



INFORMS Journal on Applied Analytics

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

Improved Collaborative Transport Planning at Dutch Logistics Service Provider Fritom

Paul Buijs, Jose Alejandro Lopez Alvarez, Marjolein Veenstra, Kees Jan Roodbergen

To cite this article:

Paul Buijs, Jose Alejandro Lopez Alvarez, Marjolein Veenstra, Kees Jan Roodbergen (2016) Improved Collaborative Transport Planning at Dutch Logistics Service Provider Fritom. INFORMS Journal on Applied Analytics 46(2):119-132. <https://doi.org/10.1287/inte.2015.0838>

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

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

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

Copyright © 2016, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Improved Collaborative Transport Planning at Dutch Logistics Service Provider Fritom

Paul Buijs, Jose Alejandro Lopez Alvarez, Marjolein Veenstra, Kees Jan Roodbergen

Department of Operations, University of Groningen, 9700 AV Groningen, Netherlands
{p.buijs@rug.nl, j.a.lopez.alvarez@rug.nl, marjolein.veenstra@rug.nl, k.j.roodbergen@rug.nl}

We study the collaborative transport planning for two autonomous business units of Fritom, a Dutch logistics service provider. This difficult planning problem does not fit any existing type of vehicle routing problem proposed in the academic literature; therefore, we define a new problem class, the *generalized pickup and delivery problem*. We also propose and evaluate methods to structure and improve Fritom's existing collaborative transport planning process. Using an extensive real-world data set from Fritom, our experiments show that our proposed methods significantly outperform Fritom's collaborative planning approach in total travel distance. We formulate managerial recommendations for Fritom and describe the implementation process for these recommendations. Finally, we present several recommendations for future research.

Keywords: collaborative transport planning; decision support systems; information systems; vehicle routing; technology; simulation; queues; innovation; research & development.

History: This paper has been refereed. Published online in *Articles in Advance* February 9, 2016.

In this paper, we study the collaboration between two business units of Fritom, a Dutch logistics service provider, and propose alternatives to improve its collaborative transport planning. Fritom consists of seven autonomous business units that provide a range of logistics services, such as warehousing, value-added logistics, liquid-food logistics, freight forwarding, and freight transportation. Each business unit has its own profit and loss responsibilities to Fritom, the holding company, which has a combined annual turnover of approximately €150 million. Domestic freight transportation, which is fulfilled by two business units that were, until recently, independently owned freight carriers, represents a large part of Fritom's business. To date, each of these business units houses an autonomous planning and customer relations department, which receives and handles requests for transportation. The network configuration of each business unit consists of a single depot from which pickup and delivery locations are serviced.

Recently, the two business units started to collaborate in providing domestic transportation services. The result was a new joint network configuration. The joint network consists of two depots that serve as a home base for trucks and can be used for transporting

loads from pickup to delivery locations. In this paper, we show how the combination of such individual network configurations led to a new type of vehicle routing problem, which had not been addressed in the literature. The main modeling challenges for this new problem type reside in determining how to best route loads through the joint network (i.e., directly or via one or both depots). We also extensively study the environment at Fritom, provide alternatives for this collaborative planning approach, and evaluate those alternatives using one year of Fritom's data.

The two business units that fulfill all domestic transportation services of the Fritom holding company are Veenstra|Fritom (V|F) and Sanders|Fritom (S|F). Unit V|F operates a depot in the north of the Netherlands; S|F operates one in the south. In principle, because additional resources can be chartered when needed, an unlimited fleet of trucks is available at these depots. Each business unit houses an autonomous planning department, which processes transportation requests received from a large number of customers and from Fritom's other business units. An analysis of transactional data from 2012 revealed that V|F received about 300 requests each day; S|F received about 500. The requests vary substantially in load

size, time restrictions, and pickup and delivery locations; however, the majority of requests specify a small load. The average load size is less than three Euro-pallets. A Euro-pallet is a four-way pallet that measures $1,200 \times 800 \times 144$ millimeters. Most requests are characterized by loose time restrictions for visiting the pickup and delivery locations. For example, the pickup of a load might occur between 8 AM and 5 PM on day 1 and its delivery might be between 9 AM and 6 PM on day 2. A single request can specify multiple pickups and (or) deliveries. Because of the traditional regional contacts of the business units, the pickup locations of most loads are near the depot of the business unit that receives the requests. Delivery locations are geographically dispersed across the Netherlands. The average distance between a pickup and delivery location is about 100 kilometers.

The problem of planning Fritom's domestic transportation entails the construction of routes to pick up and deliver loads, as specified in the requests received. The objective is to minimize the overall transportation costs, for which the total travel distance is a key driver. At V|F and S|F, planners manually construct routes. They are supported by a transportation management system that contains information on all transportation requests and visually represents the constructed routes on a map. Route planning decisions at both V|F and S|F are made in two shifts, a day shift and an evening shift. During the evening shift, one or two planners construct routes for a set of requests that are to be fulfilled the next day. During the day shift, a group of two or three planners construct new routes to fulfill newly arriving requests. Moreover, the planners can dynamically insert loads into existing, ongoing routes. Transportation requests that are not addressed during the day shift are addressed in the subsequent evening shift.

The planners have several options to fulfill a transportation request. Specifically, a load can be transported directly from its pickup to its delivery location (i.e., without an intermediate stop at a depot), or it can be redirected to a depot. The planners redirect most loads to a depot, where they stay overnight. A redirected load usually arrives at the depot in the late afternoon and departs from that depot early the next morning. This is the default fulfillment option for three main reasons. First, it enhances the opportunities

for consolidation and enables the planners to more easily identify those opportunities. Because of the many small loads, consolidation is imperative for sustainable and cost-effective transportation. Second, by redirecting most loads to a depot, dynamically inserting loads from new requests into already planned transportation routes is easier. Consequently, Fritom can accept virtually all requests, and also uphold a high service level. Third, the time restrictions for most requests simply allow for redirection with an overnight stay at the depot (i.e., 97 percent in 2012). In 2012, 29 percent of the requests even required an overnight stay. An example of the latter type of request is a load with a single pickup and multiple geographically dispersed delivery locations, which a single truck cannot fulfill within the specified time restrictions.

Redirecting loads to a depot is the default fulfillment option; however, in 2012, three percent of Fritom's requests did not allow for transportation with an overnight stay at the depot. In addition, the planners seek opportunities to reduce transportation costs by directly transporting loads that would in principle allow for an overnight stay. The planners consider three types of vehicle routes to directly transport loads from pickup to delivery locations. First, a single request can be fulfilled using a dedicated truck. Second, a single truck can fulfill the pickup and delivery for a set of requests. Third, a load can be picked up and delivered by a single truck that also carries loads that are to be transshipped overnight at the depot. Overall, the planners can fulfill requests using many types of vehicle routes. Fritom believes that this flexibility in routing provides a strong competitive advantage because it results in low transportation costs and high service levels.

In recent years, several logistics service providers with multiple depots have adopted a centralized planning approach; all information about transportation requests, fulfillment costs, and the availability of logistics resources—and all route planning decisions—are transferred to a central planning authority (Demir et al. 2014). Many other logistics service providers, such as Fritom, have formed horizontal coalitions, where individual depots retain at least some autonomy in making planning decisions. Fritom's business units V|F and S|F successfully deployed such a collaborative structure to reduce costs and improve the sustainability

of their domestic transportation. During the day shift, planners at each business unit determine, for each arriving request, whether its pickup, or pickup and delivery, can be better executed by the other business unit; they manually identify such opportunities. Most collaborative efforts take place during the evening shift when planners seek opportunities to exchange delivery loads that were redirected to a depot. Specifically, the planners consider which of the loads that are—or will arrive soon—at their depots can be delivered better by the other business unit. These loads are then transferred between the two depots.

The exchange of delivery loads between the two depots is facilitated by a shuttle service, which currently operates according to the following decision rules. The shuttle has a fixed daily capacity of two truckloads with preset departure times in the evening so that the transferred loads can be incorporated into the next day's vehicle routes departing from the other depot. Each business unit decides autonomously which of its delivery loads are to be fulfilled by the other business unit. Planners at V|F (in the north) start considering loads with the most southern delivery location—working their way up until both shuttle trailers are fully loaded. Planners at S|F start considering loads with the most northern delivery location. Whether a load is assigned to a shuttle trailer is determined manually and from a unilateral perspective. Accordingly, loads that can be delivered cost effectively from a planner's own depot are seldom transferred—even when a transfer might contribute to Fritom's overall domestic transportation performance.

Related Literature

The collaborative aspects of Fritom's route-planning problem relate to a fast-growing body of literature on horizontal collaboration among freight carriers (Crujssen et al. 2007). Several solution methods have been proposed to support planners in reallocating transportation requests among collaborating carriers (Berger and Bierwirth 2010, Schwind et al. 2009, Wang et al. 2014). The objective of these methods is to exploit the operational synergies that exist in collaborative carrier networks and to enhance the network's overall routing efficiency. Assuming that individual freight carriers want to retain at least some autonomy in

making routing decisions, the methods allow each collaborating freight carrier to first define a subset of its transportation requests as candidates for reallocation. The carrier can then use combinatorial or Vickrey auctioning techniques to reallocate requests in this subset based on the difference between the costs of self-fulfillment and the fulfillment costs of these requests after reallocation. Prior solution methods differ in their assumptions regarding the routing problem that underlies the collaborative transport planning problem. These assumptions not only define the collaborative scenarios considered, but also influence how fulfillment costs are computed. Therefore, before positioning our alternative collaborative structures for Fritom based on prior work, we provide an overview of relevant vehicle routing literature.

The route-planning problem underlying Fritom's domestic transportation services is related to the domain of pickup and delivery problems. A vast body of operations research (OR) literature describes variants of the pickup and delivery problem and proposes solution approaches. Savelsbergh and Sol (1995) present an overview of literature in that domain. Berbeglia et al. (2007) and Parragh et al. (2006, 2008) each propose a slightly different scheme to classify pickup and delivery problems into two unique problem classes: the pickup and delivery problem (PDP) and the vehicle routing problem with backhauls (VRPB).

In the PDP, routes are constructed for a fleet of vehicles, such that loads are transported from their origin to their destination in a single vehicle route; that is, they do not visit an intermediate depot (Savelsbergh and Sol 1995). The VRPB is a special case of the PDP, where either the origin or the destination of all loads is at the depot (Goetschalckx and Jacobs-Blecha 1989). Accordingly, vehicles depart from the depot carrying all loads for delivery and arrive at the depot carrying all loads that were picked up at the customer locations. The literature includes heuristics for solving single- and multiple-depot variants of the PDP (Ropke and Pisinger 2006a) and VRPB (Nagy and Salhi 2005).

The PDP and VRPB each resemble a part of Fritom's routing problem. Those requests for which the pickup and delivery locations have to be visited in a single-vehicle route fit the PDP. Requests that must be redirected to a depot can be modeled by means of two VRPB requests in different planning periods. In the

first period, a route picks up the load at its origin and carries it to the depot. In a later period, another route starts at that depot and carries the load to its destination. The VRPB is based on the simplifying assumption that each load is available at the depot from which its delivery vehicle departs (Ropke and Pisinger 2006b). In practical settings with multiple depots, however, a planner chooses a depot to which a load is redirected, after which the load must be delivered from that specific depot.

The PDP assumes that a vehicle cannot visit a depot while carrying loads. Accordingly, it cannot consider requests that must be redirected to a depot. In contrast, the VRPB assumes that all loads for delivery are located at a depot and that all loads for pickup are destined for a depot. Requests that cannot be redirected to a depot cannot be considered in the VRPB. Furthermore, requests can have a fulfillment choice; that is, they allow for redirecting a load to the depot, but the requests could also be fulfilled in a single route. At Fritom, for example, the percentage of requests with such a fulfillment choice was 68 percent in 2012. The routes for fulfilling those requests could be constructed by modeling the problem as either a VRPB or a PDP. If modeled as a PDP, the potential benefits from redirecting a load to a depot are lost. Modeling the problem as a VRPB may result in lost opportunities, which would emerge from fulfilling a load's pickup and delivery in a single route. Thus, to find the best routes for requests with a fulfillment choice, the decision on whether to redirect a load to one or multiple depots should be part of the routing model.

A subclass of the PDP allows loads to be transferred between multiple vehicles; we refer to this as the PDP with transfers (PDPT). In the PDPT, transfers can occur only when the requests are fully serviced within a single planning period; this differs from Fritom's domestic transportation problem. Moreover, existing PDPT solution methods can only solve small problem instances. Mitrović-Minić and Laporte (2006) assess the benefit of allowing transfers, and propose a heuristic that solves instances of 50 and 100 requests with at most four transfer points. Cortés et al. (2010) propose a branch-and-cut algorithm to solve to optimality instances of up to six requests, two vehicles, and one transfer point. Rais et al. (2014) present mixed-integer programming formulations for the PDPT with and without time

windows, and solve to optimality instances of five and seven requests, and as many vehicles. In these instances, transshipment is allowed at all nodes in the network. Solution approaches for larger instances exist for the dial-a-ride problem (DARP), which considers the transportation of persons instead of goods. Masson et al. (2013) propose a heuristic for instances up to 193 requests. These instances include a small number of destinations compared to the number of origins. In general, the DARP focuses on the transport service quality, particularly with regard to the service time and maximum number of transfers. Therefore, the time windows are very strict, which considerably reduces the number of feasible transfers.

In summary, existing pickup and delivery literature only addresses aspects of the planning problems typically faced by collaborating freight carriers, such as Fritom. A formal definition addressing the full extent of this class of practical problems has not yet been formulated in the literature, nor have solution approaches been proposed.

Defining the Generalized Pickup and Delivery Problem

In this section, we present a definition for a new class of vehicle routing problems. This definition is inspired by real-world routing problems as faced by collaborating road-freight carriers in the less-than-truckload industry, and particularly by Fritom. We refer to this class as the generalized pickup and delivery problem (GPDP), which is a further generalization of the general pickup and delivery problem proposed in Savelsbergh and Sol (1995). In particular, our definition relaxes the constraints that a load must be transported from its origin to its destination within a single planning period and using one vehicle.

In the GPDP, vehicle routes have to be constructed to satisfy a set of transportation requests subject to an objective, such as minimizing the total travel distance. For each load, a transportation request specifies its size, origin, destination, and time windows for visiting the origin and destination. Each of the multiple depots has a fleet of vehicles to operate the routes. Each vehicle has a given capacity, which may differ among vehicles. A vehicle starts and ends its route at a specified depot and visits a number of locations to pick up and deliver

loads. During a route, the vehicle is allowed to visit one or multiple depots for the transshipment of loads (i.e., to drop off loads for pickup by other vehicles or to pick up loads that were dropped off by other vehicles). As long as the time windows are satisfied, the pickup of a load can be performed in a planning period other than its delivery period. Therefore, a vehicle can end its route by dropping off loads at a depot for final delivery in a future planning period. Similarly, a vehicle route can start at a depot by picking up loads that were dropped off in a previous planning period.

The above problem implies the following fulfillment decisions for each request. A load can be transported by means of either a single route or multiple routes. If a load is to be transported from its origin to its destination through multiple routes, decisions must be made about the number of depots and the specific depots to which the load should be redirected, and the planning period in which parts of the transportation are performed. Moreover, by considering multiple periods, the GPDP takes into account the best depot to which the loads should be redirected. We emphasize that the characteristics of a particular request (e.g., time windows) can result in multiple feasible fulfillment options.

Proposed Alternative Collaborative Configurations

The sustainability and cost effectiveness of Fritom's domestic transportation has considerably improved as a result of the recent collaboration between the two business units. Based on their positive experience with collaboration thus far, Fritom's managers expect additional improvements as they continue to develop their collaborative structure. They seek quantitative evidence supporting this expectation and anticipate innovative OR tools will be required to successfully expand their collaborative efforts. We began our research with this as a business objective.

An optimal solution to the joint routing problem of $V|F$ and $S|F$ would implicitly allocate the fulfillment of requests to depots. As we explain in the *Related Literature* section, this problem does not match any existing pickup and delivery formulation in the literature, and the design of a solution approach entails tackling several new modeling challenges. Finding a good solution

for the routing problem of each business unit separately is less complex. Accordingly, academic literature presents several solution methods for problems related to Fritom's joint routing problem. These methods first allocate requests to depots and then separately solve a standard vehicle routing problem for each depot. They rely on the premise that reallocating requests among collaborative carriers is an effective first step toward finding a good solution to their underlying joint routing problem.

Krajewska and Kopfer (2006) propose a combinatorial auctioning method to reallocate requests, assuming that the fulfillment costs for each combination of potentially reallocated requests is known. Others have proposed route-based request exchange methods. Schwind et al. (2009), for example, propose an auctioning architecture that evaluates the operational gains associated with the reallocation of each combination of requests based on the delivery routes of the individual carriers. Most of the route-based exchange methods use a variant of the PDP to derive fulfillment costs (Berger and Bierwirth 2010, Krajewska et al. 2008, Wang et al. 2014). By considering the PDP as underlying routing problem, only the reallocation of full transportation requests is addressed (i.e., requests fulfilled using a pickup and delivery route without visiting an intermediary depot).

Within the joint routing problem of $V|F$ and $S|F$, we limit our problem scope to the exchange of delivery loads. Hence, we consider all loads that are redirected to one of the depots, but can be transferred to the other depot for delivery to its destination. Appendix A includes the mathematical model for this problem. We adopted this specific problem scope for three reasons. First, the majority of loads are currently redirected to a depot. Therefore, we consider most loads for collaboration in our proposed alternatives. Moreover, the scope we selected does not impede the manual identification of collaboration opportunities for loads that are not redirected to a depot. Second, in Fritom's current organizational structure and procedures, the data of all considered requests are available at the moment the collaborative decisions are made. Third, the scope is compatible with Fritom's current method for collaboration. Overall, our scope increases the probability that Fritom will implement the proposed alternatives; this implementation is an important objective of our research. We acknowledge that solving this problem

will not yield an optimal or near-optimal solution to the GPDP. Rather, our objective is to progress in solving the joint routing problem that will occur when the individual routing problems of V|F and S|F are fully integrated.

Proposed Alternatives

Here, we propose and evaluate two alternative configurations for Fritom's current collaboration. Appendix B includes the pseudo codes for the alternatives.

Configuration 1: Fixed Geographical Division. Configuration 1 divides the Netherlands into two parts, in which V|F and S|F only deliver loads in their respective part of the country. Such a division into dedicated delivery areas is common in transportation network design literature, where each depot in the network is responsible for all transportation activities within a fixed geographical area. Our basis for configuration 1 stems from existing solution approaches in the literature (e.g., Wasner and Zäpfel 2004).

Configuration 1 first allocates each delivery area (a two-digit postcode area) to either V|F or S|F, and then separately solves the resulting routing problems for the deliveries from each depot. We adopted a distance-based policy for the initial delivery-area allocation. That is, we allocated each delivery area to a depot by finding the minimum of the distances between the delivery area and each depot. We then improved the initial solution by quasi randomly reallocating individual, and groups of, delivery areas around the geographical division (i.e., in the middle of the Netherlands). We evaluated the impact of each reallocation by simulating the total travel distance related to delivering loads over a full year of operations. The total travel distance consists of the length of the vehicle routes delivering loads from each depot and the kilometers driven by the shuttle trucks required to transfer loads between those depots. We simulated the lengths of V|F's and S|F's delivery routes by constructing vehicle routes using a modified version of the nearest-neighbor heuristic. Our version of this heuristic includes parameters that specify the capacity of the vehicles, the speed of the vehicles, the stopping time at pickup and delivery locations, and the maximum driving time. A load must be transferred between depots when a particular delivery-area allocation specifies that the delivery of that load is to be fulfilled by a depot other than the one

that picked it up. Similar to the real-world situation at Fritom, a shuttle trailer always makes a round-trip between depots. As we discuss in *Experimental Results*, we selected the delivery-area allocation with the lowest total travel distance as representative in this configuration.

Configuration 2: Variable Geographical Division.

Configuration 2 reallocates individual delivery loads between V|F and S|F, while considering the impact on daily total travel distance. The corresponding geographical division can vary from day to day. Moreover, in contrast to configuration 1, this configuration allows delivery areas to be visited from both depots. Configuration 2 is based on the existing collaborative transport planning methods mentioned earlier. In particular, we follow the logic of Schwind et al. (2009) regarding the automatic route-based evaluation of the request reallocation. Upon implementation of that approach, they noticed it is computationally intractable to evaluate every potential request reallocation into all potential routes. Accordingly, Schwind et al. (2009) developed a simple distance-based measure to generate input for the reallocation decision.

For configuration 2, we consider an alternative means to keep the problem computationally tractable. Specifically, our approach extends the current collaborative structure at Fritom, in which reallocation is considered for batches of delivery requests that make up the load in a full shuttle trailer. Currently, the planners manually identify loads that qualify for delivery by the other business unit—and consider the potential benefits for their business unit alone. We systematize this procedure. Moreover, we allow daily changes in shuttle capacity; the current shuttle capacity between V|F and S|F is restricted to two truckloads. Configuration 2 assigns loads to a shuttle trailer according to a prioritization list for each business unit. These lists are constructed by calculating, for each load, the difference between the distances from each depot to the delivery location. Hence, a load with a destination that is farther from the depot of the business unit that received the request than from the other depot receives high priority for reallocation. Starting with the initial request allocation, as received by each business unit, the volume of reallocated loads is increased iteratively in batches of one full shuttle trailer. For each day, we calculate the total travel distance associated with

delivering loads (i.e., V|F's and S|F's delivery-route lengths and the shuttle route length) for 0 to 10 shuttle trailers—again using our nearest-neighbor heuristic. In *Experimental Results*, we select the number of shuttle trailers and corresponding reallocation of requests with the lowest number of kilometers as representative in this configuration.

Experimental Setting

We developed an experimental setting in which we simulate Fritom's operations to consistently compare the above configurations. We conducted the experiments using one year (2012) of operational data obtained from both business units. We merged the two data sets into a single data set that specifies for each request the load size, delivery location, delivery date, and originating depot (i.e., the depot to which a load was actually redirected). We cleaned the data set by eliminating loads that were transported directly between their origin and destination, and deleted requests with missing information. Moreover, we translated the various measurement units for load size in the original data sets into a single unit. Finally, days at which the number of requests and load volumes for one or both business units are extremely high or low (e.g., because of holidays) are excluded using a range of two times the standard deviation from the mean. As a result, our experimental analysis included 220 of 242 days of operations.

Experimental Results

Tables 1 and 2 display the experimental results. Table 1 lists the total travel distance associated with the two proposed alternative configurations and the initial situation (i.e., where our method constructed vehicle routes for the initial request portfolios of V|F and S|F) over the length of the entire experiment (220 days). We express the total travel distance, which consists of

delivery routes and shuttle routes, in kilometers (km). Table 2 lists the relative total travel-distance reductions. We conducted a paired *t*-test (with a degree of freedom [Df] of 219) to assess whether the differences were significant (at a significance level [Sig.] of 0.01) and calculated Pearson's *r* to determine the effect size.

The experimental results show that each proposed collaborative configuration improves transport efficiency compared to the situation without collaboration. Configuration 1 results in the largest delivery-route length reduction. Nonetheless, this configuration is not appropriate for Fritom's request characteristics because of the inevitable inefficiencies associated with transferring loads between the fixed geographical areas. Many requests concern a load with an origin in the geographical area of one depot and a destination in the geographical area of the other. Therefore, the number and volume of loads requiring a transfer is high. On average, 7.4 shuttle trailers are required to transfer loads from the south to the north and 6.4 trailers from the north to the south. The inefficiency of the shuttle trailer in this configuration resides mainly in the imbalance between the volumes of loads to be transferred. This imbalance results in hauling empty shuttle trailers. Each year, configuration 1 includes 477 empty backhaul shuttle trips. As a consequence, the delivery-route costs savings are partially diminished by shuttle costs. Configuration 2 outperforms configuration 1 by 5.2 percent (i.e., 202,573 km). This configuration imposes a balanced volume of transferred loads, and allows the number of shuttle trailers to vary from day to day to improve the overall routing efficiency. Figure 1 displays the corresponding number of shuttle trailers for each operational day we considered, which varies between two and nine days.

We assessed the performance of Fritom's current collaborative structure by running an experiment with

	Initial request allocation: No collaboration	Configuration 1: Fixed geographical division	Configuration 2: Variable geographical division
Delivery routes (km)	4,837,040	3,087,941	3,238,798
Shuttle routes (km)	—	787,752	434,322
Total travel distance (km)	4837,040	3,875,693	3,673,120

Table 1: Each of our proposed configurations considerably reduces the yearly total travel distances for both delivery and shuttle routes.

Configurations compared	Total km saved	Savings (%)	Average km saved per day	<i>t</i> -value	Df	Sig.	Effect size (<i>r</i>)
Configuration 1 vs. Initial request allocation	961,347	19.9	4,369	43.9	219	$p < 0.01$	0.83
Configuration 2 vs. Initial request allocation	1,163,920	24.0	5,290	54.9	219	$p < 0.01$	0.88
Configuration 2 vs. Configuration 1	202,573	5.2	921	29.9	219	$p < 0.01$	0.71

Table 2: Our proposed alternative collaborative configurations can considerably improve the efficiency of Fritom's domestic transportation. Configuration 2 significantly outperforms configuration 1.

two shuttle trailers using the method presented in configuration 2. This experiment results in a total travel distance of 3,933,541 km, of which 166,320 km stem from operating the shuttle trailers. Compared to this estimation of Fritom's current collaborative configuration, configuration 1 improves the total route length by 1.4 percent ($t(219) = 4$, $p < 0.01$, r , 0.13), and configuration 2 improves the total route length by 6.6 percent ($t(219) = 25$, $p < 0.01$, r , 0.65). We expect the actual savings to be larger, because our approach for assessing the current performance underestimates the total travel distance. It applies our structured approach for reallocating requests; in contrast, the planners reallocate requests ad hoc, manually, and independently of planners in the other depot. Nonetheless, our experimental results provide quantitative empirical evidence supporting the expectation that additional enhancements to Fritom's collaborative structure would yield additional transport efficiency improvements. From that perspective, configuration 2, with a variable geographical division that extends Fritom's current collaborative logic, considerably outperforms configuration 1, with its fixed geographical division—a division often found in transportation network designs in the literature.




Figure 1: In Configuration 2, the number of shuttle trailers required to facilitate the exchange of loads between V | F and S | F varies considerably from day to day.

We also investigated the degree of specialization associated with our proposed alternatives. Figures 2(a) and 2(b) highlight delivery areas if they are visited by only one business unit for more than 98 percent of the operational days. In contrast to the fixed geographical division of configuration 1, a considerable number of delivery areas are visited by trucks from both depots in configuration 2. Some delivery areas are visited by trucks from both depots, which results in less efficient delivery routes compared to configuration 1. Nonetheless, because of the more efficient utilization of shuttle trailer capacity, the overall savings are higher.

Academic literature has stressed the important role that gain sharing has in collaborative transport



Figure 2(a): In configuration 1, the Netherlands is divided into two parts, where each delivery area is visited only by trucks from either V | F (in the north of the Netherlands) or S | F (in the south).




Figure 2(b): In configuration 2, many delivery areas are visited by trucks from both depots (depicted in solid light grey). The highlighted delivery areas in the north and south are almost always visited by trucks from only one depot.

initiatives. Mistrust about the fairness of the gain-sharing policy in place has torn apart many collaborative initiatives (Crujssen et al. 2007). Fritom’s business units V|F and S|F have their own profit and loss responsibilities to the holding company; therefore, understanding the cost and revenue implications of the proposed collaborative configurations for each business unit is important. Table 3 shows that V|F receives a larger share of the overall benefits. A more detailed analysis shows that the day-to-day variation in cost savings for each business unit individually is considerable—particularly for V|F; however, they both gain from the collaboration every day. The degree of specialization for each business unit appears to be a significant driver for the differences in daily cost savings. On average, the density of delivery loads is higher in the north, which leads to a smaller specialization area and more cost savings for V|F after the load exchange.

Currently, no benefit-sharing policy is in place at Fritom, and designing such a policy was beyond the scope of our research. Based on the results shown in Table 3, however, V|F and S|F could equally share the total cost savings. The implementation of a simple

	V F	S F
Total km saved (% of total savings)	638,911 (54.9%)	525,009 (45.1%)
Average km saved per day	2,904	2,386
Standard deviation km saved per day	1,422	575
Minimum daily km saved	280	867
Maximum daily km saved	7,056	4,282

Table 3: When the cost savings associated with configuration 2 are split over the two business units, one can observe that V|F receives a larger share of the total gain. Nevertheless, both V|F and S|F gain from the enhanced collaboration.

gain-sharing policy should be feasible because V|F and S|F are part of the same holding company. In several other studies, where the fulfillment of a request is completely executed by another carrier, the division of revenues was a more complicating factor, because it requires appropriating a value only to the carrier that obtained the request—but has no involvement in its fulfillment (Krajewska et al. 2008). Our collaborative configurations consider only the delivery of loads; therefore, they have no impact on the revenues of each business unit. Regardless of which business unit finally delivers the load, all revenues from the transportation request remain at the business unit that received the transportation request.

Implementation Process

We frequently discussed the (intermediary) results of the above experiments with Fritom’s management. During those meetings, the managers understood and confirmed the limitations associated with the inability of configuration 1 to cope with the large imbalance in shuttle trailer freight flows. More important, the managers corroborated the estimated improvement potential of configuration 2. They indicated that a cost of €1 per km is a valid approximation for the actual transportation costs associated with delivering loads domestically—including fuel, labor, and depreciation costs. Accordingly, the expected yearly cost saving of configuration 2 is approximately €255,788, compared to the company’s current collaborative configuration. Based on our experimental results, Fritom has started implementing configuration 2. As the first step, it gradually increased the number of shuttle trailers from two to four each day. Implementing this step corresponds to an approximate yearly cost saving

of €172,203. Fritom's management decided to start the implementation with a stable number of shuttle trailers because this was easy to implement while a significant part of the total achievable efficiency gains could already be attained. Based on our experiments, operating a fifth daily shuttle trailer would reduce the efficiency gap to €53,848.

Fritom's objective is to fully implement configuration 2 in the short term. This requires operating a dynamic number of shuttle trailers in response to the daily variations in load volumes and imbalance between both business units' sets of requests. Fritom stressed the need for additional IT support to realize this implementation. First, all information regarding the delivery loads with a due date of the next day should be available simultaneously when the load exchange decision is made. This information is needed to determine which and how many loads are to be assigned to the shuttle trailers. In principle, the synchronous availability of this information is embedded in Fritom's current collaborative process. Nevertheless, at the start of our research project, V|F and S|F had different—weakly integrated—transport management systems, which hindered automatic extraction of the required data. Fritom is currently installing a single transport management system in which all business units can record their transaction data, such as transportation requests. Moreover, we created an IT tool to implement the algorithms from configuration 2. This tool structures and strengthens Fritom's current collaborative planning process by considerably increasing the number of loads exchanged. Its implementation can start when the new transport management system is fully implemented. During several multiday visits to the business units, we identified and reported potential challenges to the IT tool's implementation. Based on these visits, we formulated the following recommendations to Fritom's management.

Recommendation 1: At both planning departments, we observed that certain types of loads are not considered for a transfer even when their delivery area can be serviced more efficiently from the other depot. For example, the current policy dictates that loads bigger than six Euro-pallets are not transferred between depots. Moreover, a planner might want to retain full control over the delivery of a load for a particular reason; for example, the planner wants to maintain

customer intimacy, the delivery time window is narrow, or the delivery area requires a specialized driver. To efficiently service these exception loads, many nonexception loads, which could be transferred, are placed in routes with exception loads. This enables the fulfillment of exception loads by means of routes with acceptable efficiency; however, even a few exception loads each day may result in large collateral damage. We observed that the planners sometimes have difficulty in identifying sufficient loads to fill two shuttle truckloads. Our experimental results show that the optimal number of shuttle trailers is five on average when no exceptions are considered. Consequently, we recommend analyzing the various types of exception loads and determining, for each type, the need to consider it as exceptional. Fritom's managers can use our proposed method to calculate the transportation costs for each exception type and weigh the costs against their reasons for considering particular exceptions.

Recommendation 2: The manual, ad hoc, and unilateral transfer decisions in the current collaborative structure leave many operational synergies unexploited. We observed that the planners at both business units generally use their ability to transfer loads to eliminate those loads that they cannot deliver efficiently from their depots. The planners currently identify few loads that can be fulfilled more efficiently by means of a transfer if those loads can be fulfilled at an acceptable efficiency level from their own depots. We note that, up to a certain number of shuttle trailers, for each load that is transferred to the other depot, a "better" load returns. Without OR support, identifying which and how many loads should be transferred is difficult. The objective of the IT tool we developed to implement configuration 2 is to automatically identify those loads. We recommend using a tool that can systematize the transfer decisions at both business units.

Recommendation 3: We considered the impact of implementing configuration 2 on the depot operations (i.e., where the loads from incoming shuttle trailers are transshipped to delivery vehicles), although this was not part of our initial setup. The number of transferred loads usually increases with the implementation of configuration 2, which leads to increased material handling at the depot. Fritom's managers have said that they expect the transportation cost reductions to substantially exceed the expected additional material

handling costs. Nonetheless, any negative impact of the new collaborative transportation setting on the depots' operational performance should be mitigated. We recommend a minor redesign of the depots' internal layout and workforce planning to cope with handling loads from the shuttle trailers. The new layout should enable material handlers to quickly identify loads that need to be placed in the shuttle. Implementing configuration 2 will also require a shift in the number of available material handlers over time.

We finalized our research by presenting the above recommendations to Fritom's managers and by discussing our experimental and observational analyses. The managers validated our findings and agreed to implement the recommendations accordingly.

Future Outlook

We conclude this paper with an overview of promising areas for future academic research and innovations. We consider the academic state of the art, Fritom's current practices, and the contribution and limitations of the solution techniques we propose in this paper.

Figure 3 plots the efficiency improvements in domestic transportation that Fritom can realize through various collaborative configurations. We proposed and compared two alternative collaborative structures that focus on the delivery of loads that are redirected to a depot. Our experiments showed that configuration 2 (i.e., variable geographical division) significantly outperforms configuration 1 (i.e., fixed geographical division).

Our study provides several opportunities for future research and innovation. First, the proposed solutions for Fritom's collaborative transport planning can be refined to provide more transport efficiency improvements. Currently, our best alternative configuration identifies loads that qualify for a transfer by using a simple distance-based prioritization and assigns loads to the shuttle trailers in batches of one full trailer. Our experimental results indicate that Fritom can achieve substantial cost savings by focusing on the assignment of delivery loads to shuttle trailers, while taking the delivery routes from each depot into account—even when it uses simple solution techniques. This highlights the opportunities for future research to develop meta heuristics for the problem (Appendix A). The development of such heuristics would also enable Fritom to more clearly understand the loss of simplicity that results from using methods such as those we present in this paper; that is, a simple method will typically generate less improvement than a sophisticated method. Second, the search for alternative collaborative configurations can be expanded by considering aspects that were beyond the scope of our study. For example, Fritom could consider the assignment of pickups to depots. We note that such an expansion would dramatically increase the problem complexity because of the need to consider multiple planning periods.

Although collaborative transport planning can considerably improve transport efficiency, any collaborative approach can be outperformed by centralized joint route planning (Berger and Bierwirth 2010, Wang et al. 2014). In this paper, we defined the GPDP as a new




Figure 3: Fritom has realized large efficiency improvements in domestic transportation through the collaboration between V|F and S|F. It can realize additional improvements by restructuring and strengthening this collaboration.

class of pickup and delivery problems that matches this joint planning route planning problem. Promising areas for future research include many aspects of the GPDP that the literature does not yet address. One example is the development of solution methods to construct routes that consider all GPDP fulfillment choices.

Appendix A

We model the problem as follows. Let $G = (V, A)$ denote a directed graph, where $V = V^{DEP} \cup V^{CST}$ is the set of nodes and A is the set of arcs. The set of depots is represented by $V^{DEP} = \{v_1, \dots, v_d\}$, and V^{CST} represents the set of customers, where $v_i \in V^{CST}$ correspond to customer i , $i = d+1, \dots, d+n$. Each customer i has a demand equal to $q_i > 0$, which is located at a given depot. We define α_{io} as a binary constant equal to one if and only if the load of customer i , $v_i \in V^{CST}$ is located at depot o , $v_o \in V^{DEP}$. Let m be the number of vehicles available at each depot, where each vehicle has a given capacity Q . Let c_{ij} and t_{ij} denote, respectively, the travel cost and the travel time from node $v_i \in V$ to node $v_j \in V$. The service time of a node is equal to τ_i , $\forall v_i \in V$, where $\tau_i = 0$, $\forall v_o \in V^{DEP}$. The route duration, which is the sum of the total travel time and service time, may not exceed T . Let p be the number of shuttle trailers, where the maximum total load allowed to enter a shuttle trailer at a depot o , $v_o \in V^{DEP}$ is given by $\tilde{Q}$. Each shuttle trailer visits all depots, and the corresponding routing costs are given by $\tilde{c}$.

We define the following binary variables. The variable x_{ijk} is equal to one if and only if vehicle $k \in \{1, \dots, m\}$ coming from depot o , $v_o \in V^{DEP}$ visits node $v_j \in V$ immediately after node $v_i \in V$. Let y_{io} be equal to one if and only if customer i , $v_i \in V^{CST}$, is assigned to depot o , $v_o \in V^{DEP}$. Variable r_{il} is equal to one if and only if customer i , $v_i \in V^{CST}$, is assigned to shuttle trailer l , $l \in \{1, \dots, p\}$. Variable $\tilde{r}_l$ is equal to one if and only if shuttle trailer l , $l \in \{1, \dots, p\}$, is used in the solution. Then, we can formulate the problem as follows:

$$\min \sum_{v_i \in V} \sum_{v_j \in V} \sum_{k=1}^m \sum_{v_o \in V^{DEP}} c_{ij} x_{ijk} + \sum_{l=1}^p \tilde{c} \tilde{r}_l \quad (1)$$

$$\text{subject to } \sum_{v_o \in V^{DEP}} y_{io} = 1, \quad v_i \in V^{CST}, \quad (2)$$

$$\sum_{v_j \in V} \sum_{k=1}^m x_{ijk} - y_{io} = 0, \quad v_i \in V^{CST}; \quad v_o \in V^{DEP}, \quad (3)$$

$$\sum_{v_j \in V} x_{ojk} \leq 1, \quad v_o \in V^{DEP}; \quad k = 1, \dots, m, \quad (4)$$

$$\sum_{v_j \in V} x_{ijk} = 0, \quad v_o \in V^{DEP}; \quad v_o \neq v_i; \quad k = 1, \dots, m, \quad (5)$$

$$\sum_{v_j \in V} x_{jiko} - \sum_{v_j \in V} x_{ijk} = 0, \quad v_i \in V; \quad v_o \in V^{DEP}; \quad k = 1, \dots, m, \quad (6)$$

$$\sum_{v_i \in V} \sum_{v_j \in V} q_i x_{ijk} \leq Q, \quad v_o \in V^{DEP}; \quad k = 1, \dots, m, \quad (7)$$

$$\sum_{v_i \in V} \sum_{v_j \in V} (t_{ij} + \tau_i) x_{ijk} \leq T, \quad v_o \in V^{DEP}; \quad k = 1, \dots, m, \quad (8)$$

$$\sum_{v_i \in S} \sum_{v_j \in S} x_{ijk} \leq |S| - 1 \quad S \in V^{CST}; \quad |S| \geq 2; \quad v_o \in V^{DEP}; \quad k = 1, \dots, m, \quad (9)$$

$$\sum_{v_o \in V^{DEP}} \alpha_{io} y_{io} + \sum_{l=1}^p r_{il} = 1 \quad v_i \in V^{CST}, \quad (10)$$

$$\sum_{i \in V^{CST}} \alpha_{io} q_i r_{il} \leq \tilde{Q} \tilde{r}_l, \quad v_o \in V^{DEP}; \quad l = 1, \dots, p, \quad (11)$$

$$x_{ijk} \in \{0, 1\}, \quad v_i \in V; \quad v_j \in V; \quad v_o \in V^{DEP}; \quad k = 1, \dots, m, \quad (12)$$

$$y_{io} \in \{0, 1\}, \quad v_i \in V; \quad v_o \in V^{DEP}, \quad (13)$$

$$r_{il} \in \{0, 1\}, \quad v_i \in V^{CST}; \quad l = 1, \dots, p, \quad (14)$$

$$\tilde{r}_l \in \{0, 1\}, \quad l = 1, \dots, p. \quad (15)$$

The objective (1) minimizes the total routing cost of the vehicles and the shuttle trailers. Constraints (2)–(9) come from the formulation of Vidal et al. (2012) and model the assignment of customers to depots and the routing of vehicles from the depots. Because of constraints (2) and (3), customers are assigned to exactly one depot and visited by a vehicle assigned to that depot. Constraints (4) guarantee that a vehicle is used at most once. Constraints (5) ensure that a vehicle never visits a depot to which it is not assigned. Constraints (6) are the flow conservation constraints. Constraints (7) and (8) guarantee that, respectively, the capacity of the vehicles and the maximum route duration are never exceeded. Constraints (9) are the subtour elimination constraints. If and only if an item is served by a vehicle from a depot at which the load is not located, it is assigned to exactly one shuttle trailer because of constraints (10). An item can only be assigned to a used shuttle trailer, and the maximum load allowed to enter a shuttle trailer at a depot is never exceeded because of constraints (11).

Appendix B

Configuration 1

- 1 For each postcode area A in the Netherlands do
- 2 Compute the distance to each depot
- 3 Assign area A to the closest depot
- 4 Next
- 5 For each Depot o do

```

6   Compute the number of shuttle trailers
    needed to transfer loads between the depots
7   While  $|Ro| > 0$  do
8       Construct route using the
        Nearest-neighbor heuristic
9   Loop
10  Next
11  Record total routing cost (in km)
12  Create set M containing the central postcode
    areas of the Netherlands
13  For each area A in M
14      Randomly switch assigned depot
15  For each depot o do
16      Compute the number of shuttle trailers
        needed to transfer loads
        between the depots
17      While  $|Ro| > 0$  do
18          Construct route using the
            Nearest-neighbor heuristic
19      Loop
20  Next
21  Record total routing cost (in km) for iteration A
22  Next
23  Return the solution with the lowest routing cost.
    
```

Configuration 2

```

1  For each Depot o do
2      Compute score Z for all loads in Ro
        (i.e., the difference between the distances
        from the delivery location
        of load  $k$  to each depot)
3  next
4  For  $p = 0$  to 10 shuttle trailers do
5      For each depot o do
6          Fill shuttle trailer  $p$  with loads
            with the highest Z scores
7          While  $|Ro| > 0$  do
8              Construct route using the
                Nearest-neighbor heuristic
9          Loop
10         Next
11         Record the total routing cost
            (in km) for iteration  $p$ 
12         next
13     Return the solution with the
        lowest routing cost.
    
```

Acknowledgments

This research was supported by Dinalog, the Dutch Institute for Advanced Logistics.

References

- Berbeglia G, Cordeau JF, Gribkovskaia I, Laporte G (2007) Static pickup and delivery problems: A classification scheme and survey. *Top* 15(1):1–31.
- Berger S, Bierwirth C (2010) Solutions to the request reassignment problem in collaborative carrier networks. *Transport Res. E* 46(5):627–638.
- Cortés C, Matamala M, Contardo C (2010) The pickup and delivery problem with transfers: Formulation and a branch-and-cut solution method. *Eur. J. Oper. Res.* 200(3):711–724.
- Crujssens F, Dulleart W, Fleuren H (2007) Horizontal cooperation in transport and logistics: A literature review. *Transportation J.* 46(3):22–39.
- Demir E, Van Woensel T, de Kok T (2014) Multidepot distribution planning at logistics service provider Nabuurs B.V. *Interfaces* 44(6):591–604.
- Goetschalckx M, Jacobs-Blecha C (1989) The vehicle routing problem with backhauls. *Eur. J. Oper. Res.* 42(1):39–51.
- Krajewska M, Kopfer H (2006) Collaborating freight forwarding enterprises—Request allocation and profit sharing. *OR Spectrum* 28(3):301–317.
- Krajewska M, Kopfer H, Laporte G, Ropke S, Zoccoro G (2008) Horizontal cooperation of freight carriers: Request allocation and profit sharing. *J. Oper. Res. Soc.* 59(11):1483–1491.
- Masson R, Lehuédé F, Péton O (2013) An adaptive large neighborhood search for the pickup and delivery problem with transfers. *Transportation Sci.* 47(3):344–355.
- Mitrović-Minić S, Laporte G (2006) The pickup and delivery problem with time windows and transshipment. *INFOR* 44(3):217–228.
- Nagy G, Salhi S (2005) Heuristic algorithms for single and multiple depot vehicle routing problems with pickups and deliveries. *Eur. J. Oper. Res.* 162(1):126–141.
- Parragh S, Doerner K, Hartl R (2006) A survey on pickup and delivery models. Part II: Transportation between pickup and delivery locations. Working paper, University of Vienna, Vienna.
- Parragh S, Doerner K, Hartl R (2008) A survey on pickup and delivery problems. Part I: Transportation between customers and depot. *J. Betriebswirtschaft* 58(1):21–51.
- Rais A, Alvelos F, Carvalho MS (2014) New mixed integer-programming model for the pickup-and-delivery problem with transshipment. *Eur. J. Oper. Res.* 235(3):530–539.
- Ropke S, Pisinger D (2006a) An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Sci.* 40(4):455–472.
- Ropke S, Pisinger D (2006b) A unified heuristic for a large class of vehicle routing problems with backhauls. *Eur. J. Oper. Res.* 171(3):750–775.
- Savelsbergh MWP, Sol M (1995) The general pickup and delivery problem. *Transportation Sci.* 29(1):17–29.
- Schwind M, Gujo O, Vykoukal J (2009) A combinatorial intra-enterprise exchange for logistics services. *Inform. Systems E-Bus. Management* 7(4):447–471.
- Vidal T, Crainic TG, Gendreau M, Lahrichi N, Rei W (2012) Supplemental to: A hybrid genetic algorithm for multidepot and periodic vehicle routing problems. *Oper. Res.* 60(3):611–624.
- Wang X, Kopfer H, Gendreau M (2014) Operational transportation planning of freight forwarding companies in horizontal coalitions. *Eur. J. Oper. Res.* 237(3):1133–1141.
- Wasner M, Zäpfel G (2004) An integrated multi-depot hub-location vehicle routing model for network planning of parcel service. *Internat. J. Production Econom.* 90(3):403–419.

Verification Letter

Joris Wijnja, Operations Manager Veenstra, Fritom, P.O. Box 6, 8620 AA Heeg, Netherlands, writes:

“In this letter, I verify the study reported in the article ‘Improved collaborative transport planning at Dutch logistics service provider Fritom’ carried out in collaboration with a team of researchers (Paul Buijs, Jose Lopez, Marjolein Veenstra and Kees Jan Roodbergen) from the Department of Operations at the University of Groningen. The study is part of a long-term ongoing research effort towards improved collaborative transportation within the Fritom group, which we started in November 2012. The article covers a project that was conducted between January 2013 and March 2014, during which operations managers and planners from several Fritom business units actively collaborated with the team of researchers. The aim of the project was to improve the overall efficiency of our domestic freight transportation services. Specifically, the project was focused at structuring the decisions planners make when determining which delivery loads are exchanged between the two depots that are used in domestic transportation.

“The article accurately describes the current collaborative transport planning processes and network configurations of the two business units involved in our domestic freight transportation. Each of the two business units provided the researchers with full access to their operational data. The researchers transparently communicated how they cleaned and used the data for their analyses as presented in the paper. Moreover, there were several formal and informal meetings during which (intermediary) results were presented and discussed. Accordingly, I confidently verify the reported estimated cost savings associated with implementing the proposed alternatives—over €260.000 per annum for the best alternative.

“I strongly believe in the need of Operations Research (OR) tools to break through our currently unilateral perspective on collaborative transport planning. Despite our positive experience with the current collaboration, I expect that little additional benefits can be accrued from our collaborative endeavors without the support of OR tools. The tool developed and tested within this study confirms this expectation. The tool structures the load exchange decisions that are

currently made manually and ad-hoc. Moreover, the tool allows the planners at each business unit to gain insights in the overall performance of our domestic transportation. With those insights, we can further exploit the operational synergies that exist between the two business units and improve the service levels and cost-effectiveness of our operations. Lastly, I verify that we intend to implement and fully use the tool in the near future. The recommendations put forward in the paper are important first steps in that regard.”

Paul Buijs is a post-doctoral researcher at the University of Groningen, Netherlands. He holds a PhD degree in economics and business from that same university. His research interests are in sustainable logistics, including topics such as horizontal and vertical collaboration in distribution networks and the design and planning of alternative fuel transportation systems. Most of his research projects are designed and executed in close collaboration with industry and public partners.

Jose Alejandro Lopez Alvarez is a PhD candidate in the Department of Operations at the University of Groningen, Netherlands. He has a BSc degree in industrial engineering from Universidad de los Andes (Colombia) and a MSc degree in supply chain management from the University of Groningen. His main research interest relates to the development of quantitative models to support the design of efficient supply chain networks in various contexts, such as for alternative fuels.

Marjolein Veenstra is a PhD candidate in the Department of Operations at the University of Groningen, Netherlands. She holds a MSc in econometrics, operations research, and actuarial studies. Her research mainly focuses on the development of exact and heuristic methods in the area of vehicle routing problems.

Kees Jan Roodbergen is a professor of quantitative logistics at the University of Groningen, Netherlands. His research focuses on the application of OR techniques in transport, warehousing, and e-commerce logistics. He has been a visiting professor at Virginia Tech and Université Laval, and serves on the Scientific Advisory Council of the World Food Logistics Organization. He received his PhD degree from Rotterdam School of Management.