

Applying revenue management to agent-based transportation planning

Albert Douma, Peter Schuur, Matthieu van der Heijden

University of Twente

May 19, 2006

Abstract

We consider a multi-company, less-than-truckload, dynamic VRP based on the concept of multi-agent systems. We focus on the intelligence of one vehicle agent and especially on its bidding strategy. We address the problem how to price loads that are offered in real-time such that available capacity is used in the most profitable way taking into account possible future revenues. We develop methods to price loads dynamically based on revenue management concepts.

We consider a one leg problem, i.e., a vehicle travels from i to j and can wait at most τ time units in which it can get additional loads from i to j . We develop a DP to price loads given a certain amount of remaining capacity and an expected number of auctions in the time-to-go. Because a DP might be impractical if parameters change frequently and bids has to be determined in real-time, we derived two approximations to speed up calculations. The performance of these approximations are compared with the performance of the DP.

Besides we introduce a new measure to calculate the average vehicle utilisation in consolidated shipments. This measure can be calculated based on a limited amount of data and gives an indication of the efficiency of schedules and the performance of vehicles.

1 Introduction

We consider a network where loads are transported by multiple companies having a fleet of vehicles. Every vehicle can carry more than one load at a time. A vehicle is considered as a nomadic driver which means that the vehicle itself has no time restrictions and no place (home base) where it has to return after some time. Preemption and time window violation are not allowed.

An important question we will discuss is how to price loads in case of freight consolidation when loads are offered in real-time and companies have to determine their bid immediately after they

receive an announcement from the auctioneer. This is a hard question since a company is not sure how many and what kind of loads will be offered in the near future and how those loads can be consolidated. Moreover, the probability of winning a load depends on the bids of competitors.

In fact companies face the problem that not only the fixed costs, but also the variable transportation costs have to be assigned to a load. The number of loads a company can get is not known in advance and therefore it is difficult to decide which part of the costs is to be assigned to which load. In fact the question can be rephrased as how to value the remaining capacity before and after acceptance of a load.

We address the problem how to determine the price of a load when the variable transportation costs can generally be shared between loads and it is not sure how many loads a vehicle will get on a route. The paper is composed as follows. In the second section, we consider relevant literature and we explain the contribution of our study to the field. In Section 3, we describe our model, we explain why we apply revenue management concepts and we provide a utilization measure. In Section 4 we develop a methodology to price additional loads on a single arc and we give some numerical results in Section 5. In the final section we present our conclusions and propose directions for further research.

2 Related literature

The problem we study is related to several topics in the literature. We can first classify our problem to be in the area of Vehicle Routing Problems (VRPs) but there is also a relation with agent theories and economic theories.

2.1 Transportation planning

The deterministic version of our problem is described in the literature in the area of Vehicle Routing Problems (VRPs), and is generally known as the multi vehicle pickup and delivery problem with time windows. We consider the case in which vehicles can simultaneously pick up and deliver goods to customers as long as a vehicle has sufficient capacity. It is well-known that most variants of the VRP are \mathcal{NP} -hard (Garey and Johnson (1979)).

A lot of research has been devoted to VRPs in the last thirty to forty years (see Toth and Vigo (2002) for a survey). The focus of these studies was mainly on the static and deterministic cases of vehicle routing in which all information is known at the time of the planning of the routes. Little research has been done into real-time transportation planning in which information is both

dynamic and stochastic. In this paper we consider a dynamic multi agent-based VRP (MAB-VRP) in which loads are dynamically allocated to different (possibly competing) companies. We found only a few contributions in this field. We mention some examples like Weerd (2003) who approaches the problem from an Artificial Intelligence perspective providing a framework and methods for multi-agent planning. Two other interesting contributions come from Figliozzi, Mahmassani and Jaillet (2004) and Hoen and Poutré (2004). Figliozzi, Mahmassani and Jaillet (2004) present among others a framework to study carriers' strategies in an auction marketplace for dynamic, full truck load vehicle routing with time windows. However, they do not propose a pricing mechanism for pricing in less-than-truckload vehicle routing problems. The contribution of Hoen and Poutré (2004) is about a multi-company real-time vehicle routing problem with consolidation where vehicle agents can break a contract after they have accepted a load. However, this contribution does not pay much attention to the determination of a bid. It simply assumes that a bid is equal to the additional costs to transport a load, not taking into account opportunity income of future loads. The present paper focusses on how to determine a bid in case of freight consolidation such that a company maximizes profits.

2.2 Agent technology and agent-based planning

For many years, most companies focused their attention on optimization of the business processes within their organization. Over the last decade, companies have increasingly realized the strategic importance of the linkages among supply chain activities, since these could lead to competitive advantage (see Porter (1985)). As a result, some companies started to integrate and coordinate the intricate network of business relationships among supply chain members. However, these networks are often open and lowly structured. Parunak (1998) states that multi-agent systems are well-suited to applications that are modular, decentralized, changeable, ill-structured and complex.

Agent technology first emerged in the mid to late 1980's in the field of Artificial Intelligence. Since then, a lot of interest has been given to the agent concept by a variety of disciplines in computer science. According to Wooldridge and Jennings (1995), an agent is usually defined as a hardware or (more often) software based object with as key properties autonomy, social ability, reactivity and pro-activeness. This means that an agent can operate without direct intervention of humans or others (autonomy), can interact with other agents (social ability), can perceive its environment and respond to changes in it (reactiveness) and is able to exhibit goal-directed behavior by *taking the initiative* (pro-activeness). A system with more than one agent is called a multi-agent system (MAS). The basic idea behind the application of a MAS is that through the interaction of local agents, coherent global behavior is achieved.

Research into this has increased in the field of logistics and operations research over the last

years. Several researchers have attempted to apply agent technology to manufacturing systems (see, e.g., Arbib and Rossi (2000); Dewan and Joshi (2002)), supply chain management (see, e.g., Ertogral and Wu (2000); Qinghe, Kumar and Shuang (2001)) and transportation networks (see, e.g., Figliozzi, Mahmassani and Jaillet (2004), Hoen and Poutré (2004) and Mes, van der Heijden and van Harten (2006)). Publications in the latter area are, however, very limited. We mention a contribution from Mes, van der Heijden and van Harten (2006) in particular, because they implemented a dynamic vehicle routing problem with time windows and stochastic travel and handling times, and compared it with traditional OR heuristics (look-ahead rules, serial scheduling). They showed that an agent-based transportation planning and scheduling system can indeed result in coherent global behavior and that such an approach can result in high performance in terms of vehicle utilization and service levels.

2.3 Economic theory

A major source of inspiration for our research comes from free-market economics where systems are strongly influenced by pricing mechanisms. The ability of markets to efficiently allocate resources over (groups of) opportunistic individuals, is especially and directly relevant to our research. Much of the early work on market based systems was collected and published by Clearwater (1996). Markets can be well modelled in multi-agent systems, since they are naturally decentralized, flexible, robust and allow a natural decomposition.

To automatically assign resources to agents based on a market approach, it is necessary that agents i) use a common protocol by which they are able to come to an agreement, and ii) are willing to fix their agreement in a formal or “legally binding” contract (Brandt, Brauer and Weiß (2000)). The protocol concerns the process of finding an appropriate task assignment and the contract concerns the consequences and commitments resulting from the assignment on which agents agreed. Two standard types of task assignment protocols are negotiation-based protocols and auction-based protocols. We choose for an auction protocol, which has several distinct and advantageous features compared to negotiation-based protocols (see Brandt, Brauer and Weiß (2000)). Moreover, according to Sandholm (1996), auctions can be used successfully in open systems consisting of self-interested agents.

In the literature we can find a lot of different auction types, like the Dutch, English, first-price sealed bid and second-price sealed bid auction that all have specific properties. Especially the Vickrey-auction (second-price sealed bid) is a widely adopted protocol in multi-agent systems because of three main characteristics, i) the exchange of information and time consumption is low, ii) it possesses a dominant strategy, namely, to bid one’s ‘true value estimation’ and iii) it is a sealed bid auction, which means that private valuations remain secret (Brandt and Weiß 2001).

We decided not to use it for several reasons. First reason is that one's true valuation in case of freight consolidation is a very vague concept, since it is not possible to determine exactly the costs to transport a load in advance. Secondly, although bidding one's true valuation is the dominant strategy in a one-shot Vickrey auction, it is not in sequential interrelated Vickrey auctions (see, (Sandholm 1996)). Here we assume that the bidder's true valuation means that the value of the auctioned object does not depend on the value other bidders attach to the object and that bidding the true valuation is dominant in the sense that it leads to the highest pay-off ((Sandholm 1996)). This means that one of the main reasons to apply Vickrey-auctions does not hold in sequential Vickrey auctions, since it becomes profitable to model the lowest bid of other players. Thirdly, (looking at practice) also in the human world Vickrey auctions are rarely applied, mainly because bidders have to trust the auctioneer and (s)he must be trustworthy (see e.g. Rothkopf, Teisberg and Kahn (1990)). Finally, profits in a Vickrey auction depend highly on the amount of competitors and the way they bid. This means that the results of a Vickrey auction become harder to analyze, because it is a mix of strategic bidding and profits. We therefore apply the first-price sealed bid auction, since this protocol has the advantage that it makes the analysis a bit simpler by omitting the profits that occur as a result of differences in bids. Nevertheless, our model can be extended to Vickrey-auctions without much effort.

Another source of inspiration comes from the field of revenue (or yield) management as a means of pricing resources dynamically. There is a lot of literature on this topic with well-known applications in air transport (see for an overview, e.g., McGill and van Ryzin (1999), Talluri and Ryzin (2005)). We did not find an application of this idea in vehicle routing as a means to create routes. Most papers on revenue management (especially in the air transport) assume that routes are given (e.g. Amsterdam - New York) and that the question is how to maximize the revenues on these given routes. We combine vehicle routing with revenue management techniques to be able to create profitable routes, to give means for determining an appropriate waiting strategy and to provide a basis for pricing loads in a realistic way.

2.4 Contribution to the literature

From the previous discussion we conclude that it is worthwhile to investigate the benefits of a multi-agent approach for transportation planning. This paper is the first in a series of papers in which we develop a methodology to price loads for consolidated freights, such that vehicle capacity is used efficiently and profitably. For a vehicle it is hard to determine the true valuation of a load, since the value of a load depends on i) the number of loads that can be consolidated on a route and ii) the price the vehicle agent can get. In this paper we present the fundamentals of a method, based on revenue management techniques, to address this problem by adjusting bids dynamically.

We consider a single arc problem, i.e., we consider a vehicle which has to travel from node i to node j and can wait at maximum τ time units. During these τ time units we assume that only loads are offered that also have to be transported between i and j and can possibly be consolidated with previously accepted loads. All loads have unit weights. We develop a method to optimize revenues during the τ time units. The single arc revenue optimization described in the present paper constitutes one of the building blocks of a future paper in which we develop a methodology that is able maximize the turnover of a vehicle facing loads with different weights, picked-up and delivered at different locations with different time windows. The main differences of this paper with standard revenue management techniques is first that a vehicle agent does not use classes (i.e. setting a price and expect some demand), but can bid every price instead (i.e. set a price in respons to demand). Second, a vehicle knows pretty well the total supply of loads since they are all publically auctioned. A vehicle agent does not influence supply by changing its price, but can change its winning probability by changing its price.

3 Model

In this section we first introduce our agent system, we introduce our model and we explain the added value of revenue management techniques in vehicle routing.

3.1 Agent system

The global structure of our multi-agent system (MAS) to model a real-time transportation problem is presented in Figure 1. In the MAS we distinguish one or more transport companies, one or more shipping companies and a market mechanism, which support the assignment of loads to vehicles.




Figure 1: *Global structure of a multi-agent transport model*

In the MAS shipper agents offer loads to the auctioneer. The auctioneer in turn auctions the load and collects the bids of all the transport companies. Transport companies react by means of their fleet owner agents, after consulting their vehicle agents about a reasonable bid.

After the auctioneer receives all the bids of the transport companies, it determines a winner and makes up a contract, unless the shipper agent disagrees on the price. Loads are offered in real-time and are not known beforehand by the transport companies. Transport companies have to bid immediately after they receive an announcement of the auctioneer. As auction we use a first-price sealed-bid auction. In Figure 1 the arrows represent communication lines.

We focus on the intelligence of one vehicle agent and especially on its bidding strategy. The vehicle agent is expected to act in the “best interest” of the vehicle, which in this case means maximizing profits. The task of a vehicle agent is to plan and schedule activities and to bid on loads. The vehicle agent can learn from the past and apply this knowledge in making its decisions. For the vehicle agent its competitors are black boxes that can adopt every possible bidding strategy, although we assume all companies to act rationally with respect to their objectives.

We assume that loads arrive one-by-one (no batch arrivals) and that every load is auctioned separately and chronologically, i.e., in the sequence they are offered to the system.

3.2 Model definition

We define the transport network as a complete directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{A})$ with \mathcal{N} the set of nodes and \mathcal{A} the set of arcs (i, j) between $i, j \in \mathcal{N}$ with $i \neq j$. In the network a vehicle can travel between any two distinct nodes $i, j \in \mathcal{N}$. The shortest distance to travel from i to j is denoted by d_{ij} . We assume that loads are transported between nodes and can be loaded and unloaded at every node. Both actions, loading and unloading, we call handling. For convenience we assume that both loading and unloading at node i take the same time h_i for each load. We assume that h_i is deterministic and depends only on the node (not on the weight of a load). We assume that the loading and unloading docks at every node have infinite capacity.

Loads arrive one by one (no batch arrivals) according to a Poisson process. Every load l has five attributes, namely a weight w_l , an origin i_l and destination j_l , a time window for loading $[a_{i_l}, b_{i_l}]$ and unloading $[a_{j_l}, b_{j_l}]$. Time window violation is not allowed. Every time a new load is offered for auctioning, the process as described in Section 3.1 takes place. Recall that the auction protocol we adopted is the first-price sealed bid auction. Winning prices are not published, which means that a vehicle agent only has information about the number of times it won an auction with a certain bid as fraction of the number of times it bid the same price. Based on this information it can derive information about the probability of winning, bidding a certain bid. Historical information can of course be stored and analyzed at different levels of detail by e.g. aggregating data of locations located in the same regions.

On the network several transport companies operate, which can adopt different bidding strategies. Information about the network, historical data and other relevant information is not shared among transport companies. Within a transport company this information is only available to the company vehicles. Transport companies can communicate with each other and with the auctioneer by their fleet manager. Communication between vehicles is only possible within the same transport company.

We consider the situation of one specific transport company at time t_0 . We want to develop a methodology that enables this company to deal with events happening in the next T time units, i.e., in the time frame $[t_0, t_0 + T]$, with time horizon $T > 0$. In this time frame new loads may arrive. Let us define \tilde{t}_n as the arrival time of the n^{th} load and suppose $t_0 = \tilde{t}_0$, i.e., time t_0 is equal to the arrival time of the first load. We consider a specific vehicle v which is part of the fleet \mathcal{V} of a certain company. The objective of every vehicle is to maximize its profit, i.e., to maximize revenues minus costs. The vehicle has a fixed weight capacity q_v and an average velocity of ζ_v meters per second. We assume that vehicles can combine a load with every other load (no specific shape restrictions) as long as capacity suffices. Loads can be unloaded irrespective of the loading sequence. Preemption is not allowed. Every vehicle maintains a schedule \mathcal{S}_v which is an ordered set of pairs (l, i) , with l a specific load and i the pick-up or delivery location. Schedule \mathcal{S}_v is updated at discrete moments in time through a) completion of a move, b) insertion of a new load or c) a control action resulting in a change in the schedule. Specifically, a control action may be the decision to move pro-actively to another node if more freight can be expected there. A control action also be to wait some time compatible with the time windows. Reason to wait is either there are no loads left in the list of loads or the vehicle expects to retrieve more revenue waiting first than driving first.

Every time a new load l is offered the vehicle agent is asked to calculate a bid for this load. To do so, it inserts the new load into its current schedule \mathcal{S}_v resulting in a temporary schedule \mathcal{S}_v^* which is feasible and results in the lowest decrease of direct costs (the distance the vehicle drives). If the vehicle agent wins the auction, schedule \mathcal{S}_v^* is used from then on. Otherwise the agent maintains its current schedule \mathcal{S}_v and forgets about \mathcal{S}_v^* .

In a future paper we develop a methodology that is able to maximize the turnover of a vehicle facing loads with different weights, picked-up and delivered at different locations with different time windows. This is rather complex however. To provide preliminary insight we present in this paper a methodology for a fundamental, so to say atomic subproblem, namely: to maximize turnover over a single arc between two nodes, once a vehicle agent has committed itself to move over this arc.

3.3 Revenue and utilization in the general MAB-VRP

Before we enter the discussion on single arc revenue optimization, let us elaborate upon some essential aspects of our general MAB-VRP. In this section we first explain why we apply revenue management concepts, we refine our model and we present a utilization measure for consolidated freight transportation problems.

3.3.1 Reasons for applying revenue management concepts

Revenue management focuses on ways to convert existing demand into higher revenues. The reason we use this approach is more or less the same as for airlines. The reason airlines focus on revenue maximization is that they have high fixed costs and relatively low variable costs. When they would, e.g., focus on maximizing profit per passenger they can easily make wrong decisions, because the cost (and profit) per passenger depends for a major part on the fixed costs. The fixed costs per passenger in turn depend on the number of passengers with which these costs can be shared. Moreover, the number of passengers depends on the price airlines ask for a certain trip. Airlines therefore focus on maximizing revenues as a way to maximize profits. In vehicle routing we face the same problem when a vehicle fixates a route by accepting loads. Fixing a route results in fixed costs for the vehicle, because it now has to travel that route for that load. To maximize profits the vehicle can set high prices, but then there is a risk to lose customers and profits. However, setting the price very low results in a lot of customers, possibly more than the vehicle can transport.

An important difference between revenue management in airlines and its application in vehicle routing, is that in the former situation the flight schedule is predetermined. For the airplane it is no question any more which 'route' a plane will fly, but how to maximize the revenues at the planned flights. In vehicle routing we first have to decide which route a vehicle will travel and derived from that the question how to maximize revenues given this route. Moreover, in airlines the flight schedule (not a single passenger) determines the time window an airplane is at an airport, whereas in vehicle routing every load can have an different time window for pick-up or delivery. Another important distinction, is that in airlines a plane cannot pick-up passengers during its flight from city A to city B as easy as e.g. a bus can do. This means that demand is concentrated in a limited number of cities. In vehicle routing not all loads need to have the same origin location and destination location when they travel together between A and B in the same vehicle. Finally, a difference with airlines is that planes have less freedom in determining their waiting strategy. Usually their flight schedules are fixed months in advance whereas the schedule of a vehicle is constructed in real-time. This means that a vehicle can decide during operations when and where to wait, if that leaves more opportunity to gain additional turnover.

3.3.2 Seed and additional loads

We make a distinction between loads dictating the route a vehicle travels and loads that can easily be inserted in the current route. The first type of load we call *seed loads* and the second type of load we call *additional loads*. A load is either a seed load or an additional load. Let us consider load l that is offered at t_0 . Let $TD(\mathcal{S})$ be a function that returns the length of schedule \mathcal{S} in kilometers and $LD(l)$ the distance between the origin i_l and destination j_l of load l . Recall that \mathcal{S}^* is the temporary schedule the vehicle obtains after inserting l in \mathcal{S} . Then we define:

$$\text{load } l \text{ is a } \begin{cases} \text{seed load,} & \text{if } \frac{TD(\mathcal{S}^*) - TD(\mathcal{S}_v)}{LD(l)} \geq \alpha \\ \text{additional load,} & \text{if } \frac{TD(\mathcal{S}^*) - TD(\mathcal{S}_v)}{LD(l)} < \alpha \end{cases}$$

with α the so-called *increase kilometers driven threshold*. In words it says that if the length of the temporal schedule in kilometers (compared to the current schedule) has increased with more than or exactly α times the distance between the origin and destination of load l , the load is considered as seed load, otherwise, it is an additional load. E.g. $\alpha = 1$ might be a useful value. The pricing of these loads has to be done differently, where seed load pricing is by far the most complicated task for several reasons. Firstly, a seed load not only has direct revenues but also limits to a great extent possible future revenues, i.e., by choosing to travel to a certain city one cannot carry loads at the same moment to a city that is in the opposite direction. Secondly, not all loads that are offered during a certain time interval are equally attractive, especially not when they do not fit well in the current schedule. For the additional loads we use revenue management techniques. For seed loads more advanced methods have to be applied.

Of course one might apply another definition to distinguish seed from additional loads. For the moment this definition is convenient to our opinion. Further research is necessary to find out what definition serves our purposes best. In this paper we assume for convenience that time windows of additional loads are large enough with respect to arrival and departure times in the current schedule, i.e., they do not create a need to change arrival and departure times in the current schedule.

In this paper we only consider the pricing of additional loads. Considering both the pricing of seed and additional loads at once is rather complex, since they mutually influence each other in an intricate play which has a great impact on the aim to maximize turnover. In a future paper we will study the pricing of seed loads which needs theory about pricing additional loads to understand the intricate relations between the two.

3.3.3 Measuring vehicle utilization in consolidated shipments

The question may arise why to focus on revenue maximization instead of vehicle utilization. The latter could possibly result in more efficient plans. We think however that to optimize a transportation planning for consolidated freight in a system in which multiple companies participate, we have to take both revenues and vehicle utilization into account, or stated more generally: revenues and costs. The reason is simply that a company aims to make profit by means of transportation, which means to maximize turnover and minimize costs. By applying revenue management ideas to transportation planning an agent can weigh possible revenues on the one hand with costs for transportation on the other hand. Moreover, revenue management based techniques can support a suitable way to set prices for customers and to distribute gains over the participating companies. Although in revenue management it might be obvious that it results in higher utilization of the resource, it is not in vehicle routing since a vehicle has to weigh the additional effort (e.g. loading and unloading time) of transporting an additional load with the additional revenues. To investigate whether focussing on revenue maximization will indeed results in higher vehicle utilization, we first have to define what we mean with vehicle utilization.

Measuring the average utilization degree of a vehicle in the case of freight consolidation is however not trivial. First reason is that the number of loads a vehicle transports can change every moment in time. To measure the average utilization per kilometer one can register for every kilometer driven how many loads a vehicle was transporting. However, and this is the second reason, this way of measuring the utilization does not take into account the necessity of the transportation. Vehicles can transport e.g. a kind of dummy load, filling half of their capacity, to get a high utilization degree. We therefore propose the following definition which takes the previously mentioned issues into account. This definition is new, as far as we know, and has the advantage that it can be calculated very easily and very fast. In practice this measure can be very helpful as well.

Definition: Let $[T_1, T_2]$ be a time interval and v a vehicle. For each load l let:

D_l = the time at which load l is picked-up

$t_{i_l j_l}$ = the shortest time to travel from origin i_l to destination j_l of load l

τ_l = the time load l is actually in the vehicle during time interval $[T_1, T_2]$

Let $I_l = [D_l, D_l + t_{i_l j_l}] \cap [T_1, T_2]$. Let m be the standard Euclidian measure on \mathbb{R} . Note that $m(I_l) \leq \tau_l$. We now define transport utilization degree $U_v^{[T_1, T_2]}$ for vehicle v over time interval $[T_1, T_2]$ as follows:

$$U_v^{[T_1, T_2]} = \frac{\sum_l w_l \cdot m(I_l)}{q_v \cdot (T_2 - T_1)}$$

Theorem 1 Let v be a vehicle and $[T_1, T_2]$ a time interval and assume that the schedule of v was feasible w.r.t. capacity constraints, then $U_v^{[T_1, T_2]} \in [0, 1]$.

Proof Let t_1, t_2, \dots, t_K be the times within $[T_1, T_2]$ at which load changes take place, i.e., either a load is loaded or a load is unloaded. Let $t_0 = T_1$ and $t_{K+1} = T_2$. Then for each l :

$$\tau_l = \sum_{k=0}^{K+1} \alpha_k^l \cdot (t_{k+1} - t_k) \text{ where } \alpha_k^l = \begin{cases} 1, & \text{if load } l \text{ is on board during } [t_{k+1} - t_k] \\ 0, & \text{else} \end{cases}$$

Hence,

$$\begin{aligned} \sum_l w_l \cdot m(I_l) &\leq \sum_l w_l \cdot \tau_l = \sum_l w_l \left(\sum_{k=0}^{K+1} \alpha_k^l \cdot (t_{k+1} - t_k) \right) = \sum_{k=0}^{K+1} \left(\sum_l w_l \cdot \alpha_k^l \right) \cdot (t_{k+1} - t_k) \leq \\ &\sum_{k=0}^{K+1} q_v \cdot (t_{k+1} - t_k) = q_v \cdot (T_2 - T_1) \blacksquare \end{aligned}$$

The measure assumes in fact that a vehicle can drive ideally by considering necessary handling and waiting time as time that should be minimized as much as possible to use the vehicle as efficiently as possible. This assumption is a bit unrealistic, since a utilization degree of 100% can only be realized in a utopian situation or in a very small time interval.

4 Pricing additional loads

In this section we investigate the way additional loads can be priced to maximize revenues which in turn means maximizing profits. We consider a single arc problem, i.e., consider a vehicle which has to travel from node i to node j and can wait at maximum τ time units. During these τ time units we assume that only loads are offered that also have to be transported between i and j and can possibly be consolidated with previously accepted loads. All loads have unit weights. If we use the term *cost price* we mean the travel costs travelling from i to j and the costs for loading at i and unloading at j .

4.1 Introduction

In this section we first elaborate upon modelling the winning probability of a bidder in a single auction. Then we derive the optimal prices in sequential first-price sealed bid auctions by means of a dynamic programming (DP) recursion. For the DP we discretize time to in intervals of length ϵ . However, evaluating a dynamic programming (DP) recursion in real-time to find the optimal price to bid can be very time consuming. Especially when parameters change frequently, a DP recursion might become impractical. Parameters change e.g. as result of changes in the strategy of competitors. In this section we therefore derive an approximation (in continuous time) of the DP, to perform calculations quickly without losing too much of the solution quality. To derive

the approximations we use the equal price assumption, which we explain in a moment. In this paper we only consider the setting mentioned at the start of this section, which is probably not very realistic. In a future paper we aim to extend these approximations to deal with settings as mentioned in Section 3.2.

4.2 Modelling winning probabilities

In this paper we focus on one of the transport companies trying to create an intelligent mass to gain competitive advantage. For this transport company its competitors are black boxes that can adopt every possible bidding strategy, although we assume all companies to act rationally with respect to their objectives. Our transport company is assumed to be a *price taker*, which means that its actions do not influence the price on the market. We also assume that the lowest price of competitors is distributed according to a Weibull distribution and that auctions are i.i.d. We choose for this distribution for several reasons. First of all, we are mainly interested in the left side of the distribution, not in the tail. It is well-known that in that case, a lot of stochastic models are not very sensitive for more than the first to moments of the underlying probability distribution (see, Tijms (1986)). The reason is that the difference between most distributions (e.g. gamma, log-normal or Weibull) becomes most significant in their tails. Second reason is that the theory of the so-called Extreme Value Distributions (EVD) supports our choice. The EVD is the limiting distribution for the minimum or the maximum value of a large collection of random observations from the same arbitrary distribution. We can use this distribution independent of the bidding distribution of a single company and independent of the type of auction we are using. For further reading we refer to Reiss and Thomas (1997) and Mes, van der Heijden and Schuur (2006).

4.3 Optimal prices in sequential first-price sealed bid auctions

We divide the time-to-go in sufficiently small intervals with length ϵ . We assume that in such an interval of length ϵ at most one load is offered. The interarrival time between two jobs is assumed to be exponentially distributed. The probability that a load is offered within this interval is equal to the probability that the interarrival time between two jobs is smaller than ϵ . So the probability of one load arrival is

$$\begin{aligned} \text{one arrival} &:= p_\epsilon = F(\epsilon) = 1 - e^{-\lambda\epsilon} \\ \text{no arrival} &:= 1 - p_\epsilon = 1 - F(\epsilon) = e^{-\lambda\epsilon} \end{aligned}$$

At every decision moment j either no load or one load is offered. If a load is offered, the optimal price has to be determined based on the time-to-go and the remaining capacity. We define a state (b, c) by 1) the remaining capacity c and 2) whether a load is offered ($b = 1$) or no load is offered ($b = 0$).

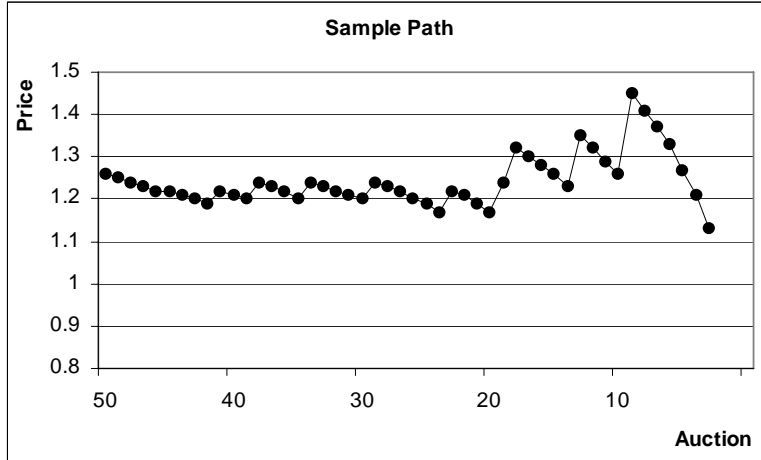


Figure 2: *Sample price course when a vehicle agent bids successively on loads*

We can now derive the value function $V^j(b, c)$ for every state (b, c) in stage j :

$$\begin{aligned}
 V^j(1, c) &= \max_{x \in S} [p(x) \cdot [x + p_\epsilon \cdot V^{j-1}(1, c-1) + (1-p_\epsilon) \cdot V^{j-1}(0, c-1)] + (1-p(x)) \cdot V^{j-1}(0, c)] \\
 V^j(0, c) &= p_\epsilon \cdot V^{j-1}(1, c) + (1-p_\epsilon) \cdot V^{j-1}(0, c)
 \end{aligned}$$

Remark that in the states $(1, c)$ the vehicle agent has additional information in contrast to $(0, c)$ states, namely that in these states an auction is held for sure and auctions in the time-to-go appear with certain probability. The states $(1, c)$ we call active states, the states $(0, c)$ are called passive states.

To get a feeling for the prices a vehicle agent would bid in successive auctions using the presented DP, we draw a sample path for a Weibull distribution $e^{-\left(\frac{x}{\eta}\right)^\gamma}$ with parameters $\eta = 1$ and $\gamma = 2$ and with $\lambda = 1$ and $\tau = 50$ and $c = 10$. The length of the intervals $\epsilon = 0.001$, so every time unit (in which on average one load is expected) is divided in 1000 smaller intervals. The y-axis depicts prices normalized to one and the x-axis shows the number of auctions in the time-to-go.

In Figure 2 we see that the course of the bid sequence is clearly indented. Increases in the bid price follows directly after the vehicle agent wins an auction and hence has less remaining capacity. The indented curve is in accordance with results obtained by Talluri and Ryzin (2005). It means that it is optimal to bid a different price in every auction and that the price course is not necessarily convex. This is of interest for the development of our approximations.

4.4 Optimal price approximation using equal price assumption

The approximations are based on the so-called *equal price* assumption. With this assumption we mean that a company determines in a certain state (b, c) the optimal price in this auction,

assuming that it bids this price in all the auctions that possibly come during τ . In reality it bids in every auction a different price, since it in every auction evaluates the best price assuming that in all the auctions that possible come during τ it bids this price as well.

To simplify the analysis, we first consider the passive states as stepping stone for the active states. We discuss them successively.

4.4.1 Analytical derivation for passive states

First we make a derivation for the expected turnover considering the $(0, c)$ states and based on that we give the expression to find the optimal price in a specific state.

Recall that c denotes the remaining capacity of the vehicle and τ is the time-to-go. We write $\xi_k(\tau) = \frac{(\lambda\tau)^k e^{-\lambda\tau}}{k!}$ for the (Poisson) probability that during τ exactly k auctions will be held. Suppose the agent bids price x . The probability that during τ exactly m auctions will be won if each bid equals x is given by

$$\begin{aligned} P(m, x, \tau) &= e^{-\lambda\tau} \cdot p(x)^m \sum_{s=0}^{\infty} \frac{(\lambda\tau)^{m+s}}{(m+s)!} \cdot \binom{m+s}{m} \cdot (1-p(x))^s \\ &= e^{-\lambda\tau} \frac{(\lambda\tau p(x))^m}{m!} \sum_{s=0}^{\infty} \frac{(\lambda\tau)^s}{s!} \cdot (1-p(x))^s = \frac{(\lambda\tau p(x))^m}{m!} e^{-\lambda\tau p(x)} \end{aligned}$$

Note that $P(m, x, \tau)$ defines a Poisson arrival process with intensity $\lambda p(x)$. Now, the expected turnover $\mathbb{E}^{<c}[x]$ corresponding with a remaining capacity usage less than c is equal to:

$$\mathbb{E}^{<c}[x] = \sum_{m=0}^{c-1} mxP(m, x, \tau)$$

And the expected turnover $\mathbb{E}^c[x]$ corresponding with a remaining capacity usage equal to c is given by:

$$\mathbb{E}^c[x] = \sum_{m=c}^{\infty} cxP(m, x, \tau)$$

Combining the terms we find for the total turnover:

$$\mathbb{E}[x] = \mathbb{E}^{<c}[x] + \mathbb{E}^c[x] = x \sum_{m=0}^{\infty} \min\{m, c\} P(m, x, \tau) \quad (1)$$

To find the optimal price maximizing turnover in a specific state we take the derivative of $\mathbb{E}[x]$ and set it to zero. To use the chain rule for derivation we first define for $\alpha = \lambda\tau p(x)$

$$\begin{aligned} \Upsilon(\alpha) &\equiv e^{-\alpha} \sum_{m=0}^{\infty} \min\{m, c\} \frac{\alpha^m}{m!} \\ &= c - e^{-\alpha} \sum_{m=0}^{c-1} (c-m) \frac{\alpha^m}{m!} \end{aligned}$$

For the derivative we obtain

$$\frac{d}{d\alpha} \Upsilon(\alpha) = e^{-\alpha} \sum_{m=0}^{c-1} \frac{\alpha^m}{m!}$$

The expected turnover is equal to $x \cdot \Upsilon(\alpha)$, with $\alpha = \lambda\tau p(x)$. The derivative is equal to:

$$\frac{d}{dx} (x \cdot \Upsilon(\alpha)) = c + e^{-\alpha} \sum_{m=0}^{c-1} \left(x \frac{d\alpha}{dx} - c + m \right) \frac{\alpha^m}{m!}$$

It is hard to find analytically the value x for which $\mathbb{E}[x]$ is maximized and we therefore propose to use e.g. Regula Falsi to find it numerically. Applying a numerical optimization method, one has to be aware that the derivative has a lot of roots, since it goes to zero when x goes to infinity.

4.4.2 Analytical derivation for active states

Again we first derive the expected turnover considering the $(1, c)$ states and based on that we give the expression to find the optimal price in a specific state. The expected turnover is now given by:

$$\mathbb{E}[x] = xp(x) + x \cdot e^{-\lambda\tau p(x)} \cdot \sum_{k=0}^{\infty} A_{k,c,x} \frac{(\lambda\tau p(x))^k}{k!}$$

where $A_{k,c,x} = [p(x) \cdot \min\{k, c-1\} + (1-p(x)) \cdot \min\{k, c\}]$. We can now distinguish two cases:

1. $k \leq c-1 \Rightarrow A_{k,c,x} = p(x)k + (1-p(x))k = k$
2. $k \geq c \Rightarrow A_{k,c,x} = p(x)(c-1) + (1-p(x))c = c - p(x)$

This results in:

$$\mathbb{E}[x] = xp(x) + x \cdot e^{-\lambda\tau p(x)} \cdot \left[\sum_{k=0}^{c-1} k \frac{(\lambda\tau p(x))^k}{k!} + \sum_{k=c}^{\infty} (c-p(x)) \frac{(\lambda\tau p(x))^k}{k!} \right]$$

However, we know that $0 \leq p(x) \leq 1$ which means that

$$\mathbb{E}[x] = x \cdot \left\{ p(x) + e^{-\lambda\tau p(x)} \cdot \sum_{k=0}^{\infty} \min(k, c-p(x)) \frac{(\lambda\tau p(x))^k}{k!} \right\} \quad (2)$$

Let us again define $\Upsilon(\alpha)$ as

$$\Upsilon(\alpha) \equiv \frac{\alpha}{\lambda\tau} + e^{-\alpha} \sum_{k=0}^{\infty} \min\left\{k, c - \frac{\alpha}{\lambda\tau}\right\} \frac{\alpha^k}{k!}, \text{ where } \alpha = \lambda\tau p(x)$$

Note that the function $\Upsilon(\alpha)$ can also be written as:

$$\Upsilon(\alpha) = c - e^{-\alpha} \sum_{k=0}^{c-1} \left(c - \frac{\alpha}{\lambda\tau} - k \right) \frac{\alpha^k}{k!}$$

A somewhat lengthy, but straightforward calculation yields:

$$\frac{d}{d\alpha} \Upsilon(\alpha) = e^{-\alpha} \sum_{k=0}^{c-1} \frac{\alpha^k}{k!} - e^{-\alpha} \frac{\alpha^c}{\lambda\tau(c-1)!}$$

The expected turnover is equal to $x \cdot \Upsilon(\alpha)$, where $\alpha = \lambda\tau p(x)$. The derivative is equal to

$$\frac{d}{dx} (x \cdot \Upsilon(\alpha)) = \Upsilon(\alpha) + x \left[e^{-\alpha} \sum_{k=0}^{c-1} \frac{\alpha^k}{k!} - e^{-\alpha} \frac{\alpha^c}{\lambda\tau(c-1)!} \right] \frac{d\alpha}{dx}$$

Again it is hard to find analytically the value x for which $\mathbb{E}[x]$ is maximized and we therefore propose to use e.g. Regula Falsi in this case as well to find the optimum numerically. Applying a numerical optimization method, one has to be aware again that the derivative has a lot of roots, since it goes to zero when x goes to infinity.

4.4.3 Approximation for active states

A simple approximation of equation 2 we be derived as follows:

$$\begin{aligned} \mathbb{E}^{approx}[x] &= xp(x) + xe^{-\lambda\tau p(x)} \cdot \sum_{k=0}^{\infty} \min(k, c-p(x)) \frac{(\lambda\tau p(x))^k}{k!} \\ &\leq xp(x) + xe^{-\lambda\tau p(x)} \cdot \min \left\{ \sum_{k=0}^{\infty} k \frac{(\lambda\tau p(x))^k}{k!}, \sum_{k=0}^{\infty} (c-p(x)) \frac{(\lambda\tau p(x))^k}{k!} \right\} \\ &= xp(x) + xe^{-\lambda\tau p(x)} \cdot \min \left\{ \lambda\tau p(x) e^{\lambda\tau p(x)}, (c-p(x)) e^{\lambda\tau p(x)} \right\} \\ &= x \min \{ (\lambda\tau + 1) p(x), c \} \end{aligned}$$

The price x for which $\mathbb{E}^{approx}[x]$ is maximized when $(\lambda\tau + 1)p(x) = c$. Recall that we assume $p(x)$ to be Weibull distributed with $p(x) = e^{-\left(\frac{x}{\eta}\right)^\gamma}$ which means that

$$\arg \max_x \mathbb{E}^{approx}[x] = \begin{cases} \gamma^{-\left(\frac{1}{\gamma}\right)} \cdot \eta, & \text{if } e^{\frac{1}{\gamma}} \geq \frac{(\lambda\tau+1)}{c} \\ \sqrt[\gamma]{-\ln \frac{c}{(\lambda\tau+1)}} \cdot \eta, & \text{otherwise} \end{cases}$$

In this case we in fact use the expectation instead of the distribution to find the price for which $\mathbb{E}^{approx}[x]$ is maximized. This a more deterministic approach. In the next section we compare the results of this approach with the results of the DP of Section 4.3.

5 Numerical results

To compare the performance of the approximations in the previous section with the DP we can use two methods. First, we can draw a lot of sample paths and see what the difference in the average expected turnover is. However, due to the number of parameters $(\lambda, \eta, \gamma, c, \tau)$ we would

get a large table with results. The second approach is to perform a DP like recursion for the approximations as well. We divide time in intervals of length ϵ and perform the recursion, but instead of choosing the optimal price to bid in every stage j and every state $(1, c)$ (as in the DP) we use the approximations determine the bid. In this way we in fact determine all possible sample paths. After this DP-like recursion, we compare the expected turnover in every stage j and every state $(1, c)$ of the DP with one of the approximations, resulting in the relative performance of the DP compared to the approximations. We refer to the approximation of Section 4.4.2 as the Analytical Equal Price Formula (Analytical EPF) and to the approximation of Section 4.4.3 as the Approximated Equal Price Formula (Approx EPF).

In this section we evaluate the performance of the DP and the approximations by considering three scenarios. The three scenarios are defined by the function describing the lowest bid of all competitors. As mentioned in Section 3.2 we assume the lowest bid of all competitors to be Weibull distributed as $e^{-\left(\frac{x}{\eta}\right)^\gamma}$, with scale parameter η and shape parameter γ both being strictly positive. We choose the parameters η and γ such that the sensitivity of the turnover for small deviations from the optimal price in a single auction varies from lower to higher. The optimal price (denoted by x^*) is in all scenario's more or less the same to facilitate comparison.

	η	γ	x^*
Scenario 1 (S1)	1	2	0.71
Scenario 2 (S2)	1	3	0.69
Scenario 3 (S3)	1	5	0.72

In a plot the probability functions and the expected turnover functions can be depicted as in Figure 3.

We use the same settings for λ and ϵ as in section 4.3. We let τ vary between 0..50 and the remaining capacity c between 0..50.

5.1 Sample path

To get a feeling of the corresponding price course derived from our EPFs, we depicted the approximated price in every state besides the DP price in the sample path of Figure 2. Note that this is just for comparison of the course of the prices, not to say anything about revenues. If the agent would indeed have used the approximation for bid price determination, it would have influenced its winning probabilities and so its revenues.

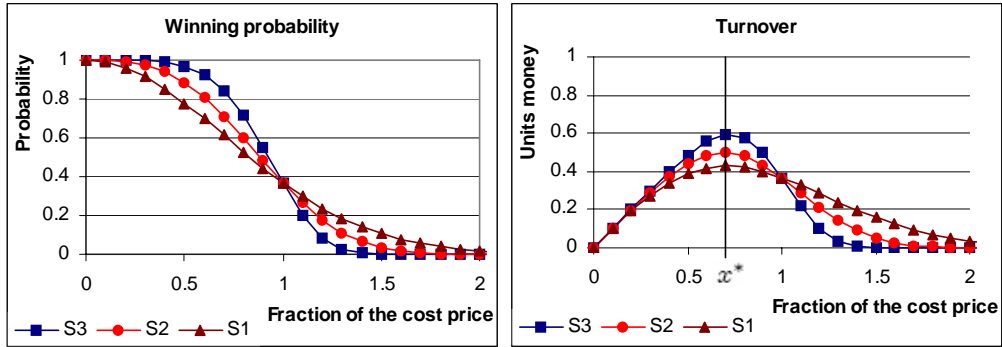


Figure 3: *The three scenario's depicted with respect to the winning probability (resp. the turnover) related to the fraction of the cost price*

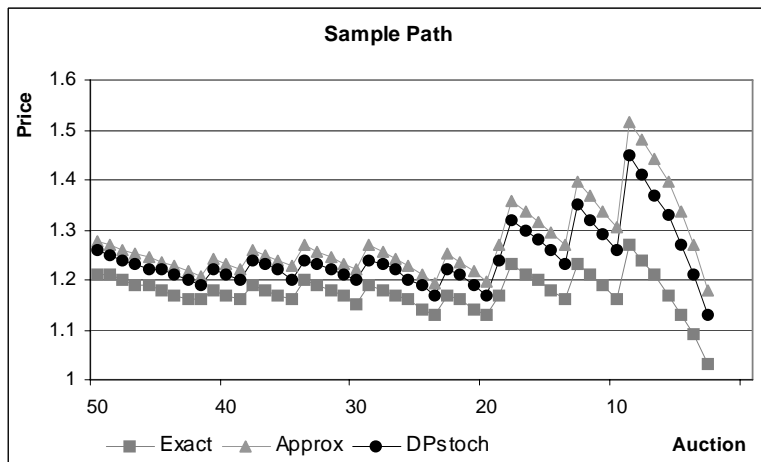


Figure 4: *Sample price course when a vehicle agent bids successively on loads*




Figure 5: *The expected turnover of the exact DP-recursion divided by the expected turnover using the Analytical EPF in every state.*

From Figure 4 we see several things. First, the course of the sequence for both the approximations is indented as well. Second, notice that the sequence of bid prices is not linear in the end, but follows a bend curve. One can see that our approximations follows the price course, though they both deviate a bit from the DP-curve.

Concluding, we see that our EPFs, despite the equal-price assumption, shows the same price course as the DP-recursion. Moreover, the picture gives rise to the thought that the Approx EPF approximates the DP recursion pretty well, especially for large n and c .

5.2 Scenario comparison

In this section we investigate to what extend the EPFs approximate the DP. To do so, we divided the expected turnover in every stage j and every state $(1, c)$ in the DP by the expected turnover in the same stage j and state $(1, c)$ using one of the EPFs. First consider the Analytical EPF.

From Figure 5 we firstly that the flatter the expected revenue curve is, the lesser the difference in expected turnover is with respect to the optimum. Secondly, the smaller the remaining capacity the greater the difference in turnover becomes.

Considering the difference in expected turnover between the Approx EPF and the DP, as it is depicted in Figure 6, we see that depending on the scenario the Approx EPF approximates the DP better than the Analytical EPF (considering all stages and states). Note first that the Approx EPF deviates more from the DP when the corresponding the Weibull-distribution is more sensitive for the optimum. Secondly, in every picture we see a kind of hill of about 1% high that goes in diagonal direction and decreases slowly. This hill is the result of the deterministic approach, i.e.,




Figure 6: *The expected turnover of the exact DP-recursion divided by the expected turnover using the Approx EPF in every state.*

the fact that we use expectations instead of distributions to optimize the bid prices. Consider for instance the situation of $c = 4$ and suppose we expect to win four auctions in the time-to-go. In the approximation we assume to win exactly four auctions, although in reality it can be more (which are impracticable) and on average it will probably be less than four. This means that the utilization of the vehicle is probably less than one, just because we were too optimistic about the future.

6 Conclusions and directions for further research

In this paper we consider an agent-based less-than-full truckload transportation problem. In this transportation problem loads are assigned to vehicles by means of a first-price sealed-bid auction. A supplier offers a load to the auctioneer, which in turn invites all the vehicle agents in the system to bid on the load. Loads are assigned solely based on price.

We consider a single vehicle v at time t_0 which has to travel from i to j and can wait at maximum τ time units. During these time units we assume that only loads are offered that also have to be transported between i and j and can possibly be consolidated with previous accepted loads. All loads have unit weights. For this specific vehicle, which can be part of a larger fleet of a certain company, all competitors are black-boxes and the vehicle agent is assumed to be *price-taker*. We assume the lowest bid of all competitors to be Weibull distributed.

The question we address is how to price a load in real time when this load can probably be consolidated with other loads in the future. Bidding in real-time means that a vehicle agent has to respond to an auction directly and for the auctioned load exclusively. Pricing should be done

dynamically, to be able to survive in a changing competitive environment. Our approach consists of several steps. First the distinction between two types of loads, namely so-called *seed loads* and *additional loads*. Additional loads are loads that are 'easily' included in the current schedule, i.e., result in only small changes in the current plan of the vehicle. Seed loads, on the other hand, determine to a great extent the route the vehicle drives and so have a great impact to be *potentially* profitable. The next step is to find a model to price loads such that expected profits are maximized. In this paper we consider the pricing of additional loads, since the pricing of seed loads is more complicated and depends on theory about pricing additional loads.

We develop a DP-recursion which provides the optimal price in a certain state, which is defined by the remaining vehicle capacity and the expected number of auctions in the time-to-go. However, evaluating a dynamic programming (DP) recursion in real-time to find the optimal price to bid can be very time consuming. We therefore derive two approximations which speed up the calculations and to analyze whether they provide a basis for more sophisticated models incorporating e.g. different weights of load, different time windows, different pick-up and delivery locations et cetera.

The question may arise whether the derived approximations of the optimum are useful and applicable in practise. A difference in turnover, which are profits as well in this case, of about five percent can be pretty much in today's competitive world. Although this is true, we like to mention firstly that the difference in the Approx EPF has a clear maximum for small c and n . Secondly, it would be interesting to investigate to what extent the assumption under which these approximations are derived (namely using the expectation instead of the distribution) has a great impact when more complexity is added to the model. Considering the benefits of fast calculations, we think the approximations can be useful in pricing additional loads. As alternative, we can also apply a mixed model, i.e., using the DP-recursion to determine the price for small c and n and using the approximation for large c and n .

The single arc revenue optimization described in the present paper constitutes one of the building blocks of a future paper in which we develop a methodology that is able maximize the turnover of a vehicle facing loads with different weights, picked-up and delivered at different locations with different time windows. This includes more sophisticated additional and seed load pricing models and a profound theory about the relation between additional and seed load pricing.

Acknowledgement. *This research is performed within a TRANSUMO program. TRANSUMO is a Dutch foundation which support research to realize a transition to sustainable mobility. We are grateful to TRANSUMO for financing this research.*

References

- Arbib, C. and F. Rossi (2000). Optimal resource assignment through negotiation in a multi-agent manufacturing system. *IIE Transactions* 32, 963–974.
- Brandt, F., W. Brauer, and G. Weiß (2000). Task assignment in multiagent systems based on vickrey-type auctioning and leveled commitment contracting. In *In Proceedings of the Fourth International Workshop on Cooperative Information Agents (CIA-2000)*, Boston, MA.
- Brandt, F. and G. Weiß (2001). Antisocial agents and vickrey auctions. In J. M. Tambe and M. (Eds.), *Intelligent Agents VIII*, Volume 2333 of *Lecture Notes in Artificial Intelligence*, pp. 335–347. Berlin Heidelberg: Springer-Verlag.
- Clearwater, S. (1996). *Market-based control: a paradigm for distributed resource allocation*. Singapore: World Scientific.
- Dewan, P. and S. Joshi (2002). Auction-based distributed scheduling in a dynamic job shop environment. *International Journal of Production Research* 40(5), 1173–1191.
- Ertogral, K. and S.D. Wu (2000). Auction-theoretic coordination of production planning in the supply chain. *IIE Transactions* 32, 931–940.
- Figliozzi, M.A., H.S. Mahmassani, and P. Jaillet (2004). Competitive performance assessment of dynamic vehicle routing technologies using sequential auctions. *Transportation Research Record* 1882, 10–18.
- Garey, M. and D. Johnson (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Fransico.
- Hoen, P.J. 't and J.A. la Poutré (2004). A decommitment strategy in a competitive multi-agent transportation setting. In D. P. P. Faratin and J. Rodriguez-Aguilar (Eds.), *Agent Mediated Electronic Commerce V (AMEC-V)*, pp. 56–72. Berlin: Springer - Verlag.
- McGill, J.I. and G.J. van Ryzin (1999, May). Revenue management: research overview and prospects. *Transportation Science* 33(2), 233–256.
- Mes, M.R.K., M.C. van der Heijden, and Peter Schuur (2006). Opportunity costs calculation in agent-based vehicle routing and scheduling. *Beta Working Paper Series, WP-168*.
- Mes, M.R.K., M.C. van der Heijden, and A. van Harten (2006). Comparison of agent-based scheduling to look-ahead heuristics for real-time transportation problems. *Accepted for publication in European Journal of Operation Research*.
- Parunak, H. Van Dyke (1998). Practical and industrial applications of agent-based systems. *Environmental Research Institute of Michigan (ERIM)*.

- Porter, M.E. (1985). *Competitive Advantage: Creating and Sustaining Superior Performance*. New York: Free Press.
- Qinghe, H., A. Kumar, and Z. Shuang (2001). A bidding decision model in multiagent supply chain planning. *International Journal of Production Research* 39(15), 3291–3301.
- Reiss, R.K. and M. Thomas (1997). *Statistical Analysis of Extreme Values*. Basel, Switzerland: Birkhauser Verlag.
- Rothkopf, Michael H, Thomas J Teisberg, and Edward P Kahn (1990). Why are vickrey auctions rare? *Journal of Political Economy* 98(1), 94–109. available at <http://ideas.repec.org/a/ucp/jpolec/v98y1990i1p94-109.html>.
- Sandholm, T. (1996). Limitations of the vickrey auction in computational multiagent systems. In *In Proceedings of the 2nd International Conference on Multiagent Systems (ICMAS-96)*, Menlo Park, CA, pp. 299–306. AAAI Press.
- Talluri and Van Ryzin (2005, februari). *The theory and practise of revenue management* (first ed.). Springer.
- Tijms, H.C. (1986). *Stochastic Modelling and Analysis: A Computational Approach* (1 ed.). John Wiley & Sons Ltd.
- Toth, P. and D. Vigo (2002). *The Vehicle Routing Problem*. SIAM Monographs on Discrete Mathematics and Applications. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics.
- Weerdt, M. de (2003). *Plan Merging in Multi-Agent Systems*. Ph. D. thesis, Technische Universiteit Delft.
- Wooldridge, M. and N.R. Jennings (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review* 10(2), 115–152.