The value of information sharing for platform-based collaborative vehicle routing

Los, Johan; Schulte, Frederik; Spaan, Matthijs T.J.; Negenborn, Rudy R.

DOI
10.1016/j.tre.2020.102011

Publication date
2020

Document Version
Final published version

Published in
Transportation Research Part E: Logistics and Transportation Review

Citation (APA)
Los, J., Schulte, F., Spaan, M. T. J., & Negenborn, R. R. (2020). The value of information sharing for platform-based collaborative vehicle routing. *Transportation Research Part E: Logistics and Transportation Review*, *141*, Article 102011. https://doi.org/10.1016/j.tre.2020.102011

Transportation Research Part E 141 (2020) 102011

Contents lists available at ScienceDirect

Transportation Research Part E

journal homepage: www.elsevier.com/locate/tre

The value of information sharing for platform-based collaborative vehicle routing

Johan Los[a,*], Frederik Schulte[a], Matthijs T.J. Spaan[b], Rudy R. Negenborn[a]

[a] Department of Maritime and Transport Technology, Delft University of Technology, Mekelweg 2, 2628 CD Delft, the Netherlands
[b] Department of Software Technology, Delft University of Technology, Van Mourik Broekmanweg 6, 2628 XE Delft, the Netherlands

ARTICLE INFO

Keywords:
Collaborative vehicle routing
Collaborative transportation
Information sharing
Platform-based transportation
Multi-agent system
Logistics
Dynamic pickup and delivery problem

ABSTRACT

Cooperation is important in order to find efficient vehicle routing solutions for the growing transportation market. Increasingly, platforms emerge as facilitators for this kind of collaborative transportation. However, individual actors connected to a platform might refuse to share (parts of) their information due to reasons of competition. Though the need for realistic information sharing models is widely acknowledged by transportation researchers and practitioners, the precise value of such information is mostly unknown. We investigate the quality of solutions that can be obtained when different types and levels of carrier data are available. We consider an auction-based Multi-Agent System to solve large-scale, dynamic pickup and delivery problems, and vary whether carriers' positions or route plans are available, and whether carriers are fully cooperative or more competitive in placing their bids by sharing or hiding their marginal costs. In total, we evaluate 9 different information sharing policies. The availability of vehicle position and route plan information turns out to be important for decreasing total route costs, and sharing marginal costs has a positive impact on service level. We provide detailed insights into trade-offs of carriers' confidentiality concerns and a range of system performance objectives (service level, travel costs, and carrier profits) under different circumstances (various numbers of auctions, penalties for rejected orders, emission or congestion fees, and different problem characteristics). Based on these results, platform providers can stimulate sharing of certain information to improve the total system efficiency.

1. Introduction

The increasing demand for transportation requires efficient and robust vehicle routing. Providing a reliable and quick transportation service while minimizing costs, energy consumption, and pollution is important for carriers, customers, and society in general. A significant gain can be obtained when different carriers cooperate: combining deliveries can reduce travel costs, pollution, and congestion. Various studies have shown improvements of about 20–30% in comparison with non-collaborative situations (Gansterer and Hartl, 2018b). For individual carriers, even profit improvements of up to 800% may be achieved if they cooperate in larger coalitions (Schulte et al., 2019).

With the rise of platform companies (e.g., UberFreight, BlackBuck, and Quicargo), cooperative planning is becoming even more important. Such platforms act as intermediary between carriers and customers, to connect transport supply and demand dynamically without having direct control over these actors. Hence, incentives for cooperation need to be provided to both customers and carriers.

A limitation of the types and amounts of operational information to be shared with the platform might be one of these incentives. Carriers can be requested to reveal part of their private information, such as vehicle availability, intended routes, or costs for

* Corresponding author.
E-mail addresses: J.Los@tudelft.nl (J. Los), F.Schulte@tudelft.nl (F. Schulte), M.T.J.Spaan@tudelft.nl (M.T.J. Spaan), R.R.Negenborn@tudelft.nl (R.R. Negenborn).

https://doi.org/10.1016/j.tre.2020.102011
Received 7 October 2019; Received in revised form 27 April 2020; Accepted 13 June 2020
1366-5545/

Table 1
Overview of information sharing properties in collaborative transportation. Explicit comparisons of different levels of information sharing are given in boldface.

		Problem characteristics							Dynamism		Information sharing				
Category	Reference	IA[1]	TW[2]	PD[3]	LTL[4]	#Ord[5]	#Carr[6]	#Veh[7]	Ord[8]	Veh[9]	A[10]	N[11]	It[12]	Bu[13]	Description
Centralized optimization	Dahl and Derigs (2011)	✓	✓	✓	✓	NAv[14]	~50	8905	✓		F	C		NAp[15]	full information; system suggests exchanges
	Molenbruch et al. (2017)	✓	✓	✓	✓	400	4	32			F	C		NAp	full information; system computes solution
	Montoya-Torres et al. (2016)	✓			NAp	61	3	3			F	C		NAp	full information; system computes solution
	Schulte et al. (2017)	✓	✓	✓		10–75	4–50	4–50			F	C		NAp	full information; system computes solution
Centralized support	Berger and Bierwirth (2010)	✓		✓	NAp	< 100	3	3			P	C		✓	marginal profit for (bundles of) orders
	Dai et al. (2014)	✓	✓	✓	✓	15–24	3	3–30			P	C	✓	✓	yes–no bids + **variable: profit**
	Gansterer and Hartl (2018a)	✓		✓	✓	30–210	3	3			P	C		✓	marginal profit for bundles
	Gansterer et al. (2020)	✓		✓	✓	30–90	3–6	9–18			P	C		✓	marginal profit for bundles + **variable: aggregate geographical data, marginal profit for requests**
	Lai et al. (2017)	✓		✓		30–245	3–24	∞			P	C	✓		request portfolio, marginal profit for order
	Li et al. (2015)	✓		✓	✓	9–15	3	6			P	C	✓		marginal profits for outsourcing and sourcing orders
	Lyu et al. (2019)	✓	✓	✓	✓	9–45	3	9			P	C	✓	✓	outsourcing prices, yes–no bids + **variable: profit**
	Wang and Kopfer (2014)	✓	✓	✓	✓	104–266	2–5	19–61			P	C	✓	✓	number of vehicles, costs of candidate routes
	Wang and Kopfer (2015)	✓	✓	✓		~1767	NAv	NAv	✓		P	C	✓	✓	number of vehicles, costs of candidate routes
Distributed development	Dai and Chen (2011)	✓	✓	✓	✓	9	3	3–30	✓		P	D	✓		outsourcing prices, yes–no bids
	Figliozzi et al. (2004)		✓	✓		NAv	4	8	✓		P	D			marginal cost for order
	Figliozzi et al. (2005)		✓	✓		NAv	NAv	4	✓		P	D			adapted marginal cost for order (also broadcasted)
	Gath (2016)		✓		✓	100–200	NAp	3–29			P	D	✓		marginal cost for order
	Máhr et al. (2010)		✓	✓		65	NAp	40	✓	✓	P	D	✓		marginal cost for order
	Mes et al. (2013)		✓	✓		NAv	10	10	✓		P	D	✓		marginal cost for order (including opportunity costs)
	This article		✓	✓	✓	1000	150	150	✓	✓	P	D	✓		**variable: marginal cost for order, location information**

[1] **IA**: Initial assignment of orders to carriers.
[2] **TW**: Time windows.
[3] **PD**: Pickups and deliveries.
[4] **LTL**: Less than truckload problem.
[5] **#Ord**: Number of orders.
[6] **#Carr**: Number of carriers.
[7] **#Veh**: Number of vehicles.
[8] **Ord**: Dynamic orders.
[9] **Veh**: Dynamic vehicles.
[10] **A**: Amount of information sharing (**F**: full information sharing; **P**: partial information sharing).
[11] **N**: Nature of information sharing (**C**: centralized information sharing, i.e., with central coordinator; **D**: distributed information sharing, e.g., with other carriers or with customers).
[12] **It**: Iterative information sharing (multiple rounds).
[13] **Bu**: Information regarding bundles of orders (combinatorial approach).
[14] **NAv**: Not available.
[15] **NAp**: Not applicable.

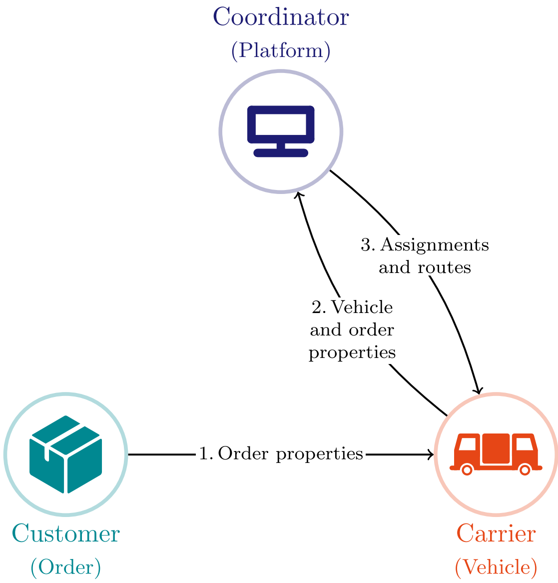

Fig. 1. Communication between coordinator, customer, and carrier in a centralized optimization scenario.

satisfying a request to enable cooperation on a platform. However, they might be reluctant to share all confidential information due to competition or legislation – a problem referred to as coopetition (Cruijssen et al., 2007; Cleophas et al., 2019). On the other hand, all carrier information may not always be needed. In highly dynamic scenarios, there is generally no time for computing optimal solutions. Reasonable approximations based on partial information may be of great value in these situations.

The precise effect of available carrier information on solution quality, however, is not known. Current research on collaborative transportation mainly applies centralized models or combinatorial auctions (Gansterer and Hartl, 2018b), in which a central coordinator collects information and defines a solution. These centralized approaches (which are also applied in current transportation platforms) require certain fixed levels of information sharing with a central authority. Multi-agent approaches, on the other hand, allow to model different information sharing levels in a flexible manner, even without an active central authority (customers and carriers might need the platform mainly for getting to know each other). Nonetheless, current multi-agent approaches have considered either fully cooperative carriers, i.e., they accept orders that are unprofitable for themselves if the global system is better off (Máhr et al., 2010; Gath, 2016), or fully competitive carriers, i.e., they just try to maximize their own profit (Figliozzi et al., 2004; Figliozzi et al., 2005). In a realistic system, on the other hand, carriers might want to find a balance between cooperation and competition by sharing a limited amount of information.

In this article, we investigate the value of sharing different types and amounts of carrier information for solution quality in distributed dynamic vehicle routing problems. We focus on the allocation of orders that are submitted to a platform among cooperating carriers associated to the platform. We develop a Multi-Agent System (MAS) in which order agents and vehicle agents interact through auctions for this Dynamic Pickup and Delivery Problem (DPDP). In a series of simulations, we focus on two types of information: we vary what is known about a vehicle's position and route plan, and what cost information a vehicle agent includes in its bid. The results show trade-offs between carriers' tendency to protect information and system performance based on service level, travel costs, and profits. Although the approach abstracts from direct incentives and explicit gain sharing, a system-wide improvement in solution quality is observed with increased information sharing, leaving room for the design of incentives to share the relevant information.

2. Related work

Horizontal collaboration between carriers has extensively been considered in the literature (see, e.g., the reviews by Cruijssen et al., 2007; Verdonck et al., 2013; Gansterer and Hartl, 2018b; Cleophas et al., 2019; Pan et al., 2019). Although several collaboration aspects (e.g., types of collaboration, request exchange mechanisms, gain sharing) have been investigated in depth, the issue of information sharing has hardly gotten any attention. Most authors mention that information sharing might be difficult in practice due to competition and privacy regulations, but disregard the implications for their approach because they focus on other aspects. Some authors emphasize that their approach does not require full information sharing, but do not state explicitly what information is being exchanged in their approach and what can be inferred from this information.

We give an overview of the information sharing properties of several recent articles on collaborative vehicle routing and compare them in Table 1. The table is divided into three parts. The first part contains articles assuming full information sharing with a central coordinator. Generally, these studies assume that centralized optimization is possible, and compare this scenario with a non-cooperative scenario.

The second part contains approaches assuming partial information exchange between carriers and a central coordinator. Although not all information is shared, the coordinator can contribute to a solution by, e.g., proposing order exchanges or iteratively updating prices. Hence, the central coordinator supports the process of cooperation but has no full control.

The third part considers partial information exchange on a local level, e.g., between carriers among each other or between carriers

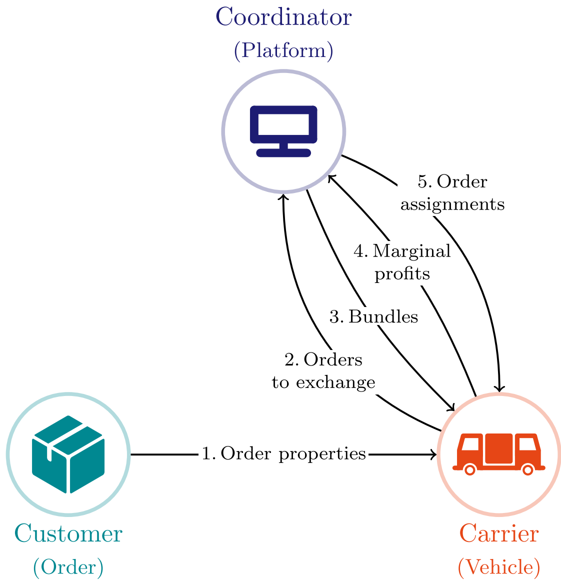

Fig. 2. Communication between coordinator, customer, and carrier in a combinatorial auction, a typical example of centralized support.

and customers. Generally, solutions are developed in a distributed way, while a platform is mainly needed for connecting the different participants and for providing a reliable communication framework. Still, the platform can be in charge of minor but important decisions.

Throughout the three categories, we furthermore distinguish between methods with iterative information sharing and methods with a single round of information exchange, and between methods where aggregate information is being exchanged (information regarding bundles of requests) and methods with information exchange on the level of a single request.

2.1. Full information sharing with central coordinator

In approaches with full information sharing, a central coordinator is aware of all information about orders and vehicles (see Fig. 1). While this coordinator computes and imposes a global solution in most approaches (Montoya-Torres et al., 2016; Schulte et al., 2017; Molenbruch et al., 2017), it suggests order exchanges in the approach of Dahl and Derigs (2011). Individual carriers might accept or deny these exchanges, and hence stay autonomous.[1] The majority of studies in this area focuses on the difference in profit between a scenario without cooperation (i.e., each carrier serves its own orders), and a fully cooperative scenario for which full information sharing is necessary. Although the initial assignment of orders to carriers is fixed in general, Molenbruch et al. (2017) compare different initial assignments for dial-a-ride problems.

2.2. Partial information sharing with central coordinator

Several methods apply partial information sharing with a central coordinator (see Fig. 2). The coordinator has no full control, but has a high-level overview over (part of) the problem since it collects information from several carriers. It supports the carriers by proposing or establishing profitable collaborations.

In combinatorial auctions with a central auctioneer, each carrier can submit some requests to a request pool, and subsequently bid its marginal profit for the bundles of requests that the auctioneer generated (Berger and Bierwirth, 2010; Gansterer and Hartl, 2018a; Gansterer et al., 2020).[2] Hence, aggregated profit information is exchanged only once with a central coordinator.

Li et al. (2015) also consider a central coordinator, but focus on single request exchange. Carriers submit their request with lowest marginal profit to a pool, and subsequently bid their marginal profit for all requests in the pool. In each iteration, the coordinator selects the request with highest profit gain for exchange. Hence, carriers iteratively need to share marginal profit information with the coordinator.

Lai et al. (2017) reverse the order: carriers may bid their marginal profit on single requests that they want to obtain, and based on the demand and price, every carrier decides which bundle of requests it is willing to sell. The coordinator then tries to find the best match. A drawback of this approach is that all carriers have to disclose their full request portfolio.

Dai et al. (2014) consider a central auctioneer that proposes outsourcing prices for requests, and iteratively updates the prices based on the bidding behaviour of the carriers. The auctioneer may adapt the prices with or without knowing profit information from the carriers.

Lyu et al. (2019) propose a similar system: carriers determine outsourcing prices for bundles of requests, and may bid on all bundles in the request pool. An auctioneer then solves the winner determination problem, either based on profit information from the

[1] Hence, the approach of Dahl and Derigs (2011) tends to the category of centralized support.

[2] Berger and Bierwirth (2010) also consider the situation where each carrier bids for each individual request in the pool, i.e., a setting without bundles.

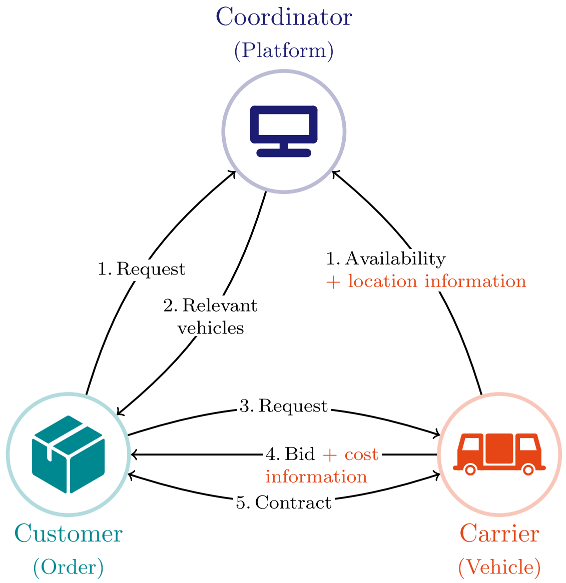

Fig. 3. Communication between coordinator, customer, and carrier in our multi-agent approach, a typical example of distributed solution development. Location information sharing and cost information sharing are optional within our approach.

carriers, or without this information.

In the approach of Wang and Kopfer (2014, 2015), carriers submit requests that they want to exchange to a pool, and then iteratively propose prices for bundles of requests (to be served in one route) that they want to obtain. The information that is being exchanged is comparable to that in combinatorial auctions, but carriers themselves can propose the bundles now. However, they also need to reveal their number of vehicles to the central coordinator, since the assigned number of bundles per carrier may not exceed the number of available vehicles.

2.3. Distributed partial information sharing

Whereas all communications go via a central coordinator in the previously described approaches, a third category of information sharing is fundamentally different: interaction takes place directly between the cooperating participants, without the intervention of a central coordinator (see Fig. 3). The main role of the coordinator is to connect the participants to each other, while the participants themselves locally develop a solution. However, the coordinator can have a crucial influence in the selection of participants in large-scale scenarios.

In the approach of Dai and Chen (2011), carriers offer their requests at an outsourcing price to all other carriers. Based on the bids of other carriers, the outsourcing price is adapted, or the request is exchanged. Hence, carriers do not include profit information in their bids, but they reveal private information by iteratively communicating outsourcing prices.

Although most approaches assume initial assignment of requests to carriers, Máhr et al. (2010) and Gath (2016) consider cooperative approaches where orders are not initially assigned. In their approaches, orders act as individual agents that try to find the most appropriate vehicle by organizing auctions, and carriers share their marginal costs only with the order agents. Although no direct cost information is being exchanged between carriers, they might infer information of other carriers from iterative auctions. Furthermore, Máhr et al. (2010) apply an improvement step in which carriers directly share uncompleted parts of their routes.

Mes et al. (2013) consider a comparable mechanism, but order agents are allowed to delay commitments if the best bid is still higher than a reserve price. Furthermore, vehicle agents can reject already accepted orders, and consider the possible impact of accepting an order on future revenues. Hence, opportunity costs are included within their bids.

Figliozzi et al. (2004, 2005) also consider a marketplace without initial assignment of orders, but focus on competitive carrier behaviour. Carriers do not need to reveal their real costs, but may bid strategically, based on experience from earlier auctions or even based on all previous bids of all other carriers (Figliozzi et al., 2005).

2.4. Comparing information sharing policies

In all considered approaches, the necessary information is simply assumed to be shared by the carriers. Although this information is slightly different from approach to approach, there is hardly any assessment of the value of information sharing, as considered important by Gansterer and Hartl (2018b). To the best of our knowledge, an explicit assessment of the value of information sharing has only been done by Gansterer et al. (2020), Lyu et al. (2019), and Dai et al. (2014).

Gansterer et al. (2020) consider a combinatorial auction and show that it is beneficial for a carrier to share aggregate geographical information about its requests, since other carriers might consider to submit requests to the pool that are close to these requests. Furthermore, it is beneficial to share the marginal profit of individual requests within the pool with the auctioneer, since the auctioneer can generate more adequate bundles based on this information. The experiments, however, were performed on rather small static instances.

Table 2
Nomenclature for problem instances and solutions

Category	Notation	Description
General instance properties	C	Set of orders
	K	Set of vehicles
	τ	Time horizon
Properties of order $c \in C$	r_c	Release time
	p_c	Pickup location
	d_c	Delivery location
	q_c	Quantity
	g_c	Price
Properties of vehicle $k \in K$	r_k	Release time
	e_k	Earliest availability time
	l_k	Latest availability time
	α_k	Start location
	ω_k	End location
	c_k	Capacity
Sets of locations	P	All pickup locations
	D	All delivery locations
	A	All vehicle start locations
	Ω	All vehicle end locations
	I	Possible interim locations
Properties of locations $i, j \in P \cup D \cup A \cup \Omega \cup I$	e_i	Earliest service time
	l_i	Latest service time
	d_i	Service duration
	c_{ij}	Travel cost
	t_{ij}	Travel time
Properties of route plan for vehicle k at time t	m^{kt}	Number of stops
	ρ_i^{kt}	Location of i-th stop
	$a(\rho_i^{kt})$	Arrival time
	$d(\rho_i^{kt})$	Departure time
	$s(\rho_i^{kt})$	Service start time
	$l(\rho_i^{kt})$	Load after service
Properties of final solution R^τ	C_S	Set of served orders
	C_R	Set of rejected orders

Lyu et al. (2019) compare a setting in which carriers do not reveal any profit information to the auctioneer with a setting in which carriers share their marginal profit for both outsourced and demanded bundles of requests. In the first case, the winner determination problem maximizes the number of bundles being exchanged. In the second case, the total profit increase is maximized. Lyu et al. (2019) do not find a significant difference between the two scenarios, but they test on very small instances of 3 carriers with 5 requests each. In fact, both scenarios obtain the same results as a centralized planning approach in most cases.

The variable information sharing in the approach of Dai et al. (2014) is similar: within the conflict resolution step of the winner determination problem, the auctioneer favours carriers with the largest profits if profits are known. Otherwise, it selects carriers randomly. Dai et al. (2014) conclude that the scenario with profit information sharing yields better results than the scenario without this information. However, their instances are rather small (3 carriers, 5 or 8 requests each), and in more than half of the tests, both methods find the centrally computed optimum.

In this section, we have classified the collaborative vehicle routing literature based on information sharing properties. We identified full information sharing with a central coordinator (generally used for central optimization), partial information sharing with a central coordinator (generally used for centralized support), and distributed partial information sharing (generally used for distributed solution development). The literature review has shown that there is a lack of systematical assessments of types and levels of information sharing within a collaborative routing approach for complex, large-scale, dynamic problems. This article contributes to filling this gap by explicitly proposing different levels of location and cost information sharing, and analyzing their influence on solution quality.

3. Problem description

We consider a transportation platform to which thousands of vehicles and millions of requests may connect. Companies and individuals can submit their (autonomous) vehicles to the fleet at any moment, to make some money by contributing to the transportation service. Customers with specific requests can ask for transportation at any time, and are partly autonomous in selecting a service. Hence, we assume a platform that connects transport supply and demand in a complex, dynamic world. The efficiency of such a platform where all orders could be assigned to all carriers is, due to a larger solution space, expected to be higher than that of traditional collaborative approaches where carriers exchange only part of their initially contracted customers.

The problem builds on the well-known Pickup and Delivery Problem (PDP, see Parragh et al., 2008). In addition, we assume the following:

- The problem is dynamic. Orders can be added, changed, or canceled at any time. Furthermore, vehicles can join or leave the platform at any moment. The system has to deal with these dynamics in real time.
- The number of orders might be very large. Hence, central reoptimization after arrival of an order will not be possible.
- The problem is inherently decentralized. The availability time windows, route plans, and marginal costs of vehicles might not be known to the platform. Classical approaches for (dynamic) PDPs are hence not appropriate.

3.1. Notation and objectives

A problem instance (see Table 2) consists of a set of orders C, each representing a customer, and a set of vehicles K, each representing a different carrier. Each order $c \in C$ has a release time r_c (i.e., the time the order becomes known to the system), a pickup location p_c, a delivery location d_c, a load quantity q_c, and a price g_c that it pays for transportation. Each vehicle $k \in K$ has a release time r_k, a capacity c_k, an availability time window $[e_k, l_k]$, a start location α_k, and an end location ω_k. The release time specifies when the vehicle's availability becomes known to the system, whereas the time window $[e_k, l_k]$ specifies its actual availability. Hence, $r_k \leqslant e_k$.

Each location i has an earliest service start time e_i and a latest service start time l_i, that together determine the time window $[e_i, l_i]$ in which the service can start. Furthermore, d_i is the service duration, determining the time that is needed for the service at i. (For each vehicle $k \in K$, we set $e_{\alpha_k} = e_{\omega_k} = e_k$, $l_{\alpha_k} = l_{\omega_k} = l_k$, and $d_{\alpha_k} = d_{\omega_k} = 0$ to model that it only operates during its availability time window.) For each pair of locations (i, j), we denote the travel time and travel costs from i to j by t_{ij} and c_{ij}, respectively. All times are assumed to be before the time horizon τ.

A (temporary) solution at time t for a problem instance is defined as a set of routes $R^t = \{\langle\rho^{1t}\rangle, \cdots, \langle\rho^{|K|t}\rangle\}$, where each route (plan) $\langle\rho_i^{kt}\rangle_{i=1}^{m^{kt}}$ is a sequence of m^{kt} locations representing the (partially completed) path of vehicle k at time t. Due to the dynamic nature of the problem, the (partially completed) route plans can change during operational time. Only orders $c \in C$ for which p_c and d_c occur in the definitive route $\langle\rho^{k\tau}\rangle$ are served by vehicle k. Let C_S and C_R denote the sets of served and rejected orders in the final solution R^τ, i.e.:

$$C_S = \{c \in C \mid \exists\, k \in K \;\; \exists\, i \;\; \rho_i^{k\tau} = p_c\}; \text{ and} \tag{1}$$

$$C_R = C \setminus C_S. \tag{2}$$

The global goal is to obtain a final solution R^τ with several objectives:

- maximize the service level

$$SL = \frac{|C_S|}{|C|}, \tag{3}$$

  i.e., serve as many orders as possible;
- minimize the total travel costs

$$TC = \sum_{k \in K} \sum_{h=2}^{m^{k\tau}} c_{ij}, \;\; \text{where } i = \rho_{h-1}^{k\tau} \text{ and } j = \rho_h^{k\tau}; \text{ and} \tag{4}$$

- maximize the total profit

$$PR = \sum_{c \in C_S} g_c - TC - \gamma|C_R|, \tag{5}$$

  where γ is a parameter representing the fine per rejected order.

Since our goal is to get insight into the performance of the system under different information-sharing scenarios, we do not prioritize the objectives. The total profit, however, might be used as an aggregate objective, in which the relative importance of service level and travel costs can be controlled by the values of γ and c_{ij}.

Each stakeholder might have a different focus: customers might mainly be concerned about the service level, while a platform owner or authority might also be concerned about environmental aspects, as represented by travel costs. Individual carriers, however, might have different goals. Without an effective profit sharing mechanism, they likely want to maximize their own profit at the expense of the global system. Define the equivalents of the global objectives on vehicle level as follows:

$$TC^k = \sum_{h=2}^{m^{k\tau}} c_{ij}, \;\; \text{where } i = \rho_{h-1}^{k\tau} \text{ and } j = \rho_h^{k\tau}; \text{ and} \tag{6}$$

$$PR^k = \sum_{c \in C_S^k} g_c - TC^k, \;\; \text{where } C_S^k = \{c \in C \mid \exists\, i \;\; \rho_i^{k\tau} = p_c\}. \tag{7}$$

Carriers might completely focus on their local profit by maximizing PR^k, or might be concerned with the global system, for example, if all carriers are charged equally for the term $-\gamma|C_R|$ in the global profit function. In even more cooperative situations, carriers do not want to optimize their local profit, but focus completely on the global profit PR, and rely on a fair profit allocation system. We elaborate on the different possibilities in Sections 4 and 5.

3.2. Constraints and dynamics

In our dynamic setting, we assume that vehicles might change direction at any time t (except during service time, i.e., loading or unloading needs to be fully completed). Hence, we define the DPDP on the Euclidean plane, with travel times and travel costs proportional to Euclidean distances between two locations. By doing so, the location of a vehicle can be determined at every time t if we keep track of departure and arrival times, and travel times and travel costs are defined for each pair of locations.

Let $P = \bigcup_{c \in C}\{p_c\}$ and $D = \bigcup_{c \in C}\{d_c\}$ be the sets of all pickup and all delivery locations, respectively. Also, let $A = \bigcup_{k \in K}\{\alpha_k\}$ and $\Omega = \bigcup_{k \in K}\{\omega_k\}$ be the sets of all start and end positions of the vehicles. Furthermore, by I we denote a set of additional interim locations at which the vehicles might change direction. We require P, D, A, Ω, and I to be disjoint sets, but they can contain duplicates of the same physical locations. For $i \in I$, we always have $d_i = 0$.

A route (plan) $\langle \rho_i^{kt} \rangle_{i=1}^{m^{kt}}$ for a vehicle $k \in K$ at time $t \geqslant r_k$ consists of at least two locations ($m^{kt} \geqslant 2$), is associated with four functions a: $\bigcup_{i=2}^{m^{kt}}\{\rho_i^{kt}\} \to \mathbb{R}$, d: $\bigcup_{i=1}^{m^{kt}-1}\{\rho_i^{kt}\} \to \mathbb{R}$ (representing arrival and departure times, respectively), s: $\bigcup_{i=1}^{m^{kt}}\{\rho_i^{kt}\} \to \mathbb{R}$ (representing service start time), and l: $\bigcup_{i=1}^{m^{kt}}\{\rho_i^{kt}\} \to \mathbb{R}$ (representing vehicle load after service), and has the following constraints:

- the vehicle starts at its start location and stops at its end location, i.e.:

$$\rho_1^{kt} = \alpha_k \text{ and } \rho_{m^{kt}}^{kt} = \omega_k; \tag{8}$$

- all time constraints are respected, i.e.:

$$e_{\rho_i^{kt}} \leqslant s(\rho_i^{kt}) \leqslant l_{\rho_i^{kt}} \ \forall\, i \in \{1, \cdots, m^{kt}\}; \tag{9}$$

$$a(\rho_i^{kt}) \leqslant s(\rho_i^{kt}) \ \forall\, i \in \{2, \cdots, m^{kt}\}; \tag{10}$$

$$s(\rho_i^{kt}) + d_{\rho_i^{kt}} \leqslant d(\rho_i^{kt}) \ \forall\, i \in \{1, \cdots, m^{kt} - 1\}; \tag{11}$$

$$d(\rho_{i-1}^{kt}) + t_{\rho_{i-1}^{kt}, \rho_i^{kt}} = a(\rho_i^{kt}) \ \forall\, i \in \{2, \cdots, m^{kt}\}; \tag{12}$$

- all capacity constraints are respected, i.e.:

$$l(\rho_1^{kt}) = l(\rho_{m^{kt}}^{kt}) = 0; \tag{13}$$

$$\rho_i^{kt} = p_c \Rightarrow l(\rho_i^{kt}) = l(\rho_{i-1}^{kt}) + q_c \ \forall\, i \in \{2, \cdots, m^{kt} - 1\} \ \ \forall\, c \in C; \tag{14}$$

$$\rho_i^{kt} = d_c \Rightarrow l(\rho_i^{kt}) = l(\rho_{i-1}^{kt}) - q_c \ \forall\, i \in \{2, \cdots, m^{kt} - 1\} \ \ \forall\, c \in C; \tag{15}$$

$$l(\rho_i^{kt}) \leqslant c_k \ \forall\, i \in \{2, \cdots, m^{kt} - 1\}; \tag{16}$$

- pickups and deliveries are paired and respect precedence constraints, i.e.:

$$\rho_i^{kt} = p_c \Rightarrow \exists\, j > i \ \rho_j^{kt} = d_c \ \forall\, i \in \{2, \cdots, m^{kt} - 1\} \ \ \forall\, c \in C; \tag{17}$$

$$\rho_i^{kt} = d_c \Rightarrow \exists\, j < i \ \rho_j^{kt} = p_c \ \forall\, i \in \{2, \cdots, m^{kt} - 1\} \ \ \forall\, c \in C; \text{ and} \tag{18}$$

- start and stop locations occur only at the beginning and end of a route, i.e.:

$$\rho_i^{kt} \in P \cup D \cup I \ \forall\, i \in \{2, \cdots, m^{kt} - 1\}. \tag{19}$$

For each vehicle $k \in K$ and each time $t < r_k$, we have an empty route $\langle \rho^{kt} \rangle = \langle \rangle$.

A (temporary) solution $R^t = \{\langle \rho^{1t} \rangle, \cdots, \langle \rho^{|K|t} \rangle\}$ for a problem instance has the following restrictions:

- the pickup and delivery corresponding to each order $c \in C$ occur within the route of at most one vehicle, i.e.:

$$\rho_i^{kt} = p_c \Rightarrow (\rho_j^{\bar{k}t} \neq p_c \vee (\bar{k} = k \wedge i = j)) \;\forall\, c \in C \;\; \forall\, k, \bar{k} \in K \;\; \forall\, i \in \{2, \cdots, m^{kt} - 1\} \;\; \forall\, j \in \{2, \cdots, m^{\bar{k}t} - 1\}; \text{ and} \tag{20}$$

- no elements unknown at time t appear, i.e.:

$$(\rho_i^{kt} = p_c \vee \rho_i^{kt} = d_c) \Rightarrow r_c \leqslant t \;\forall\, c \in C \;\; \forall\, k \in K \;\; \forall\, i \in \{2, \cdots, m^{kt} - 1\}. \tag{21}$$

Due to the dynamic nature of the problem, the solution needs to be iteratively updated during run time. For each vehicle $k \in K$, an initial route $\langle \rho^{kr_k} \rangle = \langle \alpha_k, \omega_k \rangle$ is given at time r_k. Given a vehicle k and two time points u and t (with $r_k \leqslant u < t$), the route $\langle \rho^{ku} \rangle$ might be changed into a route $\langle \rho^{kt} \rangle$ only if the part until time t remains unaffected. Formally, $\rho_1^{kt} = \rho_1^{ku}$, and we have the following requirements for each ρ_i^{ku} $(i \neq 1)$ within $\langle \rho^{ku} \rangle$:

- if the service at ρ_i^{ku} has already started at time t (i.e., $s(\rho_i^{ku}) \leqslant t$), then the location is fixed and arrival and service time cannot be changed anymore, i.e.:

$$\rho_i^{kt} = \rho_i^{ku},\; a(\rho_i^{kt}) = a(\rho_i^{ku}), \;\text{ and } s(\rho_i^{kt}) = s(\rho_i^{ku}); \tag{22}$$

  the departure time, however, may be changed according to $d(\rho_i^{kt}) \geqslant t$ if $d(\rho_i^{ku}) > t$;
- if the vehicle has arrived at ρ_i^{ku} at time t but has not yet started the service (i.e., $a(\rho_i^{ku}) \leqslant t < s(\rho_i^{ku})$), there are two options:
  - the vehicle can still process the planned pickup or delivery, i.e.:

$$\rho_i^{kt} = \rho_i^{ku}; \text{ or} \tag{23}$$

  - the vehicle may withdraw from service and leave the location by changing it into an interim location, i.e.:

$$\rho_i^{kt} = \hat{\rho} \;\text{ for } \hat{\rho} \in I \text{ a duplicate of } \rho_i^{ku} \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = t; \tag{24}$$

  in either case, $a(\rho_i^{kt}) = a(\rho_i^{ku})$, but the service start time and departure time may be changed according to $s(\rho_i^{kt}) \geqslant t$ and $d(\rho_i^{kt}) \geqslant t$;
- if the vehicle is driving towards ρ_i^{ku} at time t (i.e., $d(\rho_{i-1}^{ku}) \leqslant t < a(\rho_i^{ku})$), there are again two options:
  - the vehicle may stick to the original plan, i.e.,

$$\rho_i^{kt} = \rho_i^{ku}; \tag{25}$$

  in this case, the arrival time stays the same, but service and depart times may be changed, if possible, i.e.:

$$a(\rho_i^{kt}) = a(\rho_i^{ku}), \;\; s(\rho_i^{kt}) \geqslant a(\rho_i^{ku}), \;\text{ and } d(\rho_i^{kt}) \geqslant s(\rho_i^{ku}); \text{ or} \tag{26}$$

  - the vehicle may change direction immediately at time t, i.e.:

$$\rho_i^{kt} = \hat{\rho} \;\text{ for } \hat{\rho} \in I \text{ the actual location of the vehicle at time } t \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = t, \tag{27}$$

  and arrival, service, and departure times are set accordingly:

$$a(\rho_i^{kt}) = s(\rho_i^{kt}) = e_{\hat{\rho}}, \;\text{ and } d(\rho_i^{kt}) \geqslant e_{\hat{\rho}}; \tag{28}$$

- if none of the tree above conditions hold, there are no restrictions for ρ_i^{kt}, i.e., ρ_i^{ku} may be changed or removed from the route (subject to the definitions of a route and a solution, of course).

3.3. Example

In Fig. 4, we show how a solution for a dynamic PDP could develop over time. Initially, two vehicles and two orders are known to the system. Although vehicle 1 is only available from $t = 4$ on, it can already plan its route at $t = 0$. Since the capacity of vehicle 1 is not sufficient for order 4, the best plan is to assign order 4 to vehicle 2 and order 5 to vehicle 1. Vehicle 2 immediately starts driving towards p_4. At $t = 4$, vehicle 1 becomes available and starts driving towards p_5.

Then, at $t = 7$, vehicle 3 connects to the platform and announces its availability from $t = 8.5$ on. The total driven distance can be decreased if vehicle 3 serves order 4 and vehicle 2 serves order 5, since vehicle 2 is already quite close to p_5 at $t = 7$. Hence, the plan is updated at $t = 7$ as follows: vehicle 1 stops driving towards p_5 and stays idle, vehicle 2 starts driving towards p_5, and vehicle 3 plans to

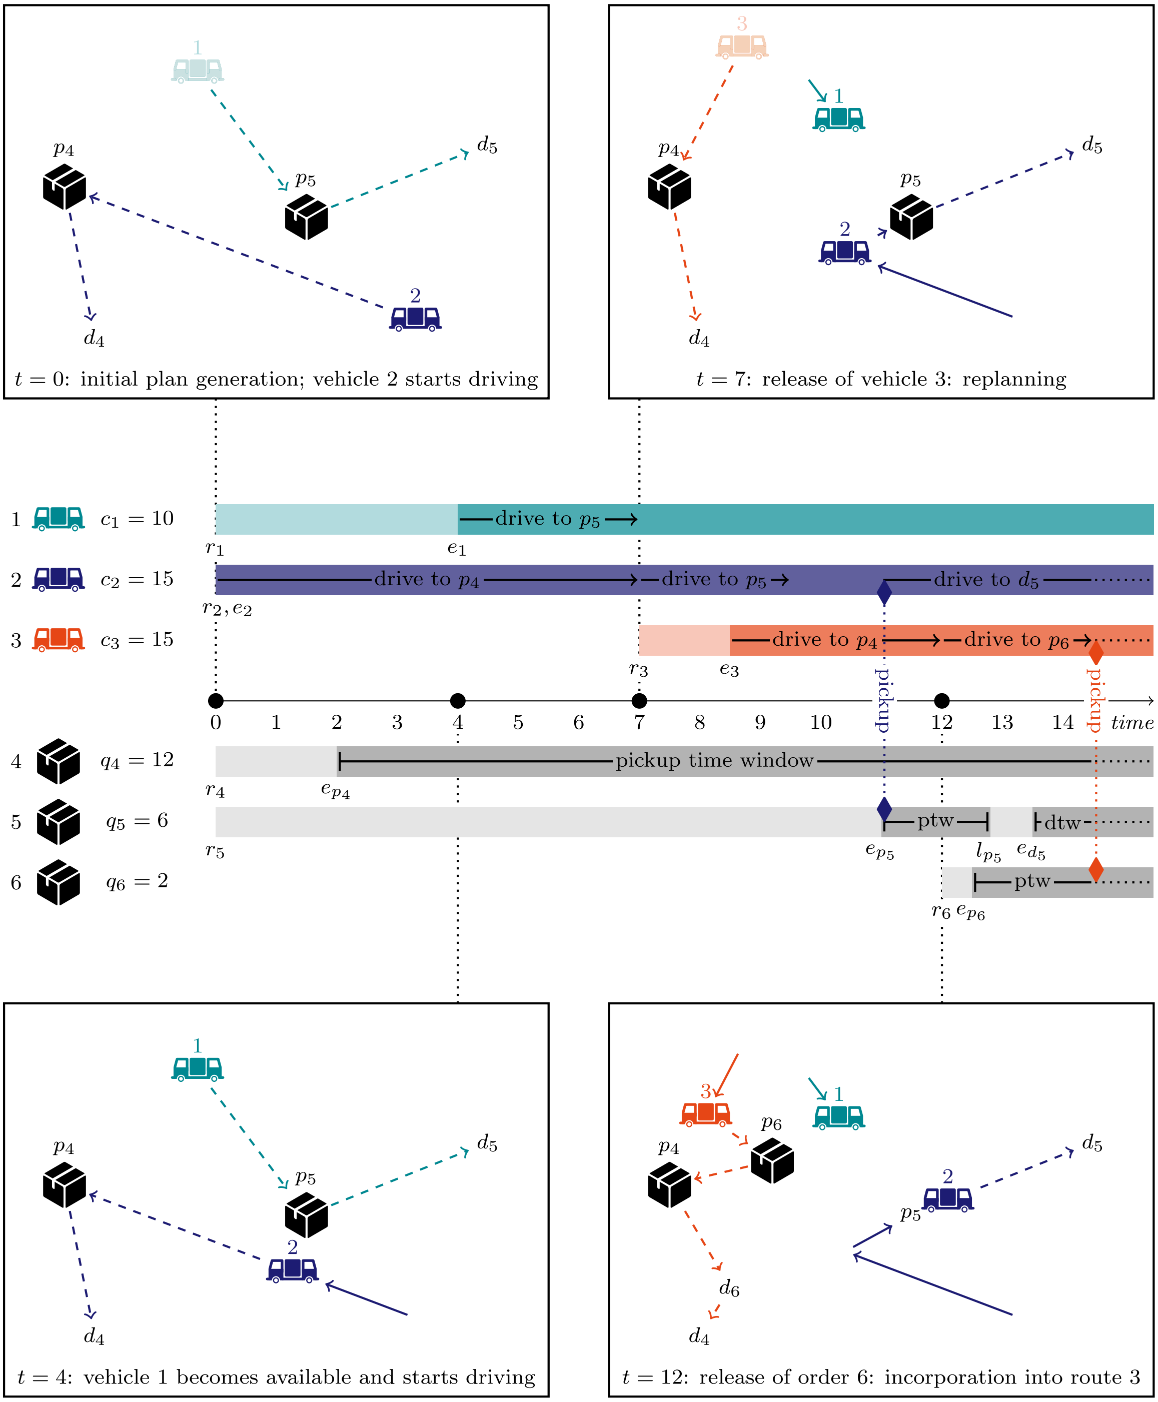

Fig. 4. First part of the development of a solution for a DPDP instance with 3 orders and 3 vehicles. The time line shows the vehicle and order properties, as well as the vehicle actions. The four panels show the positions, planned routes (dashed lines), and completed routes (solid lines) of the vehicles at $t = 0$, $t = 4$, $t = 7$, and $t = 12$.

go to p_4 immediately when it becomes available at $t = 8.5$. At $t = 9.5$, vehicle 2 reaches p_5, but cannot yet start the pickup, since the pickup time window of order 5 has not yet opened. Hence, vehicle 2 waits until $t = 11$, performs the pickup, and next drives towards d_5.

At $t = 12$, a new order is submitted to the platform. Although its pickup location is close to vehicle 1, a better option is to assign it to vehicle 3 since vehicle 3 has enough capacity for carrying order 4 and 6 together and d_6 is close to d_4. Hence, at $t = 12$, vehicle 3 includes order 6 into its route, and changes direction towards p_6.

Note that the final solution is not optimal: with hindsight it would have been better if vehicle 2 had stayed idle and vehicle 1 had served order 5.

4. Information sharing

If carriers are concerned with their privacy, they might not want to share (part of) their properties with the platform. Therefore, we propose to construct vehicle routes on the local level: transportation contracts are made through local interactions between carriers and customers (see Fig. 3). To determine the influence of different levels of information sharing in such an approach, we define several sharing policies for carriers. We distinguish between information about locations and route plans of vehicles (Section 4.1), and information about travel costs of the vehicles (Section 4.2). Although the information on locations and travel costs is not necessary for the mechanism to work, we expect that more information will improve solution quality.

4.1. Vehicle route information

Information about the current and future locations of vehicles might be of importance for customers, to be able to select the most appropriate vehicles. We consider the following three (incrementing) levels of location information that carriers might share with the platform (see Fig. 3), and discuss the potential for customers when this information is known.

- **No position sharing (NPS)** If carriers do not share any information about their vehicle positions, customers cannot infer which vehicles will be most appropriate.
- **Current position sharing (CPS)** If carriers share their actual vehicle positions, customers can select the vehicles that are closest to their pickup locations for interaction. Although there is no guarantee that the closest vehicle will transport the order in an optimal solution, nearer vehicles are in general more likely to have lower detour costs for the order.
- **Full plan sharing (FPS)** If carriers share their complete route plans, i.e., give access to $\langle \rho^{kt} \rangle$, customers might construct even better estimates of the appropriateness of a vehicle. The spatiotemporal distance of the vehicle route to either the pickup or the delivery location can be used as a heuristic: if the vehicle has already planned to be in the neighborhood of the pickup or delivery location of an order, just at a time that is within or close to the pickup or delivery time window, the order might be included into the route at relatively low extra costs. Formally, the spatiotemporal distance of a pickup or delivery location i to a vehicle route $\langle \rho^{kt} \rangle$ is given by $\min_{u \in [t, l_k], v \in [e_i, l_i]} \theta |u - v| + c_{i \hat{\rho}^{ku}}$, where $\hat{\rho}^{ku}$ is the location of vehicle k at time u, and θ is a parameter controlling the relative importance of time.

4.2. Marginal cost information

Information about the marginal travel costs of a vehicle for an individual order is relevant for the global solution: if an order can be served by another vehicle with lower marginal travel costs, this will improve *TC*. Although there are not always direct advantages of sharing marginal travel costs for a carrier, system-wide costs may be reduced on the longer term. If gains are redistributed, this can be beneficial for individuals as well (see Section 3.1). In other scenarios, however, carriers might want to keep their cost information confidential, or want to provide their marginal costs only if the order will contribute to an increase of the local vehicle profit PR^k. We consider the following three (decreasing) levels of cost sharing.

- **Full cost sharing (FCS)** A carrier is always willing to share cost information, even if this does not improve the local vehicle profit PR^k directly.
- **Partial cost sharing (PCS)** A carrier is only willing to share cost information if this can contribute to an increase of the local vehicle profit PR^k.
- **No cost sharing (NCS)** A carrier is never willing to share cost information.

Note that cost information is not shared with the platform, but only with the individual customers (see Fig. 3), within an auction system that will be described in the next section.

5. Multi-agent approach

To investigate the value of different types and amounts of (locally) shared information, we introduce a Multi-Agent System for solving the DPDP in a distributed manner, in accordance with the assumption that vehicles and orders can independently attach to the platform without a central authority that is aware of all information. For each order, an order agent is introduced that needs to make a contract for transportation. Each vehicle is represented by a vehicle agent, that has the goal of making some profit by efficiently serving orders. Order agents and vehicle agents try to make contracts by interacting in multi-agent auctions (Wooldridge, 2009): order agents act as auctioneers by offering a transportation task; vehicle agents respond by proposing a bid (see Fig. 5)[3]. This approach allows fast responses within a dynamic supply and demand market, in which we can vary the information that is being exchanged. Although there is no guarantee on optimal solutions, our MAS is competitive with standard centralized heuristics for the DPDP: we found improvements of up to 4% on the benchmark instances of Mitrović-Minić et al. (2004) (see Appendix A).

[3] In line with other literature on MASs for logistics (e.g., Van Lon and Holvoet, 2017), we use the auction terminology, although we do not focus on a real competitive setting where we aim for incentive compatibility.

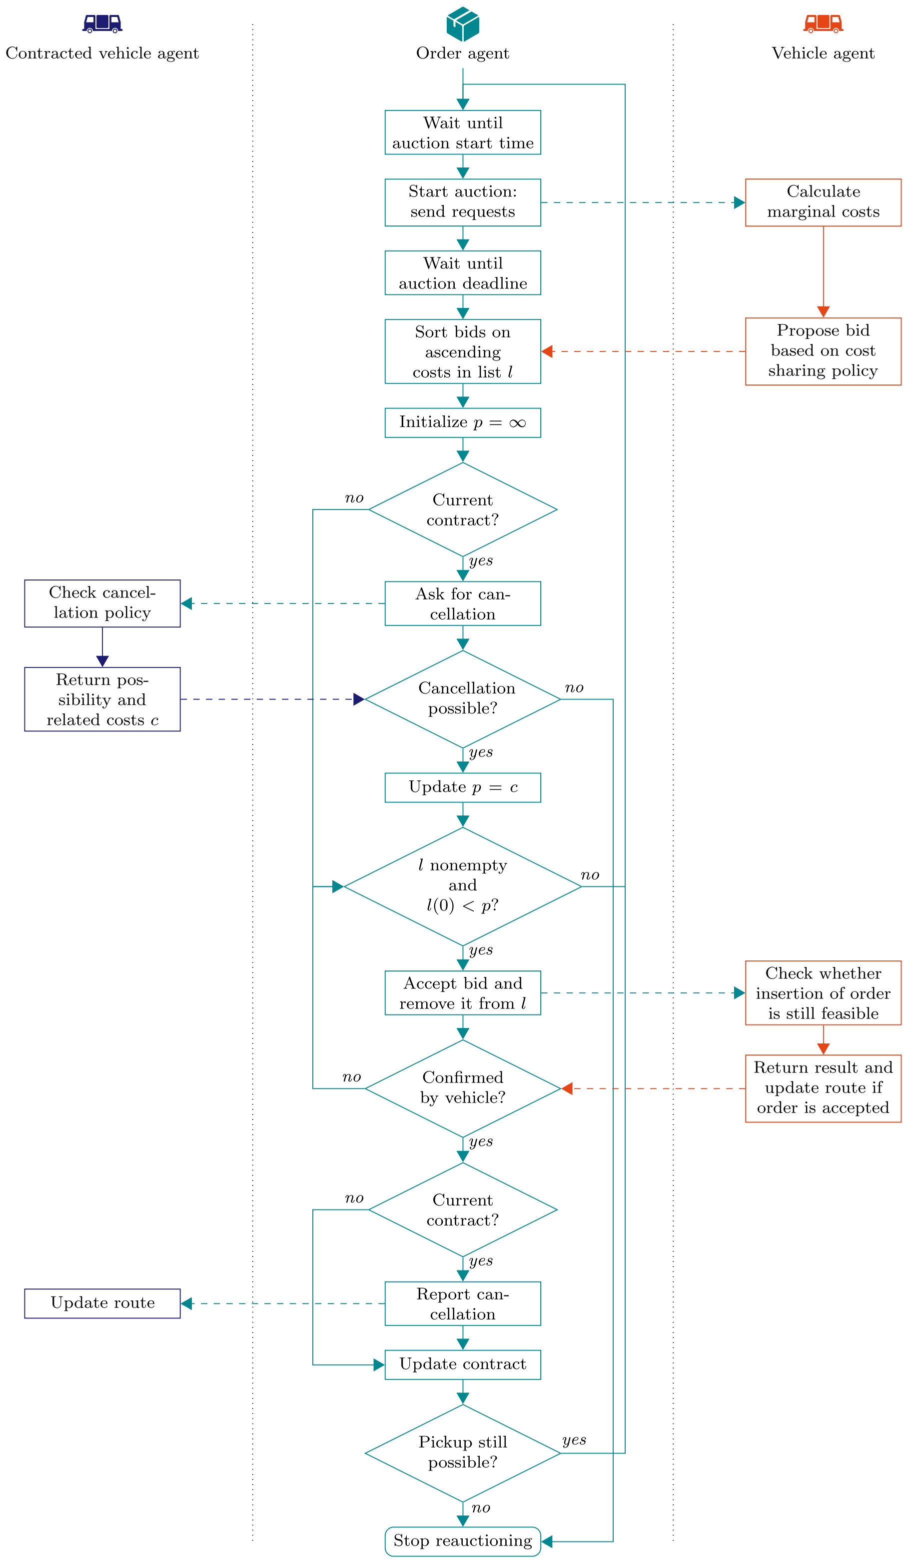

Fig. 5. Flowchart of the iterative auction procedure within the Multi-Agent System. Dashed arrows represent exchange of information between different actors.

5.1. Order agent

An order agent for an order $c \in C$ is responsible for making and, if possible, improving a transportation contract with a vehicle agent. At time r_c, it sends a request, containing p_c and d_c (with the corresponding time windows and service durations), the quantity q_c, the price g_c, and an auction end time, to a well-selected group of vehicle agents. While in current MAS approaches a request is sent to all vehicle agents (Máhr et al., 2010), a lot of computational effort might be saved if the order agent interacts only with a restricted group of most promising vehicles, especially if the number of vehicles is large. Therefore, the platform selects a certain percentage of the currently available vehicles for the order agent, based on the available vehicle location information (see Section 4.1):

- In case of NPS, the platform cannot do better than selecting vehicles randomly (see Fig. 6a).
- With CPS, the platform selects the vehicles that are closest to p_c at the time of the request (see Fig. 6b).
- For FPS, the platform selects the vehicles with lowest spatiotemporal distance of the vehicle route to either p_c or d_c (see Fig. 6c).

After communicating the request to the selected vehicles, an order agent waits for bids from the vehicles. At the auction end time, the received bid with lowest value is chosen (this decision might be random if the bids do not contain values; see Section 5.2). The order agent asks the vehicle agent that submitted this bid to incorporate the task into its route. Since the route plan of the vehicle has possibly been changed between proposing the bid and receiving the acceptance message, the vehicle agent needs to check if the bid has not become outdated. If insertion is still feasible, they establish a contract. Otherwise, the order agent chooses the vehicle agent of the second-best bid, and so on, until a contract is made or no feasible bid remains.

After each auction, the order agent schedules to start a new auction after some time, since a more suitable vehicle might appear in the dynamic world. If an order agent already has a contract, it asks the contracted vehicle agent at the end of the auction if cancellation of the contract is still possible, and what the actual costs are. Only if cancellation is still possible and the costs of the best new bid are lower than these actual costs, the contract will be replaced. If contract cancellation is no longer possible (since the pickup has already taken place or it is too expensive for the vehicle agent) or the order agent did not succeed in making a contract at time l_{p_c}, the order agent stops scheduling new auctions.

5.2. Vehicle agent

A vehicle agent for a vehicle $k \in K$ is responsible for constructing a feasible route, while maximizing profit by accepting or rejecting orders. As discussed in Section 3.1, vehicle agents can merely focus on their own profit, or be more cooperative and focus on the global profit. This difference is expressed in the bidding strategy of the vehicle agent.

When the vehicle agent of vehicle k receives a request with auction end time t from the order agent for order c, it computes the marginal cost $MC^{kt}(c)$ for this order with respect to its current route plan at time t, i.e., the additional travel costs for including c into its route after time t. A formal description is given in Appendix B.[4] Furthermore, the marginal profit $MP^{kt}(c) = g_c - MC^{kt}(c)$ is computed. If including the order is feasible, the vehicle agent returns a bid to the order agent, based on the specific cost sharing policy (see Section 4.2):

- With FCS, the vehicle agent always submits a bid with $MC^{kt}(c)$, even when the marginal profit $MP^{kt}(c)$ for including c is negative (see Fig. 7a). The vehicle agent is fully cooperative in the sense that it allows the order agent to make a contract with the vehicle agent with lowest marginal cost.
- With PCS, the vehicle agent only submits a bid with $MC^{kt}(c)$ if $MP^{kt}(c) > 0$ (see Fig. 7b). If the marginal profit is negative, no response is given. Hence, the vehicle agent focuses more on PR^k than with FCS, but is still cooperative in the sense that it allows vehicle agents with lower marginal cost to make a contract with the order agent.
- For NCS, the vehicle agent submits a bid if $MP^{kt}(c) > 0$, but does not include the marginal cost within the bid (see Fig. 7c). This means that the order agent cannot compare the different bids and needs to select one at random.

When the vehicle agent of vehicle k has a contract with the order agent of order c, and receives a cancellation request of this order agent for time t, it checks whether the order has not yet been picked up at time t, and if its cancellation policy allows cancellation at time t. Only in this case, the contract may be terminated. Cancellation policies depend on the cost sharing policy as follows:

- With FCS, the vehicle agent returns the actual marginal cost $MC^{kt}(c)$ and allows cancellation if the marginal cost of another vehicle for the order is lower.
- With PCS, the vehicle agent computes $MC^{kt}(c) - \phi \cdot g_c$, for $\phi \in [0, 1]$. If this value is negative, the vehicle agent will not allow cancellation. Otherwise, this value is returned as actual cost. For low values of ϕ, the vehicle agent is rather cooperative, but probably shares too much information. For high values of ϕ, the vehicle agent is more reticent regarding cancellation and shares in general less information. Only if the new vehicle has a significantly lower marginal cost, a cancellation is allowed. If $\phi = 1$, cancellation is never possible, since the vehicle agent would not have submitted a bid if the marginal cost is higher than g_c.

[4] Although an exact or metaheuristic approach for computing the marginal cost might be used in real-world applications (where each vehicle agent has its own computational resources), we approximate it by an elementary insertion heuristic in our simulations: the current route of the vehicle is kept, while p_c and d_c are inserted at the feasible positions that minimize additional travel costs.

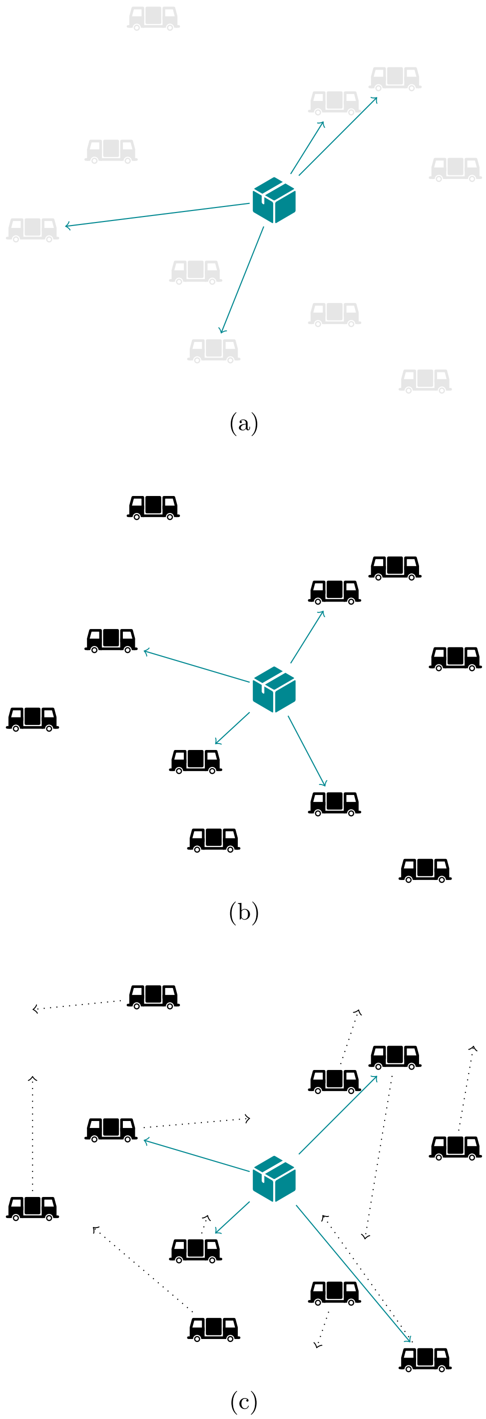

Fig. 6. Selection of 40% of the available vehicles under different location information sharing policies: NPS (a), CPS (b), and FPS (c). A solid arrow represents selection of a vehicle; a dotted arrow represents the route plan of a vehicle.

- With NCS, cancellation is only allowed if $MC^{kt}(c) - \psi g_c > 0$, for $\psi \in [0, 1]$. The vehicle agent does not share any value. Since a new vehicle is selected randomly, the value of ψ needs to be not too low: this would result in multiple contract cancellations, although removing an order is likely to decrease the quality of a route. On the other hand, if $\psi = 1$, cancellations are never allowed, while some successfull reauctions might improve the global solution quality.

If a vehicle agent receives an acceptance message from an order agent, it checks whether the order still can be placed into the route. If so, a contract is made, and the vehicle agent updates the route plan for the vehicle accordingly.

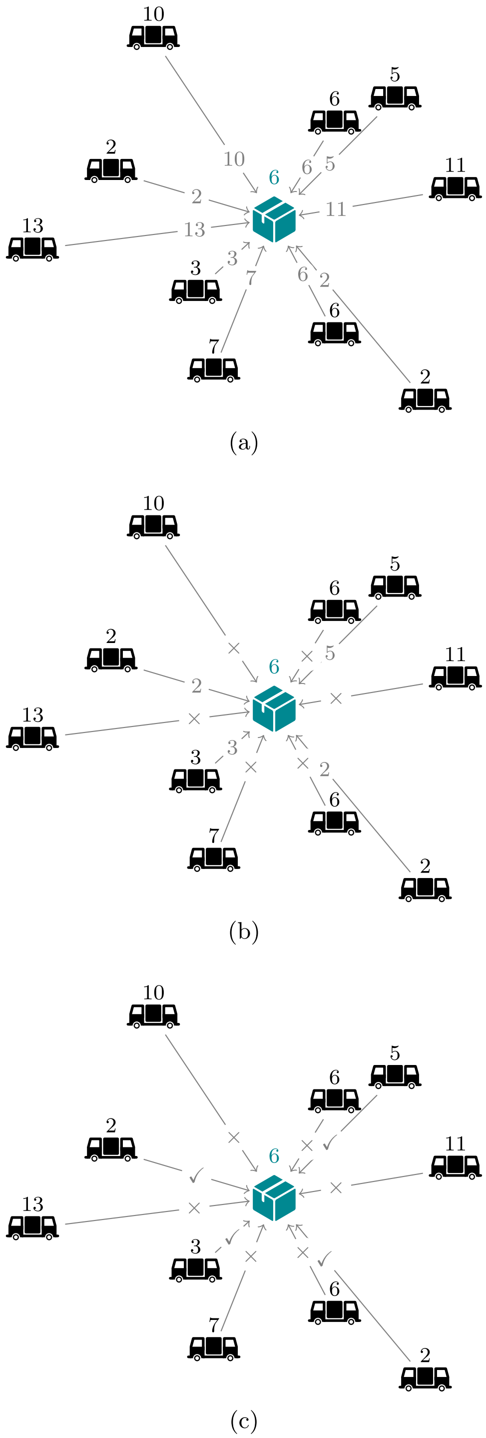

Fig. 7. Vehicle bids given their marginal cost (above the vehicles) and the order price (above the order) under different cost sharing policies: FCS (a), PCS (b), and NCS (c). An arrow with a number, checkmark or cross represents a bid with marginal cost, a bid without any value, and no bid at all, respectively.

6. Computational study

6.1. Problem instances

To compare the different information sharing policies, we generated a set of 10 DPDP instances with time windows, quantities, and order prices.[5] The instances were defined on an approximately 1000×1000 area with a 10h time horizon, and contain each 1000 orders and 150 vehicles. Recall that each vehicle represents a different carrier. Half of the vehicles is available during the complete time horizon; the other half has randomly selected r_k, e_k, and l_k. All vehicle capacities were set to 100.

Order quantities were sampled from $\mathcal{N}(20, 5^2)$, service durations (in seconds) for pickups and deliveries from $\mathcal{N}(120, 30^2)$, and time window lengths (in seconds) for pickups and deliveries from $\mathcal{N}(9000, 1800^2)$. The x and y coordinates of all pickup and delivery locations and vehicle start locations were sampled from $\mathcal{N}(500, 75^2)$ with probability 0.5, from $\mathcal{N}(200, 75^2)$ with probability 0.25, or from $\mathcal{N}(800, 75^2)$ with probability 0.25, resulting in 9 clusters of locations. All instances represent an open DPDP, i.e., end locations for vehicles are not prescribed.

Travel times t_{ij} equal the Euclidean distance between locations. Furthermore, $c_{ij} = 0.011 \cdot t_{ij}$ for all locations i and j, and $g_c = 0.014 \cdot t_{p_c d_c}$ for all orders c.

We refer to the instance set described above as the *base set*. In Section 6.6, we check whether results are generalizable to other scenarios. For this, we created three related problem instance sets. First, we created two sets with lower transport capacity by reducing the number of vehicles. These *low capacity set* and *medium capacity set* are copies of the base set in which only 75 or 100 of the vehicles per instance are kept, respectively. The base set acts as a high capacity set where all orders could easily be served.

Second, we created a set in which the requests are more urgent. This *urgent set* is a copy of the base set in which the latest service times for pickups and deliveries have been changed according to a sampling from $\mathcal{N}(900, 300^2)$ instead of $\mathcal{N}(9000, 1800^2)$ for the time window lengths. Earliest service times are kept the same. The base set acts as corresponding non-urgent set with more time between release times and deadlines of orders.

6.2. Experimental settings

The MAS was implemented in Go and all experiments were performed single-threaded on a 64-bit machine running Linux with a 3.30 GHz Intel i5-4590 CPU and 8 GB of RAM. Unless otherwise specified, we used the following default settings.

The time between the start of an auction and the auction deadline is 10s; vehicle agents compute results for 1s after the actual auction deadline to make sure that the bid is still feasible after auction processing and acceptance communication. If an order agent tries to make a contract with a vehicle agent but the bid is no longer feasible, the new auction deadline is set to 1ms after the current auction deadline. The time between an auction deadline and the start time of a reauction for an order c is $(l_{p_c} - r_c)/m$, where m denotes a maximum number of auctions per order, which was set to 10.

The standard number of vehicles that an order agent interacts with is 10% of the currently known vehicles. For computational purposes, we aggregate over regions with $t = 1800$s, $x = 50$, and $y = 50$ instead of computing the exact spatiotemporal distances when selecting vehicles.

If a vehicle agent determines a new route plan and has some temporal flexibility within its schedule, it waits for 20% of the available time at its current position. By this, orders that appear in the neighbourhood of the vehicle soon after plan determination might still be included into the route, but most of the available flexibility can be used to change plans when the vehicle is already driving.

The values of ϕ and ψ were set to 0.2 and 0.6, respectively.

Instead of processing the auctions in real time, we keep track of the times and give all processes enough time to fully complete. In general, the total time needed for the experiments with this approach is much less than the real total time (one run with FCS and the settings described above takes about 25s, one run with interaction with 100% of the vehicles takes about 180s, and computation times for PCS and NCS are even lower). Furthermore, when running in real time, vehicle agents do not always have enough time to place a bid, since they all need to compute it at the same moment but share a single processor. In real-world applications, this will not be a problem, since computations can be distributed over multiple processors.

6.3. Comparing different information sharing scenarios

To compare the influence of the different information sharing policies, we combine each of the three vehicle position sharing policies (NPS, CPS, and FPS) with each of the three cost sharing policies (FCS, PCS, and NCS), resulting in 9 scenarios. Furthermore, we want to compare the influence of the number of vehicles that each order interacts with, and hence run each scenario with different vehicle interaction percentages of the total known fleet.

To be able to compute average results over the different instances, we normalize the travel costs and profit as follows. For each instance, we compute 10 solutions by running the MAS with parameter settings that likely give a high quality solution, i.e., we use FCS, 100% vehicle interaction, and a fast reauctioning time of 60s between the deadline of an auction and the start time of a reauction for the order. Per instance, the solution out of 10 runs with highest service level (and, in case of a tie, lowest travel costs) is used as best known solution (BKS). Let R be the solution for a run of the MAS on a problem instance, then normalized travel costs and

[5] The data set is available at http://doi.org/10.4121/uuid:a363bf41-ab09-4201-81de-4da49bae281d.

normalized profit for R are defined by $TC(R)/TC(\text{BKS})$ and $PR(R)/PR(\text{BKS})$, respectively.

The results in terms of service level, normalized travel costs, and normalized profit for the different vehicle location and travel cost information sharing scenarios are shown in Fig. 8, where we vary the percentage of (currently known) vehicles that each order interacts with (VIP). All results are averages over 10 instances and 10 runs per instance. The response rate in Fig. 8d is the average number of times that a vehicle agent sends a positive response (i.e., a bid) upon a single request of an order agent.

Fig. 8a shows that the service level is almost always 1 if costs are always shared (FCS). For a VIP of 5% or 10%, the service level is lower if positions are shared (CPS or FPS) than if positions are not shared (NPS). This could be caused by a myopic vehicle selection heuristic: for CPS and FPS, the order agent selects only a limited group of closest vehicles within each reauction, while this group might not differ that much from auction to auction. With NPS, however, the group of contacted vehicles could be much larger since it is randomly selected each auction.

The policies in which costs are not fully shared, PCS and NCS, result in significantly lower service levels than FCS, especially for the lower VIPs, as expected: the bid submission rate is much lower than for FCS (see Fig. 8d), and it is likely that some orders do not receive any feasible bid. Interestingly, the service level is higher for NCS than for PCS. An explanation can be that some advantageous vehicles remain unused and still are available for subsequent auction rounds with NCS (since order agents choose vehicles randomly), while order agents immediately select the best vehicles with PCS. This greediness can cause a local shortage of feasible vehicles on the long term. This explanation corresponds with the higher number of bids per request for NCS than for PCS (see Fig. 8d).

Fig. 8b and Fig. 9 show decreasing travel costs with increasing VIP for the policies where costs are shared (FCS and PCS). Since the service level stays almost equal for FCS, we can conclude that a higher VIP results in better solutions. The communicational and computational load, however, is also higher. Without cost sharing (NCS), the travel costs increase with VIP, but the service level also does. Only for NCS with NPS, we can observe that a higher VIP is helpful, since the service level increases while the travel costs are constant.

The same figures show a clear advantage of route information sharing: the travel costs are in general lower for FPS than for CPS, and lower for CPS than for NPS (given the same cost sharing policy). Fig. 8c shows the same pattern for the profit: FPS gives higher profit than CPS, and CPS results in higher profit than NPS, indicating that position sharing is profitable. This difference is the largest for lower VIPs. (Indeed, with a VIP of 100, position sharing is not relevant.) However, there is one exception: for FCS, NPS performs slightly better than CPS at a VIP of 40% and 70%. Again, this might be due to the larger range of vehicles that NPS selects within the subsequent reauctions. Certainly, the received number of bids per request is a bit lower for CPS than for NPS with FCS (see Fig. 8d), while this is not the case with PCS and NCS.

In general, the policy in which vehicle agents always share their costs (FCS) yields the best results. Since the profit is highest for FCS with the higher VIP values, vehicle agents have an incentive to use this cost sharing policy. Furthermore, sharing complete route plans (FPS) is beneficial: for a VIP of 40% with FPS, service level and profit are about the same as for a VIP of 100%, while the communicational and computational load is about 40% of it. Only with lower VIPs, using PCS is beneficial for vehicle agents compared to FCS.

6.4. Influence of number of reauctions

A larger number of reauctions per order can increase the solution quality: due to the dynamic nature of the problem, combinatorial advantages might arise after some time. In Fig. 10, we show the results of varying the maximum number of auctions per order (MNA) for different cost sharing methods and different fines for rejected orders (γ). We used a vehicle interaction percentage of 10% and no position sharing.

Omitting reauctions seriously deteriorates the service level if costs are not always shared: Fig. 10a shows a steep increase in service level for PCS and NCS when MNA increases from 1 to 5. After 10 auctions, this increase is much smaller. When costs are always shared (FCS), however, the service level is always almost 1, even when reauctions are omitted.

For the travel costs, a similar pattern is visible (see Fig. 10b): when MNA increases from 1 to 5, the travel costs increase a lot for PCS and NCS (due to the need of including an extra 20% of the orders). From an MNA of 5 on, the travel costs decrease a little for PCS and NCS, while the service level still increases. Hence, the higher number of auctions results in slightly better solutions. For FCS, the costs decrease when increasing MNA, while the service level is almost the same. Hence, solutions are better with higher MNA, but the effect is largest for low MNA values.

The relevance of a high number of auctions is also reflected in the profit (see Fig. 10c): for all cost sharing policies, a higher MNA yields higher profits. While PCS always outperforms NCS in terms of profit, it depends on the fine per rejected order (the value of γ) and MNA whether FCS outperforms PCS. In general, with a higher MNA, a lower value of γ is needed to let FCS be superior. For an MNA of 5, PCS performs better than FCS, even for $\gamma = 2$. For an MNA of 10, however, PCS performs better than FCS for $\gamma \leq 1$, but worse for $\gamma \geqslant 1.5$. For an MNA of 15, PCS is only better than FCS for $\gamma = 0$, and for an MNA of 20, FCS outperforms PCS irrespective of the value of γ.

Hence, by enabling a larger number of auctions or setting a high fine for rejected orders, individual vehicles might prefer the FCS policy over the PCS policy, resulting in a higher service level. A drawback, however, is that FCS in combination with a large number of auctions comes with the highest communicational load (see Fig. 10d).

Furthermore, an increase in allowed auctions might result in an increase of sequential route changes for vehicles. This generally is no problem with autonomous vehicles but might be uncomfortable for drivers. In Fig. 11, we show the average number of route plan changes per vehicle under different circumstances. From Fig. 11a, we observe that vehicles on average have more route updates with FCS than with NCS, and even less with PCS. In Fig. 11b, the number of route updates that require an immediate change of direction is shown. This number is very low compared to the total number of route updates. Interestingly, it is generally highest for CPS, followed by NPS and FPS, and decreases with increasing VIP. Fig. 11c shows that the number of route updates does increase with MNA, but not linearly. In Fig. 11d, we observe that the number of turnoffs is very low when only one auction is allowed (MNA = 1), and slightly higher (but not strictly increasing) for larger values of MNA.

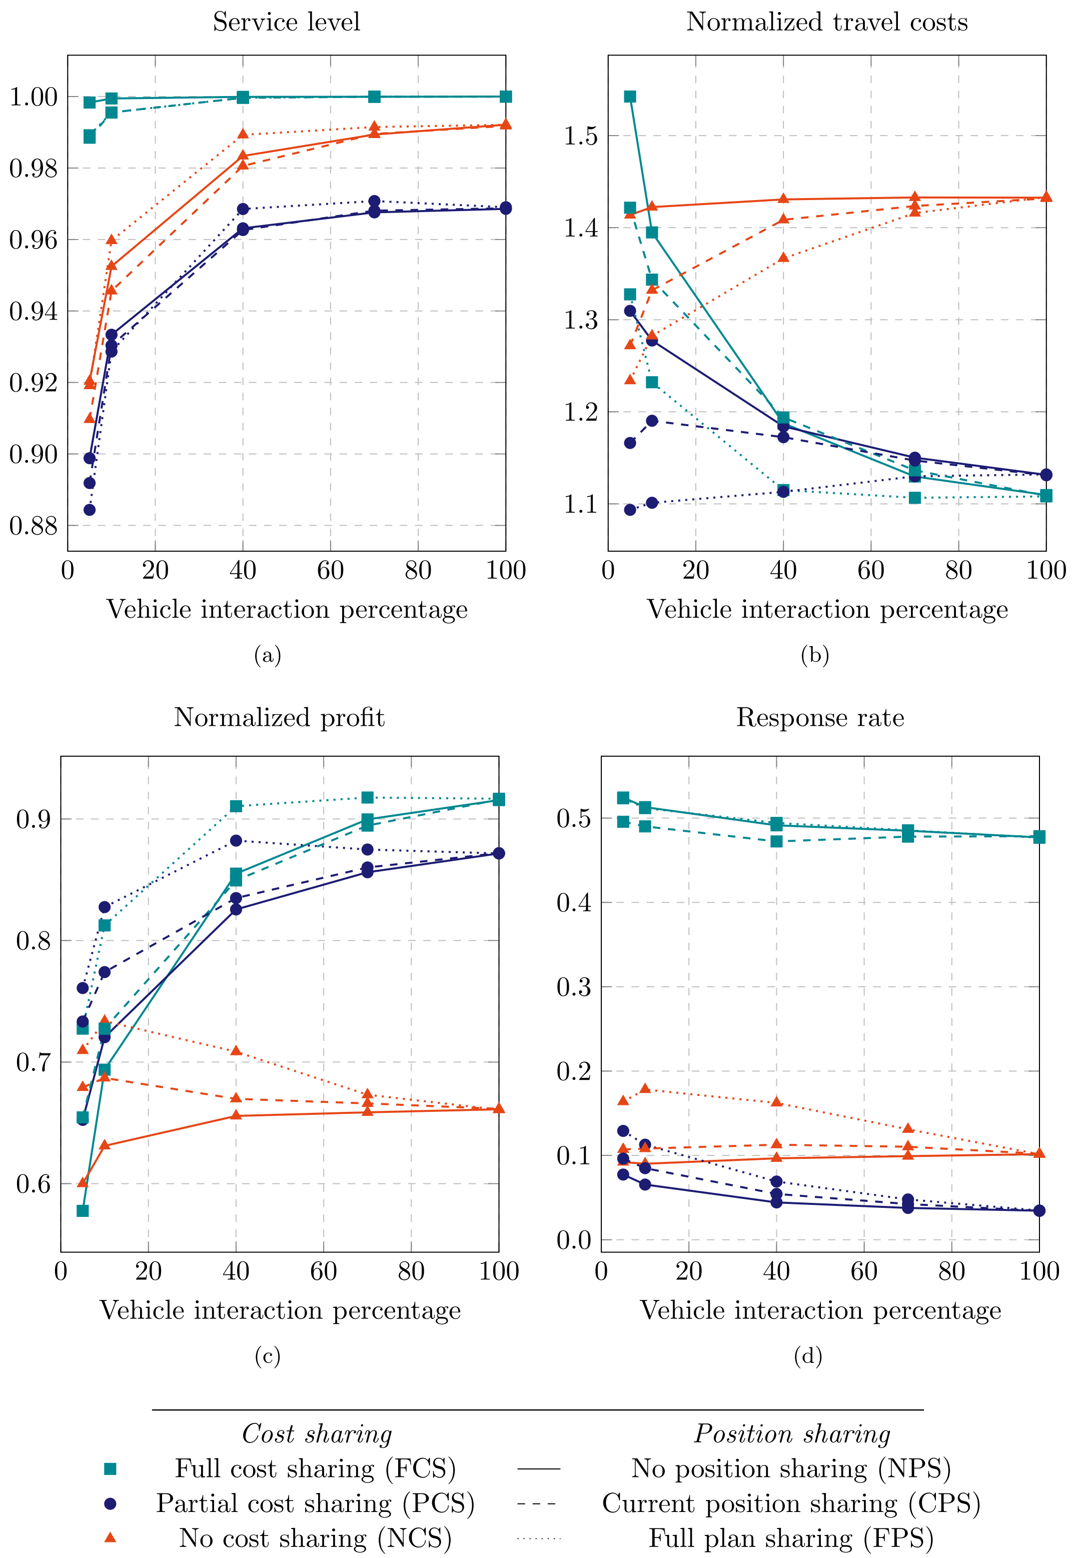

Fig. 8. Mean results on varying VIP for the three cost sharing and the three position sharing methods. In (c), there is no fine for rejected orders ($\gamma = 0$).

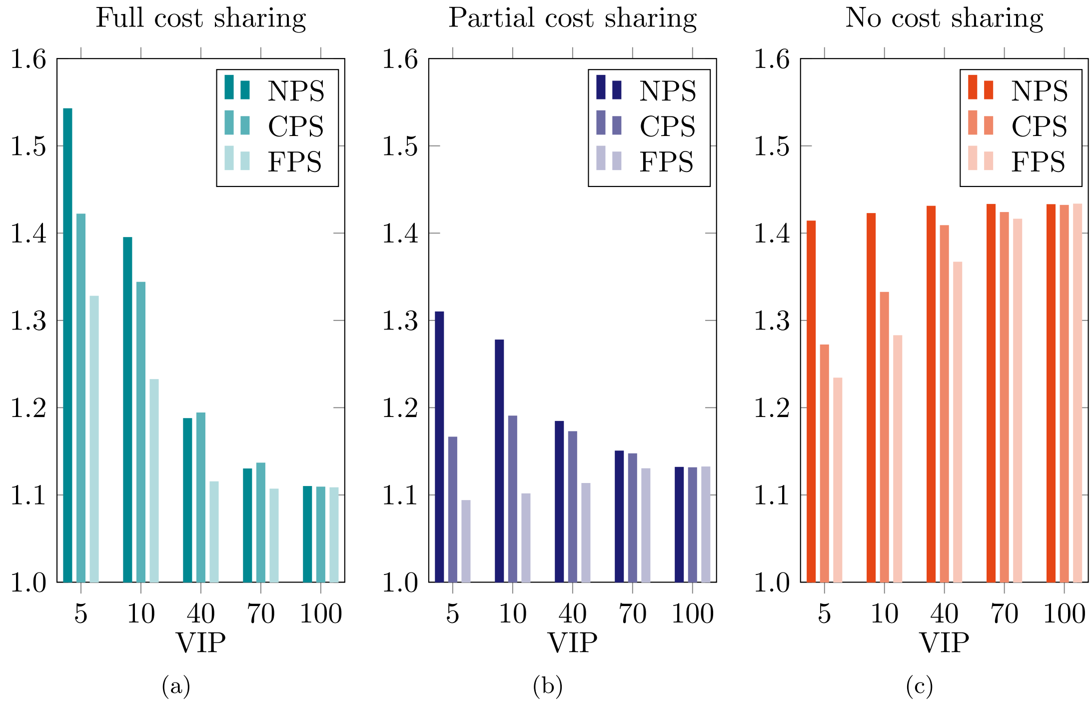

Fig. 9. Mean normalized travel costs on varying vehicle interaction percentage (VIP) for the three cost sharing and the three position sharing methods.

6.5. Influence of emission or congestion penalties

Since the cost information sharing policy (FCS) provides higher service levels than the policies where costs are not (always) shared (PCS and NCS), a platform owner or government might want to stimulate vehicle agents to share their costs. An extra charge on driven distance (next to the regular travel costs, e.g., by applying taxes or specific emission or congestion penalties) can be expected to be helpful, since the bidding behaviour with FCS is not affected by travel costs, while the bidding behaviour with PCS and NCS is highly affected by travel costs.

In Fig. 12, we show the results for varying values of the travel costs with penalties per unit (TCPU) for different cost sharing methods and different fines for rejected orders (γ). Again, we used a vehicle interaction percentage of 10% and no position sharing. Fig. 12a shows that the service level decreases for PCS and NCS if the TCPU increases, as expected: the marginal profit for an order decreases if the marginal costs increase, and hence, offering a bid is attractive for less vehicles. Fig. 12d confirms this observation: the number of bids that an order receives upon a request is lower for higher TCPU values. The service level with FCS, on the other hand, is not dependent on the TCPU, as expected. Also, the travel costs for FCS increase linearly with the TCPU (see Fig. 12b), and the profit decreases accordingly (see Fig. 12c).

The profits for PCS and NCS decrease especially when TCPU increases from 0.012 to 0.013 (see Fig. 12c). This is as expected, because TCPU is close to 0.014, the price per unit that orders pay for transportation.

An interesting trade-off between FCS and PCS is mainly visible for these higher TCPU values. While PCS results in higher profit than FCS if TCPU $\leqslant 0.012$ and $\gamma \leqslant 0.5$ or even $\gamma \leqslant 1.0$, FCS results in highest profits when TCPU $= 0.013$, irrespective of the value of γ.

Setting the TCPU and γ wisely may result in an increased incentive for vehicle agents to fully share their costs, and hence improve the service level, without putting an extra communicational load upon the system (as it was the case with increasing the MNA in Section 6.4).

6.6. Influence of fleet capacity and order urgency

To analyze whether the results hold in other scenarios, we repeat the experiments on the low capacity set, the medium capacity set, and the urgent set (see Section 6.1). Here, we describe the main findings; detailed results are given in Appendices C–E.

In Fig. 13, we show the results on the high, medium, and low capacity set for different cost sharing policies. The position sharing policy is fixed as NPS, but the results for other position sharing policies are comparable.

The service level descreases if the available number of vehicles decreases, as expected. In all instance sets, FCS results in a higher service level than PCS, and PCS performs better than NCS.

With respect to the travel costs, FCS performs worse if the fleet capacity descreases: where FCS obtains lowest travel costs for a VIP of 70% or 100% in the high capacity set, it ends up between PCS and NCS for the medium capacity set, and even results in highest travel costs with the low capacity set. Note that the relative travel costs of PCS and NCS are not dependent on the ratio of supply and demand: PCS always results in lower travel costs than NCS.

A similar pattern is visible for the profit: with the high capacity set, FCS results in highest profits, whereas PCS outperforms FCS with the medium capacity set. Even NCS results in higher profits than FCS in most cases with the low capacity set.

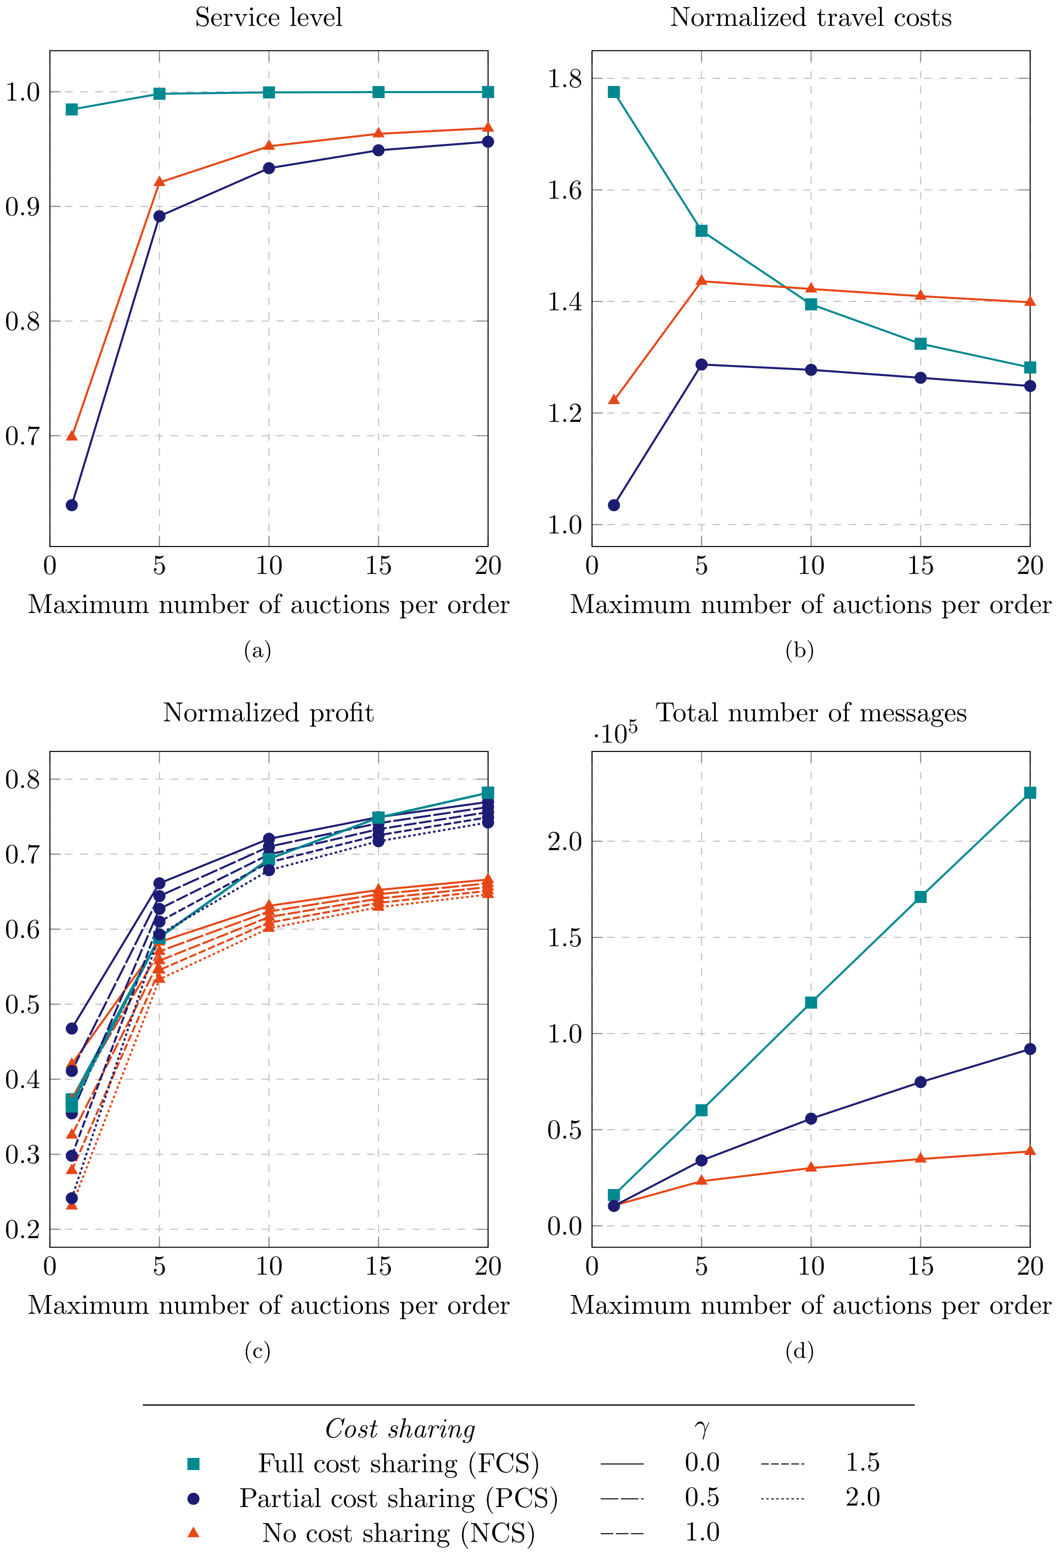

Fig. 10. Mean results on varying maximum number of auctions per order for the three cost sharing methods and different fines per rejected order (γ).

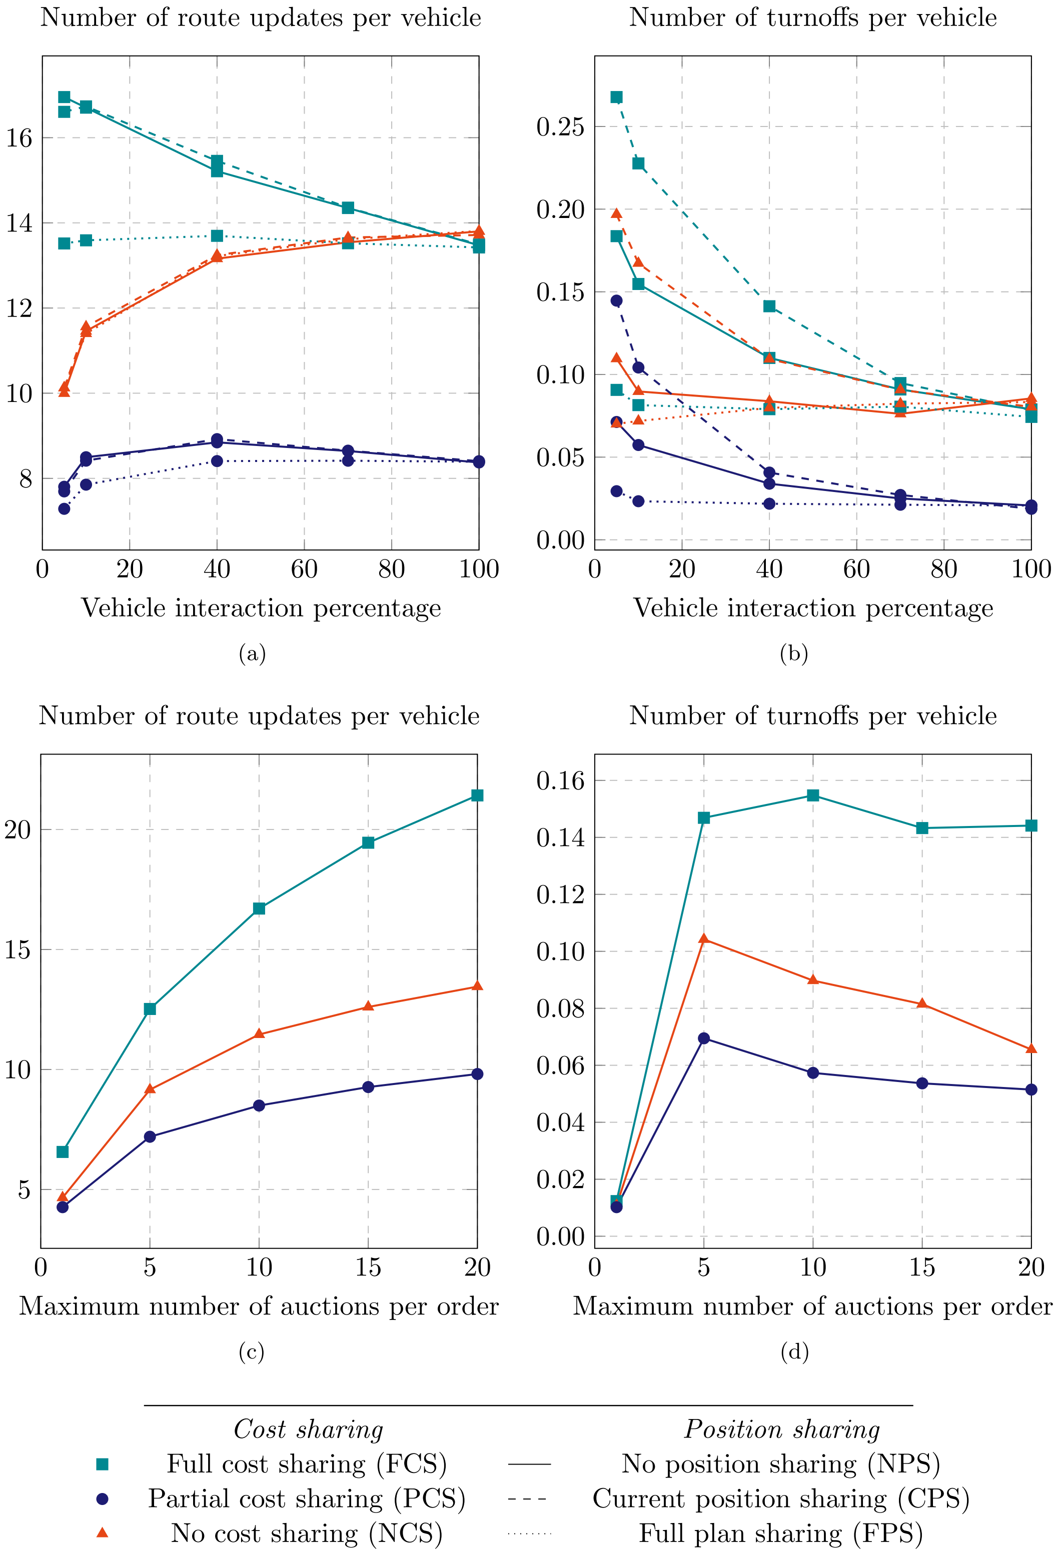

Fig. 11. Mean numbers of plan updates per vehicle. (a) and (c) show the total average number of plan updates (i.e., the number of times a vehicle agent inserts or removes a task), whereas (b) and (d) show the average number of plan updates resulting in an immediate change in driving direction.

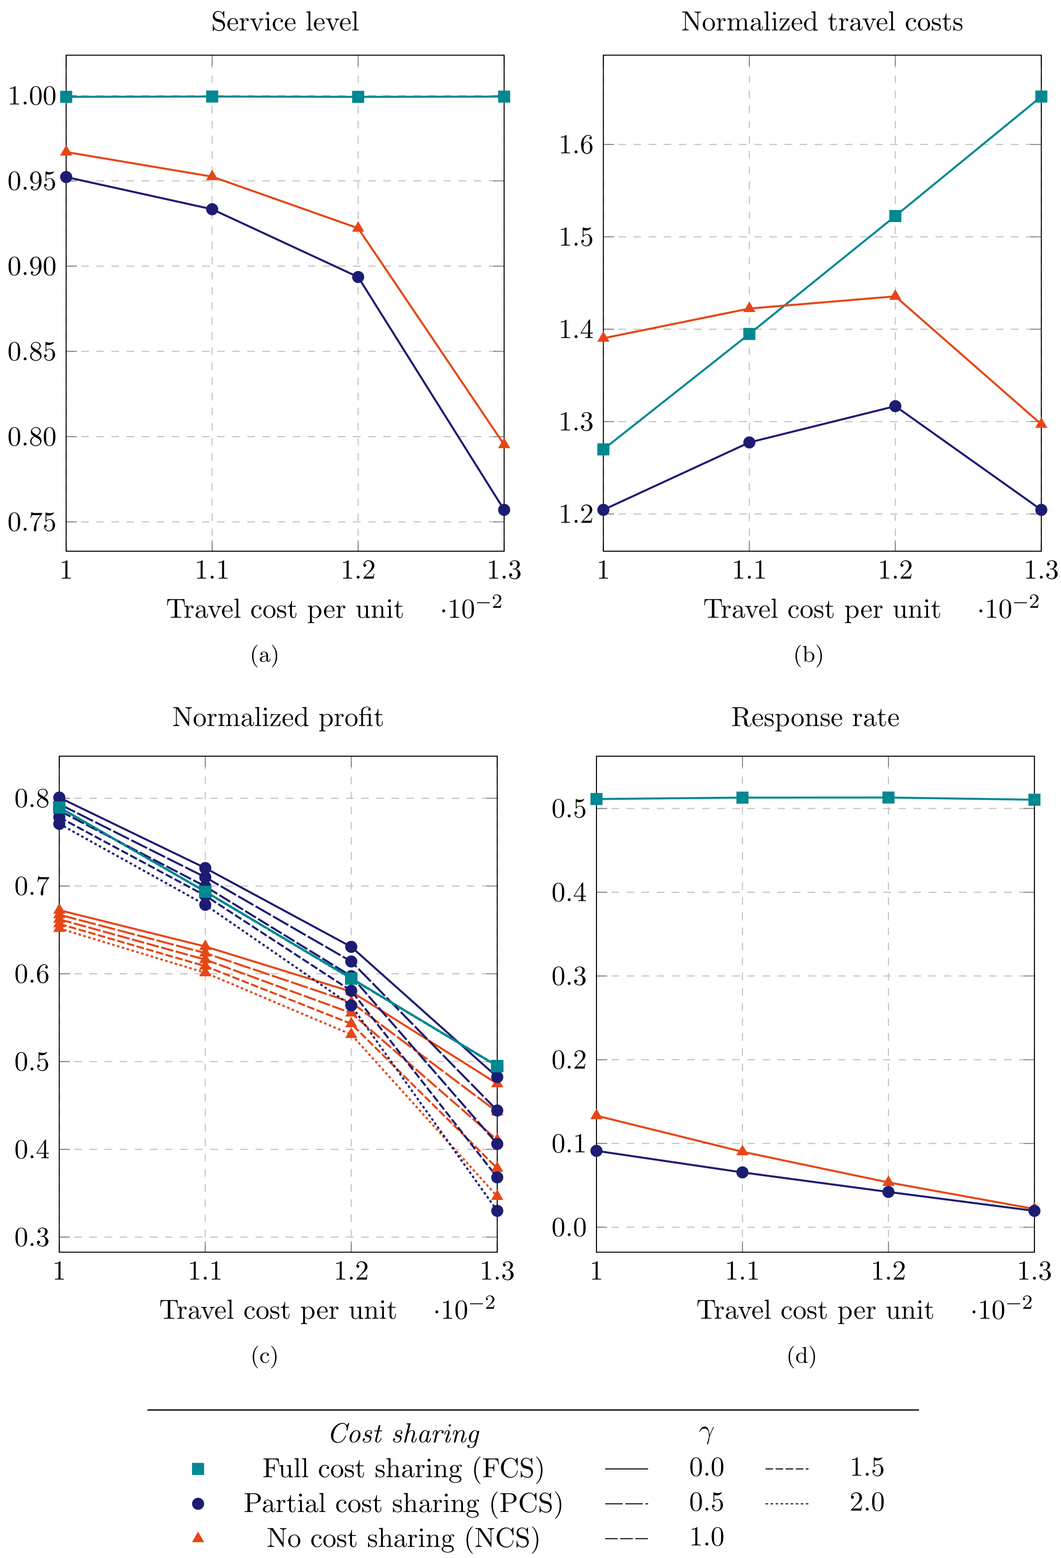

Fig. 12. Mean results on varying travel costs for the three cost sharing methods and different fines per rejected order (γ).

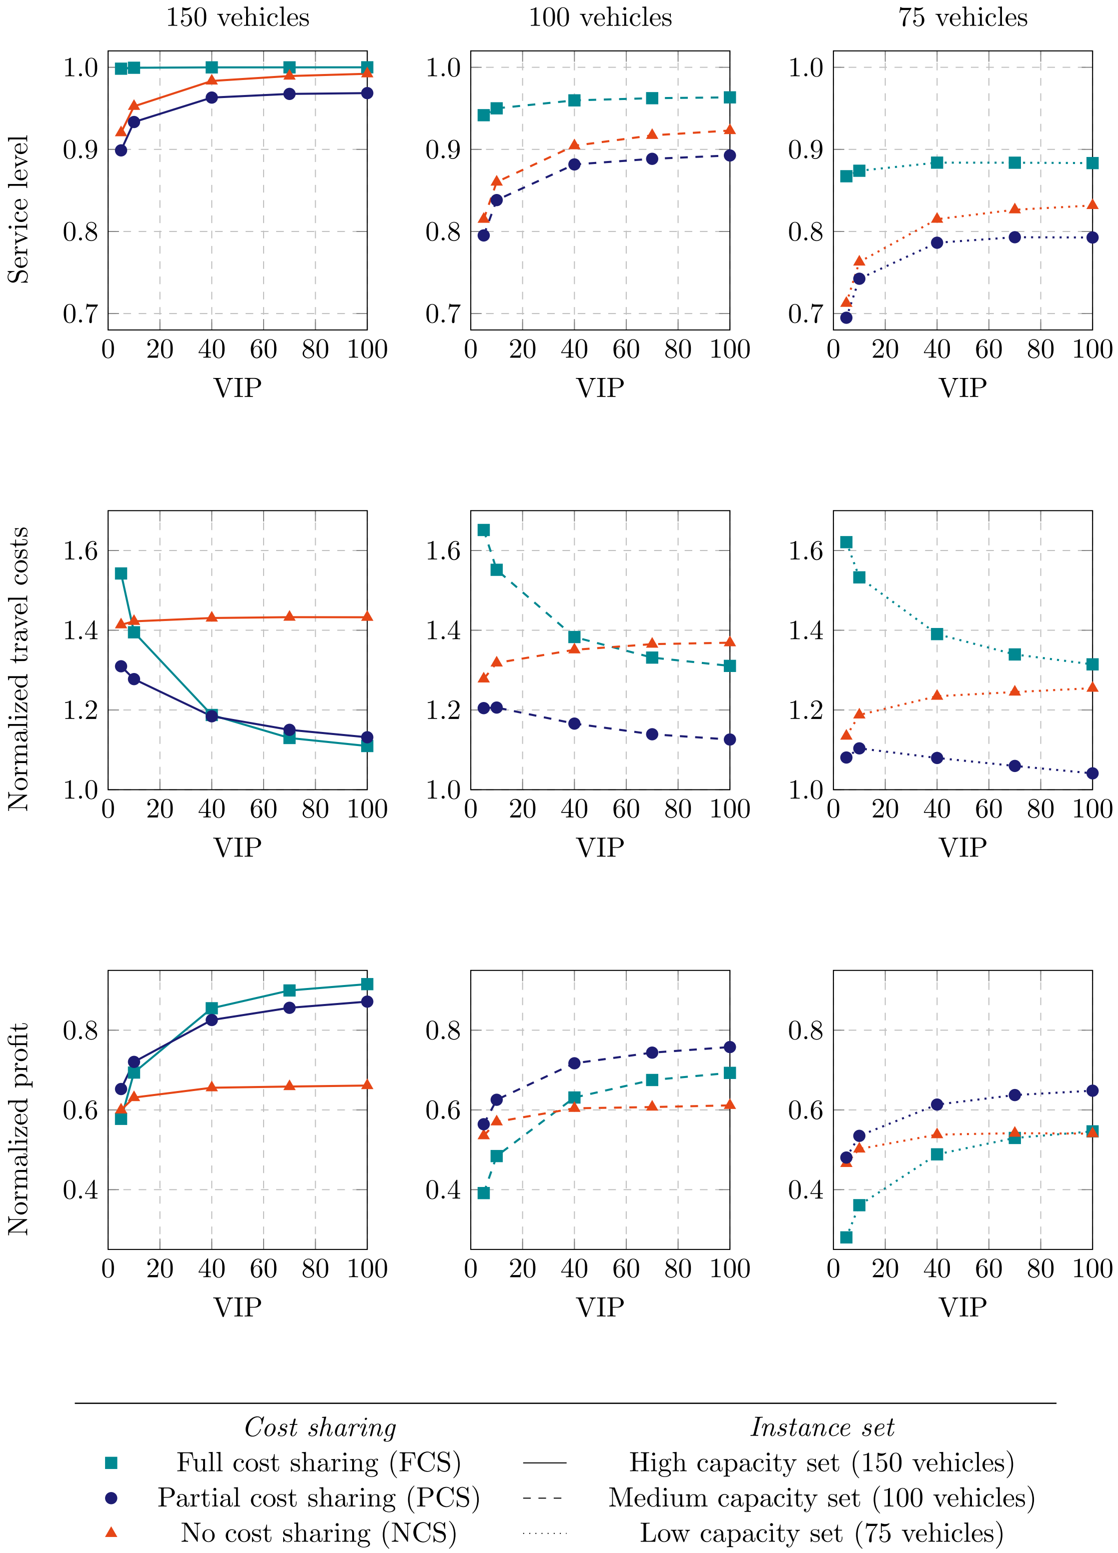

Fig. 13. Mean results on varying VIP for the three cost sharing methods on the different instance sets. All results are normalized with respect to the BKSs of the high capacity set. There is no fine for rejected orders ($\gamma = 0$) and NPS is used everywhere.

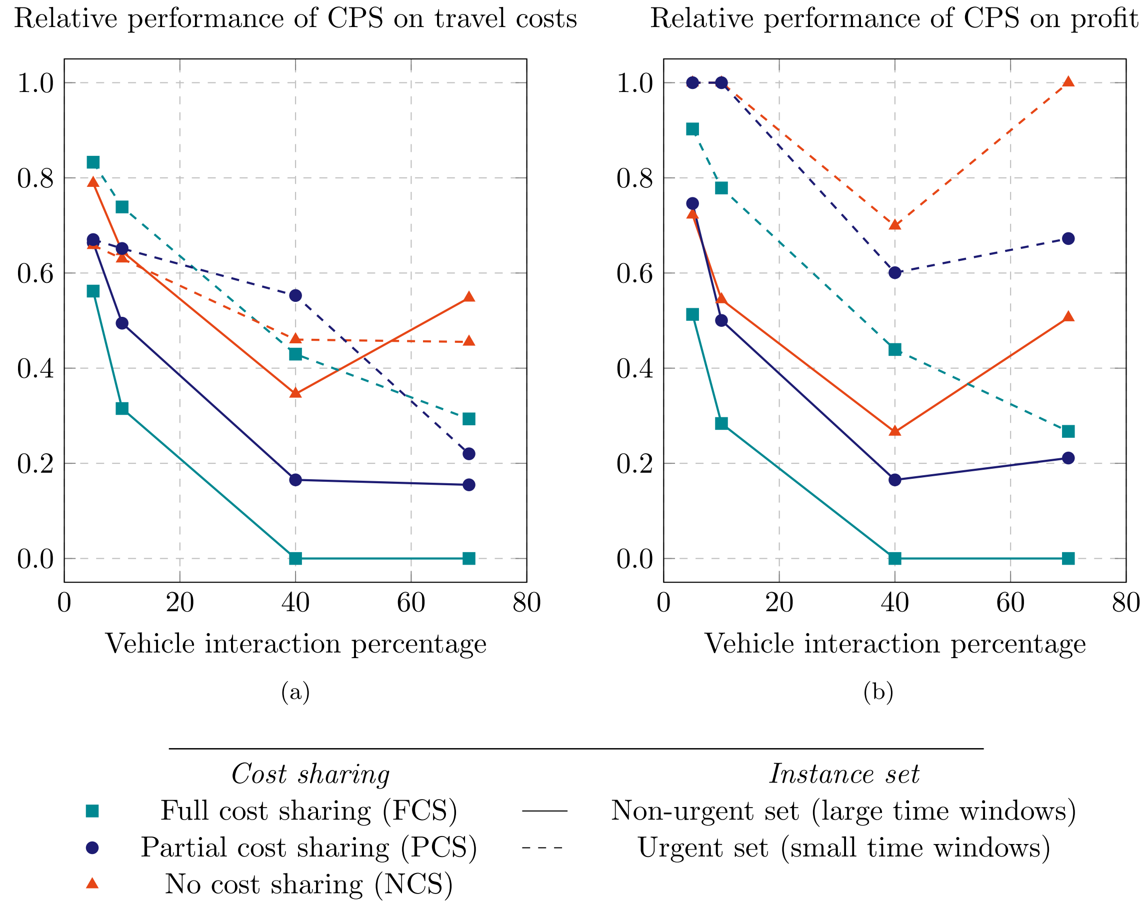

Fig. 14. Performance of CPS on travel costs and profit relative to all three position sharing methods (FPS, CPS, and NPS) on varying VIP for the urgent and the non-urgent instance set. A performance value of 1 indicates that CPS performs best (has lowest travel costs or highest profit) among FPS, CPS, and NPS, whereas a performance value of 0 indicates that CPS performs worst (has highest travel costs or lowest profit) among FPS, CPS, and NPS. (For a VIP of 100%, all three position sharing policies perform equally).

For the urgent instances, it is expected that current position information is more important than for non-urgent instances. For the later ones, there might be enough time for efficient planning and replanning, while urgent orders need a vehicle relatively soon to be picked up before the latest service time. Only vehicles that are close enough to the pickup location are feasible ones.

In Fig. 14, we show the relative performance of the CPS policy on travel costs and profit with respect to the values obtained with the other position sharing policies. Per cost sharing policy and per interaction percentage, the relative performance of CPS on travel costs is calculated by $(TC_{\max} - TC(\text{CPS}))/(TC_{\max} - TC_{\min})$, where $TC_{\max}$ denotes the maximum of $TC(\text{FPS})$, $TC(\text{CPS})$, and $TC(\text{NPS})$, and $TC_{\min}$ denotes the minimum of these three values. Equivalently, the relative performance of CPS on profit is given by $(PR(\text{CPS}) - PR_{\min})/(PR_{\max} - PR_{\min})$.

From Fig. 14a, we observe that CPS obtains relatively lower travel costs for the urgent set than for the non-urgent set. Only for NCS and a VIP of 5% or 70%, CPS performs better on the non-urgent set than on the urgent set. Note also that CPS performs worst (and has even higher travel costs than NPS, see Fig. 8b) with FCS and higher VIPs for the non-urgent set, while this is not the case for the urgent set.

The effect is even larger on profits, as shown in Fig. 14b. CPS always performs better on the urgent set than on the non-urgent set. Furthermore, CPS reaches maximal performance on profit for a VIP of 5% or 10% with PCS or NCS on the urgent set. (Indeed, CPS outperforms FPS in these cases; see Fig. E.1c).

Besides the differences in the value of position sharing, the results for the urgent instances are similar to the results on the instances with lower capacity: the service level is lower than for the non-urgent set, and the FCS policy results in relatively higher travel costs and lower profits (see Fig. E.1).

7. Insights for platform providers

We have compared different levels of information sharing that can be used in platform-based collaborative vehicle routing and have demonstrated the influence of the different information sharing policies on service levels, total travel costs, and total profits. We have also discussed interaction effects of the different information sharing policies with vehicle interaction percentages, number of auctions, and travel cost penalties. Finally, we have investigated the effects of urgency and the ratio of supply and demand on the value of information. Our findings contain various insights for platform providers.

- **Impact of information sharing:** Sharing cost information and route information improves solutions in terms of all observed criteria, but the effect of cost information is generally the dominant factor. Thus, if carriers are hesitant to share route information with the central platform, platform providers may want to encourage them to share at least their cost information. Since this is only shared on a local level (between individual customers and carriers) it may more easily be accepted. Moreover, partial information sharing leads to significant improvements in all conducted experiments and should therefore be considered when full information sharing is not accepted.
- **Route information sharing:** Sharing route information becomes more important if only low vehicle interaction percentages are possible. Significant improvements in travel costs and profits are then obtained. In this situation, it is important for a platform provider to convince carriers of the need of route information sharing. The increased profits resulting from sharing position information may be used by platform providers to create specific incentives for sharing this kind of information.
- **Cost information sharing:** Sharing cost information is mainly important for obtaining high service levels, i.e., a high customer acceptance rate. In most cases, no orders are rejected if costs are shared, although no or partial cost sharing results in several order rejections. To offer a good service to customers, a platform provider may want to set strong incentives for carriers to share cost information. If the platform does not only provide contact details of carriers to customers, but actually acts as facilitator for local communication between individual carriers and customers, it can be designed in such a way that costs need to be provided.
- **No cost sharing:** Travel costs (and the related congestion and emission penalties) are significantly larger when no costs are shared at all. If carriers do not want to share their costs every time, occasional sharing already results in a notable improvement. Again, this is also serving the self-interest of carriers since profits are significantly lower if they do not share costs at all.
- **Impact of carrier interaction:** It is not necessary that customers interact with all available carriers. An interaction percentage of approximately 40% generally already leads to a significant share of the possible gains from information sharing, especially if costs are fully or partially shared and full route plans are shared. For these situations, service levels as well as profits do not increase if the vehicle interaction percentage is increased above 40%.
- **Impact of reauctions:** Allowing reauctions (i.e., iterative interactions between customers and carriers) is important. With more reauctions, the service levels can significantly be improved (if costs are not or not always shared), or the travel costs can be decreased (if costs are shared). It is therefore important that a platform allows multiple interaction rounds for each customer.
  By allowing multiple reauctions, a platform provider could encourage carriers to always share costs, since profits are larger in this setting than with a partial cost sharing policy. The advantage for customers (and indirectly for the platform) is that the service levels are higher with a full cost sharing policy.
  A potential problem is the communicational load that increases for larger numbers of reauctions, especially when more information is shared. If many reauctions are not possible, full cost information sharing can be induced by charging the coalition of carriers a high fine for rejected orders.
- **Impact of emission and congestion penalties:** Imposing emission or congestion penalties on driven distance is another means to stimulate cost information sharing. Strategies in which cost information is not or only partially shared, are more significantly impacted by these penalties than a full cost sharing policy.
- **Impact of problem characteristics:** When transportation problems get more difficult (due to a limited number of vehicles or due to more urgent orders), cost information sharing is still relevant for obtaining high service levels. However, total profits might end up higher if cost information is not or not always shared.
  Sharing information about actual vehicle positions becomes much more relevant in scenarios with urgent orders than in situations where longer replanning times are available. Under certain circumstances, sharing full route plans is unnecessary since sharing only current positions already leads to the same profits.
- **Cooperation incentives:** In conventional central optimization approaches, it is generally assumed that carriers trust on a (fair) profit sharing mechanism by the platform. In reality, however, carriers often want to be given a certain leeway in their own decision-making rather than being dominated by a central decision maker. The approach proposed in this work takes this aspect into account and lets carriers make their individual decisions. Moreover, individual carriers may not be convinced by average profit gains due to collaboration. The proposed MAS approach enables alternative incentives, e.g., active design of the bidding process such that carriers learn trough experimentation that they benefit from information sharing.
- **Manipulation:** In particular situations, the proposed auctioning approach might be vulnerable with respect to manipulation. Carriers can, for instance, speculate on untruthful bidding. If a carrier reports lower marginal costs than it actually has, a customer might select this carrier instead of a carrier that contributes most to the global objectives. There are several ways to deal with this (e.g., relaxing the assumption of fixed customer prices), and platform providers should find the most suitable method for their specific business model.

8. Conclusions and future research

In this work, we have investigated the value of sharing different types and amounts of carrier information for solving real-world platform-based collaborative vehicle routing problems. While the need for realistic information sharing is consistently emphasized in research on collaborative transportation, dedicated information sharing policies have hardly been investigated. To model various information sharing policies without a central coordinator, we proposed a Multi-Agent System in which order agents locally organize auctions and vehicle agents respond with proposals for transportation. In our approach, vehicle agents shared two types of information. First, we let vehicles share details about their location by revealing their full route plan, only their position, or no location information at all. Second, we modeled three different bidding strategies that allow vehicles to hide their marginal costs for transporting the auctioned order, to share their marginal cost information only if the order is profitable, or to always share their marginal costs. Thus, the approach explicitly models the information exchange of different actors in the considered collaborative pickup and delivery problem.

Based on computational experiments with various instances, we find that the solutions generally improve in terms of service level, travel costs, and profit when information sharing is increased. Marginal cost information and position information both contribute to solution quality, but cost information has the largest impact and is crucial for high service levels. Generally, when full cost information is available, travel costs can be lower and profits higher than when only partial cost information is available. However, with rather restricted problems (when fleet capacity is limited or orders are urgent), full cost information is not necessary to obtain maximal profits.

Position information is most important if only small groups of vehicles are involved within an auction. Profits can significantly be improved if vehicles' positions are known, and even more if their full route plans are available. With urgent orders, however, information about current vehicle positions is sufficient – knowing full plans then does not improve the solutions.

Our results provide detailed insight into the prevailing trade-off between privacy and protection of competitive advantages, on the one hand, and service levels, total travel costs, carriers' profits, and communicational requirements, on the other hand. Sharing information plays a major role in improving routing solutions, although individual carriers may be hesitant to reveal it. Hence, depending on the specific application or desired outcomes, platform providers could stimulate sharing of certain kinds or parts of information. We show that allowing multiple interaction rounds and imposing emission or congestion penalties to carriers can stimulate them to share more information.

Quite remarkably, an interaction with only about 40% of the vehicles enables the largest part of possible improvements. Platform providers may thus not need to require full information sharing – as widely assumed in collaborative transportation research – from their customers and could possibly convince hesitant individuals to engage in limited information sharing. Moreover, the proposed multi-agent approach does not center the decision-making power and gives individual carriers a certain leeway to make decisions on their own which allows the platform providers to create respective incentives for each of these decisions. The results also demonstrate that an iterative interaction process, for instance, several auction rounds, may reduce costs as well as improve service levels and let carriers realize the benefits from information sharing.

In this work, we modeled vehicles as autonomous decision makers. As part of future work, this approach can also be extended by considering full information sharing and optimization within groups of vehicles belonging to the same carrier. Moreover, studying heterogeneous carriers with different attitudes towards information sharing policies is important since information sharing implies different risks and benefits for different carriers. Larger companies might be more concerned with their privacy or competitive advantage, while smaller firms can obtain larger collaboration advantages. Finally, future research will also address the issue of potential manipulation and further investigate how incentives have to be designed such that platform customers can learn about information sharing benefits through experimentation.

Acknowledgement

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).

Appendix A. Comparison of MAS and centralized heuristics

To compare the MAS with standard centralized heuristics, we have run it on the benchmark DPDP instances of Mitrović-Minić et al. (2004), which were also considered by Holborn (2013). Mitrović-Minić et al. (2004) used cheapest insertion heuristics and tabu search improvement while they focus both on short-term and long-term goals. Holborn (2013) used more advanced insertion and improvement heuristics. Both use a central approach and process orders from certain intervals in batches.

The 90 instances of Mitrović-Minić et al. (2004), containing 100, 500, or 1000 dynamic pickup and delivery requests with time windows, were defined on a complete directed graph. Hence, we restricted I to contain only duplicates of locations within $P \cup D \cup \Omega$. Furthermore, Eq. (27) was replaced by

$$\rho_i^{kt} = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ a duplicate of } \rho_i^{ku} \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = a(\rho_i^{ku}), \tag{29}$$

and Eq. (32) was replaced by

$$\bar{\alpha}_k = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ a duplicate of } \rho_i^{kt} \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = a(\rho_i^{kt}). \tag{30}$$

We applied our MAS using FCS, since prices of orders are not defined on the instances, and did not use position sharing since we always interacted with all vehicles. A maximum number of 100 auctions was allowed for each order. Durations for processing auctions were set to 0.6 times the standard durations.

We ran our MAS 10 times for all instances. Orders were never rejected. Tables A.1–A.3 show the results of the MAS on the 100, 500 and 1000 order instances and compare them with previously published results.

Our MAS improved upon previously published results by about 4% on the 100 order instances, and about 0.5% on average for the 500 and 1000 order instances, with mean improvements of about 2% for the best MAS results.

Hence, we can conclude that our MAS performs as well as current approaches. Although the decentralized approach does not directly take into account combinatorial properties of orders, the orders' active requests immediately after their release time and the frequent reauctioning turn out to be competitive with current centralized approaches.

Table A.1
Comparison of results for the instances with 100 requests from Mitrović-Minić et al. (2004). The columns *M-M* and *H* show the total travel costs found by Mitrović-Minić et al. (2004) and Holborn (2013), respectively. The columns *Min* and *Avg* ± *Std* show the minimum travel costs and the average travel costs plus standard deviation that our MAS obtained within 10 runs. The column I_H shows the improvement of *H* relative to *M-M*, the column I_{Min} shows the improvement of *Min* relative to the best of *M-M* and *H*, and the column I_{Avg} shows the improvement of *Avg* relative to the best of *M-M* and *H*.

Instance	M-M	H	I_H(%)	Min	I_{Min}(%)	Avg ± Std	I_{Avg}(%)
0	2,656.41	2,642.97	0.51	2,432.29	7.97	2,465.86 ± 11.80	6.70
1	2,700.60	2,605.27	3.53	2,302.93	11.60	2,302.93 ± 0.00	11.60
2	2,774.64	2,797.31	−0.82	2,622.01	5.50	2,622.01 ± 0.00	5.50
3	2,853.89	2,695.67	5.54	2,454.04	8.96	2,454.04 ± 0.00	8.96
4	2,787.88	2,727.12	2.18	2,623.02	3.82	2,623.02 ± 0.00	3.82
5	2,965.55	2,790.13	5.92	2,604.05	6.67	2,604.05 ± 0.00	6.67
6	2,631.34	2,596.22	1.33	2,444.16	5.86	2,444.16 ± 0.00	5.86
7	2,674.47	2,725.43	−1.91	2,516.61	5.90	2,516.61 ± 0.00	5.90
8	2,888.39	2,726.56	5.60	2,745.16	−0.68	2,745.16 ± 0.00	−0.68
9	2,978.87	2,778.56	6.72	2,671.48	3.85	2,671.48 ± 0.00	3.85
10	2,576.58	2,523.94	2.04	2,412.21	4.43	2,412.21 ± 0.00	4.43
11	2,812.76	2,638.94	6.18	2,584.79	2.05	2,584.79 ± 0.00	2.05
12	2,677.90	2,658.15	0.74	2,597.43	2.28	2,603.71 ± 2.21	2.05
13	2,703.01	2,594.72	4.01	2,587.62	0.27	2,610.13 ± 11.86	−0.59
14	3,016.79	2,740.93	9.14	2,747.01	−0.22	2,747.01 ± 0.00	−0.22
15	2,759.91	2,684.19	2.74	2,607.91	2.84	2,607.91 ± 0.00	2.84
16	2,694.01	2,539.09	5.75	2,516.35	0.90	2,516.35 ± 0.00	0.90
17	2,894.00	2,698.98	6.74	2,619.64	2.94	2,619.64 ± 0.00	2.94
18	2,696.56	2,703.61	−0.26	2,613.42	3.08	2,613.42 ± 0.00	3.08
19	2,537.65	2,508.39	1.15	2,414.55	3.74	2,414.55 ± 0.00	3.74
20	2,819.49	2,561.16	9.16	2,527.43	1.32	2,527.43 ± 0.00	1.32
21	2,704.22	2,693.24	0.41	2,655.24	1.41	2,655.24 ± 0.00	1.41
22	2,860.85	2,848.82	0.42	2,608.83	8.42	2,608.83 ± 0.00	8.42
23	2,479.15	2,388.92	3.64	2,473.74	−3.55	2,473.74 ± 0.00	−3.55
24	2,894.79	2,816.63	2.70	2,616.85	7.09	2,616.85 ± 0.00	7.09
25	2,543.57	2,396.13	5.80	2,398.44	−0.10	2,398.44 ± 0.00	−0.10
26	2,889.89	2,788.19	3.52	2,651.95	4.89	2,651.95 ± 0.00	4.89
27	2,780.39	2,629.16	5.44	2,482.36	5.58	2,482.36 ± 0.00	5.58
28	2,653.85	2,454.65	7.51	2,280.93	7.08	2,282.77 ± 1.27	7.00
29	2,763.60	2,690.50	2.65	2,377.00	11.65	2,377.00 ± 0.00	11.65
Average	2,755.70	2,654.79	3.60	2,539.65	4.19	2,541.79	4.10

Table A.2
Comparison of results for the instances with 500 requests from Mitrović-Minić et al. (2004). The columns *M-M* and *H* show the total travel costs found by Mitrović-Minić et al. (2004) and Holborn (2013), respectively. The columns *Min* and *Avg* ± *Std* show the minimum travel costs and the average travel costs plus standard deviation that our MAS obtained within 10 runs. The column I_H shows the improvement of *H* relative to *M-M*, the column I_{Min} shows the improvement of *Min* relative to *H*, and the column I_{Avg} shows the improvement of *Avg* relative to *H*.

Instance	M-M	H	I_H(%)	Min	I_{Min}(%)	Avg ± Std	I_{Avg}(%)
0	10,053.62	8,739.43	13.07	8,624.81	1.31	8,642.96 ± 39.79	1.10
1	9,699.48	8,349.54	13.92	8,348.31	0.01	8,426.96 ± 27.65	−0.93
2	9,608.40	8,202.41	14.63	8,053.33	1.82	8,078.43 ± 22.75	1.51
3	9,807.06	8,350.71	14.85	8,253.04	1.17	8,253.35 ± 0.40	1.17
4	10,176.05	8,832.41	13.20	8,564.35	3.03	8,692.37 ± 85.38	1.59
5	10,133.55	8,733.24	13.82	8,535.39	2.27	8,548.53 ± 5.05	2.12
6	10,045.82	8,485.01	15.54	8,452.11	0.39	8,541.94 ± 83.63	−0.67
7	9,978.97	8,753.23	12.28	8,326.42	4.88	8,520.84 ± 144.00	2.65
8	9,651.25	8,513.46	11.79	8,172.69	4.00	8,422.74 ± 150.06	1.07
9	9,707.42	8,865.12	8.68	8,316.24	6.19	8,707.15 ± 190.44	1.78
10	9,200.16	8,458.44	8.06	8,195.19	3.11	8,445.93 ± 100.02	0.15
11	9,710.40	8,586.22	11.58	8,248.30	3.94	8,477.08 ± 125.44	1.27
12	9,748.16	8,600.62	11.77	8,420.31	2.10	8,487.48 ± 79.83	1.32
13	9,961.84	8,380.88	15.87	8,557.82	−2.11	8,642.67 ± 35.88	−3.12
14	9,560.35	8,390.46	12.24	8,154.69	2.81	8,200.68 ± 65.42	2.26
15	9,296.75	8,448.59	9.12	8,225.50	2.64	8,352.47 ± 92.73	1.14
16	9,784.43	8,500.53	13.12	8,337.16	1.92	8,479.57 ± 121.81	0.25
17	9,917.51	8,411.73	15.18	8,451.29	−0.47	8,534.18 ± 60.07	−1.46
18	9,729.92	8,554.13	12.08	8,355.77	2.32	8,424.19 ± 48.33	1.52
19	9,721.48	8,297.99	14.64	8,184.58	1.37	8,257.46 ± 51.87	0.49
20	10,118.79	8,742.17	13.60	8,572.05	1.95	8,765.13 ± 154.12	−0.26

(*continued on next page*)

Table A.2 (*continued*)

Instance	M-M	H	I_H(%)	Min	I_{Min}(%)	Avg ± Std	I_{Avg}(%)
21	9,458.99	8,742.40	7.58	8,220.11	5.97	8,323.56 ± 132.64	4.79
22	10,126.10	8,739.42	13.69	8,695.30	0.50	8,821.53 ± 85.22	−0.94
23	9,879.78	8,533.37	13.63	8,500.39	0.39	8,500.39 ± 0.00	0.39
24	9,313.77	8,572.88	7.95	8,380.59	2.24	8,478.14 ± 100.30	1.11
25	9,637.84	8,323.20	13.64	8,268.16	0.66	8,428.31 ± 99.25	−1.26
26	10,349.09	8,684.06	16.09	8,296.99	4.46	8,399.11 ± 53.83	3.28
27	9,925.99	8,411.79	15.25	8,143.39	3.19	8,338.56 ± 194.02	0.87
28	9,823.70	8,572.82	12.73	8,549.65	0.27	8,687.32 ± 97.74	−1.34
29	9,997.84	8,066.93	19.31	8,416.19	−4.33	8,527.18 ± 174.76	−5.71
Average	9,804.15	8,528.11	12.96	8,360.67	1.93	8,480.21	0.54

Table A.3
Comparison of results for the instances with 1000 requests from Mitrović-Minić et al. (2004). The column *H* shows the total travel costs found by Holborn (2013). For the instances with 1000 requests, only the average result of Mitrović-Minić et al. (2004) is known, given in column *M-M*. The columns *Min* and *Avg* ± *Std* show the minimum travel costs and the average travel costs plus standard deviation that our MAS obtained within 10 runs. The column I_H shows the average improvement of *H* relative to *M-M*. The column I_{Min} shows the improvement of *Min* relative to *H*, and the column I_{Avg} shows the improvement of *Avg* relative to *H*.

Instance	M-M	H	I_H(%)	Min	I_{Min}(%)	Avg ± Std	I_{Avg}(%)
0	–	13,909.70	–	13,478.84	3.10	13,938.26 ± 225.96	−0.21
1	–	14,545.70	–	13,813.08	5.04	14,150.75 ± 155.98	2.72
2	–	13,997.10	–	13,672.09	2.32	13,905.62 ± 133.24	0.65
3	–	14,606.10	–	14,139.90	3.19	14,516.25 ± 249.31	0.62
4	–	14,257.40	–	13,898.35	2.52	14,163.47 ± 167.71	0.66
5	–	14,312.60	–	13,840.33	3.30	14,101.84 ± 204.04	1.47
6	–	13,754.10	–	13,461.04	2.13	13,828.27 ± 206.05	–0.54
7	–	14,202.80	–	13,842.95	2.53	14,077.49 ± 144.50	0.88
8	–	14,003.10	–	13,762.18	1.72	14,121.62 ± 227.92	−0.85
9	–	14,408.00	–	13,871.93	3.72	14,182.39 ± 223.79	1.57
10	–	14,222.80	–	14,076.89	1.03	14,238.56 ± 139.87	−0.11
11	–	14,496.50	–	14,004.43	3.39	14,225.29 ± 158.21	1.87
12	–	14,324.20	–	13,741.95	4.06	14,176.29 ± 214.21	1.03
13	–	14,079.10	–	13,587.32	3.49	13,923.32 ± 148.97	1.11
14	–	13,923.80	–	13,649.60	1.97	14,073.83 ± 257.04	−1.08
15	–	14,463.90	–	14,254.98	1.44	14,450.65 ± 141.55	0.09
16	–	14,398.40	–	14,228.08	1.18	14,450.56 ± 129.29	−0.36
17	–	14,626.00	–	13,897.35	4.98	14,202.78 ± 163.18	2.89
18	–	14,288.40	–	13,586.74	4.91	13,747.76 ± 123.50	3.78
19	–	13,366.20	–	13,206.87	1.19	13,503.14 ± 212.28	−1.02
20	–	14,111.50	–	13,732.49	2.69	13,956.56 ± 182.77	1.10
21	–	14,140.10	–	13,817.91	2.28	14,214.95 ± 245.40	−0.53
22	–	14,010.20	–	13,732.70	1.98	14,085.22 ± 198.89	−0.54
23	–	14,077.90	–	14,037.16	0.29	14,277.14 ± 129.58	−1.42
24	–	14,203.80	–	13,936.68	1.88	14,156.15 ± 180.83	0.34
25	–	13,709.20	–	13,270.70	3.20	13,626.73 ± 214.01	0.60
26	–	14,080.90	–	13,545.05	3.81	13,972.07 ± 215.28	0.77
27	–	13,907.80	–	13,950.05	−0.30	14,151.44 ± 89.46	−1.75
28	–	14,655.20	–	14,307.43	2.37	14,463.46 ± 93.16	1.31
29	–	14,562.10	–	13,950.44	4.20	14,182.80 ± 139.83	2.60
Average	17,610.45	14,188.15	19.43	13,809.85	2.65	14,102.16	0.59

Appendix B. Marginal cost for including an order

Given a vehicle k and an order $\bar{c}$ that is requested to be included into the route at time t, define the following:

- $\bar{\alpha_k}$ is the location of vehicle k at time t, i.e.,
  - if $t < d(\alpha_k)$, then $\bar{\alpha_k} = \alpha_k$;
  - if $\exists\ i\ a(\rho_i^{kt}) \leqslant t < d(\rho_i^{kt})$, then

$$\bar{\alpha_k} = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ a duplicate of } \rho_i^{kt} \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = t; \tag{31}$$

  - if $\exists\ i\ d(\rho_{i-1}^{kt}) \leqslant t < a(\rho_i^{kt})$, then

$$\bar{\alpha_k} = \hat{\rho} \text{ for } \hat{\rho} \in I \text{ the actual location of the vehicle at time } t, \text{ with } e_{\hat{\rho}} = l_{\hat{\rho}} = t; \tag{32}$$

- $C^k = \{c \in C \mid \exists\ i(\rho_i^{kt} = p_c \wedge s(\rho_i^{kt}) > t)\} \cup \{\bar{c}\}$ is the set of current orders of k for which the service has not yet started at t together with the new order;
- $D^k = \{c \in C \mid \exists\ i, j(\rho_i^{kt} = p_c \wedge s(\rho_i^{kt}) > t \wedge \rho_j^{kt} = d_c \wedge s(\rho_j^{kt}) \leqslant t)\}$ is the set of orders of k for which the delivery has not yet started, but the pickup has already started;
- $N = \{\bar{\alpha_k}\} \cup \{\omega_k\} \cup \{p_c | c \in C^k\} \cup \{d_c | c \in C^k\} \cup \{d_c | c \in D^k\}$ is the set of all relevant locations;
- $E = \{(i, j) | i, j \in N,\ i \neq \omega_k, j \neq \bar{\alpha_k},\ i \neq j\}$ is the set of all relevant arcs;
- $\bar{q_i} = \begin{cases} q_c & \text{if } i = p_c \text{ for } c \in C^k \\ -q_c & \text{if } i = d_c \text{ for } c \in C^k \cup D^k \\ 0 & \text{if } i = \omega_k \\ \sum_{c \in D^k} q_c & \text{if } i = \bar{\alpha_k} \end{cases}$ represents the quantity to pick up at location i.

The vehicle then needs to solve the following mathematical program $M^{\bar{c}}$, which is a slight adaptation of the single-vehicle PDP.

$$\min \sum_{(i,j) \in E} c_{ij} x_{ij} \tag{33}$$

subject to

$$\sum_{i | (i,j) \in E} x_{ij} = 1\ \forall\ j \in N \tag{34}$$

$$\sum_{j | (i,j) \in E} x_{ij} = 1\ \forall\ i \in N \tag{35}$$

$$y_j - y_i - W_i x_{ij} \geqslant \bar{q_j} - W_i\ \forall\ (i, j) \in E \tag{36}$$

$$y_i \geqslant \max(0, \bar{q_i})\ \forall\ i \in N \tag{37}$$

$$y_i \leqslant \min(c_k, c_k + \bar{q_i})\ \forall\ i \in N \tag{38}$$

$$z_{d_c} - z_{p_c} \geqslant 0\ \forall\ c \in C^k \tag{39}$$

$$z_j - z_i - M_{ij} x_{ij} \geqslant d_i + t_{ij} - M_{ij}\ \forall\ (i, j) \in E \tag{40}$$

$$z_i \geqslant e_i\ \forall\ i \in N \tag{41}$$

$$z_i \leqslant l_i\ \forall\ i \in N \tag{42}$$

where

$$W_i = \min(c_k, c_k + \bar{q_i})\ \forall\ i \in N \tag{43}$$

$$M_{ij} = \max(0, l_i + d_i + t_{ij} - e_j)\ \forall\ (i, j) \in E \tag{44}$$

Constraints 36 and 40 are the linearized forms of $y_j \geqslant (y_i + \bar{q_j}) x_{ij}$ and $z_j \geqslant (z_i + d_i + t_{ij}) x_{ij}$, see Parragh et al. (2008, p. 86) and Cordeau (2006, p. 575).

Equivalently, define the mathematical program M where $\bar{c}$ is omitted from C^k. The marginal cost $MC^{kt}(\bar{c})$ for order $\bar{c}$ is defined as the solution of $M^{\bar{c}}$ minus the solution of M.

Appendix C. Results for low capacity set (75 vehicle instances)

The full results for the low capacity set are given in Fig. C.1, Fig. C.2, and Fig. C.3.

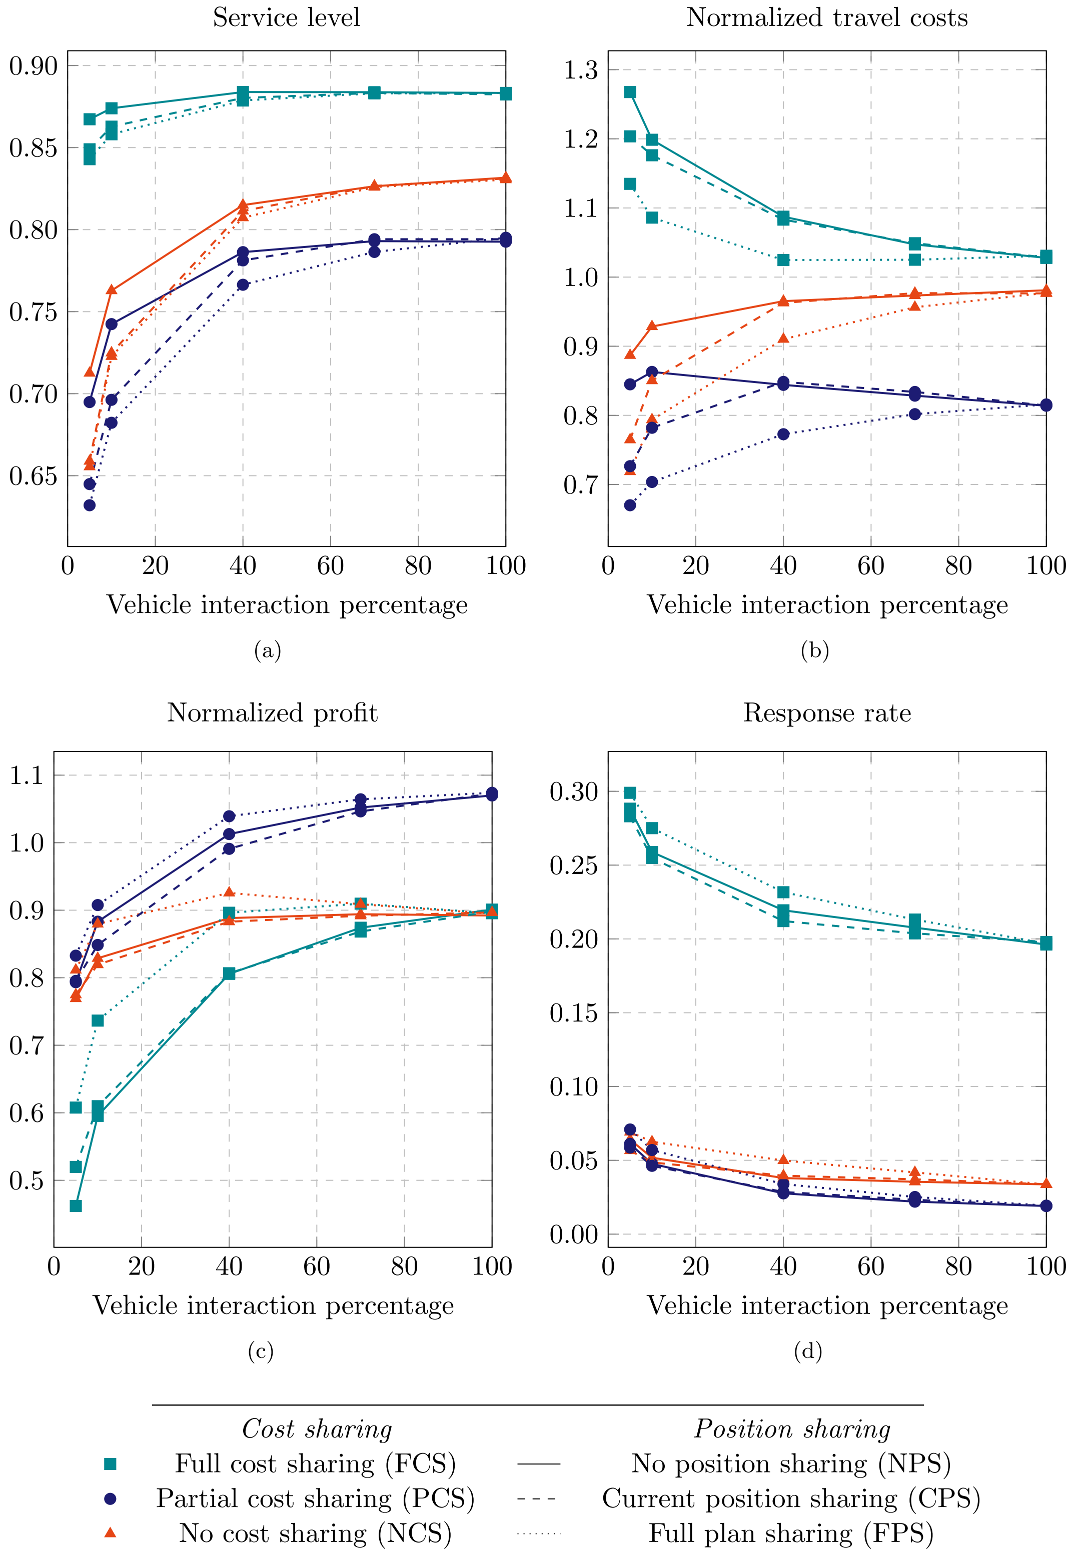

Fig. C.1. Mean results on varying VIP for the three cost sharing and the three position sharing methods on the low capacity set. In (c), there is no fine for rejected orders ($\gamma = 0$).

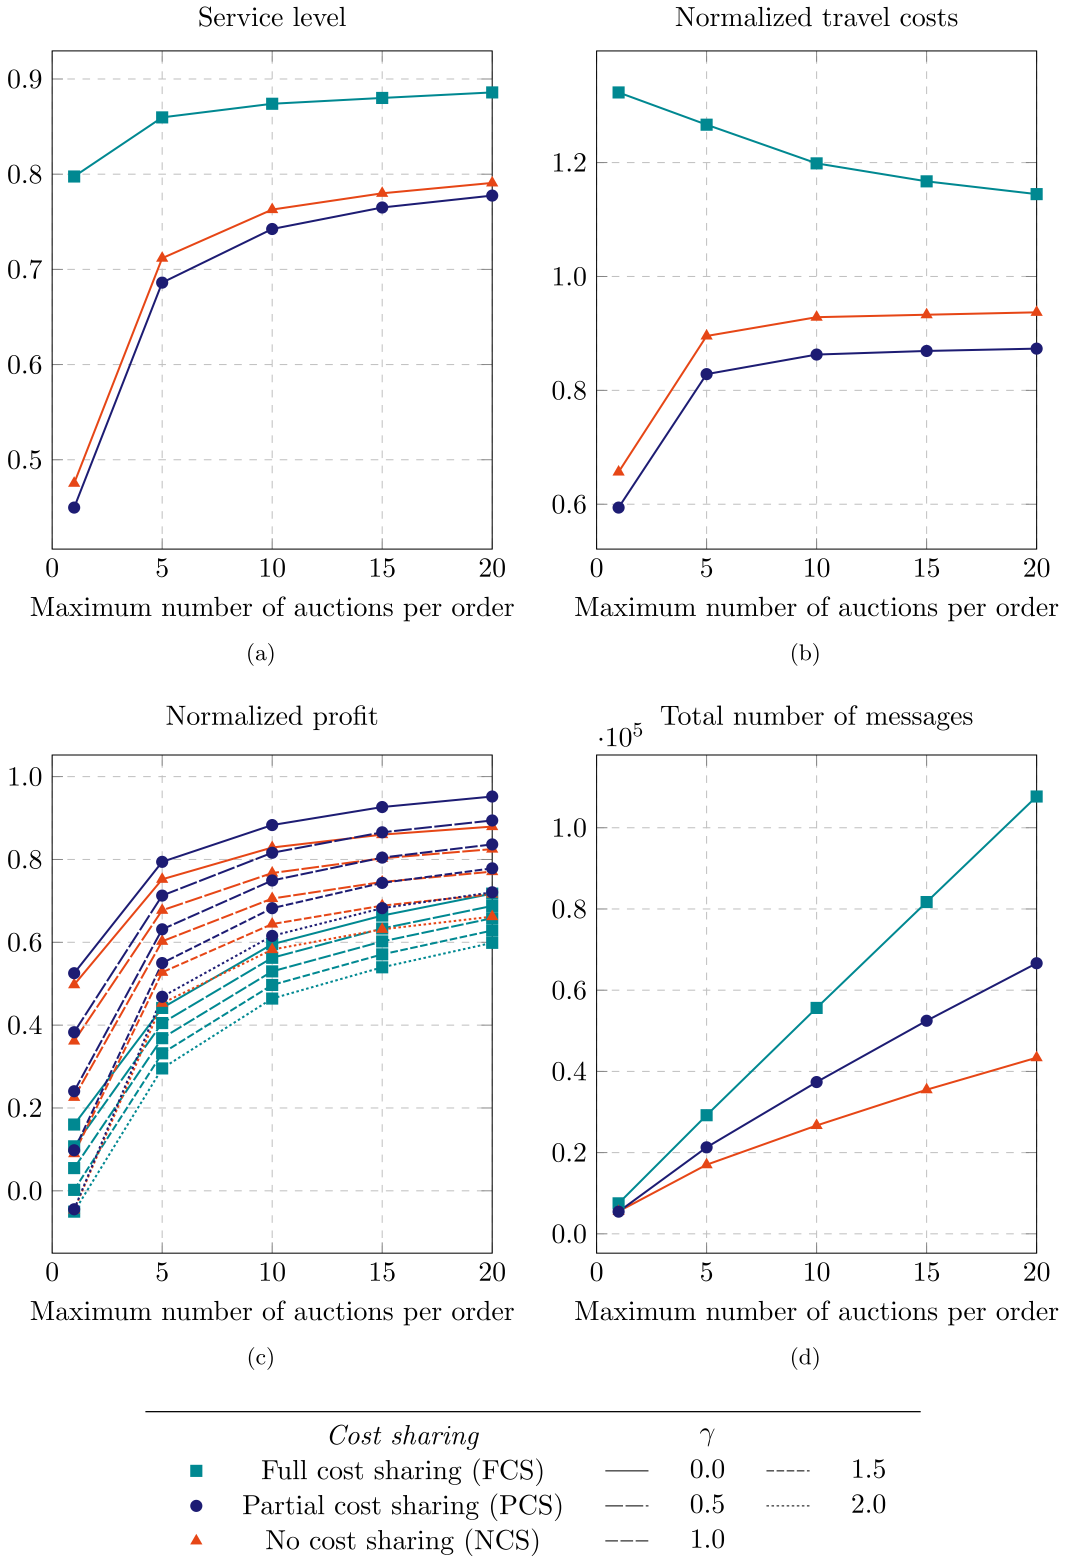

Fig. C.2. Mean results on varying maximum number of auctions per order for the three cost sharing methods and different fines per rejected order (γ) on the low capacity set.

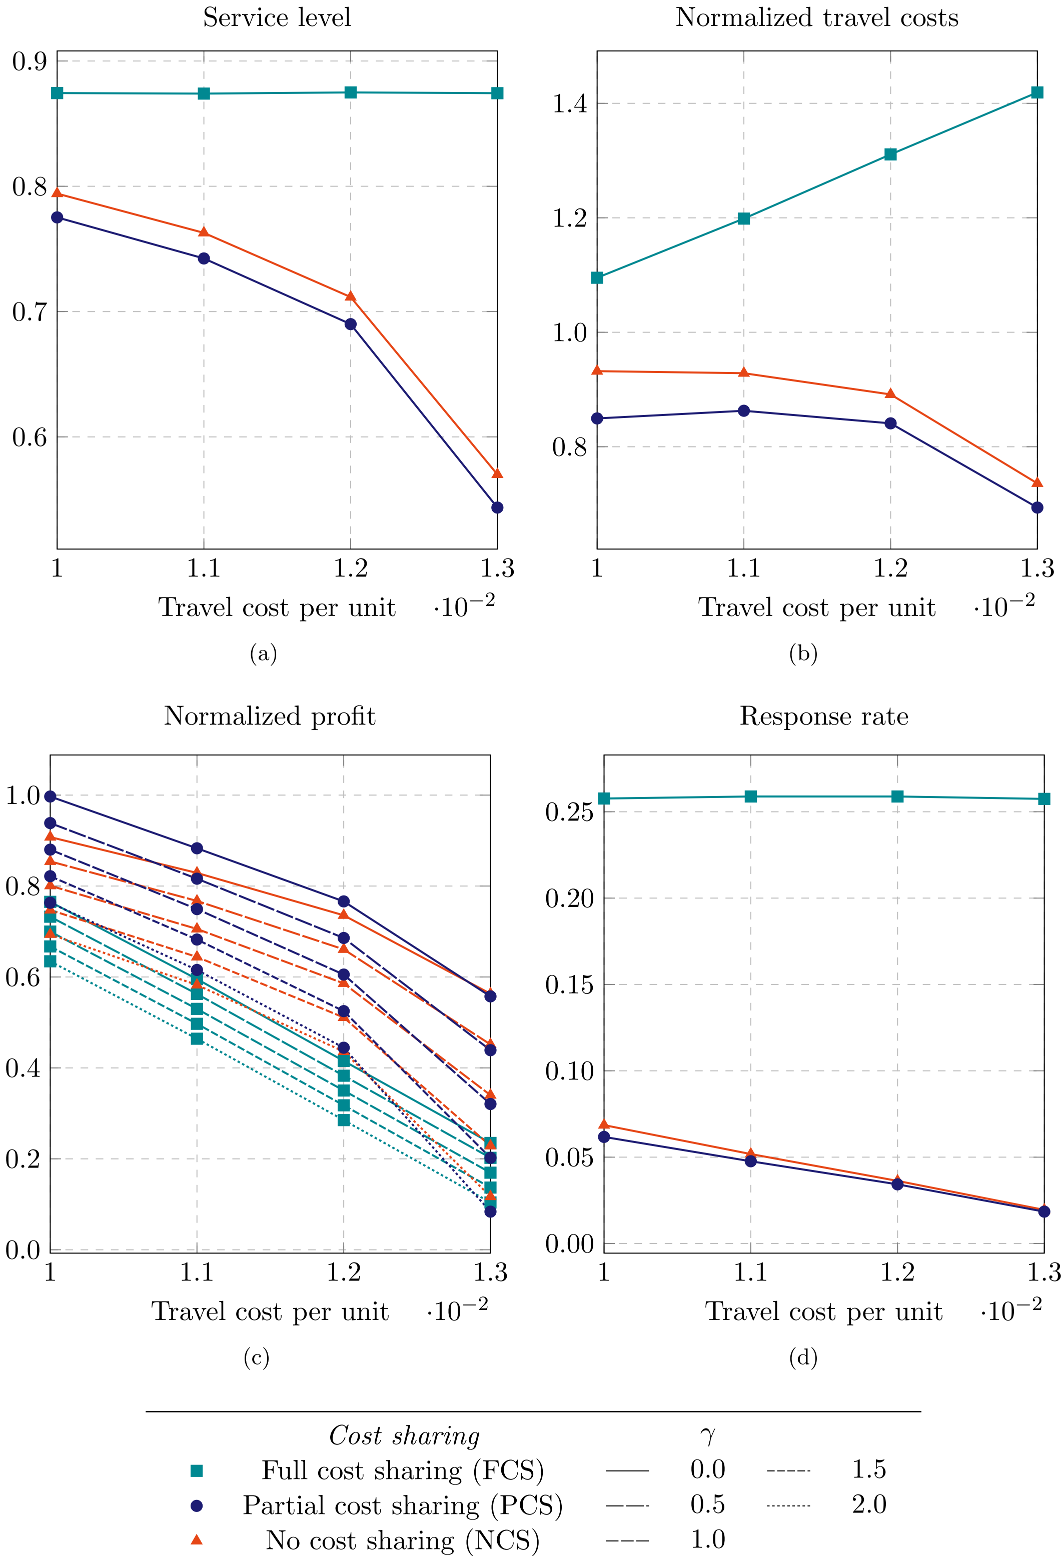

Fig. C.3. Mean results on varying travel costs for the three cost sharing methods and different fines per rejected order (γ) on the low capacity set.

Appendix D. Results for medium capacity set (100 vehicle instances)

The full results for the medium capacity set are given in Fig. D.1, Fig. D.2, and Fig. D.3.

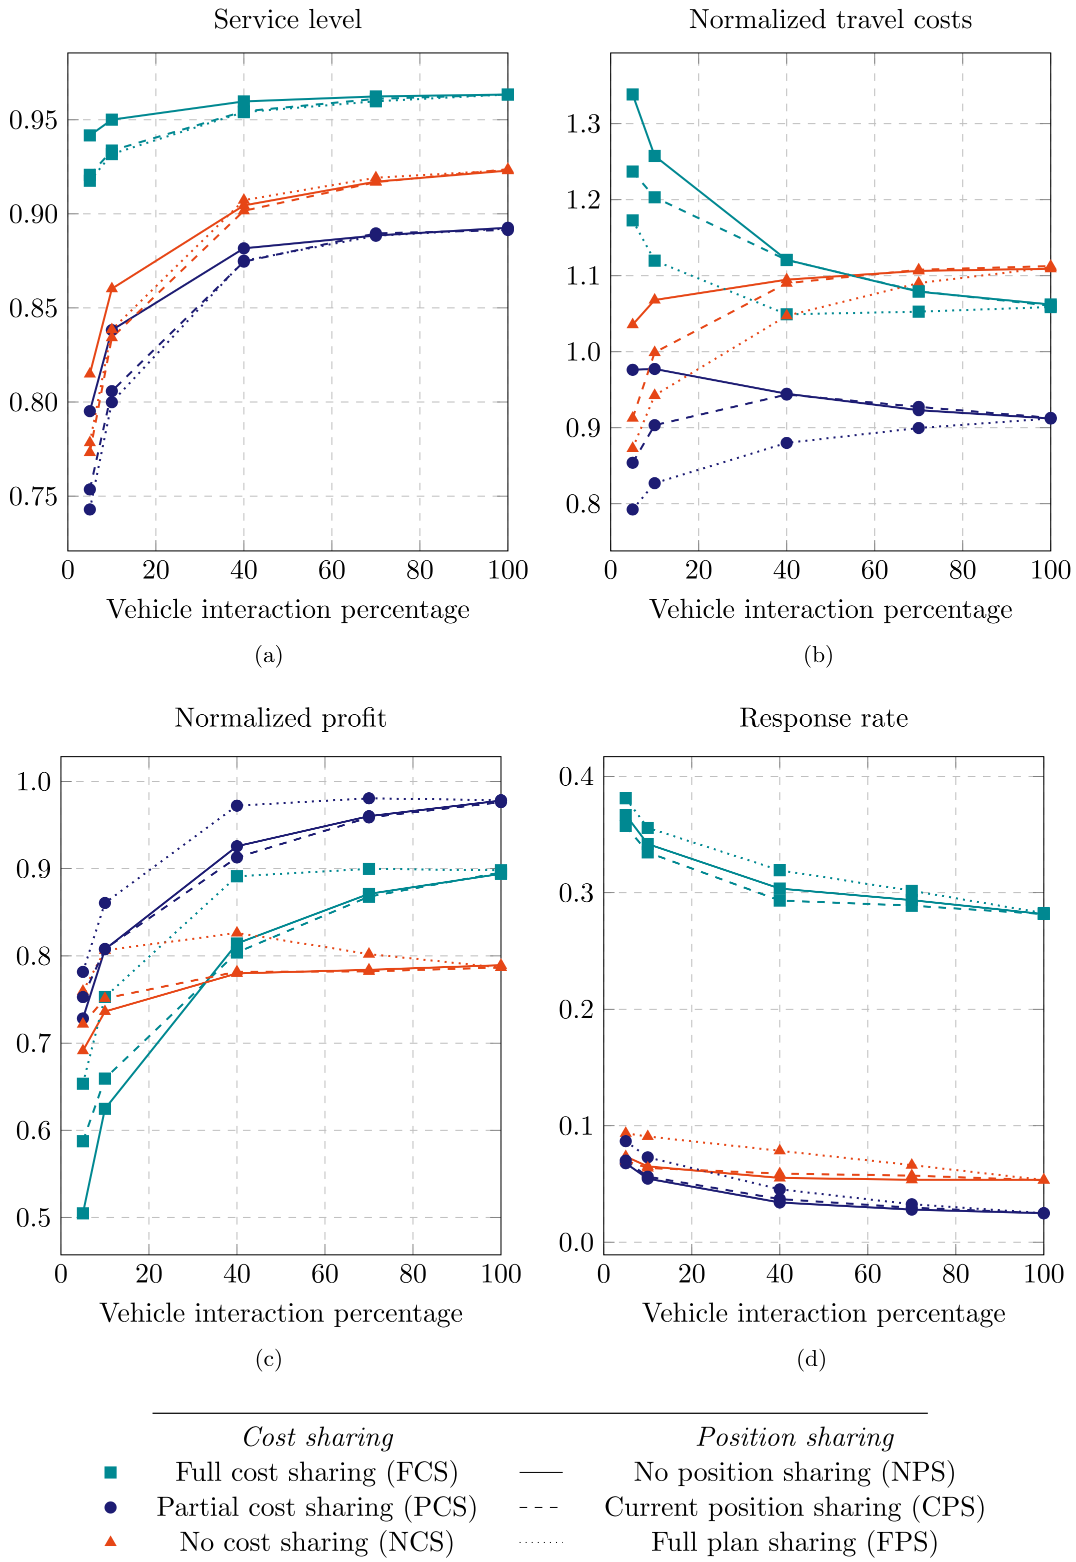

Fig. D.1. Mean results on varying VIP for the three cost sharing and the three position sharing methods on the medium capacity set. In (c), there is no fine for rejected orders ($\gamma = 0$).

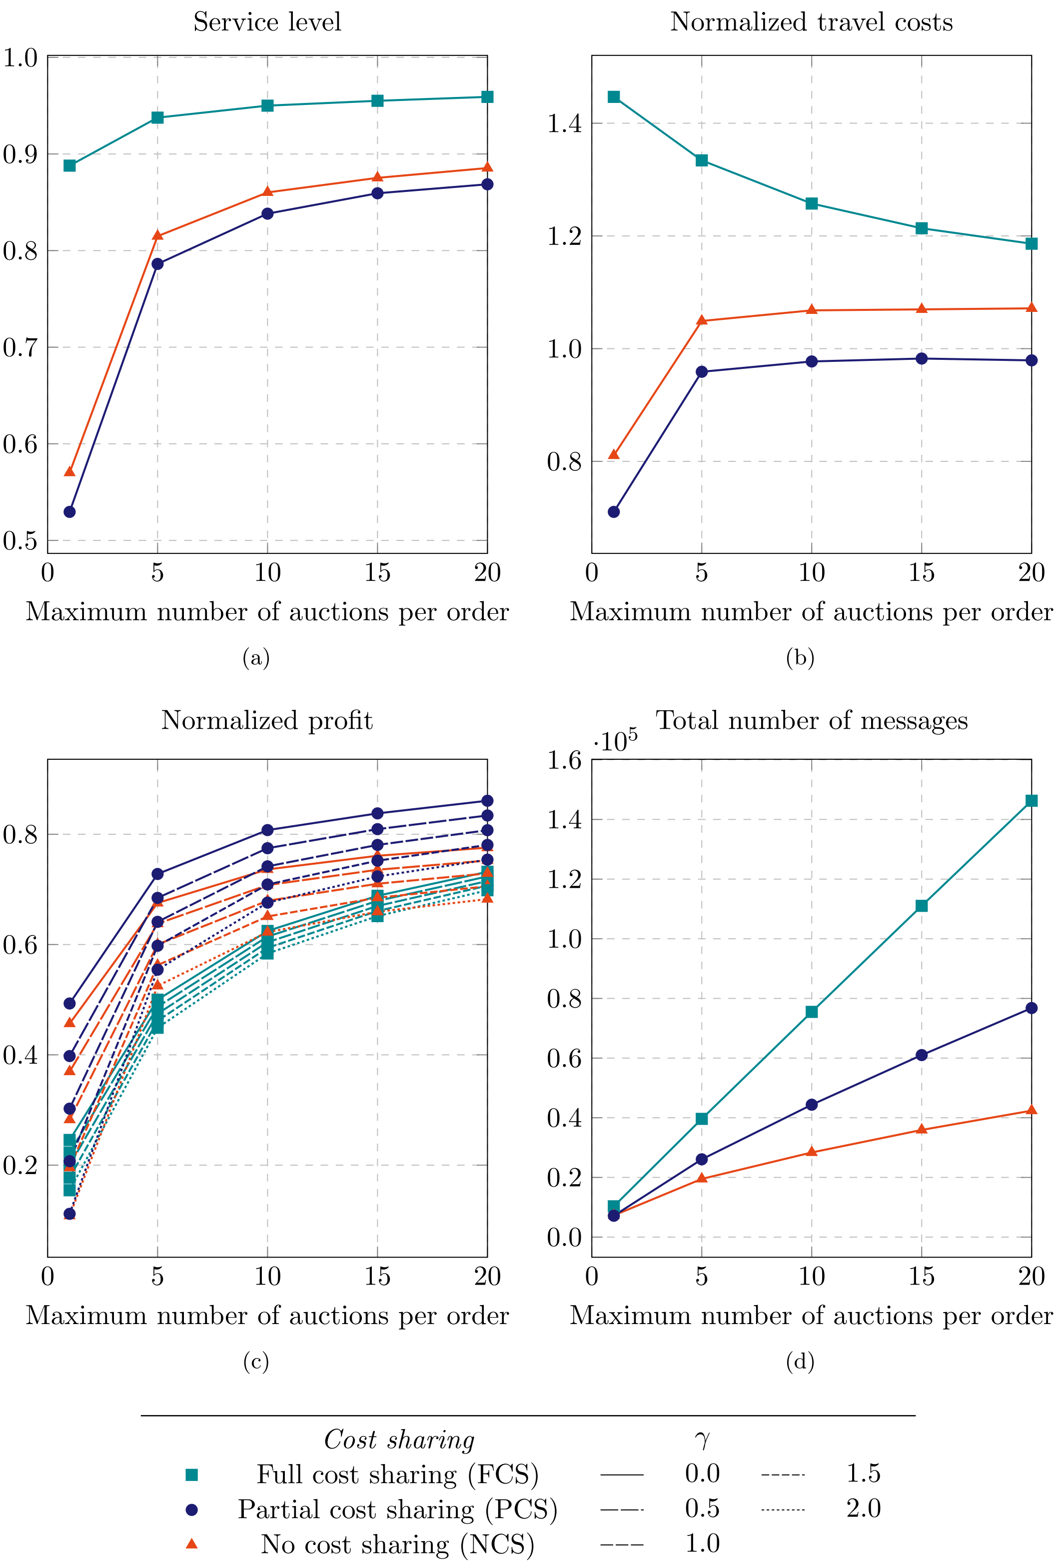

Fig. D.2. Mean results on varying maximum number of auctions per order for the three cost sharing methods and different fines per rejected order (γ) on the medium capacity set.

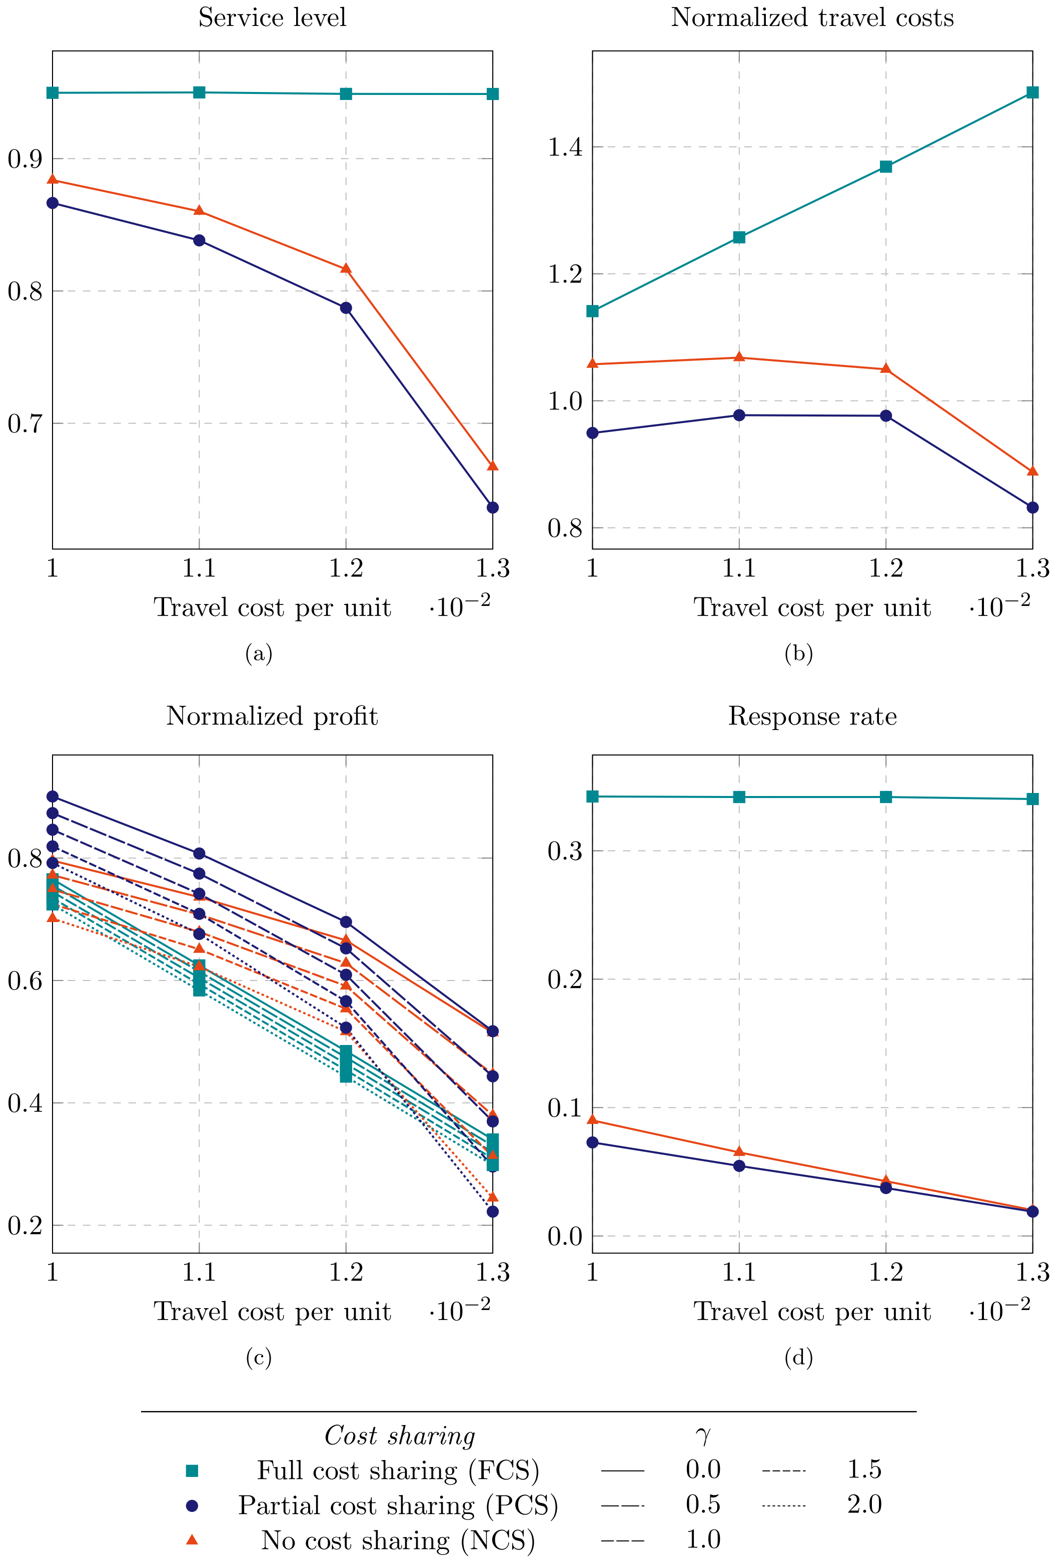

Fig. D.3. Mean results on varying travel costs for the three cost sharing methods and different fines per rejected order (γ) on the medium capacity set.

Appendix E. Results for urgent set (small time window instances)

The full results for the urgent set are given in Fig. E.1, Fig. E.2 and Fig. E.3.

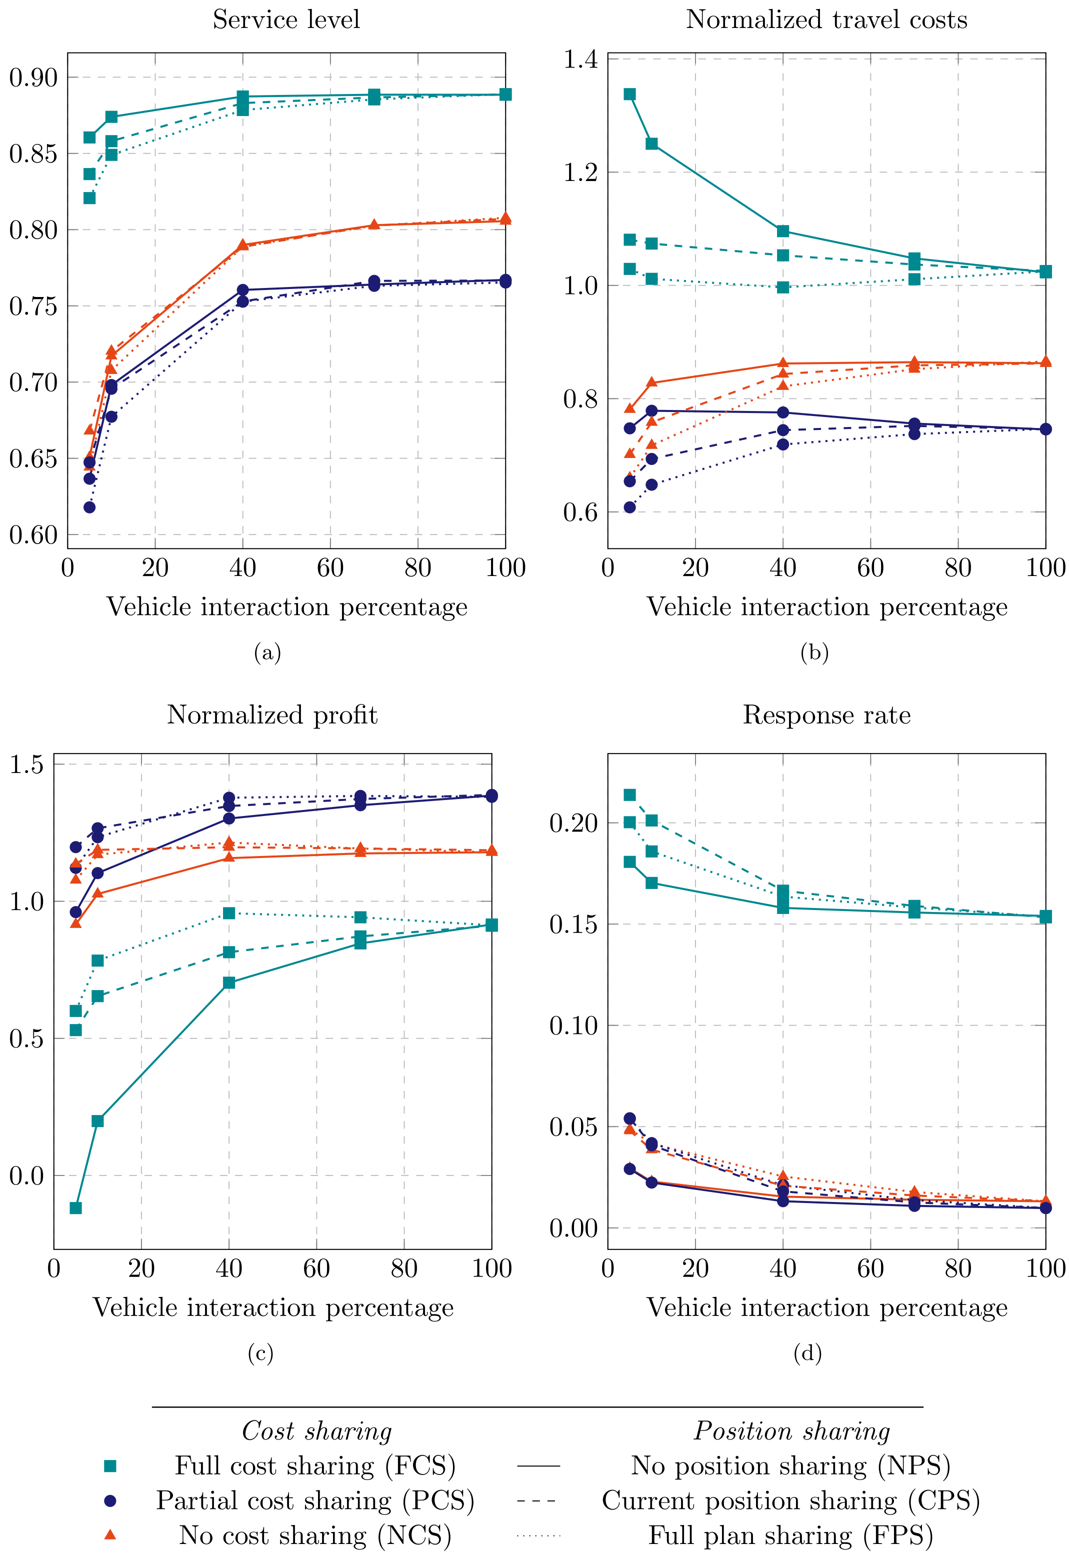

Fig. E.1. Mean results on varying VIP for the three cost sharing and the three position sharing methods on the urgent set. In (c), there is no fine for rejected orders ($\gamma = 0$).

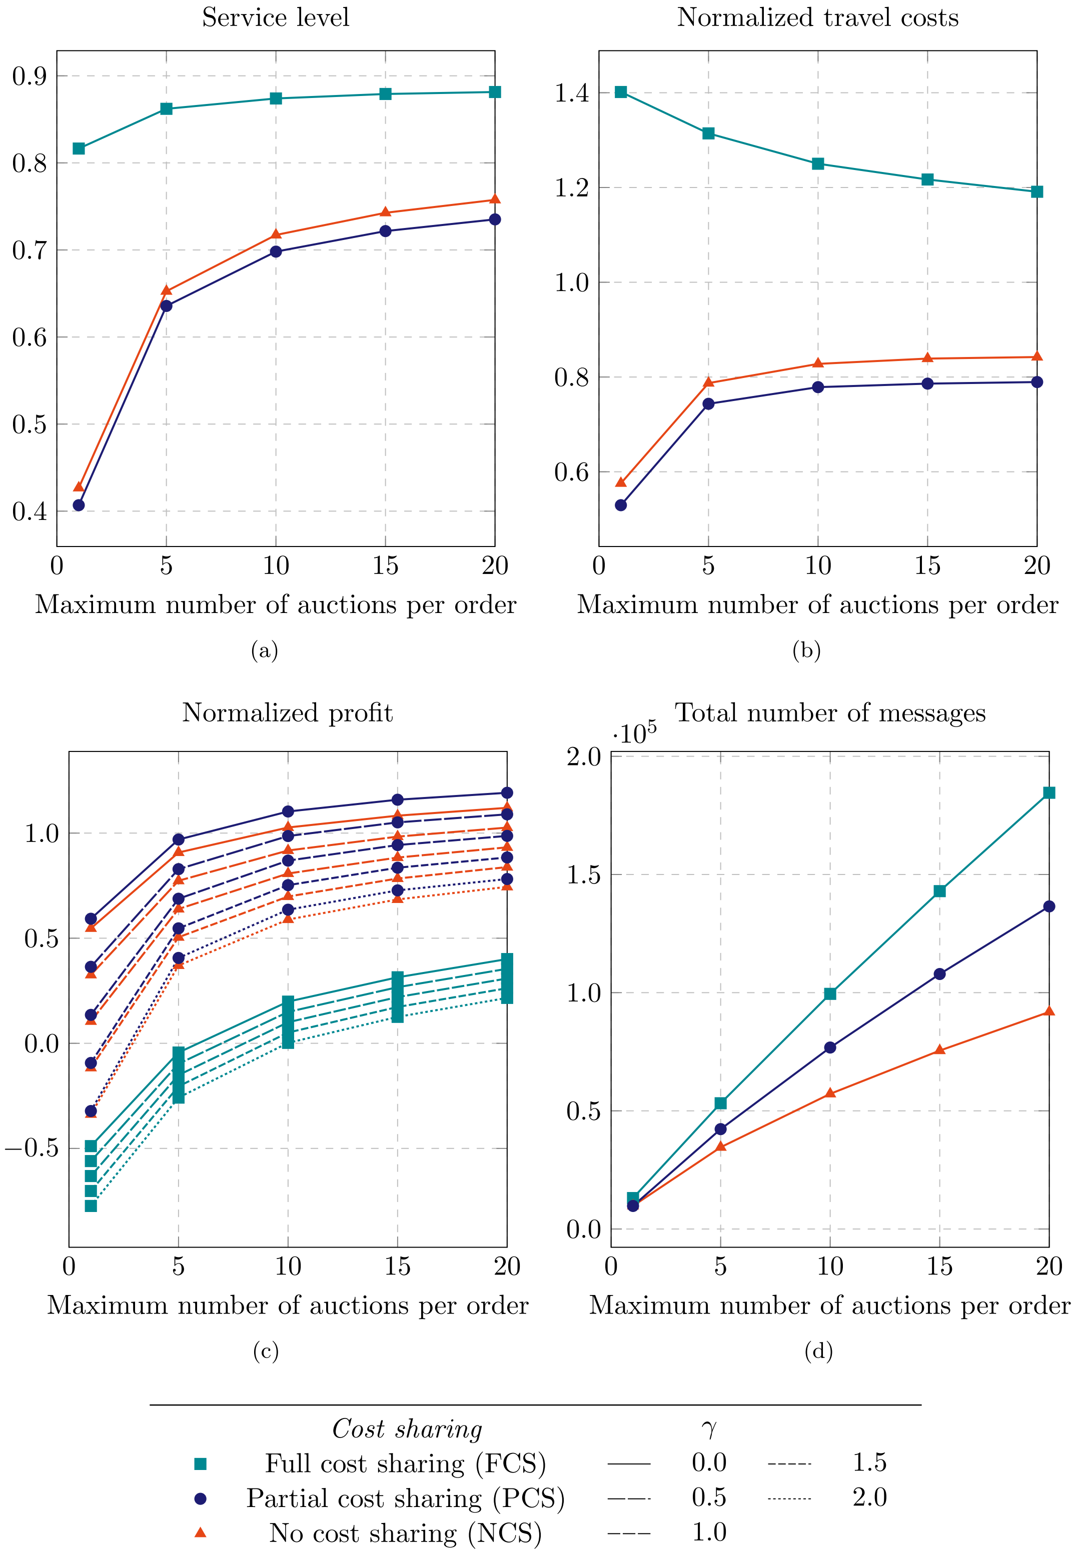

Fig. E.2. Mean results on varying maximum number of auctions per order for the three cost sharing methods and different fines per rejected order (γ) on the urgent set.

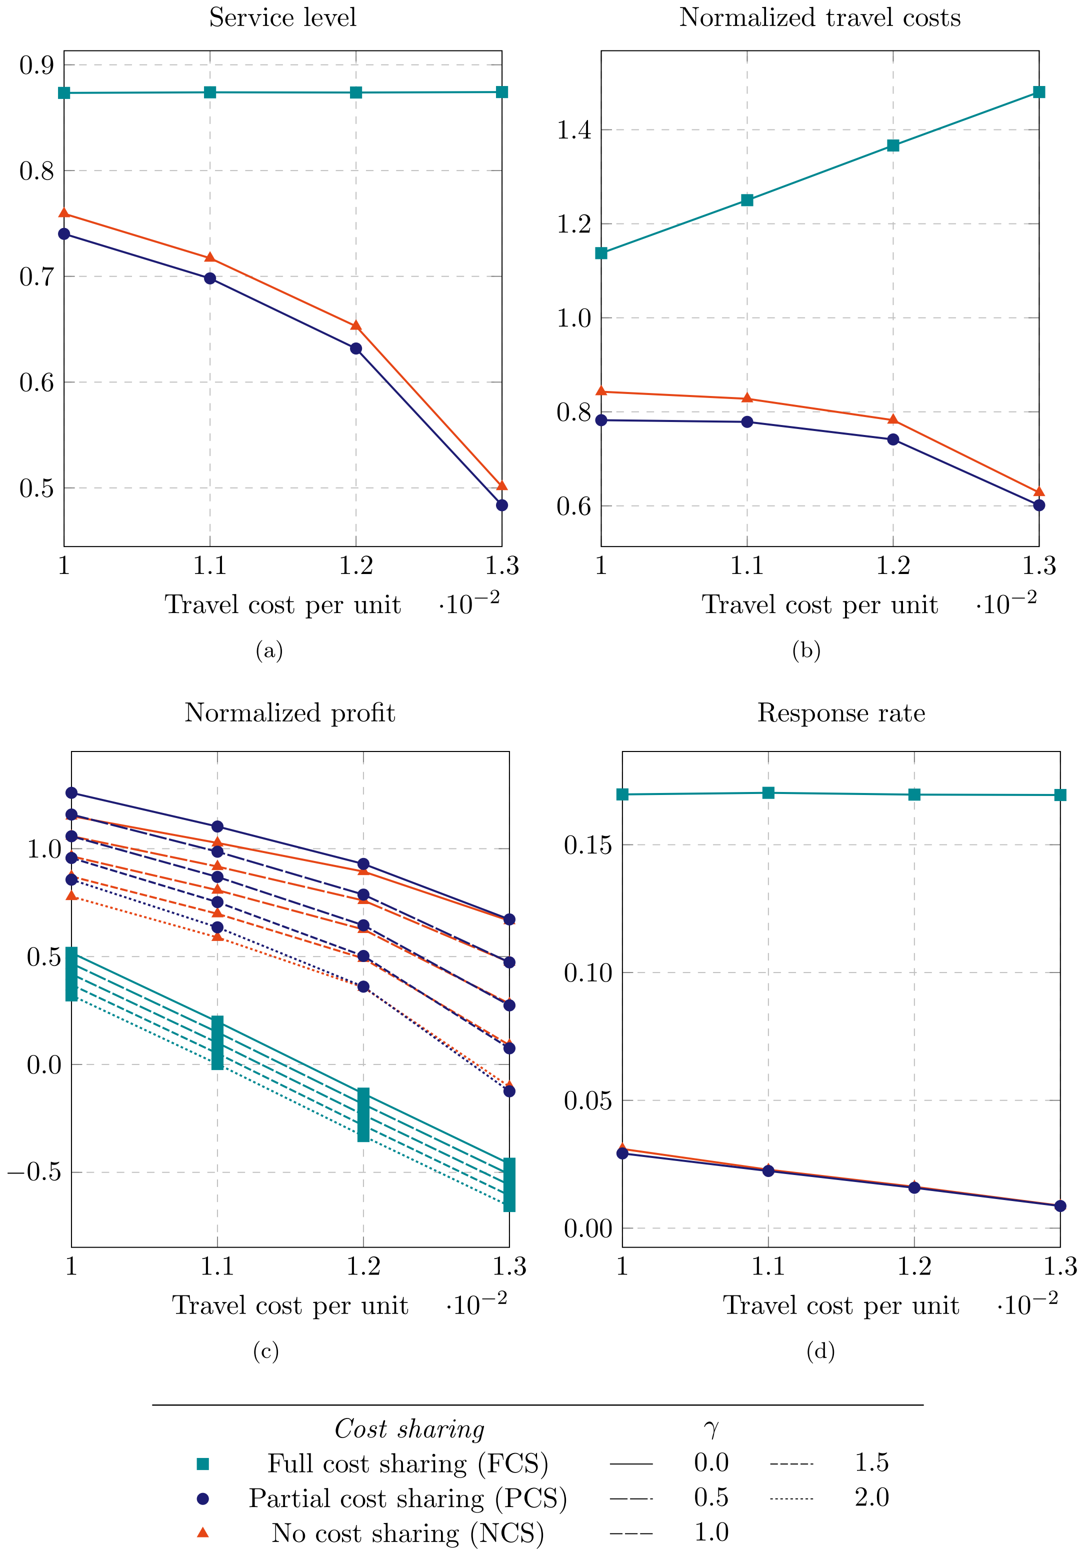

Fig. E.3. Mean results on varying travel costs for the three cost sharing methods and different fines per rejected order (γ) on the urgent set.

References

Berger, S., Bierwirth, C., 2010. Solutions to the request reassignment problem in collaborative carrier networks. Transp. Res. Part E 46, 627–638.
Cleophas, C., Cottrill, C., Ehmke, J.F., Tierney, K., 2019. Collaborative urban transportation: Recent advances in theory and practice. Eur. J. Oper. Res. 273, 801–816.
Cordeau, J.-F., 2006. A branch-and-cut algorithm for the dial-a-ride problem. Oper. Res. 54, 573–586.
Cruijssen, F., Dullaert, W., Fleuren, H., 2007. Horizontal cooperation in transport and logistics: A literature review. Transport. J. 46, 22–39.
Dahl, S., Derigs, U., 2011. Cooperative planning in express carrier networks – an empirical study on the effectiveness of a real-time decision support system. Decis. Support Syst. 51, 620–626.
Dai, B., Chen, H., 2011. A multi-agent and auction-based framework and approach for carrier collaboration. Logist. Res. 3, 101–120.
Dai, B., Chen, H., Yang, G., 2014. Price-setting based combinatorial auction approach for carrier collaboration with pickup and delivery requests. Oper. Res. Int. J. 14, 361–386.
Figliozzi, M.A., Mahmassani, H.S., Jaillet, P., 2004. Competitive performance assessment of dynamic vehicle routing technologies using sequential auctions. Transp. Res. Rec. 1882, 10–18.
Figliozzi, M.A., Mahmassani, H.S., Jaillet, P., 2005. Impacts of auction settings on the performance of truckload transportation marketplaces. Transp. Res. Rec. 1906, 89–96.
Gansterer, M., Hartl, R.F., 2018a. Centralized bundle generation in auction-based collaborative transportation. OR Spectrum 40, 613–635.
Gansterer, M., Hartl, R.F., 2018b. Collaborative vehicle routing: A survey. Eur. J. Oper. Res. 268, 1–12.
Gansterer, M., Hartl, R.F., Savelsbergh, M., 2020. The value of information in auction-based carrier collaborations. Int. J. Prod. Econ. 221, 107485.
Gath, M., 2016. Optimizing Transport Logistics Processes with Multiagent Planning and Control. Springer, Wiesbaden.
Holborn, P.L. (2013). Heuristics for dynamic vehicle routing problems with pickups and deliveries and time windows. PhD thesis, Cardiff University.
Lai, M., Cai, X., Hu, Q., 2017. An iterative auction for carrier collaboration in truckload pickup and delivery. Transp. Res. Part E 107, 60–80.
Li, J., Rong, G., Feng, Y., 2015. Request selection and exchange approach for carrier collaboration based on auction of a single request. Transp. Res. Part E 84, 23–39.
Lyu, X., Chen, H., Wang, N., Yang, Z., 2019. A multi-round exchange mechanism for carrier collaboration in less than truckload transportation. Transp. Res. Part E 129, 38–59.
Máhr, T., Srour, J., de Weerdt, M., Zuidwijk, R., 2010. 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.
Mes, M., van der Heijden, M., Schuur, P., 2013. Interaction between intelligent agent strategies for real-time transportation planning. CEJOR 21, 337–358.
Mitrović-Minić, S., Krishnamurti, R., Laporte, G., 2004. Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows. Transp. Res. Part B 38, 669–685.
Molenbruch, Y., Braekers, K., Caris, A., 2017. Benefits of horizontal cooperation in dial-a-ride services. Transp. Res. Part E 107, 97–119.
Montoya-Torres, J.R., Muñoz-Villamizar, A., Vega-Mejía, C.A., 2016. On the impact of collaborative strategies for goods delivery in city logistics. Prod. Plan. Control 27, 443–455.
Pan, S., Trentesaux, D., Ballot, E., Huang, G.Q., 2019. Horizontal collaborative transport: Survey of solutions and practical implementation issues. Int. J. Prod. Res. 57, 5340–5361.
Parragh, S.N., Doerner, K.F., Hartl, R.F., 2008. A survey on pickup and delivery problems: Part II: Transportation between pickup and delivery locations. J. für Betriebswirtschaft 58, 81–117.
Schulte, F., Lalla-Ruiz, E., González-Ramírez, R.G., Voß, S., 2017. Reducing port-related empty truck emissions: A mathematical approach for truck appointments with collaboration. Transp. Res. Part E 105, 195–212.
Schulte, F., Lalla-Ruiz, E., Schwarze, S., González-Ramírez, R.G., Voß, S., 2019. Scalable core and Shapley value allocation methods for collaborative transportation. Technical report. University of Hamburg.
Van Lon, R.R.S., Holvoet, T., 2017. When do agents outperform centralized algorithms? A systematic empirical evaluation in logistics. Auton. Agent. Multi-Agent Syst. 31, 1578–1609.
Verdonck, L., Caris, A., Ramaekers, K., Janssens, G.K., 2013. Collaborative logistics from the perspective of road transportation companies. Transp. Rev. 33, 700–719.
Wang, X., Kopfer, H., 2014. Collaborative transportation planning of less-than-truckload freight. OR Spectrum 36, 357–380.
Wang, X., Kopfer, H., 2015. Rolling horizon planning for a dynamic collaborative routing problem with full-truckload pickup and delivery requests. Flexible Serv. Manuf. J. 27, 509–533.
Wooldridge, M., 2009. An Introduction to Multi-Agent Systems. John Wiley & Sons, Chichester.