informs
http://pubsonline.informs.org/journal/inte/

INTERFACES
Articles in Advance, pp. 1–16
ISSN 0092-2102 (print), ISSN 1526-551X (online)

Tri-Vizor Uses an Efficient Algorithm to Identify Collaborative Shipping Opportunities

Stefan Creemers,[a] **Gert Woumans,**[a] **Robert Boute,**[b] **Jeroen Beliën**[c]

[a] IÉSEG School of Management, 59000 Lille, France; and Research Center for Operations Management, KU Leuven, 3000 Leuven, Belgium; [b] Technology and Operations Management Area, Vlerick Business School, 9000 Ghent, Belgium; and Research Center for Operations Management, KU Leuven, 3000 Leuven, Belgium; [c] Research Center for Information Management, Modelling and Simulation, KU Leuven, 1000 Brussels, Belgium

Contact: s.creemers@ieseg.fr (SC); woumansgert@telenet.be (GW); robert.boute@kuleuven.be (RB); jeroen.belien@kuleuven.be (JB)

Received: November 17, 2015
Revised: May 19, 2016; July 26, 2016; September 14, 2016; September 28, 2016
Accepted: September 30, 2016
Published Online in Articles in Advance:

https://doi.org/10.1287/inte.2016.0878

Abstract. Collaborative shipping programs, whereby companies bundle their transport loads, are a growing trend in logistics. By bundling shipments with other partners, available space in truck hauls for one company can be used to transport shipments for other companies. The benefits are reduced logistics costs and a lower carbon footprint. Although the advantages of collaborative shipping are clear, finding suitable collaboration partners is a major impediment. In this article we present a tool that enables the quick identification of potential partners based on their geographical compatibility, even when the database of shipment lanes is very large. The tool allows the detection of bundling, back-hauling, and round-trip opportunities, as well as "collect-and-or-drop" opportunities in which shipments are collected and (or) dropped off en route. Tri-Vizor, a facilitator and orchestrator of horizontal logistics partnerships, is currently using this tool. Any company that is looking for collaborative shipping partners would also find it valuable. For Tri-Vizor, whose database has grown to over 130,000 shipment lanes, this tool has become an indispensable asset in detecting collaborative shipping opportunities.

History: This paper was refereed.

Keywords: transport: general • collaborative shipping • optimization

A growing trend in improving logistics efficiency is to set up logistics partnerships with other companies. One can distinguish between *vertical* and *horizontal* supply chain collaborations. Vertical collaborations are established between suppliers and buyers. An example of vertical collaboration is sharing information on customer orders upstream of the supply chain to reduce demand uncertainty for the suppliers. Horizontal collaborations are established between companies that operate at the same level in different supply chains (e.g., between suppliers or between buyers). One example of horizontal collaboration is sharing transportation capacity when moving freight, an option that benefits the environment and yields substantial network efficiencies (Saenz 2012). Even coopetitors (i.e., competitors who occasionally cooperate) sometimes set up a horizontal cooperation (Leitner et al. 2011). Horizontal partnerships in logistics have the potential to generate substantial gains for the partners by leveraging the overlaps in transport networks (Cruijssen et al. 2007a, Leitner et al. 2011). Vertical collaborations have been successfully established for many years; however, horizontal collaboration initiatives are more recent and are expected to become more prevalent in the near future.

Bundling freight is not a new concept; it is essentially what logistics providers do. When companies outsource their logistics to a logistics service provider, the provider can combine freight loads of its customers if shipment timings are identical and if there is a geographical match. Collaborative shipping is different; opportunities must be detected prior to shipment, and if desirable, plans can be changed and shipments delayed or moved forward to benefit from the joint transport. The consolidation can be based on geography and (or) time (i.e., a shipment might be rescheduled if it creates synergies). Flexibility on the part of each partner allows the exploitation of more opportunities for bundling and permits the creation of better and cheaper distribution plans (Vanovermeire and Sörensen 2014). Boute et al. (2011) report on the collaboration of two pharmaceutical companies, Baxter and UCB. Flexible planning by each partner enabled synergies; Baxter was willing to postpone shipping some of its orders, thereby freeing up space for UCB who was shipping lower freight volumes and with a lower frequency. This contrasts with traditional freight groupage, which is mainly reactive: in groupage shipping, the logistics provider decides to bundle less-than-container loads (LTL) in the execution phase, rather than in the planning phase, and the consolidation is only geographical (i.e., the timing of the shipments cannot change).

In a recent World Economic Forum report, collaborative shipping was identified as one of the processes in driving shared value (World Economic Forum 2015). The European Commission reported that 27.3 percent of national road freight kilometers (km) were empty hauls in 2010; in addition, vehicles carrying loads are typically loaded at only 57 percent of their maximum gross weight (Doherty and Hoyle 2009). Increasing the efficiency of the European road freight is, Therefore, one of the main goals of the European Commission. Horizontal collaborations by bundling transport help to increase the utilization rate in transport, thus reducing the number of transports. Companies benefit as their transportation costs go down, and when the number of truck hauls decreases, some harmful external effects directly related to road freight transport, such as greenhouse gas (GHG) emissions, pollution, and congestion, are also mitigated. Freight consolidation across companies can also lead to increased scale effects, facilitating a modal shift. Multimodal transport requires a certain volume to be economically viable; for small to medium-sized companies, such transport is often not possible unless they can set up a horizontal collaboration with other companies. Pan et al. (2013) show that collaborative shipping may reduce CO_2 emissions by up to 14 percent without a modal shift, and by 52 percent when allowing a modal shift to include transport by train.

Until recently, the potential of horizontal supply chain collaboration has remained largely untapped. However, establishing horizontal partnerships is not straightforward. Even companies that are willing to cooperate may encounter practical impediments. A survey by Cruijssen et al. (2007b) shows that finding suitable partners is viewed as the third-largest impediment (after the allocation of the gains and the identification of partners that are able to coordinate the activities). Suitability depends on both tangible (e.g., companies with similar transport lanes) and nontangible factors (e.g., trust between companies). In this article we focus on the tangible aspects and evaluate the geographical compatibility of a partnership. Potential partners need to have transport routes that are in sufficiently close proximity such that their trucks and (or) empty space can be shared. Our tool allows a given company to identify all relevant collaborative shipping opportunities to enable that company to bundle transports that have roughly the same origin and destination, use an empty back-hauling trip for another transport, or avoid empty back-hauling trips by making a round trip that consists of three or more stops. In addition, our tool also detects collect-and-or-drop opportunities whereby shipments are collected and (or) dropped off en route. We refer to our tool, which Tri-Vizor and other companies have implemented, as the bundling, back-hauling, and round-trip tool (BBaRT).

Tri-Vizor is a facilitator and orchestrator of logistics horizontal collaboration partnerships. It identifies potential collaborative shipping partnerships and has responsibility for the operational coordination and synchronization of the shipments. To do so, Tri-Vizor relies on the geographical shipping data of these companies to analyze their compatibility. Over time, its database has become very large and thus time consuming to analyze manually. BBaRT helps Tri-Vizor to automate this analysis process and allows it to quickly detect promising partnerships that are compatible with respect to cargo and routing.

In the remainder of this article, we describe the problem in detail, present the working dynamics of BBaRT, and use an example to illustrate our methodology. We conclude by summing up the benefits accrued from using BBaRT.

Setting Up Collaborative Shipping Partnerships

Verstrepen et al. (2009) developed a conceptual framework to set up and to maintain collaborative shipping partnerships. When companies become aware of

Figure 1. The BBaRT Algorithm Enables the Identification of Collaboration Partners

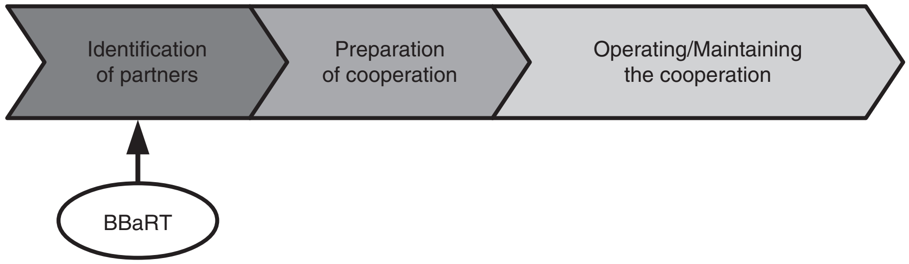

Note. The preparation (i.e., planning and synchronization of shipments) and the implementation of the partners' cooperation happens in later stages.

the need (and the benefits) of collaborative shipping, they start to search for potential collaboration partners. The BBaRT tool operates in this arena by identifying potential partners based on their geographical compatibility. When the collaboration partners have been identified, the cooperation, including the planning and synchronization of the shipments, the choice of a joint carrier, and the gain sharing, can be defined. The final stage addresses the effective implementation and operation of the collaboration using a control tower. Figure 1 illustrates this process and positions BBaRT's role within it.

The overlap between transportation networks provides opportunities for collaborative shipping. Over the years, Tri-Vizor has collected shipment lane data of thousands of companies. This database currently contains more than 130,000 shipment lanes, including company order data, GPS coordinates of the origins and destinations of the shipments, the types of transport used, and the yearly transport quantities expressed in truck equivalents. The initiative to detect potential collaboration partners can originate from either companies (i.e., they ask Tri-Vizor to set up and (or) operate a collaborative shipping partnership) or from Tri-Vizor who proactively approaches companies to propose a partnership based on the shipment lanes in its database. Our tool focuses on finding geographically compatible shipments that can share the same means of transportation.

Prior to the introduction of BBaRT, Tri-Vizor relied on data in Excel pivot tables, which were based on country and region codes of the origins and destinations. The shipments that had matching regional codes for origins and destinations were grouped and were manually checked for feasible bundling combinations. As the number of shipment lanes in its database grew over time, using pivot tables became cumbersome and time consuming. BBaRT immensely reduces manual efforts by automating the checking of geographical compatibility, and thus allows Tri-Vizor to identify more and better partnerships (e.g., more partners and more route overlaps).

BBaRT detects three types of collaborative shipping opportunities: (1) bundling of shipments in the same direction (bundling opportunities), (2) using shipments to utilize the (empty) back-hauling trip (back-hauling opportunities), and (3) round trips in which subsequent shipments can form a round trip as an alternative to back-hauling (round-trip) opportunities. Figures 2(a) to 2(c) show a graphical representation of the following collaborative opportunities:

- Bundling opportunities are found when two or more shipments have their origins and their destinations within a radius r of each other.
- Back-hauling opportunities require that the origins and destinations of two shipments lie within a radius r of each other, and vice versa.
- Round-trip opportunities require the destination and origin to be within a distance r of the respective origin and destination of other partnering lanes.

BBaRT also detects more complex opportunities that include multiple stops and (or) return trips. Figure 2(d) illustrates an example of such an opportunity. In addition, BBaRT identifies collect-and-or-drop opportunities in which shipments are collected and (or) dropped off en route; collect-and-or-drop opportunities can only be found in combination with a bundling, back-hauling, and (or) round-trip opportunity. Collect-and-or-drop opportunities can be found if a shipment can be collected and (or) dropped off at a

Figure 2. BBaRT Identifies Bundling, Back-Hauling, and Round-Trip Opportunities Where the Origin or Destination Points of Different Trips Are Within a Small Distance (r) of One Another

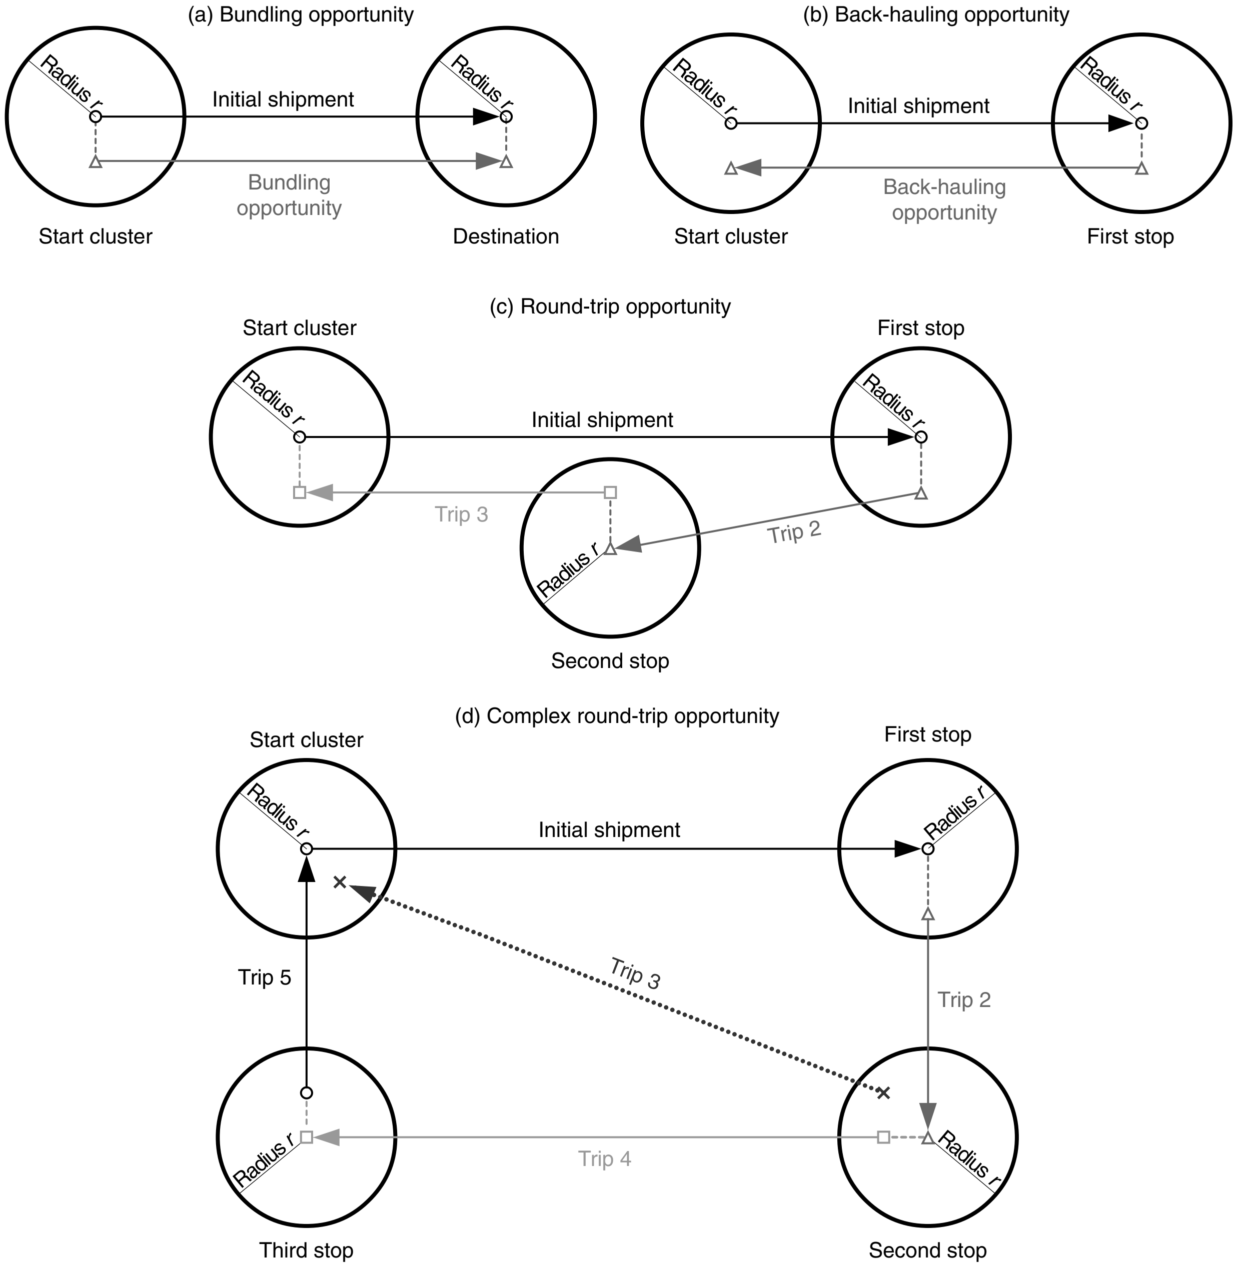

location that is within a distance r_2 of an existing route. Figures 3(a)–3(c) illustrate simple collect-and-or-drop opportunities. BBaRT, however, can also detect more complex collect-and-or-drop opportunities, such as the one illustrated in Figure 3(d). The user defines radius r and r_2: a smaller (larger) radius means that a shorter (longer) detour is needed to accommodate the collaborative shipping. As the radius increases, the number of proposed collaborative shipping opportunities also rises accordingly.

In essence, the problem is the identification of shipments with similar origins and destinations. More

Figure 3. BBaRT Identifies Opportunities Where Shipments Are Dropped Off En Route, Are Collected En Route, or Are Both Collected and Dropped Off En Route

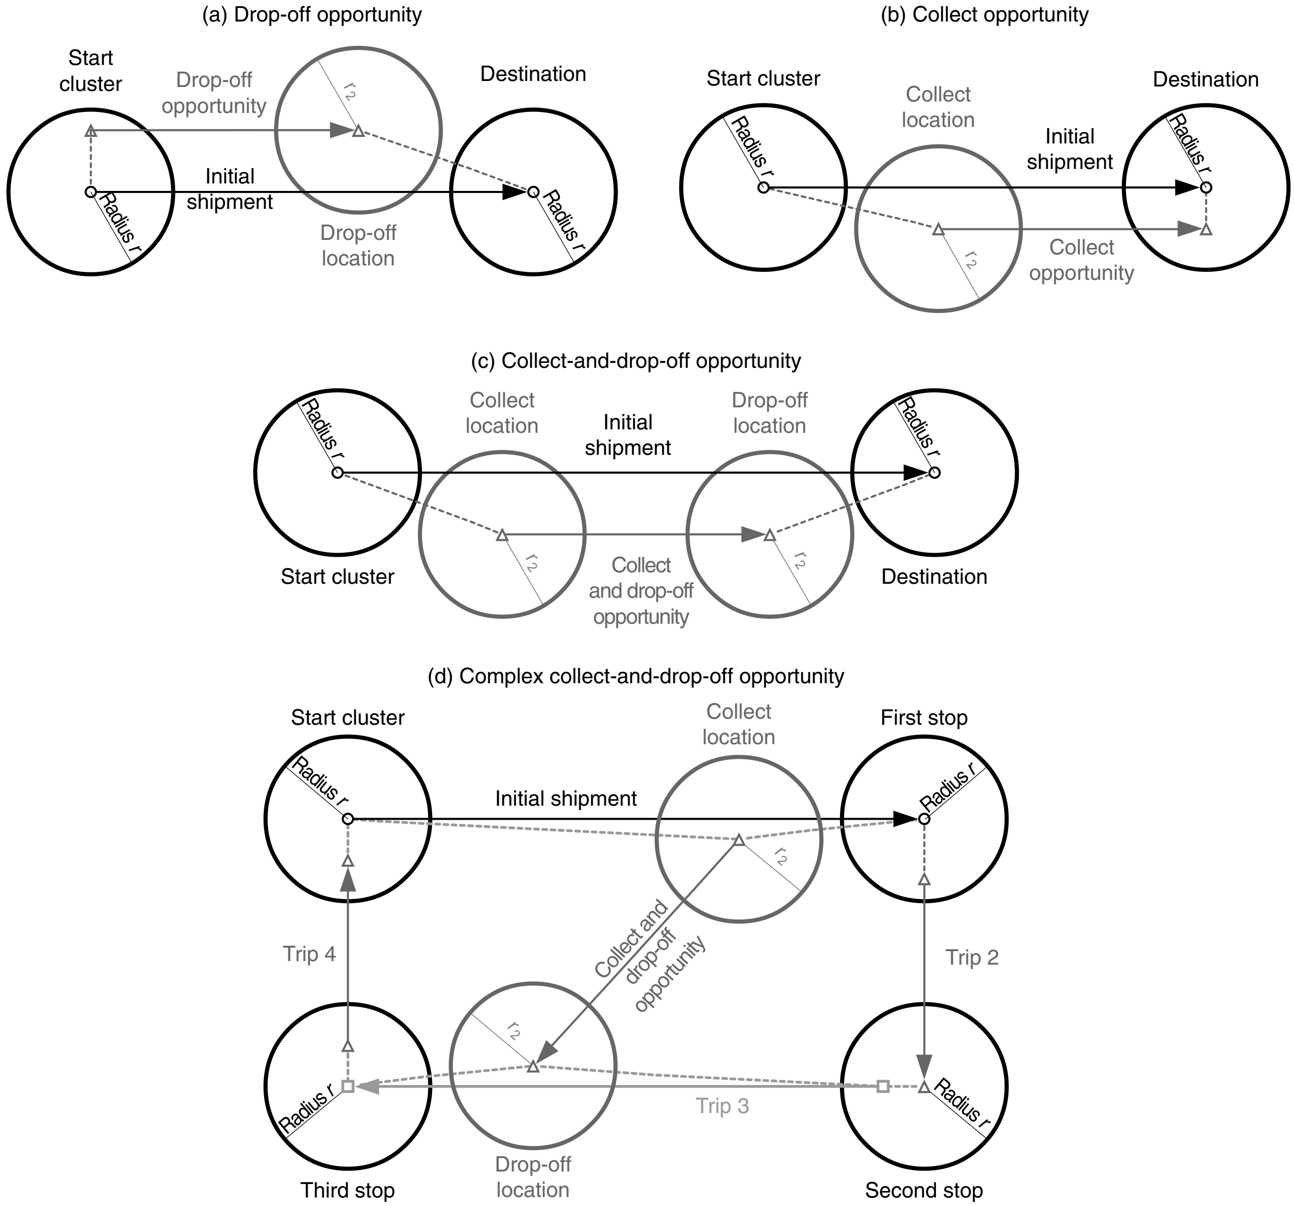

Note. To collect or drop off a shipment, its origin or destination needs to be within a distance r_2 of an existing route.

formally, we are looking for neighbors in a multidimensional space. Each shipment has four coordinates, resulting in a four-dimensional space: the latitude of the origin, the longitude of the origin, the latitude of the destination, and the longitude of the destination. The searching of a multidimensional space is a known problem in the literature. Solutions require the data to be in a specific data structure so that the data can be searched efficiently. For data with a high number of dimensions, more complex structures are advised (e.g., metric trees, R-trees, and k-dimensional trees); Bentley (1975), Bentley and Friedman (1979), Guttman (1984), Uhlmann (1991), Yi (2008), Moro (2009), Lakemond et al. (2013) provide more details on these structures. In this article, we adopt a different approach. Our rationale is threefold:

- A sorted list of lanes (even if they are sorted on multiple dimensions) does not permit the quick detection of collaborative shipping opportunities. To determine whether two lanes can be bundled, their geographical compatibility must be assessed (i.e., distances need to be calculated). As such, sorting is only part of the solution. To quickly detect bundling opportunities, BBaRT combines sorting with a bounding-box approach. Whereas sorting allows the quick lookup of potential partner lanes, a bounding-box approach allows the quick filtering of these lanes based on their geographical compatibility. To the best of our knowledge, we are the first to combine sorting and a bounding-box approach to find collaborative shipping opportunities.
- To detect collect-and-or-drop opportunities, we calculate the rotated coordinates of the origin and the destination of each lane (see also infra for more details). A list of lanes that is sorted on the nonrotated coordinates is of no use here. Again, we need to combine sorting with a bounding-box approach.
- Sorting in itself does not permit finding round trips or more complex collaborative shipping opportunities (e.g., those with two or more stops). To detect these opportunities, a fast queue-based search algorithm is required (see also infra for more details).

Methodology

In this section we discuss the rationale of the algorithms we developed to identify collaborative shipping opportunities. We relegate the details of our algorithms to the appendix. BBaRT applies three steps: (1) preparation of the database, (2) identification of collaborative opportunities, and (3) ranking of the opportunities. Next, we discuss each of these steps and illustrate them using a small sample database. For each shipment lane, the sample database contains the origin and destination coordinates, the yearly transport quantities (in truck equivalents), the distance between the origin and destination, and the identification (ID) of the company that uses the shipment lane.

Data Preparation

First, we reorganize the shipment data into a useable data structure. All shipments that have identical origins and destinations (e.g., because they are from the same company) are grouped into one-way lanes with the coordinates of the origin and the destination, and the yearly number of truck equivalents that are shipped through this lane, which is the sum of the individual shipments with the same origin and destination. The data structure links the lanes to the original shipments so that all other information (order data, the type of transport, and the yearly transport quantities) is preserved. Note that the timings of the shipments do not have to be taken into account, because identical timings of the shipments are not necessarily required to find good collaborative shipping opportunities. Padilla Tinoco et al. (2015) demonstrate that collaborative shipping is always beneficial, even if companies have to adjust their order and (or) transport frequencies and the timings of their shipments. Preparing the data in this manner reduces the number of candidate lanes (and prevents generating obvious bundling opportunities between identical lanes) without any loss of information. Table 1 illustrates the result of this data-preparation step.

Subsequently, the data set is sorted on the latitudinal coordinate. Sorting on one dimension only permits a quick lookup of candidate lanes. A bounding-box approach can then be used to filter the candidate lanes based on their geographical proximity.

Identification of Collaborative Shipping Opportunities

Figure 4 plots the lanes shown in Table 1 on a map. As this figure shows, four clusters have shipments with

Table 1. We Aggregate the Data in the Shipment Database to Obtain a Useable Data Structure with Candidate Lanes That Have a Unique Origin and Destination

			Origin		Destination	
Lane	Truck equivalents	Trip distance	Latitude	Longitude	Latitude	Longitude
1	50	725	29.11	19.78	52.96	88.2
2	50	694	75.45	72.33	30.05	19.8
3	10	716	30.05	19.8	53.5	87.4
4	30	266	53.5	87.4	75.45	72.33
5	30	719	29.45	20.04	53.17	87.88
6	20	275	53.17	87.88	76.08	72.73
7	40	704	76.08	72.73	29.45	20.04
8	90	204	84.21	19.81	63.77	19.85
9	90	467	77.41	71.95	75.04	25.35
10	90	448	74.73	24.84	30.18	19.79
11	10	165	47.62	71.78	53.05	87.38
12	10	260	75.22	70.33	57.35	51.48
13	10	151	57.91	84.56	70.44	76.12
14	20	193	46.39	68.24	60.54	55.12

Figure 4. The BBaRT Tool Clusters Data in Table 1 to Identify a Number of Collaboration Opportunities Including Bundling (e.g., Lane 1 and Lane 5), Round-Trip (e.g., Lanes 5, 6, and 7), and Collect-and-or-Drop Opportunities (e.g., Lanes 4 and 13)

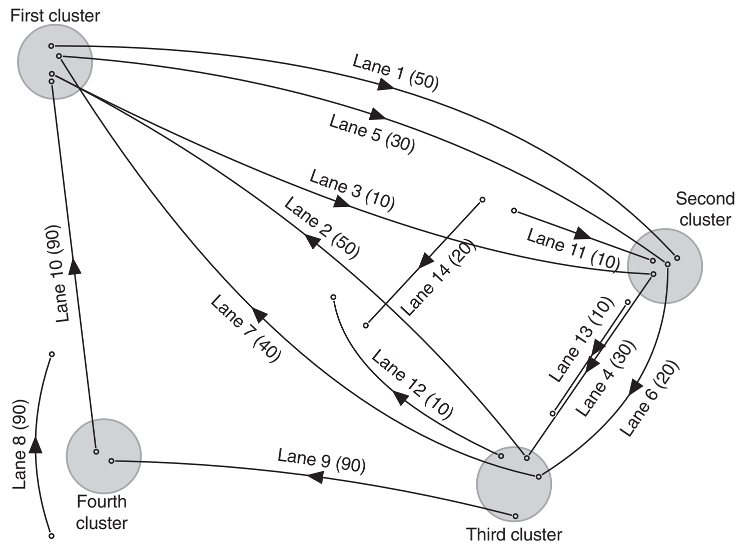

similar origins and destinations. We illustrate the use of our algorithms by applying it to our sample database.

To find bundling, back-hauling, and round-trip opportunities, we need to be able to quickly determine the neighboring lanes in the area around a point P (i.e., the origin and (or) destination of a lane). The sorted list of lanes allows us to quickly look up potential candidate lanes. However, because the lanes are sorted only on the latitudinal coordinate, it is possible (and even likely) that some of these candidate lanes are not close to point P at all. To filter the lanes that are close to point P, we use a bounding-box approach. For each lane that arrives in and (or) departs from a location in the bounding box for point P, we calculate the distance toward point P. We retain only those lanes for which the distance is less than r; r is defined by the user as a prespecified parameter that denotes the maximum detour the user is willing to accommodate to benefit from joint transport. All other lanes can be discarded. Note that if the number of lanes is large, calculating the pairwise distance between each of the lane coordinates to determine whether the origin and (or) destination of a lane is in the proximity of the origin and (or) destination of another lane, would not be a practical (or feasible) solution. In the case of Tri-Vizor for example, calculating or storing the distances between the origins and (or) destinations of the 130,000 shipment lanes in its database would be impossible.

The above approach permits the identification of bundling, back-hauling, and round-trip opportunities; detecting collect-and-or-drop opportunities requires a slightly different approach. For example, consider Lanes 4 and 13. From Figure 4 we can see an additional opportunity to bundle: the shipment associated with Lane 13 can be picked up and dropped off en route to the destination of Lane 4. In this case the bounding box for the origin and (or) destination of Lane 4 is not useful; we need a bounding box for Lane 4 itself. Unfortunately, however, this bounding box is not perpendicular to the coordinate system; therefore, we cannot quickly determine which points are located inside the bounding box. One possible solution is to rotate the coordinate system with an angle that corresponds to the angle of Lane 4. By doing so, the bounding box around Lane 4 becomes perpendicular to the rotated coordinate system. To detect lanes that are close to Lane 4, we again use a sorted list of lanes. This time, however, the list is sorted on the *rotated* latitudinal coordinate. The bounding box around Lane 4 serves as a filter to determine whether a candidate lane is within a distance r_2 of Lane 4 itself (where r_2 is specified by the user). Figure 5 illustrates the process. From Figure 5(b), we can see that the shortest distance between the origin of Lane 13 and Lane 4 equals 1.6 km (i.e., 102.49 minus 102.33). The shortest distance between the destination of Lane 13 and Lane 4 is 2.9 km. Therefore, if r_2

Figure 5. To Easily Detect the Collect-and-or-Drop Opportunity of Lane 13 with Lane 4, BBaRT Rotates the Coordinate System to Make Lane 4 Perpendicular to One of the Axes in the New Coordinate System

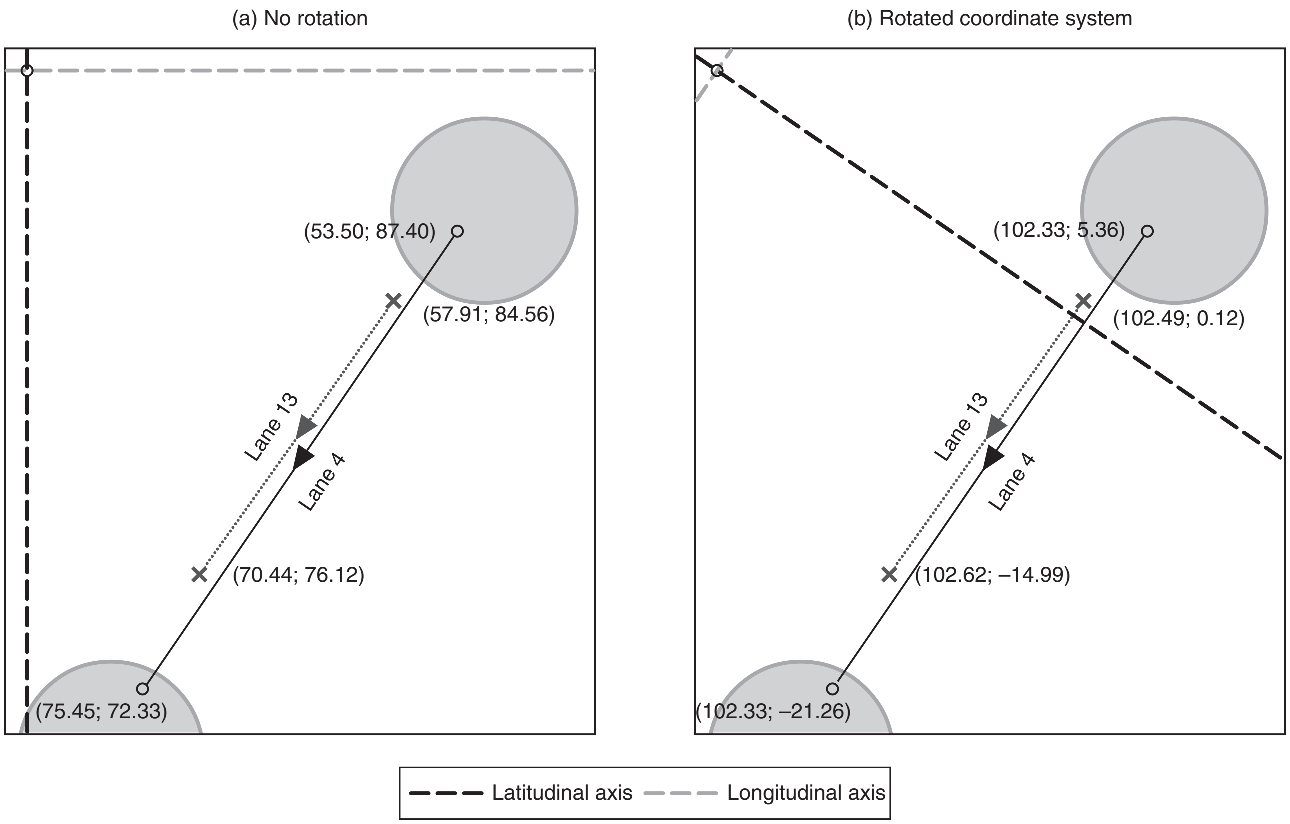

is set larger than 2.9 km, Lane 13 is detected as a collect-and-or-drop opportunity for Lane 4.

The combination of a sorted list and a bounding-box approach proves to be of critical importance in determining bundling, back-hauling, round-trip, and collect-and-or-drop opportunities. To filter out the opportunities with the largest savings potential (in terms of transport cost and carbon footprint), we rank them. We discuss this ranking in section *Ranking of Opportunities*.

Ranking of Opportunities

The output of the algorithms is a list of potential partnerships. To evaluate these collaborative partnerships, we assess the economic and environmental benefits of each collaborative shipping opportunity. The financial and environmental cost for each transport lane is represented on a per-volume, per-distance basis; therefore, we report on the distance traveled and the number of tonne-kilometers (tkm) as a proxy for the transportation costs and GHG emissions. The degree of transportation cost savings and GHG reductions then depend on the joint (shared) distance traveled and on the number of shipments that are bundled over this shared distance, and thus on the number of trucks that can be reduced. As a proxy for these savings, we measure the joint (shared) distance, the total volume that is shipped over this shared distance, and (or) the joint (shared) number of tonne-kilometers. Hence, we obtain the following set of key performance indicators (KPIs) for each collaborative shipping opportunity:

- The total distance traveled (note that the total distance depends on the selected routing).
- The total shared distance (i.e., the distance over which shipments are bundled and the transport is joint).
- The total volume that is shipped by the identified opportunity.
- The total shared volume (i.e., the joint volume that is shipped over the shared distance).

• The total number of tkms (i.e., volume times distance that is traveled by that volume) of the identified opportunity.

• The total number of shared tkms (i.e., the combination of shared volume and shared distance).

Based on these KPIs, the following ratios assess the economical and environmental benefits of a collaborative shipping opportunity; for each of these ratios, the user can specify a minimum required value:

• The ratio of shared distance to total distance (representing the overlap of the lanes).

• The ratio of shared volume to total volume (the volumes shipped over the shared distance compared to the volumes shipped over the total distance).

• The ratio of shared tkms to total tkms.

Typically the total number of shared tkms will be the most relevant, because it is the best proxy for the savings in transportation costs and GHG emissions. However, by giving a weight to each of these KPIs, the user can decide which KPI is most relevant. For example, a user can be interested in bundling shipments only over long distances, or prefer bundling high volumes over shorter distances, thus enabling the potential partnerships to be ranked in accordance with the preferences of the user.

To find the total distance of the collaborative shipping opportunity, we need to determine the routing of the collaborative shipping. This problem is known as the clustered traveling salesman problem (CTSP); Chisman (1975) provides details. The CTSP is an extension of the traveling salesman problem (TSP) in which a set of cities is partitioned into clusters, and the salesman has to visit the cities of each cluster consecutively (Helsgaun 2014). Because the CTSP is an NP-hard problem, and because computation speed is important for the user, we use a simple closest-neighbor heuristic to determine the routing of the collaborative shipping. In addition, upon arrival at a cluster, we first drop off shipments before collecting any new shipments. Note that because we are not solving the routing problem at the operational level, we do not incorporate truck capacities. Once a collaborative shipping opportunity is selected, its feasibility should be tested by solving a pickup and delivery problem; see Savelsbergh and Sol (1995) for details.

The best collaborative bundling opportunity in our sample database visits four clusters and comprises 13 lanes. Table 2 reports the details of all KPIs for this collaborative opportunity. The collaborative shipping process is as follows:

• First, we collect the shipments of Lanes 5, 1, and 3 in the first cluster. Next, we depart for the second cluster. On the way to the second cluster, we also collect the shipments of Lanes 14 and 11.

• Upon arrival at the second cluster, we drop off the shipments of Lanes 11, 3, 5, and 1. Next, we collect the shipments of Lanes 6 and 4, and we depart for the third cluster. On the way to the third cluster, we collect and drop off the shipment of Lane 13.

• At the third cluster, we drop off the shipments of Lanes 4 and 6. At this point, the routing splits into two paths.

• A first path (Path B in Table 2) collects the shipments of Lanes 7, 2, and 12. On the way back to the first cluster, we drop off the shipments of Lanes 14 and 12. Upon arrival at the first cluster, we drop off the shipments of Lanes 2 and 7.

• The second path (Path C in Table 2) collects the shipment of Lane 9 and moves on to the fourth cluster. Upon arrival at the fourth cluster, we drop off the shipment of Lane 9 and collect the shipment of Lane 10. Finally, we return to the first cluster and drop off the shipment of Lane 10.

This collaborative shipping opportunity ranges over 2,674 km. Of these, 1,723 km are shared between two or more lanes; that is, 64.43 percent of the distance is shared. In total, we ship 460 truck equivalents—280 are over the shared distance; that is, 60.87 percent of the volume is shared. Finally, the total number of tkms equals 246,054, 66.35 percent of which is shared; that is, 163,245 tkms are shared. The bulk of the nonshared transport originates from Lanes 9 and 10. Therefore, a user who wants to increase the shared ratios can choose to omit those lanes from the collaborative shipping opportunity.

Use of BBaRT

Tri-Vizor has used the BBaRT tool since October 2013. For Tri-Vizor, BBaRT has become an indispensable asset to identify new collaborative shipping opportunities and to bring together potential collaboration partners. As the number of shipment lanes in its database grew over time and its networks grew in complexity, manually detecting bundling or back-hauling opportunities using Excel pivot tables became more difficult and

Table 2. The Routing of the Example Bundling Opportunity Consists of Three Parts: A Common Path A, a Path B, and a Path C

Departure location	Arrival location	Trip distance	Volume before departure	Volume picked up upon departure	Volume during trip	Volume dropped off upon arrival	Volume after arrival	Cumulative distance	Cumulative volume	Cumulative tkm	Cumulative shared distance	Cumulative shared volume	Cumulative shared tkm
Path A													
5o	1*o*	4	0	30	30	0	30	4	30	128	0	0	0
1o	3*o*	9	30	50	80	0	80	14	80	881	9	80	752
3o	14*o*	511	80	10	90	0	90	525	90	46,890	521	90	46,762
14o	11*o*	37	90	20	110	0	110	562	110	51,013	558	110	50,884
11o	11*d*	165	110	10	120	10	110	728	120	70,834	723	120	70,706
11d	3*d*	5	110	0	110	10	100	732	120	71,330	728	120	71,201
3d	5*d*	6	100	0	100	30	70	738	120	71,912	734	120	71,784
5d	1*d*	4	70	0	70	50	20	742	120	72,180	737	120	72,052
1d	6*o*	4	20	0	20	0	20	746	120	72,257	737	120	72,052
6o	4*o*	6	20	20	40	0	40	751	140	72,490	743	140	72,285
4o	13*o*	52	40	30	70	0	70	804	170	76,161	796	170	75,956
13o	13*d*	151	70	10	80	10	70	955	180	88,247	947	180	88,042
13d	4*d*	63	70	0	70	30	40	1,018	180	92,645	1,010	180	92,440
4d	6*d*	7	40	0	40	20	20	1,025	180	92,943	1,017	180	92,738
Path B (first return trip)													
6d	7*o*	15	20	0	20	0	20	1,041	270	92,943	1,017	180	92,738
7o	2*o*	467	20	40	60	0	60	1,507	310	134,937	1,017	220	92,738
2o	12*o*	0	60	50	110	0	110	1,507	360	134,937	1,017	270	92,738
12o	14*d*	7	110	10	120	20	100	1,515	370	135,385	1,025	280	93,186
14d	12*d*	20	100	0	100	10	90	1,535	370	137,600	1,045	280	95,401
12d	2*d*	211	90	0	90	50	40	1,746	370	162,966	1,256	280	120,767
2d	7*d*	48	40	0	40	40	0	1,795	370	167,806	1,304	280	125,607
Path C (second return trip)													
6d	9*o*	418	0	0	0	0	0	2,213	180	205,444	1,723	180	163,245
9o	9*d*	6	0	90	90	90	0	2,219	270	205,703	1,723	180	163,245
9d	10*o*	6	0	0	0	0	0	2,225	370	205,703	1,723	280	163,245
10o	10*d*	448	0	90	90	90	0	2,674	460	246,054	1,723	280	163,245
Shared distance ratio		0.6443											
Shared volume ratio		0.6087											
Shared tkm ratio		0.6635											
Xo is the origin of Lane X and Xd is the destination of Lane X													

Notes. Each path consists of a number of locations that are visited. For each path, the table shows the distance and volume traveled between two locations, and its impact on the KPIs.

time consuming, and using these tables to find round-trips or collect-and-or-drop opportunities was simply impossible.

When a new company contacts Tri-Vizor to search for potential partnerships, Tri-Vizor uses BBaRT to match that company's lanes with the current database of over 130,000 shipment lanes and BBaRT rapidly provides a long list of collaborative shipping opportunities. Since BBaRT's introduction at Tri-Vizor, it has been used to identify at least one or two collaborative shipping opportunities per month.

One example of an opportunity that could not have been detected without using BBaRT is the regional freight-flow bundling in the Province of Gelderland in the Netherlands, where 13 companies, including a mix of small and medium-sized enterprises (SMEs) and multinationals, showed an interest in bundling transports. The user input consisted of an Excel template in which the user collected all structural full truckload (FTL) and LTL flows, including the coordinates of origin and destination, the yearly shipped volume (in truck equivalents), the current freight mode, and the specific transport conditions if applicable (e.g., temperature controlled, dangerous goods, contaminated risks). These data were collected under a nondisclosure agreement. The collected flow data were cleaned and converted into standard format, and a geo-coding was performed on the origin and destination locations (i.e., they were converted to XY-coordinates) to enable additional automated processing. The shipper data were then merged into one project database with 4,512 transport lanes and 2,426 unique origin and destination locations.

Synergies were analyzed internally (between the 13 Gelderland shippers) and externally (with Tri-Vizor's entire database). The BBaRT tool detected approximately 1,000 internal collaborative shipping opportunities between the Gelderland shippers (bundling + round trips); of these, 14 were identified as opportunities with high potential (i.e., high volumes and high shipment frequency). The external analysis, matching the Gelderland shipper flows against Tri-Vizor's database of more than 130,000 EMEA shipment lanes, revealed approximately 4,500 collaborative shipping opportunities; of these, 300 had a transport frequency of more than once per week. From this set, Tri-Vizor used BBaRT to retain 16 interesting combinations. This was the basis for starting a discussion with the companies involved. Examples of topics discussed included the synchronization of the shipments, the joint cost drivers, the choice of the joint transport company, and the transparent sharing of the gains.

A major Belgian retailer has also used the BBaRT tool to identify collaborative shipping opportunities in its supplier base. Due to strict confidentiality reasons, we are not able to report on the details of this collaboration. However, as with Tri-Vizor, as the number of this retailer's shipment lanes grew, the company's ability to manually identify collaborative shipping opportunities became an impossible task.

Note that BBaRT is used in the first stage of setting up a collaborative shipping partnership; that is, it is used to identify potential bundling partners (Figure 1). The additional optimization and (or) implementation of the bundling occurs in later stages. For example, after the identification of the potential partners, topics that must be discussed include shipment synchronization (i.e., how to synchronize the joint replenishment orders), planning details (e.g., fixed weekly departure days or dynamic planning), choice of the joint carrier, and gain sharing among the collaboration partners. At this stage, the collaboration may cease to exist. However, without identification of the opportunity, the discussion would never have been initiated in the first place.

We expect the collaboration gains found by using BBaRT to exceed those of previous opportunities identified by Tri-Vizor, where only two companies bundled their transport. We refer to two pilot reports. The first discusses collaborative shipping agreements between two manufacturing companies, JSP and Hammerwerk, for the bundling of road transport (Verstrepen and Jacobs 2012). The result was a 20–30 percent reduction in CO_2 emissions. The second report addresses the horizontal collaboration in fresh and chilled retail distribution between the fast-moving consumer goods (FMCG) shippers Nestle & PepsiCo (Jacobs et al. 2014). The result was transport cost savings of 10–15 percent. When the number of participating companies is greater than two, we expect higher gains, because there are more collaborative shipping opportunities to exploit and thus greater savings to reap.

Directions for Future Research

Although BBaRT in its current version has proven to be an extremely valuable tool for Tri-Vizor, we envision several directions for future research. First, BBaRT is currently limited to finding bundling opportunities that are geographically compatible. Future research should focus on extending BBaRT to include cargo compatibility, order frequencies, and other important characteristics. At this moment, the participating companies do not provide such information; however, we believe that this will happen in the near future. Second, BBaRT currently uses a weighted function to determine the score of a collaborative shipping opportunity. An interesting enhancement would be the examination of more advanced methods of multiobjective optimization, such as finding Pareto-efficient bundles. Third, BBaRT is limited to identifying collaborative shipping opportunities. Following the identification phase, the implementation of collaborative shipping opportunities at the operational level requires modeling and solving a pickup and delivery routing problem. Therefore, future research could study the type of pickup and delivery problem that arises from each type of collaborative shipping opportunity that BBaRT identifies. This would enable BBaRT to go beyond the identification phase and support implementation at the operational level.

Appendix

In the appendix, we explain how to obtain sets of matching lanes (i.e., geographically compatible lanes). We then show how these sets of matching lanes are used to detect collaborative shipping opportunities. Finally, we discuss how to calculate the routing and the KPIs of a collaborative shipping opportunity.

Sets of Matching Lanes

We define the following sets of matching lanes:

- $\mathbf{M}_i^{OO}$: the set of lanes that has its origin within a radius r of the origin of Lane i.
- $\mathbf{M}_i^{DO}$: the set of lanes that has its origin within a radius r of the destination of Lane i.
- $\mathbf{M}_i^{OD}$: the set of lanes that has its destination within a radius r of the origin of Lane i.
- $\mathbf{M}_i^{DD}$: the set of lanes that has its destination within a radius r of the destination of Lane i.
- $\mathbf{M}_i^{LO}$: the set of lanes that has its origin within a distance r_2 of Lane i.
- $\mathbf{M}_i^{LD}$: the set of lanes that has its destination within a distance r_2 of Lane i.

Using these sets of matching lanes, we can quickly identify bundling, back-hauling, round-trip, and collect-and-or-drop opportunities. In addition, let $\mathbf{A}_t^O$ denote the list of lanes that is sorted on the latitudinal coordinate of the origin, which is rotated by t degrees (i.e., $\mathbf{A}_0^O$ is the sorted list of original coordinates). We define $\mathbf{A}_t^D$ analogously, however, it is sorted on the latitudinal coordinate of the destination. We also define $X_{i,t}^O$, $X_{i,t}^D$, $Y_{i,t}^O$, and $Y_{i,t}^D$ as the longitudinal and latitudinal coordinates of the origin and the destination of Lane i that are rotated by t degrees. Note that, if $t=0$, $X_{i,t}^O$, $X_{i,t}^D$, $Y_{i,t}^O$, and $Y_{i,t}^D$ represent the original longitudinal and latitudinal coordinates of the origin and the destination of Lane i.

Algorithm 1 outlines how to obtain set $\mathbf{M}_i^{OO}$. Sets $\mathbf{M}_i^{DO}$, $\mathbf{M}_i^{OD}$, and $\mathbf{M}_i^{DD}$ are obtained analogously. To obtain sets $\mathbf{M}_i^{OO}$, $\mathbf{M}_i^{DO}$, $\mathbf{M}_i^{OD}$, and $\mathbf{M}_i^{DD}$, we first use a list of lanes that is sorted on the latitudinal coordinate of the origin (destination) (depending on whether we are looking for a matching origin or a matching destination) to create a set of candidate lanes (i.e., $\mathbf{L}_i^C$), each of which may have its origin (destination) close to the origin (destination) of Lane i. More specifically, the set of candidate lanes contains all lanes that have latitudinal coordinate of origin (destination) within the interval $]Y_{i,0}-r;\, Y_{i,0}+r[$. Next, we filter out those lanes whose longitudinal coordinate of origin (destination) falls within the interval $]X_{i,0}-r;\, X_{i,0}+r[$. All resulting lanes (i.e., the set of filtered candidate lanes $\mathbf{L}_i^F$) have their origin (destination) within the bounding box around the origin (destination) of Lane i. To ensure that the origin (destination) of a filtered candidate Lane j is within a distance of r km from the origin (destination) of Lane i, we need to calculate distance d_{ij}. If d_{ij} is smaller than r, Lane j is a matching lane, and it is added to the set of matching lanes of Lane i.

Algorithm 1 (Finding set of matching lanes $\mathbf{M}_i^{OO}$).

```
for All Lanes i do
  Use sorted list A_0^O to obtain set of candidate lanes L_i^COO
    that may have origin close to the origin of Lane i
  for All Lanes j in L_i^COO do
    if X_{j,0}^O > (X_{i,0}^O − r) then
      if X_{j,0}^O < (X_{i,0}^O + r) then
        Lane j has origin inside the bounding box around
          the origin of Lane i
        Add Lane j to set of filtered candidate lanes L_i^FOO
          that may have origin close to the origin of Lane i
      end if
    end if
  end for
  for All Lanes j in L_i^FOO do
    Calculate distance d_ij^OO between the origin of Lane i
      and the origin of Lane j
    if d_ij^OO < r then
      Add Lane j to set of matching lanes M_i^OO
    end if
  end for
```

end for.

Let θ_i denote the smallest positive angle of Lane i (expressed in degrees):

$$\theta_i = \begin{cases} \arctan\left[\dfrac{Y_{i,0}^D - Y_{i,0}^O}{X_{i,0}^D - X_{i,0}^O}\dfrac{180}{\pi}\right] & \text{if } \arctan\dfrac{Y_{i,0}^D - Y_{i,0}^O}{X_{i,0}^D - X_{i,0}^O} \geq 0, \\ 360 + \arctan\left[\dfrac{Y_{i,0}^D - Y_{i,0}^O}{X_{i,0}^D - X_{i,0}^O}\dfrac{180}{\pi}\right] & \text{if } \arctan\dfrac{Y_{i,0}^D - Y_{i,0}^O}{X_{i,0}^D - X_{i,0}^O} < 0. \end{cases}$$

When rotating the coordinate system by θ_i degrees, Lane i is perpendicular to the longitudinal axis (i.e., in the rotated coordinate system $Y_{i,\theta_i}^O = Y_{i,\theta_i}^D$). Rotating the coordinate system allows us to detect collect-and-or-drop opportunities. Algorithm 2 outlines how to obtain set $\mathbf{M}_i^{LO}$. Set $\mathbf{M}_i^{LD}$ is obtained analogously. To obtain sets $\mathbf{M}_i^{LO}$ and $\mathbf{M}_i^{LD}$, we iterate over all integer degrees between 0 and 360. For each degree (1) we determine the rotated coordinates of origin and (or) destination for each lane, (2) we sort all lanes on the rotated latitudinal coordinate of the origin and (or) destination, and (3) we iterate over all lanes. If the angle of Lane i rounds down to the current integer degree, we try to identify all collect-and-or-drop opportunities for Lane i. To detect the collect-and-or-drop opportunities for Lane i, we apply a logic similar to the logic that we use in Algorithm 1. First, we obtain a set of candidate lanes by using a list of lanes that is sorted on the rotated latitude of the origin and (or) destination. Note that we rotate by an integer number of degrees. Most likely, however, the angle of Lane i is not an integer number. As such, Lane i will only be approximately perpendicular to the longitudinal axis of the coordinate system. To take this deviation into account, we need to inflate the bounding box around Lane i itself. The inflation is captured by ε_i:

$$\varepsilon_i = 0.0175 d_i,$$

where d_i is the length of Lane i and 0.0175 is the maximum error incurred per km, given that the maximum difference in angle is 1 degree (i.e., if Lane i is 100 km long, the maximum error due to inaccurate rotation is 1.75 km; the slope of a line that is tilted by 1 degree is 1.75 percent). When rotating by t degrees, the set of candidate lanes contains all lanes that have latitudinal coordinate of origin and (or) destination within the interval $]Y_{i,t} - (r + \varepsilon_i); Y_{i,t} + (r + \varepsilon_i)[$. Next, we filter out those lanes whose longitudinal coordinate of origin and (or) destination falls within the interval $]X_{i,t} - (r_2 + \varepsilon_i); X_{i,t} + (r_2 + \varepsilon_i)[$. All resulting lanes (i.e., the set of filtered candidate lanes $\mathbf{L}_i^F$) have their origins and (or) destinations within the inflated bounding box around Lane i itself. To ensure that the origin and (or) destination of a filtered candidate Lane j is within a distance of r_2 km from Lane i, we need to rotate the latitudinal coordinate of the origin and (or) destination of the filtered candidate lanes by θ_i degrees. These rotated coordinates can then be used to determine the shortest distance between the origin and (or) destination of the filtered candidate lanes and Lane i. If this distance is smaller than r_2, we have found a matching lane and add it to the set of matching lanes.

Algorithm 2 (Finding set of matching lanes $\mathbf{M}_i^{LO}$).

```
for All integer degrees t = 0 up to t = 360 do
  for All Lanes i do
    Calculate rotated coordinates X^O_{i,t} and Y^O_{i,t} of the
      origin of Lane i.
  end for
  Obtain A^O_t by sorting all lanes on their latitudinal
    coordinate of the origin that is rotated with t degrees.
  for All Lanes i do
    if ⌊θ_i⌋ = t then
      Calculate rotated coordinate Y^O_{i,θ_i} of the origin
        of Lane i
      Use sorted list A^O_t to obtain set of candidate lanes
        L^{CLO}_i that may have origin close to Lane i itself
      Use inflated bounding box to obtain filtered
        set of candidate lanes L^{FLO}_i that may have origin
        close to Lane i itself
      for All Lanes j in L^{FLO}_i do
        Calculate rotated coordinates Y^O_{j,θ_i} of the origin
          of Lane j
        if |Y^O_{i,θ_i} − Y^O_{j,θ_i}| < r_2 then
          Add Lane j to set of matching lanes M^{LO}_i
        end if
      end for
    end if
  end for
end for.
```

Identifying Collaborative Shipping Opportunities

We use a queue-based approach to determine all collaborative shipping opportunities. Algorithm 3 outlines the approach. We first create a set of initial queue elements, each of which contains a single lane. Next, we process and generate new queue elements until we have evaluated all collaborative shipping opportunities. Note that each queue element has an active cluster that contains a single lane, which connects the active cluster with the next cluster. When processing a queue element, we first try to find additional bundling opportunities for the lane that departs from the active cluster (using sets $\mathbf{M}_i^{OO}$ and $\mathbf{M}_i^{DD}$). In the example of a collaborative shipping opportunity in Figure 1, Lane 1 is bundled with Lanes 3 and 5. After all bundling opportunities are found at the active cluster, we increment the active cluster (i.e., the next cluster becomes the active cluster). Next, we try to identify all return trips to previous clusters using sets $\mathbf{M}_i^{DO}$ and $\mathbf{M}_i^{DO}$. If the maximum number of clusters has not yet been reached (the maximum number of clusters can be specified by the user), we initiate a new queue element by adding a new lane that departs from the active cluster. If, however, the maximum number of clusters has been reached, no more lanes can be added, and we perform the following steps:

• Detection of collect-and-or-drop opportunities using sets $\mathbf{M}_i^{LO}$ and $\mathbf{M}_i^{LD}$. Note that a lane is included as a collect-and-or-drop opportunity if its origin is close to any location and (or) lane in the bundling opportunity and its destination is close to another location and (or) lane that is visited afterwards.

• Determination of the routing (i.e., the sequence in which the locations will be visited) using Algorithm 4.

• Calculation of KPIs given the routing.

We record the collaborative shipping opportunity only if: (1) it has a better weighted KPI score than the previously recorded opportunities, and (2) it differs sufficiently from previously recorded opportunities. To determine whether the opportunity is sufficiently different, we evaluate the number of lanes that are shared between recorded opportunities. The user can specify that a maximum percentage of lanes are shared. As a result, each recorded opportunity has at most a given percentage of shared lanes.

Algorithm 3 (Identification of collaborative shipping opportunities).

for All Lanes i **do**
  Initialize new queue element where: (1) Lane i connects cluster 1 and cluster 2, and (2) the first cluster is the active cluster
end for
while There are queue elements left to be processed **do**
  Increment the active cluster
  for All clusters visited before the active cluster **do**
    Find bundling opportunities towards the active cluster using sets $\mathbf{M}_i^{OO}$ and $\mathbf{M}_i^{DD}$
  end for
  Find all return trips to previous clusters using sets $\mathbf{M}_i^{DO}$ and $\mathbf{M}_i^{OD}$
  if The maximum number of clusters has been reached **then**
    No more lanes can be added: (1) detect all collect-and-or-drop opportunities using sets $\mathbf{M}_i^{LO}$ and $\mathbf{M}_i^{LD}$, (2) determine routing using Algorithm 4, (3) calculate KPIs, and (4) record the collaborative shipping opportunity if it is better than and sufficiently different from already recorded opportunities
  else
    Initialize a new queue element where a new lane is added that has origin in the active cluster and destination that will form a new cluster
  end if
end while.

Calculation of KPIs of Collaborative Shipping Opportunities

To calculate the KPIs, we must first determine the routing (i.e., the sequence in which the locations of the collaborative bundling opportunity are visited). Algorithm 4 outlines how the routing is obtained using a closest-neighbor heuristic. In the algorithm, *forward* clusters are defined as clusters that have not yet been visited and *return* clusters are clusters that have been visited. Note that, upon arrival at a cluster, shipments on transport are first dropped off before any new shipments are collected.

Algorithm 4 (Algorithm to determine the routing of the collaborative shipping opportunity).

Initialize entry location of the first cluster
for All cluster **do**
  Start from the entry location of the cluster and use closest-neighbor heuristic to drop off all shipments bound for this cluster
  The *pivot* location is set as the last location where a shipment was dropped off
  for All *forward* clusters **do**
    Start from the *pivot* location and use closest-neighbor heuristic to collect all shipments that are bound for the *forward* cluster
    Use closest-neighbor heuristic to collect and (or) drop off all shipments en route to the *forward* cluster
    Use closest-neighbor heuristic to determine the entry point of the *forward* cluster
  end for
  for All *return* clusters **do**
    Start from the *pivot* location and use closest-neighbor heuristic to collect all shipments that are bound for the *return* cluster
    Use closest-neighbor heuristic to collect and (or) drop off all shipments en route to the *return* cluster
    Use closest-neighbor heuristic to determine the reentry location of the *return* cluster
    Start from the reentry location and use closest-neighbor heuristic to drop off all shipments that are bound for the *return* cluster
  end for
end for.

To calculate the KPIs, we keep track of the shipments that arrive and depart at each location. This way, we can update the KPIs at each location. For example, consider the second location in the example of a collaborative shipping opportunity illustrated in Table 2. The second location has coordinates (29.11; 19.78), and we arrive there after we have collected the shipment of Lane 5. At the second location, we pick up the shipment of Lane 1 and depart for the next location (transporting the shipments of both Lanes 1 and 5). The next location has coordinates (30.05; 19.80) and is 9 km away. Given this information, we update the KPIs: (1) the cumulative total volume increases by 50, (2) the cumulative total distance increases by 9 km, (3) the cumulative total tkm increases by 752, (4) the cumulative shared volume increases by 80, (5) the cumulative shared distance increases by 9 km, and (6) the cumulative shared tkm increases by 752. After all locations have been processed, we have obtained the KPIs and can calculate the shared volume, distance, and tkm ratio.

Note that if a return trip is made, the route may be split into two paths, and the transported volume is also split. For example, consider the *pivot* location with coordinates (76.08; 72.73). At this location, the route splits into two paths (Path B and Path C). Path B returns to the first cluster, and Path C moves forward to the fourth cluster before also returning to the first cluster. All shipments that are still on transport at the *pivot* location are to be dropped off on the way to the first cluster (i.e., Shipment 14). No shipments on the transport are bound for the fourth cluster. Therefore, all volume is assigned to Path B; no volume is assigned to Path C.

References

Bentley JL (1975) Multidimensional binary search trees used for associative searching. *Commun. ACM* 18(9):509–517.

Bentley JL, Friedman JH (1979) Data structures for range searching. *ACM Comput. Surveys* 11(4):397–409.

Boute R, Robberecht R, Mostinckx M, Verstrepen S (2011) Car pooling for cargo. Kersten W, Blecker T, Jahn C, eds. *Proc.* 2011 *Internat. Supply Chain Conf.* (Books on Demand, Hamburg, Germany).

Chisman JA (1975) The clustered traveling salesman problem. *Comput. Oper. Res.* 2(2):115–119.

Cruijssen F, Cools M, Dullaert W (2007b) Horizontal cooperation in logistics: Opportunities and impediments. *Transportation Res. Part E Logist. Transportation Rev.* 43(2):129–142.

Cruijssen F, Dullaert W, Fleuren H (2007a) Horizontal cooperation in transport and logistics: A literature review. *Transportation J.* 46(3):22–39.

Doherty S, Hoyle S (2009) Supply chain decarbonization: The role of logistics and transport in reducing supply chain carbon emissions. Report, World Economic Forum, Geneva, Switzerland.

Guttman A (1984) R-trees: A dynamic index structure for spatial searching. Yormark B, ed. *Proc. 1984 ACM SIGMOD Internat. Conf. Management Data* (ACM Press, New York), 47–57.

Helsgaun K (2014) Solving the clustered traveling salesman problem using the Lin-Kernighan-Helsgaun algorithm. Technical report, Roskilde University, Roskilde, Denmark.

Jacobs K, Van Lent C, Verstrepen S, Bogen M (2014) CO^3 test project report: Horizontal collaboration in fresh and chilled retail distribution. Accessed October 5, 2016, http://www.co3-project.eu/wo3/wp-content/uploads/2011/12/CO3-Deliverable-Nestlé-Pepsico-STEF-case-study-1.pdf.

Lakemond R, Fookes C, Sridharan S (2013) Fast exact nearest neighbour matching in high dimensions using d-D sort. Accessed October 21, 2016, https://www.hindawi.com/journals/isrn/2013/405680/.

Leitner R, Meizer F, Prochazka M, Sihn W (2011) Structural concepts for horizontal cooperation to increase efficiency in logistics. *CIRP J. Manufacturing Sci. Tech.* 4(3):332–337.

Moro MM (2009) Range query. Liu L, Özsu MT, eds. *Encyclopedia of Database Systems* (Springer, New York), 2324–2325.

Padilla Tinoco S, Boute R, Creemers S (2015) Bundling of transport and its impact on inventories and cost allocation in horizontal supply chain partnership. Narayanan S, Bhageria R, eds. *26th Ann. POMS Conf.* (Production and Operations Management Society, Washington, DC).

Pan S, Ballot E, Fontane F (2013) The reduction of greenhouse gas emissions from freight transport by pooling supply chains. *Internat. J. Production Econom.* 143(1):86–94.

Saenz MJ (2012) Driving more efficient logistics networks through horizontal collaboration. *Supply Chain Management Rev.* (October 24), http://www.scmr.com/article/driving_more_efficient_logistics_networks_through_horizontal_collaboration/.

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

Uhlmann J (1991) Satisfying general proximity/Similarity queries with metric trees. *Inform. Processing Lett.* 40(4):175–179.

Vanovermeire C, Sörensen K (2014) Measuring and rewarding flexibility in collaborative distribution, including two-partner coalitions. *Eur. J. Oper. Res.* 239(1):157–165.

Verstrepen S, Jacobs K (2012) CO^3 test project report: Creation of an orchestrated horizontal collaboration for road bundling between 2 shippers. Accessed October 5, 2016, http://www.co3-project.eu/wo3/wp-content/uploads/2011/12/CO^3-D4.1-Test_Case_road_bundling_JSP_Hammerwerk-Executive-Summary.pdf.

Verstrepen S, Cools M, Cruijssen F, Dullaert W (2009) A dynamic framework for managing horizontal cooperation in logistics. *Internat. J. Logist. Systems Management* 5(3–4):228–248.

World Economic Forum (2015) Beyond supply chains: Empowering responsible value chains. Report, World Economic Forum, Geneva, Switzerland.

Yi K (2008) R-Trees. Kao MY, ed. *Encyclopedia of Algorithms* (Springer, New York), 800–802.

Verification Letter

Alexis Van Breedam, CEO TRI-VIZOR NV, Tri-Vizor, Galileilaan 15, 2845 Niel, Belgium, writes:

"With this letter I confirm that my company, TRI-VIZOR, uses the BBaRT tool for detecting new transport bundling opportunities for collaborative shipping. BBaRT allows us to efficiently search our large database of over 130,000 shipment lanes. Before BBaRT this database needed to be searched manually, which highly restricted our capabilities to fully explore the bundling opportunities in our growing database. The introduction of BBaRT has not only automated, but has also significantly speeded up the search for collaborative shipping opportunities. As a result, several transport bundling collaborations, suggested by BBaRT, have been initiated and even more are expected to follow in the (near) future."

Stefan Creemers is a professor at the IÉSEG School of Management, a visiting professor at KU Leuven, and a board member of PICS Belgium. He received his PhD in operations management from KU Leuven in 2009 and has published award-winning research in the fields of project management, queueing theory, and supply chain management. He is an area editor for *INFORMS Transactions on Education*.

Gert Woumans received his MSc as a business engineer in 2012 from the Hogeschool-Universiteit Brussel in Brussels, Belgium. He taught at the IÉSEG School of Management as a research and teaching assistant in the fields of operations management and strategy, and performed research as a PhD candidate at the Faculty of Economics at KU Leuven, Belgium. His research interests are optimization, scheduling, and horizontal collaboration.

Robert Boute is a professor of operations and supply chain management at Vlerick Business School and University KU Leuven. He was visiting professor at Kellogg School of Management (Northwestern University), Peking University (BiMBA), and ISM in St. Petersburg. He received the best paper award for his research on supply chain coordination and has won several best teachers' awards. His research focuses on global supply chain management and is often conducted in collaboration with companies.

Jeroen Beliën is a professor at the Faculty of Business and Economics at KU Leuven campus Brussels, Belgium. He received his MSc degree in business engineering from KU Leuven in 2001 and his PhD degree in applied economics from the KU Leuven in 2006. His research is mainly focused on combinatorial optimization and decomposition algorithms with applications in healthcare, maintenance, energy systems, and scheduling. Dr. Beliën is a board member of PICS Belgium and editor-in-chief of *INFORMS Transactions on Education*.