doi:10.1287/inte.2016.0878 • 🌐 www.stefancreemers.be • ✉ info@stefancreemers.be

Tri-Vizor uses an efficient algorithm to identify collaborative shipping opportunities

Stefan Creemers
Gert Woumans
Robert Boute
Jeroen Beliën

Abstract - **Collaborative shipping, where companies bundle their transport loads, is 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. This comes at the benefit of reduced logistics costs, as well as a lower carbon footprint. Although its advantages are clear, a major impediment is to find suitable collaboration partners. 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 where shipments are collected and/or dropped off en route. The tool is currently being used (among others) by Tri-Vizor, a facilitator and orchestrator of horizontal logistics partnerships, but is also applicable for any company that is looking for collaborative shipping partners. For Tri-Vizor, the tool has become an indispensable asset to detect collaborative shipping opportunities as their database has grown to over 130,000 shipment lanes.**

1 Introduction

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 the supply chain in order to reduce demand uncertainty for the suppliers. Horizontal collaborations are established between companies that operate at the same level in different supply chains, i.e., between suppliers or between buyers. Sharing transportation capacity when moving freight is an example of horizontal collaboration, an option that benefits the environment and yields substantial network efficiencies (Saenz, 2012). It is even possible that two "co-opetitors" set up a horizontal cooperation (European Union, 2001; Leitner et al., 2011). Horizontal partnerships in logistics have the potential to generate substantial gains by leveraging the overlaps in transport networks (Cruijssen et al., 2007a; Leitner et al., 2011). Whereas vertical collaborations have already been successfully established for many years, horizontal collaboration initiatives are more recent and are expected to become more widespread in the near future.

The bundling of freight is nothing new, since this is essentially what logistics providers do. When companies outsource their logistics to a logistics service provider, the provider can combine freight loads of their customer base if shipment timings are identical and if there is a geographical match. Collaborative shipping is different: opportunities are detected prior to shipment, and if desired, plans are changed and shipments are delayed or moved forward in order to benefit from joint transport. The consolidation is both in geography and in time, i.e., a shipment might be rescheduled if it creates synergies. More flexibility of each partner allows to exploit more opportunities for bundling, and allows to create better and cheaper distribution plans (Vanovermeire and Sörensen, 2014). Boute et al. (2011) report on the collaboration of two pharmaceutical companies, Baxter and UCB, where synergies are generated by flexible planning: Baxter has the possibility to postpone some of its orders, which frees up space for UCB who was shipping low volumes with a lower frequency. This contrasts with traditional freight groupage, which is mainly reactive: in groupage shipping, the logistics provider decides upon bundling LTL (less than container loads) in the execution phase, rather than in the planning phase, and the consolidation is only geographical, i.e., the timing of the shipments is not allowed to change.

In the most recent World Economic Forum report, collaborative shipping has been identified as one of the recipes to drive shared value (World Economic Forum, 2015). The European Commission reported that 27.3 percent of national road freight kilometers were empty hauls in 2010, and when carrying a load, vehicles are typically loaded for 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 helps 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 can be reduced, some harmful external effects directly related to road freight transport are also mitigated (greenhouse gas (GHG) emission, pollution, congestion, etc.). Freight consolidation across companies can also lead to increased scale effects, facilitating a modal shift. As multi-modal transport requires a certain volume to be economically viable, for small to medium-sized companies it is often not possible without setting 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 today, the potential of horizontal supply chain collaboration remains largely untapped. Establishing horizontal partnerships is also not straightforward. Even when companies are willing to cooperate, there are still many practical impediments. A survey by Cruijssen et al. (2007b) shows that finding suitable partners is seen 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 non-tangible aspects (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 at a close enough distance so that trucks/empty space can be shared. Our tool allows to identify all relevant collaborative shipping opportunities for a given company: (1) bundling transports that have roughly the same origin and destination, (2) using an empty back-hauling trip for another transport, or (3) avoiding 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 where shipments are collected and/or dropped off en route. We refer to our tool as "BBaRT": Bundling, Back-hauling, and Round-trip Tool. BBaRT has among others been implemented by the company Tri-Vizor.

Tri-Vizor is a facilitator and orchestrator of logistics horizontal collaboration partnerships. It identifies potential collaborative shipping partnerships and is in charge of 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, their database has become very large and thus very time consuming to analyze on a manual basis. BBaRT helps to automate the process and allows to quickly detect promising partnerships that are compatible with respect to cargo and routing.

In the remainder of this article, we first describe the problem in more detail, after which we present the working dynamics of BBaRT. We illustrate our methodology by means of an elaborate example. We conclude by summing up the benefits that are obtained by using our tool.

2 Setting up collaborative shipping partnerships

Verstrepen et al. (2009) developed a conceptual framework to set up and to maintain collaborative shipping partnerships. As soon as companies are aware of the need (and the benefits) of collaborative shipping, they start looking for potential collaboration partners. Our BBaRT tool operates in this stage and identifies potential partners based on their geographical compatibility. As soon as the collaboration partners are identified, the cooperation can be prepared: the planning and synchronization of the shipments, the choice of the joint carrier, the gain sharing, etc. The final stage is concerned with the effective implementation and operation of the collaboration using a control tower. Figure 1 illustrates this process and situates the role of BBaRT herein.

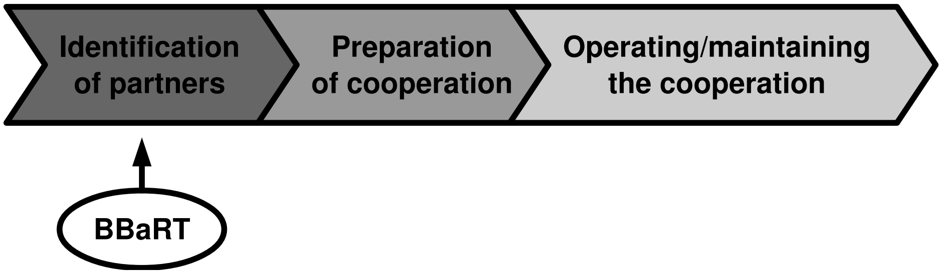

Figure 1: The BBaRT algorithm enables the identification of collaboration partners. The preparation (planning and synchronization of shipments) and the effective implementation of the cooperation happens in later stages.

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

Prior to the introduction of BBaRT, Tri-Vizor relied on Excel pivot tables based on country and region codes of the origin and destination. The shipments that had matching regional codes for origin and destination were grouped and were manually checked for feasible bundling combinations. As the number of shipment lanes in their database grew over time, using Excel pivot tables became cumbersome and too time consuming. Our tool reduces the manual efforts immensely by automating the checking of the geographical compatibility. BBaRT allows to identify not only more, but also better partnerships (with more partners, more route overlaps, etc.).

BBaRT detects three different 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-c) show a graphical representation of these 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 origin and the destination of two shipments lie within a radius r of each other, and vice versa.
- For round-trip opportunities, we require that the destination and the origin are within a distance r of the respective origin and destination of other partnering lanes.

BBaRT also detects more complex opportunities where multiple stops and/or return trips are made. An example of such an opportunity is illustrated in Figure 2d. In addition to that, BBaRT also identifies collect-and/or-drop opportunities in which shipments are collected and/or dropped off en route (note that collect-and/or-drop opportunities can only be found in combination with a bundling, back-hauling, and/or round-trip opportunity). BBaRT detects collect-and/or-drop opportunities that can be found if a shipment can be collected and/or dropped off at a location which is within a distance r_2 of an existing route. Figures 3(a-c) illustrate simple collect-and/or-drop opportunities. BBaRT, however, can also detect more complex collect-and/or-drop opportunities, such as the one provided in Figure 3d. Radius r and r_2 are defined by the user: the smaller the radius, the less detour is needed to accommodate the collaborative shipping. As the radius increases, the number of proposed collaborative shipping opportunities will also go up accordingly.

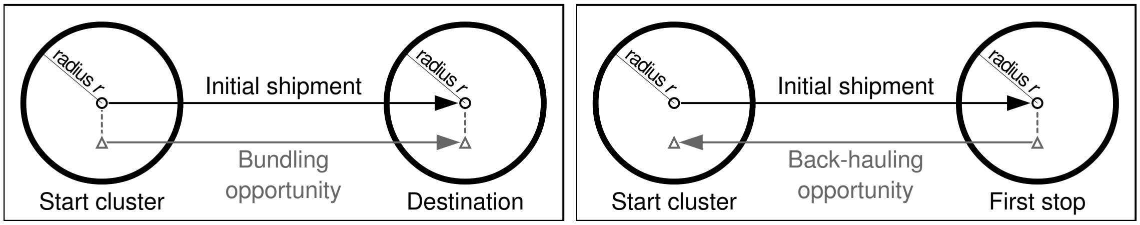

(a) Bundling opportunity (b) Back-hauling opportunity

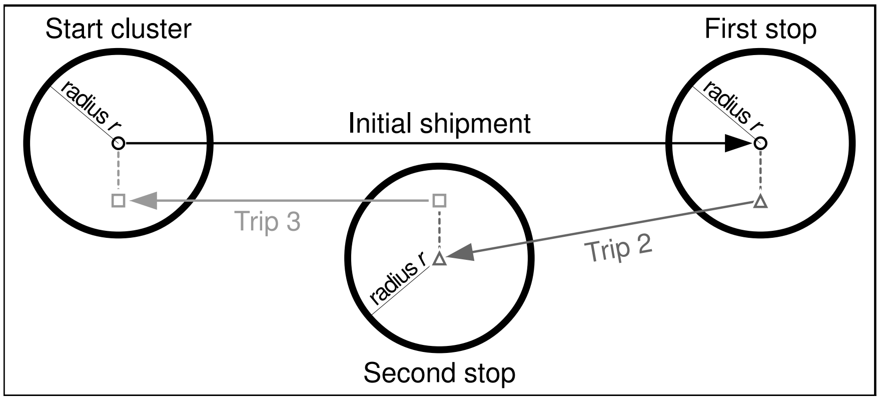

(c) Round-trip opportunity

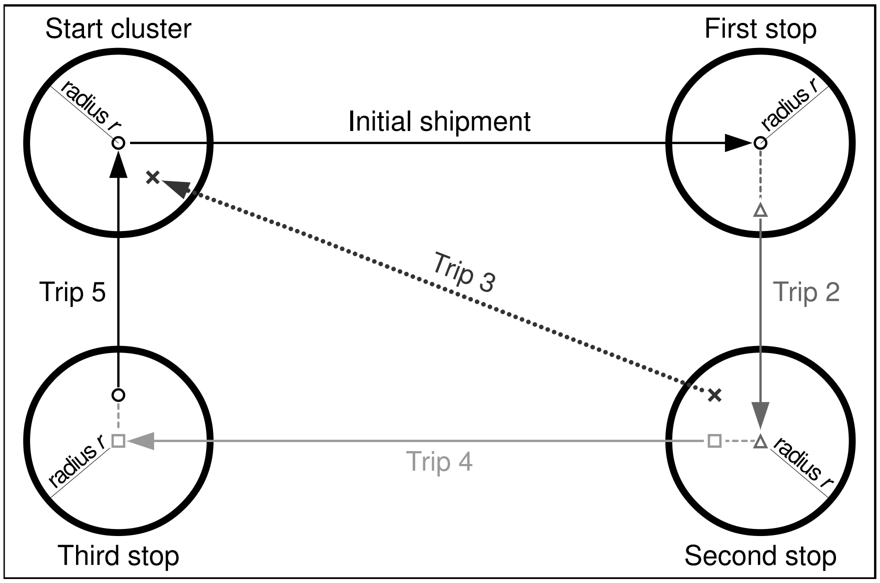

(d) Complex round-trip opportunity

Figure 2: BBart identifies bundling, back-hauling and round-trip opportunities where the origin/destination points of different trips are within a small distance (r) of one another.

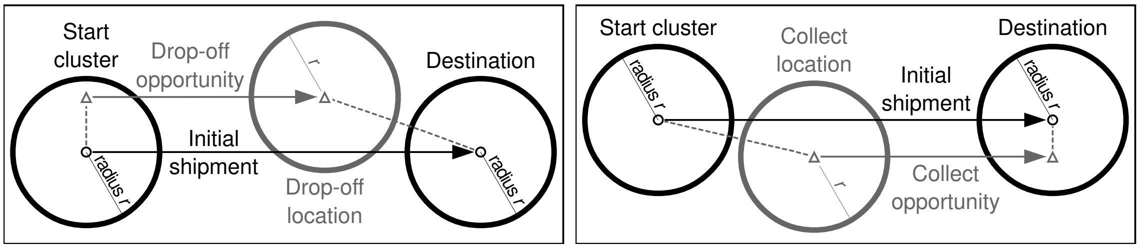

(a) Drop-off opportunity (b) Collect opportunity

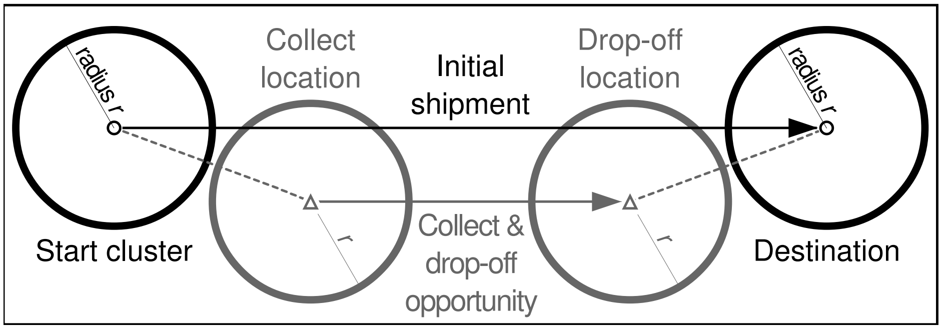

(c) Collect-and-drop-off opportunity

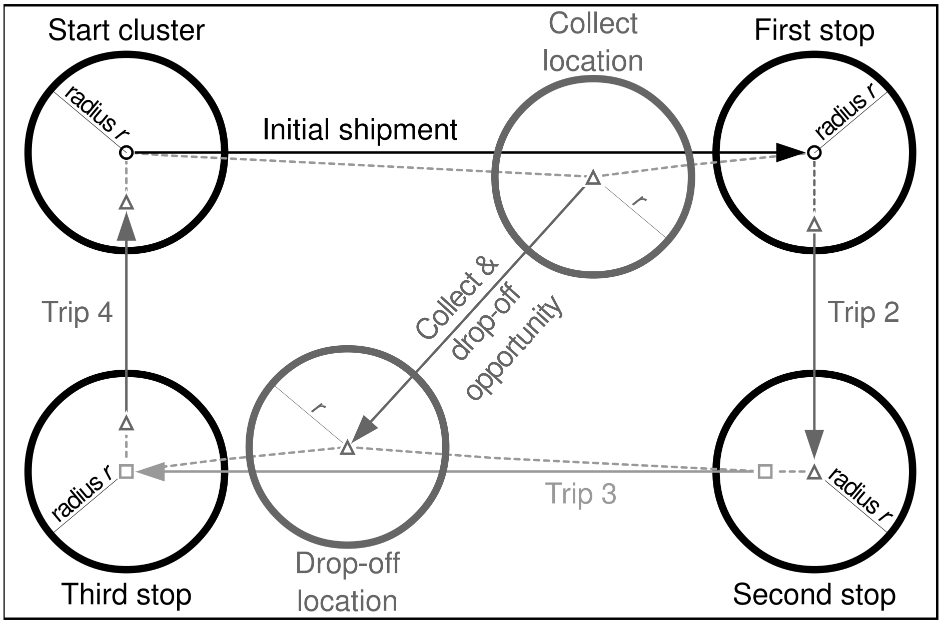

(d) Complex collect-and-drop-off opportunity

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. In order to collect/drop off a shipment, its origin/destination needs to be in a distance r_2 of an existing route.

In essence, the problem boils down to the identification of shipments with similar origins and destinations. More formally, we are looking for neighbors in a multidimensional space. Every shipment has four coordinates making it 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 (see Bentley (1975), Bentley and Friedman (1979), Guttman (1984), Uhlmann (1991), Yi (2008), Moro (2009), and Lakemond et al. (2013) for more details). In this article, we adopt a different approach. The rationale why we choose for a different approach is threefold:

- A sorted list of lanes (even if they are sorted on multiple dimensions) does not allow the quick detection of collaborative shipping opportunities. In order to determine whether or not two lanes can be bundled, their geographical compatibility needs to be assessed (i.e., distances need to be calculated). As such, sorting in itself is not sufficient – sorting is only part of the solution. To quickly detect bundling opportunities, BBaRT combines sorting with a bounding-box approach (see also infra). 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.

- In order to detect collect-and/or-drop opportunities, we calculate the rotated coordinates of the origin and the destination of each lane (see also infra). A list of lanes that is sorted on the non-rotated coordinates is of no use here. Again, we need to combine sorting with a bounding-box approach.

- Sorting in itself does not allow to find round trips or more complex collaborative shipping opportunities with two or more stops. In order to detect these opportunities, a fast queue-based search algorithm is required (see also infra).

3 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. In what follows, we discuss each of these steps and illustrate with a small sample database. For each shipment lane, the sample database contains the coordinates of origin and destination, the yearly transport quantities (in truck equivalents), the distance between origin and destination, and the ID of the company that uses the shipment lane.

3.1 Data preparation

First, we re-organize the shipment data into a useable data structure. All shipments that have identical origin and destination – for instance, 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 does not have to be taken into account, as identical timings of the shipments are not necessarily required in order 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/transport frequencies and the timing of their shipments). By preparing the data in this manner, the number of candidate lanes can be reduced (and it prevents generating obvious bundling opportunities between identical lanes), without losing any information. Table 1 illustrates the result of this data-preparation step.

Lane	Truck equivalents	Trip distance	Origin		Destination	
			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

Table 1: We aggregate the data in the shipment database to obtain a useable data structure with candidate lanes that have unique origin and destination.

Subsequently, the dataset is sorted on the latitudinal coordinate. Sorting on one dimension only allows to quickly lookup candidate lanes. A bounding-box approach can then be used to filter the candidate lanes based on their geographical proximity.

3.2 Identification of collaborative shipping opportunities

Figure 4 plots the different lanes in Table 1 on a map. As can be seen in Figure 4, there are four clusters with shipments that have similar origin and destination. We illustrate the use of our algorithms by applying it to our sample database.

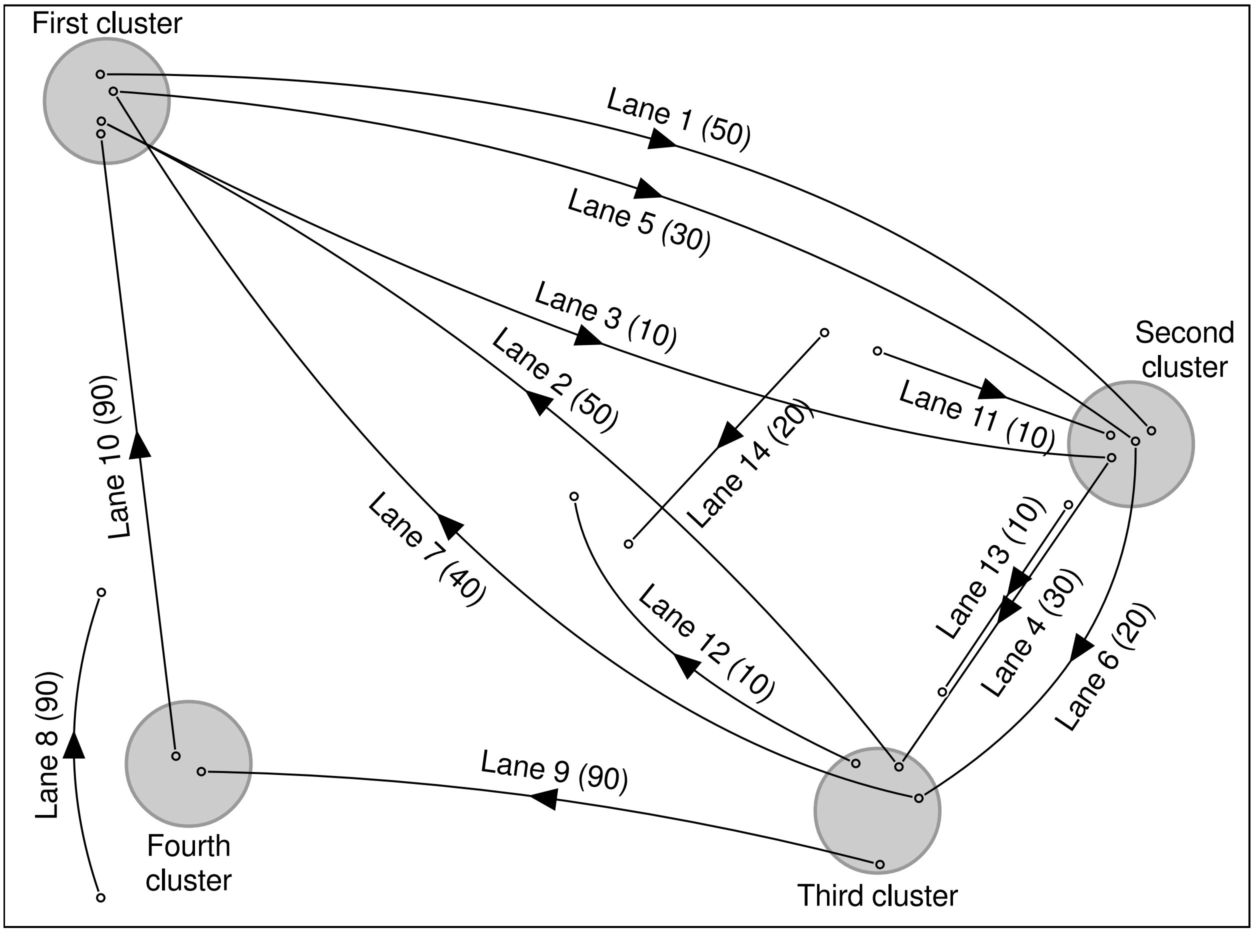

Figure 4: The BBart tool clusters data in Table 1 to identify a number of collaboration opportunities including bundling (e.g., lane 1 and lane 5), round-trip (e.g., lanes 5, 6, and 7), and collect-and/or-drop opportunities (e.g., lanes 4 and 13).

In order 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/destination of a lane). The sorted list of lanes allows to quickly lookup potential candidate lanes. However, because the lanes are sorted on the latitudinal coordinate only, it is possible (and even likely) that some of these candidate lanes are in fact not close to point P at all. To filter those lanes that are close to point P, we use a bounding-box approach. For each lane that arrives in/departs from a location in the bounding box for point P, we calculate the distance towards point P. We retain only those lanes for which the distance is smaller than r, with r a prespecified parameter by the user denoting the maximal detour that you are willing to accommodate in order to benefit from joint transport. All other lanes can be discarded. Note that calculating the pairwise distance between each of the lane coordinates to determine whether or not the origin/destination of a lane is in the proximity of the origin/destination of another lane, would not be a practical (nor feasible) solution if the number of lanes is large. In the case of Tri-Vizor for instance, it would not be possible to calculate/store the distances between the origins/destinations of the 130,000 shipment lanes in their database.

Whereas the above approach allows to identify bundling, back-hauling, and round-trip opportunities, detecting collect-and/or-drop opportunities requires a slightly different approach. Take for example lanes 4 and 13. From Figure 4, we can see that there is 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/destination of lane 4 is of no use. In fact, we need a bounding box for lane 4 itself. Unfortunately, however, this bounding box is not perpendicular to the coordinate system, and therefore it is no longer possible to quickly determine what 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. In order to lookup 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 or not a candidate lane is within a distance r_2 of lane 4 itself (where r_2 is specified by the user). This process is illustrated in Figure 5. From Figure 5b, we can see that the shortest distance between the origin of lane 13 and lane 4 itself equals 1.6 km (i.e., 102.49 minus 102.33). The shortest distance between the destination of lane 13 and lane 4 itself amounts to 2.9 km. Therefore, if r_2 is set larger than 2.9 km, lane 13 is detected as a collect-and/or-drop opportunity for lane 4.

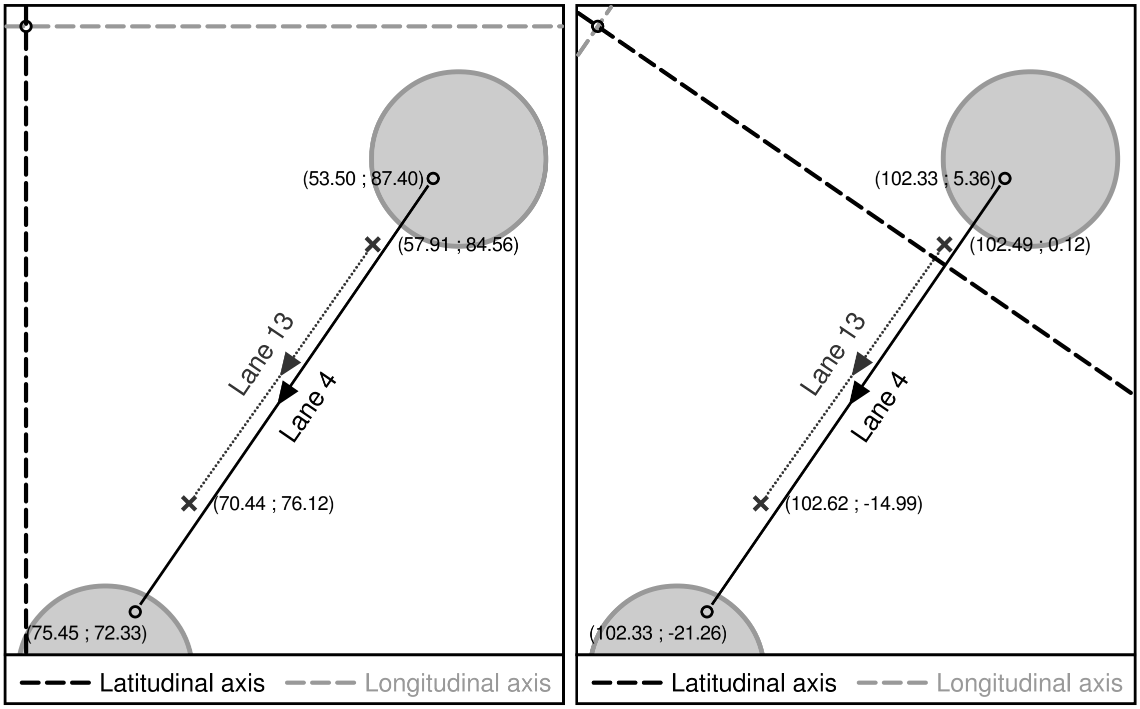

(a) No rotation (b) Rotated coordinate system

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.

The combination of a sorted list and a bounding-box approach proves to be of critical importance in order to determine bundling, back-hauling, round-trip, and collect-and/or-drop opportunities. In order to filter out the opportunities with the largest saving potential (in terms of transport cost and carbon footprint), we rank them, which is discussed next.

3.3 Ranking of opportunities

The output of the algorithms is a list of potential partnerships. In order to evaluate these collaborative partnerships, we assess the economical and environmental benefits of each collaborative shipping opportunity. As the financial and environmental cost for each transport lane is represented on a per-volume, per-distance basis, 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, as well as the number of shipments that are bundled over this shared distance, and thus 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 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 tonne-kilometers (tkm) (volume times distance that is traveled by that volume) of the identified opportunity.
- The total number of *shared* tkm (combining both 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 over total distance (representing the overlap of the lanes).
- The ratio of shared volume over total volume (the volumes shipped over the shared distance compared to the volumes shipped over the total distance).
- The ratio of shared tkm over total tkm.

Typically the total number of *shared tkm* will be most relevant, as 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 instance, a user can be interested in bundling shipments only over long distances, or rather bundling high volumes over shorter distances. This way, the potential partnerships can be ranked in line with the preferences of the user.

In order 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) (see Chisman, 1975). The CTSP is an extension of the Traveling Salesman Problem (TSP) where the 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 in order 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, since we are not solving the routing problem at 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).

The best collaborative bundling opportunity in our sample database visits 4 clusters and comprises 13 lanes. Table 2 reports the details of all KPIs for this collaborative opportunity. The collaborative shipping goes 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 first drop off the shipments of lanes 4 and 6. At this point, the routing splits in 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 2,674 km, 1,723 km were shared between two or more lanes, meaning that 64.43% of the distance is shared. In total, we ship 460 truck equivalents of which 280 over the shared distance, i.e., 60.87% of the volume is shared. Last but not least, the total number of tkm equals 246,054, of which 66.35% is shared (163,245 tkm are shared). The bulk of the non-shared transport originates from lanes 9 and 10. Therefore, if the user wants to increase the shared ratios, he/she can choose to omit those lanes from the collaborative shipping opportunity.

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	1o	4	0	30	30	0	30	4	30	128	0	0	0
1o	3o	9	30	50	80	0	80	14	80	881	9	80	752
3o	14o	511	80	10	90	0	90	525	90	46890	521	90	46762
14o	11o	37	90	20	110	0	110	562	110	51013	558	110	50884
11o	11d	165	110	10	120	10	110	728	120	70834	723	120	70706
11d	3d	5	110	0	110	10	100	732	120	71330	728	120	71201
3d	5d	6	100	0	100	30	70	738	120	71912	734	120	71784
5d	1d	4	70	0	70	50	20	742	120	72180	737	120	72052
1d	6o	4	20	0	20	0	20	746	120	72257	737	120	72052
6o	4o	6	20	20	40	0	40	751	140	72490	743	140	72285
4o	13o	52	40	30	70	0	70	804	170	76161	796	170	75956
13o	13d	151	70	10	80	10	70	955	180	88247	947	180	88042
13d	4d	63	70	0	70	30	40	1018	180	92645	1010	180	92440
4d	6d	7	40	0	40	20	20	1025	180	92943	1017	180	92738
Path B (first return trip)													
6d	7o	15	20	0	20	0	20	1041	270	92943	1017	180	92738
7o	2o	467	20	40	60	0	60	1507	310	134937	1017	220	92738
2o	12o	0	60	50	110	0	110	1507	360	134937	1017	270	92738
12o	14d	7	110	10	120	20	100	1515	370	135385	1025	280	93186
14d	12d	20	100	0	100	10	90	1535	370	137600	1045	280	95401
12d	2d	211	90	0	90	50	40	1746	370	162966	1256	280	120767
2d	7d	48	40	0	40	40	0	1795	370	167806	1304	280	125607
Path C (second return trip)													
6d	9o	418	0	0	0	0	0	2213	180	205444	1723	180	163245
9o	9d	6	0	90	90	90	0	2219	270	205703	1723	180	163245
9d	10o	6	0	0	0	0	0	2225	370	205703	1723	280	163245
10o	10d	448	0	90	90	90	0	2674	460	246054	1723	280	163245
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													

Table 2: The routing of the example bundling opportunity consists of three parts: a common path A, a path B, and a path C. 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.

4 Use of BBaRT

The BBaRT tool is used by Tri-Vizor since October 2013. For Tri-Vizor, BBaRT has become an indispensable asset to identify new collaborative shipping opportunities and to bring potential collaboration partners together. As the number of shipment lanes in their database grew over time, and the networks grew in complexity, it became more difficult and time-consuming to detect bundling or back-hauling opportunities on a manual basis using Excel pivot tables; finding round trips or collect-and/or-drop opportunities was even not possible using pivot tables.

When a new company contacts Tri-Vizor to search potential partnerships, BBaRT is used to match their lanes with the current database of over 130,000 shipment lanes and a long-list of collaborative shipping opportunities is provided in no time. Since the introduction of BBaRT, at least one or two collaborative shipping opportunities per month have been identified and proposed based on the recommendations made by BBaRT.

One example of such an opportunity that could not have been detected without BBaRT, is the regional freight flow bundling in the Province of Gelderland in the Netherlands, where 13 companies (mix of SMEs and multinationals) showed interest to bundle transports. The user input consisted of an Excel template collecting all structural FTL/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, etc.). These data were collected under a non-disclosure agreement. The collected flow data were cleaned and converted into standard format and a geo-coding was performed of the origin/destination locations (i.e., conversion to XY-coordinates) to enable further automatic processing. The shipper data were then merged into 1 project database with 4,512 transport lanes and 2,426 unique origin/destination locations.

Synergies were analyzed internally (between the 13 Gelderland shippers) and externally (with Tri-Vizor's entire database). The BBaRT tool detected ca. 1,000 internal collaborative shipping opportunities between the Gelderland shippers (bundling + round trips), of which 14 were filtered out as opportunities with high potential (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 ca. 4,500 collaborative shipping opportunities, of which 300 with a transport frequency of more than once per week. From this set, 16 interesting combinations were retained using BBaRT. This was the basis to start a discussion with the companies involved (on the synchronization of the shipments, the joint cost drivers, the choice of the joint transport company, the transparent sharing of the gains, etc.).

The BBaRT tool has also been used by a major Belgian retailer 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 the number of shipment lanes grows, the identification of collaborative shipping opportunities is an impossible task to do manually.

Note that BBaRT is used in the first stage when setting up a collaborative shipping partnership, i.e., the identification of potential bundling partners (see also Figure 1). The further optimization/implementation of the bundling itself is done in later stages. For instance, how will the shipments be synchronized (who will drive the joint "rhythm"), how will the planning occur (fixed weekly departure days vs. dynamic planning), choice of the joint carrier, how will the gains be shared among the collaboration partners, etc. are subject to further discussion after identification of the potential partners. It goes without saying that at this stage the collaboration may seize to exist. However, without proper identification, the discussion would never have been initiated in the first place.

It is expected that the collaboration gains found using BBaRT will exceed previous pilots identified by Tri-Vizor, where only two companies bundled their transport. We refer to the reporting of collaborative shipping agreements with two companies to the bundling of road transport by the manufacturing companies JSP and Hammerwerk, and the horizontal collaboration in fresh & chilled retail distribution between the FMCG shippers Nestle & PepsiCo, resulting in $10-15\%$ transport cost savings and a $20-30\%$ reduction in CO_2 emissions (see Verstrepen and Jacobs 2012, and Jacobs et al. 2014 for a summary report of these pilots). When the number of participating companies is larger than two, it is expected that the gains may be even higher, as there are more collaborative shipping opportunities to exploit and thus even more savings to reap.

5 Directions for future research

Although BBaRT in its current version has proven to be an extremely valuable tool for Tri-Vizor, we see several directions for future research. First of all, BBaRT is currently limited to finding bundling opportunities that are geographically compatible. Future research should focus on extending BBaRT to also include cargo compatibility, order frequencies, and other important characteristics. At this moment, the participating companies do not provide such information, but we strongly 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. It would be interesting to examine more advanced methods of multi-objective optimization, for instance, finding Pareto efficient bundles. Thirdly, BBaRT is limited to the identification of collaborative shipping opportunities. After the identification phase, the implementation of collaborative shipping opportunities at the operational level requires to model and solve 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 is identified by BBaRT. This would enable BBaRT to go beyond the identification phase and support implementation at the operational level.

References

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

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

[3] 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).

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

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

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

[7] Doherty S, Hoyle S (2009) Supply Chain Decarbonisation: The role of logistics and transport in reducing supply chain carbon emissions. World Economic Forum, Geneva, Switzerland.

[8] Guttman A (1984) R-trees. Yormark B, eds. *Proc. 1984 ACM SIGMOD Internat. Conf. on Management of Data* (ACM Press, New York, NY, USA), 47–57.

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

[10] Jacobs K, Van Lent C, Verstrepen S, Bogen M (2014) Horizontal collaboration in fresh & chilled retail distribution. www.co3-project.eu.

[11] Lakemond R, Fookes C, Sridharan S (2013) Fast exact nearest neighbour matching in high dimensions using -D Sort. *ISRN Mach. Vis.* 2013:1–8.

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

[13] Moro MM (2009) Range query. Liu L, ÖZSU MT, eds. *Encyclopedia of Database Systems* (Springer US), 2324–2325.

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

[15] 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., *2015 Product. and Oper. Management Society Annual Conf.* (POMS, Washington, DC).

[16] Saenz MJ (2012) Driving more efficient logistics networks through horizontal collaboration. *Supply Chain Management Rev.*, October.

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

[18] Uhlmann J (1991) Satisfying general proximity/similarity queries with metric trees. *Inf. Process. Lett.* 40(4):175–179.

[19] 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.

[20] Verstrepen S, Jacobs K (2012) Creation of an orchestrated horizontal collaboration for road bundling between two shippers. www.co3-project.eu.

[21] Yi K (2008) R-Trees. Kao MY, eds. *Encyclopedia of Algorithms* (Springer US), 800–802.

[22] World Economic Forum (2015) Beyond Supply Chains: Empowering responsible value chains. World Economic Forum, Geneva, Switzerland.

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

Appendix A

In what follows, we explain how to obtain sets of matching lanes (i.e., geographically compatible lanes). Next, we 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 itself.
- $\mathbf{M}_i^{LD}$: the set of lanes that has its destination within a distance r_2 of lane i itself.

Using these sets of matching lanes, it is possible to 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 that is rotated with t degrees (i.e., $\mathbf{A}_0^O$ is the sorted list of original coordinates). $\mathbf{A}_t^D$ is defined analogously, however, 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 with 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. In order 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., $\mathbf{L}_i^C$) that may have their 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. In order to be sure 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 $\mathbf{A}_0^O$ to obtain set of candidate lanes $\mathbf{L}_i^{COO}$ that may have origin close to the origin of lane $i$
  for All lanes $j$ in $\mathbf{L}_i^{COO}$ 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 $\mathbf{L}_i^{FOO}$ that may have origin close to the origin of lane $i$
      end if
    end if
  end for
  for All lanes $j$ in $\mathbf{L}_i^{FOO}$ 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 $\mathbf{M}_i^{OO}$
    end if
  end for
end for
```

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

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

When rotating the coordinate system with 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. In order to obtain sets $\mathbf{M}_i^{LO}$ and $\mathbf{M}_i^{LD}$, we iterate over all integer degrees between 0 and 360. For every degree: (1) we determine the rotated coordinates of origin/destination for each lane, (2) we sort all lanes on the rotated latitudinal coordinate of the origin/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. In order to detect the collect-and/or-drop opportunities for lane i, we apply a similar logic as the one that is used 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/destination. Note that we rotate with 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. In order 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 with 1 degree is 1.75 percent). When rotating with t degrees, the set of candidate lanes contains all lanes that have latitudinal coordinate of origin/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/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 origin/destination within the inflated bounding box around lane i itself. In order to be sure that the origin/destination of a filtered candidate lane j is within a distance of r_2 km from lane i itself, we need to rotate the latitudinal coordinate of the origin/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/destination of the filtered candidate lanes and lane i itself. 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 $X_{i,t}^O$ and $Y_{i,t}^O$ of the origin of lane $i$.
  end for
  Obtain $\mathbf{A}_t^O$ by sorting all lanes on their latitudinal coordinate of the origin that is rotated with $t$ degrees.
  for All lanes $i$ do
    if $\lfloor\theta_i\rfloor = t$ then
      Calculate rotated coordinate $Y_{i,\theta_i}^O$ of the origin of lane $i$
      Use sorted list $\mathbf{A}_t^O$ to obtain set of candidate lanes $\mathbf{L}_i^{CLO}$ that may have origin close to lane $i$ itself
      Use inflated bounding box to obtain filtered set of candidate lanes $\mathbf{L}_i^{FLO}$ that may have origin close to lane $i$ itself
      for All lanes $j$ in $\mathbf{L}_i^{FLO}$ do
        Calculate rotated coordinates $Y_{j,\theta_i}^O$ of the origin of lane $j$
        if $|Y_{i,\theta_i}^O - Y_{j,\theta_i}^O| < r_2$ then
          Add lane $j$ to set of matching lanes $\mathbf{M}_i^{LO}$
        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 that each contain a single lane. Next, we will process and generate new queue elements until all collaborative shipping opportunities have been evaluated. Note that each queue element has an active cluster that contains a single lane that 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 $\mathbf{M}_i^{DD}$). For instance, in the example collaborative shipping opportunity, lane 1 is bundled with lanes 3 and 5 (see Figure 4). 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 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, on the other hand, the maximum number of clusters has been reached, no more lanes can be added, and the following steps need to be performed:

- 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/lane in the bundling opportunity and its destination is close to another location/lane that is visited afterwards.
- Determining 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 already recorded opportunities, and (2) if it is sufficiently different from already recorded opportunities. In order to determine whether the opportunity is sufficiently different, we evaluate the number lanes that are shared between recorded opportunities. The user can specify that a maximum percentage of lanes are shared. As a result, each of the recorded opportunities has at most a given percentage of shared lanes.

Calculation of KPIs of collaborative shipping opportunities

In order to calculate the KPIs, we first need to 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. *Return* clusters, on the other hand, are clusters that have already been visited. Note that, upon arrival at a cluster, shipments on transport are first dropped off before any new shipments are collected.

In order to calculate the KPIs, we keep track of the shipments that arrive and depart at every location. This way, we can update the KPIs at every location. For instance, consider the second location in the example collaborative shipping opportunity (see Table 2). The second location has coordinates $(29.11; 19.78)$, and we arrive there after having 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 shipments of both lanes 1 and 5). The next location has coordinates $(30.05; 19.80)$ and is 9 km away. Given this information, the KPIs are updated:

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
```

(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 different paths, and the transported volume is split as well. Consider for instance 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 transport are bound for the fourth cluster. As a result, all volume is assigned to path B whereas path C is not assigned any volume at all.

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 re-entry location of the return cluster

        Start from the re-entry location and use closest-neighbor heuristic to drop off all shipments that are bound for the return cluster
    end for
end for
```