

Interaction between intelligent agent strategies for real-time transportation planning

Martijn Mes · Matthieu van der Heijden · Peter Schuur

© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract In this paper we study the real-time scheduling of time-sensitive full truck-load pickup-and-delivery jobs. The problem involves the allocation of jobs to a fixed set of vehicles which might belong to different collaborating transportation agencies. A recently proposed solution methodology for this problem is the use of a multi-agent system where shipper agents offer jobs through sequential auctions and vehicle agents bid on these jobs. In this paper we consider such a system where *both* the vehicle agents and the shipper agents are using profit maximizing look-ahead strategies. Our main contribution is that we study the interrelation of these strategies and their impact on the system-wide logistical costs. From our simulation results, we conclude that the system-wide logistical costs (i) are always reduced by using the look-ahead strategies instead of a myopic strategy (10–20%) and (ii) the joint effect of two look-ahead strategies is larger than the effect of an individual strategy. To provide an indication of the savings that might be realized under centralized decision making, we benchmark our results against an integer programming approach.

Keywords Multi-agent systems · Collaborative planning · Auctions/bidding · Transportation

1 Introduction

During the last decade there has been a growing interest in collaborative logistics as a result of increasing pressure on shippers and carriers to operate more efficiently. Cooperation among transportation agencies takes place on various organizational and

M. Mes (✉) · M. van der Heijden · P. Schuur
Department of Operational Methods for Production and Logistics, School of Management and Governance, University of Twente, Enschede, The Netherlands
e-mail: m.r.k.mes@utwente.nl

institutional levels, and in various forms. These forms range from spot markets to private fleets. In spot markets, a large number of shippers and carriers exchange loads and vehicle capacity. In private fleets, a shipper has exclusive and direct control over a fleet of vehicles. Situated between these extremes are contractual agreement structures where carriers and/or shippers form a coalition to increase operational efficiency. Potential costs reductions within these coalitions are commonly estimated around 15% of the total transportation cost (Crujssen and Salomon 2004; Ergun et al. 2007a; Schwind et al. 2009).

The focus of this paper is on networks where full truckload shipments (FTL) are offered to a coalition of independent carriers, with the objective to improve the planning solution for the entire system, i.e., to reduce the system-wide logistical costs consisting of travel costs and costs for late deliveries. Given this focus, we distinguish several application areas: (i) shippers that form a core carrier program in which they form partnerships with a few large carriers, (ii) private fleets where individual vehicles are modeled as autonomous entities (Mes et al. 2008; Böhnlein et al. 2010), (iii) freight forwarders with profit centers that operate as independent carriers (Gomber et al. 1997; Krajewska et al. 2008), and (iv) the so-called groupage systems as introduced by Kopfer and Pankratz (1999) which enable the exchange of transportation requests between independent carriers to achieve an equilibrium between requested and available transport resources (Krajewska and Kopfer 2006). In the remainder we refer to these types of possible application areas as *collaborative networks*.

The common denominator in the aforementioned collaborative networks is the need for decentralized planning where partners conduct their planning autonomously and only exchange limited information. A frequently proposed solution concept here are auction mechanisms and more specifically Multi-Agent Systems (MAS). A MAS consists of a group of intelligent and autonomous computational entities (agents) which coordinate their capacities in order to achieve certain (local or global) goals (Wooldridge 1999). These systems are particularly useful to model autonomous decision making in transportation logistics. As noted in Hülsmann et al. (2009), this modeling can be done at various levels of detail, ranging from an agent for a fleet of vehicles to individual vehicle agents, and from freight forwarding agents to agents representing individual transportation jobs. For the purpose of this paper, we limit ourselves to shipper agents and vehicle agents. From an abstract point of view, a shipper agent is responsible for finding transport capacity for an individual transportation job at lowest possible costs and a vehicle agent is responsible for acquiring transportation jobs for a single vehicle and to perform these jobs in an efficient way. Depending on the application area, the shipper agent might represent a shipper but also a carrier that wants to outsource (a part of) his transportation requests, and the vehicle agent might represent an individual truck, an automated guided vehicle, or even a carrier. The main decisions here are the allocation of transportation jobs to vehicles and the timing of these jobs. The allocation of jobs is typically done using a sequential auction procedure where a shipper agent starts an auction for each incoming job and vehicle agents bid on these jobs.

The idea of an auction-based allocation mechanism raises a problem: since jobs arrive in real-time, an optimal allocation can only be derived afterwards, i.e., when all jobs are known. This means that a certain allocation may become unfavorable when

new jobs appear. To overcome this, the individual agents have to take into account future job arrivals in their current decision making. In the literature, several look-ahead policies have been proposed for vehicle agents and shipper agents (see Sect. 2 for an overview). To the best of our knowledge, the *interaction* of intelligent behavior of vehicle agents and shipper agents has never been studied: does the behavior of the individual agents counteract or strengthen each other in terms of the overall system performance? This is the focus of the present paper.

We base our look-ahead policies on the results of two earlier papers. In [Mes et al. \(2008\)](#) two auction strategies for shipper agents are proposed, namely the use of reserve prices and decommitment penalties. In [Mes et al. \(2010\)](#) pricing and scheduling strategies for vehicle agents are proposed where not only the direct costs of jobs are taken into account, but also the impact on future opportunities. The policies in both papers have been designed for individual players operating in spot markets whereas we now consider collaborative networks. This change of application area causes two difficulties. First, our objective differs: in the spot markets we focus on the revenues of a single player whereas in collaborative networks we aim to minimize the system-wide costs. Second, learning might become more difficult: in the spot markets we consider one player which anticipates to the more or less constant behavior of the other players. In collaborative networks all players might adapt themselves to each other which might not converge to a stable behavior. In this paper we study the emergent behavior within such a MAS. The goal of this paper is to study the interrelation of the individual strategies and to benchmark their performance to centralized planning where the individual agencies have to give up their autonomy. Although the latter is not realistic in practice, it provides an upper bound of the performance that can be reached using a MAS.

The remainder of this paper is structured as follows. In Sect. 2 we give a brief overview of the relevant literature and state our contribution. In Sect. 3 we present our model of the transportation market. We present the various look-ahead policies in Sect. 4. We present the experimental settings and numerical results in Sects. 5 and 6 respectively. We close with conclusions in Sect. 7.

2 Literature

In collaborative logistics, carriers and/or shippers may form a coalition in which they exchange requests from various partners to form more efficient routes in terms of utilization rates and idle trips. During the last decade, a number of articles have appeared on this topic. We mention a few examples. [Ergun et al. \(2007b\)](#) considers the case where shippers collaborate to form efficient tours that are offered to carriers. [Puettmann and Stadtler \(2010\)](#) introduce a collaborative planning approach for intermodal freight transportation. [Kopfer and Pankratz \(1999\)](#) introduced the so-called groupage systems which are defined as a logistic interorganisation system which exchanges information and manages capacity balancing by the cooperation between several independent carriers. Further studies on these groupage systems can be found in [Krajewska and Kopfer \(2006\)](#), [Hülsmann et al. \(2009\)](#). Other examples of horizontal collaboration between carriers can be found in

Schönberger (2005), Berger and Bierwirth (2010), Liu et al. (2010), Dahl and Derigs (2011). In the broader field of supply chain management, several collaborative planning approaches have been developed. For example, Dudek and Stadtler (2005) study negotiation-based collaborative planning between two independent supply chain partners and Berger and Schröder (2011) study a decentralized approach for collaborative forwarding of air cargo freight. For an overview on the state-of-the-art of collaborative planning in supply chains we refer to Stadtler (2009).

An important element of collaborative planning is the mechanism used to exchange transportation jobs. A frequently proposed solution for this is to use auction mechanisms. Krajewska and Kopfer (2006) propose a decentralized combinatorial auction model for the collaboration among independent freight forwarding entities. Lee et al. (2007) propose a combinatorial auction mechanism for transportation procurement of a shipper from carriers. Schwind et al. (2009) present an exchange mechanism for intra-enterprise order exchange among profit centers with the purpose of reducing total costs of the entire company. A framework for the comparison between a decentralized auction-based collaborative planning approach and a central planning approach can be found in Berger and Bierwirth (2010).

To facilitate collaboration between carriers and shippers, an often proposed methodology is the use of a MAS because such a system explicitly addresses the autonomy and the specific knowledge of the individual agencies. In two early contributions (Sandholm 1993; Fischer et al. 1996), a network of independent carriers is modeled as a MAS where carriers are represented by agents that communicate and act on a market platform. Kopfer and Pankratz (1999) discuss the modeling of groupage systems as MAS. Gomber et al. (1997) considers a freight forwarder with several independent profit centers. The profit centers are represented by agents and collaboration takes place through auction mechanisms. In Hoen and La Poutre (2004) a MAS is presented for real-time vehicle routing problems with consolidation in a multi-company setting where cargo is assigned to vehicles using a Vickrey auction. A framework for the study of carriers' strategies in an auction marketplace for dynamic full truckload vehicle routing problems with time-windows can be found in Figliozzi et al. (2003). A similar problem is considered in Mes et al. (2007) where a comparison is made between the agent-based approach and centralized optimization methods. Böhnlein et al. (2010) propose a MAS for synchronizing production and distribution within the news paper industry. For a literature survey on MAS in the area of transportation (and traffic) we refer to Chen and Cheng (2010).

As with collaborative planning in general, auction mechanisms are also used in MAS as a mechanism to enable cooperation between the agents, i.e., to enable the exchange transportation jobs. In the transportation area, these agents typically represent resources and/or jobs. Resource agents may strive for utilization and/or profit maximization whereas job agents may focus on on-time delivery against the lowest possible costs. The main challenges with auction mechanisms are bid generation, bid pricing, bid evaluation, and profit sharing. Bid generation deals with the selection of combinations of items to bid on (see e.g., Lee et al. 2007). Profit sharing deals with the fair distribution of the additional profit generated through the collaboration process among the partners (see e.g., Krajewska and Kopfer 2006; Krajewska et al. 2008; Ozener and Ergun 2008). In this paper we do not to focus on bid generation

and profit sharing, but instead we focus on the impact of the individual bid pricing and bid evaluation strategies on the additional profits that can be achieved through collaboration. Below we elaborate on these issues.

For the bid pricing decisions of the vehicle agents, we may rely on solutions for the dynamic vehicle routing problem (DVRP). Here a number of vehicles has to satisfy transportation requests that arrive dynamically over time. This requires an online planning approach in order to include the new jobs in the vehicle schedules. Although many papers have been devoted to the dynamic vehicle routing problem, there are still some issues that have not been addressed yet (Ghiani et al. 2003). Especially regarding look-ahead policies that incorporate the future consequences of certain decisions. Here we distinguish between two types of look-ahead policies: waiting strategies (i.e., where to wait and for how long) and scheduling strategies in anticipation of future job arrivals. Examples of waiting strategies can be found in Larsen et al. (2004), Ichoua et al. (2006), Thomas (2007). Examples of look-ahead scheduling strategies can be found in Mitrović-Minić and Laporte (2004), Yang et al. (2004), Branke et al. (2005). In this paper we use a combination of look-ahead waiting and scheduling strategies as introduced in Mes et al. (2010).

For the decision making capabilities of the shipper agents our focus is on auction strategies. A commonly used auction protocol in MAS is the sequential Vickrey auction where jobs are allocated one-by-one. The difficulty with such a system is that subsequent jobs are dependent: serving one job is greatly affected by the opportunity to serve another job. To cope with the interdependencies among jobs, we may use reserve prices and/or decommitment penalties. As shown in Myerson (1981), the reserve price increases the expected revenue of the seller by preventing the object from being sold at a low price. For an extensive literature survey on this topic we refer to McAfee and McMillan (1987). Decommitment penalties are introduced in Sandholm and Lesser (2001). Here an agent can decommit (for whatever reason) simply by paying a decommitment fee to the other agent. It is shown, through game-theoretic analysis, that the option to decommit increases the Pareto efficiency of contracts and can make contracts more beneficial for both parties. In Hoen and La Poutré (2004) the decommitment concept is applied to a multi-agent transportation setting. They conclude that significant increases in profit can be achieved when the agents can decommit and postpone the transportation of a load to a more suitable time. In this paper we use a combined policy for the use of reserve prices and decommitment penalties as introduced in Mes et al. (2008).

Closely related work can be found in Berger and Bierwirth (2010), Dai and Chen (2011) where horizontal collaboration between carriers is considered and where each of the carrier agents may act as an auctioneer to outsource transportation requests while the other carrier agents act as bidders. Dai and Chen (2011) explicitly formulates this as a two decision problem, one for the auctioneer and one for the bidder. By using simulation, they evaluate collaborative approach with (i) the individual planning approach (no job exchange) and (ii) a centralized planning approach.

The *main contribution* of this paper is to study the interaction between carriers and shippers, each using look-ahead profit maximizing strategies, which has not been studied before. Most papers focus on models for a single agent type where the behavior of the counterpart is considered as exogenous. Because we include models for both

shipper agents and vehicle agents, we can study the emergent behavior of this combined system. Our study differs from [Dai and Chen \(2011\)](#) in the sense that we focus on the added value of anticipatory behavior of the two agent types. As shown in Sect. 6, we find that the *joint* effect of two policies is *larger* than the effects of the individual policies. Hence, the intelligence of both agent types *strengthens* each other rather than counteract. In addition, we provide a benchmark of the agent-based approach with a centralized mixed-integer programming approach where the combined multi-vehicle pickup and delivery problem is solved to optimality (with respect to the system-wide logistical costs) at each new job arrival. Even though centralization in a multi-actor environment is seldom feasible in practice, it provides an upper bound on the performance that could be achieved by the MAS.

3 Model of the transportation market

Jobs to transport unit loads (full truckload) arrive one-by-one. These jobs are characterized by an origin i , a destination j , a latest pickup time e of the load at the origin, and a time a at which the job becomes known in the network $a \leq e$. To introduce unbalanced transportation networks, i.e., a network in which some areas are more popular than others, we divide the network into disjoint regions a priori. We denote the set of regions by \mathcal{N} .

Within the network, all jobs have to be transported by a fixed set of vehicles that might possibly belong to different collaborating transportation agencies. The overall goal is to minimize the system-wide costs. We consider two cost components, namely the driving costs $c^d(t)$ as function of the travel time t and the penalty costs $c^p(t)$ in case of tardiness t with respect to the latest pickup time e . The time to transport a load from node i to node j is given by τ_{ij}^f (driving full). This includes travel time, and the handling time to load- and unload the job. The time to drive empty from location i to location j is given by τ_{ij}^e .

To model the transportation market we use a multi-agent system. We represent each player by an agent that acts as a decision maker. Here we restrict ourselves to vehicle agents and shipper agents. The shipper agents submit transportation jobs to the market according to some stochastic process. Vehicle agents bid on these jobs and maintain a schedule of the jobs they have won. The vehicle agents and shipper agents have individual objectives. Objective of the shipper agents is to minimize the costs for transportation given by the sum of all prices paid to the vehicles for transporting their loads. Objective of the vehicle agents is to maximize their profits given by the income from all transportation jobs minus the costs for doing these jobs. Yet, the objective of the collaboration is to minimize the system-wide costs as defined before.

To match transportation jobs with vehicle capacity we use auctions. Auction protocols that are specifically designed to deal with complementary goods are, for example, simultaneous auctions (or parallel auctions) where bidders participate in multiple auctions at the same time and combinatorial auctions where bidders may bid on combinations of items. However, combinatorial auctions involve many inherently difficult problems. As mentioned in [Song and Regan \(2005\)](#), we face the bid construction problem, where bidders have to compute bids over different job combinations; and

the winner determination problem, where jobs have to be allocated among a group of bidders. In addition, we face the profit sharing problem (Krajewska and Kopfer 2006) and these procedures might not be directly applicable in situations where jobs arrive at different points in time. In this paper we choose for sequential reverse Vickrey auctions, i.e., for each job we use a single auction round in which the lowest bidder wins the auction at the price of the second-lowest bid. The Vickrey auction has been widely used for MAS because (i) it requires a single bidding round and (ii) it forces bidders—under some mild conditions, see Vickrey (1961)—to bid their true valuation of the object, thereby avoiding many bidding problems (e.g. speculation on profit margins). However, this property no longer holds in sequential auctions where the valuation of bundles of items, acquired in separate auctions, differs from the sum of the valuations of individual items. This certainly is the case in sequential transportation procurement auctions, where bundles form efficient routes consisting of multiple pickup and/or delivery locations. To cope with the interdependencies between jobs we focus on the use of look-ahead strategies for bid pricing and bid evaluation. For clarity of exposition, we make the following assumptions:

- All jobs have to be transported eventually.
- The total transportation capacity is sufficient to handle all jobs in the long run.
- A job in process cannot be interrupted (no preemption); i.e., a vehicle may not temporarily drop a load in order to handle a more profitable job and return later on.
- The handling times and travel times are deterministic.

Further, given our focus on bid pricing and bid evaluation, we take the profit sharing phase for granted and abstract from organizational aspects by ignoring the ownership of vehicles (basically assuming each carrier has one vehicle) and by introducing a single shipper agent that receives and auctions all jobs.

4 Auction and bidding strategies

A job is allocated to a vehicle whenever the shipper decides which vehicle agent will win the auction, if any. After the arrival of new jobs, it may appear that the job assignment is not optimal anymore, i.e., we have a misallocation. Especially when jobs are complementary (e.g. transportation jobs that can be served sequentially by the same vehicle) or substitutable (e.g. transportation jobs that are available at the same time), a certain allocation may become unfavorable when new jobs appear. To improve the allocation of jobs, we take the sequential transportation procurement auction as given, and focus on strategies for the participants. We consider the following options:

1. Delaying commitments: the shipper agent may delay commitments by refusing the current lowest bid based on a reserve price and to start a new auction for the same job later on.
2. Breaking commitments: the vehicle agents are allowed to reject an accepted job in favor of another job. The shipper reconsiders the decommitted job by starting a new auction for this job.
3. Valuation of opportunities: the vehicle agents include opportunity costs in their bids.

For *delaying commitments*, a shipper uses reserve prices in the auctions. When all bids are higher than the reserve price, the shipper rejects them and starts a new auction later on. This way, shippers avoid misallocations by postponing commitments for which they expect to make a better allocation in the future. So if the shipper has plenty of time to auction a certain job, it will not agree with a relatively high bid. When the time for dispatch comes nearer, the price it is willing to accept will rise. We call this a *dynamic threshold policy*.

The idea of *breaking commitments* is that the shipper allows a vehicle to decommit from an agreement against a certain penalty. These penalties are chosen such, that whenever a vehicle decommits a job, they cover the expected extra costs for finding a new vehicle. This way, potential misallocations can be corrected. After a vehicle has decommitted a job, the shipper re-auctions the job in order to find a new vehicle that is willing to do this job. We call this a *decommitment policy*. Note that such a policy is only reasonable in case of collaborative networks, because shippers operating in spot markets certainly would add a risk premium to the decommitment penalties.

In the third option, vehicle agents try to avoid misallocations by not only taking into account the direct impact of doing a certain job, but also its impact on the expected future revenue. This impact on future revenues is captured using the concept of *opportunity costs*. The opportunity costs are affected by job characteristics, such as the destination of the new job, but also by the order and timing of jobs in a schedule. These opportunity costs play a role in the bid pricing decisions of vehicles, but also in their scheduling decisions.

We implement the market-based multi-agent system as follows. When a job arrives at the shipper, it starts an auction by sending an announcement with job requirements to all vehicles. In return, each vehicle calculates a bid considering the marginal costs of doing this job and its impact on future opportunities (Sect. 4.1). Next, the shipper has to decide whether to accept the lowest bid (Sect. 4.2). A shipper may decide to reject all bids and start a new auction later on. Otherwise, the winning vehicle is informed and all vehicles receive information on the lowest bid, which can be extended to cases with less information transparency as shown in Mes et al. (2010).

If the shipper allows decommitment, it also calculates the time-dependent decommitment penalty for the new job and sends this to the winning vehicle (Sect. 4.3). The winning vehicle implements the schedule change. If the winning vehicle decided to decommit from another job, then this decommitment is announced to the shipper, which in turn immediately starts a new auction for this job. After each auction, both the shipper and the vehicles store information of the lowest bid together with the job characteristics. They use this information to periodically update their beliefs about other players (see Sect. 4.1–4.3). A general impression of the situation is given in Fig. 1.

In the next sections (Sect. 4.1–4.3) we describe the three policies in more detail and present some small modifications to adapt these policies to collaborative networks.

4.1 Opportunity valuation policy

In this section we briefly describe the opportunity valuation policy as introduced in Mes et al. (2010). Also, we present a minor modification to this policy to apply it to collaborative networks.




Fig. 1 Transportation procurement market

To support job sequencing decisions and bid pricing decisions, vehicles maintain a job schedule. Vehicles are not restricted by the scheduled pickup times, but can simply decide to insert new jobs or to wait at some location after delivery of a job. The vehicles use an insertion heuristic, see [Campbell and Savelsbergh \(2004\)](#). Here a vehicle contemplates the insertion of a new job at any position in the current schedule without altering the order of execution for the other jobs.

At each point in time, a vehicle v has a job schedule Ψ_v , i.e., a list of jobs with scheduled pickup times. These pickup times are scheduled as early as possible, taking into account the required times for empty moves. In the remainder we denote (i) the number of jobs in a schedule by M , (ii) the destination region of the last job in the schedule Ψ by *schedule destination* $d(\Psi)$, and (iii) the time until the expected arrival time at the schedule destination by *length of a schedule* $l(\Psi)$.

To capture the impact of a schedule on future opportunities, we use an end-value $V(i, t)$ which provides an indication of the attractiveness of a schedule destination i . Specifically, $V(i, t)$ gives the expected profit during a period t after arrival at schedule destination i . The end-values are calculated using Stochastic Dynamic Programming. The information required consists of the job arrival patterns and the distribution of the lowest bid for various job characteristics. This information can be collected from the auctions. For more details on this we refer to [Mes et al. \(2010\)](#).

The end-values are used by the vehicle agents (i) to calculate a bid price for a new job, (ii) to choose an appropriate insertion position for a new job, and (iii) to support so-called *pro-active move decisions*, i.e, moving empty in anticipation of future job requests. Below we elaborate on these decisions.

Consider vehicle v with M jobs in its current schedule Ψ_v . If $M = 0$, there is only one way to schedule a new job. If $M > 0$, a new job can be scheduled in M possible ways, since the first job cannot be interrupted. We write $\Psi_{v\varphi}^m$ for schedule alternative m , where the new job φ is inserted after job m . The direct costs for vehicle v for inserting a new job φ after the m th job in its current schedule are given by (i) the costs for the expected additional travel time $\Delta T_{v\varphi}^m$ and (ii) the expected additional tardiness $\Delta D_{v\varphi}^m$. Besides these direct costs, a vehicle also faces *opportunity costs*. The opportunity costs of schedule alternative $\Psi_{v\varphi}^m$ of vehicle v within a given planning horizon T are given by the difference in end-value of the schedule alternative $\Psi_{v\varphi}^m$ compared to the current schedule Ψ_v . These opportunity costs are given by

$$OC(\Psi_{v\varphi}^m) = V(d(\Psi_v), T - l(\Psi_v)) - V(d(\Psi_{v\varphi}^m), T - l(\Psi_{v\varphi}^m)).$$

The bid price of vehicle v , for inserting a new job φ in its current schedule Ψ_v , is given by the direct costs of the cheapest insertion plus the opportunity costs

$$b(v, \varphi) = \min_{m=1, \dots, M} \left\{ c^d(\Delta T_{v\varphi}^m) + c^p(\Delta D_{v\varphi}^m) + OC(\Psi_{v\varphi}^m) \right\}.$$

We denote the schedule $\Psi_{v\varphi}^m$ with the lowest costs by $\Psi_{v\varphi}^*$. A vehicle agent updates its schedule when (i) an auction for a new job φ is won and (ii) the first loaded move in a schedule has been completed. In the first case, the vehicle agent replaces its current schedule Ψ_v with $\Psi_{v\varphi}^*$. In the second case, the vehicle agent has to decide upon its next move. Here we assume that if the vehicle schedule is not empty, it will drive immediately to the origin of the next job. Otherwise, the vehicle agent has to decide whether to stay or to move pro-actively to another node in anticipation of future demand. For a given node i , the decision to move to node δ will result in an empty move with travel time $\tau_{i\delta}^e$ and costs $c^d(\tau_{i\delta}^e)$. The pro-active move decision is then given by the node δ that maximizes the revenue within the remaining planning horizon $T - \tau_{i\delta}^e$ after arrival at node δ , minus the cost for this empty move

$$\delta(i) = \arg \max_{\delta \in \mathcal{N}} \left\{ -c^d(\tau_{i\delta}^e) + V(\delta, T - \tau_{i\delta}^e) \right\}.$$

Note that more complicated decisions are involved when vehicles not always start the next job as early as possible. Extending the start of the next job might be beneficial in anticipation of new job arrivals (see e.g., [Ichoua et al. 2006](#); [Thomas 2007](#); [Mes et al. 2010](#)).

The opportunity valuation policy has originally been designed for spot markets where we are dealing with a large number of vehicles each applying their own policy. In collaborative networks, all vehicles include opportunity costs in their bid pricing and scheduling decisions. As a consequence, the performance of each individual player is influenced by (i) other vehicle agents charging opportunity costs and (ii) the shipper agent that employs reserve prices or allows decommitment of jobs. When all players use exactly the same end-values the system might become unstable with ever increasing prices. To illustrate this, suppose all players update their end-values at the same time periodically. As mentioned earlier, the end-values describe the expected profit of a vehicle within a certain period depending on its schedule destination. The profit of a vehicle is given by the prices of the jobs it won minus the transportation and penalty costs for serving these jobs. Given the Vickrey auction, the price of a job is given by the second lowest bid which includes opportunity costs. Since the opportunity costs will typically (and at least on average) be greater than zero, the expected profits of the vehicles increase. Because profits increase, the end-values in the next period will also increase. Hence the opportunity costs vehicle charge in their bid prices also increase. As a result, the prices for jobs increase with each periodic update of the end-values.

To prevent the increase in bid prices, we slightly modify the opportunity valuation policy for the use in collaborative networks. The expected rewards are calculated similarly as before by taking the difference between the lowest and second lowest bid. However, because both bids include opportunity costs, we subtract the opportunity costs from the expected rewards.

4.2 Dynamic threshold policy

In this section we briefly describe the dynamic threshold policy as introduced in Mes et al. (2008). Also, we present a minor modification to this policy to apply it to collaborative networks.

By using the dynamic threshold policy, a shipper has the opportunity to auction a job multiple times. We assume that the time between subsequent auction rounds is fixed and equal to R . After each auction, the shipper agent has to decide whether to accept the lowest bid. This decision can be supported by a threshold value $\alpha(n, d, b)$ which is given by the expected price a shipper has to pay in the auction rounds $n + 1, \dots, N$, given that it rejects the current lowest bid b for a job with distance d . We added the current bid b in the state space, because sequential bids for the same jobs are correlated. For R relatively small, the vehicle schedules at the next auction round will not be that different and the same probably holds for the lowest bid.

The optimal policy is to accept the current bid b in auction round n for a job with distance d , only when this value b is below a threshold value $\alpha(n, d, b)$. To calculate the threshold values we introduce a probability density function $P_{n,d}(b)$ of the lowest bid b at auction round n for a job with distance d . Here we discretize the possible bid prices in K classes. We further introduce B_n as the stochastic variable for the lowest bid at auction round n , q^u being the probability that the lowest bid is updated between two auction rounds, and ϕ being the slope of the linear regression between pairs of lowest bids in subsequent auction rounds. We use this slope to include correlations in price deviations (difference between the expected lowest bid and realized lowest bid) in subsequent auctions rounds.

We calculate the threshold values backwards, starting from the last auction round N having a threshold value $\alpha_N = \infty$, i.e., in the last auction round we accept the lowest bid. As in Mes et al. (2008), the recursive relation for the threshold values is given by

$$\begin{aligned} \alpha(n, d, b) = & (1 - q^u) \min \{b, \alpha(n + 1, d, b)\} \\ & + q^u \sum_{k=0}^K P_{n+1,d}(k) \min \{k + \phi [b - \mathbb{E}[B_n]], \\ & \alpha(n + 1, d, k + \phi [b - \mathbb{E}[B_n]])\}. \end{aligned}$$

To calculate the threshold values, the shipper has to learn the values of $P_{n,d}(b)$, q^u , and ϕ . Learning is based on historical observations of the lowest bid, see Mes et al. (2008) for details. The use of historical observations assumes a more or less stable system. This will not be the case in collaborative networks where players adapt themselves to others. As a result, the system might not converge to a stable situation, just like we saw with the opportunity valuation policy (see Sect. 4.1). To see whether the behavior of all players converges to some stable level, and if so, how long this takes, we introduce *learning periods*. The idea of a learning period is that players do not change their behavior during this period. At the end of such a period, players use the observations from this period to update their policies (i.e., recalculate $P_{n,d}(b)$, q^u , ϕ) which they use throughout the next learning period.

Besides using multiple learning periods, we make one additional modification. Because we consider unbalanced networks where some regions are more popular than others, we have to include the origin region i and destination region j in the threshold values $\alpha_{ij}(n, d, b)$. To calculate the threshold values, the shipper estimates the probability density function $P_{n,d}(b)$ using multiple linear regression. Hence, we also have to include the origin and destination region in the regression functions. We simply do this, by adding $|\mathcal{N}| - 1$ indicator functions for both the origin and destination region.

4.3 Decommitment policy

In this section we briefly describe the decommitment policy as introduced in [Mes et al. \(2008\)](#). At the end of this section we present a minor modification to this policy to apply it to collaborative networks.

By using the decommitment policy, the shipper agent allows vehicles to decommit from an agreement (a job) against a predetermined time-dependent penalty. This penalty for a given job, as a function of the remaining time until the latest pickup time, is calculated by the shipper directly at the start of an auction for this job and is announced to the vehicles together with the other job characteristics. Whenever a vehicle decommits, (i) it will not receive the agreed price for the decommitted job, (ii) it has to pay the shipper the time-dependent decommitment penalty, and (iii) the shipper immediately starts a new auction for this job.

The decommitment penalty equals the expected extra costs for a shipper to find a new carrier (so we assume risk neutral shippers). The decommitment penalty $D_{s,t}$ is given by the expected lowest bid at the decommitment time t minus the expected lowest bid at the initial commitment time s : $D_{s,t} = \mathbb{E}[B_t] - \mathbb{E}[B_s]$. However, when the shipper uses the decommitment policy in combination with the dynamic threshold policy, the decommitment penalties $D_{s,t}$ are given by the difference in threshold prices between the initial commitment time s and the decommitment time t . We modify threshold values by letting them depend on the remaining time t instead of on the auction round n . This is a minor modification which can be done rather easily, see [Mes et al. \(2008\)](#). The decommitment penalties are then given by: $D_{s,t} = \alpha(t, d, b) - \alpha(s, d, b)$.

To adjust the decommitment policy to collaborative networks we perform the same modifications as with the dynamic threshold policy, i.e., we include the origin and destination region in the threshold values and we use multiple learning periods.

5 Experimental settings

The goal of this simulation study is to evaluate the impact of combinations of shipper's and vehicles' look-ahead strategies on the system-wide logistical costs. To use the local look-ahead strategies, the agents have to learn the behavior of others. In this study, we want to distinguish the effects of learning from the interrelation of the policies themselves. To do this, each simulation run consist of a learning phase where agents learn from their environment and a steady state phase where agents use the information gathered from the learning phase. During the learning phase, learning takes place periodically (i) by estimation of all required parameters using observations from the

Table 1 Origin probabilities

Degree of balance	Origin probabilities for node/region i ($i = 1, \dots, 4$)
Balanced	$\frac{1}{4}(1 + (i - 1) * 0.0)$
Slightly unbalanced	$\frac{1}{7}(1 + (i - 1) * 0.5)$
Unbalanced	$\frac{1}{10}(1 + (i - 1) * 1.0)$

past period and (ii) by updating the policies in accordance with this. We set the length of a learning period to 10 days, which is sufficient to allow a reasonable amount of observations for various job characteristics. To study the interrelation of the policies, we only consider data from the steady state phase and regard the learning phase as a warm-up period.

We consider a transportation area where locations are distributed within a 100×100 km square area with Euclidean distances. To distinguish between more and less attractive locations, we divide the area into four equal-sized square regions. The regions are numbered consecutively per row, starting in the upper left corner and ending in the lower right corner. To adjust the transportation flow, we set for each region an origin probability, which is the probability that this region becomes the origin of a new job. For a given job, we first draw an origin region using the given origin probabilities and next draw a destination region randomly from the remaining regions. Within a given origin/destination region, we draw an (x, y) coordinate randomly from the square area. The different origin probabilities are shown in Table 1.

We use 10 vehicles, each having a travel speed of 50 km/h. The travel costs and penalty costs are 1 and 10 per minute respectively. The loading- and unloading times are 5 min each. For the dynamic programming recursions on the end-values, we discretize time into periods of 1 min and use a planning horizon T of 12,000 min (≈ 192 times the travel time between two random points in the transportation area). Jobs arrive according to a Poisson process.

For the vehicle agents we consider the following policies:

- MY** Myopic policy: the vehicle agents do not use opportunity costs in their bid pricing and scheduling decisions, i.e., they use the equations from Sect. 4.1 with $OC(\Psi) = 0$ for all schedules Ψ and they do not make pro-active moves.
- OV** Opportunity valuation policy: the vehicle agents do use opportunity costs in their bid pricing and scheduling decisions, and also make pro-active moves.

For the shipper agent we consider the following policies:

- MY** Myopic policy: the shipper agent always accepts the lowest bid and does not allow decommitment.
- DEC** Decommitment policy: the shipper agent allows decommitment of jobs.
- RES** Dynamic threshold policy: the shipper agent uses reserve prices.

For the shipper agent, we decided not to consider the combination of DEC and RES given it results in a relative minor improvement at the expense of a major increase in computation time, see Mes et al. (2008). Given the policies mentioned above, we end up with $2 \times 3 = 6$ possible agent-based control structures (combination of individual policies). We denote a control structure by A/B where A refers to the policy used by all the vehicle agents and B to the policy used by the shipper agent.

A problem with online planning is that we generally compare different heuristics without having any benchmark for the effectiveness in terms of total relevant costs. To have an indication of the quality of our multi-agent approach, we would like to have a reasonable lower bound for the minimum costs. One option is to perform central optimization afterwards when all jobs are known. This is the optimal solution, but we use more information than is available during online execution. Therefore, this lower bound is usually far off the performance of online heuristics; so this is not a realistic bound. The option we consider here is central re-optimization of the problem each time new information arrives. Although it is not guaranteed that we find the minimum costs in this way, it gives a reasonable estimate of the performance that could be achieved based on the information we actually have under central planning. Specifically, we consider a reoptimization policy where the offline multi-vehicle pickup and delivery problem is solved at each new job arrival. Obviously, this policy is not practical for (i) real-time planning purposes of problems of realistic size and (ii) situations in which we are dealing with multiple collaborative transportation agencies that want to maintain a certain level of autonomy. As a benchmark, we use a slightly modified version of the mixed-integer programming formulation given in [Yang et al. \(2004\)](#). In this formulation, the problem is modeled as an assignment problem with timing constraints. The assignment problem consists of finding a least-cost set of cycles describing the order in which each truck should serve the jobs. We slightly modified the formulation in the sense that all jobs have to be carried out and have to be accepted immediately once they are known. We denote the benchmarking policy by BENCH.

As experimental factors we choose the following: control policy, degree of balance, time-window length, and time between jobs. We decided not to vary the number of vehicles, travel speed, handling times, and network size, since their impact on the vehicle capacity (i.e., the average amount of jobs that can be transported per time unit) is already captured by the factor time between jobs. We also fix the travel costs and penalty costs per time unit, but we do provide insight into the different cost components by showing the realized service levels and transportation costs. The ranges of the experimental factors are shown in [Table 2](#). A full factorial experiment with respect to these factors would require $7 \times 3 \times 4 \times 4 = 336$ experiments. For clarity of exposition, and to reduce computation time, we consider (i) all combinations of the factors Control, Degree of balance, and Time-window length; with as fixed settings a time between jobs of 800 s and (ii) all combinations of the factors Policies, Degree of balance, and Time between jobs; with as fixed settings a time-window length of 600 m. As a consequence, we consider $2 \times 7 \times 3 \times 4 = 168$ experiments. In the learning phase we only consider the unbalanced network with a time-window length of 600 m and a time between jobs of 800 s. In the remainder we refer to this setting of time-windows and time between jobs as default configuration.

As primary performance indicator we consider the average costs per job which consists of empty travel costs and penalty costs. The loaded travel costs are excluded because they do not depend on the decisions to be taken. In addition, we consider the *relative savings* of a certain policy which are defined as the relative difference in average costs of this policy compared to the average costs of the myopic policy. In mathematical form this would be

Table 2 Experimental factors

Factor	Values
Control (vehicle/shipper)	MY/MY, MY/DEC, MY/RES OV/MY, OV/DEC, OV/RES, BENCH
Degree of balance	Balanced, slightly unbalanced, unbalanced
Time-window length (min)	300, 400, 500, 600
Time between jobs (s)	700, 800, 900, 1,000

$$\text{relative savings} = 100 \times \left(\frac{\text{average costs of myopic policy} - \text{average costs of policy}}{\text{average costs of myopic policy}} \right).$$

For our simulations, we use a replication/deletion approach, see [Law \(2007\)](#), where each experiment consists of a number of replications (each with different seeds) and a warm-up period. The warm-up period consists of a number of learning periods times the length of a learning period (10 days). The length of each simulation run, excluding the warm-up period, is 100 days. For all experiments, we use 5 replications, which appear to be sufficient for a confidence level of 95% with a relative error of 5% with respect to the average costs per job.

6 Numerical results

First we present the results from the learning period (Sect. 6.1) after which we present the steady state performance of the various policies (Sect. 6.2).

6.1 Learning behavior

Here we evaluate the impact of the number of learning periods (1–9) on the average costs per job, see Fig. 2. Obviously, the policies MY/MY and BENCH do not require learning. The individual policies OV and DEC (policies MY/DEC, OV/MY, and OV/DEC) do not need many learning periods, i.e., one period seems to be enough. The major advantage of this is that they are suitable for nonstationary environments. For the individual policy RES (policies MY/RES and OV/RES) we see that it takes some time to come up with reasonable relative savings; with one learning period we even see that the average costs per job increase compared to the myopic policy MY/MY. For the remainder of this section we use a warm-up period consisting of 5 learning periods. From Fig. 2 we see that this number is sufficient for most policies to converge to a relatively stable performance.

6.2 Steady state comparison of policies

Here we evaluate the interrelation of the shipper's and vehicles' look-ahead strategies. We use the experimental factors as shown in Table 2. All figures in this section display




Fig. 2 Average costs per job as a function of the number of learning periods




Fig. 3 Simulation results for balanced networks

the costs of the agent-based policies relative to the performance of the myopic policy given in percentages. For the unbalanced network we also show performance data with respect to the absolute costs and some additional performance indicators. These data can be found in the “Appendix”.

First we consider the balanced network. Given that the opportunity valuation policy only benefits from imbalance in a network, we omit the policies OV/MY, OV/DEC, and OV/RES. The results for the remaining policies can be found in Fig. 3.

From this figure we draw the following conclusions. First, the relative savings of all policies increase with increasing time-window length. This means that with increasing time-window length, the differences between the myopic policy and the other policies are getting larger. The shippers’ policies RES and DEC benefit from increasing time-windows because there is simply more time to delay (RES) or break (DEC) commitments. The benchmarking policy also takes advantage of increasing time-windows since there will be an increasing probability that the policy will find a better set of




Fig. 4 Simulation results for slightly unbalanced networks

vehicle schedules compared with the myopic policy (see “Appendix”, Table 3 for the unbalanced network).

We further see that the relative savings of MY/DEC and BENCH decrease with increasing time between jobs (decreasing number of jobs). The reason for this is the following. The advantage of MY/DEC and BENCH over the myopic policy is that they allow exchange of jobs between vehicle schedules (by means of swapping jobs or completely reassigning all jobs). However, with increasing time between jobs, the average schedule length of the vehicle will become shorter. So, there will be less to gain by exchanging jobs. The relative savings of MY/RES increase with increasing time between jobs since the probability of late delivery decreases, which is really an issue with this policy (see “Appendix”, Table 4 for the unbalanced network). Also, with increasing time between jobs, the probability of finding better vehicle schedules in future auction rounds (the principle behind MY/RES) will increase.

A final observation here is that the gap between the agent-based policies and our benchmark remains relatively large. We come back to this issue at the end of this section.

Next we consider the case of slightly unbalanced networks, see Fig. 4. The travel distances (empty as well as loaded) are getting longer with increasing imbalance. As a consequence, it will be harder to deliver all jobs on time. This has a similar effect as the decreasing time-between jobs in the balanced case. This explains why within the slightly unbalanced networks the relative savings of the MY/DEC and BENCH are higher most of the time whereas the savings for the policy MY/RES are lower in most cases. We further see that performances of OV/MY and MY/DEC are close to each other. Finally, the combination of shipper’s and vehicles’ strategies (OV/RES and OV/DEC) always increases the performance. The best combination of local policies here is OV/DEC.

Next we consider the case of unbalanced networks. The results can be found in Fig. 5. By introducing more imbalance, it becomes even harder to deliver all jobs on time. Using a similar argumentation as given above, the relative savings of MY/DEC and BENCH in the unbalanced network become slightly higher and the relative savings of MY/RES slightly lower. However, the differences between the unbalanced and the slightly unbalanced case are smaller than the differences between the slightly unbalanced and balanced case. The reason for this is that within the unbalanced network,




Fig. 5 Simulation results for unbalanced networks

the majority of transport takes place within one region. This region can be regarded as balanced because the origin and destination coordinates are drawn randomly within this region.

We further see that the gap between the best agent-based policy and our benchmark becomes smaller; especially with increasing time between jobs. Again we see that from the local policies, the combination OV/DEC performs best in most cases. Only in the situation with relative low job arrival intensity (time between jobs of 1,000), the policy OV/RES performs best.

Finally, we observe a peculiarity: the relative savings of our benchmark decreases with increasing time-window length, something which was not the case in the balanced and slightly unbalanced case. The reason for this is that with increasing imbalance, more capacity will be lost on empty movements. As a result, it becomes relatively more difficult for the myopic policy MY/MY, compared with the policy BENCH, to achieve a high service level with short time-windows. Therefore, the policy MY/MY will benefit relatively more from an increase in time-windows. Hence, the relative savings for BENCH, compared to those of MY/MY, decreases as can be seen from Table 3 in the “Appendix”.

To summarize the results, we have seen that combinations of vehicle and shipper strategies always improve the performance compared to one of the individual policies. In almost all cases, the combination of the opportunity valuation policy and the decommitment policy (OV/DEC) works best. Only in settings with long time-windows or few jobs, the combination of the opportunity valuation policy and the dynamic threshold policy comes in favor. In almost all cases the relative savings of these policies lie between 10 and 20%.

The results of the local policies seem promising, but there is still a gap with our benchmark. For example, in the unbalanced networks, the policy OV/DEC (the best combination of local policies) is only able to achieve, on average, 53% of the savings from our benchmark. There are two extenuating circumstances here. First, the problem under consideration is relative simple and clean in the sense that only one type of decision is involved (assigning jobs to certain positions in truck schedules) and only one type of uncertainty is involved (the job arrival process). In earlier work, see Mes et al. (2007), where we considered a problem involving many more decision types and also uncertainty in handling and travel times, we drew an opposite conclusion

in favor of the local policies. Second, the benchmarking policy requires considerably more computation time compared to the local policies, and these computation times will explode with increasing problem size. For this reason we only considered relatively small problem sizes. With more realistic problem sizes, the multi-agent approach would still be able to perform real-time whereas the central approach would require approximations or has to be replaced by heuristic procedures. Some details regarding the computation times can be found in the “Appendix”.

7 Conclusions

In this paper we studied the interaction between vehicle agents and shipper agents in a market-based multi-agent system for full truckload transportation. Shipper agents offer the transport jobs through sequential auctions. A set of vehicle agents compete with each other to serve these jobs. For the shipper agent we considered two auction strategies, namely a dynamic threshold policy and a decommitment policy. For the vehicle agents we considered opportunity valuation policies where not only the direct costs of jobs are taken into account, but also the impact on future opportunities. We used simulation to evaluate the benefits of the different strategies and to study their interrelation. Our main conclusions are the following:

- The combination of vehicle and shipper strategies performs better than the individual policies. On average we observe a reduction of 10–20% in the costs for tardiness and repositioning of the vehicles.
- The combination of the opportunity valuation policy and the decommitment policy works best in almost all cases and requires relatively limited computation time. The combination of the opportunity valuation policy with the dynamic threshold policy comes in favor in settings with long time-windows or fewer jobs.
- The performance of the individual policies depends a lot on the network structure and job characteristics. The opportunity valuation policies of the vehicles benefit from the imbalance in the network where some regions are more popular than others. These policies are therefore especially suitable for unbalanced networks. The dynamic threshold policy and decommitment policy of the shipper benefit from fluctuations in bid prices due to the possibilities of combining jobs. The decommitment policy is especially suitable for balanced networks. The dynamic threshold policy is especially suitable for settings with long time-windows or fewer jobs.
- There is still a gap between the agent-based policies and our benchmarking policy which reoptimizes the multi-vehicle pickup and delivery problem at each new job arrival. For example, in the unbalanced network, the control OV/DEC achieves on average 53% of the savings from our benchmark. However, the benchmarking policy might simply not always be applicable due to its computational complexity and because it ignores the autonomy of the different actors. To use the benchmarking policy, we only considered small problem instances in this paper; larger instances certainly would require approximations or other solution methodologies. Furthermore, the agent-based approach might come in favor with increasing uncertainty as shown in [Mes et al. \(2007\)](#).

The gap between the agent-based policies and our benchmark gives rise to further research. Specifically we focus on two issues. First, the improvement of the local policies by using approximate dynamic programming where we try to learn the value functions without using a detailed model of the environment's dynamics. Second, the integration of the concepts opportunity costs, threshold values, and decommitment penalties, in a mathematical programming approach that is used on a central level.

Appendix

In this section we show additional performance data with respect to the unbalanced network. As performance indicators we consider the average costs per job (Costs), the percentage of the total driving distance that is driven loaded (DL), and the service level (SL) defined by the percentage of jobs that are delivered on time. The results for varying time-window length (TW) can be found in Table 3 and the results for varying time between jobs (TBJ) in Table 4.

For our experiments we used the simulation software Plant Simulation 8.2 and an Intel Pentium 4 processor at 3.4 GHz. To speed up the simulations, we programmed the dynamic threshold policy in Delphi 7 as a dynamic link library which we included in our simulation environment. We solved the mixed-integer programming formulation using CPLEX 11. The average computation times per job for the decentralized policies under the default configuration range from 0.019 s in case of MY/MY and 3.974 s in

Table 3 Simulation results for varying TW for unbalanced networks

TW	300			400			500			600		
	Costs	DL	SL	Costs	DL	SL	Costs	DL	SL	Costs	DL	SL
MY/MY	37.5	66.5	95.8	32.4	67.0	97.4	31.3	67.5	97.3	30.6	67.8	97.9
MY/DEC	33.3	67.2	97.9	29.2	68.2	98.8	28.4	68.7	99.0	27.8	69.1	99.2
MY/RES	42.0	67.0	93.2	34.5	68.0	95.9	32.9	68.7	94.7	31.8	69.1	94.7
OV/MY	30.6	68.1	97.5	28.5	68.9	98.3	28.0	69.3	98.7	27.4	69.8	98.7
OV/DEC	30.0	68.5	98.5	27.6	69.5	99.1	26.4	70.3	99.4	26.0	70.6	99.4
OV/RES	35.8	68.5	94.7	29.8	69.7	96.3	28.1	70.4	96.4	26.8	70.7	96.7
BENCH	24.7	71.8	99.6	22.6	73.2	99.7	22.0	74.0	99.8	21.5	74.5	99.9

Table 4 Simulation results for varying TBJ for unbalanced networks

TBJ	700			800			900			1,000		
	Costs	DL	SL	Costs	DL	SL	Costs	DL	SL	Costs	DL	SL
MY/MY	32.8	67.7	96.5	30.6	67.8	97.9	30.6	67.7	98.4	30.5	67.7	98.0
MY/DEC	30.8	68.4	97.5	27.8	69.1	99.2	27.6	69.2	99.3	28.1	68.7	99.4
MY/RES	40.0	68.1	93.7	31.8	69.1	94.7	28.9	69.3	95.9	28.0	69.2	96.7
OV/MY	28.9	69.2	97.5	27.4	69.8	98.7	27.2	69.9	98.9	27.3	69.8	99.2
OV/DEC	27.5	70.1	98.3	26.0	70.6	99.4	26.1	70.4	99.6	26.4	70.1	99.7
OV/RES	29.7	70.2	96.8	26.8	70.7	96.7	26.0	70.6	97.2	25.9	70.5	97.9
BENCH	22.8	73.3	99.7	21.5	74.5	99.9	22.0	74.1	99.9	22.5	73.5	99.9

case of OV/RES. The policy RES requires relatively more computation time because the shipper has to calculate the dynamic threshold recursion at each auction whereas with OV and DEC some computations can be done offline. The computation time for OV/DEC, the best agent-based policy on average, is 0.249 s and for BENCH this is 22.048 s. Keeping in mind that these differences in computation times will increase with increasing problem size, the achievement of the agent-based policy of typically 50% of the savings from the benchmark policy can be regarded as impressive.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Berger M, Schröder M (2011) Multicriteria decentralized decision making in logistic chains: a dynamic programming approach for collaborative forwarding of air cargo freight. *Logist Res* 3(2-3):121–132
- Berger S, Bierwirth C (2010) Solutions to the request reassignment problem in collaborative carrier networks. *Transp Res Part E Logist Transp Rev* 46(5):627–638
- Böhnlein D, Schweiger K, Tuma A (2010) Multi-agent-based transport planning in the newspaper industry. *International Journal of Production Economics*
- Branke J, Middendorf M, Noeth G, Dessouky M (2005) Waiting strategies for dynamic vehicle routing. *Transp Sci* 39(3):298–312
- Campbell AM, Savelsbergh M (2004) Efficient insertion heuristics for vehicle routing and scheduling problems. *Transp Sci* 38:369–378
- Chen B, Cheng HH (2010) A review of the applications of agent technology in traffic and transportation systems. *IEEE Trans on Intell Transp Syst* 11(2):485–497
- Cruijssen F, Salomon M (2004) Empirical study: order sharing between transportation companies may result in cost reductions between 5 to 15 percent, Technical report, CentER discussion paper no. 2004-80
- 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(3):620–626
- Dai B, Chen H (2011) A multi-agent and auction-based framework and approach for carrier collaboration. *Logist Res* 3(2):101–120
- Dudek G, Stadler H (2005) Negotiation-based collaborative planning between supply chains partners. *Eur J Oper Res* 163(3):668–687
- Ergun O, Kuyzu G, Savelsbergh M (2007a) Reducing truckload transportation costs through collaboration. *Transp Sci* 41(2):206–221
- Ergun O, Kuyzu G, Savelsbergh M (2007b) Shipper collaboration. *Comput Oper Res* 34(6):1551–1560
- Figliozzi MA, Mahmassani HS, Jaillet P (2003) Framework for study of carrier strategies in auction-based transportation marketplace. *Transp Res Rec* 1854:162–170
- Fischer K, Muller JP, Pischel M (1996) Cooperative transportation scheduling: an application domain for DAI. *J Appl Artif Intell. Special issue on Intelligent Agents* 10(1):1–33
- Ghiani G, Guerriero F, Laporte G, Musmanno R (2003) Real time vehicle routing: solution concepts, algorithms and parallel computing strategies. *Eur J Oper Res* 151(1):1–11
- Gomber P, Schmidt C, Weinhardt C (1997) Elektronische märkte für die dezentrale transportplanung. *Wirtschaftsinformatik* 39(2):137–145
- Hoen PJ 't, La Poutré JA (2004) A decommitment strategy in a competitive multi-agent transportation setting. In: Faratin P, Parkes DC, Rodriguez-Aguilar JA, Walsh WE (eds) *Agent mediated electronic commerce V (AMEC-V)*, vol 3048 of lecture notes in artificial intelligence LNAI. Springer, Berlin, pp 56–72
- Hülsmann M, Kopfer H, Cordes P, Bloos M (2009) Collaborative transportation planning in complex adaptive logistics systems: a complexity science-based analysis of decision-making problems of groupage systems. In: Zhou J (ed) *Complex sciences*, vol 4 of lecture notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Springer, Berlin, pp 1160–1166

- Ichoua S, Gendreau M, Potvin JY (2006) Exploiting knowledge about future demands for real-time vehicle dispatching. *Transp Sci* 40(2):211–225
- Kopfer H, Pankratz G (1999) Das groupage-problem kooperierender verkehrsträger. In: Kall P, Lüthi HJ (eds) *Operations research proceedings 1998*. Springer, Berlin, pp 453–462
- Krajewska MA, Kopfer H (2006) Collaborating freight forwarding enterprises. *OR Spectrum* 28(3): 301–317
- Krajewska MA, Kopfer H, Laporte G, Ropke S, Zaccour G (2008) Horizontal cooperation among freight carriers: request allocation and profit sharing. *J Oper Res Soc* 59(11):1483–1491
- Larsen A, Madsen OBG, Solomon MM (2004) The a priori dynamic traveling salesman problem with time windows. *Transp Sci* 38(4):459–472
- Law AM (2007) *Simulation modeling and analysis*, 4th edn. McGraw-Hill Education, New York
- Lee C-G, Kwon RH, Ma Z (2007) A carrier's optimal bid generation problem in combinatorial auctions for transportation procurement. *Transp Res Part E Logist Transp Rev* 43(2):173–191
- Liu R, Jiang Z, Fung RYK, Chen F, Liu X (2010) Two-phase heuristic algorithms for full truckloads multi-depot capacitated vehicle routing problem in carrier collaboration. *Comput Oper Res* 37(5):950–959
- McAfee RP, McMillan J (1987) Auctions and bidding. *J Econ Lit* 25(2):699–738
- Mes MRK, van der Heijden MC, Schuur PC (2008) Dynamic threshold policy for delaying and breaking commitments in transportation auctions. *Transp Res Part C* 17(2):208–223
- Mes MRK, van der Heijden MC, Schuur PC (2010) Look-ahead strategies for dynamic pickup and delivery problems. *OR Spectrum* 32(2):395–421
- Mes MRK, van der Heijden MC, van Harten A (2007) Comparison of agent-based scheduling to look-ahead heuristics for real-time transportation problems. *Eur J Oper Res* 181(1):59–75
- Mes MRK, van der Heijden MC, van Hillegersberg J (2008) Design choices for agent-based control of AGVs in the dough making process. *Decis Support Syst* 44(4):983–999
- Mitrović-Minić S, Laporte G (2004) Waiting strategies for the dynamic pickup and delivery problem with timewindows. *Transp Res Part B* 38(7):635–655
- Myerson RB (1981) Optimal auction design. *Math Oper Res* 6(1):58–73
- Ozener O, Ergun O (2008) Allocating costs in a collaborative transportation procurement network. *Transp Sci* 42(2):146–165
- Puettmann C, Stadler H (2010) A collaborative planning approach for intermodal freight transportation. *OR Spectrum* 32(3):809–830
- Sandholm T (1993) An implementation of the contract net protocol based on marginal cost calculations. In: *Proceedings of the eleventh national conference on Artificial intelligence, AAAI'93*, pp 256–262
- Sandholm T, Lesser V (2001) Leveled commitment contracts and strategic breach. *Games Econ Behav* 35(1-2):212–270
- Schönberger J (2005) *Operational freight carrier planning: basic concepts, optimization models and advanced memetic algorithms*. GOR-publications, Springer, Berlin
- Schwind M, Gujo O, Vykoukal J (2009) A combinatorial intra-enterprise exchange for logistics services. *Inf Syst E-Bus Manag* 7(4):447–471
- Song J, Regan AC (2005) Approximation algorithms for the bid construction problem in combinatorial auctions for the procurement of freight transportation contracts. *Transp Res Part B* 39(10):914–933
- Stadler H (2009) A framework for collaborative planning and state-of-the-art. *OR Spectrum* 31(1):5–30
- Thomas BW (2007) Waiting strategies for anticipating service requests from known customer locations. *Transp Sci* 41(3):319–331
- Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J Finance* 16(1):8–37
- Wooldridge MJ (1999) *Intelligent agents*. In: Weiss G (ed) *Multiagent Systems*. The MIT Press, Cambridge pp 27–77
- Yang J, Jaillet P, Mahmassani HS (2004) Real-time multivehicle truckload pickup and delivery problems. *Transp Sci* 38(2):135–148