Combining pickups and deliveries in vehicle routing – An assessment of carbon emission effects

Marcel Turkensteen[(1)], Geir Hasle[(2)]

[(1)] CORAL, Department of Economics and Business Economics,
School of Business and Social Sciences, Aarhus University,
Fuglesangs Alle 4, 8210 Aarhus V, Denmark
matu@econ.au.dk

[(2)] Mathematics and Cybernetics, SINTEF Digital,
P.O. Box 124 Blindern, NO-0314 Oslo, Norway
geir.hasle@sintef.no

Abstract

This paper studies the effect on carbon emissions of consolidation of shipments on trucks. New positioning and communication technologies, as well as decision support systems for vehicle routing, enable better utilization of vehicle capacity, reduced travel distance, and thereby carbon emission reductions. We present a novel carbon emission analysis method that determines the emission savings obtained by an individual transport provider, who receives customer orders for outbound deliveries as well as pickup orders from supply locations. The transport provider can improve vehicle utilization by performing pickups and deliveries jointly instead of using separate trucks. In our model we assume that the transport provider minimizes costs by use of a tool that calculates detailed vehicle routing plans, i.e., an assignment of each transport order to a specific vehicle in the fleet, and the sequence of customer visit for each vehicle. We compare a basic set-up, in which pickups and deliveries are segregated and performed with separate vehicles, with two consolidation set-ups where pickups and deliveries may be mixed more or less freely on a single vehicle. By allowing mixing, the average vehicle load will increase and the total driven distance will decrease. To compare carbon emissions for the three set-ups, we use a carbon assessment method that uses the distance driven and the average load factor. An increase in the load factor can reduce part of the emission savings from consolidation. We find that emission savings are relatively large in case of small vehicles and for delivery and pickup locations that are relatively far from the depot. However, if a truck visits many demand and supply locations before returning to the depot, we observe negligible carbon emission decreases or even emission increases for consolidation set-ups, meaning that in such cases investing in consolidation through joint pickups and deliveries may not be effective. The results of our study will be useful for transport users and providers, policymakers, as well as vehicle routing technology vendors.

Keywords: Pickup and Delivery, Consolidation, Carbon emissions

1. Introduction

Freight transport has a considerable negative environmental impact in the form of local pollutants, such as particular matter, and global pollutants, such as greenhouse gases. Worldwide, the movement of freight is responsible for about 10% of energy related carbon emissions; In domestic situations, much of these emissions are due to road transportation (92% in the United Kingdom) (McKinnon et al., 2015). One way to reduce these emissions is by using existing vehicle capacity better: McKinnon et al. (2015) report that around 20 to 25% of the hauls are performed with empty vehicles and that the average degree of vehicle utilization in the European Union ranges from 28% in Ireland to 45% in Denmark.

Our study is motivated by the current aim of reducing carbon emissions through more efficient routing in freight transport. Léonardi and Baumgartner (2004) note that four freight transport efficiencies should be improved in order to reduce emissions: routing, logistics, driving and vehicle. Routing efficiencies refer to the route that the vehicles follow with the implication that a shorter route yields lower emissions. Logistic efficiencies refer to capacity utilization on-board a vehicle with the implication that greater load consolidation will reduce emissions. Driving efficiencies refer to the way the vehicle is driven in terms of speed and idling. Finally, vehicle efficiencies relate to the design of the vehicle itself in terms of fuel efficiency or alternative technologies. In this paper, we focus on the interrelated routing and logistics efficiencies. Specifically, we study the routing and load consolidation implications on carbon emissions in two pickup and delivery scenarios - one in which the vehicles perform all deliveries prior to commencing the pickups and one in which deliveries and pickups can occur freely.

There are studies on the effect of improving vehicle utilization on a large scale, e.g., in a region or a country, as in Walnum and Simonsen (2015). Such studies typically measure fuel usage or carbon emissions and relate the observed levels to input variables such as the load on the vehicle and the speed. However, it is then difficult to disentangle the improvements from better vehicle utilization from improvements in vehicle technology or routing tools. In our approach we model the decisions of an individual transport provider and determine the carbon emission savings resulting from consolidation of shipments. This allows us to assess the impact of vehicle utilization in isolation.

We can draw on techniques from the field of *green logistics*, where one minimizes fuel consumption or carbon emissions, often alongside other objectives such as costs; see Demir et al. (2014a). Even though we do not optimize carbon emissions or fuel consumption, we compare the emissions of set-ups with different load factors, and need a similar method for computing emissions. In this paper, we use the simple but reasonable carbon emission computation method from Turkensteen (2016b), see Section 3.

If shipments are consolidated in such a way that the vehicle route does not change, vehicle utilization is increased. One consequence is that we need fewer hauls to serve the transport demand, thus reducing the distance driven. However, the effect of higher average payload will counteract the resulting carbon emission savings. We call this the *payload effect*.

An interesting side-effect occurs when consolidation is performed in such a way that new locations are added to the original route. In that case, the average distance traveled by each item can increase, for example because items destined to the end of a route have to go on a detour through locations that are not in the set-up without consolidation. In such cases, the emissions due to the payload on the vehicle can increase. We call this the *detour effect*.

As indicated above, we consider a specific case of consolidation, namely the possible combination of deliveries from a depot and pickups destined to the same depot on the same vehicle. This case is interesting for two reasons. First, both the payload and the detour effect can occur: When pickup locations are added to a delivery route, the items to be delivered may travel over a longer distance. The same applies for the addition of delivery locations to a pickup route. Second, there are several situations where combined deliveries and pickups on the same vehicle are attractive in practice.

In an overview, McLeod et al. (2008) focus on urban areas in the United Kingdom, mainly of the collection of packaging waste. They mention that retailers such as ASDA, Sainsbury, and Next use trucks to return recyclable packaging materials. Anily and Federgruen (1990) argue that grocery stores have discovered cost-cutting potential by allowing vehicles to collect large volumes of inbound materials on their delivery routes. In forestry, studies such as Carlgren et al. (2006) consider the usage of a return haul from a factory to a forest to carry a load in the opposite direction for another forestry company.

Another interesting application arises from a trend towards the 'circular economy' where materials and products are reused (CircularEconomy.com, 2016). One manifestation of this could be that a company sells a service, e.g. clothing or printing, to companies rather than physical products such as clothes or printers. The company would then be responsible for replacing or refilling products at given points in time. A well-known example is Eastman Kodak (Krumweide and Sheu, 2002). In clothing, the upcoming Danish company Vigga.us leases children's clothes to customers, takes them back after usage, and replaces them with larger size clothes (Vigga.us, 2015). The examples from McLeod et al. (2008) fall in the same category. In all these cases, used materials or products could be collected by the vehicles that perform the deliveries.

In order to perform our analysis we construct a model of the decisions taken by the transport provider. This model and its underlying assumptions are presented in Section 2. We consider two consolidation options, namely *backhauling*, a set-up in which all deliveries should take place before any pickup, and *mixing*, a set-up in which deliveries and pickups can be mixed freely, as long as the vehicle's capacity is not exceeded. We analyze the carbon emission effects of the three different levels of flexibility regarding consolidation. Through computational experiments on a diverse set of instances, we consider the effect of different characteristics of the situation, such as the number of delivery and pickup locations and their distribution in an area.

The environmental effect of combining deliveries and pickups has not been given much attention in the literature. To the best of our knowledge, the only study that provides numerical results on carbon emission savings is the one by Ubeda et al. (2011). The authors consider the case of backhauling within a case company during one week, and compare this to a (current) set-up with separate delivery and pickups and the set-up with the lowest cost, as well as a set-up with minimal carbon emissions. It is found that the backhauling option yields about 15% less carbon emissions than the current set-up and around 5% less than the set-up that minimizes total costs. A more integrated set-up with mixing is not considered, and the results are confined to a single case study, with a single vehicle. As far as we know, no general study of carbon emission effects of combining deliveries and pickups has been reported.

The results of our study can be of use for companies and policymakers. A company or policymaker may wish to consider investments in the tools and vehicles necessary to consolidate inbound and outbound shipments, partially to reduce carbon emissions or fuel usage. Our results can be used to determine the case where such investments would be most successful. If it turns out that the emission savings are minimal in a given situation, e.g. for transport in a rural area, investments could be redirected to new vehicle technology or routing decisions: In fact, we show that there may be distribution situations where such consolidation can lead to emission increases. Further, providers of vehicle routing tools (Bräysy and Hasle, 2014) may utilize our results in the future to enhance their products with better functionality for assessing carbon emission effects.

The rest of the paper is organized as follows. Section 2 describes our model of the distribution situation with deliveries and pickups, including the decisions to be made by the transport provider, and the model assumptions. Section 3 describes the method for computing carbon emissions and Section 4 describes how we compute route lengths (distances) and load factors. Section 5 describes our experimental set-up, and Section 6 presents results and accompanying analysis. Finally, we draw conclusions and point to further research in Section 7.

2. Our model of combining deliveries and pickups

We wish to isolate the effect on carbon emissions of combining two separate flows, namely one flow from the depot to delivery locations, and one flow from the pickup locations (which could overlap with delivery locations) to the depot. In this section, we describe how a transport provider would operate such vehicles in our model, and specify the assumptions.

The transport provider has a homogeneous fleet of vehicles. There is a set of *delivery locations*, to which items of specified size should be delivered, and a set of *pickup locations*, from which items of given size are picked up. A location with both demand and supply is called a *joint location*. The goal of the transport provider is to find a cost minimal *routing plan*: a set of routes from the depot through all locations such that the required quantities are delivered to or picked up from each location, and vehicle capacity is obeyed. Figure 1 presents an example with three delivery and three pickup locations.

In the basic *separate set-up*, a given vehicle can only be used to perform either deliveries or pickups on a route, but not both. As a consequence, there will be separate routes for the deliveries and the pickups. Now we assume that the transport provider has the option to use a given vehicle to perform both deliveries and pickups. We investigate and compare three set-ups corresponding to different levels of consolidation flexibility:

- the *separate set-up*, with separate fleets and routing plans for pickups and deliveries

- the *backhauling set-up*, where deliveries and pickups can be combined on a route, but all deliveries must take place before the first pickup

- the *mixing set-up*, where pickups and deliveries may be combined freely on a vehicle.

Insert Fig. 1 about here.

In Figure 1, we illustrate an example with deliveries from the depot to 3 locations with demand of 3 t (ton) each, and pickups from 3 other locations with supply 2 t, each destined for the depot. Vehicles with a capacity of 10 t are available for deliveries and pickups. The driven distance in the separate set-up equals 77 km. If, however, pickup and delivery items can be mixed freely on a single vehicle as long as their total weight never exceeds the 10 t capacity, the length of the combined delivery and pickup route is 42 km compared to the total length of the separate delivery and pickup routes of 77 km.

In order to illustrate how the detour effect may counteract the emission savings from reduced travel distance, we use the example from Figure 1. In the mixing set-up, the number of ton-kilometers (tkm)[1] equals 317, but in the separate set-up, the total tkm is only 248. Implementation of combined

[1]This is calculated by multiplying the load on each leg with the length of that leg, and summing over all legs in the route.

delivery and pickups reduces the distance by 35 km to 42 km, but due to the detour effect, the number of tkm increases with 79, because items are on average transported over longer distances. These additional tkm may offset the emission savings.

A question is then how the routing plans for the three set-ups are obtained. We assume that the transport provider minimizes total distance as a proxy for routing costs[2] in each of the three set-ups. The obtained solutions should be of reasonable quality that would result from an industrial vehicle routing solver; see Section 4. An alternative objective would be to explicitly minimize carbon emissions or fuel consumption, as is done in Ubeda et al. (2011). The reason for not selecting this type of objective is that we would like to consider the routing solutions that a transport provider would construct in practice. We believe that transport providers of today will not accept routes with higher costs in return for lower carbon emissions. Moreover, many models that minimize fuel consumption or carbon emissions often depend on complicated computations; see Section 3.

There are several key assumptions to our model. We assume that the quantities to be picked up and delivered are given. The fleet is homogeneous, i.e., the vehicles all have the same capacity and the same emission characteristics. We note that, in order to combine deliveries and pickups, it may be necessary to invest in new vehicles with different characteristics, but this issue is outside the scope of our study. We limit ourselves to situations where many pickups and/or deliveries can be visited by a vehicle. Routes originate and terminate at a single location: the depot. Therefore, we do not study the case where a vehicle is filled on a return leg by taking a transport haul to a different location than the depot. These assumptions keep the model and its solution tractable. Removing them will be topics for future research.

3. Computation of carbon emissions

In order to compare the carbon emissions of the three set-ups, we need a way to calculate the emissions related to the distance driven and the load on the vehicle. To illustrate, we use the running example from Figure 1. In this example, the mixing set-up reduces the distance from 77 to 42 km compared

[2]We assume that the vehicles all drive with the same constant speed, so time is proportional to distance.

to the separate set-up, but increases the number of tkm from 248 to 317. In this Section, we describe a simplified but reasonable method for comparing the carbon emissions of different set-ups.

For our comparative study, we need the marginal emissions of adding a ton of load to the vehicle: average emissions per km or per tkm are of no use. A standard approach is to compute the emissions per km of a vehicle with a given load. In our example, the truck starts by driving 15 km with 9 t of goods. It then drops off 3 t and drives 1 km with a load of 6 t, proceeds to pick up 2 t, and so on. We wish to determine the carbon emissions of such a routing plan, or, rather, compare the emissions resulting from different routing plans for different set-ups.

In the green vehicle routing literature, two different approaches are used to compute the impact of the load on the vehicle, namely engine emission models and actual measurements on given vehicles (Turkensteen, 2016b). *Engine emission models* relate emissions or fuel consumption to factors such as the load on the vehicle, the driven distance, the speed of the vehicle, and other aspects of the driving situation (Demir et al., 2011). For example, the study by Franceschetti et al. (2013) measures the fuel consumption of a vehicle with a weight of 6 t when empty and a maximum capacity of 6.25 t for driving at different fixed speeds. *Actual measurements* are available for specific vehicles with different load factors, e.g., fully loaded and empty. For example, Ubeda et al. (2011) use previously measured fuel consumption of a given (but further unspecified) vehicle at a case company. However, both approaches have the disadvantage that they are only valid for specific vehicles, and possibly under specific conditions. Engine emission models have the additional disadvantage that computed carbon emission levels depend on the input parameters, such as the speed profile of a given transportation haul. This makes it difficult to obtain generally valid carbon emission levels.

We choose to simplify our computations by making the following two assumptions:

- We do not differentiate between driving conditions in the considered area. In reality, there are different driving conditions on, e.g., urban roads, rural roads, highways, and roads in mountain regions. This would have an effect on carbon emissions: For example, the load on the vehicle has the largest impact on highways (Turkensteen, 2016a) and driving a heavily loaded vehicle uphill would contribute strongly to carbon emissions. However, since our goal is to investigate the effect of various delivery and pickup schemes, it is not unreasonable to assume that driving conditions are the same on all considered paths.

- The marginal carbon emissions are the same for each added ton of load. Widely accepted engine emission models, such as CMEM by Barth et al. (2005) and PHEM by Hausberger et al. (2009), relate emissions and fuel consumption linearly to the mass of the load on the vehicle, given that all other factors (e.g., driving conditions or the vehicle) are constant. Many papers in transport science confirm this relation; see e.g. Walnum and Simonsen (2015) and Chapter 3 of McKinnon et al. (2015).

Based on these two assumptions, we may simplify the emission computations. As a consequence of the linear relationship, a vehicle that drives precisely 30% loaded (weight-laden) has the same expected carbon emissions as the same vehicle that drives 30% loaded on average. The *load factor (LF)* is defined as the average weight of the load compared to the maximum capacity that can be taken by the vehicle. It can also be computed as the total number of tkm on a given haul divided by the maximum number of tkm for a fully loaded vehicle over the same distance. In McKinnon et al. (2015), this is known as the *weight based lading factor*.

Given our assumptions, it suffices to know the proportion of emissions that is due to the maximum net payload on the vehicle, the *load-based emission percentage (LBEP)*. This measure is introduced and explained Turkensteen (2016b), and the range of realistic LBEP values for different sized vehicles are presented. The emissions of a vehicle are computed as follows:

$$Emission\ units = [(1 - LBEP) + LBEP \times LF]d, \tag{1}$$

where LF is the average load factor, and d is the driven distance. The term $(1 - LBEP)$ denotes the emissions of an empty vehicle and the term $LBEP \times LF$ the emissions due to the load on the vehicle[3].

We illustrate the emission computations in (1) with the example from Figure 1, where the vehicle capacity is 10 t, the distances in the separate and combined pickup and delivery set-ups are 77 and 42 km, and the number of

[3]The emission units measure the number of km that a fully loaded vehicle should cover to obtain the same emissions. Absolute emissions are obtained by multiplying this with the emissions per km of a fully loaded vehicle.

tkm are 248 and 317, respectively. The average load factors are for separate delivery and pickup $248/(77 \times 10) = 32.08\%$ (where a fully loaded vehicle would transport 10 t over 77 km) and for the mixing set-up $317/(42 \times 10) = 74.48\%$. Now suppose that the maximum load of the vehicle in isolation contributes 20% of the carbon emissions. In other words: a fully loaded vehicle gives emissions of 1 unit and an empty one 0.8 units; if the vehicle is 32.21% loaded on average, it emits on average $(0.8 + 0.2 \times 0.3221) = 0.8644$ units[4] per km. For our example, the computed emissions are $(0.8 + 0.2 \times 0.3221) \times 77 = 66.56$ units in the separate set-up and $(0.8+0.2\times0.7448)\times42 = 39.86$ units in the mixing set-up. Thus, if the LBEP value of the vehicle is 20%, the usage of the mixing set-up reduces carbon emissions by around 40% in our example.

The paper by Turkensteen (2016b) contains a survey study on the LBEP obtained in different databases and studies. It shows that for relatively light vehicles, such as vans, the percentage lies between 10 and 30%, increasing to between 30 and 50% for trucks. One of the factors causing this large degree of variation is formed by the driving conditions. The gross vehicle weight (GVW), the weight of a fully loaded truck, is often used to classify vehicles; see Campbell (1995). McLeod et al. (2008) find that the vehicles used in reverse logistics vary in GVW from 5 t vans to large 40 t trucks. Therefore, the entire range of LBEP values can apply to our study. Table 1 describes to which case each of the selected percentages apply.

Insert Table 1 about here.

Our experiments are based on a set of problem instances. Each instance specifies locations and sizes of the delivery and pickup orders, and the vehicle capacity. For each such instance, we first calculate routing plans for each set-up, compute the associated distance and average load factor, and then calculate the carbon emissions for LBEP values between 10% and 50%. In the following section we give a detailed description of the distance and load computations.

[4] This does not give how many kg each unit corresponds to, but this is not relevant for our further analysis: the obtained data is sufficient to compare the different set-ups.

4. Distance and load computations

As described in Section 2, a transport provider has the option to use separate vehicles for deliveries and pickups, or to consolidate with backhauling or mixing. In each of these set-ups, the transport provider wishes to minimize costs by minimizing total driven distance. Hence, the problem of the transport service provider is to find a routing plan for serving a given set of delivery and pickup orders such that the total cost is minimized. The plan must adhere to the vehicle's capacity and possibly other operational restrictions such as time windows. The class to which this problem belongs is known as Vehicle Routing Problems (VRPs). In fact, many variants exist to match the specific conditions of a given situation.

The most basic VRP variant is the Capacitated VRP (CVRP). Here, a fleet of identical, capacitated vehicles based at a depot is available for serving a set of customer requests of a given size and location. For the separate set-up, we may solve two CVRPs: one for the delivery requests and one for the pickup requests. The Vehicle Routing Problem with Backhauls (VRPB) minimizes the length of a route in which combining delivery and pickup orders on a vehicle is allowed but where all deliveries must take place before any pickup. This VRP variant corresponds to our backhauling set-up. The problem of finding the best route where deliveries and pickups can be mixed freely on a vehicle (as long as the vehicle's capacity is not exceeded), is the VRP with Pickups and Deliveries (VRPPD). For an extensive overview of these and other VRP variants, we refer to Toth and Vigo (2014).

4.1. Route length estimates by continuous approximation

First, we consider an analytical approach that yields an estimate of the length of routes through a number of locations, namely *continuous approximation*; see e.g., Daganzo (2004). The advantage of this approach is that it directly provides the driven distance as a function of some key parameters in a distribution situation, such as the size of the area and the number of locations to be visited.

A VRP solution can contain multiple routes. Each route consists of *headways*, i.e., the legs from the depot to the first location and from the last location back to the depot, denoted by $\bar{l}$, and a *detour*, the leg on a route between the first and the last customer. The parameter C specifies the maximum number of stops per route. The route length d_{VRP} is computed as (Daganzo, 2004; Larsen and Turkensteen, 2014):

$$d_{VRP} = 2(N/C)\bar{l} + [(C-1)/C]K\sqrt{AN}, \tag{2}$$

where A is the surface of the area, and K describes the type of distance (e.g. $K = 0.57$ for Euclidean distances and 0.72 for Manhattan distances). The term $[(C-1)/C]$ is often set to 1. The headway length $\bar{l}$ can be computed as $\bar{l} = (2N/C)E(r)$, where $E(r)$ is the average distance from the depot to any location in the area.

The paper by Beullens et al. (2004) provides distance estimates to be used in our set-ups with backhauling and mixing. We denote the number of pickup locations by N_c, the number of delivery locations by N_d, and the number of joint locations by N_e. We then have the following VRPs for the different set-ups:

- The CVRP for separate pickups contains $N_c + N_e$ locations;
- The CVRP for separate deliveries contains $N_d + N_e$ locations;
- The backhauling VRPB combines outward routes through $N_d + N_e$ locations with inward routes through $N_c + N_e$ locations;
- The mixing VRPPD contains $N_c + N_d + N_e$ locations.

The total distances from the CVRP instances with separate pickups and with separate deliveries can be computed by inserting $N = N_c + N_e$ and $N = N_d + N_e$ locations in Eq. (2), respectively, and determining d_{VRP}. The total distances for the VRP variants with mixing and backhauling are formulated in Beullens et al. (2004). Instead of the maximum number of stops C, the approximation uses the vehicle capacity of W, the amount of Q_d delivered per location, and the amount Q_c picked up per location (all three are measured in the same unit). This is necessary as the rationale behind the backhauling and mixing set-ups is that more locations can be visited than in the separate pickup and delivery routes, namely both supply and demand locations. Thus the separate pickup route has at most $C = W/Q_c$ stops, and the delivery route $C = W/Q_d$.

The total distance d_b in the backhauling set-up is:

$$\begin{aligned} d_b = 2(\bar{l}/W) \max[(N_d + N_e)Q_d, (N_c + N_e)Q_c] + \\ K[\sqrt{A(N_d + N_e)} + \sqrt{A(N_c + N_e)}] \end{aligned} \tag{3}$$

The total distance d_m in the mixing set-up is:

$$d_m = 2(\bar{l}/W)\max[(N_d + N_e)Q_d, (N_c + N_e)Q_c] + K\sqrt{A(N_c + N_d + N_e)} \tag{4}$$

From the results in Eq. (2-4), we can derive that the following factors influence the relative distances in the three set-ups. The *headway length* ($\bar{l}$) has a larger impact in the separate set-up as there are more routes. The *vehicle capacity* W influences the number of routes and their lengths. The expected detour lengths are roughly equal ($K\sqrt{A(N_e + N_c)} + K\sqrt{A(N_e + N_d)}$ in the separate and the backhauling set-ups, but clearly shorter in the mixing set-up ($K\sqrt{A(N_c + N_d + N_e)}$, in particular for large values of N_e. An important determinant of the relative distances in the three set-ups is the *number of locations* N_c compared to N_d, and also on the *the number of joint locations* N_e. These factors are varied in our numerical experiments.

4.2. Routing plan and distance calculation by solving VRPs

The second method involves solving the above-mentioned variants of the VRP, namely the CVRP, the VRPB, and the VRPPD algorithmically. The VRPB has originally been treated as a separate problem; see Goetschalckx and Jacobs-Blecha (1989); Toth and Vigo (2001). However, more recently, it has also been treated as a special case of the VRPPD and is solved as such; see Wassan and Nagy (2014). A survey on methods for the VRPPD is given in Parragh et al. (2008), where it is stated that exact methods for (static) pickup and delivery problems have not solved instances with more than 96 requests and that for larger instances heuristics were needed. A more recent exact method, the Branch and Cut and Price algorithm of Ropke and Cordeau (2009), solves instances of size 100 to optimality within around 3 minutes. An overview of meta-heuristics is given in Wassan and Nagy (2014). Popular solution approaches include meta-heuristics such as Adaptive Large Neighborhood Search (Ropke and Pisinger, 2006) and the modified savings heuristics described in Dethloff (2001). Strikingly, none of the overview papers Berbeglia et al. (2007) or Parragh et al. (2008) present an overview of actual applications of VRPPDs.

Recently, environmental aspects have been treated explicitly in the VRP literature: the objectives of the problems called *Green Vehicle Routing* and *Pollution Routing* include the minimization of fuel consumption and carbon emissions; see the overview papers by Demir et al. (2014b) and Lin et al.

(2014). The argument for such approaches is that minimization of distance does not always lead to minimal carbon emissions or fuel consumption, as the load on and the speed of the vehicle can play an important role (Bektas and Laporte, 2011). In Ubeda et al. (2011), it is found that the greenest solution can have significantly lower carbon emissions further than the set-up with backhauling.

In our experiments we should avoid that the results are due to quality differences in the solutions to the different set-ups. We therefore wish to obtain solutions of sufficiently high quality to all types of considered routing instances. For our purpose, we need a solver that is able to find high quality solutions to the three types of VRP in reasonable time. Moreover, a solver that is used in industry adds to the realism of our experiments. The solutions should have a quality close to the best known to ensure that striking results are not due to poor solutions. Based on these criteria, we select the Spider industrial VRP solver developed by SINTEF in computational experiments. Spider is built on a rich, generic model of VRPs and has been used in a variety of applications. Spider has also yielded good results on standard benchmarks for several VRP variants. For details, we refer to Hasle and Kloster (2007). It seems reasonable to assume that the routing solutions produced by Spider are representative of the routing plans used by a modern transport provider that utilizes VRP software. We use Spider to solve the CVRP, VRPB, and VRPPD instances in our numerical experiments. For all experiments, Spider is run until there have been 50 iterations without improvements. We combine the CVRP solutions with delivery orders and with pickup orders into one solution for the separate delivery and pickup set-up.

4.3. Calculation of load factor

The next step is to compute a vehicle's average *load factor*. Beullens et al. (2004) present no continuous approximation results on these. Some studies in the field such as Burns et al. (1985) assume that vehicles deliver at a constant rate. If the same holds for supply, we can expect that the load factor in backhauling is roughly similar to the ones observed in the separate set-up and that the load factor in mixing is about twice as high. However, there are several complicating factors, such as a difference between demand and supply that make it necessary to determine the load factors in the set-ups experimentally.

For algorithmic VRP methods, the computation of the load factor is straightforward from the generated routing plan: For each vehicle, we obtain the order in which locations are visited. From the demand and supply quantities at each location, we find how much is on the vehicle during each part of a route. As in the example from Figure 1, one can compute the number of ton-kilometers and from that, the average load factor.

We can obtain analytical estimates of expected driven distances from the field of continuous approximation and algorithmic approaches to find routes of high quality in each of the set-ups. The continuous approximation results provide the key determinants of the driven distance used in Section 5 but in order to compute load factors and to establish the variation within and between instances, we use a commercial VRP solver.

5. Design of numerical experiments

The purpose of our numerical experiments is to find the carbon emission savings from combining deliveries and pickups. In accordance with our model of the transport provider from Section 2, we determine the routing plans, i.e. the sets of routes that minimize the cost (total distance), for the separate, the backhauling, and the mixing set-ups. The goal is to determine the driven distances and load factors for all set-ups. The routing plans are determined with the Spider industrial solver presented in Section 4, and the resulting carbon emissions are computed using the LBEP values from Section 3. The basic, separate set-up serves as the baseline. Emission savings from backhauling and mixing are reported relative to the separate set-up. In our experiments we assume that the distances between all locations are Euclidean.

Our study on the continuous approximation results in Section 4 indicates that the following factors have an impact on the driven distance and the load factor: the headway length, the vehicle capacity, the number of joint locations, and inequality in the number of demand and supply points. To this we can add the LBEP value from Section 3, which plays a role when the load factors in set-ups differ. The vehicle capacity (relative to the quantities at each delivery or pickup location), the LBEP value, and the relative number of demand and supply locations, as well as the number of joint locations, are varied in our experiments, typically between a low level and a high level; see Table 2.

In order to assess the impact of the factors, we use a diverse set of instances. An instance in our experimental set-up is formed by a set of locations with their pickup and/or delivery demand and the depot location, the distances between each pair of locations, and the vehicle capacity. We distances are Euclidean, they can be derived from the coordinates of the locations.

The instances used are derived from the well-known *Li and Lim* instances that contain coupled demand and supply locations in a two-dimensional plane. Distances are Euclidean. At each pickup location, goods of a specified size must be picked up, and delivered later by the same vehicle to the corresponding delivery location. A homogeneous fleet of vehicles with a given capacity is based at a depot, available for servicing the pickup/delivery requests. There are time windows on the pickups and deliveries during which service has to start. These instances are introduced in Li and Lim (2003) and can be retrieved from `http://www.sintef.no/top`.

To suit our three set-ups, we disregard time windows and the coupling (same vehicle) and precedence constraints on the pickup and delivery task pairs in the Li and Lim instances. All collected items are destined for the depot and all deliveries originate from the depot. With no time windows, route lengths are only limited by the capacity of the vehicle.

Regarding the *headways*, we can separate between instances the `lr` instances with short headways, the `lrc` instances with intermediate headways, and the `lc` instances with long headways. In `lr` instances, the locations are uniformly distributed in an area, whereas in `lc` instances the locations are clustered and generally relatively far away from the depot, as is illustrated for the `lc101` instance in Figure 2. Regarding *vehicle capacity*, there are so-called type 1 instances (numbered 101-109) with low capacity and around five or six short routes, and type 2 instances (numbered 201-208) with large vehicle capacity and typically one or two long routes.

In order to measure the effect of having *joint locations*, we create modified Li and Lim instances (the `joint` instances) such that half of the locations are joint locations. In order to measure the impact of *different numbers of pickup and delivery locations*, we construct the `uneq` instances where half of the delivery locations are removed, so that the number of delivery locations is half the number of pickup locations (the results are similar if we do the reverse). Finally, the *LBEP value* is not dependent on the setting of instance but can be applied when the instance is solved and the distance and load factor is obtained. We set the LBEP to a low value of 10% and a high value of 50%.

Insert Fig. 2 about here.

Table 2 describes which type of instances correspond to low, intermediate, and high values of the five parameters in our experimental design. For example, the instance `lr202` has short headways but, since the vehicle capacity is high, relatively long routes. The corresponding instances `lr202uneq` and `lr202joint`, with different numbers of pickup and delivery locations and joint locations, respectively, are obtained from the regular instance `lr202`.

Insert Table 2 about here.

We first determine the distances and load factors for all set-ups and all instances. The distances and load factors of the separate delivery pickup routes are combined. Based on the distances and load factors, we compute the carbon emission savings relative to the separate set-up for LBEP values of 10% and 50%. The results for all individual instances are reported in Tables 8, 9, 10, 11 in the Appendix.

6. Results of the numerical experiments

In this section we present the key findings of our computational experiments. The distance and average load factors of all instances are reported in Tables 8, 9, 10, and 11. These detailed results form the basis of the treatment in this section. Based on the observed distances and load factors we establish the resulting emission savings and their relation with the key factors listed in Section 5. Finally, we characterize the cases in which backhauling and mixing are most and least effective for carbon emission savings from consolidating shipments.

Recall that the names `lc`, `lr`, and `lrc` denote instances with locations on a plane that are uniformly random, clustered, and a combination of both, respectively. Type 1 / type 2 instances have small / large vehicle capacity. Suffixes `uneq` and `joint` denote instances with unequal demand and supply, and added joint locations, respectively.

6.1. Observed distances and load factors

The key determinants of carbon emission savings are the driven distance and the average load factor in each set-up. In order to compare results across instance classes we show the average load factor and the distance in each instance class and for all three set-ups in Table 3.

In general, the distances for backhauling are about 10% to 20% smaller than for the separate set-up. For mixing, the corresponding distance reduction is between 20% and 40%. The load factor in the backhauling set-up fluctuates around that of the separate set-up; the load factor of the mixing set-up is a factor 1.5 to 1.8 higher.

Insert Table 3 about here.

We consider the impact of parameters on the load factor and driven distance. The impact of the *vehicle capacity* is that large load factors are observed low for the type 2 instances. The cause is that in some cases two vehicles are needed where the second vehicle is poorly utilized, but this is not a property that can generally be expected. Distance savings from backhauling in particular are quite small for the `lr2` instances. The long *headway lengths* in the `lc` instances mean the distance savings largest here. Another consequence is that the load factor of backhauling is higher than in `lr` instances; this is most pronounced in the type 1 instances with low vehicle capacity, where the headways constitute a relatively large share of routes. Other factors, such as the presence of joint locations, have little impact on load factors, so the results are not reported here.

The proportion of delivery, pickup, and joint locations appears to have an impact on the distance and the load factor. For both backhauling and mixing, distance savings are small for the `uneq` instances, and load factors are low compared to the separate set-up. As expected, this is so because it is not possible to find similar delivery and pickup quantities for all trucks (the option to use vehicles only for pickups or deliveries is unattractive when the total distance is minimized). For the `joint` instances, load factors are quite high and distance savings large: Deliveries and pickups being at the same location has the consequence that the vehicle can be well utilized. For backhauling, load factors are particularly high for `lc1uneq` instances, where few locations are visited and one can save distances and increase load factor by starting the pickup part of the route at the location where the delivery part terminates.

Finally, we can determine the size of the *detour effect* from the distances and the load factors in the set-ups for each instance. As the demand and supply quantities in all set-ups are equal, the driven distances times the load factor in the respective set-ups give a term that is proportional to the number of tkm. If this amount is consistently larger for the mixing or backhauling set-up than for the separate set-up, this is due to the detour effect. From the results in the individual instances in Tables 8 to 11, we find sizeable variations in the number of tkm for most instance types. These are due to specific routing decisions, e.g., the decision to visit a delivery location with much demand first in a certain set-up. However, these variations even out for almost all instance classes. The detour effect is only significant for the regular `lr2` instances (P value of 0.006), where the number of tkm is about 50% higher in the mixing set-up than in the separate set-up.

6.2. Carbon emission savings

Now we address the carbon emission savings from backhauling and mixing relative to the separate pickup and delivery set-up. Recall that our carbon emission assessment tool is not designed to measure absolute emissions but to compare routing strategies with different distances and load factors. We illustrate the computation of carbon emission savings from the distance and the average load factor using the instance `lc102`. We find that the separate set-up has a total distance of 1187.7 km (over 6 pickup and 6 delivery routes) with an average load factor of 50.8% and the mixing set-up has a total distance of 693.6 km with an average load factor of 65.6%; see Table 8. If the LBEP value is 10%, it holds according to (1) that the emissions in the backhauling set-up are $(0.9 + 0.1 \times 0.508) \times 1187.7 = 1129.3$ units, whereas those in the mixing set-up are $(0.9 + 0.1 \times 0.656) \times 693.6 = 669.7$ units. Thus, emissions in the mixing set-up are 40.4% lower than in the separate set-up. All carbon emission savings reported here are the result of such computations. Below, we describe the impact of the key factors on the observed emission savings.

The question is: Are the emission savings consistent with the distance savings reported in Table 3? We find that for an LBEP value of 10%, emission savings and distance savings are almost proportional. However, for the LBEP value of 50%, these savings diverge, as we show graphically in Figures 3 and 4. In these figures, the horizontal axis represents the distance savings of each instance and the observed emission savings on the vertical axis. If distance savings and emission savings were equal, all instances would be on the dotted diagonal line. For backhauling, it is almost equally likely that emission savings are higher or lower than distance savings. For mixing, emission savings are generally clearly smaller than distance savings. Two highlighted instance types are the `uneq` instances, where emission savings are relatively high compared to distance savings due to the low observed load factors, and the `joint` instances, where the reverse is true and emission savings are low compared to distance savings.

Insert Fig. 3 about here.

Insert Fig. 4 about here.

Insert Table 4 about here.

The first impact that we consider in Table 4 is the *headway length*, where we compare the average, the maximum, and the minimum emission savings of the `lc` and `lr` instances. The largest emission savings are, as expected, observed for the clustered `lc` instances with long headways. The influence of headway length appears to be strong for both mixing and backhauling.

In order to evaluate the impact of the *vehicle capacity*, we separate between the type 1 and 2 instances with low and high vehicle capacities, respectively, in Table 5. The type 1 instances, with relatively many headways and short detours, have the largest emission savings, in particular for the backhauling set-up. Interestingly, the vehicle capacity appears to have little influence on the distance savings from mixing, which can be observed from the small difference between emission savings for type 1 and 2 instances for the LBEP value of 10%. For large LBEP values, however, the increased detour effect in mixing reduces the emission savings for type 2 instances. In fact, the combination of a large vehicle capacity and short headways can lead to larger emissions for the consolidated set-ups than for the separate set-up; see Section 6.3.

Insert Table 5 about here.

The effect of having *unequal numbers of supply and demand locations* in Table 6 is ambiguous. Mixing can accommodate the unequal number of demand and supply points quite well and achieve large emission savings, in particular for the LBEP value of 10%. In the backhauling set-up, on the other hand, it is necessary to construct a separate return haul through relatively few pickup locations, which makes that distance savings are close to or less than those for regular instances. Only for the LBEP value of 50%, the decreased load factor leads to some emission savings.

Insert Table 6 about here.

Next we compare the instances that have joint locations (`lcjoint` and `lrjoint`) to the corresponding instances that do not (`lc` and `lr`) in Table 7. As expected, the mixing set-up benefits most from the addition of joint locations with carbon emission savings of up to 42.4% for the instance `lc105joint`. The impact of backhauling is small and opposite to that in the case with `uneq` instances as both the load factor and the distance are slightly higher than for regular instances.

Insert Table 7 about here.

One finding is that there is interaction between the different parameters. The interaction between vehicle capacity and headway length is as follows: The smaller the vehicle capacity, the more routes are needed to visit all locations, and the larger the benefit from mixing and backhauling. A more subtle interaction effect takes place between the LBEP value and the number of delivery, pickup, and joint locations. Having more joint locations increases the load factor in the mixing set-up and thereby the impact of the LBEP value, whereas the effect can go in the opposite direction when demand and supply quantities becomes unequal. When the right interaction of factors take place, extremely high or low emission savings can occur. We discuss these cases in Section 6.3.

6.3. Extreme cases of carbon emission savings

In our experiments the mixing set-up can achieve emission savings of 42% and backhauling up to 25% compared to the separate set-up. The mixing set-up achieves its largest emission savings for instances with many joint locations, a small vehicle capacity, and a low LBEP value (since the load factor in mixing is generally high). The largest emission savings of around 42% are achieved for the instance `lc105joint` and an LBEP value of 10%. Backhauling yields the largest emission savings for instances with long headways and small vehicle capacities; As the load factor is similar to that of separate pickups and deliveries, the LBEP value generally has little impact. The largest emission savings of around 30% are attained for the instance `lc101`.

Interestingly, there are also cases for which a set-up with backhauling or mixing can lead to emission increases over separate pickups and deliveries. We call these situations the *backhauling paradox* and the *mixing paradox*, respectively. Both set-ups achieve distance savings for all our instances and emission increases can only occur if items are on average transported over a longer distance than in the separate set-up, i.e., the detour effect occurs.

The backhauling paradox occurs for some `lr2` instances (with short headways), namely `lr208` and `lr103joint`, but also for the `lc205` instance, where headways are long in absolute terms but only constitute a small part of the total route length. For these instances the solver happens to find a solutions in which items are transported over longer distances in the backhauling set-up than in the separate set-up.

The mixing paradox occurs for the instances `lr202` and `lr206` for an LBEP value of 50%. In both cases distance savings are around 20%. The payload effect reduces some of the savings but the detour effect can turn these into emission increases. The addition of joint locations makes the mixing paradox less likely as the distance savings are larger and the detour effect is smaller for such instances.

7. Conclusions and future research

This paper assesses the carbon emissions impact of consolidating shipments from and to a centralized depot. To that end we model a transport provider, compute cost-optimized solutions (by minimizing total distance) in separate and consolidated set-ups, and determine the carbon emission savings to be achieved by consolidation from the driven distances and the load factors. Combining deliveries and pickups can be attractive, not only because costs can be reduced through shorter distances, but also because the distance savings can entail environmental impact reductions. We consider the consolidation set-up in which all deliveries are made before the pickups, called backhauling, and the set-up where one can freely mix pickups and deliveries while not violating the vehicle's capacity, called mixing.

We find emission savings from backhauling and mixing can be up to 35% and 40%, respectively. These savings are attained if the distances between the depot and the nearest locations are long and the vehicle capacity is relatively small, for example, in an urban area where the depot is located on the edge of the city and vans or small trucks are used. The emission savings from mixing are highest in case of many locations with both demand and supply, e.g., shops that both sell new items and return used items.

However, there are also conditions under which consolidating outbound deliveries and inbound pickups can increase the total emissions from the vehicles compared to separate collections and deliveries. Mixing can lead to emission increases if the following conditions hold: 1) many locations are evenly distributed between supply and demand locations (but without joint locations), 2) the distance from the depot to the nearest locations is small, 3) the vehicle capacity is large, and 4) the payload causes a large share of carbon emissions (this can occur for heavy vehicles). Backhauling can occasionally lead to emission increases when the headway lengths are short and when the solution at hand has a higher load factor than the separate set-up.

These findings can provide some guidance to policymakers and possibly transport providers. Increasing the degree of vehicle utilization through backhauling may be attractive, not only because costs can be saved through shorter distances but also because the distance savings can entail environmental impact reductions. However, better vehicle utilization is not beneficial for the environment if it is achieved by moving items over longer distances. Our results indicate that these environmental gains can be quite significant in the aforementioned situations with small vehicles and routes with few stops, as could be expected in urban situations. However, if one uses a large vehicle (where the load can cause much of the carbon emissions), these gains may be much less significant and therefore less attractive from an environmental perspective.

One explanation for our results is a mismatch between the chosen objective, the distance driven (and the number of vehicles used) and the environmental impact, which partially depends on the load on the vehicle. Some studies argue that one should also minimize carbon emissions or fuel usage (which are, arguably, proportional to emissions), as has been done in Ubeda et al. (2011). This may allow transporters to find greener mixing solutions, for example, by reducing the average load on the vehicle; our results indicates that this may be particularly relevant in a set-up with combined pickups and deliveries.

In our study, we do not compare *realistic* operational costs between different pickup and delivery set-ups, as we use distance as a proxy for cost. To the best of our knowledge, studies into the relative operational cost of such set-ups are rare. An interesting direction for future research is to determine the cost savings from combining pickups and deliveries. Another issue that we do not address in this paper is the fact that collection of used materials is often environmentally friendly: the waste materials from the product do not end up on a landfill, raw materials may be preserved, and the pollution and emissions from making the product or its materials are prevented. By enabling collection against little extra costs, backhauling and mixing may make such recycling or re-manufacturing cost-effective.

We have made limiting assumptions to the different set-ups, giving rise to directions for future research. As has been pointed out in Section 6, an interesting direction for future research is to use road distances and driving times drawn from areas with different geographies. Secondly, the route length is limited by the vehicle's capacity rather than driving time. The impact of driving time constraints would also be interesting to investigate. A consequence of such constraints could be a reduction in the total distance savings from mixing and backhauling, but also a decrease in the load factor.

A potential weakness of our study is that we restrict ourselves to instances with Euclidean distance. The question is whether the usage of road distances leads to different results than in our experiments. The results from Berens and Körling (1985) suggest that road distances are often well approximated if the Euclidean distances between locations are multiplied by a constant. Moreover, Cooper (1983) conduct a study to determine the fit of a linear function of straight-line distance of actual costs in the British East Midlands, and find a very high R^2-value of 0.97. These results imply that in some cases Euclidean distances are quite representative. We therefore expect that the experiments on instances with actual road distances give similar results to those obtained with Euclidean distances. We have performed initial experiments that indicate that this is indeed so. However, using real road distances and driving times is an interesting direction for future research, as it enables us to measure the impact of various types of geography and topography.

Another assumption is that mixing of deliveries and pickups on a truck does not lead to additional time consumption. If this is so, the costs of a mixing set-up may not only be related to distance or fuel consumption, but also to the time consumption due to moving items on the truck. Thirdly, items to be delivered and items that have been picked up can be freely mixed on the vehicle, and it may limit the number of stops within a time constraint on each route. It could be interesting to determine distance and emission savings if only partial mixing is possible: the detour effect may be reduced compared to full mixing but the same holds for the distance savings. Finally, there are no time windows on our deliveries and pickups. Time windows can deteriorate backhauling and mixing solutions if they make that the pickups and deliveries in the same area cannot be combined in the vehicle. In general, we expect that additional constraints on the routes would have an impact on the quality of the mixing solution in particular, since the capacity of the vehicle can be well utilized during most of the route and the room for modifications may be small. A characteristic of our set-ups is that the depot is both the origin of the deliveries and the destination of the pickups. As we describe in our introduction, more general forms of consolidation can be considered.

Acknowledgement

The work presented in this paper was partly funded by the Research Council of Norway as a part of the DynamITe project [Contract 246825/O70, SMARTRANS].

References

Anily, S., Federgruen, A., 1990. One warehouse multiple retailer systems with vehicle routing costs. Management Science 36, 1.

Barth, M., Younglove, T., Scora, G., 01-01 2005. Development of a heavy-duty diesel modal emissions and fuel consumption model. Tech. rep., Research report California Partners for Advanced Transit and Highways (PATH), UC Berkeley.

Bektas, T., Laporte, G., 2011. The pollution-routing problem. Transportation Research B 45 (8), 1332–1252.

Berbeglia, G., Cordeau, J., Gribkovskaia, I., Laporte, G., 2007. Static pickup and delivery problems: a classification scheme and survey. TOP 15, 1–31.

Berens, W., Körling, F.-J., 1985. Estimating road distances by mathematical functions. European Journal of Operational Research 21, 54–56.

Beullens, P., Van Oudheusden, D., Van Wassenhove, L., 2004. Collection and vehicle routing issues in reverse logistics. In: Dekker, R., Fleischmann, M., Inderfurth, K., van Wassenhove, L. (Eds.), Reverse Logistics – Quantitative Models for Closed-Loop Supply Chains. Springer Verlag, Ch. 5, pp. 95–134.

Bräysy, O., Hasle, G., 2014. Software tools and emerging technologies for vehicle routing and intermodal transportation. In: Toth, P., Vigo, D. (Eds.), Vehicle Routing: Problems, Methods, and Applications. SIAM, Ch. 12, pp. 351–380.

Burns, L., Hall, R., Blumenfeld, D., Daganzo, C., 1985. Distribution strategies that minimize transportation and inventory costs. Operations Research 33, 469–490.

Campbell, J., 1995. Using small trucks to circumvent large truck restrictions: Impacts on truck emissions and performance measures. Transportation Research Part A. 29A (6), 445–458.

Carlgren, C.-G., Carlsson, D., Rönnqvist, M., 2006. Log sorting in forest harvest areas integrated with transportation planning using backhauling. Scandinavian Journal of Forest Research 21 (3), 260–271.

CircularEconomy.com, 2016. Online portal with initiatives towards a circular economy.

Cooper, J., 1983. The use of straight line distances in solutions to the vehicle scheduling problem. The Journal of the Operational Research Society, 34 (5), 419–424.

Daganzo, C. (Ed.), 2004. Logistics Systems Analysis, 4th Edition. Springer Verlag.

Demir, E., Bektas, T., Laporte, G., 2011. A comparative analysis of several vehicle emission models for road freight transportation. Transportation Research D 16, 347–357.

Demir, E., Bektas, T., Laporte, G., 2014a. The bi-objective pollution routing problem. European Journal of Operational Research 232, 464–478.

Demir, E., Bektas, T., Laporte, G., 2014b. A review of recent research on green road freight transportation. European Journal of Operational Research 237 (3), 775–793.

Dethloff, J., 2001. Vehicle routing and reverse logistics: the vehicle routing problem with simultaneous delivery and pick-up. OR Spektrum 23, 79–96.

Franceschetti, A., Honhon, D., Van Woensel, T., Bektaş, T., Laporte, G., 2013. The time-dependent pollution-routing problem. Transportation Research Part B: Methodological 56, 265–293.

Goetschalckx, M., Jacobs-Blecha, C., 1989. The vehicle routing problem with backhauls. European Journal of Operational Research 42, 39–51.

Hasle, G., Kloster, O., 2007. Industrial vehicle routing. In: Hasle, G., Lie, K.-A., Quak, E. (Eds.), Geometric Modelling, Numerical Simulation, and Optimization – Applied Mathematics at SINTEF. Springer, pp. 397–435.

Hausberger, S., Rexeis, M., Zallinger, M., Luz, R., 2009. Emission Factors from the Model PHEM for the HBEFA Version 3. Report Nr. I-20/2009 Haus-Em 33 (08), 679.

Krumweide, D., Sheu, C., 2002. A model for reverse logistics by third-party providers. Omega 30 (5), 325–333.

Larsen, C., Turkensteen, M., 2014. A vendor managed inventory model using continuous approximations for route length estimates and Markov chain modeling for cost estimates. International Journal of Production Economics 157, 120–132.

Léonardi, J., Baumgartner, M., 2004. CO_2 efficiency in road freight transportation: Status quo, measures and potential. Transportation Research Part D 9, 451–464.

Li, H., Lim, A., 2003. A metaheuristic for the pickup and delivery problem with time windows. International Journal on Artificial Intelligence Tools 12 (173), 173–186.

Lin, C., Choy, K., Ho, G., Chung, S., Lam, H., 2014. Survey of Green Vehicle Routing Problem: past and future trends. Expert Systems with Applications 41, 1118–1138.

McKinnon, A., Cullinane, S., Browne, M., Whiteing, A., 2015. Green Logistics, 3rd Edition. Kogan Page.

McLeod, F., Hickford, A., Maynard, S., Cherrett, T., Allen, J., 2008. Developing innovative and more sustainable approaches to reverse logistics for the collection, recycling and disposal of waste products from urban centres. Tech. rep., Green Logistics Project.

Parragh, S., Doerner, K., Hartl, R., 2008. A survey on pickup and delivery problems, part I: Transportation between customers and depot. JfB 58, 21–51.

Ropke, S., Cordeau, J.-F., 2009. Branch and cut and price for the pickup and delivery problem with time windows. Transportation Science 43 (3), 267–286.

Ropke, S., Pisinger, D., 2006. A unified heuristic for a large class of vehicle routing problems with backhauls. European Journal of Operational Research 171, 750–775.

Toth, P., Vigo, D., 2001. Vrp with backhauls. In: Toth, P., Vigo, D. (Eds.), The Vehicle Routing Problem. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, pp. 195–224.
URL `http://dl.acm.org/citation.cfm?id=505847.505855`

Toth, P., Vigo, D. (Eds.), 2014. Vehicle Routing: Problems, Methods, and Applications, 2nd Edition. SIAM Monographs on Discrete Mathematics and Applications, Philadelphia.

Turkensteen, M., 2016a. The accuracy of carbon emission and fuel consumption computations in green vehicle routing. Submitted to European Journal of Operational Research.

Turkensteen, M., 2016b. The computation of carbon emissions due to the net payload on a truck.
URL `http://pure.au.dk/portal/files/100551507/MatuEmissionsfromLoad.pdf`

Ubeda, S., Arcelus, F., Faulin, J., 2011. Green logistics at Eroski: a case study. International Journal of Production Economics 131, 44–51.

Vigga.us, 2015. Homepage of the company Vigga.

Walnum, H., Simonsen, M., 2015. Does driving behavior matter? An analysis of fuel consumption data from heavy-duty trucks. Transportation Research Part D 36, 107–120.

Wassan, N., Nagy, G., 2014. Vehicle routing problem with deliveries and pickups: Modelling issues and meta-heuristics solution approaches. International Journal of Transportation 2 (1), 95–110.

Glossary

Backhauling: A set-up where items can be taken on a truck after all deliveries have been made;

Mixing: A set-up where items to be delivered and picked up items can be combined on a truck;

VRP: Vehicle Routing Problem – a transport optimization problem where the goal is to find a cost minimal set of routes for a fleet of capacitated vehicles such that every customer demand is serviced, and vehicle capacity is obeyed. The VRP comes in many guises.

CVRP: The basic version of the VRP in the fleet is homogeneous, all routes start and stop at the depot, and the only constraint is vehicle capacity.

VRPB: An extension of the CVRP with both deliveries and pickups that may be serviced by the same vehicle, but where deliveries must take place before pickups;

VRPPD: An extension of the CVRP where each customer may have both deliveries and pickups. Pickup and delivery items may be mixed on a vehicle.

Load factor: The share of the vehicle that is used. In our case, we use the weight-based lading factor: the actual weight on the truck divided by the truck's maximum weight capacity.

Appendix

In the subsequent Tables 8, 9, 10, 11, we present the load factors, driven distances, and number of vehicles, for all instances and for all set-ups: separate, backhauling, and mixing.

Insert Table 8 about here.

Insert Table 9 about here.

Insert Table 10 about here.

Insert Table 11 about here.

Figure captions:

Figure 1: A pickup, a delivery, and a combined pickup and delivery route through a set of delivery and pickup locations.

Figure 2: Locations (pickup tasks, delivery tasks, depot) for the Li and Lim instance `lc101`.

Figure 3: The emission savings plotted against distance savings in the backhauling set-up.

Figure 4: The emission savings plotted against distance savings in the mixing set-up.

Table captions:

Table 1: LBEPs for different vehicle types and conditions, based on Turkensteen (2016b)

Table 2: The experimental set-up with the type of instances used in low, intermediate (where applicable), and high values of the relevant parameters

Table 3: Average load factors, distances, and relative distance savings compared to separate set-up

Table 4: Carbon emission savings (minimum, average, maximum) for the backhauling and mixing set-ups aggregated for type 1 and 2 instances

Table 5: Carbon emission savings (minimum, average, maximum) for the backhauling and mixing set-ups aggregated for type 1 and 2 instances

Table 6: Carbon emission savings (minimum, average, maximum) for the backhauling and mixing set-ups for regular instances (`lc` and `lr`) versus `uneq` instances

Table 7: Carbon emission savings (minimum, average, maximum) for the backhauling and mixing set-ups for regular instances (`lr` and `lr`) versus `joint` instances

Table 8: Individual results of `lc` and `lr` instances for delivery, pickup, backhauling, and mixing set-ups

Table 9: Individual results of `lc` and `lr` instances with joint demand and supply locations for delivery, pickup, backhauling, and mixing set-ups

Table 10: Individual results of `lrc` instances for delivery, pickup, backhauling, and mixing set-ups

Table 11: Individual results of **lc** and **lr** instances with unequal demand and supply for delivery, pickup, backhauling, and mixing set-ups

Abstract:

This paper studies the effect on carbon emissions of consolidation of shipments on trucks. New positioning and communication technologies, as well as decision support systems for vehicle routing, enable better utilization of vehicle capacity, reduced travel distance, and thereby carbon emission reductions. We present a novel carbon emission analysis method that determines the emission savings obtained by an individual transport provider, who receives customer orders for outbound deliveries as well as pickup orders from supply locations. The transport provider can improve vehicle utilization by performing pickups and deliveries jointly instead of using separate trucks. In our model we assume that the transport provider minimizes costs by use of a tool that calculates detailed vehicle routing plans, i.e., an assignment of each transport order to a specific vehicle in the fleet, and the sequence of customer visit for each vehicle. We compare a basic set-up, in which pickups and deliveries are segregated and performed with separate vehicles, with two consolidation set-ups where pickups and deliveries may be mixed more or less freely on a single vehicle. By allowing mixing, the average vehicle load will increase and the total driven distance will decrease. To compare carbon emissions for the three set-ups, we use a carbon assessment method that uses the distance driven and the average load factor. An increase in the load factor can reduce part of the emission savings from consolidation. We find that emission savings are relatively large in case of small vehicles and for delivery and pickup locations that are relatively far from the depot. However, if a truck visits many demand and supply locations before returning to the depot, we observe negligible carbon emission decreases or even emission increases for consolidation set-ups, meaning that in such cases investing in consolidation through joint pickups and deliveries may not be effective. The results of our study will be useful for transport users and providers, policymakers, as well as vehicle routing technology vendors.

Highlights:

- First systematic study on the emission effects from combining pickups and deliveries.
- Free mixing of pickups and deliveries often gives the largest emission savings of between 20 and 40%.
- In some set-ups, the impact of the heavier load outweighs the impact of distance savings.

Keywords:

- Pickup and Delivery
- Consolidation
- Carbon emissions

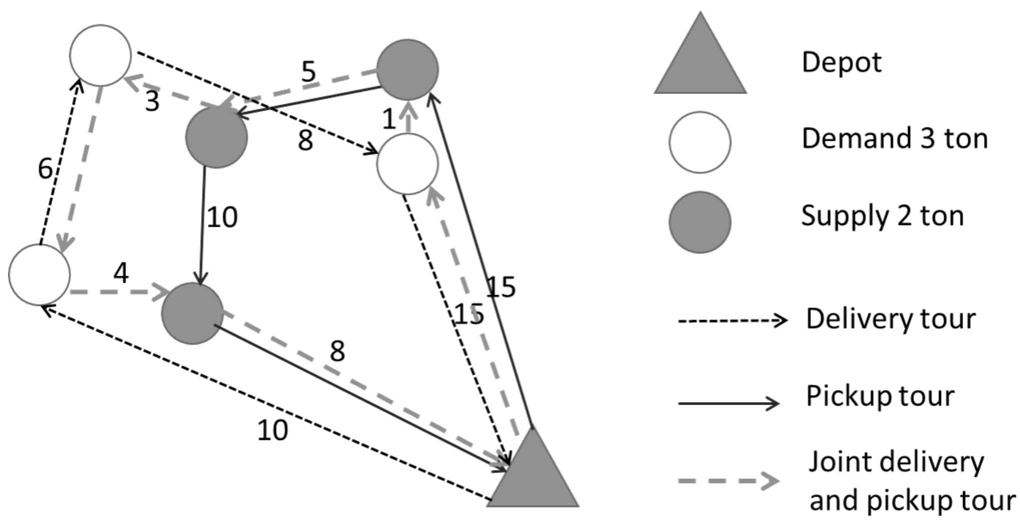

Figure 1: A pickup, a delivery, and a combined pickup and delivery route through a set of delivery and pickup locations.

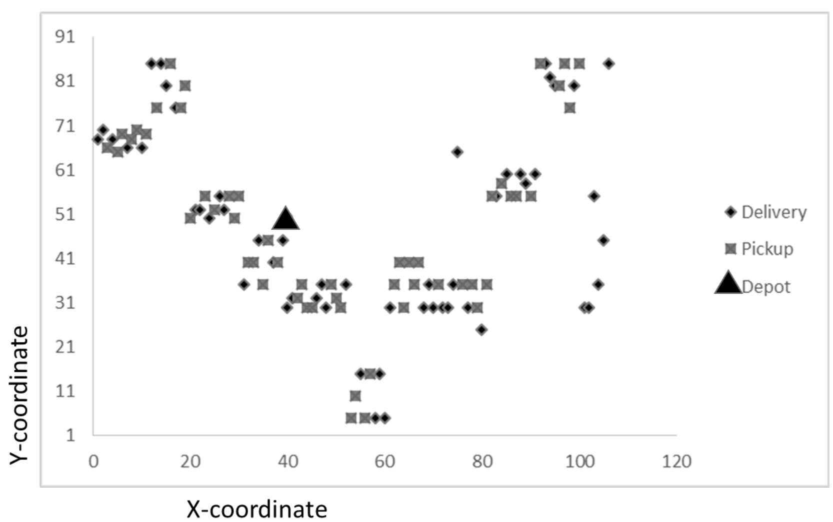

Figure 2: Locations (pickup tasks, delivery tasks, depot) for the Li and Lim instance `lc101`.

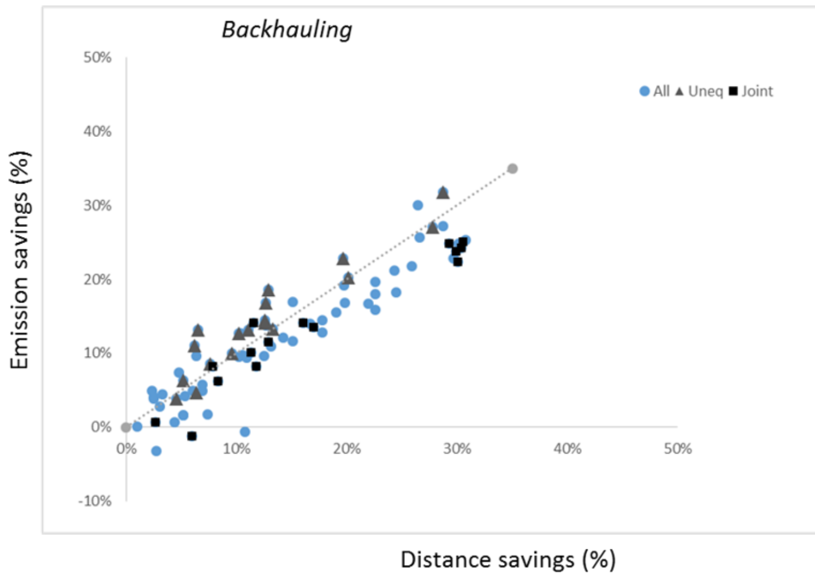

Figure 3: The emission savings plotted against distance savings in the backhauling set-up.

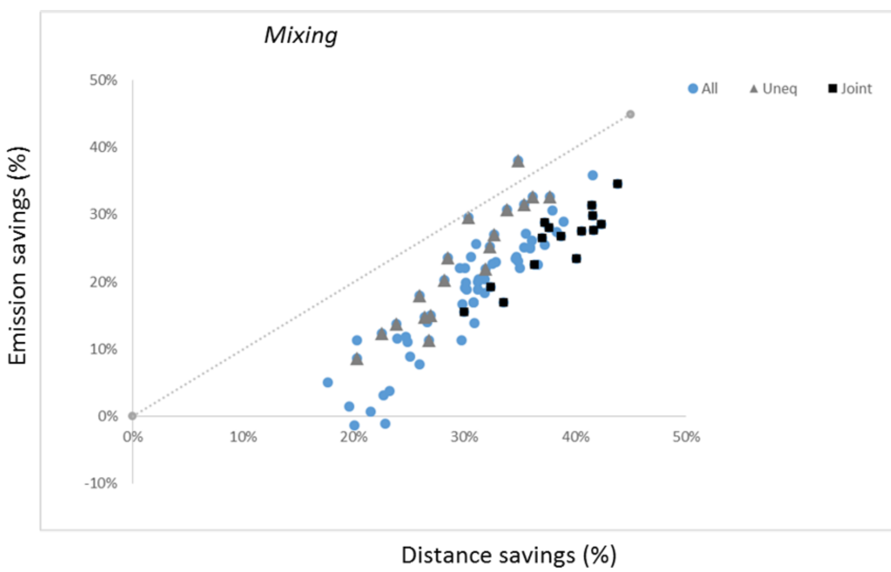

Figure 4: The emission savings plotted against distance savings in the mixing set-up.

Table 1: LBEPs for different vehicle types and conditions, based on Turkensteen (2016b)

LBEP	Vehicle type / Condition
10%	Van, urban conditions
20%	Small trucks
30%	Medium trucks
40%	Heavy trucks
50%	Heavy trucks under free-flowing conditions

Table 2: The experimental set-up with the type of instances used in low, intermediate (where applicable), and high values of the relevant parameters

	Level		
Factor	Low	Intermediate	High
Headway length	lr	lrc	lc
Vehicle capacity	type 1		type 2
Joint locations	regular		joint
Unequal delivery / pickup	regular		uneq
LBEP	LBEP of 10%		LBEP of 50%

Table 3: Average load factors, distances, and relative distance savings compared to separate set-up

		Delivery and pickup type			Rel. distance savings	
Instance	Result	Del. + Pickup	Backh.	Mixing	Backh.	Mixing
`lc1`	Distance	1144.4	842.5	753.7	26.4%	34.1%
	Load factor	47.8%	57.9%	71.6%		
`lc2`	Distance	893.9	773.0	584.6	8.0%	30.4%
	Load factor	37.4%	39.8%	65.7%		
`lr1`	Distance	1096.3	891.8	795.2	18.7%	27.5%
	Load factor	45.3%	46.4%	68.3%		
`lr2`	Distance	844.2	824.2	653.0	2.4%	22.6%
	Load factor	39.2%	39.2%	74.6%		
`lrc1`	Distance	1096.3	891.8	795.2	20.4%	32.5%
	Load factor	48.6%	55.9%	72.7%		
`lrc2`	Distance	814.6	772.0	562.0	5.2%	31.0%
	Load factor	42.1%	46.2%	63.9%		
`lc1uneq`	Distance	945.0	716.2	615.4	24.2%	34.9%
	Load factor	49.30%	46.72%	53.03%		
`lc2uneq`	Distance	773.2	703.1	516.3	9.1%	33.2%
	Load factor	35.27%	29.99%	47.29%		
`lr1uneq`	Distance	875.8	775.6	650.1	11.4%	25.8%
	Load factor	44.26%	40.81%	60.19%		
`lr2uneq`	Distance	721.6	669.7	540.0	7.2%	25.2%
	Load factor	31.41%	31.10%	53.81%		
`lc1joint`	Distance	1359.4	950.8	823.6	30.1%	39.4%
	Load factor	46.91%	59.53%	70.10%		
`lr1joint`	Distance	1349	1176.4	886.6	12.8%	34.3%
	Load factor	44.05%	49.72%	74.61%		
`lc2joint`	Distance	1157.8	1032.6	723	10.8%	37.6%
	Load factor	44.19%	46.68%	77.08%		

Table 4: Carbon emission savings (minimum, average, maximum) for the backhauling and mixing set-ups aggregated for type 1 and 2 instances

	Backhauling			Mixing		
LBEP 50%	Min	Average	Max	Min	Average	Max
All `lc` instances	-0.60%	13.48%	25.71%	8.58%	23.81%	35.85%
All `lr` instances	-3.18%	10.02%	30.05%	-1.42%	13.73%	30.65%
LBEP 10%	Min	Average	Max	Min	Average	Max
All `lc` instances	2.31%	14.85%	29.93%	18.64%	33.13%	42.40%
All `lr` instances	-0.17%	10.33%	28.52%	15.66%	24.99%	36.91%

Table 5: Carbon emission savings (minimum, average, maximum) for the backhauling and mixing set-ups aggregated for type 1 and 2 instances

	Backhauling			Mixing		
LBEP 50%	Min	Average	Max	Min	Average	Max
All type 1 instances	9.97%	19.70%	31.75%	5.05%	23.76%	38.09%
All type 2 instances	-3.18%	7.16%	16.84%	-1.42%	19.01%	31.47%
LBEP 10%	Min	Average	Max	Min	Average	Max
All type 1 instances	9.62%	21.47%	29.93%	15.66%	31.09%	42.40%
All type 2 instances	-0.17%	7.11%	13.25%	16.92%	29.48%	40.24%

Table 6: Carbon emission savings (minimum, average, maximum) for the backhauling and mixing set-ups for regular instances (`lc` and `lr`) versus `uneq` instances

	Backhauling			Mixing		
LBEP 50%	Min	Average	Max	Min	Average	Max
Regular instances	-3.18%	12.21%	30.05%	-1.42%	16.08%	35.85%
`uneq` instances	3.89%	14.78%	31.75%	8.58%	22.64%	38.09%
LBEP 10%	Min	Average	Max	Min	Average	Max
Regular instances	-0.17%	13.86%	29.93%	15.66%	27.76%	40.68%
`uneq` instances	4.43%	13.22%	29.22%	18.64%	28.69%	36.89%

Table 7: Carbon emission savings (minimum, average, maximum) for the backhauling and mixing set-ups for regular instances (`lr` and `lr`) versus `joint` instances

	Backhauling			Mixing		
LBEP 50%	Min	Average	Max	Min	Average	Max
Regular instances	-3.18%	13.76%	30.05%	-1.42%	14.40%	35.85%
Joint instances	-1.25%	13.70%	25.09%	15.52%	25.84%	34.57%
LBEP 10%	Min	Average	Max	Min	Average	Max
Regular instances	-0.17%	15.36%	29.93%	15.66%	26.19%	40.68%
Joint instances	2.31%	16.53%	29.71%	27.77%	36.38%	42.40%

Table 8: Individual results of `lc` and `lr` instances for delivery, pickup, backhauling, and mixing set-ups

	Distance in set-up				*Load factor in set-up*			
Instance	Delivery	Pickup	Backh.	Mixing	Delivery	Pickup	Backh.	Mixing
lc101	610.7	571.6	826.0	813.2	50.51%	48.92%	61.07%	74.24%
lc102	604.4	583.3	871.9	693.7	51.96%	49.63%	52.63%	65.64%
lc103	555.1	530.1	763.3	709.7	47.98%	43.36%	59.95%	70.49%
lc104	553.3	544.9	857.1	670.8	46.76%	46.77%	56.55%	70.79%
lc105	606.1	594.7	831.3	767.4	40.73%	46.11%	54.82%	65.74%
lc106	568.1	560.4	873.8	769.3	47.15%	55.09%	60.02%	76.44%
lc107	568.1	560.4	873.8	769.3	47.15%	47.30%	60.02%	76.44%
lr101	501.4	697.2	881.2	743.7	46.09%	57.62%	45.37%	70.79%
lr102	517.7	747.3	901.3	824.6	46.72%	42.05%	47.04%	69.98%
lr103	547.3	497.9	887.0	860.7	50.27%	48.61%	46.27%	72.34%
lr104	483.6	520.6	872.7	754.9	42.45%	44.20%	46.93%	69.60%
lr105	544.6	552.9	941.7	770.1	44.22%	43.71%	47.42%	70.88%
lr106	549.6	549.0	903.2	826.9	40.92%	42.09%	47.21%	65.67%
lr107	543.3	540.9	870.3	758.2	51.43%	41.47%	47.43%	69.36%
lr108	492.3	517.0	881.2	768.1	41.52%	43.60%	40.93%	65.65%
lr109	534.6	530.2	887.1	849.2	45.87%	43.14%	49.19%	60.80%
lc201	456.4	442.8	837.1	587.1	35.64%	45.64%	42.28%	64.21%
lc202	453.7	418.0	780.4	602.7	41.70%	36.28%	40.26%	73.41%
lc203	450.7	432.3	791.5	554.4	35.32%	36.84%	32.44%	61.52%
lc204	464.7	448.6	798.9	578.5	39.14%	36.25%	42.20%	68.36%
lc205	451.0	444.1	798.9	582.1	36.08%	36.91%	53.85%	63.60%
lc206	440.9	437.6	832.1	617.1	33.48%	36.38%	36.49%	70.37%
lc207	453.6	451.9	806.7	580.5	37.37%	34.06%	37.95%	58.65%
lc208	453.6	444.6	806.7	580.5	37.37%	36.34%	37.95%	58.65%
lr201	440.9	421.4	821.2	645.9	48.52%	45.87%	43.21%	79.10%
lr202	433.4	417.3	827.8	656.3	33.67%	38.79%	44.41%	78.54%
lr203	402.0	434.0	815.5	641.9	42.70%	36.46%	37.15%	74.86%
lr204	415.5	428.4	824.0	662.2	36.79%	39.54%	34.62%	74.87%
lr205	466.8	407.4	847.6	647.8	46.45%	44.41%	45.93%	81.20%
lr206	420.8	406.4	819.0	661.4	31.95%	40.05%	37.24%	72.41%
lr207	395.7	439.0	814.1	645.6	37.04%	34.31%	33.64%	69.94%
lr208	393.0	431.7	824.7	663.4	33.38%	37.37%	37.01%	65.90%

Table 9: Individual results of `lc` and `lr` instances with joint demand and supply locations for delivery, pickup, backhauling, and mixing set-ups

	Distance in set-up				*Load factor in set-up*			
Instance	Delivery	Pickup	Backh.	Mixing	Delivery	Pickup	Backh.	Mixing
lc101joint	673	691	949	851	49.96%	44.89%	60.60%	69.90%
lc102joint	709	721	1001	901	46.66%	46.91%	62.93%	71.01%
lc103joint	695	635	931	834	46.08%	44.08%	57.81%	64.62%
lc104joint	656	648	922	763	47.91%	49.63%	58.26%	74.29%
lc105joint	704	665	951	769	48.10%	44.80%	58.04%	70.67%
lr101joint	658	635	1140	875	47.64%	47.80%	53.73%	76.18%
lr102joint	718	689	1181	862	47.23%	42.42%	48.28%	73.20%
lr103joint	632	722	1273	899	45.51%	34.64%	50.46%	74.75%
lr104joint	609	694	1136	913	41.71%	47.77%	47.19%	74.88%
lr105joint	707	681	1152	884	44.76%	41.16%	48.94%	74.02%
lc201joint	530	512	1015	624	46.15%	45.68%	48.81%	86.29%
lc202joint	533	508	924	600	46.18%	42.35%	46.21%	78.87%
lc203joint	529	486	936	594	38.21%	47.74%	42.27%	71.32%
lc204joint	527	513	954	607	42.68%	44.83%	47.19%	78.32%
lc205joint	521	519	920	618	40.47%	43.11%	37.82%	72.90%

Table 10: Individual results of `lrc` instances for delivery, pickup, backhauling, and mixing set-ups

	Distance in set-up				*Load factor in set-up*			
Instance	Delivery	Pickup	Backh.	Mixing	Delivery	Pickup	Backh.	Mixing
lrc101	661	636	961	914	47.02%	48.97%	56.23%	63.88%
lrc102	677	650	1091	974	43.57%	51.95%	56.58%	73.04%
lrc103	631	634	1014	871	44.38%	52.30%	53.99%	74.90%
lrc104	646	634	1088	789	38.92%	49.90%	50.23%	69.85%
lrc105	668	628	979	830	46.38%	52.03%	61.46%	74.27%
lrc106	667	604	1029	888	46.45%	48.61%	53.95%	71.50%
lrc107	635	639	986	890	50.66%	59.38%	60.78%	77.89%
lrc201	482	480	900	663	46.87%	47.74%	42.09%	58.91%
lrc202	491	485	908	655	34.00%	45.30%	42.66%	60.26%
lrc203	508	434	886	654	37.44%	43.13%	41.66%	53.88%
lrc204	465	472	906	655	44.77%	43.88%	42.56%	60.83%
lrc205	486	469	884	656	36.80%	41.21%	47.43%	60.96%
lrc206	507	465	929	656	40.75%	49.37%	50.54%	65.95%
lrc207	467	482	900	656	41.98%	42.38%	47.56%	70.78%

Table 11: Individual results of **lc** and **lr** instances with unequal demand and supply for delivery, pickup, backhauling, and mixing set-ups

	Distance in set-up				*Load factor in set-up*			
Instance	Delivery	Pickup	Backh.	Mixing	Delivery	Pickup	Backh.	Mixing
lc101uneq	407	572	781	625	39.67%	57.44%	49.94%	58.24%
lc102uneq	428	530	692	597	50.90%	50.98%	52.53%	63.20%
lc104uneq	330	539	698	605	45.23%	51.24%	43.12%	50.49%
lc105uneq	402	572	694	635	39.30%	53.23%	41.28%	40.18%
lc201uneq	296	443	645	497	22.55%	45.64%	29.83%	47.92%
lc202uneq	357	418	727	525	32.98%	38.63%	28.99%	50.13%
lc203uneq	358	432	739	523	29.65%	38.94%	25.20%	40.87%
lc204uneq	327	449	678	528	28.86%	39.13%	32.42%	54.35%
lc205uneq	342	444	726	508	29.63%	39.08%	33.51%	43.17%
lr101uneq	299	530	737	593	48.58%	46.80%	43.97%	57.57%
lr102uneq	356	557	792	696	36.26%	46.11%	42.24%	61.21%
lr103uneq	376	498	785	677	53.05%	42.05%	42.76%	66.04%
lr104uneq	343	521	781	639	33.51%	44.20%	39.32%	55.04%
lr105uneq	346	553	783	646	45.56%	44.95%	35.76%	61.08%
lr202uneq	290	456	713	545	14.09%	45.17%	33.98%	54.94%
lr203uneq	271	428	663	558	17.18%	40.08%	29.65%	50.48%
lr204uneq	273	472	651	548	16.98%	37.65%	27.25%	50.66%
lr205uneq	292	403	652	509	19.76%	39.66%	33.51%	59.15%