

Decentralized Combinatorial Auctions for Dynamic and Large-Scale Collaborative Vehicle Routing

Los, Johan; Schulte, Frederik; Gansterer, Margaretha; Hartl, Richard F.; Spaan, Matthijs T.J.; Negenborn, Rudy R.

DOI

[10.1007/978-3-030-59747-4_14](https://doi.org/10.1007/978-3-030-59747-4_14)

Publication date

2020

Document Version

Final published version

Published in

Computational Logistics

Citation (APA)

Los, J., Schulte, F., Gansterer, M., Hartl, R. F., Spaan, M. T. J., & Negenborn, R. R. (2020). Decentralized Combinatorial Auctions for Dynamic and Large-Scale Collaborative Vehicle Routing. In E. Lalla-Ruiz, M. Mes, & S. Voß (Eds.), *Computational Logistics: Proceedings of the 11th International Conference, ICCL 2020* (pp. 215-230). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 12433 LNCS). Springer. https://doi.org/10.1007/978-3-030-59747-4_14

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Decentralized Combinatorial Auctions for Dynamic and Large-Scale Collaborative Vehicle Routing

Johan Los^{1(✉)}, Frederik Schulte¹, Margaretha Gansterer², Richard F. Hartl³,
Matthijs T. J. Spaan⁴, and Rudy R. Negenborn¹

¹ Department of Maritime and Transport Technology,
Delft University of Technology, Mekelweg 2, 2628 CD Delft, The Netherlands
{J.Los,F.Schulte,R.R.Negenborn}@tudelft.nl

² Department of Operations, Energy, and Environmental Management,
University of Klagenfurt, Universitätsstraße 65–67, 9020 Klagenfurt, Austria
Margaretha.Gansterer@aau.at

³ Department of Business Decisions and Analytics, University of Vienna,
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria
Richard.Hartl@univie.ac.at

⁴ Department of Software Technology, Delft University of Technology,
Van Mourik Broekmanweg 6, 2628 XE Delft, The Netherlands
M.T.J.Spaan@tudelft.nl

Abstract. While collaborative vehicle routing has a significant potential to reduce transportation costs and emissions, current approaches are limited in terms of applicability, unrealistic assumptions, and low scalability. Centralized planning generally assumes full information and full control, which is often unacceptable for individual carriers. Combinatorial auctions with one central auctioneer overcome this problem and provide good results, but are limited to small static problems. Multi-agent approaches have been proposed for large dynamic problems, but do not directly take the advantages of bundling into account. We propose an approach where participants can individually outsource orders, while a platform can suggest bundles of the offered requests to improve solutions. We consider bundles of size 2 and 3 and show that travel costs can be decreased with 1.7% compared to the scenario with only single order auctions. Moreover, experiments on data from a Dutch transportation platform company show that large-scale collaboration through a platform results in system-wide savings of up to 79% for 1000 carriers.

Keywords: Collaborative vehicle routing · Collaborative transportation · Platform-based transportation · Combinatorial auctions · Multi-Agent System · Bundling · Logistics · Dynamic Pickup and Delivery Problem

1 Introduction

Horizontal collaboration is an effective approach to increase transportation efficiency (Verdonck et al. 2013; Gansterer and Hartl 2018b; Pan et al. 2019). While

traditional collaborative vehicle routing focuses on exchange of orders between limited numbers of carriers, the rise of transportation platform companies allows large-scale cooperations, where shippers could directly connect to the platform as well. This raises the need for systems that can assist in both allocating and reallocating orders to carriers in real time without having direct control over the cooperative (but nevertheless rational) participants.

Centralized collaboration approaches have been studied to assess the possible gains of collaboration (Fernández et al. 2018; Molenbruch et al. 2017; Schulte et al. 2017), but these generally make the assumption of complete control and full information availability—which cannot always be assumed in real-world applications. Decentralized approaches with a central auctioneer, and combinatorial auctions in particular (Berger and Bierwirth 2010; Gansterer and Hartl 2018a), overcome these problems, but available computational studies are limited to static problems with small numbers of carriers and orders. For order allocation in larger dynamic problems, Multi-Agent Systems (MASs) have been used, where shippers iteratively offer jobs in auctions, and carriers can bid on them (Máhr et al. 2010; Mes et al. 2013; Los et al. 2020). Although different improvements to the basic auction system have been proposed, interaction effects of orders have not been considered, to the best of our knowledge. Offering bundles of orders, however, is relevant if the individual orders could not be accepted when they are offered in sequence, but are profitable if they are combined. We expect that offering bundles within a MAS could improve results, while the extra effort for carriers to compute a bid on a bundle is limited.

Further, MASs are generally used for allocation, rather than reallocation of orders, although they are suitable for both. By focusing completely on reallocation of orders among carriers, we are able to investigate the possible gains of cooperation among a large number of carriers, while such gains have only been investigated for cooperation between a few carriers so far.

Hence, the contribution of this article is twofold. First, we explore whether and to which extent applying bundling principles to multi-agent approaches can improve results for large-scale dynamic settings. Second, we investigate the possible gains of cooperation among a large number of carriers using this approach.

2 Related Work

Two main collaborative vehicle routing research areas have been distinguished: centralized collaboration and decentralized collaboration (Gansterer and Hartl 2018b, 2020).

Centralized collaboration models mainly assume a set of orders for each carrier and compute what gains could theoretically be obtained if orders are shared. Approximation algorithms are used to compare the solution where each carrier performs only its own orders and the solution where (part of) the orders can be exchanged. It is assumed that all required information is known, which might be difficult in practice. Centralized collaboration models have been developed for different applications: Fernández et al. (2018) consider a problem where

customers request service from different companies and will be attended by only a subset of these companies. Molenbruch et al. (2017) study cooperation of different dial-a-ride providers. Montoya-Torres et al. (2016) compare a non-cooperative and a cooperative scenario for a specific case of city logistics. The number of cooperating carriers in the computational studies, however, ranges from 2–4 (see Table 1). Schulte et al. (2017) use larger instances of up to 50 carriers to investigate emission reductions by carrier cooperation in port-related truck operations.

Table 1. Overview of collaborative transportation approaches.

Cat	Reference	R	A	T	P	L	#Ord	#Carr	#Veh	I	B
CC	Fernández et al. (2018)	✓				✓	18–30	2	∞		
	Molenbruch et al. (2017)	✓	✓	✓	✓	✓	400	4	32		
	Montoya-Torres et al. (2016)	✓					61	3	3		
	Schulte et al. (2017)	✓		✓	✓		10–75	4–50	4–50		
DC	Berger and Bierwirth (2010)	✓			✓	✓	<100	3	3		✓
	Dai et al. (2014)	✓		✓	✓	✓	15–24	3	3–30		✓
	Gansterer and Hartl (2018a)	✓			✓	✓	30–210	3	3		✓
	Gansterer et al. (2020)	✓			✓	✓	30–90	3–6	9–18		✓
	Lai et al. (2017)	✓			✓		30–245	3–24	∞		✓
	Li et al. (2015)	✓			✓	✓	9–15	3	6		✓
	Lyu et al. (2019)	✓		✓	✓	✓	9–45	3	9		✓
	Wang and Kopfer (2014)	✓		✓	✓	✓	104–266	2–5	19–61		✓
	Wang and Kopfer (2015)	✓		✓	✓		~1767	NAv	NAv		✓
DL	Dai and Chen (2011)	✓		✓	✓	✓	9	3	3–30		✓
	Figliozzi et al. (2004)		✓	✓	✓		NAv	4	8		
	Figliozzi et al. (2005)		✓	✓	✓		NAv	NAv	4		
	Van Lon and Holvoet (2017)		✓	✓	✓		120–1200	NAp	10–100		✓
	Los et al. (2020)		✓	✓	✓	✓	1000	150	150		✓
	Máhr et al. (2010)		✓	✓	✓		65	NAp	40		✓
	Mes et al. (2013)		✓	✓	✓		NAv	10	10		✓
This article		✓	✓	✓	✓	✓	1000–2000	1–1000	150–1000	✓	✓

CC: Centralized collaboration; **DC:** Decentralized collaboration with central auctions; **DL:** Decentralized collaboration with local auctions; **R:** Reallocation of orders; **A:** Allocation of unassigned orders; **T:** Time windows; **P:** Pickups and deliveries; **L:** Less than truckload; **#Ord:** Number of orders; **#Carr:** Number of carriers; **#Veh:** Number of vehicles; **I:** Iterative auctions; **B:** Bundling of orders.

Within the literature on decentralized collaboration, two approaches can be distinguished: decentralized collaboration with central auctions and decentralized collaboration with local auctions (see Table 1).

Decentralized collaboration with central auctions assumes that one central auctioneer interacts with all carriers but does not have complete information. An advantage is that the auctioneer can give some guarantees, e.g., it can ensure that all orders are assigned by solving the winner determination problem. The generally large complexity of such subproblems for the coordinator, however, restricts the size of instances that can be solved. In combinatorial auctions (Berger and Bierwirth 2010; Gansterer and Hartl 2018a; Gansterer et al. 2020), each carrier submits unprofitable orders to the auctioneer. To reduce complexity, the auctioneer proposes only a limited subset of attractive bundles of these orders, and

all carriers bid on them. The auctioneer then computes the optimal assignment. Different iterative variants where bundles of orders are considered and the auctioneer finally determines a solution based on the information of different carriers have been studied by Dai et al. (2014), Lyu et al. (2019) and Wang and Kopfer (2014, 2015) (see Table 1). Other variants where bids are made only for single orders have been considered by Lai et al. (2017) and Li et al. (2015).

In decentralized collaboration with local auctions, no central auctioneer is considered. In contrast, any actor can act as auctioneer at any time by starting an auction on (part of) the order(s) that it is responsible for. Hence, local improvements can be made without guarantees on the feasibility of other orders and on global solution quality. Consequently, quick adjustments in dynamic large-scale problems are possible. Generally, this approach is used for allocation of orders to carriers (or even to separate vehicles of one carrier), but Dai and Chen (2011) apply it for reallocation as well (see Table 1). Máhr et al. (2010) and Van Lon and Holvoet (2017) consider MASs with local auctions to examine whether such a decentralized approach can outperform centralized approaches, without focusing on incentives for different carriers. Different carrier strategies and learning mechanisms are considered by Figliozzi et al. (2004, 2005). Mes et al. (2013) investigate the interaction of several look-ahead policies for shippers and carriers, namely delaying commitments, breaking commitments, and valuation of opportunities with respect to future orders. Los et al. (2020) examine the value of information sharing in a MAS.

While bundling of orders is often considered in decentralized collaboration approaches with central auctions, this is not the case for decentralized collaboration approaches with local auctions. Since interaction effects could be relevant in these approaches as well, the focus of this article is on including combinatorial aspects within local auctions. Furthermore, we approximate possible cooperation gains for up to 1000 carriers with our decentralized approach, while the current centralized collaboration literature considers significantly lower carrier numbers.

3 Problem Description

We consider a transportation platform that connects shippers and carriers, and improves routes by allowing carriers to outsource orders. We focus on a dynamic Pickup and Delivery Problem (PDP) where an order either is submitted to the platform by the shipper, or is directly assigned to a specific carrier. In the later case, the contracted carrier can be seen as the owner—the original shipper is then irrelevant. The platform organizes auctions to contract carriers for the unassigned orders. Furthermore, carriers cooperate in the sense that already contracted orders can be sold to other carriers that can deliver them cheaper.

A problem instance consists of a set of shippers S , a set of carriers C , a set of orders O_s for each shipper $s \in S$, a set of initially assigned orders O_c for each carrier $c \in C$ (with $O_S = \bigcup_{s \in S} O_s$ the total set of unassigned orders, $O_C = \bigcup_{c \in C} O_c$ the total set of assigned orders, and $O = O_S \cup O_C$ the total set of orders), and a set of capacitated vehicles V_c for each carrier $c \in C$ (with $V = \bigcup_{c \in C} V_c$ the total set of vehicles).

Each order $o \in O$ represents a load of a certain quantity that must be transported from a pickup location p_o to a delivery location d_o . The pickup or delivery, taking a certain service duration, must start in a time window $[e_i, l_i]$, for $i \in \{p_o, d_o\}$. The release time r_o denotes when the order becomes known to the system. For $o \in O_S$, a reservation price f_o is defined, i.e., a maximum value that the shipper is willing to pay for transportation.

Each vehicle $v \in V$ has an availability time window $[e_v, l_v]$; it becomes available at the initial location α_v at e_v and needs to be at the end location ω_v at l_v . All vehicles $v \in V_c$ become known to the system at the release time r_c of the carrier.

Travel time and travel costs from location i to location j are assumed to be identical for all vehicles, and are denoted by t_{ij} and z_{ij} , respectively. By τ , we denote the time horizon of the problem instance.

A (temporary) solution at time t for a problem instance is given by a set of routes $R^t = \{\langle \rho^{1t} \rangle, \dots, \langle \rho^{|V|t} \rangle\}$, where each route (plan) $\langle \rho_i^{vt} \rangle_{i=1}^{m^{vt}}$ is a sequence of m^{vt} locations representing the (partially completed) path of vehicle v at time t , respecting time, capacity, and precedence constraints. A formal description can be found in Los et al. (2020).

Individual shippers have the goal of outsourcing their orders at a price as low as possible, but not exceeding their reservation prices. Carriers have the goal of maximizing profit, and do this by accepting and outsourcing orders such that the differences between the payment (made to them in case the order is accepted, or paid by them in case the order is outsourced) and the marginal travel costs for the orders are maximized. Together, they contribute to the global goal of obtaining a final solution R^τ with minimal total travel costs.

4 Auction Approach

We propose a multi-agent approach where orders are iteratively offered in reverse auctions. All available carriers (acting as sellers of service) can bid for them, and the carrier with lowest bid wins the auction: it pays the price of its bid, and becomes responsible for filling the order. In contrast to previous approaches (Máhr et al. 2010; Mes et al. 2013; Los et al. 2020), we do not restrict an auctioneer to be a shipper or carrier offering a separate order: we introduce bundle auctioneers as well, denoted by A_B , offering a group of orders $B \subseteq O$. The orders within a bundle are not necessarily owned by the same shipper or carrier, since bundle auctioneers can be generated by the platform.

4.1 Auction Procedure

When order $o \in O$ becomes available at r_o , auctioneer $A_{\{o\}}$ (acting on behalf of shipper s if $o \in O_s$ or acting on behalf of carrier c if $o \in O_c$, but operated by the platform) is initialized and becomes active. Furthermore, the platform immediately generates, if possible, bundle auctioneers A_B with $o \in B$ and $|B| > 1$ (based on similarity of o and previously released orders that are known to the

platform, as we will define in Sect. 4.2) and activates them shortly after $A_{\{o\}}$ has been activated.

When active, auctioneer A_B iteratively organizes auctions. Given a maximum number of auctions a per auctioneer and its activation time r_{A_B} , the time between subsequent auctions is set to $(\min_{o \in B} l_{p_o} - r_{A_B})/a$. The auction at time t then is as follows:

- Each carrier $c \in C$ bids its marginal costs $MC_c^t(B)$ for bundle B at time t , i.e., the extra travel costs for inserting all orders in B into its routes, given the situation at time t . If one or more of the orders in B are already planned in the routes of the carrier, the marginal costs are computed as if these orders were not yet planned.
- A_B compares the bids; let b_0 be the lowest bid provided by carrier c_0 .
- A_B examines the current costs $CC^t(B)$ for bundle B at time t , given by the sum of the marginal costs for assigned orders and the reservation prices for unassigned orders at time t :

$$CC^t(B) = \sum_{c \in C} MC_c^t(B \cap O_c^t) + \sum_{o \in B \cap O_S^t} f_o,$$

where $O_c^t = \{o \in O \mid \exists v \in V_c \exists i \in \{1, \dots, m^{vt}\} \rho_i^{vt} = p_o\}$ is the total set of orders that carrier c has in its route plans at time t and $O_S^t = \{o \in O \mid \neg \exists v \in V \exists i \in \{1, \dots, m^{vt}\} \rho_i^{vt} = p_o\}$ is the set of unassigned orders at time t .

- If $b_0 < CC^t(B)$, the bid is accepted. The platform informs all involved shippers and carriers, who update their contracts and routing plans. Furthermore, the platform receives in total $CC^t(B)$ from the outsourcing shippers and carriers and pays b_0 to the winning carrier c_0 . The gain of $CC^t(B) - b_0$ is divided over the participants as incentive to cooperate, following some profit distribution function.

To avoid unserved orders, the bid is accepted as well if $b_0 \geq CC^t(B)$ and one or more of the orders in B are (due to initial assignment) owned by a carrier who cannot serve them. In this case, they have to cover the difference between b_0 and $CC^t(B)$ with shares proportional to the reservation prices. In other cases when $b_0 \geq CC^t(B)$, no (re)allocations and no payments take place.

When transportation of one of the orders in B starts or the latest pickup time of one of the orders has passed without a contract for that order, A_B becomes inactive.

4.2 Bundle Generation

Selling bundles of orders within a MAS is relevant if for (some of the) individual orders, the best bid is higher than the current costs, while the best bid for the bundle is below the current costs for the bundle. This is likely to happen if orders are close to each other (both in space and time) since they might be combined within the same vehicle route with lower marginal costs.

Relatedness of orders has been defined by Ropke and Pisinger (2006) for PDPs in the context of Large Neighborhood Search (LNS). Since the goal there is to select orders from routes that can be reinserted at each other's places, both pickup locations and delivery locations need to be similar and actual visiting times are compared. For our application, it is already sufficient if one of the locations of one order is similar to one of the locations of the other order and the time windows are not too different. Gansterer and Hartl (2018a) have investigated bundle criteria based on isolation, density and tour length. Isolation, however, is not useful in our context (since we do not require partitions of the complete set of requests) and time windows are not considered in their approach. Hence, we propose a new relatedness measure and bundling procedure that can be applied in the MAS.

We define a relatedness measure $R(o, \hat{o})$ for two orders o and \hat{o} as follows:

$$R(o, \hat{o}) = \min(\text{sim}(p_o, d_{\hat{o}}), \text{sim}(d_o, p_{\hat{o}}), 0.5(\text{sim}(p_o, p_{\hat{o}}) + \text{sim}(d_o, d_{\hat{o}}))), \quad (1)$$

where the similarity of two pickup or delivery locations i and j is defined based on both travel time and time windows:

$$\text{sim}(i, j) = \gamma t_{ij} + W(i, j). \quad (2)$$

Here, W represents the minimal waiting time (due to time window restrictions) at one of the locations if a vehicle serves both locations immediately after each other. Formally,

$$W(i, j) = \max(0, \min(W_D(i, j), W_D(j, i))), \quad (3)$$

where

$$W_D(i, j) = \begin{cases} \infty & \text{if } e_i + t_{ij} > l_j; \\ \max(e_i + t_{ij}, e_j) - \min(l_i + t_{ij}, l_j) & \text{otherwise.} \end{cases} \quad (4)$$

In Eq. 2, γ is a parameter (generally $\gamma > 1$) representing the cost of travel time relative to waiting time. In this article, we use $\gamma = 2$. In Eq. 1, the minimum over three terms is taken. If the pickup of one of the orders is similar to the delivery of the other order, the orders might form a good match, irrespective of the other pickup and delivery locations and times. If, however, both pickup locations are similar, it does matter whether the delivery locations are similar. If they are at opposite directions, combining the orders might appear less useful than if they are similar as well. Hence, the third term in Eq. 1 involves similarity of both pickup and delivery locations.

The platform dynamically generates bundles based on the relatedness measure R . Given a new order o at release time t and the pool of not yet being transported orders O^t , x bundles of size 2 and y bundles of size 3 are generated as follows:

- **Bundles of size 2.** The platform generates bundles $\{o, \hat{o}\}$ for $\hat{o} \in O^t$ and keeps the x bundles with minimal $R(o, \hat{o})$.

- **Bundles of size 3.** The platform generates bundles $\{o, \hat{o}, \check{o}\}$ for $\hat{o}, \check{o} \in O^t$ and keeps the y bundles for which $\min(R(o, \hat{o}) + R(\hat{o}, \check{o}), R(o, \check{o}) + R(\check{o}, \hat{o}), R(o, \hat{o}) + R(o, \check{o}))$ is minimal. Not all three orders have to be highly related to each other to form an attractive bundle, but each order in the bundle needs to be highly related to at least another order in the bundle.

4.3 Theoretical Analysis

By applying bundling, the relatively easy subproblems of the carriers become more difficult. For a carrier $c \in C$ approximating its marginal costs for a bundle B , a basic insertion heuristic that iteratively inserts the order that can be inserted at least costs has a complexity of $\mathcal{O}(|B|^2 |V_c| (l + |B|)^3)$, with l the maximum length of a vehicle route.¹ For single orders, this reduces to $\mathcal{O}(|V_c| l^3)$. In practice, however, a lot of options might be quickly pruned due to time, precedence and capacity constraints.

To get insights into the possible impact of bundling within a MAS, we assume that (estimates of the) real marginal costs are always reported. Although strategic bidding might occur in practice, it is not straightforward (Gansterer and Hartl 2018a). For carriers or shippers mentioning the marginal costs or reservation prices for outsourcing orders, we make the following observations.

- They do not report a value above the true value, since they need to pay this.
- They might report a lower value, but this comes with the risk that the lowest bid b_0 is not lower than $CC^t(B)$, hindering the trade. Indeed, they might report lower values and slightly increase them in next auction rounds, but due to the dynamic environment, there is no guarantee on success.

For carriers placing a bid to acquire a bundle B , we can reason as follows, where $MC_c^t(B)$ denotes the carrier's marginal costs, b_0 denotes the carrier's bid, and g denotes the profit that a winning carrier makes, i.e., g is a fraction of $CC^t(B) - b_0$, dependent on the used profit distribution function.

- They will not bid a value $b_0 < MC_c^t(B)$ if g is expected to be relatively small, since the compensation $b_0 + g$ will not cover the extra costs $MC_c^t(B)$.
- They might place a bid $b_0 < MC_c^t(B)$ if g is expected to be relatively high. If $b_0 + g > MC_c^t(B)$, lowering the bid is a good strategy to outbid another carrier with a bid between b_0 and $MC_c^t(B)$.
- They might bid a value $b_0 > MC_c^t(B)$ to get a higher compensation, but this comes at the risk of not winning the auction anymore.

¹ Per main iteration ($|B|$ in total), the insertion costs for all resulting orders (at most $|B|$) at all routes ($|V_c|$ in total) need to be checked. Insertion of both the pickup and the delivery needs to be checked for each position in the route (which can be up to $l + 2|B| - 2$ positions when the last order of the bundle must be inserted), and a chain of time consistency checks might be necessary along the complete route in the worst case as well.

The approach guarantees that no carrier is worse off per auction. They might, however, be worse off on the long term if they get dynamically revealed yet assigned tasks that produce bad interactions with the tasks they acquired before. Nevertheless, individual rationality is guaranteed if all assigned tasks are known by the carriers beforehand.

5 Computational Study

Our experiments are based on a real-world data set consisting of well over 12000 orders that have been received by the Dutch transportation platform company Quicargo. To investigate the impact of bundling, we created 6 instances of 2000 orders where half of the orders are assigned to carriers (to generate initial routes), and the other half is unassigned and released during operations. All orders were available for (re)allocation. To examine the possible gains of cooperation among a large number of carriers, we created 6 instances of 1000 orders and varied the number of carriers that own them.

The MAS was implemented in Go and all experiments were performed on a single machine with Intel Core i5-750 CPU at 2.67 GHz and 6 GB of RAM. Within our experiments, we used the insertion heuristic described by Campbell and Savelsbergh (2004), adapted to pickup and delivery problems, for computation of marginal costs and construction of initial routes.

5.1 Impact of Bundling

We generated 6 instances of 2000 orders and 150 carriers. Each carrier has a single randomly chosen depot with 1 to 3 vehicles of capacity 13.6 (loading meters). Half of the orders were randomly assigned to one of the 10% closest carriers in terms of distance between pickup location and depot. About 66% of the carriers are always available, and the others have restricted availability times based on the assigned orders. Release times of assigned orders were always set equal to release times of the corresponding carriers, to make sure that routes can be constructed immediately. Pickup and delivery locations (in and close to the Netherlands) and estimated load quantities are as in the original data set. Pickup and delivery time windows have been kept, as well as release times for the unassigned orders, but random shifts of a whole number of days have been made such that all orders fell within a time horizon of 10 days. Travel speed was set to 0.015 km/s and reservation prices were set to 2.5 times the travel costs between pickup and delivery locations.

We ran experiments where only single orders were auctioned, and compared them with runs in which 3 bundles of size 2 and 1 of size 3 were introduced per order. First, we allowed a maximum of 10 auctions per auctioneer. The results in terms of travel costs, the number of rejected orders, and the increase in computation time are given in Table 2 (set 1). Addition of the bundles decreases the travel costs by about 4% on average, but increases the computation time by a factor 10.

Table 2. Results for single order auctions (S) and bundle auctions (B).

Instance			Settings		Results					
I	C	RP	A _S	A _B	TC _S	TC _B	D(%)	R _S	R _B	T
1	150	2.5	10	10 + 4 × 10	9705.58	8955.25	7.73	1	2	6.43
2	150	2.5	10	10 + 4 × 10	9440.72	8997.33	4.70	1	1	9.03
3	150	2.5	10	10 + 4 × 10	9383.03	9151.61	2.47	0	1	10.34
4	150	2.5	10	10 + 4 × 10	9405.44	8923.82	5.12	1	1	11.72
5	150	2.5	10	10 + 4 × 10	8806.27	8716.37	1.02	0	1	11.72
6	150	2.5	10	10 + 4 × 10	8945.61	8639.87	3.42	2	3	12.91
Average of set 1							4.08	0.83	1.5	10.36
1	150	2.5	50	10 + 4 × 10	9253.74	8955.25	3.23	1	2	0.66
2	150	2.5	50	10 + 4 × 10	9169.23	8997.33	1.87	0	1	0.99
3	150	2.5	50	10 + 4 × 10	9149.17	9151.61	-0.03	0	1	1.15
4	150	2.5	50	10 + 4 × 10	9179.95	8923.82	2.79	1	1	1.14
5	150	2.5	50	10 + 4 × 10	8658.51	8716.37	-0.67	0	1	1.23
6	150	2.5	50	10 + 4 × 10	8856.09	8639.87	2.44	1	3	1.35
Average of set 2							1.61	0.5	1.5	1.09
1*	500	2	30	10 + 4 × 5	9185.33	8934.01	2.74	3	0	1.39
2*	500	2	30	10 + 4 × 5	9014.15	9002.89	0.12	1	0	1.85
3*	500	2	30	10 + 4 × 5	9231.21	9058.68	1.87	2	2	1.92
4*	500	2	30	10 + 4 × 5	9273.16	9027.70	2.65	2	3	2.18
5*	500	2	30	10 + 4 × 5	8864.34	8706.58	1.78	1	0	2.19
6*	500	2	30	10 + 4 × 5	8917.12	8828.61	0.99	0	0	2.55
Average of set 3							1.69	1.5	0.83	2.01

I: Instance number (instances marked with a * differ only in number of carriers and reservation prices); **C:** Number of carriers; **RP:** Reservation price factor (relative to travel costs); **A_{S,B}:** Maximum number of auctions per order in the experiments without (S) or with (B) bundling ($a + bc$ denotes a maximum of a auctions for the single order and b times a maximum of c auctions for bundles with the order); **TC_{S,B}:** Total travel costs obtained without (S) or with (B) bundling; **D(%):** Decrease in travel costs for bundling compared to no bundling. **R_{S,B}:** Number of rejected orders without (S) or with (B) bundling; **T:** Computation time increase factor for bundling compared to no bundling.

To check whether the improvements are caused by bundling or rather by the increased total number of auctions, we ran a second experiment where we allowed a maximum of 50 auctions per single order auctioneer, such that both the scenario without bundling and the scenario with bundling have the same maximum total number of auctions (see set 2 in Table 2). Although there is a small reduction in travel costs on average (1.6%) when bundles are used, two individual cases have slightly higher travel costs (instance 3 and 5). Furthermore,

the number of rejected orders is generally larger with bundling than with single order auctions. Hence, no clear improvement due to bundling can be observed within this set.

One might conjecture, however, that the reservation prices are too high to take any advantage of bundling. If reservation prices are lower, it is more likely that serving some individual orders is too expensive, but combinations still could be advantageous. Hence, we ran a third set of experiments where we lowered the reservation price of the orders to 2 times the travel costs between pickup and delivery locations. At the same moment, we increased the number of carriers to 500 to avoid large numbers of rejected orders. The results (Table 2, set 3) show that travel costs can be reduced by about 1.7%, and the number of rejected orders generally does not increase if bundles are used. The computation time for the bundling scenario (about 4 h on average) is twice as high as that of the single order auction scenario (about 2 h).

5.2 Cooperation Gains

To assess the possible gains from cooperation, we created another series of instances where each order is initially assigned to a carrier. We created 6 different instance sets of 1000 orders, and used 4 carrier configurations (1000, 100, 10 carriers, or 1 carrier) and 2 assignment configurations (close or random) per set. Within each set, 50 depots with 20 vehicles each were defined, which were distributed equally among 1000, 100, 10 carriers, or 1 carrier, such that each carrier has 1, 10, 100, or 1000 vehicles, respectively. Each order is always associated with the same depot—the depot closest to the pickup location of the order in case of close assignment, and a random depot in case of random assignment. Then, the order was assigned to a random carrier having vehicles in that depot. All carriers are continuously available, but original order release times were kept. Other properties are the same as described in Sect. 5.1.

Because carriers do not optimize their individual routes (but only use the insertion heuristic), the MAS can serve as a system that reassigns orders to the same carrier if the selling carrier is the same as the lowest bidder, i.e., carriers can improve their own routes through the auction system. To avoid this behaviour in our cooperation gain assessment, we added an individual improvement phase consisting of 50 hill-climbing LNS iterations (see Pisinger and Ropke 2019) after each insertion of an order into a carrier’s plan.

In Fig. 1a, we show the decrease in total travel costs if carriers collaborate compared to the scenario where each carrier only serves its own orders. Savings of up to 79% and 70% (for 1000 carriers, with random and close assignment, respectively) are observed. The savings with close assignments are lower than the savings with random assignments, as expected, but are still rather large.

Note that a 22% gain through self-reassignment can be obtained if only one carrier is present that does not apply individual improvement phases. With individual improvements, however, this is not possible, as expected. Furthermore, gains through self-reassignment are lower for larger numbers of carriers, since the number of orders per carrier is smaller. With only 1 order per carrier (which




Fig. 1. Results for varying numbers of cooperating carriers: (a) decrease in total travel costs compared to a non-cooperative setting; (b) computation times.

is the average for the 1000 carrier instances) the insertion heuristic will already find the best individual solution. Indeed, the individual improvement phase does not significantly change the gains if 1000 carriers collaborate.

When LNS improvement is applied, the computation time (see Fig. 1b) slightly increases if carriers’ subproblems are relatively small (for large numbers of carriers) but increases a lot if the subproblems are more complex (for lower numbers of carriers). In practice, a suitable individual optimization method should be selected, based on the individual carriers’ available resources and time.

To give an indication of advantages for individual carriers, we show the distribution of the gains among the platform and the carriers in Fig. 2, and the minimum and maximum shares of the total gain that are obtained by individual carriers in Table 3. The used profit distribution function in each auction (see Sect. 4.1) is as follows: 30% of the gain was assigned to the winning carrier, 60% to the selling carriers and shippers (proportional to their current costs), and 10% was kept by the platform itself. In case of pure self-reassignment (which sometimes occurred, despite the LNS improvement), no payments were made. The platform’s share of total gains is larger when lower numbers of carriers participate (except for the trivial case of 1 carrier where no payments are made) and larger for close assignments than for random assignments. In extreme cases, a single carrier can obtain 3.9% (1000 carrier instances) to 60.9% (10 carrier instances) of the total profits made. Maximum losses (due to the fact that car-

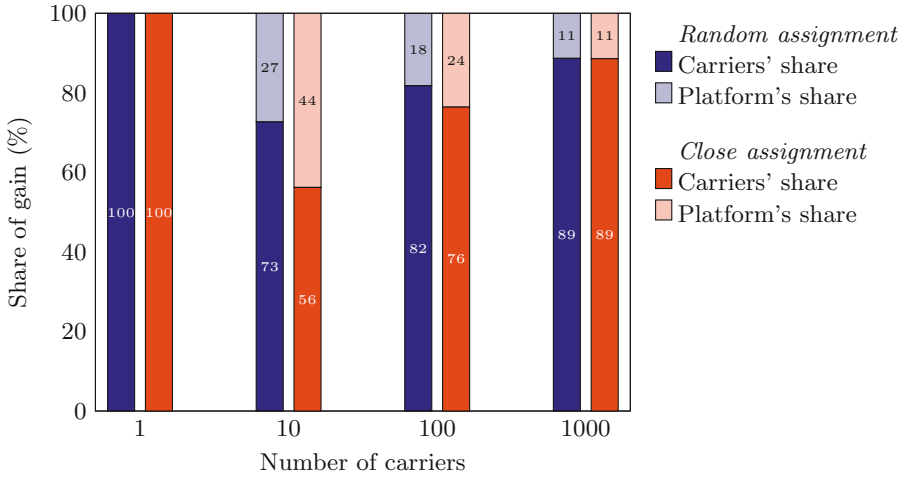


Fig. 2. Shares of the cooperation gains that are obtained by the carrier collective and by the platform.

riers acquire or outsource orders without knowing or even considering future orders that are assigned to them) are considerably lower when orders are randomly assigned (0.1% of the total gains for 1000 carriers up to 16.3% for 10 carriers) than when orders are assigned to the closest carrier (0.3% for 1000 carriers up to 59.1% for 10 carriers). This might be attributed to the high interaction effects of future orders for a carrier: outsourcing an early order might be relatively costly, given similar jobs that later on still appear to be executed.

Table 3. Minimum and maximum percentages of the total gains that are obtained by an individual carrier with single order auctions only (S) and with bundle auctions (B).

Carriers	Setting	1		10		100		1000	
		S	B	S	B	S	B	S	B
Random	Minimum	100.0	100.0	-16.3	-12.0	-1.1	-1.0	-0.1	-0.1
	Maximum	100.0	100.0	52.1	43.6	10.5	14.6	3.5	3.3
Close	Minimum	100.0	100.0	-59.1	-50.7	-14.4	-14.0	-0.3	-0.3
	Maximum	100.0	100.0	54.1	60.9	23.9	22.9	3.9	3.7

6 Conclusions

We have investigated the potential savings by auctioning bundles of orders within a multi-agent approach for dynamic large-scale collaborative vehicle routing.

While combinatorial auctions with a central auctioneer provide good results on small-size static problems, similar bundling approaches have not been applied within multi-agent systems to solve dynamic problems. We considered a platform that dynamically creates bundles of 2 or 3 orders, and auctions them, as well as separate orders, to all carriers.

A computational study based on a real-world data set shows that applying bundling can save 1.7% of travel costs, but that the applicability is highly dependent on the problem characteristics. A structured analysis applying different bundle configurations on instances with various properties should clarify the full potential of local bundling in distributed vehicle routing. In this work, we used one policy to define bundles. Alternatively, different compositions, sizes, and numbers of bundles could also be used.

Moreover, the approach and data set allowed us to investigate the potential of collaboration of a large number of carriers through a platform, while current literature hardly considers more than 3 or 4 carriers. Our preliminary results show that cost reductions increase with more participants. Up to 79% of travel costs can be saved with 1000 cooperating carriers. However, besides the assumptions in problem properties, the quality of carriers' local route optimization could have an impact on the exact savings, especially if their subproblems become more complex. Future work should examine what quality of subproblem solutions is desirable and what resources and time to obtain them are acceptable.

In this article, we assumed true value reporting by carriers and shippers. Although false bidding is not trivial in the proposed setting, an experimental study on possible individual gains by reporting false values is planned as part of future work. Furthermore, second price auctions can be used to eliminate incentives of false bidding on the short term. An interesting question is whether the extra payments could be covered by bundling gains. Finally, we plan to investigate a scenario with mixed levels of autonomy, where shippers and carriers either can be in charge of auctioning orders themselves, or outsource this process to the platform.

Acknowledgements. This research is supported by the project “Dynamic Fleet Management (P14-18 – project 3)” (project 14894) of the Netherlands Organization for Scientific Research (NWO), domain Applied and Engineering Sciences (TTW).

References

- Berger, S., Bierwirth, C.: Solutions to the request reassignment problem in collaborative carrier networks. *Transp. Res. Part E* **46**, 627–638 (2010)
- Campbell, A.M., Savelsbergh, M.: Efficient insertion heuristics for vehicle routing and scheduling problems. *Transp. Sci.* **38**, 369–378 (2004)
- Dai, B., Chen, H.: A multi-agent and auction-based framework and approach for carrier collaboration. *Logist. Res.* **3**, 101–120 (2011)
- Dai, B., Chen, H., Yang, G.: Price-setting based combinatorial auction approach for carrier collaboration with pickup and delivery requests. *Oper. Res. Int. J.* **14**(3), 361–386 (2014). <https://doi.org/10.1007/s12351-014-0141-1>

- Fernández, E., Roca-Riu, M., Speranza, M.G.: The shared customer collaboration vehicle routing problem. *Eur. J. Oper. Res.* **265**, 1078–1093 (2018)
- Figliozzi, M.A., Mahmassani, H.S., Jaillet, P.: Competitive performance assessment of dynamic vehicle routing technologies using sequential auctions. *Transp. Res. Rec.* **1882**, 10–18 (2004)
- Figliozzi, M.A., Mahmassani, H.S., Jaillet, P.: Impacts of auction settings on the performance of truckload transportation marketplaces. *Transp. Res. Rec.* **1906**, 89–96 (2005)
- Gansterer, M., Hartl, R.F.: Centralized bundle generation in auction-based collaborative transportation. *OR Spectr.* **40**(3), 613–635 (2018a). <https://doi.org/10.1007/s00291-018-0516-4>
- Gansterer, M., Hartl, R.F.: Collaborative vehicle routing: a survey. *Eur. J. Oper. Res.* **268**, 1–12 (2018b)
- Gansterer, M., Hartl, R.F.: Shared resources in collaborative vehicle routing. *TOP* **28**(1), 1–20 (2020). <https://doi.org/10.1007/s11750-020-00541-6>
- Gansterer, M., Hartl, R.F., Savelsbergh, M.: The value of information in auction-based carrier collaborations. *Int. J. Prod. Econ.* **221**, 107485 (2020)
- Lai, M., Cai, X., Hu, Q.: An iterative auction for carrier collaboration in truckload pickup and delivery. *Transp. Res. Part E* **107**, 60–80 (2017)
- Li, J., Rong, G., Feng, Y.: Request selection and exchange approach for carrier collaboration based on auction of a single request. *Transp. Res. Part E* **84**, 23–39 (2015)
- Van Lon, R.R.S., Holvoet, T.: When do agents outperform centralized algorithms? A systematic empirical evaluation in logistics. *Auton. Agent. Multi-Agent Syst.* **31**, 1578–1609 (2017)
- Los, J., Schulte, F., Spaan, M.T.J., Negenborn, R.R.: The value of information sharing for platform-based collaborative vehicle routing. *Transp. Res. Part E* **141**, 102011 (2020)
- Lyu, X., Chen, H., Wang, N., Yang, Z.: A multi-round exchange mechanism for carrier collaboration in less than truckload transportation. *Transp. Res. Part E* **129**, 38–59 (2019)
- Máhr, T., Srouf, J., de Weerd, M., Zuidwijk, R.: Can agents measure up? A comparative study of an agent-based and on-line optimization approach for a drayage problem with uncertainty. *Transp. Res. Part C* **18**, 99–119 (2010)
- Mes, M., van der Heijden, M., Schuur, P.: Interaction between intelligent agent strategies for real-time transportation planning. *CEJOR* **21**, 337–358 (2013)
- Molenbruch, Y., Braekers, K., Caris, A.: Benefits of horizontal cooperation in dial-a-ride services. *Transp. Res. Part E* **107**, 97–119 (2017)
- Montoya-Torres, J.R., Muñoz-Villamizar, A., Vega-Mejía, C.A.: On the impact of collaborative strategies for goods delivery in city logistics. *Prod. Plann. Control* **27**, 443–455 (2016)
- Pan, S., Trentesaux, D., Ballot, E., Huang, G.Q.: Horizontal collaborative transport: survey of solutions and practical implementation issues. *Int. J. Prod. Res.* **57**, 5340–5361 (2019)
- Pisinger, D., Ropke, S.: Large neighborhood search. In: Gendreau, M., Potvin, J.-Y. (eds.) *Handbook of Metaheuristics*. ISORMS, vol. 272, pp. 99–127. Springer, Cham (2019). https://doi.org/10.1007/978-3-319-91086-4_4
- Ropke, S., Pisinger, D.: An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transp. Sci.* **40**, 455–472 (2006)
- Schulte, F., Lalla-Ruiz, E., González-Ramírez, R.G., Voß, S.: Reducing port-related empty truck emissions: a mathematical approach for truck appointments with collaboration. *Transp. Res. Part E* **105**, 195–212 (2017)

- Verdonck, L., Caris, A., Ramaekers, K., Janssens, G.K.: Collaborative logistics from the perspective of road transportation companies. *Transp. Rev.* **33**, 700–719 (2013)
- Wang, X., Kopfer, H.: Collaborative transportation planning of less-than-truckload freight. *OR Spectr.* **36**(2), 357–380 (2014). <https://doi.org/10.1007/s00291-013-0331-x>
- Wang, X., Kopfer, H.: Rolling horizon planning for a dynamic collaborative routing problem with full-truckload pickup and delivery requests. *Flex. Serv. Manuf. J.* **27**(4), 509–533 (2015). <https://doi.org/10.1007/s10696-015-9212-8>