



A multi-agent platform for auction-based allocation of loads in transportation logistics

Valentin Robu^{a,*}, Han Noot^a, Han La Poutré^a, Willem-Jan van Schijndel^b

^a CWI, National Research Center for Mathematics and Computer Science, Amsterdam, The Netherlands

^b VOS Logistics Organizing, Nijmegen, The Netherlands

ARTICLE INFO

Keywords:

Transportation logistics
Multi-agent systems
Supply chain management
Software platform
Auction platforms
Human-computer interaction

ABSTRACT

This paper describes an agent-based platform for the allocation of loads in distributed transportation logistics, developed as a collaboration between CWI, Dutch National Center for Mathematics and Computer Science, Amsterdam and Vos Logistics Organizing, Nijmegen, The Netherlands.

The platform follows a real business scenario proposed by Vos, and it involves a set of agents bidding for transportation loads to be distributed from a central depot in the Netherlands to different locations across Germany. The platform supports both human agents (i.e. transportation planners), who can bid through specialized planning and bidding interfaces, as well as automated, software agents. We exemplify how the proposed platform can be used to test both the bidding behaviour of human logistics planners, as well as the performance of automated auction bidding strategies, developed for such settings.

The paper first introduces the business problem setting and then describes the architecture and main characteristics of our auction platform. We conclude with a preliminary discussion of our experience from a human bidding experiment, involving Vos planners competing for orders both against each other and against some (simple) automated strategies.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

Decentralized markets and distributed auctions have received a lot of research interest, as important coordination mechanisms between self-interested agents. Recent research, both by the authors of this paper and many others (Reeves, Wellman, MacKie-Mason, & Osepayshvili, 2005; Robu & La Poutré, 2008; 't Hoen & La Poutré, 2006; 't Hoen, Robu, & La Poutré, 2005) etc. proposes increasingly complex algorithms for bidding and modeling decisions in such agent-mediated auction environments.

Transportation logistics and supply chain management represent challenging and potentially very fruitful areas for the application of agent-based electronic market techniques, such as auctions. The increasing complexity and shifting structure of modern supply chains, as well as increasing competitive pressures in this market has led to an increasing demand and interest in such distributed optimization techniques, involving multiple parties. The practical impact of improved allocation which can be achieved through such techniques can be significant. For example, in the Netherlands, the average transport performance is between 40% and 60%. Improving this utilization rate is also the goal of the DEAL (Distributed Engine

for Advanced Logistics) project, which groups together several universities and large logistics service providers in the Netherlands. The work reported here was also carried out in the framework of this project, involving two of the main partners, namely CWI, Amsterdam and Vos Logistics Organizing, Nijmegen, the Netherlands.

1.1. The multi-party logistics domain

Several trends have recently produced a significant impact on the area of transportation logistics. One of these is an increase in competition, with the continual entry of new carriers in the market pushing down expected profit margins. Another one is the increasing complexity and sophistication of modern supply chains. In fact, due to increasing and shifting trade patterns, not only transportation chains have become more dynamic, but also their structure has become increasingly complex.

For example, nowadays it is no longer the case that the company that accepts a transportation order also owns the actual capacity (i.e. trucks) to carry it. Often, multinational companies with large, regular amounts of cargo to be delivered prefer to outsource these orders to other companies that undertake to find convenient delivery options, within a set of pre-negotiated terms. These intermediary logistic companies then negotiate the distribution of these orders with other smaller companies who have the

* Corresponding author.

E-mail addresses: vr2@ecs.soton.ac.uk (V. Robu), han@cwi.nl (H. Noot), hl@cwi.nl (H. La Poutré).

actual transportation capacity (which own the actual trucks and hire the drivers). This can be actually a cheaper option in many cases, as smaller transportation companies often do not have the complex cost structure that larger companies have (Robu et al., 2008; van Amstel & Goor, 2003; van der Putten, Robu, La Poutré, Jorritsma, & Gal, 2006).

In standard transportation management literature (van Amstel & Goor, 2003) such distributed supply chains are called multi-party logistics. Existing literature (van Amstel & Goor, 2003) identifies several classes of logistic provider companies, based on the type of services they offer. Although there is some disagreement about the exact usage of the terms, in our approach (and the remainder of this paper) we use the term 3PL company (third-party logistics providers) to denote those that have their own transport capacity (i.e. truck fleet) and plan this own capacity and 4PL company (i.e. fourth-party logistics provider) to denote those companies which “orchestrate” the supply chain, i.e. acquire large sets of orders from large shippers and then re-distribute these orders among a set of other companies with actual transport capacity.

1.2. Company profile

Founded in 1944 as a one-truck company, transporting loads between Oss and Nijmegen in The Netherlands, Vos Logistics has grown into one of the larger logistics service providers in Europe. It has over 3000 trucks, 10000 trailers and containers, 325 storage silos and 2 rail service centers. Vos employs 5000 people working at more than 45 locations throughout Europe, while annual turn-over approaches 1 billion euro.¹ The increasing complexity of transportation chains has induced Vos Logistics to offer new solutions to its large corporate customers (shippers), which can now outsource all of their transportation activities to Vos. This lets them avoid the problem of finding and negotiating with individual suppliers, billing, following up orders etc. Another advantage of using this outsourcing service for large shippers is that Vos Logistics has a much better knowledge of the transportation market, so it is better positioned to find suitable sub-contractors. Vos Logistics Organizing from Nijmegen (henceforth abbreviated VLO in this paper) is a subsidiary of Vos Logistics B.V. that was set up in order to handle such complex supply chain orchestration activities. Based on the taxonomy above, VLO (the subsidiary) can be seen as a 4PL company, though its parent company, Vos Logistics was founded as a 3PL company and does have its own trucks. Hence, VLO acts as an intermediary company that acquires large (sets of) orders from suppliers and negotiates the allocation of the orders, the terms of transportation (i.e. delivery deadlines, destination) as well as the price at which other carrier companies subcontract these orders.

1.3. Automating multi-party logistics using agents

The focus of this work is on automating, through an agent system the second part of the market interaction, i.e. the daily outsourcing of transportation orders to carrier companies who will actually transport them. The first part, which is actually acquiring these orders from large shippers presents less opportunities for automation through a multi-agent system. The reason is that these contracts are usually fewer, larger and closed over a longer time horizon (e.g. a company based in the US may delegate to Vos Logistics Organizing the delivery of the goods imported into Europe over a period of one year). Such large, complex type of decisions cannot yet be expected to be delegated to software agents.

However, allocation of orders on a daily basis to different 3PL carriers was identified as an area with clear potential to benefit

from more automated techniques (our previous AAMAS'06 survey paper (van der Putten et al., 2006) examined this potential). This automation would involve decision support systems for human planners in the first stage, and next some of the decisions could be delegated to software agents.

A final note about how the allocation occurs in current practice: in the Vos case, negotiation over most orders occurs in a small group of companies who are invited to submit bids for different orders as they arrive in the system. In some cases in which no reasonably priced offer is made, Vos may also solicit other outside companies and carriers to submit a bid (this includes multimodal options, such as rail or water transportation carriers). However, these cases are mostly exceptions (they account for less than 20% of the total value of the orders (van der Putten et al., 2006)), so most business is conducted in a group of (up to) 10 companies that can submit bids for a given set of orders. This is the case we are interested in automating through the auction platform presented in this paper.

1.4. Goals of this work

Over the years, several successful agent-based auction platforms have been developed in order to allow comparison and evaluation of automated trading strategies. The Trading Agent Competition is, perhaps, the most well-known example of this (see (Wellman, Greenwald, & Stone, 2007) for an overview) – most related to this work being its supply chain version (Sadeh, Eriksson, Finne, & Janson, 2003). These platforms are, however, simply not suitable for our basic goal, which is to convince the Vos Logistics Organizing management (and their partner carrier companies) that agent-mediated electronic auctions can actually be used in practice to automate their daily outsourcing of transportation orders. For this purpose, a custom-based platform was required, modeled around a business case which the planners that actually perform these operations daily can easily recognize and use.

Since the final system is to be used by logistics planners, such a system should closely resemble a real-world case, and allow its users to identify the bidding and planning decisions to be taken in this platform as decisions they would usually also take in real life. It should have an interactive, intuitive interface and, moreover, it should seamlessly integrate human agents who take planning and bidding decisions with automated agents implementing an algorithmic strategy or heuristic. This point is especially important for acceptance, since during operational adoption of such a system, it is not realistic to expect that a company would immediately delegate all market decisions to a piece of software, without being confident that such decisions closely model those their human planners would make. To summarize, the goals of this project (and corresponding platform) are:

- The overall goal of the project is to demonstrate the feasibility of applying such an auction system in the day-to-day transportation outsourcing activities of Vos Logistics Organizing (VLO), Nijmegen.
- As a more detailed goal, the platform should allow us to illustrate how different mechanism choices, such as allowing flexible pick-up/delivery times or decommitment ('t Hoen et al., 2005) (with or without a penalty) can improve efficiency and participant profits.

From an AI or agent researcher's point of view, the developed system can also form a platform to test different aspects of distributed decision making in logistics auctions, more specifically:

¹ These figures are valid for the 2006/2007 financial year.

- Testing increasingly complex automated trading strategies. At this stage, some very simple strategies have been developed, whose role is mostly to stabilize the market, to make it more realistic. However, more intelligent strategies for this setting can be easily added to the existing platform.
- The demonstrator can also be seen as a platform for analyzing and testing the behaviour of human planners taking part in such an auction.

We wish to emphasize that this paper is not concerned with proving that any particular bidding strategy, mechanism or scheduling method is superior to others. The readers can consult work by (some of) the authors of this paper which presents and evaluates such strategies, at a more abstract level, in 't Hoen et al. (2005), Robu and La Poutré (2008), Mous, Robu, and La Poutré (2008) and 't Hoen and La Poutré (2006). Rather, our goal in this project is to build an environment which directly models current business practice in transportation logistics (more specifically, a real business case provided by Vos Logistics Organizing, Nijmegen) and in which different analytically-developed strategies can be adapted and tested.

The rest of this paper is organized as follows. Section 2 provides a high-level overview of our platform and the business case on which it is based. Section 3 describes in more detail the auctioneer agent, as well as the auction protocol used. Section 4 describes the functionality and behaviour of the automated agents that are currently part of the proposed platform, while Section 5 describes the human agent interface and functionality. Section 5 also introduces the cost structure that was used for the agents and the planning assistance interface that was built to assist human planners in taking bidding decisions. Section 6 presents some (very preliminary) results and impressions from a study conducted at Vos Logistics, involving 6 human planners bidding against each other and against our agents, while Section 7 concludes the paper with a discussion.

2. Overview of the business case and our platform

The demonstration takes its starting point in a real-world case of how transportation loads from a depot in the south of the Netherlands can be distributed across Germany. In order to preserve the privacy of Vos Logistics Organizing, as well as their customers and business partners, some parts of the model are purposely left unspecified or details have been slightly changed, without really affecting how realistic our model is. This especially holds for the names of the customer companies and some specific details about the data used. The main reason for this is that our platform is intended for evaluation not only by planners employed by Vos Logistics, but also by those of some partner companies. The main parts of the problem setting can be summarized as:

- All orders used in the demonstration will be fictive (i.e. randomly generated, not real orders), but, in order to assure the platform is realistic, their destination postcodes, weights, times of delivery etc. are based on real-world distributions.
- All outgoing orders are assumed to be delivered starting from a depot near Maastricht (a town in the south of the Netherlands), while possible return freight (i.e. pick-up) orders appear at destinations across Germany.
- There are n players playing the role of the carriers (this can vary, we estimated that in our setting it will be around 10) and one player in the role of VLO (i.e. the auctioneer).
- Each carrier has k trucks to plan (in our demonstration, in order to allow the players to follow all the details simultaneously, we agreed k could be relatively small, e.g. $k = 5 \dots 10$). Each truck has a standard capacity of 26 pallets, where pallets are all assumed to have a standard weight of 1000 kg/pallet.

2.1. Generating transportation orders

A data set of about 4000 orders was supplied by Vos Logistics, corresponding to orders for a period of time from a real case. These real orders never actually appear in our simulated platform, since that might violate confidentiality agreements between VLO and the shipper company. However, the orders actually appearing in our platform very closely resemble real orders, as follows.

The German destination (or origin) postcode for each order, which is a two-digit number, was generated as follows. The first digit (corresponding to the broad geographical region), was generated at random using the probability distributions extracted from real data. The weight of the order (expressed in the number of pallets from 1...26), was also generated at random, again from a distribution extracted from the data. In general, some order weights are much more common than others and, furthermore, this also varies by delivery region: some regions receive larger cargo orders, while for some smaller, more frequent orders are the norm.² Therefore, the distribution for generating the weight is also dependent on the delivery region (corresponding to the first digit of the postcode). Finally, the second digit of the postcode (which corresponds to a specific town within this general postcode region) was generated at random, but 50% of the weight was given to the 2–3 most important second digits for the area (usually corresponding to a larger town or population center).

In order to have a closed loop demonstration, we assume that the carriers also have return orders available. The return orders are, conceptually, offered by sellers from different areas – although in our demo they will be sold through the same auction mechanism. Outgoing and return orders have asymmetric distributions (60% of all orders are outgoing and only 40% are return orders). This is also realistic for this business scenario, given available data. In real life there are two types of orders: “ON” orders (which must be delivered exactly on their target delivery date) and “BY” orders (which are to be delivered by a certain deadline date, where early delivery is allowed). To simplify the setting, and also allow more flexibility in planning, all orders in our simulated platform will be considered “BY” orders.

Another very important parameter in such a platform is the lead-time of an order, which, roughly defined, represents the difference in days between the time when an order is to be delivered (i.e. the delivery time or deadline) and the time when the order actually appears in the platform (is put up for auction). Here, we also follow a pattern extracted from the real data, as described in the following.

Each order is assigned a random lead-time, produced using a series of adapted, lognormal distributions. The peak of these lognormals will be the first acceptable lead-time day for the order (see Fig. 1 for an illustration). Because such distributions have a long tail, this means that orders to be delivered 3, 4 days or even a week after the minimum lead-time can still appear, albeit with exponentially decreasing probability. For example, most orders to be generated with a minimum lead-time of 1 are to be delivered in the next two days.

The reasons why we need several lognormal distributions is that different types of orders have different lead-times (we identified 3 categories, according to the order data supplied). Thus, orders that are to be delivered to postcode regions in the west of Germany (places closed to the Dutch border) and whose delivery and return trip can be completed within the same day have, in general, shorter lead-times than orders that require a minimum of two days travel (including the return trip).

² While we cannot give the full details, a statistically weak, but still significant correlation coefficient of $\eta = 0.4$ was found between the delivery area postcode and the size of an order.

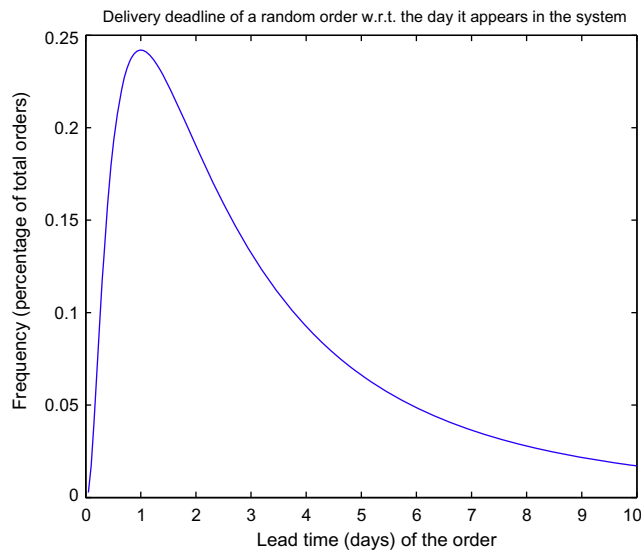


Fig. 1. Example distribution of delivery deadlines for orders, in number of days from the present time, for a load with minimum lead-time of 1. The date when an order appears in the system corresponds to the origin.

2.2. Computing prices and costs

As one would expect in any auction platform, the final price for each order will be determined by the bidding in the open market. However, in an interactive demonstration, we had to build in a mechanism to assure that prices for the orders quickly converge to actual prices (in euro) that human planners would expect to see. Fortunately, also in current practice there is a mechanism to assure this. There is a partner company of Vos Logistics (the name of which, again, we cannot give for privacy reasons), that can transport orders to any destination in Germany. They do provide a standard price scheme which quotes a delivery price for any combination of order size (in number of pallets) and German post-code region. It is very important to stress that these are maximal prices: in general VLO expects to get (and usually gets) much better delivery prices from their closed group negotiation with the partner carriers, otherwise it would be unable to make a profit. The services of this company are only considered if Vos fails to attract a realistic bid for an order from any of the carriers in their closed group (which can sometimes happen, though rather seldom).

However, having such a set of prices is useful in our system, because it provides a benchmark of what kind of prices are realistic. The way we use this information is in designing the bidding strategy of our automated agents, whose bidding strategy will depend on this standard prices (an exact description of the functionality of these automated agents is provided in Section 4). The point of these agents, in this version of the software, is not to beat the human planners, but to assure that the competition bids they see (and implicitly, the bids they have to submit to beat them), are around actual market prices they would encounter in real life. Henceforth in this paper, we will refer to this set of prices as the standard industry price table.

Finally, a word should be said about cost data. We have also obtained and incorporated in the cost structure of the bidders, detailed information tables about the exact driving times and distances to any postcode location in Germany, as well as realistic estimations of the fixed costs (e.g. driver salaries, truck maintenance) and variable costs per km (including driving tax and fuel costs). These were incorporated in the cost structures of the bidders when planning their routes (a thorough description is provided in Section 5).

3. Auction protocol and design of the auctioneer agent

This section describes the main characteristics of the auction protocols used, as well as other characteristics of the auctioneer agent. To allow more planning flexibility, but also to follow current tendering practices, orders with different lead-times are auctioned with slightly different auction protocols, as described below.

3.1. Auction set-up

Loads are auctioned sequentially (or in 3–5 small batches distributed throughout the day). This resembles current transportation practice. Often, loads are offered by different shippers, who have different deadlines throughout the day for placing their orders.

For the current set-up, all auctions are descending English auctions, but adapted to better fit the actual tendering process, as it is currently performed. There are two main types of auctions, differentiated by their closing protocol.

3.2. Auctions for loads with a short lead-time

This protocol (more similar to ascending English auctions³), is applied to orders with delivery deadlines which are 1 or 2 days away from the current time. The auction is incrementally descending (lowest offer wins). After the last offer has been placed, the other bidders are given at least 1/2 h to respond with a new offer, after which the auction closes and the lowest bidder so far is awarded the order. Of course, in our simulated environment 1/2 h is replaced by 30 s to 1 min. The actual delay to be used (in number of seconds) can be specified by the user controlling the auctioneer agent, through an interface. Therefore, our auctions have a “soft” closing time (deadline), i.e. they are extended for a short time after the last bid is received, in order to allow other bidders the chance to respond to it.

3.3. Auctions for orders with a longer time horizon

For orders with delivery deadlines over 3 days into the future, the simplified protocol cannot be applied, since most bidders do not plan so far in advance. Additionally, some flexibility must be added in the simulations, in order for us to observe the benefits of allowing time window relaxation and adding a penalty for delayed orders (we return to discuss this issue in Section 5.2).

Therefore, for such orders we use the following decision procedure. For each order, we set a reservation threshold (visible or invisible to the bidders themselves), which gives a reasonable market cost of the order which a shipper would accept in order to have a commitment (without waiting until the last moment to go through the auctions). In our demonstration, the threshold could be set as a percentage below the standard industry price table (as described above) for this configuration of load and destination postcode.

When the order appears in the system, all bidders are informed and can make offers. If a carrier makes a bid that is higher than the reservation price (i.e. not acceptable), then the offer is rejected, the carrier is informed of this and can bid again. A rejected offer (above the reservation price) is thus non-binding to either party, i.e. no commitment exists. If any carrier makes a bid that is below the reservation threshold, and thus acceptable, then all carriers are informed and the auction is moved to the “usual” auction queue (i.e. sold through the auction protocol described in Section 3.2).

³ To be more precise, this extending deadline protocol resembles the most the protocol used by the e-commerce site Amazon.com.

This means, bidders will have sufficient time to respond after the first offer is made, otherwise the contract is awarded to the initial bidder. If, by 2 days before the deadline, no carrier made a bid in the “acceptable” range (i.e. below the reservation price), then the load is still auctioned using the “usual” procedure, described in Section 3.2.

This protocol ensures that bidders that wish to plan in advance are given the chance to do so, but only if they make a reasonable offer, where by “reasonable” we mean considerably below the price that could be expected to be achieved by waiting closer to the actual deadline. An optional alternative, that could be of interest here, is to allow the human playing the VLO side to change the acceptable reservation threshold during the game, if time passes and an order does not appear to attract enough attention and thus risks remaining undelivered.

Finally, as a future research idea, the reservation threshold could be made dynamic (i.e. automatically increasing), according to a discount function. This function would balance the shipper's desire of getting a better price for his delivery and the risk of not getting his load delivered in time, as the deadline approaches. This is relatively easy to implement in the current demonstration tool but, at least for the moment, we prefer to focus on usability studies using the simpler setting.

3.4. Total capacity of loads to be generated per day

A problem that arises in designing such an agent trading platform is to choose the total capacity of orders which should be generated per day. This choice is an important one, because it gives the player an impression of how “competitive” the whole scenario feels. We propose an estimation model for the total capacity of the market that depends on several parameters:

- n – number of participants representing carriers
- k – number of trucks/participant (our case, e.g. $k = 5$)
- $p = 26$ – number of standard pallets/truck
- s – a coefficient representing the “saturation” of the market. This is an important parameter, which allows us to control the market balance between demand (i.e. coming from outstanding orders) and available supply of transportation capacity.

A rough heuristic evaluation of the total capacity of the simulated market we consider will be given by:

$$s * n * k * p$$

Thus, orders will be generated at random using the above distributions, until the total capacity reaches the above value (after choosing the saturation parameter s). This will necessarily be only a rough estimation: because orders are at random and there are time window constraints, there is no real way to know what is the true capacity of the market – unless we would centrally compute, in advance, the best possible plan for the day for all available trucks. This is not really feasible and it is also not required, because in practice not all capacity of the trucks of a carrier company is allocated in the “closed group” auction. In practice, trucks taking part in such an auction may also acquire loads elsewhere – and they only fill up using the current auction. Furthermore, there could be some differentiation between the capacities of different players.

In order to account for this, we could make the following choice: of the total estimated market capacity, we consider that $i\%$ is filled from other sources (“ i ” stands for the initial fill percentage). Thus, an estimated $s * n * k * p * (1 - \frac{i}{100})$ in total capacity will be filled through the auctions, and $s * n * k * p * \frac{i}{100}$ will be pre-filled, through a heuristic, before the auction starts.

3.5. Auctioneer user interface

A screen shot of the auctioneer interface is presented in Fig. 2. As is readily apparent from this figure, all interfaces were designed in the Dutch language, to ease understanding in a logistics business environment, which is our target audience. However, we believe this paper should provide sufficient description to allow a non-Dutch speaker to understand our platform's functionality and capabilities.

Basically, both the order generation and awarding of orders (i.e. auction closing process) executed by the auctioneer platform can be run in two possible ways:

- **Automated control:** In automatic order generation, the user only specifies the parameters of the generation process (as described above) and the arrival rates of orders in the platform. In automatic tendering mode, the auctioneer waits a number of seconds after the last received bid (which the user specifies through the interface) before making the decision to award the order. This actually varies based on the order lead-time, as described in Section 3.2. Orders with longer lead-times, which remain open for bids until a few days before the delivery deadline, are temporarily shown in a different list and are moved to the “active bidding” queue two days before expected delivery.
- **Human control:** In our interface, a human auctioneer (representing the 4PL company, in this case Vos Logistics Organizing) can make, change or correct any of the decisions taken by the system (either order generation or tendering of orders). We found this is a very useful feature in any live, interactive simulation with several human planners, who first are required to get used to the interface etc. This lets the human auctioneer feel firmly in control of the process, even if he chooses to let the software agent take some of the decisions on his/her behalf.

The switch between these modes can be performed dynamically (and online), by simply checking/unchecking a multi-option box.

4. Automated bidders: description and user interface

The role of the automated bidding agents is to ensure the stability of the market and that prices in the demonstrator converge to a realistic level. Therefore, it is enough in a first implementation, if the automated agents use a simple, myopic bidding strategy. The bids are simply based on a standard industry price table (c.f. Section 2.2), which gives a rate for each combination of load/delivery region.

Since this is an English auction, there are two levels, which are randomly determined for each bidding agent: the level of the initial bid and the reservation level (i.e. the lowest the agent will go with his/her bids). Both are generated at random from normal distributions, which are centered at certain levels above and below those taken from our industry price table, as supplied by Vos Logistics. The parameters to be set for automated strategies are:

- Percentage of mean mark-up of the initial bid over the industry price table (and the corresponding dispersion).
- Percentage of the reservation price vs. standard industry price table, for that postcode region and weight (again, this is the mean of the distribution, and a dispersion is also chosen).
- Concession speed (giving how fast the agent's bids go down from his initial price to the reservation price, i.e. frequency of bidding).
- Number of automated bidders and percentage of orders the automated agents bid on. This gives the pressure that independent bidders apply on the market.

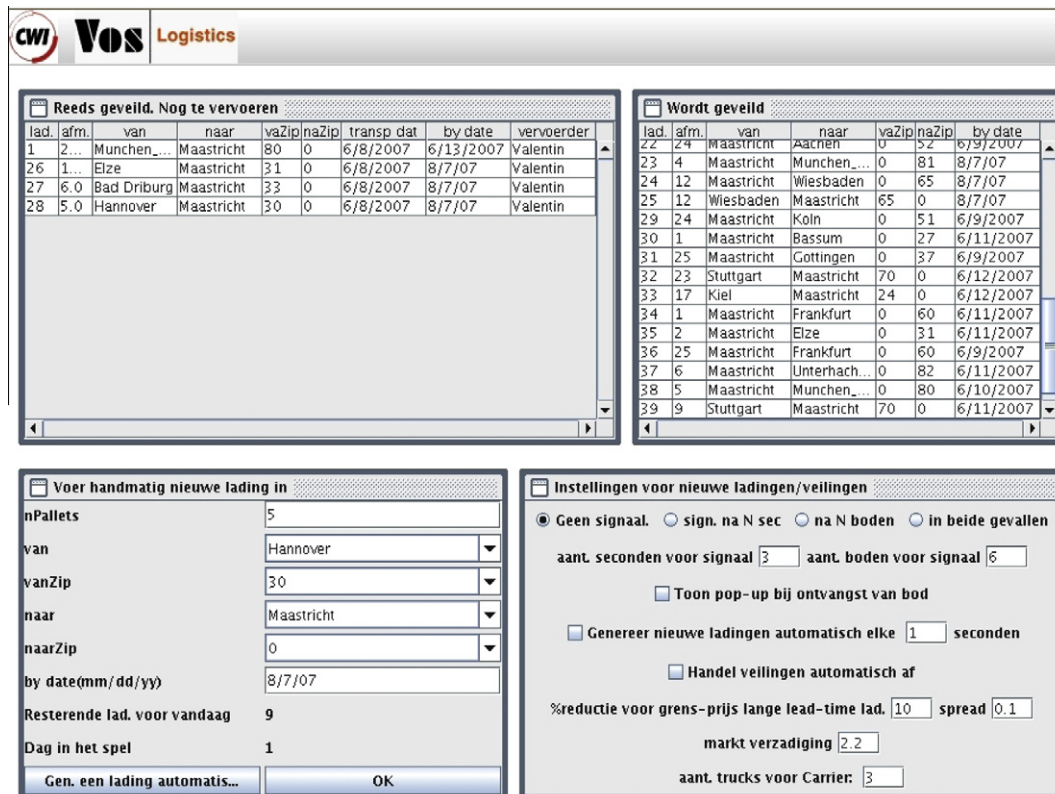


Fig. 2. Part of the interface of the auctioneer, showing 4 windows (in clockwise order): 1. List of transportation orders already auctioned, but still to be transported; 2. List of transportation orders currently being auctioned; 3. Control panel for setting game and simulation parameters (such as automatic or manual allocation of orders); 4. Manual control panel for load parameters.

5. The carrier agents: description and user interfaces

This section aims to give a technical description of the problem faced by the human carriers in our model and the interface available to them in the demonstrator. More precisely, two distinct interface windows are available to human carriers:

- One for visualizing ongoing auctions for loads and bidding
- One for planning assistance, in which human planners are given a stylized impression of their transportation capacity (i.e. trucks) and can visualize and automatically determine the way acquired orders (or orders the user currently considers bidding on) fit in their already planned routes, as well as the expected costs.

The interface for visualizing the list of auctions and placing of bids is not illustrated here, as it is similar (in appearance, not necessarily in functionality), to the auctioneer's interface screen from Fig. 2. Figs. 3–5 all show different situations for the planning support interface, which is the most important for the planners, since it helps them decide how to bid and which kind of transportation order combinations would be profitable for them.

5.1. Transportation model and carrier costs

The transportation planning problem, is, in itself, a complex one to solve. The bidding decisions which the carrier takes are directly dependent on the way the carrier can fit the loads he is bidding on in his/her already existing plans (i.e. how well he/she can form profitable bundles of loads during planning). In turn, this depends on the cost model. Our tool does provide planning assistance, by

computing the costs for each combination of loads considered. In our model, costs of each carrier are of two types:

- Fixed costs, per day and per truck. These are expressed as a fixed amount (in euro).
- Variable costs: all these costs are assumed to be proportional to the distance travelled. These are expressed as a cost in euro per kilometer travelled.

Both of these are set to a realistic level, after discussions with Vos Logistics. The distances within Germany, as well as from Maastricht to/from destination postcodes in Germany are computed based on a supplied distance table. This distance table contains, for each pair of first two digits of German postcodes, a distance in km, as well as a distance in km from any German postcode to/from Maastricht.

Our planning tool enables the carrier to visualize how filled the trucks are at each time point, the time windows in which loads can be delivered, as well as any violation of constraints. There are several types of constraints that need to be met in transportation settings. First there are obvious capacity constraints: a truck cannot be filled at any one time with more than 26 pallets. Second, there is a strict legal constraint about the maximal driving time any driver can actually drive per day – in the EU, this is fixed at 9 h. Any driving plan has to satisfy these constraints to be feasible.

The tool also provides decision support (see below), by computing the length of the route for the partial daily plan – and, thus, the costs incurred so far, for each possible bid the human planner chooses to make. The length of the route is computed (given the distance table available), through a simple *insertion heuristic*. Insertion heuristics are known to provide a very good approximation of the optimum in small settings – and are known to be

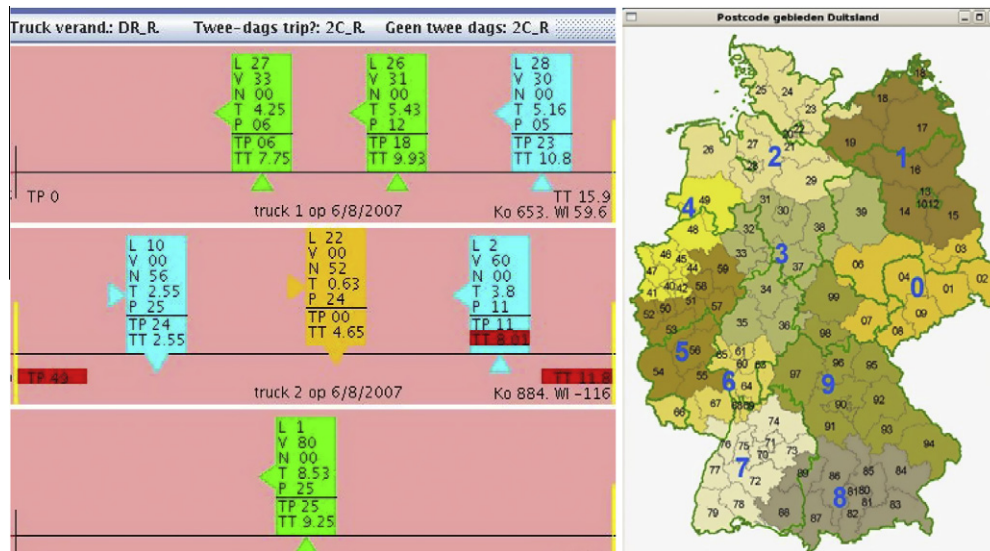


Fig. 3. Basic layout of the planning support window, showing 3 trucks and one transportation day horizon. Each line represents a truck, and each colored container a load (see below for a description of the symbols on each load). For each day, the costs (Ko), profits (Wi) and total traveling times (TT) are computed by the system. For the second truck, the interface shows that inserting an additional load (shown in brownish yellow) in the driving schedule would lead to a violation of the legal driving hours constraints for this truck. The right-hand side shows a window with the map of Germany (shown here reduced), which helps human planners in exploring load combinations. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

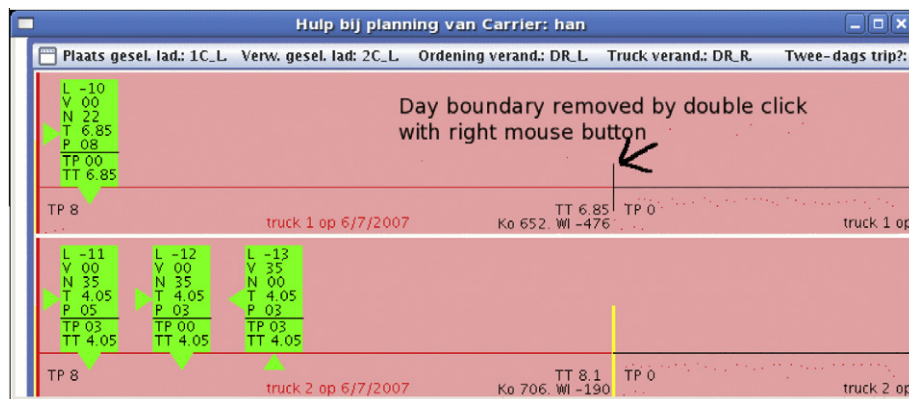


Fig. 4. Layout of the planning support window, showing a set of [green] loads, which have already been contracted. As before, for each day, the costs (Ko), profits (Wi) and total traveling times (TT) are computed by the system. Vertical yellow lines represent day boundaries, which can be removed by the user to allow for multi-day planning. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

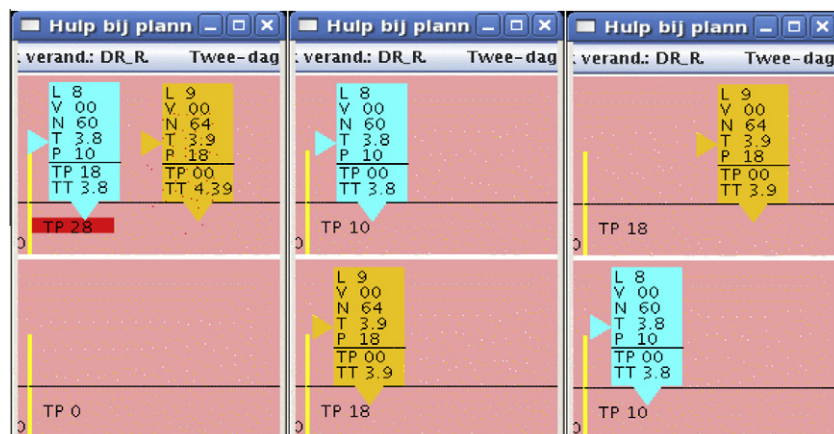


Fig. 5. Left: A number of pallets constraint violation (maximum admitted, 26 pallets/truck), and two possible solutions (center and right), with loads being moved to different days.

computationally more efficient than solving the TSP problem with a more advanced method. Thus, at each point, the expected profit the agent can make so far can also be computed.

5.2. Penalty for late deliveries

An issue of relative importance in actual applications is what happens if delivery is (slightly) late, compared to the agreed date.⁴ In real life, this does happen to a very small percentage of accepted orders, because profit margins in transportation logistics are tight and carriers have to try to make use of all possible bundling options. Given the business of the underlying customer company, we have decided not to treat slight delays as a strict, inviolable constraint, but to allow orders to be maximum one day late, against payment of a penalty. There are two ways to model the penalty in our system:

- Fixed costs/day of delay (e.g. 50–100 euro for each day the truck is late).
- Proportional, as a percentage of the total value of the transportation order.

In our setting, we currently implement a fixed penalty/day of delay – as opposed to a penalty which is proportional to the value of the order. This is a realistic model, since any delay can be seen as a loss in the reputation of the carrier, regardless of the size or actual value of the order. It is up to the bidding carrier if he chooses to incur this penalty in his planning, but in the current set-up only exceptionally profitable planning configurations would justify the chosen level of penalty for an order. Future versions of the system could consider allowing for differentiated bidding, based on the exact date when the order is delivered (an option discussed in van der Putten et al. (2006)).

5.3. Information supplied about other carriers during the competition

An important point to be discussed is what kind of information should be available to human bidders (carriers) in the tool, regarding the activity of the other bidding carriers. This represents a trade-off decision, since on one hand we need to model real life and not compromise the privacy of competing parties, on the other hand in a dynamic simulation environment, agents can be expected to have a reasonable idea about their competition. The following choices have been made:

- Regarding other bids made on existing orders (which the agent is also interested in), the agent should be able to visualize the amounts of the competing bids for the loads he/she is also interested in, but not the identity of the other bidders. Otherwise said, he can see how far he needs to lower his prices to win, but not where the competition for the orders is coming from.
- At the end of each day, a “leader board” is displayed, giving the gross profits rates so far, for all human carriers in the game. We recognize this information about the competition may not be known in real life, but it may be important in an interactive, game-like simulation scenario for the participants to have a signal of how well they are doing, by comparison to their competition. Also, only knowing the profit margins does not reveal much (if anything) about the bidding strategy and underlying planning of the competing carriers.

5.4. Planning and bidding decision support interface

The software developed for human carrier agents has two distinct interfaces: the bidding and the planning support interfaces. In this paper we only illustrate (in Figs. 3–5) some of the features of the planning support interface, as the bidding interface contains relatively straightforward lists of orders one which one can place bids.

The planning interface (see Figs. 3–5) consists of several horizontal lines, one per each truck that the carrier owns. All trips are assumed to be return trips to/from a depot in Maastricht, for any postcode address in Germany. These trips can be one-day trips, for short-distance orders or two-day trips, for destinations further away (the choice is made by simply clicking a yellow vertical bar).

The interface is a drag-and-drop one, which makes it intuitive and very easy to use. Loads are marked in the system by colored rectangular shapes, marked by two arrows. The side arrows represent pick-up, respectively drop-off points, within the schedule of that day. Each load is marked with: its load no (L), the 2-digit German postcodes of the source (V) and destination (T), number of pallets (P) and time it takes to transport this load (T). The total number of pallets and total traveling time are shown below a black line. Constraint violations will automatically be highlighted in red.

Load symbols can have 3 possible colours:

- **Green:** Loads which have been already acquired (and awarded to the carrier) in auction and which need to be planned for transportation.
- **Light blue:** Loads for which a bid has been placed (thus the carrier agent is bound by the bid he made, since bids are binding), but which have not been won yet by the carrier at the price he offered.
- **Yellow-brown:** Loads which are only placed for tentative planning to see if the planning constraints (total driving time, number of pallets etc.) can still be satisfied given already acquired loads, as well as an estimate of expected profits.

For each truck timeline and day, the system automatically computes the total driving time and the number of pallets loaded and automatically signals (by highlighting in red) if any constraints are being violated. The most useful feature for deciding the minimum bid level is, however, the online computation of the potential profit and loss to be made by inserting a load in the current route. This is basically the difference between the current bid made for the load and the cost of the extra travel detour for delivery/picking up that load. Empty scheduled already start with a negative profit associated to them, equalling the fixed costs per day and truck.

For loads that have not been bid on yet, but are tentatively dragged and dropped into the schedule, the information about changes in pricing provides very useful information about what is the minimum bid that can be placed if the carrier decides to acquire that load.

6. Industrial adoption case study

A large-scale, realistic demonstration of the platform involving 6 experienced Vos transportation planners was performed at Vos Logistics. In this test, planners were asked to bid against each other and against our software agents for loads, and their strategies as well as the profit they made with the acquired loads was recorded. After this demonstration, several conclusions were highlighted regarding the performance of our automated auction system:

- First, the bidding and planning support interfaces were considered very helpful and realistic by all the planners involved. Some participants even claimed they were superior to the planning system currently being used in everyday planning.

⁴ As already discussed in Section 2, early deliveries are allowed, since we consider all our orders “BY” type of orders.

- The presence of automated bidding agents (although they currently only bid based on a randomly perturbed set of industry prices), is crucial for the stability of the market and the convergence of prices to realistic levels.
- The profit levels in the simulation do, very roughly, commensurate with the skill of the bidder. However, in order to ensure that the profit rates actually match current practice, the pricing scheme and other system parameters require some further refinement.
- The planning scenarios considered in the simulation could be expanded to consider some other situations appearing in real life (e.g. multiple one-day return trips).
- Other, more advanced functionality could be built into the platform, such as support for combinatorial bidding (Sandholm, 2002) or allowing the possibility of decommitment for loads already acquired (a possibility analytically studied by us in 't Hoen et al. (2005)).

Overall, the planners and managers present were quite impressed with the faithfulness to reality of our platform, and it was agreed that a larger test will be conducted, as well as more concrete steps will be taken towards operational use of such techniques.

7. Discussion

Transportation logistics represents an important application area for multi-agent systems, due to its inherently distributed and dynamic nature. Several approaches have been presented in recent years to this problem, some leading to commercially successful, operational systems. The LS/AT system, presented in Dorer and Calisti (2005) is one of the most well-known systems that uses agent techniques (mostly constraint-reasoning type techniques) for dynamic transport optimization. The Magenta system (Skobeelev, Glaschenko, Grachev, & Inozemstev, 2007) is another such system, which explores the use of swarm-based optimization techniques in this setting.

By contrast to these systems, the emphasis in our approach is not directly on optimization of the planning (though that remains, of course, the final goal), but on automating the market interaction between several companies in a multi-party logistics negotiation. Our approach can be seen as creating a testbed, in which each company or carrier can then apply its own optimization and bidding techniques, the performance of these techniques can then be easily measured and compared.

The approach we take is most similar to the work which proposes different trading platforms to test different aspects of bidding and decision making in electronic markets. There are many such platforms proposed in multi-agent literature, the most well known being the Trading Agent Competition (TAC); the most similar TAC to our approach is, probably, the supply chain TAC version (Sadeh et al., 2003). Of course, our platform may not have all the sophisticated features of the TAC platforms, but unlike TAC, the starting point of our work was in the applicability of the market setting to a real business case, rather than scientific curiosity or relevance. To the best of our knowledge, it is the first paper to describe an agent-mediated auction platform that is modeled around a real life business scenario, where the orders characteristics, costs, profit margins etc. all resemble those encountered in real life.

Another important aspect of our platform is the ability to integrate human bidders and automated trading strategies in the same platform. We feel this is crucial for real business adoption of agent-mediated electronic market techniques because, at least for some of the interacting parties, the human owners will want to remain in control, before delegating any financial decision (e.g. bidding) to a software agent. In multi-agent literature there are some games

specifically developed to test human decision making in negotiation and auctions (a good example is the Colored Trails game (Grosz et al.)), but again our platform has the advantage of allowing us to assess such decisions in a real business environment. Finally, somewhat related to our approach is work on designing stock market trading platforms to test automated bidding strategies (of which PLAT (Kearns & Ortiz, 2003) is a well-known example). While this line of work also uses real financial order data to design a realistic market, the characteristics of stock markets (i.e. double auction setting) is very different from the transportation business case we consider.

We conclude that, overall, our platform did achieve the scope it was built for: to convince Vos Logistics Organizing that the agent-based approach is a valid solution for their business problem. Nevertheless, there are still many aspects open for further research. The first would be to conduct a (set of) larger scale experiments to get more detailed human bidding data, and to develop better techniques to analyze this data. The second is to adapt some of the bidding strategies developed analytically in our more theoretical lines of work (Robu & La Poutré, 2008; 't Hoen & La Poutré, 2006; 't Hoen et al., 2005), and test their performance in this environment, both against other strategies and against human planners.

Acknowledgement

The research for this paper was done in the framework of the SenterNovem Distributed Engine for Advanced Logistics (DEAL) project.

References

- Dorer, K. & Calisti, M. (2005). An adaptive solution to dynamic transport optimization. In *Proceedings of AAMAS'05 industry track, Utrecht* (pp. 45–51).
- Grosz, B., Kraus, S., Talman, S., Stossel, & B., Halvin, M. (xxxx) The influence of social dependencies on decision making: Investigations with a new game. In *AAMAS'04* (pp. 782–791).
- Kearns, M., & Ortiz, L. (2003). The Penn-Lehman automated trading project. *IEEE Intelligent Systems* (November–December).
- Mous, L., Robu, V., & La Poutré, H. (2008). Using priced options to solve the exposure problem in sequential auctions. In *Proceedings of AMEC'08 Workshop, Estoril, Portugal* [Post-proceedings to appear in Springer LNBI].
- Reeves, D. M., Wellman, M. P., MacKie-Mason, J. K., & Osepashvili, A. (2005). Exploring bidding strategies for market based scheduling. *Decision Support Systems*, 39, 67–85.
- Robu, V., & La Poutré, H. (2008). Designing bidding strategies in sequential auctions for risk averse agents. In *Proceedings of AMEC'07 workshop on agent-mediated electronic commerce and trading agent design and analysis, Honolulu, Hawai'i*. Springer lecture notes in business information processing (LNBI) (Vol. 13, pp. 76–89).
- Robu, V., Noot, H. La Poutré, H., & van Schijndel, W. J. (2008). An interactive platform for auction-based allocation of loads in transportation logistics. In *Proceedings of the seventh international joint conference on autonomous agent and multi-agent systems (AAMAS'08), Industry Track, Estoril* (pp. 3–10).
- Sadeh, N., Eriksson, J., Finne, N., & Janson, S. (2003). TAC-03: A supply-chain trading competition. *AI Magazine*, 24(1), 92–94.
- Sandholm, T. (2002). Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1–2), 1–54.
- Skobeelev, P., Glaschenko, A., Grachev, I., & Inozemstev, S. (2007). MAGENTA Technology: Case studies of Magenta i-Scheduler for Road Transportation. In *Proceedings AAMAS'07 (Industry Track), Honolulu, USA* (pp. 1385–1392).
- 't Hoen, P. J., & La Poutré, J. A. (2006). Repeated auctions with complementarities. In *Proceedings of AMEC'05, Utrecht, 2005*. Springer lecture notes in artificial intelligence (Vol. 3937, pp. 16–29).
- 't Hoen, P. J., Robu, V., & La Poutré, J. A. (2005). Decommittment in a competitive multi-agent transportation setting. *Agent-based applications, platforms & development kits*. *Whitestein Series*. Birkhauser, Springer.
- van Amstel, P. W., & Goor, A. R. (2003). *Van logistiek naar supply chain management*. Kluwer Publishers.
- van der Putten, S., Robu, V., La Poutré, J. A., Jorritsma, A., & Gal, M. (2006). Automating supply chain negotiations using autonomous agents: A case study in transportation logistics. In *Proceeding of fifth international joint conference on autonomous agent and multi-agent systems (AAMAS 2006), Hakodate, Japan* (pp. 1506–1513).
- Wellman, M. P., Greenwald, A., & Stone, P. (2007). *Autonomous bidding agents: Strategies and lessons from the trading agent competition*. MIT Press.