Centrum voor Wiskunde en Informatica

**REPORT***RAPPORT*

SEN

Software Engineering

*Software ENgineering*

SEN

# Bidding with decommitment in a multi-agent transportation model

P.J. 't Hoen, D.D.B. van Bragt, J.A. La Poutré

**REPORT SEN-R0220 NOVEMBER 30, 2002**

CWI is the National Research Institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organization for Scientific Research (NWO).
CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)


P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-369X

# Bidding with Decommitment in a Multi-Agent Transportation Model

P.J. 't Hoen[1] D.D.B. van Bragt[1] J.A. La Poutré [1,2]
hoen@cwi.nl David.van.Bragt@cwi.nl Han.La.Poutre@cwi.nl
[1] *CWI, Centre for Mathematics and Computer Science*
*P.O. Box 94079, 1090 GB Amsterdam, The Netherlands*
[2] *School of Technology Management; Eindhoven University of Technology*
*De Lismortel 2, 5600 MB Eindhoven, The Netherlands*

ABSTRACT
Decommitment is the action of foregoing a previous contract for another (superior) offer. It has been shown that, using decommitment, agents can reach higher utility levels in case of negotiations with uncertainty about future prospects. In this paper, we study the decommitment concept for the novel setting of a large-scale logistics setting with multiple, competing companies. Orders for transportation of loads are acquired by agents of the (competing) companies by bidding in parallel auctions. We find significant increases in profit when the agents can decommit and postpone the transportation of a load to a more suitable time. Furthermore, we analyze the circumstances for which decommitment has a positive impact if agents (trucks) are capable of handling multiple contracts simultaneously.

*2000 Mathematics Subject Classification:* 68T05, 68T20, 68T37, 68U20
*1998 ACM Computing Classification System:* C.2.4, H.4.2, I.2.11
*Keywords and Phrases:* Multi-Agent Systems, Decommitment, Auctions, Logistics
*Note:* Work carried out under theme SEN4 "Evolutionary Systems and Applied Algorithmics". This research has been performed within the framework of the project "Distributed Engine for Advanced Logistics (DEAL)" funded by the E.E.T. program in the Netherlands.

## 1. Introduction

A levelled commitment protocol for negotiations between agents is presented in [20, 22, 1, 21]. In this protocol, agents have the opportunity to unilaterally *decommit* from contracts. That is, they can forego a previous contract for another (superior) offer. Sandholm *et al.* have shown formally that by incorporating this decommitment option the degree of Pareto efficiency of the reached agreements can increase. Decommitment improves the performance as agents can escape from premature local minima by adjusting their contracts to new offers that were not known when the earlier contracts were negotiated.

A recent development is the use of multi-agent systems (MASs) [7, 2, 14, 25]. This research field aims at the development of robust, distributed market mechanisms [5, 15]. In this paper, we consider the role of decommitment within a MAS logistics setting. In our model, companies compete in an open and dynamic market for the transportation of cargo. Each truck is equipped with an agent. These agents are are able to bid for cargo which is offered in distributed, online auctions. In this setting, decommitment is the possibility of an agent to forego a previously won contract in favor of a more profitable load.

We show in a series of computer experiments that a significant increase in performance can be realized by a company which allows its agents to decommit loads (as opposed to a company with agents who only employ the option of regular, binding bidding). In particular, the experiments show that decommitment is a clearly superior strategy if an agent is close to the limit of its capacity. This is a general result for agents capable of simultaneous tasks who use a decommitment strategy.

We study the impact of a decommitment strategy for various settings. For example, we vary the number of companies and trucks, the distribution of cargo, the shape of the transportation grid, etc. Furthermore, we investigate the impact of decommitment when cargo transport is from distribution center to distribution center (a setting which is relevant for international transport). We also examine the impact of decommitment for different price functions and in case of "myopic" bidding agents (who search for new cargo within a limited distance).

The remainder of this paper is organized as follows. Section 2 discusses the transportation model that we use in this paper. The market mechanism is described in Section 3. Section 4 discusses our application of decommitment in a market setting. The computer experiments are presented in Section 5. Section 6 contains some concluding remarks.

## 2. The Transportation Model

In this section, we present the transportation model that is used in this paper.[1] We have kept the transportation model, the market mechanism, and the structure of the bidding agents relatively simple to keep the analysis as transparent as possible. Some extensions of the basic model are further discussed in Section 5.

We show in Section 5 that performance can increase significantly when a decommitment strategy is used. We expect the (positive) effect of decommitment to increase when the complexity of the transportation model increases (because the uncertainty of possible future events also increases in this case). In Section 5.7 we investigate some venues to substantiate this claim.

### *2.1 Outline*

We consider a model with competing companies. Each company has a collection of trucks which are placed on a grid and which together add to the company's profit by transporting cargo. We assume that the companies try to optimize their net profit. Each truck is coupled with an agent that bids for cargo for its "own" truck. The goal of such an agent is to maximize the profit made by the truck.[2] Unless stated otherwise, we assume (for simplicity and to facilitate the analysis of the model's results) that all companies consist of the same number of (homogeneous) trucks.

Poot *et al.* [16] give an exhaustive list of performance measures for the transportation of cargo found in literature. The basic performance measures that we consider are (i) the profit made as a function of the total number of transported loads, (ii) the profit as a function of the bulk of the transported loads, and (iii) the costs as a function of the distance travelled for the made deliveries.

The grid is filled with a number of depots. Loads for pickup are locally aggregated at these depots. Such an aggregation procedure is for example used by UPS,[3] where cargo is first delivered to one of the nearby distribution centers. Warehousing, where goods from multiple companies are collected for bundled transport, is another example. This aggregation can take place over relatively short distances or over more substantial distances (e.g., in case of international transport). In general, the origin of loads will not be randomly distributed but clustered, depending on population centers and business locations [13]. We can thus also consider depots as abstractions of important population or business centers. Section 5 presents such a model.

In the experiments, the delivery point for a load is randomly chosen.[4] In Section 5, we also experiment with depot-to-depot delivery of cargo (to simulate international transport). In that case, depots double as aggregation centers for international cargo (before national distribution takes place).

A new load should be delivered no later than the end of the next day. Most regular mail services (e.g., UPS) and many wholesale suppliers employ such a model of "next day delivery". In the simulations,

[1] The computer model has been programmed in the Java programming language (version 1.4). We thank Stefan Blom for allowing us to use the STW cluster at CWI.

[2] In the text, we sometimes blur the line between the agent and its truck and speak of "the truck bids" instead of "the agent associated with the truck bids".

[3] See www.ups.com.

[4] As long as the transportation distance is not zero.

each depot has a number of loads available for transport at the start of the day. New orders also arrive for transport in the course of the day.

Each depot serves as an independent auction site where trucks of the individual companies compete for the available cargo (see Section 3.1). The trucks can place bids on the initial orders at the start of the day. We also allow the trucks to bid on new orders which appear during the day. This "on-the-fly" bidding enhances the transportation capacity (as trucks can fit in new loads if they have sufficient capacity) and speeds up the delivery of loads. Cargo which is left over at the end of the day is treated as initial cargo for the next day. In the experiments, we observe the operations of the companies over a number of days (as actions in one given day can obviously affect the performance in consecutive days).

### *2.2 The Trucks*

The trucks drive round trips in the course of a day. Each individual truck starts from the same initial location each day, and returns to this *homebase* at the end of the day. Alternative distributions of the trucks (e.g., dynamically changing over time) can of course occur in practice. Such distributions, however, significantly complicate the analysis of the model's results, especially over multiple days. Furthermore, the distribution of (the origin of) loads is often quite similar from day to day; population and business centers do not change dramatically overnight. An identical distribution of trucks at the start of each day is thus a reasonable modelling choice. Furthermore, round trips from the same homebase are practical from the perspective of fueling and maintenance of the trucks, the driver, etc. In our simulations, the homebase of a truck coincides with one of the depots. This is in line with the tendency of companies to base their trucks close to the sources of cargo (to maximize operational profits).

Legal restrictions typically limit the number of hours that truck drivers can work per day. There may also be a maximum distance which can be driven in one day. In addition, speed limits need to be taken into account. We set the length of a working day equal to 8 hours, where each hour is treated as a single "round". We also assume (for simplicity) that the trucks travel with a constant "average" speed. A truck can start another round trip on the same day if sufficient time remains to complete the trip the same day.

According to [27], the transportation for roughly 80% of loads is dominantly limited in one dimension. In Europe, this dimension is volume; in the United States this dimension is weight.[5] We hence use a model where we characterize the cargo and the carrying capacity of the trucks in only one dimension, which we (without loss of generality) call weight.

### *2.3 The Routing Problem*

All trucks move on a finite-size square grid. We use a plain grid, where trucks can move in an arbitrary direction as long as they stay on the grid, as well as a Sugarscape-like grid [6], where the grid is a torus (to remove boundary effects) and where trucks follow the gridlines. For both grids, trucks can take the shortest route between two points. We hence do not consider the separate problem of routing through a (congested) network.

We calculate shortest routes for less than 7 intermediate locations by evaluating all possible permutations as this is still feasible computationally. Longer routes are determined by (best) insertion of the pickup and delivery point for the extra cargo [26]. In addition, the segments of the route formed by the insertions are recalculated.[6]

[5] Private communication (with E. Tempelman).

[6] Shortest routes can be calculated very fast by running parallel algorithms on multiple processors [9]. Also of interest are insertion heuristics as discussed in [8] which can be used to generate routes with an increased performance. Such approaches are beyond the scope of our efforts, but may be of interest for real-life applications.

## 3. The Market Mechanism

### 3.1 The Auction Procedure

The depots for cargo serve as market places where the available cargo is auctioned. Each piece of cargo is sold in a separate auction. After a load is sold, it awaits pickup at the depot and is no longer available for other interested parties. The auctions continue until all cargo is sold or until no further bids are placed.

Agents are not allowed to bid for bundles of cargo. Such a combinatorial auction type is not practical in our transportation framework because the number of different bidding options is huge (around 300 pieces of cargo are sometimes offered in the experiments, yielding an intractable number of bundles).[7] We also do not allow agents to participate simultaneously in multiple auctions [17, 4]. Simultaneous auctions are difficult to apply in our case because an agent's valuation for a load is typically dependent on which other loads are won, and at what cost.[8] For this reason, and for the sake of computational feasibility, we allow each agent to only place a bid for at most one load in each round of auctions. Our agents can thus be seen as computationally and rationally bounded (in the sense of [20, 30, 23]), although they repair (some of) their non-optimal local decisions through decommitment.

Each piece of cargo is sold in a separate Vickrey auction. In this auction type, the highest bidder wins the contract but pays the second-highest price.[9] In our model, neither the number of participants nor the submitted bids are revealed by the auctioneer.[10] An attractive property of the one-shot (private-value) Vickrey auction is that it is a (weakly) dominating strategy to bid the true valuation for the good [29, 10].[11] Another attractive property of the Vickrey auction is that a limited amount of communication between the auctioneer and the bidders is required (as opposed to the "open-cry" English auction).

The agents use the following strategy in each bidding round. First, they determine the valuation of each piece of cargo which is offered in an auction. We set the valuation for a load equal to the amount of money which the truck receives when the load is delivered minus the additional costs associated with the new path (which incorporates the pickup and delivery of the new load). The application of more elaborate valuation functions can also be useful for the transportation problem which is considered here. For example, the value of a load can increase when the truck, by transporting the extra load, can move cheaply to an area of the grid with a high density of depots. Another venue of research is in the line of COIN [31, 28], where the aim would be to modify the agents' valuation function to let them more efficiently cooperate as one company. Such refinements of the agent's valuation function form an interesting topic for further studies.

There is obviously an incentive for a company to avoid competition between its own trucks. As part of its strategy, each company therefore makes a pre-selection which determines which agents are allowed to bid for the company in each auction. In this pre-selection phase, the company compares the valuations of the company's agents for the available cargo. The agent with the highest valuation (overall) then bids (its valuation) in the proper auction. This auction is then closed for other agents of the same firm. In this manner, we eliminate the possibility that the no. 2 in the auction, who determines the price, is an agent from the same company. The company then repeats this procedure to select a second agent, which is allowed to bid in another auction, etc. Using this strategy, the agents of a company distribute themselves over a larger set of auctions than would otherwise be the

---

[7] Determining the winners of a combinatorial auction is NP-complete. There has recently been a surge of research in this area, however. A fast algorithm for winner determination has for instance been proposed in [24].

[8] Schillo *et al.* analyze the risk of overbidding when participating in simultaneous auctions and propose a strategy with a constrained number of decommitments (and associated penalties).

[9] Ties are broken at random.

[10] We do not use or reveal sensitive business information in our market mechanism. When extensions of the model are considered (e.g., models where companies receive information about their competitors' actions and behavior) privacy issues should be taken into account.

[11] It is important to note here that the Vickrey auction has some known deficiencies. Furthermore, limitations of the protocol may arise when the Vickrey protocol is used for automated auctions and bidding is done by computational agents [19]. These aspects deserve further attention for future implementations.

case. This, in general, increases the competition between the trucks of different companies.

## 4. The Decommitment Option

Contracts are typically binding in traditional multi-agent negotiation protocols with self interested agents. In [20, 22, 1], a more general protocol with continuous levels of commitment (based upon a monetary penalty method) is proposed and analyzed. The key ingredient of this protocol, the option to break an agreement in favor of a better deal, is maintained in our negotiation model. In the experiments, an agent with a decommitment strategy can (try to) improve its total profits by bidding for a new load with the additional possibility to discard a load to which it committed earlier.

The following decommitment strategy is used by the agents. An agent tries to fit in loads which are transported given that another load is discarded first. If the additional load fits in the truck (after discarding another load), its valuation is discounted with the loss in payment for transport of the decommitted load. The agents bid for the auction with the highest valuation, with or without a possible decommitment of a load.

For simplicity, we only allow agents to discard only one load at the time. Multiple decommitments can be incorporated in a similar way in our model. However, the number of options that should be considered by an agent can become very large in this case, so for practical reasons we do not allow multiple decommitments at the same time. Furthermore, only (contracts for) loads which have been won but are not yet picked up can be discarded, to avoid the possible extra cost of unloading. Decommitment is hence an administrative action.

Trust and reputation are of great importance in the world of (electronic) contract negotiation [12]. A bad track-record can even lead to the shunning of a party in negotiations. How an agent (auctioneer or a client) will change its attitude towards a party which, in the past, has decommitted from a negotiated contract is an open question, however.

In our market mechanism, we circumvent the problem of determining the right penalty levels for different degrees of decommitment. We also avoid the negative effects of decommitment on the reputation of a company. We achieve this by delivering decommitted cargo by a truck from the same company as the truck who decommitted the load (with consideration of delivery constraints). We thus "hide" the process of rejecting deals from the customer: a truck only postpones the transport of decommitted cargo until another truck of the same company becomes available. A company who is using this decommitment strategy retains its reputation and performs according to the contract. This way, it is possible to achieve a gain in performance without, for example, a loss of market share due to undesirable behavior.

We, however, do not allow agents to decommit cargo which must be delivered today (see Section 2.1). Loads for delivery the same day are not decommitted to minimize the chance of a too-late delivery. Furthermore, we have constrained the possible backlog of decommitted loads by only allowing a decommit by an individual truck if the total number of currently unassigned, decommitted loads does not exceed the number of trucks in the company.[12] This approach leads to good results.[13]

Decommitted cargo is once again offered in a Vickrey auction. This auction is, however, only accessible for agents of the company which should deliver the load. The auctions for decommitted cargo thus serve as internal re-sale markets for companies. Note that the bids for decommitted cargo are made in terms of "blue"(i.e., fake) money as the contract for transportation has already been won by the company. We require that new bids for decommitted cargo (in terms of blue money) exceed the original bid costs (in terms of real or "green" money). This rule is used to ensure that the original bidding costs for the decommitted load are covered. Computational experiments show that the performance of the companies deteriorates significantly when this bidding rule is not used. As an alternative, a decommitted load could be offered in a public auction to the competing companies. This procedure results in a net profit for the decommitting company if the original bidding costs for

[12] Alternative, more sophisticated heuristics are a topic of further research.

[13] In the computational experiments, at most 5% of the acquired loads were decommitted. Less than 0.2% of the decommitted loads were delivered too late.

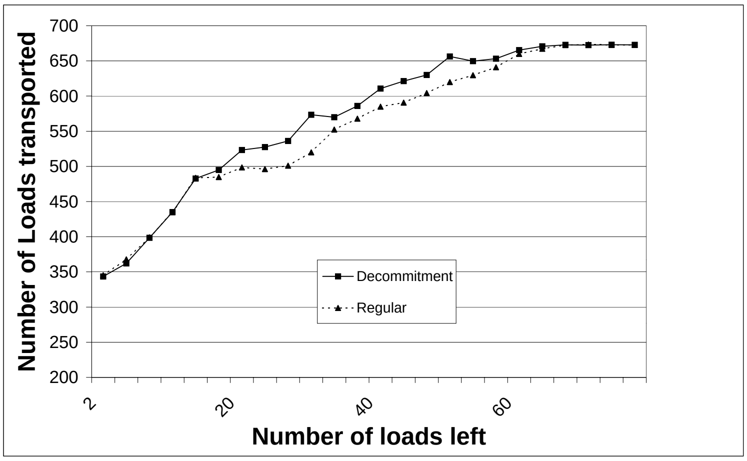


Figure 1: The added value of decommitment for a wide range of production rates.

the decommitted load are covered by the new bid.

The re-sale auctions for decommitted loads are held in parallel with the public auctions to maintain a sufficient degree of competition with the other companies on the running auctions. This way, the competition with the other companies is maintained, while the decommitted load is still won by the most appropriate truck from the decommitting company.

## 5. Results

In this section, we study the performance of a company who uses a decommitment strategy. Appendix A contains the settings of the experiments. In Section 5.1, we first investigate the impact of decommitment in a small transportation model (with only two depots). We then extend the analysis to more complex transportation models. Section 5.2 contains results for a Sugarscape-like model. In this model, the edges of the transportation grid are connected (to suppress boundary effects). In Section 5.3 we consider a finite-size model with a Gaussian distribution of the production. In Sections 5.2–5.5, we further investigate the effect of decommitment for these two models (as a function of the number of depots, the number of trucks per depot, the number of decommitting firms, etc.). Special cases of the models are presented in Sections 5.6 and 5.7.

### *5.1 When Does Decommitment Work?*

We observe in the computational experiments that decommitment of a load occurs predominantly when trucks are close to their maximum capacity. To understand this result, it is useful to first consider two extreme situations: (i) an extreme shortage of available cargo and (ii) an extreme excess of available cargo (relative to the carrying capacity of the trucks). In case of an extreme shortage of loads, a truck will not decommit current loads as it has a large excess capacity: it is more profitable to add a load to a relatively empty truck than to replace one load by another one. In the other extreme (excess of cargo), a new load, which efficiently fills the remaining capacity of the truck, is almost always available. Again, decommitment does not occur in this case, as adding a load which (exactly) fits is more profitable than fine tuning profits at the cost of another load which is dropped.

Figure 1 illustrates the role of decommitment for different production rates. In this figure, we plot the number of transported loads as a function of the number of loads which remain at the depots at the end of the experiment. We generated this figure by varying the production rate at the depots. On the far left, the production rate is low. As a consequence, the available loads are almost all picked up and transported. Note that in this case (shortage of cargo), the effect of decommitment is small (for reasons given above). If the production rate increases, we move to the right in Fig. 1. The (positive) effect of decommitment then increases, as the trucks reach their capacity limits. On the far right in Fig. 1, the production rate is very high. In this case (an excess of cargo), the role of decommitment decreases (as competition for cargo diminishes and trucks can easily fill up their capacity). When the production rate becomes very high, the maximum number of tasks that the trucks are able to handle is reached.

Hence, decommitment is most beneficial when a truck is close to reaching its maximum capacity and there is no ideal load available to fill up the remaining capacity. For a company (with multiple trucks), we expect that the use of a decommitment strategy has a strongly positive effect when a significant fraction of its trucks is in this situation. When the supply of loads approximately matches the carrying capacity of the trucks, the above condition is generally met. We note that there are often economic incentives (in real-life situations) which drive the market to a more balanced situation if supply and demand do not match. Hence, a decommitment strategy can be expected to have an impact in real markets if agents are capable of handling several tasks simultaneously.

In the following sections, we study the role of decommitment in case of more complex transportation models. We keep the number of companies and trucks constant in these simulations and observe the performance of the companies over a number of days (15). For each setting, we determine the maximum number of loads which can be transported without more than a given percentage of too late loads (5%). Following this procedure, we then determine the performance of the companies who use decommitment as part of their strategy (in comparison with companies who only employ the option of regular, binding bidding).

### *5.2 The Sugarscape Model*

We first consider a "Sugarscape-like" grid [6]. Like in Sugarscape, we connect the edges of the grid (to suppress boundary effects). In addition, trucks can only move along the gridlines (i.e., they cannot move diagonally). We place the depots with equal spacings on the grid (the distance is 2 nodes); each depot also has the same production rate. With these assumptions, we obtain a highly symmetric "transportation world".

The performance of the Sugarscape model is summarized in Table 1. We consider two companies

Table 1: Results for a Sugarscape model.

| depots | 4 | 9 | 25 |
|---|---|---|---|
| loads | 940/987 | 1826/1920 | 3794/4197 |
| profit | 91/99 | 420/446 | 585/627 |
| increase | 5/8.7% | 5.1/10.6% | 10.6/7.1% |

in these experiments; only one uses a decommitment strategy. Table 1 shows results for 4, 9, and 25 depots. Note that the grid is already filled densely in case of 25 depots (out of 100 possible locations). Competition between the two companies becomes rather intense in this case. In Table 1, we report the number of transported loads and the profit that is generated (in 1000 monetary units). The performance of the company who does not use a decommitment strategy is given before the slash (/), followed by the performance of the company who does use decommitment. The relative increase of the number of transported loads (before the slash) and the relative increase in profit (after the slash) are also given for the company who uses decommitment. Note the significant positive effect of decommitment in this transportation model (on the order of 5-10%).

*5.3 A Gaussian Distribution Model*

The above-described transportation model is highly stylized. For example, boundary effects are suppressed (by using a toroidal grid), depots are equally spaced, production is uniform, and trucks can only move along the gridlines. We investigate in this section whether the decommitment strategy also works for a transportation model which does not make these limiting assumptions.

This alternative model consists of a plain square grid. The trucks can move in arbitrary directions on the grid, as long as they do not exceed the grid's boundaries. The depots are placed at random locations on the grid. Furthermore, we do no longer assume that production is uniform. Instead, we assume that the spatial production rate follows a Gaussian distribution (with its peak in the center of the grid) and then assign each new load to the nearest depot for transportation[14]. Such a model is representative of a large city or a major business center which is surrounded by smaller cities or businesses [13].

The remainder of this paper discusses results obtained for this model. Figure 2 shows the profits made by a company (with and without the use of a decommitment strategy) as a function of the number of depots on the grid. Note the positive effect of decommitment on a company's profit.

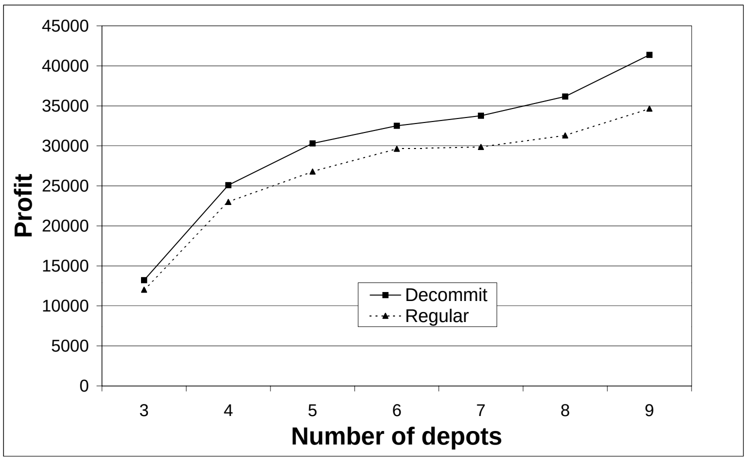


Figure 2: Profits made by a company (with and without decommitment) as a function of the number of depots on the grid.

This effect becomes especially large in case of a densely filled grid. In the experiments, we observed on average one decommitment per truck per day (increasing to a maximum of three per day for a densely filled grid). Results (not shown) for more than two companies show similar trends for the decommitting company. Figure 3 shows that the number of transported loads also increases when a company uses a decommitment strategy.

It is also important to note that the use of decommitment (by one company) can decrease the performance of the non-decommitting companies. This loss can amount to half the increase in profit of the company who uses a decommitment strategy. This effect is of importance when the margin for survival is small (and under-performing companies may be removed from the field).

[14] Production is maximized by maximizing the standard deviation of the Gaussian.

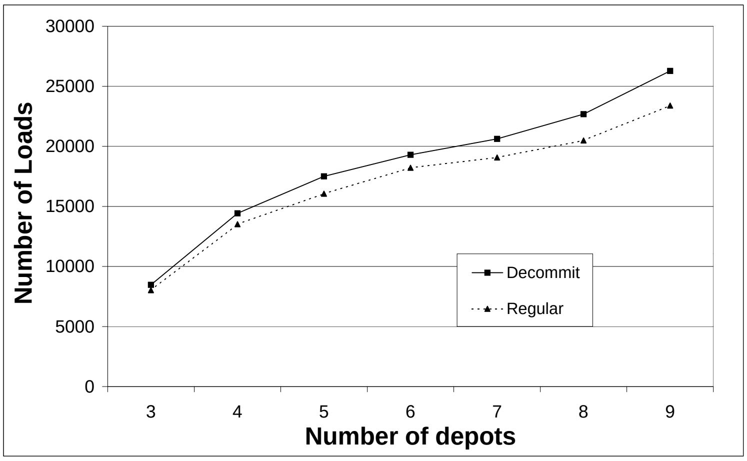


Figure 3: Number of transported loads as a function of the number of depots on the grid. Decommitment has a clear positive effect: the number of carried loads increases significantly.

*5.4 Multiple Trucks Per Depot*

In the previous experiments, only one truck per company was stationed at each depot. Figure 4 shows how a firm's profit depends on the number of trucks per depot (with and without decommitment). Note that the effect of the decommitment strategy clearly increases as the number of trucks on the

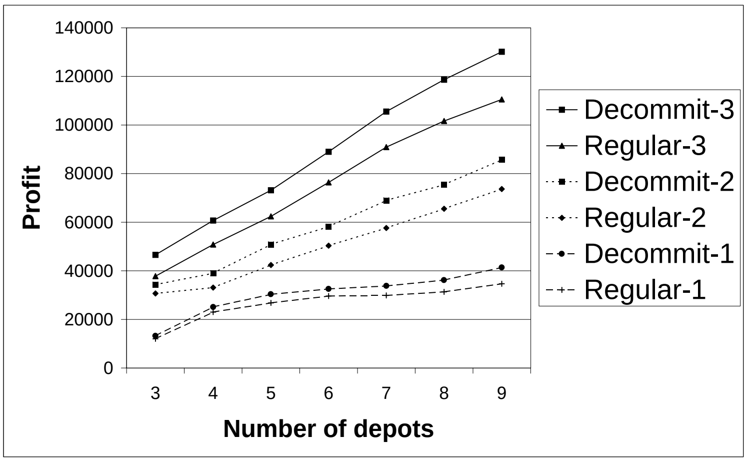


Figure 4: Influence of the number of trucks per depot on the profit made by a company (with and without decommitment). (The number of trucks per depot is indicated in the figure's key.)

grid increases.

*5.5 Multiple Decommitting Firms*

The previous results show that using a decommitment strategy can be beneficial for a company. Stated otherwise, decommitment can give a company a competitive edge in an otherwise symmetric market. Intelligent opponents are, however, not static and countermeasures can be expected if a firm uses a superior strategy [11, 18]. For instance, competitors can also adopt a decommitment strategy once this strategy has proven its usefulness.

We have studied what happens if multiple companies use a decommitment strategy. Experimental results (not shown) show improvements for each decommitting company, as in Figs. 2 and 3. In general, it is thus attractive for a company to use a decommitment strategy. The absolute performance of the decommitment strategy increases slightly as more companies adopt this tactic. However, the relative increase in performance, with respect to the competing companies, drops with a growing number of decommitters.

*5.6 Depot-To-Depot Routing*

An interesting modification of the transportation model is to restrict the delivery of loads to depots. This is a typical scenario in a factory setting (where produced items are inputs for other production processes, similar to supply chain management). Such a scenario can also be relevant in case of international transport (e.g., a layered, holistic setting as discussed in [3] and [25]).

In Fig. 5, we show results with and without depot-to-depot routing. Note that the profits increase

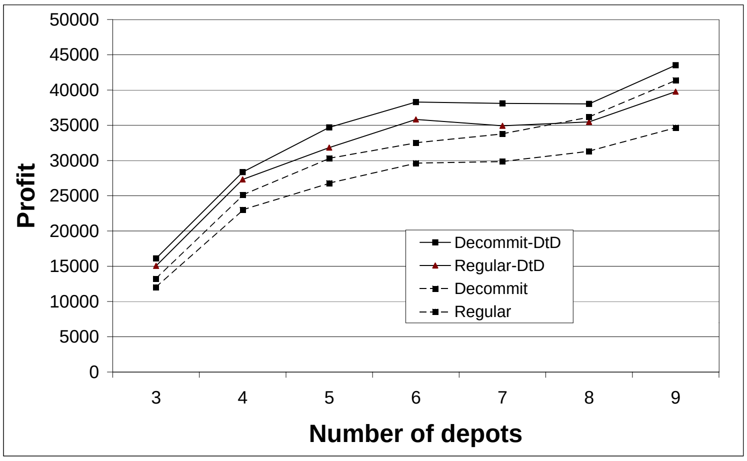


Figure 5: The effect of decommitment in case of depot-to-depot (DtD) routing.

in case of depot-to-depot routing as there is a stronger bundling in the destinations of the goods. Consequently, more efficient routes can be driven. The impact of the decommitment strategy also increases, as the shuffling of loads within an existing route is facilitated.

*5.7 Alternative Settings*

In this final section, we investigate two changes in the transportation model which further increase the impact of the decommitment strategy. We first consider a price function for which the correct prediction of future loads becomes more important (due to a greater difference in the price of individual loads). Secondly, we investigate the impact of restricting the available information to the agents (by restricting the distance over which an agent can shop for loads).

In Fig. 6, we show the strong relative increase in profits when a quadratic price function is used.[15] A similar effect as visible in Fig. 6 occurs if the price for delivery increases sharply as the deadline for

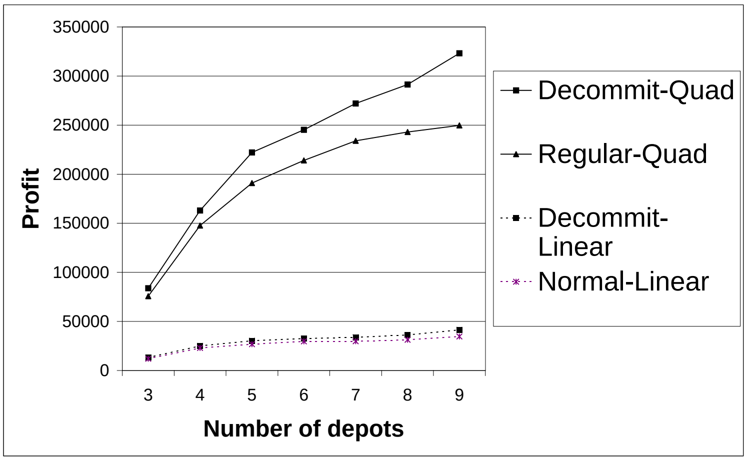


Figure 6: The effect of decommitment in case of nonlinear (quadratic) price functions.

delivery approaches. In both cases there is a strong incentive for agents to correctly anticipate which profitable loads will still appear.

Additional experiments also show that the effect of decommitment increases if the truck s' agents are more "myopic". Truck agents can decide to limit their bidding range due to communication overhead or a lack of computational resources. In Fig. 7, we show the impact of decommitment when an agent only considers loads (for pickup) which are not too far away from its current location.[16] This figure shows that the absolute and relative impact of decommitment increases in this case, as an agent is less able to observe the available loads and thus makes less optimal choices in the course of time which need to be repaired.

## 6. Conclusions and Discussion

We study the use of a decommitment strategy in case of on-line bidding for cargo in a multi-company, multi-depot transportation setting. In our model, a truck can decommit a load in lieu of a more favorable item of cargo. We observe significant increases in profit, especially when the number of companies and depots (and hence the number of trucks) becomes large. Such increases are significant in the competitive market of transport where a 4% profit is considered exceptional. For example, the average profit margin for the Dutch road transport sector (from 1989 to 1999) was only 1.6% [27].[17] Adoption of a decommitment strategy can thus give a company a significant edge.

For specific applications, the added value of decommitment, and the circumstances where it can be applied successfully, should be studied further. However, based upon our computational experiments, we expect that the positive impact of a decommitment strategy increases with the complexity of the operating domain, as it then becomes of greater importance to have the opportunity to roll-back a previous suboptimal decision [21].

[15] The price for a load is $40 + weight(l)^2 + distance(l)$, instead of the linear price function given in Appendix A.
[16] We use an operating range of one quarter of the size of the grid.
[17] Before taxes.

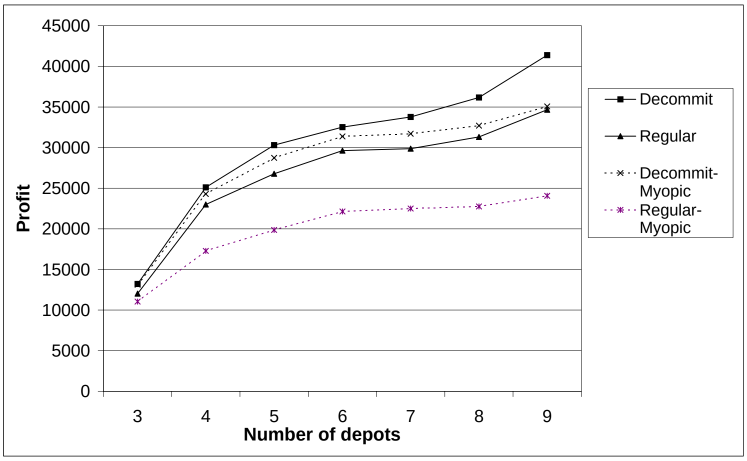


Figure 7: The role of decommitment in case of "myopic" bidding agents.

We also observed that decommitment has the highest impact when an agent is close to its maximum capacity for handling multiple contracts in parallel. With sufficient capacity, it is often more beneficial to add an extra contract than to replace a won contract in favor of a superior offer. Hence, for multi-agent systems where agents are capable of handling several tasks simultaneously, a decommitment strategy can be expected to have its largest impact when the agents are operated at (almost) full capacity.

# References


1. M. Andersson and T. Sandholm. Leveled commitment contracts with myopic and strategic agents. *Journal of Economic Dynamics and Control*, 25:615–640, 2001.
2. S. Bohte, E. Gerding, and H. L. Poutré. Competitive market-based allocation of consumer attention space. In M. Wellman, editor, *Proceedings of the 3rd ACM Conference on Electronic Commerce (EC-01)*, pages 202–206. The ACM Press, 2001.
3. H.-J. Bürckert, K. Fischer, and G. Vierke. Transportation scheduling with holonic MAS - the TELETRUCK approach. In *Third International Conference on Practical Applications of Intelligent Agents and Multiagents (PAAM 98)*, 1998.
4. A. Byde, C. Preist, and N. R. Jennings. Decision procedures for multiple auctions. In *Autonomous Agents & Multiagent Systems*, pages 613–622, part 2. ACM press, 2002.
5. S. Clearwater, editor. *Market based Control of Distributed Systems.* World Scientific Press, Singapore., 1995.
6. J. Epstein and R. Axtell. *Growing Artificial Societies: Social Science From The Bottom Up.* Brookings Institution, 1996.
7. K. Fischer, J. P. Müller, and M. Pischel. Cooperative transportation scheduling, an application domain for DAI. *Journal of Applied Artificial Intelligence, special issue on intelligent agents*, 10(1), 1996.
8. M. Gendreau and J. Potvin. Dynamic vehicle routing and dispatching. In T. Crainic and G. Laporte, editors, *Fleet Management and Logistics*, pages 115–126. Kluwer, 1998.
9. G. Kindervater, J. Lenstra, and M. Savelsbergh. Parallel local search for the time-constrained traveling salesman. *European Journal of Operational Research*, 33:65–81, 1988.
10. P. Klemperer. Auction theory: a guide to the literature. *Journal of economic surveys*, pages 227–286, 1999.
11. E. N. Luttwak. *Strategy, The Logic of War and Peace.* The Belknap Press of Harvard University Press, 1987.
12. L. Mui, A. Halberstadt, and M. Mojdeh. Notions of reputation in multi-agents systems: A review. In *Autonomous Agents & Multiagent Systems.* ACM press, 2002.
13. H. S. Otter, A. van der Veen, and H. J. de Vriend. ABLOoM: Location behaviour, spatial patterns,

and agent-based modelling. *Journal of Artificial Societies and Social Simulation*, 4(4), 2001.

14. D. C. Parkes and L. H. Ungar. An auction-based method for decentralized train scheduling. In *Proceedings 5th International Conference on Autonomous Agents (Agents'01)*, 2001.
15. P.J. 't Hoen, S. Bohte, E. Gerding, and H. La Poutré. Implementation of a competitive market-based allocation of consumer attention space. In W. Walsh, editor, *Proceedings of the 4th Workshop on Agent Mediated Electronic Commerce at AAMAS 2002. to appear in Lecture Notes in Artificial Intelligence from Springer-Verlag volume 2531*. Springer, 2002.
16. A. Poot, G. Kant, and A. Wagelmans. A savings based method for real-life vehicle routing problems. Technical Report EI 9938/A, Erasmus University Rotterdam, Econometric Institute in its series Econometric Institute Reports, 1999.
17. C. Preist, C. Bartolini, and I. Phillips. Algorithm design for agents which participate in multiple simultaneous auctions. In *Proceedings of Agent Mediated E-Commerce, LNAI 2003*, page 139 ff, 2001.
18. J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, Cambridge, MA, USA, 1994.
19. T. Sandholm. Limitations of the vickrey auction in computational multiagent systems. In *2nd International Conference on Multiagent Systems (ICMAS-96)*, pages 299–306. AAAI Press, 1996.
20. T. Sandholm and V. Lesser. Issues in automated negotiation and electronic commerce: Extending the contract net framework. In *Proceedings of the First International Conference on Multiagent Systems.*, pages 328–335, Menlo park, California, 1995. AAAI Press / MIT Press.
21. T. Sandholm and V. Lesser. Leveled-commitment contracting, a backtracking instrument for multiagent systems. *AI Magazine*, Fall 2002:89–100, 2002.
22. T. Sandholm and V. R. Lesser. Advantages of a leveled commitment contracting protocol. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, Portland, OR, 1996.
23. T. Sandholm and V. R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1-2):99–137, 1997.
24. T. Sandholm, T. Suri, S. Gilpin, and A. Levine. CABOB: A fast optimal algorithm for combinatorial auctions. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2001.
25. J. Sauer, T. Freese, and T. Teschke. Towards agent-based multi-site scheduling. In *ECAI 2000 European Conference on Artificial Intelligence 14th Workshop, New Results in Planning, Scheduling and Design (PUK2000)*, 2000.
26. M. Solomon. Algorithms for the vehicule routing and scheduling problems with time window constraints. *Operational Research 35*, pages 254–265, 1987.
27. E. Tempelman. Daf-trucks- where materials make money. In *Second Workshop on Cold and Hot Forging of Light-Weight Materials, Delft, from the ICFG (International Cold Forging Group)*, 2002.
28. K. Tumer, A. Agogino, and D. Wolpert. Learning sequences of actions in collectives of autonomous agents. In *Autonomous Agents & Multiagent Systems*, pages 378–385, part 1. ACM press, 2002.
29. W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.
30. T. Wagner and V. Lesser. Toward ubiquitous satisficing agent control. In *1998 AAAI Symposium on Satisficing Models*, 1998.
31. D. Wolpert, K.Tumer, and J. Frank. Using collective intelligence to route internet traffic. In *Advances in Neural Information Processing Systems-11*, pages 952–958, Denver, 1998.

## 1. Experimental Settings

Unless stated otherwise, we consider the situation in which two companies compete with each other. Only one of the two companies uses a decommitment strategy. For each experiment, we averaged the results over 60 runs and 15 consecutive days. We applied the Wilcoxon test to verify that the reported effects of decommitment are statistically significant. All experiments, unless stated otherwise, use one truck per company per depot (for a fair comparison of the performance of the companies).

We set the distance between two adjacent nodes of the grid equal to 20 km. By default, the grid size is equal to $10 * 10$. In the Sugarscape-like model studied in Section 5.2, the grid size is depending on the number of depots (to ensure a uniform spacing between them). The (average) speed of all trucks is equal to 70 km/hour. The carrying capacity of the trucks is set at 350 units of weight. The weight of the loads is in the range of $[10, 70]$.

The *price* for a load $l$ is set equal to $40 + 2 * weight(l) + distance(l)$, where the transportation *distance* for a load $l$ is the distance from the origin of the load to its destination, and where *weight* denotes the size of the load. This cost function is derived from UPS (see www.ups.nl). UPS uses a constant fee plus a weight-proportional term for its standard packages. We also added a cost per distance. This improved performance (with and without decommitment) as distance then became a stronger issue in bidding. The cost for movement per km is one currency unit.

It is important to note that we set the price for delivery independent of the moment that a load is offered for transport. The price does therefore not increase as the time window for delivery shrinks. This ensures that the results of our experiments are not biased to show good results for decommitment. With increasing prices (close to the delivery time) there is otherwise an incentive for trucks to favor new loads with tight deadlines due to higher profits. In Section 5.7, we study the impact of nonlinear (quadratic) price functions in more detail.