be evaluated and rated by both weight and volume constraints; chemical materials may need to be shipped in separate trailers). The model thus requires sophisticated subcomponents to handle all the parts and aspects of the shipments.

• Routes: Routes are multistop Hamiltonian cycles for transporting shipments. In the VRPTRCC model, the routes represent different transportation modes from the origin to the destination within the required time window and layovers. Routes may be direct (one arc from origin to destination) or multistop (consisting of a series of arcs between consecutive destinations visited by the route). The definition of a route includes every aspect of the trip: capacity of the trailer, time of arrival and departure at each stop, cost, modes (dedicated fleet or common carriers), layovers, shipment classes, travel distances, and so forth.

• Capacity constraints: For every arc of every route, the total weight and volume associated with the shipments routed through a multistop arc must be less than or equal to the arc's capacity. Therefore, the NMOT allows oversized shipments to be outsourced through a make–buy decision. In addition, for example, our NMOT will also evaluate and optimize the cost reduction of shipping oversized shipments to consolidation hubs by separate trailers.

• Time regulations: Each customer has a prespecified time interval for delivery. In addition, truck drivers are constrained by additional time regulations like working-hour limits and layover limits. Each time the driver hits a working-hour limit, he or she needs to take a break (layover). Thus, any feasible route must serve all the customers within their time windows, observing all the time regulations.

• Additional constraints: The NMOT involves detailed operational constraints, such as maximum allowed layovers, maximum allowed driving range between each layover, maximum allowed distance per route, weekend delivery allowance, maximum allowed intranode distance, and so on. Including all the constraints will make the combinatorial network huge and, thus, challenging to solve within acceptable time limits.

• Cost decomposition: After generating the optimal network mode solutions, the NMOT will decompose the optimal route costs by the shipment level. This process automates DHL's Cost Model (CM) rating formulas, which allows the NMOT to give its users a quoted cost for each shipment.

Figure 1 illustrates a representative map of the network. Under the structure of the VRPTRCC system, each warehouse as a depot must have enough associated freight vehicles to supply all the demands. With shipments assigned to warehouses, each driver will take charge of a route, visit the customers exactly one time, and make on-time package deliveries. After the final delivery of his or her route, the driver will return to the depot, thereby forming a so-called Hamiltonian cycle. No late delivery or extra layovers exceeding the constraints are allowed. For certain shipments or routes, outsourcing from a third-party carrier can reduce overall network costs. DHL officials will diagnose such opportunities and make decisions.

Clearly, the number of feasible solutions for this combinatorial optimization problem increases exponentially with the number of customers to be serviced. Because the problem scale may range from many hundreds to half a million shipments and large numbers of distribution centers, previous solution processes are of low efficiency. The previous solution process relied on routing with multiple commercial software programs and manual make–buy decisions. This process ordered the power of multiple software programs and the expertise from the transportation analysts in two stages. When dealing with a large shipment file, repeatedly using multiple software programs not only is time consuming, but also produces less satisfactory results. In addition, because the previous process required the intervention from analysts to finalize the mode decisions, the solution quality could also be impacted by the experience and skill of the analysts.

The primary mathematical objective is to minimize the total cost of moving all shipments from DHL warehouses to their assigned destinations. A typical planning horizon is two weeks before biddings or operations, and the problem is mostly solved by season with historical or predicted data. A built-in assumption—established many years ago by industry conventions—is that each shipment is preassigned to a depot, and the depots work independently, except for cases in which the destination for a shipment may be a depot as well. Relaxing this assumption opens up tremendous opportunities to iterate or parallel multiple depots to solve VRPTRCC.

Previous Process

Prior to our project, the steps in Table 1 were used to assist the decision-making process.

This process can be lengthy and requires the transfer of information between multiple software programs with no quality guarantee. Another long-term drawback of this process is that commercial software standardizes the VRP optimizations, which makes it impossible to add DHL's tailored objective and constraints. These disadvantages often result in highly impractical solutions (e.g., the minimized total distance

Figure 1. (Color online) The Map Shows a General Location of DCs and Mixed Types of Deliveries Within the DHL Supply Chain Transportation Network in the Midwest and Northeast Sections of the United States

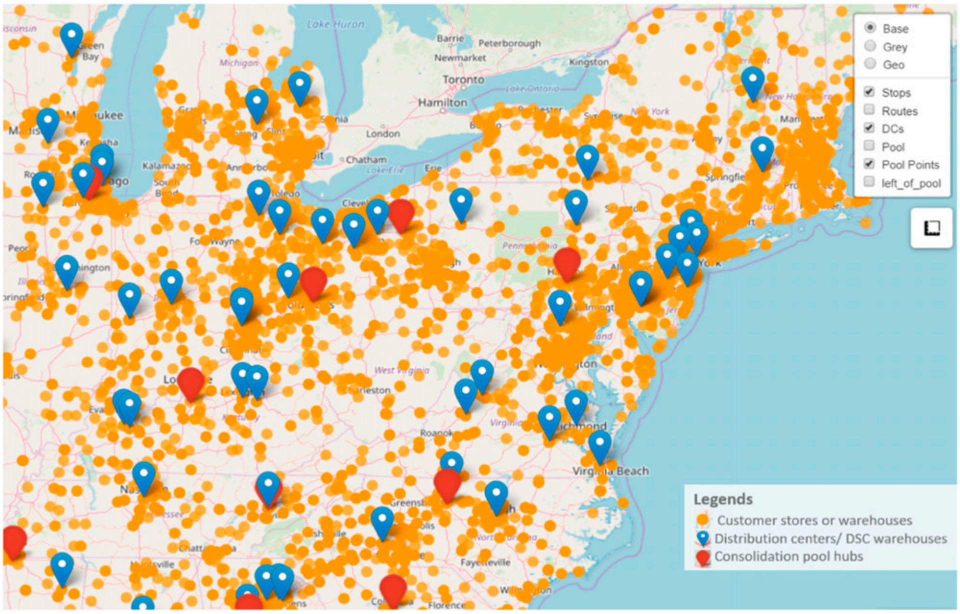

Table 1. Descriptions of the Previous Procedures at DHL Supply Chain

Steps	Previous procedures	Average time (hours)
1	Extract and clean raw shipments information from the database.	8
2	Preprocess the data, sample the subsets, and make initial transportation-mode decisions manually with specified constraints.	16
3	Route the trips and schedule the freight resources with multiple standard commercial software based on mileage ranges.	32
4	Allocate the route cost to each shipment and compare this with the transactional common carrier cost obtained from multiple sources.	24
5	Take the remaining dedicated shipments and repeat the routing procedures.	32
6	Finally, quote the solutions and make proposals to the customers.	8

does not always indicate a minimized route cost). Certain constraints (i.e., *make–buy decisions*, *layovers*, and *additional constraints*) cannot be added to the software, limiting the previous process from generating accurate solutions.

Expectations for the NMOT

Given the drawbacks of the previous process, the following key expectations for the NMOT software were articulated by the leadership of DHL Supply Chain North America:

- Can be added to the standard control report process to quickly interpret fleet potential;
- Increases the accuracy for determining DHL fleet and subcontracted common-carrier decisions;
- Decreases the time it takes to make such planning decisions—estimated savings of at least 1,200 hours per year;
- Increases the promised savings on managed transportation deals, thus improving win probability;
- Grows fleet business; and
- Maintains control of the costs of developing an in-house tool and training the users.

To achieve these goals, the cost-oriented ant colony optimization is to be developed.

Modeling and Solution Methods

The NMOT model optimizes the transportation strategies within DHL Supply Chain and helps analysts quickly make better-informed decisions to drive top-line growth and achieve bottom-line savings. Specifically, the NMOT solves the VRPTRCC with three hybrid ACO-based algorithms. The total cost of routing the shipments is minimized over all the feasible assignments from the make–buy decisions—that is, those that satisfy both the assignment and operational constraints. This cost includes the fixed cost incurred for the use of any route, as well as the variable cost incurred for multistop routing and scheduling. We can mathematically formulate the VRPTRCC model (Mixed Integer Linear Programming Model in the appendix) as an MILP. In this model, the key decisions are which transportation modes to use and how many resources to place on each distribution center. The objective function is formulated to minimize the network cost combining fleet routes and third-party carriers' services. We use this model as the benchmark to examine the solution quality of the NMOT.

From our literature review and DHL's expertise, we identified a number of methods for finding an optimal or high-quality solution for variants of VRPs; these methods include commercial routing software, routing-specific software libraries, heuristic/metaheuristic methods, and optimization routing software libraries, such as Gurobi, IOLG CPLEX, Google OR-Tools, and others. To choose among these methods, we used the requirements discussed below to evaluate the alternatives.

Firstly, *solution time* is an important metric for the performance of the NMOT. The developed tool should produce optimal or near-optimal solutions quickly. Because the scale of our problem can be as huge as half a million shipments, the required time for a medium-sized problem (e.g., thousands of shipments) should be within about two hours. The second standard is the *optimality*. Solutions must be optimal for small cases or near-optimal for the large-sized cases. We use the term *near-optimal* for results that are within approximately 5% of optimal. If a solution is not optimal, there should be no "obvious" improvements (i.e., improvements that can be recognized by simple inspection). Next, the tool should cover *flexibility* and *expandability*—that is, the functionality of the NMOT can be extended and cover additional requirements. For instance, after discussions with the transportation analysts at DHL Supply Chain, we added an option in the developer sheet to solve an *Open-Depot VRP*, which permits the vehicle to stay at the last stop without traveling back to the depot. This is a specific resource-reallocation need for DHL Supply Chain. This particular feature can be added to the model through additional constraints in the corresponding VRPTRCC model and preprocessing of the distance matrices. Also, certain stores and warehouses do not accept weekend deliveries, so the tool should satisfy the time-window constraints and consider different delivery options. While implementing the current versions of the NMOT, we are continuing to add more functionalities to the tool regarding additional constraints. Finally, the *usability* for the analyst is a priority. Most of the practitioners in DHL mentioned to the development team that they prefer a Worksheet-built tool, which can be user-friendly and connectable to their regular work reports. In view of these requirements, we decided to provide a hybrid metaheuristic–optimization solution. Specifically, we used a tailored ACO-based algorithm to be the main structure of our search algorithm, with dynamic programming to optimize the local solutions within neighboring routes.

To refine the searching performance of regular ACO methods, we adopted some optimization and data-structure improvements. One of the key optimization methods in the NMOT is dynamic programming (DP). DP treats the problem as an exponential tree and recursively applies the same reasoning to each condensed subproblem (leaves) to solve the polynomial-sized graphs until reaching the optimal (Cormen et al. 2009). DP is widely used in finding solutions to the shortest-path problem. To fully release the power of ACO and avoid unnecessary calculation waste, we adopt DP in our local search (LS) portion of the algorithm, in which only neighborhood routes are optimized.

We design three scenarios to evaluate the performance of the NMOT. First, we check the solution quality with exact optimal solutions. The VRPTRCC MILP model is solved with the Gurobi software tool kit. Second, we quantify the savings by comparing the NMOT with the previous process. Third, we compare NMOT with an alternative heuristic on selected projects. We describe the specific inputs needed to construct the model and the tool in the NMOT Implementation and Practical Application sections.

Network Mode Optimization Tool Implementation

We now discuss the NMOT implementation in more detail. In particular, we discuss the inputs required to construct the model as described in the Modeling and Solution Methods section and the preprocessing efforts we generated using the available DHL database. Similarly, we discuss the necessary lessons to guide the DHL stakeholders by solving the VRPTRCC model. In addition, we describe the stages in which we improve the NMOT algorithms and how the variants differ from each other. Finally, we discuss error handling and continuous performance improvement.

NMOT Inputs and Preprocessing Efforts

There are different settings for a problem run in the NMOT. We may define it as a set of networks of different depots in close or sparse geographic areas. We may also define it to permit a mix of shipment classes or aggregates and to split oversized shipments through consolidation hubs. Different settings may result in different preprocessing works and different constraints in the model. For example, a setting of homogeneous class on each trailer (mostly occurs for chemical and medical companies) constrained the shipment types to be routed together. If we allow splitting and consolidating the oversized shipments, then a different CM calculation may be required because the pricing is different. Once we have defined the problem, but prior to running the optimization model, it is necessary to pick the representative data and transform the initial data into the form that the tool requires as input. Inaccuracies and inconsistencies (e.g., missing digits in ZIP codes, longitude/latitude mismatch, missing layover time/delivery times, and weight/cube overcapacity) are identified and corrected. Missing values are imputed, and historical data are cherry-picked to estimate the future demands (e.g., peak and off-season, statistical tests). This preprocessing step prepares the input data files based on the optimization settings, tool structure, and format. In the past, DHL ran similar preprocessing procedures for different purposes on multiple software, such as bidding, routing, and other consolidations. Now, we integrate these procedures in the NMOT.

The goal of the shipments preprocessing step is to identify the attributes necessary for each delivery and to generate a standard input and output template that can be used in the optimization model. Beginning with the extraction of data from DHL transportation network systems, Figure 2 shows a summarized flowchart to complete the process.

Three underlying objects of the NMOT model structure are warehouses, shipments, and routes. The optimization model requires the following specific inputs to define an instance of the model:

- The origin distribution center,
- The shipment attributes (e.g., locations, weights, and volumes),
- Distance matrix for any possible combination of shipments,
- The transportation characteristics of the trailers (e.g., speed, capacity, layovers, working hours, and unloading time),
- The fixed transportation costs per day,
- The variable transportation cost (i.e., cost per mile, cost per stop, etc.), and
- Tariffs from the contracted common carriers.

In practice, different customers may have different delivery requests, which can be generalized as com-

Figure 2. Preprocessing Data Steps Used in Previous Process and the NMOT

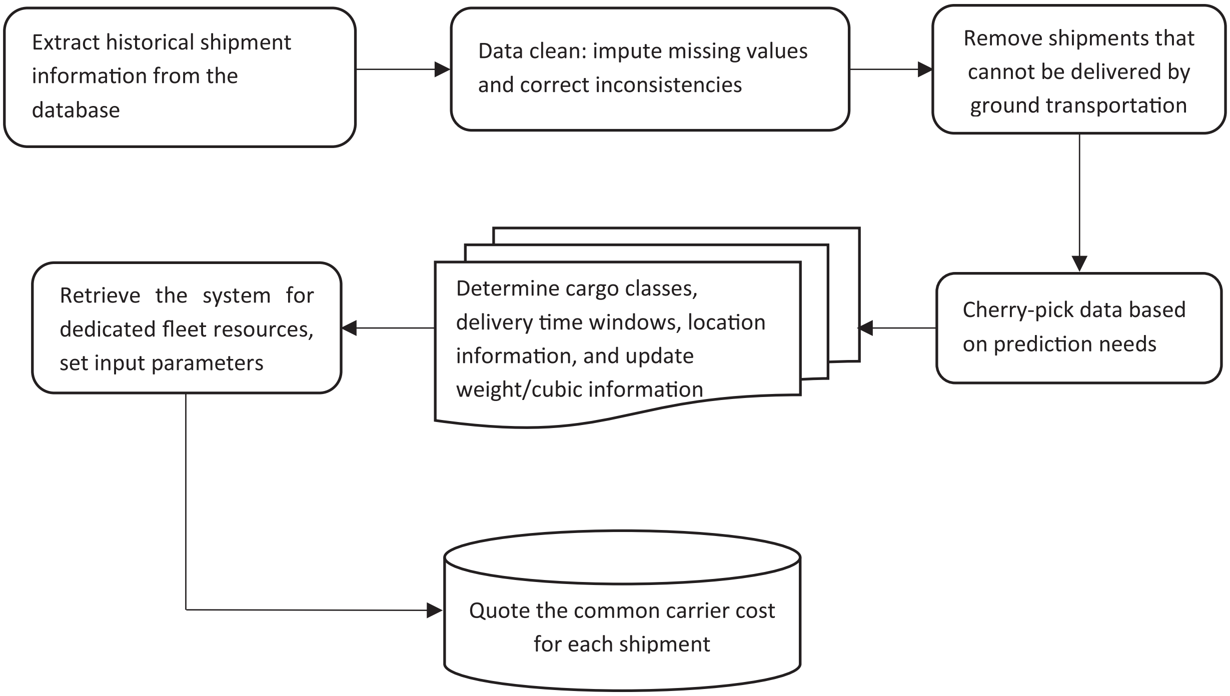

pound time-window inputs. Routes generally will have the same capacity on every arc along the trip because the truck size will remain unchanged. As previously discussed, a particular option is designed to deal with the oversized cases. To evaluate the cost-saving opportunities that the NMOT model has identified, users can use standard input files with our default formats, which contain the regular operational constraints in DHL's transportation network.

Development of the NMOT

The development and test of NMOT algorithms involved three stages, or variants. We designate these three variants as follows: Ant Colony System–Greedy (ACS-Greedy), Ant Colony System Local Searches (ACSLS), and Node-Removal Ant Colony System Local Search (NR-ACSLS). In each of the stages, we ran the algorithm on multiple historical projects, visualized the results on a map, and then communicated with the transportation team. If any impractical routes were observed, such routes and other affected routes were rerun with the NMOT. The results from NMOT reruns gave us insights to improve the algorithms by either adjusting the parameters or using new optimization techniques. Because heuristic algorithms do not guarantee optimal solutions, studying the drawbacks of the NMOT through the testing and implementation was important to the success of the project. In other words, through the three stages, we improved the tool's algorithms to better serve our internal customers.

1. The first stage of our algorithm is called the ACS-Greedy. The direct execution is to apply ACO iterations, assuming that DHL has sufficient trucks to deliver all the shipments. Then, each route is evaluated "greedily" as to whether it would be cheaper for a common-carrier delivery. The ACS-Greedy methodology is straightforward and interpretable, yet the weakness of this algorithm is also easily observed. First, we found from our implementation that the derived solution may be easily trapped into a local optimum, especially for high-density areas. Apart from the quality issue, we also observed that this version of the NMOT is time consuming. For example, sometimes there are obvious common-carrier opportunities, yet the NMOT still carries those shipments through the iterations and makes greedy make–buy decisions after the last iteration. This wastes computation time, often leading to suboptimal solutions.

2. Observing the weakness of the first version of the NMOT, we researched and evaluated multiple refinement scenarios and designed the updated algorithm called ACSLS. In this version, we involved DP as the local search-optimization model and submerged the make–buy decision to the ant-searching process. The performance of the NMOT is thus improved, as seen in a few examples in the Results Comparison section of the appendix.

3. The third version of the NMOT is called NR-ACSLS. With the ACSLS algorithm, the NMOT had assisted the routine transportation projects with satisfactory results. However, we realized that under

Figure 3. (Color online) Sample Route that Favors Node Removal: The Costs Are (Decomposed Cost, Common-Carrier Cost)

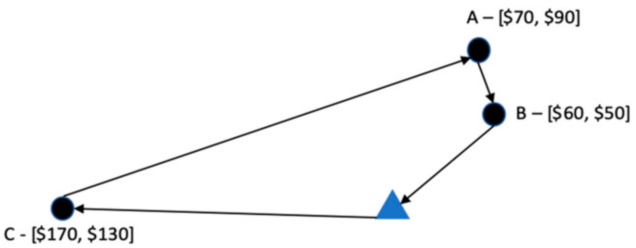

particular circumstances, ACSLS may decide to outsource an entire route, whereas outsourcing one or two nodes from the route would produce a slightly better solution. For instance, by running ACSLS, the route in Figure 3 will be outsourced to a common carrier, because the dedicated route cost is $170 + $70 + $60 = $300, and the corresponding common carrier cost is $130 + $90 + $50 = $270. However, if we consider excluding node C from the route, the network cost decreases to $130 + $70 + $60 = $260. NR-ACSLS will run a fixed number of iterations or until a prespecified convergence criterion is achieved. In each iteration, we first route the shipments by ACSLS, then compare the route costs and the common-carrier costs (e.g., step 1 in Figure 4). Next, if certain shipments/routes are not cost-efficient on dedicated fleet, they will be removed and outsourced to common carriers (e.g., step 2 in Figure 4). The node-removal (NR) ACSLS procedures are continued until no third-party shipments can be turned back to routes. Finally, if there are still some routes that are more expensive than common carriers, the nodes or the routes will be sent to common carriers.

The continuing development of NMOT algorithms is a process of error handling and performance improvement. With the repeated test and evaluation by the team and transportation analysts, we diagnose the drawbacks of the model, refine them, and rule out certain boundary conditions. The tool can even report and send feedback through error-message boxes when the user makes some mistakes in usage. Visualizing the outputs, especially the routes, is another challenge for most of commercial routing software. In the NMOT, we provided two options: If you have access to the Internet, a "Map" macro will cast your routes to Bing Map, so that you will have a timely detailed graph of the routes; otherwise, a macro called our developed R mapping application programming interface can give you a neat, straight-line route visualization.

Challenges of Integrating the NMOT in the Bidding Process

Another important aspect of the NMOT's implementation at DHL is the output processing. During initial testing and implementation stages, we received many enhancement suggestions. For example, some analysts did not agree with the cost-decomposition formulae built in the tool. Whereas corporate business analysts believed that weights should be the main capacity constraint, sector business analysts preferred to focus on the volume constraints. Other analysts suggested to enrich the NMOT's general outputs (i.e., summary statistics, route-sequence labels, transportation-mode comparisons, etc.) so that they can integrate the outputs directly to the downstream reporting software. The aforementioned feedback was collected through numerous interdepartmental meetings, analysts' interviews, and surveys. The NMOT team then studied the standard workflows of the bidding projects, rebuilt the output formats, and updated the cost-decomposition formulae. Furthermore, some analysts complained that new users who are not familiar with the NMOT may have limited ideas of how to avoid errors. To avoid this issue, we wrote and revised the NMOT user manuals. We also scheduled workshops to train new employees on the NMOT. In summary, using the NMOT will result in

Figure 4. (Color online) Algorithm Iterations: Make–Buy Decisions

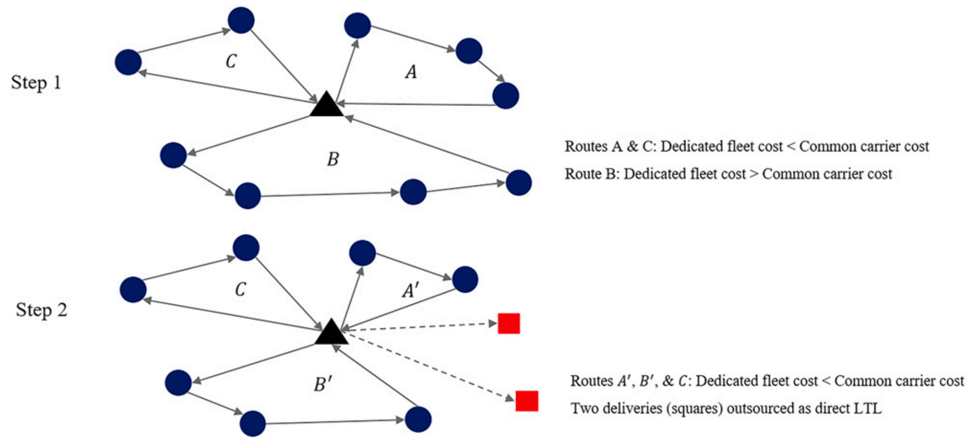

Figure 5. (Color online) Current Process Using the NMOT and Quantified Time Savings

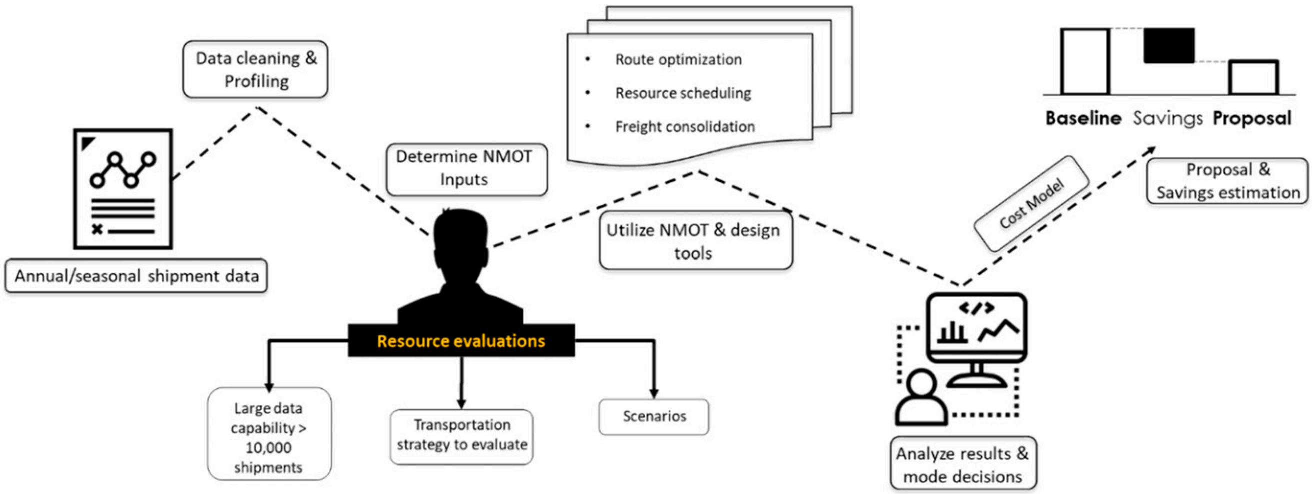

routes with low network costs that are embedded in the DHL bidding process.

Figure 5 shows the process flow for current transportation-mode decision projects at DHL Supply Chain after NMOT implementation. Previously, route optimization, resource scheduling, and freight consolidation were carried out by using multiple commercial software programs. Because these decisions were carried out sequentially, to ensure good quality of the final solution, analysts had to execute numerous iterations for the three aforementioned steps.

The NMOT automates and improves over the previous process by including the three decisions in one run. It evaluates multiple solutions using a novel ACO-based hybrid algorithm and outputs a better solution in a considerably faster time than the previous manual process in Table 1. Average time per project for the previous process was approximately 120 man-hours, within which 88 hours were spent in steps 3, 4, and 5 (route optimization, resource scheduling, and freight consolidation), as shown in Table 1. Step 3, which takes approximately eight hours, in the current process using NMOT, replaces steps 3, 4, and 5 from the previous process (see Table 2). The NMOT built-in algorithm searches across millions of possible solutions iteratively so that analysts are able to spare more time on visualizing the solutions, cost modeling, and preparing bidding proposals. According to our survey of the transportation team at DHL Supply Chain, transportation analysts work on more than 100 projects annually. Currently, the NMOT can be incorporated into the workflows of at least 15% of the 100 projects, which leads to a savings of more than 1,200 hours per year.

Practical Applications

User-Friendly VBA-Based Tool

In this section, we evaluate the quality of our Visual Basics for Applications (VBA)-built optimization tool by comparing the results with four alternatives, including the Gurobi branch-and-cut solver. To illustrate how the NMOT optimizes the solution and solving time, we implement the approaches with ongoing projects from DHL Supply Chain North America. Figure 6 gives us a vision of how our output is visualized by the tool (apart from the model-setting pages and macros). The output sheet will provide the transportation-mode decisions for each shipment—"Dedicated Fleet" or "Common Carriers"—with its decomposed cost calculated by CM. The "RoutingOutput" page shows the detailed routes information, whereas the "CommonCarrier" page lists the shipments to be served by third parties and the corresponding costs. The "Route_Detail" page contains the macros to map and visualize the routes. Finally, the "Route_Summary" page summarizes the solutions.

Table 2. Descriptions of the New Procedures at DHL Supply Chain

Steps	New procedures	Average time (hours)
1	Extract and clean raw shipments information from the database.	8
2	Preprocess the data, sample the subsets, and make initial transportation-mode decisions manually with specified constraints.	16
3	Determine mode split decisions using NMOT.	8
4	Finally, evaluate the solutions and make proposals to the customers.	8

Figure 6. (Color online) Representation of NMOT Model Solution for 3,085 Shipments

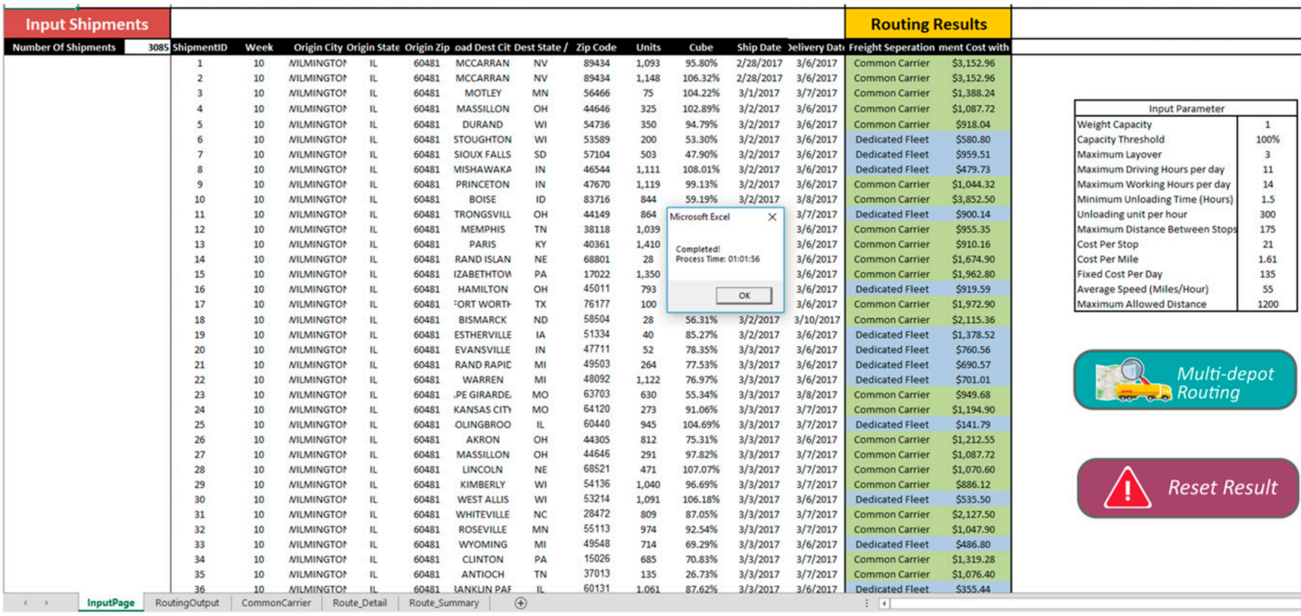

Standard Input Parameters

Each homogeneous truck has 45,000 pounds of weight and 3,000 cubic feet of capacity. The maximum allowed working time is 14 hours. If the time threshold is exceeded, truck drivers will take a 10-hour layover. There are at most three layovers allowed for a round trip. When truck drivers arrive at each customer site, it will take them at least a half-hour to complete the unloading works. The working time will be estimated based on the cargos. The fixed cost for a truck is $200/day, and the mileage charge is $2/mile. It costs $30 for each unloading stop on the route. The average speed for the truck is 55 miles/hour. If the tour distance between the destination and its depot is more than 1,200 miles, this shipment will be directly assigned to a common carrier. The main inputs of the problem and the algorithm parameters are summarized in Table 3.

Table 3. Basic Input Parameters and ACS Parameters

Panel A. Input parameters	
Capacity threshold (%)	90
Maximum number of layovers	3
Maximum working hours per day	14
Minimum unloading time (hours)	0.5
Unloading unit per hour	300
Maximum distance between stops (miles)	120
Cost per stop ($)	30
Cost per mile ($)	2
Fixed cost per day ($)	200
Average speed (miles/hour)	55
Maximum allowed distance (miles)	1,200
Panel B. ACS parameters	
Number of colonies	50
Initial pheromone π_0	1/total distance
Initial probability q_0	0.9
Visibility parameter β	0.9
Initial evaporation ρ_0	0.1

Parameter Settings and Reoptimization

Sensitivity analyses on the input parameters help NMOT practitioners draw meaningful conclusions and identify savings opportunities for consulting projects. For example, when the vehicle fixed costs increase, it can be expected that more customer deliveries are outsourced. However, long-haul transportations are usually not sensitive to the fixed cost—the Cost Per Mile and Cost Per Stop are more significant, according to our implementation. The impact of make–buy decisions on the structure of optimal routes can be observed from Figure 4. The top panel (step 1) assumes that all the shipments are delivered by dedicated fleet, whereas the bottom panel (step 2) introduces common carriers. In this case, better routes (i.e., with fewer turnovers) can be constructed by making outsourcing decisions. This might be due to the exclusion of certain customers, owing to tight time windows, from the dedicated fleet-routing plans.

In addition, users can try different algorithm settings in the NMOT. On one hand, DHL analysts may run the tool for routes within each independent geographical region. The NMOT will store the initial optimal solutions, geographical candidate lists, and arc information and then merge partial graphs with less computational efforts. On the other hand, one of the advantages of ACO-based algorithms is their interpretability, which allows the analysts to tune up the ACS parameters. For example, experienced analysts conduct extensive explorations using local search to break ties between routes with similar costs. Thus, they may increase the number of colonies and decrease the visibility parameters to perform more iterations for dense areas. For sparse areas, too many iterations may be a waste of computational efforts; therefore, increasing the value of initial pheromone π_0 may help the algorithm to converge faster.

Solution Analysis

In this section, we consider six real-life problems that were solved by using the previous DHL process. The actual shipment sizes range from 3,805 (as problem 5 in Table 4) to 147,185. We ran the NMOT on six different subareas of the large data sets, so that other computational methods (for instance, the previous process and the MILP model) can be run and compared on these problems within reasonable time. The six comparison problems in Table 4A range from 64 to 6,714 packages delivered within the planning period. Information about the packages and depots used is shown in Table 4A. Also included is the maximum number of packages for a single depot (Max N_i), because we are running the NMOT for the networks with multiple depots. Each problem covers a local area in the United States, as shown in Table 4A, although a few of the shipments are delivered to distant states.

The data timeline ranges from one week to four weeks, while weekend delivery is only allowed for problem 2, for instances shown in Table 4A. The variants of the NMOT are executed on a desktop of Intel Core-I7 2.8 GHz with 8 G of RAM. The same problems were then solved by Gurobi 8.1 with the mathematical model provided in Mixed Integer Linear Programming in the appendix, which was coded in Python and threaded across eight cores. Gurobi by default solves an MILP using its built-in branch-and-cut algorithms. For our problem, because the linear relaxation of the original formulation is extremely weak at each node, the branch-and-bound tree may not be fully enumerated within a reasonable time, even for a small problem. Instead of relying on the default solver, we generate two sets of valid inequalities that involve solving a separation problem to identify the cuts sequentially when necessary. To add the user-defined valid inequalities to the solver, two selections need to be specified. First, we turn on the *PreCrush=1* option in Gurobi 8.1 and keep the solver's cuts active. This will allow our valid inequalities to be added dynamically to each fractional node. Second, to cut off the fractional solutions that are in maximum violation of the constraints, we optimize the separation problem using the two separation heuristics presented by Bard et al. (2002). Then, we follow the default branching rule of Gurobi until there is no gap between the upper bound and lower bound or the solver fails to reach the optimal within the set time limit. Details of the cuts can be found in Valid Inequalities for VRPTRCC in the appendix. Thus, branch-and-cut methods are used as a benchmark for comparing computational quality. Despite this advantage, branch-and-cut is too inefficient to handle realistic problem sizes. Table 4B lists the experimentation results using the NMOT and human schedulers. "Previous Process" shown in Table 4B and Table A.2 compares the previous process with the current NMOT process. Results Comparison in the appendix is a comprehensive comparison between each of the proposed ACO-based algorithms and the branch-and-cut solver.

As can be observed from Table 4B, the NMOT generated considerable time savings (more than 17 days) and cost reductions for the three small-sized projects (up to 27.3%). The cost reductions were achieved by redesigning the network modes using the NMOT. The analysts saved considerable time with the NMOT on planning-mode decisions compared with the previous process—for example, more than 16 days for problem 5. The users also acknowledged that the standardized outputs increased their efficiency in reporting. The "Decision" column in Table 4B shows us the mode decisions generated by NMOT—*Dedicated Fleet or Common Carriers*. It indicates that NMOT tends to find more dedicated route opportunities than the previous manual process.

Figure 7 depicts the routing and outsourcing decisions for case 2. Six dedicated routes, including one route in the depot's city, cover 18 customers in the corresponding cities. The "triangles" in Figure 7 are common-carrier shipments. By closely looking into this solution, it is not surprising that the customers located in more remote areas tend to be outsourced.

Table 4A. The Complexity of the Six Problems by Comparing the Basic Shipment Information

Problem	No. of depots	*N*. (shipment size)	Max N_i. (max single depot size)	Areas in the United States	Time periods	Weekend delivery
1	1	80	80	Georgia	9/21/2018–9/25/2018	No
2	1	64	64	Louisiana	7/16/2018–7/22/2018	Yes
3	2	320	212	Midwest	1/13/2019–1/25/2019	No
4	3	526	387	Northeast	4/17/2018–4/26/2018	No
5	5	3,805	1,928	Midwest	1/24/2018–2/18/2018	No
6	7	6,714	3,897	West Coast	4/13/201–5/5/2018	No

Figure 7. (Color online) Routes Branching from a Depot

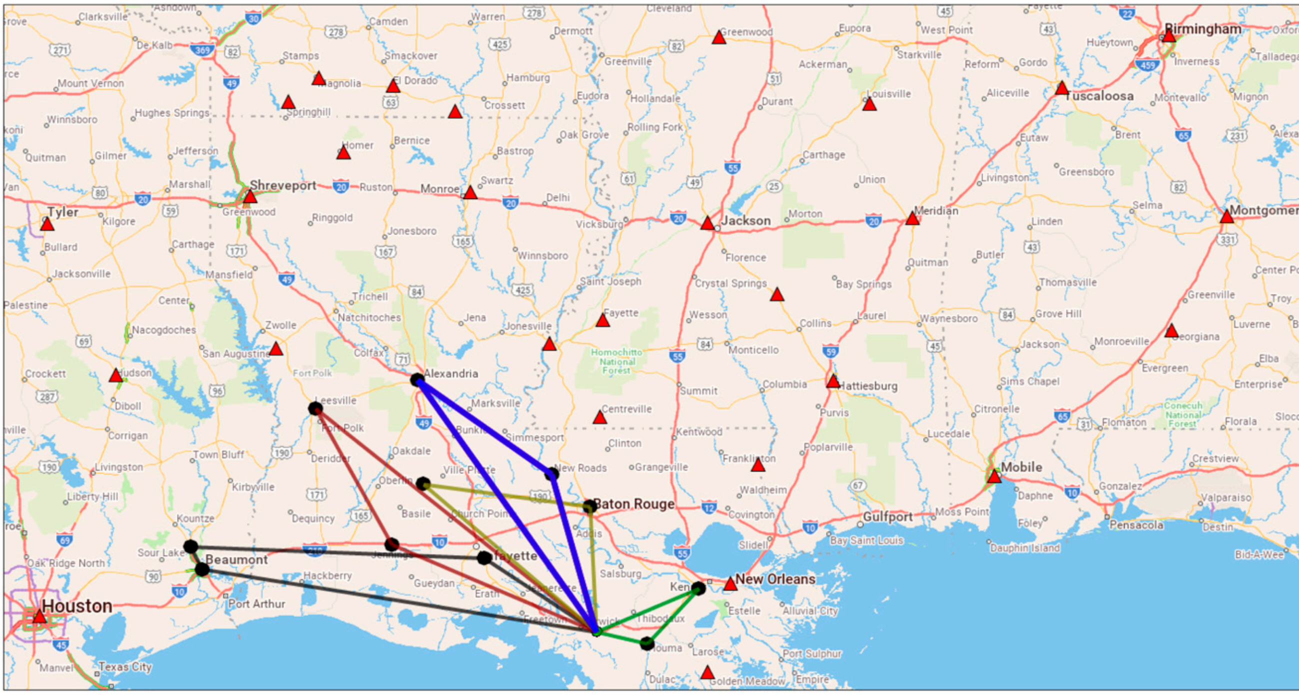

Run the NMOT on Bidding Projects

Besides the aforementioned comparison of solutions from the NMOT with previously available solutions, the NMOT can coarsely be evaluated by considering the costs if only third-party (common carrier) deliveries were performed. In fact, purely outsourcing the shipments to a common carrier may cause a high network cost, and the savings compared with this cost cannot be achieved in practice. As a real-world test, consider shipments from one of DHL's partners, a luxury car maker. Figure 8 and Table 5 document the results from an NMOT implementation. Figure 8 summarizes the savings associated with 18,460 shipments in U.S. dollars. The NMOT system heuristically apportions the costs of routes to the costs of individual package deliveries and compares the apportioned costs with the third-party quote costs. Figure 8 shows the information about routed packages for which it could save money to hire a third party, based on apportioned internal costs ("*Save*"). Also, the instances that internal deliveries would not produce savings are accounted for ("*Not Save*"). As shown in Figure 8, the third-party cost is quite high for full delivery of all packages, indicating that the internal costs are generally low. However, there are a few packages for which hiring a third party would likely save money, as evaluated by the apportionment. The detailed case-study summary is shown in Table 5. In this project, the NMOT allows a setting called *one-way multistop*, in which specific operational constraints are customized such that the vehicles stay at the last stop of the one-way route instead of returning to the depot. *Closed loop*, by its name, indicates our assumptions of the VRPTRCC formulation in the appendix that all trucks return to their depots. The third column in Table 5 denotes the savings percentage compared with the corresponding common-carrier costs. Admittedly, this recorded number of 48% in Table 5 is not the actual savings for this project, because outsourcing the entire less-than-truckload routing business to a common carrier is not a relevant option considered by planners (far too expensive). Also, the improved solutions only indirectly affect operations because they are for planning to support bids for projects and positioning resources for pickup by third parties.

Figure 8. (Color online) Cost Savings Decomposition with 18,460 Shipments in a Test Problem

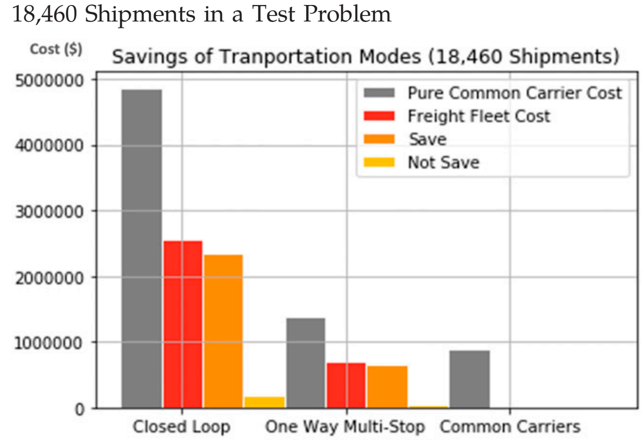

Table 4B. Quantified Savings with NMOT

	Previous process			NMOT				
Problems	Cost ($)	Time (days)	Decision	Cost ($)	Time (seconds)	Decision	Time savings (days)	Cost reductions (%)
1	9,376.58	1.5	73/7	8,751.25	56.34	80/0	>1.4	6.71
2	21,037.56	0.5	14/50	20,228.68	12.50	18/46	>0.4	3.84
5	1,387,809.86	17	1,726/2,079	1,090,118.25	4,135.00	2,597/1,208	>16.7	27.30

The actual routing of vehicles continues to be done by a series of software steps that do not involve the NMOT. According to the estimates of DHL transportation analysts, the actual savings compared with their previous process generally exceeds 15%, but no more than 20%.

Managerial Insights and Business Impacts

By the high-quality solutions provided in Tables 4, A.2, and 5, some meaningful insights can be derived. As mentioned in the Problem Descriptions section, the key step of the previous process is to route the shipments with the standard commercial software. However, the standard software only solves the problem heuristically without incorporating any business constraint—such as layovers, driving-time limits, and common-carrier options. Transportation analysts then have to tune up the routes manually by breaking and merging some routes with parameters in Table 3, panel A. Finally, those inefficient dedicated routes may be replaced by the cheaper common-carrier services. These "greedy" make–buy decisions had been applied for many years, and replacing them with optimized decisions helps bidders understand the realistic costs much more accurately. Indeed, an optimized proposal with efficient network plans will highly increase the chance for DHL to maintain their current customers and earn new businesses.

Before the implementation of the NMOT, we had many rounds of discussions with the project managers and the transportation analysts. Their feedback confirms that the NMOT can quickly help them determine cost-saving possibilities and efficient network mode structures (i.e., dedicated fleet and/or common carriers). Consider the two examples in problems 2 and 5. Figure 7 shows the NMOT solution—for instance, #2. As already discussed, customers that are isolated tend to be outsourced. Moreover, those customers either have full truckload demand or are a single customer with a relatively smaller demand—that is, each customer cannot be easily consolidated with other less-than-truckload customers in one vehicle. Analysts previously would pick up the full-truckload customers and outsource them to a common carrier before checking whether they are beyond a threshold distance from the depot. In fact, under such circumstances, a common carrier will usually provide a more competitive cost for this delivery than running DHL's dedicated fleet. However, how to identify the threshold distance is tricky for the schedulers, whereas the NMOT can easily resolve the issue through its optimization procedures. Analysts may try different fixed truck costs and variable travel costs to accommodate the customer's requirements. As a result, the NMOT automates the previous tedious mode decisions with some useful insights. In problem 5, we observed a saving of 27% compared with the previous process. The shipments are much more densely distributed, so that the previous process generated more impractical results. For instance, the longest route created from the NMOT for this problem includes 15 stops, whereas the previous process based on the commercial software gave the analyst a route with 26 stops in the same area. After considering the layovers and unloading time, the 26-stop route has been broken into three routes, which resulted in 13% higher cost.

Typically, a bidding project may take the analyst a few weeks or even months to prepare the proposal, in which the network mode optimization will occupy at least half of the time. The implementation of the NMOT

Table 5. Estimates of the Savings Associated with 18,460 Shipments in a Test Problem Compared with What It Would Cost to Deliver Everything with a Third Party (One of Our Ways to Evaluate Benefits)

Cost by transportation modes	Cost ($)	$1-\frac{\text{NMOT Cost}}{\text{Common Carriers Cost}}$ (%)	No. of routes	No. of shipments
One-way multistop	699,053.49	49.71	1,212	2,765
Closed loop	2,546,432.86	47.70	5,204	12,900
Third party (common carrier)	679,243.02	0.00	0	2,795
Total cost for DHL using NMOT	3,924,729.36	48.00	6,416	18,460

Note. Total common carrier surcharge (alternative): $7,152,504.30.

provides the power to consider the make–buy decision simultaneously in the optimization procedure. In addition, the NMOT avoids a large burden of manual tune-ups of the route orders and outsourcing decisions, although it is still necessary to check the outputs. Although the NMOT is built on metaheuristic algorithms, it can produce solutions that are believed to be high quality, but it does not guarantee optimality. DHL analysts had observed some undesirable route subsequences from the NMOT solutions, which means that the algorithms were trapped in the local minima for those problems. As we mentioned in the beginning of this paper, stochasticity has been excluded in the model, so that the expertise of transportation analysts is required to diagnose certain routes that may not be able to satisfy the time-window constraints.

The routes with intersections may yield higher costs as a result of additional travel distances. To mitigate this, outsourcing the intersection-point shipments will generate better routes. By excluding the unprofitable customers from the dedicated fleet, NMOT greatly improves work productivity, as analysts do not have to manually identify such routes. NMOT has already resulted in tangible benefits. Considering projects 1, 2, and 5 in Table A.2, the previous process gave the least favorable results. And as the number of shipments increases, the gap between the previous process and the best solution also widens, from 8.09% to 27.30%. Notice that we observed a maximum of 27% savings through our implementations. In fact, the cost reductions produced by the NMOT have helped DHL provide better transportation-mode decisions than before, leading to more competitive bidding proposals. Moreover, the time savings with the NMOT have allowed the analysts to put more efforts into downstream works—that is, visualizing mode decisions, preparing proposal reports, and actual route operations involving traffic, schedule, common-carrier rate, and so forth.

Conclusions

In this paper, we discuss a business problem encountered by DHL Supply Chain North America, which operates a large and complicated transportation network. We describe the previous processes in preparing solutions, analyze the weaknesses, and provide solid optimization models. We describe a tailored high-performance tool that we designed to solve this large-scale, NP-hard problem. Also, we propose multiple solution methods, including the NR-ACSLS method, which outperforms the alternative methods on all the test cases. The implementation of this algorithm in the DHL supply chain is estimated to save in excess of $5 million per year, mainly by permitting the project-bidding teams to recruit delivery business that improves profitability.

The optimization models were researched and developed by authors affiliated with DHL and in collaboration with The Ohio State University. The NMOT is currently in use for planning shipment allocations, vehicle routes, and purchases of third-party logistics. The methods are contributing to millions of added profits for DHL and its customers by reducing personnel, fuel, and third-party costs in North America annually. A clear extension of the work to other branches of DHL is an obvious next step. In the meantime, we are developing and implementing another improved version of the NMOT to deal with heterogeneous fleet sizes, Open-depot VRPTRCC, and delivery-pick-up VRP altogether. Another aspect of the future work is to transfer the NOMT from Excel-VBA to Python with a graphical user interface, which is under test and implementation and of high commercial value to a broader market.

Appendix. Mathematical Models and Results Comparison

Mixed-Integer Linear Programming Model

In this section, the Vehicle Routing Problem with Time Regulations and Common Carriers is formulated and is shown in Table A.1. We use the following notations for the problem parameters and variables. We formulate the problem with an arc-based model. Adulyasak et al. (2015) mentioned that the compact formulations with vehicle indices are suitable for most VRP cases, whereas if the vehicle routes are predefined or there are few possibilities to detour to other clusters, it is more appropriate to use set-partition formulations. Some efforts are also devoted to improving the performance of the model.

Sets.

I: Customers from depot d, where $i \in I = \{I_1, I_2, \ldots, I_d\}$
I_0: Nodes from a particular depot, which includes the customer set I and the depot 0
V: Vehicles from depot d, where $k \in V = R = \{1, 2, \ldots, V_d\}$.

Parameters.

e_i: Earliest service time at node i, which is the earliest acceptable delivery time
a_i: Latest service time at node i, which is defined as scheduled latest delivery time
c_{ij}: Distance cost matrix from nodes i to j (i.e., the distance multiplied by the cost per mile)
λ_i: Rated common carrier cost for shipment i
f: Fixed cost per day by using vehicle k
p: Cost per stop along the route
d^{max}: Maximum (max) allowed distance between internal nodes on the same route
q: Capacity of the homogeneous fleet vehicles
μ_i: Unloading time at customer node i
b: Layover time, which is a fixed value of 10 hours
δ: Max allowed working time per layover
g: The estimated average speed of each vehicle k, which is a fixed value of 55 mph
t^{max}: Max allowed time of duration for one route
w_i: Demand at node i.

Table A.1. Mixed-Integer Linear Programming Formulation of the VRPTRCC

Objective function

$$\text{Min} \sum_{k\in V}\left\{ \text{f}\cdot\theta_k + \sum_{i\in I_0}\sum_{j\in I_0} x_{ijk}\cdot(c_{ij}+p) - \sum_{i\in I_0} x_{i0k}\cdot p\right\} + \sum_{i\in I}\lambda_i\cdot y_i \tag{A.1}$$

Subject to:

Degree constraints

$$\sum_{i\in I_0} x_{ijk} = \sum_{h\in I_0} x_{jhk},\ \forall k\in V,\ \forall j\in I_0 \tag{A.2}$$

$$X_{iik} = 0,\ \forall i\in I_0, \forall k\in V \tag{A.3}$$

$$\sum_{k\in V}\sum_{j\in I_0} x_{ijk} + y_i = 1,\ \forall i\in I \tag{A.4}$$

$$\sum_{j\in I_0} x_{0jk} \le 1,\ \forall k\in V \tag{A.5}$$

Time-window constraints

$$a_j\cdot x_{ijk} \ge s_{ijk} \ge e_j\cdot x_{ijk},\ \forall i\in I_0, \forall j\in I, \forall k\in V \tag{A.6}$$

$$s_{ijk} \le \sum_{h\in I_0} s_{hik} + x_{ijk}\cdot\left(\frac{d_{ij}}{g}+\mu_i\right) + b\cdot l_{jk}, \forall i\in I, \forall j\in I_0, \forall k\in V \tag{A.7}$$

$$s_{ijk} \ge \sum_{h\in I_0} s_{hik} + x_{ijk}\cdot\left(\frac{d_{ij}}{g}+\mu_i\right) + b\cdot l_{jk} - (1-x_{ijk})\cdot M, \forall i\in I, \forall j\in I_0, \forall k\in V \tag{A.8}$$

Layover constraints

$$\sum_{i\in I_0}\sum_{j\in I_0} x_{ijk}\cdot\left(\frac{d_{ij}}{g}+\mu_i\right) \le \delta + 2\cdot\delta\cdot z_k,\ \forall k\in V \tag{A.9}$$

$$\sum_{i\in I_0}\sum_{j\in I_0} x_{ijk}\cdot\left(\frac{d_{ij}}{g}+\mu_i\right) \le 2\cdot\delta + \delta\cdot r_k,\ \forall k\in V \tag{A.10}$$

$$r_k + z_k = \sum_{j\in I_0} l_{jk}, \forall k\in V \tag{A.11}$$

$$l_{jk} \le \sum_{i\in I_0} x_{ijk}, \forall j\in I_0, \forall k\in V \tag{A.12}$$

Truck-capacity constraints

$$q + (w_i - q)\cdot\sum_{k\in V} x_{0ik} \ge u_i \ge w_i,\ \forall i\in I \tag{A.13}$$

$$u_i - u_j + q\cdot\left(\sum_{k\in V} x_{ijk}\right) \le q - w_j, \forall i,j\in I \tag{A.14}$$

Maximum-travel-time constraints

$$t^{max} \ge t_k \ge \sum_{i\in I_0} s_{i0k} - \sum_{j\in I_0} s_{0jk} + \sum_{j\in I_0} x_{0jk}\cdot\left(\frac{d_{0j}}{g}\right), \forall k\in V \tag{A.15}$$

$$\theta_k \ge \tfrac{t_k}{24}, \forall k\in V \tag{A.16}$$

Intranode-distance constraints

$$d_{ij}\cdot x_{ijk} \le d^{max}, \forall i,\ j\in I_0, \forall k\in V \tag{A.17}$$

Symmetry breaking inequalities

$$t_k \ge t_{k+1}, \forall k\in V \tag{A.18}$$

Integrality and nonnegativity constraints

$$x_{ijk},\ y_i,\ l_{ik},\ r_k,\ z_k \in \{0,1\},\ y_0 = 0, \forall i\in I_0, \forall j\in I_0, \forall k\in V \tag{A.19}$$

$$t_k \ge 0,\ \theta_k\in Z_+,\ s_{ijk}\ge 0,\ u_i\ge 0,\ , \forall i\in I, \forall j\in I_0, \forall k\in V \tag{A.20}$$

Variables.

s_{ijk}: An integer variable indicates the service time of node j from i on vehicle k

x_{ijk}: A binary variable indicates whether node i immediately precedes j by vehicle k

y_i: A binary variable denotes whether shipment i is assigned to a common carrier or not

l_{jk}: A binary variable indicates vehicle k taking a layover at or immediately before node j

z_k: A binary variable indicates whether vehicle k needs to have the first layover

r_k: A binary variable indicates whether vehicle k needs to have the second layover

t_k: Total time duration (in hours) of vehicle v on the route

u_i: Product weights accumulated in a vehicle up to node i, which is an integer variable

θ_k: Total number of days elapsed on route k.

The objective function (A.1) states that we want to minimize the overall cost of transporting goods. Specifically, the cost includes fixed vehicle cost on each route ($f\cdot\theta_k$), the summation of the mileage cost, and the stop cost along each arc (i, j) by each vehicle k, $\sum_{i\in I_0}\sum_{i\in I_0} x_{ijk}\cdot(c_{ij}+p)$. Because this term involves one extra stop cost charged at the depot, we need to subtract it $\sum_{i\in I_0} x_{i0k}\cdot p$. Next, we sum the above costs over the dedicated fleet and add in the shipments cost from the common carriers $\sum_{i\in I}\lambda_i\cdot y_i$. The first group of constraints are related to the flow degrees in the network. The constraints (A.2) ensure that the flows enter in node j by vehicle k should be the same as the flows leaving from j by vehicle k. Constraints (A.3) state that the flow in the same nodes is invalid—that is, $x_{iik} = 0$. Constraints (A.4) ensure that each customer i should be shipped by exactly one mode—either on a dedicated fleet or being assigned to a common carrier. Constraints (A.5) specify that each vehicle should be used at most once—that is, $\sum_{j\in I_0} x_{0jk} \le 1$.

The second group of constraints respect the time-window constraints. The service time of each customer j by vehicle k from its potential predecessor i (s_{ijk}) is constrained with the time window of customer j multiplied by x_{ijk}. If i does not precede j, then $x_{ijk} = 0$, and we have the service time equal to 0. The idea is specified in constraints (A.6). Constraints (A.7) are tighter for the service time s_{ijk} because its value should be less than or equal to the service time at node i ($\sum_{h\in I_0} s_{hik}$), plus the time spent on traveling from node i to node j ($\frac{d_{ij}}{g}$) and the unloading time of customer i (μ_i) and the layover time, if node i is the immediate predecessor of j. Notice that the term d_{ij} is the distance between i and j, while g is the estimated average speed of the vehicle. The term $b \cdot l_{jk}$ specifies whether the driver should take a layover before or once he reaches customer j. If $l_{jk} = 1$, then vehicle k will take a 10-hour-length layover immediately and continue the work afterward. Furthermore, if i and j are not connected, the service-time variable s_{ijk} is 0, which is forced by constraints (A.6). Constraints (A.8) state the same concept as (A.7) in a reverse direction. In other words, if i and j are connected, constraints (A.7) and (A.8) will give us the actual service time at customer j as $\sum_{h\in I_0} s_{hik} + x_{ijk} \cdot (\frac{d_{ij}}{g} + \mu_i) + b \cdot l_{jk}$. Otherwise, the term $(1 - x_{ijk})$ is 1, and a big number M is multiplied to relax this inequality. To accelerate the speed of the MILP solver, the value of big M should be tight; here, we use the latest delivery time among the customers (in hours) as M. Next, we use constraints (A.9,A.10) to check whether the vehicle k's driver should take layovers on the trip. If taking the first layover is necessary for driver k, which means $z_k = 1$, the total working time, $\sum_{i\in I_0}\sum_{j\in I_0} x_{ijk} \cdot (\frac{d_{ij}}{g} + \mu_i)$, must be greater than δ. While the second layover is required, the first layover should be applied ($z_k = 1$), and the total working time on vehicle k, $\sum_{i\in I_0}\sum_{j\in I_0} x_{ijk} \cdot (\frac{d_{ij}}{g} + \mu_i)$, must be greater than the double times of δ, which is 28 hours. In that case, then $z_k = 1$ and $r_k = 1$, the constraints (A.9) and the constraints (A.10) are effective. In constraints (A.11), we compute the number of layovers required on vehicle k ($r_k + z$) and let the value equal $\sum_{j\in I_0} l_{jk}$. The statement assigns the layovers to each node j if it is on vehicle k. Constraints (A.12) state that the value of the binary variable l_{jk} should be at most $\sum_{i\in I} x_{ijk}$. (A.12) and (A.13) ensure the layovers are assigned to the earliest node j on vehicle k; otherwise, if j is not on k, then $l_{jk} = 0$. Constraints (A.13) and (A.14) are weight-capacity constraints, which cut off the infeasible routes that exceed the capacity of the trucks. On the one hand, we restrict the upper bound of weight up to node i. If i is directly preceded by the depot, then u_i should be less than or equal to the demand w_i. Also, every demand should be satisfied. Then, we conclude $u_i = w_i$. On the other hand, if node i is not preceded by the depot 0, then u_i is constrained by q, and $\sum_{k\in V} x_{0ik} = 0$. Constraints (A.14) calculate the accumulated weights up to u_i by the summation of its demand w_i and the accumulated weights from its immediate predecessors by vehicle k, $\sum_{j\in I} x_{jik}$. Notice these two inequalities also ensure there must be an arc going out from the depot; otherwise, there will be an infeasible inner loop (subtour). We have proved the combo of subtour elimination constraints (A.13) and (A.14) a valid inequality. Next, we want to know the total time elapsed t_k on the vehicle k (travel + unloading + layovers) to check if the travel time limit is violated and get the days vehicle k has been used. The total time elapsed on vehicle k is evaluated by constraints (A.15), where the value should be greater or equal to the return time of vehicle k to the depot ($\sum_{i\in I_0} s_{i0k}$) minus serving time of this vehicle at the first stop ($\sum_{j\in I_0} s_{0jk}$) plus the time taken from depot to the first stop ($\sum_{j\in I_0} x_{0jk} \cdot (\frac{d_{0j}}{g})$). Constraints (A.15) also set the upper time limit for each vehicle. If the limit is violated, the route is not feasible. t_k is then divided by 24 hours to obtain the days, θ_k, of using vehicle k. Notice this value is an integer, which is the round-up of $\frac{t_k}{24}$. For any arc (i, j) to be feasible, the distance (d_{ij}) should be within the intranode distance threshold d^{max}. Constraints (A.18) break the symmetric solutions caused by the homogeneous fleet size. Finally, the integrality constraints (A.19) and the nonnegativity constrains (A.20) are given.

Valid Inequalities for VRPTRCC

All valid inequalities (VIs) for VRPTW are valid for VRPTRCC. Yet, they may need to be strengthened to capture the features of VRPTRCC, namely, the fact that the demand can be outsourced to common carriers. So, we expect the known inequalities to be tightened by reducing the right-hand side by a quantity that reflects the outsourced customer demand.

In our formulation, constraints (A.7, A.8) and (A.13, A.14) are useful to remove subtours in the solutions. Although we can show that (A.13, A.14) are stronger than (A.7, A.8), they are still weaker than the following two sets of VIs for the LP relaxation at each branching node (see Yaman 2006):

- Subtour elimination constraints: A feasible solution for VRPTRCC consists of a set of disjoint directed cycles, each containing the depot and any number of customer locations. All other cycles in the graph indicate infeasibility and need to be eliminated. Consider a customer set S, and let ϕ_S be the minimum number of trucks needed to fulfill the demand of customers in set S. Then the exponential size of the subtour elimination constraint is

$$\sum_{k\in V}\sum_{i\in S}\sum_{j\in S,\, i\neq j} x_{ijk} \le |S| - \phi_S - \sum_{i\in S} y_i,\ \forall S \subseteq I, \tag{A.21}$$

where S is can be any possible subset of the customers set I. Note that the size of S is exponential, which requires us to pick the most violated ones dynamically and add them to the solver once necessary. Inequality (A.21) specifies that the total flows in set S must be less than or equal to the cardinality of S minus the sum of minimum number of required trucks for set S and the number of outsourced customers.

- Lifted subtour elimination constraints—$\overleftarrow{D_k}$ and $\overrightarrow{D_k}$ inequalities: Fischetti and Toth (1997) designed the two lifted constraints for TSP, and Bard et al. (2002) extended them to VRPTW. Here, we modify the two constraints so that it fits for

VRPTRCC. In general, a set of l nodes $\{i_1, i_2, ..., i_l\} \subseteq I, 3 \le l \le |I|$, satisfy the following two constraints

$$\sum_{j=1}^{l-1} x_{i_j i_{j+1} k} + x_{i_l i_1 k} + 2 \cdot \sum_{j=2}^{l-1} x_{i_j i_1 k} + \sum_{j=3}^{l-1} \sum_{h=2}^{j-1} x_{i_j i_h} \le l - 1 - \sum_{i=1}^{l} y_i, \ \forall k \in V, \tag{A.22}$$

$$\sum_{j=1}^{l-1} x_{i_j i_{j+1} k} + x_{i_l i_1 k} + 2 \cdot \sum_{j=3}^{l} x_{i_1 i_j k} + \sum_{j=4}^{l} \sum_{h=3}^{j-1} x_{i_j i_h} \le l - 1 - \sum_{i=1}^{l} y_i, \ \forall k \in V, \tag{A.23}$$

where the two lifted constraints ensure the double weights, from $i_l \to i_1$ or vice versa, that are significant and there is no cycle in l. Again, we extend the two inequalities here to include the consideration of common carriers.

The number of VIs generated in the branch-and-cut algorithm greatly rely on the order of invoking VIs. The implementation first identifies the violated subtour elimination constraints at each enumeration node, followed by the search for the violated $\overrightarrow{D_k}$ and $\overleftarrow{D_k}$ inequalities. To ensure that each invoked VI cuts off the infeasible fractional solutions, it is necessary to solve the separation problems for the two types of VIs. Because the procedure of optimizing the separation problem is NP-hard (see Eisenbrand 1999), researchers adopt polynomial heuristics to speed up the process. Our work uses the two separation heuristics for subtour elimination constraints introduced by Bard et al. (2002). If no more cuts can be generated that increase the LP lower bounds, then the Gurobi solver continues to branch the next node.

Additional Symmetry-Breaking Scenarios

Finally, we briefly explore some symmetry-breaking rules in addition to the main symmetry-breaking constraint (A.18). The same trips with different indices can hopelessly slow down an MIP solution by requiring the solver to explore many alternative, equivalent solutions—so-called *symmetric solutions* (Francois 2010). To avoid this issue, we rank paths as follows:

Rule 1: If no vehicle k can service a pair of customers i, j without violating the time-window constraints, then two separate vehicles must be used to visit them. In other words, whatever is the node first visited, the service time at the other node can never begin before the latest time window unless another vehicle is used. This condition can be performed during the preprocessing procedure, such that if nodes 1 and 2 cannot be routed together, then we have the case: $x_{12k} = s_{12k} = x_{21k} = s_{21k} = 0, \ \forall k \in V$. Similarly, all the nodes that satisfy this condition are set to zero, and part of the associated constraints can be eliminated.

Rule 2: Still during the preprocessing step, suppose customer i and j are visited by the same vehicle k. Suppose the summation of the earliest start time e_i and the working time between i and j, $(\frac{d_{ij}}{g} + \mu_i)$, exceed the latest delivery time at node j a_j. Then, i cannot be visited before j, such that $x_{ijk} = 0$. For such i and j, $x_{ijk} = s_{ijk} = 0, \ \forall k \in V$. Again, part of the associated constraints can be eliminated.

Results Comparison

The ACS-Greedy without candidate sets and local improvement is included to clarify the practical benefits of those inclusions in terms of solution quality. The computational benefits of candidate set almost offset the costs of local searches (ACSLS), making ACSLS computationally outperform ACS-Greedy. As an alternative procedure, we tailored the tabu search heuristic described in Alcaraz et al. (2019) with candidate set and local search improvements. Greedy make–buy is used here to decide which shipments should be subcontracted. The length of the tabu list is set as $\frac{|I|}{2}$, which is suggested by the authors. The iteration number is chosen to be 500.

Outputs reported the best and the worst solutions produced by different methods. Owing to personnel cost considerations, calculations based on previous processes are only available for cases 1, 2, and 5. For the initial three cases, we ran 10 replications on each variant of the NMOT to get the solutions and their averaged run time. If branch-and-cut failed to give us a solution within 10 hours, we took the best solution generated by our methods and calculated the associated gaps instead. A linear relaxation was applied in cases 4–6. Numerical results with the Gurobi Barrier

Figure A.1. (Color online) Convergence Plot for Variants of NMOT Algorithms

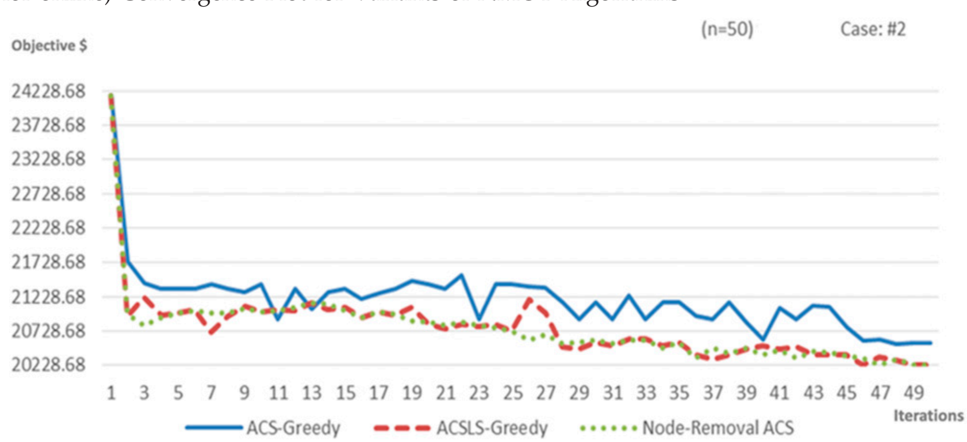

Table A.2. Computational Comparison with Gurobi Solver, Previous Process, and the Proposed ACO-Based Methods

Problem	Previous process	Tabu-LS	Gurobi B&C	ACS-Greedy	ACSLS	NR-ACSLS
Best solutions, $						
1	9,376.58	9,017.34	8,674.80	8,867.68	8,751.25	8,751.25
2	21,037.56	20,921.40	20,228.68	20,598.55	20,228.68	20,228.68
3	—	40,909.15	38,701.79	41,481.87	40,642.12	39,599.67
4	—	281,364.78	—	284,978.64	278,654.44	274,771.36
5	1,387,809.86	1,215,376.00	—	1,183,507.31	1,097,233.46	1,090,118.25
6	—	3,521,623.91	—	3,452,139.04	3,380,794.20	3,346,850.82
Worst solutions, $						
1	9,476.58	9,413.56	8,674.80	9,297.99	8,998.30	8,998.30
2	23,579.10	22,184.82	20,228.68	20,650.74	20,511.35	20,421.06
3	—	42,975.67	38,701.79	41,911.22	41,337.26	40,497.29
4	—	287,398.34	—	287,593.79	284,351.29	280,416.83
5	1,387,809.86	1,235,217.38	—	1,200,766.45	1,120,699.27	1,118,126.46
6	—	3,592,281.95	—	3,501,278.42	3,412,596.20	3,379,506.18
CPU time average (seconds)						
1	1.5 days	48.50	15,265.45	53.90	48.25	56.34
2	6 hours	15.70	8,971.30	5.07	10.65	12.50
3	—	30.60	30,389.24	27.88	22.64	15.22
4	—	428.90	>36,000	367.78	295.60	169.74
5	4.25 weeks	11,258.40	>36,000	7,567.00	6,398.00	4,135.00
6	—	14,935.20	>36,000	12,105.00	8,975.00	5,870.00
Optimal development, %						
1	8.09	3.95	0.00	2.22	0.88	0.88
2	3.99	3.42	0.00	1.87	0.00	0.00
3	—	5.70	0.00	7.18	3.01	2.32
4	—	2.40	—	3.71	1.41	Best
5	27.30	11.49	—	8.57	0.65	Best
6	—	5.22	—	3.15	1.01	Best
Best decisions: dedicated fleet/common carriers						
1	73/7	78/2	80/0	80/0	80/0	80/0
2	14/50	22/42	18/46	20/44	18/46	18/46
3	—	118/202	126/194	114/206	118/202	123/197
4	—	342/184	—	338/188	340/186	346/180
5	1,726/2,079	2,319/1,486	—	2,547/1,258	2,589/1,216	2,597/1,208
6	—	2,941/3,773	—	2,986/3,728	3,014/3,700	3,020/3,694
Best solutions: 3,600-second cutoff, $						
1	—	9,017.34	8,797.38	8,867.68	8,751.25	8,751.25
2	—	20,921.40	20,729.19	20,598.55	20,228.68	20,228.68
3	—	40,909.15	53,503.90	41,481.87	40,642.12	39,987.26
4	—	281,364.78	359,996.86	284,978.64	278,654.44	274,771.36
5	—	1,225,748.49	1,869,950.50	1,251,306.10	1,210,456.89	1,213,849.49
6	—	3,531,275.85	5,912,694.35	3,766,312.47	3,498,295.22	3,475,167.40

Algorithm (not shown) indicate that the linear bounds are weak. Another major output is the make–buy decision. Finally, a 3,600-second cutoff is set to check the quality within time limits. In many situations, it is valuable for an optimization tool to give high-quality solutions within acceptable time.

The results in Table A.2 indicate that the usage of the newer versions of the NMOT will not only help produce solutions closer to optimality, but also speed up the computations. For small cases that can be solved by branch-and-cut, ACS-Greedy is able to obtain solutions whose objective values are within 8% of the optimal, whereas ACSLS improves the solution to values within 5% of optimal. NR-ACSLS gives us the best solution for case 1 with only a 0.88% deviation, whereas both ACSLS and NR-ACSLS achieved the optimal (0.00%) for case 2 and the smallest deviation for case 3 (2.32%). For cases that cannot be solved by branch-and-cut, NR-ACSLS always obtains the best solution, whereas ACSLS gets comparable results. The solution deviation between ACS-Greedy and the best variant reaches the peak at case 5. If we consider the worst-case solutions, the quality of NR-ACSLS is the best in all of the six cases, whereas ACS-Greedy still performed the worst.

In terms of run time, branch-and-cut is obviously the worst for the test problems. In case 2 with 64 shipments, it takes more than 350 times as long as it does with primary ACS-Greedy method, and it failed to solve cases 4–6. With the increase of data sizes, we observe a pattern that the NR-ACSLS is the fastest method, whereas ACS-Greedy is the slowest.

When a 3,600-second cutoff was implemented, the branch-and-cut solver failed to produce any solution for all the instances. It should be emphasized that our objective was not to exceed the previous solution quality, but to develop a user-friendly tool for large-scale VRPTRCC.

Although ACS-Greedy lost in the solution quality, all the improved approaches proved suitable for this problem, because they can easily handle larger-size data sets. For problem 6, for instance, branch-and-cut solver produced a solution of $5,912,694.35, which is 57.0% worse than the output from ACS-Greedy.

Finally, tabu local search generated worse results than ACS-Greedy on four out of the six cases, and this method is less stable than the three ACO-based heuristics. This observation was expected because tabu search performed relatively poorly compared with population-based metaheuristics in previous studies—for example, see Alcaraz et al. (2019) for suggestions for future research.

Figure A.1 illustrates how our algorithms converge with case 2 as an example. Compared with ACS-Greedy, the other two both reached the global optimal. From our observation, Node-Removal ACS and ACSLS-Greedy are favored for a stable and near-optimal solution for most cases.

References

Adulyasak Y, Cordeau JF, Jans R (2015) The production routing problem: A review of formulations and solution algorithms. *Comput. Oper. Res.* 55:141–152.

Alcaraz J, Caballero-Arnaldos L, Vales-Alonso J (2019) Rich vehicle routing problem with last-mile outsourcing decisions. *Transportation Res. Part E* 129:263–286.

Bard JF, Kontoravdis G, Yu G (2002) A branch-and-cut procedure for the vehicle routing problem with time windows. *Transportation Sci.* 36(2):250–269.

Bettinelli A, Ceselli A, Roghini G (2011) A branch-and-cut-and-price algorithm for the multi-depot heterogeneous vehicle routing problem with time windows. *Transportation Res. Part C* 19: 723–740.

Cassettari L, Demartini M, Mosca R, Revetria R, Tonelli F (2018) A multi-stage algorithm for a capacitated vehicle routing problem with time constraints. *Algorithms* 11(5):69.

Cormen TH, Leiserson CE, Rivest RL, Stein C (2009) Dynamic programming. *Introduction to Algorithms*, 3rd ed. (The MIT Press, Cambridge, MA), 359–397.

Dantzig GB, Ramser JH (1959) The truck dispatching problem. *Management Sci.* 6(1):1–14.

Desaulniers G, Errico F, Irnich S, Schneider M (2017) Exact algorithms for electric vehicle-routing problems with time windows. *Oper. Res.* 64(6):1388–1405.

Dondo R, Cerdá J (2007) A cluster-based optimization approach for the multi-depot heterogeneous fleet vehicle routing problem with time windows. *Eur. J. Oper. Res.* 176(3):1478–1507.

Eisenbrand F (1999) On the membership problem for the elementary closure of a polyhedron. *Combinatorica* 19:297–300.

Fischetti M, Toth P (1997) A polyhedral approach to the asymmetric traveling salesman problem. *Management Sci.* 43(11): 1520–1536.

Francois M (2010) Symmetry in integer linear programming. Jünger M, Liebling ThM, Naddef D, Nemhauser GL, Pulleyblank WR, Reinelt G, Rinaldi G, Wolsey LA, eds. *50 Years of Integer Programming 1958-2008*, 6th ed. (Springer, Berlin), 663–669.

Gahm C, Brabander C, Tuma A (2017) Vehicle routing problem with private fleet, multiple common carriers offering volume discounts, and rental options. *Transportation Res. Part E.* 97(C):192–216.

Goel A (2018) Legal aspects in road transport optimization in Europe. *Transportation Res., Part E Logistics Transportation Rev.* 114:144–162.

Kallehauge B, Larsen J, Madsen O (2006) Lagrangian duality applied to the vehicle routing problem with time windows. *Comput. Oper. Res.* 33:1464–1487.

Kok AL, Meyer CM, Kopfer H, Schutten JMJ (2010) A dynamic programming heuristic for the vehicle routing problem with time windows and European community social legislation. *Transportation Sci.* 44(4):442–454.

Moon I, Lee JH, Seong J (2012) Vehicle routing problem with time windows considering overtime and outsourcing vehicles. *Expert Systems Appl.* 39(18):13202–13213.

Pellegrini P, Favaretto D, Moretti E (2007) Multiple ant colony optimization for a rich vehicle routing problem: A case study. Apolloni B, Howlett RJ, Jain L, eds. *Knowledge-Based Intelligent Inform. Engrg. Systems KES 2007*, Lecture Notes in Computer Science, vol. 4693 (Springer, Berlin), 627–634.

Rizzoli AE, Montemanni R, Lucibello E, Gambardella LM (2007) Ant colony optimization for real-world vehicle routing problems. *Swarm Intell.* 1(2):135–151.

Savelsbergh MW, Sol M (1995) The general pickup and delivery problem. *Transportation Sci.* 29(1):17–29.

Vidal T, Crainic TG, Gendreau M, Prins C (2013) Heuristics for multiattribute vehicle routing problems: A survey and synthesis. *Eur. J. Oper. Res.* 231(1):1–21.

Vidal T, Maculan N, Ochi LS, Vaz Penna PH (2016) Large neighborhoods with implicit customer selection for vehicle routing problem with profits. *Transportation Sci.* 50(2):720–734.

Wen M, Krapper E, Larsen J, Stidsen TK (2011) A multi-level variable neighborhood heuristic for a practical vehicle routing and driver scheduling problem. *Networks* 58(4):311–322.

Wu CS, Chu CW, Hsu HL (2017) A heuristic algorithm of vehicle routing problem with time windows and less than truckload carrier selection. *J. Marine Sci. Tech.* 25(2):129–141.

Yaman H (2006) Formulations and valid inequalities for the heterogeneous vehicle routing problem. *Math. Programming* 106:365–390.

Verification Letter

Jon Cox, Director, Solutions Design North America, DHL Supply Chain, 570 Polaris Parkway, Westerville, OH 43202 writes:

"I am writing this verification letter to state that the manuscript titled 'Network Mode Optimization for the DHL Supply Chain' has considered a real problem at DHL Supply Chain. Generally, transportation solutions are interconnected between subcontracted solutions and dedicated fleet solutions where subsets of shipments will be tendered to LTL/TL/intermodal carriers and the remainders of the shipments contracted to a private fleet based on cost and service.

"Previous to this project, the solving process was cumbersome and required the transfer of information between multiple software. Also, we were not guaranteed that this process would derive near optimal solutions. We collaborated with the authors and provided data to capture the real network and to conform to our operational protocols. The method that they generated has been implemented with significant increases in profitability resulting. We estimate (conservatively) that the software has increased profits for the DHL Supply Chain and its customers by $5M annually.

It has done this by adding customers won through bidding, better preparation of fleet resources, and better subcontracting decisions including related pricing."

Yibo Dang is currently a PhD candidate in operations research at Ohio State University and a research assistant at the Global Operations Science and Analytics department of DHL Supply Chain. He received an MSc in statistics from the University of Georgia and a BSc in mathematics from the Beijing Institute of Technology. His research interests include integer programming, metaheuristics, transportation, and machine learning techniques.

Manjeet Singh is the research director of the Global Operations Science and Analytics department of DHL Supply Chain. He received his PhD in industrial and systems engineering from Ohio University. His research interests focus on application of mathematical models, heuristics, algorithms, and machine learning in supply chain. He leads multiple research projects in transportation and warehouse optimization at DHL Supply Chain.

Theodore T. Allen is currently an associate professor at the Department of Integrated Systems Engineering at Ohio State University. He heads the Security and Efficiency Analytics Laboratory, whose members collaborate around data-driven optimal decision making. He is the president of www.factSpread.org, a fellow of ASQ, an officer of INFORMS, and the associate editor of *Computers & Industrial Engineering*.