

A Decommitment Strategy in a Competitive Multi-agent Transportation Setting

P.J. 't Hoen¹ and J.A. La Poutré²

¹ Center for Mathematics and Computer Science (CWI)
P.O. Box 94079, 1090 GB Amsterdam, The Netherlands

² TU Eindhoven,
De Lismortel 2, 5600 MB Eindhoven, The Netherlands

Abstract. Decommitment is the action of foregoing of a 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 online 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 are capable of handling multiple contracts simultaneously.

1 Introduction

A recent development is the investigation of application of multi-agent systems (MASs) [7, 12, 25] in the logistics of the transportation sector, a challenging area of application. The transportation sector is very competitive and profit margins are typically low. Furthermore, planning of operations is a computationally intensive task which is classically centrally organized. Such centralized solutions can however quickly become a bottleneck and do not lend themselves well to changing situations, for example incidence management, or exploiting new profitable opportunities. MASs can overcome these challenging difficulties and offer new opportunities for profit by the development of robust, distributed market mechanisms [4, 23, 14]. In this paper, we use as a model online, decentralized auctions where agents bid for cargo in a MAS logistics setting. We study a bidding strategy which is novel for such a large scale setting.

In [20, 1, 21], a leveled commitment protocol for negotiations between agents is presented. Agents have the opportunity to unilaterally *decommit* contracts. That is, they can forgo 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 as agents can escape from premature local minima by adjusting their contracts. In this paper,

decommitment is the possibility of an agent to forgo a previously won contract for a transport in favor of a more profitable load.

We show in a series of computer experiments that significant increase in performance (profit) can be realized by a company with agents who can decommit loads, as opposed to a company with agents that only employ the option of regular, binding bidding. As a necessary precondition for this gain, the experiments show that decommitment is only a clearly superior strategy for an agent close to the limit of its capacity. This is a new, general result for agents capable of handling simultaneous tasks. Furthermore, the increase in performance for our (abstract) model can be seen as a lower bound for expected increased performance in practice. We substantiate this claim through experiments that show that the relative impact of a decommitment strategy increases with the complexity of the world.

The remainder of this paper is organized as follows. Section 2 presents the transportation model that we use in this paper. The market mechanism is described in Section 3. Section 4 details our application of decommitment in a market setting. Section 5 discusses a required precondition for a successful decommitment strategy by an agent capable of handling multiple tasks concurrently. The computer experiments are presented in Section 6. Section 7 contains concluding remarks.

2 The Transportation Model

In this section, we present the transportation model that is used in this paper.³ 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 6, where we show 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 as the uncertainty of possible future events consequently increases. In Section 6.6 we investigate some venues to substantiate this claim.

2.1 Outline

The world is a simple n by n grid. This world is populated by trucks, depots with cargo, and competing companies. The trucks move over the grid and transport cargo picked up at the depots to destinations on the grid. Each truck is coupled with an agent that bids for cargo for its “own” truck.⁴ The trucks are each owned by one of the companies. The performance of a company is measured by the total profits made by its fleet of owned trucks. We consider (for simplicity and to facilitate the analysis of the model’s results) that all companies consist of the same number of (identical) trucks.

³ 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.

⁴ In the text, we sometimes blur the line between the agent and its truck.

2.2 Profits

Poot *et al.* [15] give an extensive list of performance measures for the transportation of cargo found in literature. The indicative performance measures from this list 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 traveled for the made deliveries.

2.3 Loads

Loads for pickup prior to delivery by the trucks are locally aggregated at depots. Such an aggregation procedure is for example used by UPS,⁵ 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, growing 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 [11]. We thus also consider depots as abstractions of important population or business centers. Section 6 presents such a model.

Like most regular mail services (e.g., UPS) and many wholesale suppliers, we employ a model of “next day delivery”. In the simulations, each depot has a number of loads available for transport at the start of the day. Furthermore, new orders can also arrive for transport in the course of the day.

According to [27], transportation is dominantly limited in one dimension for roughly 80% of the loads. In Europe, this dimension is volume; in the United States this dimension is weight.[17] 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.

2.4 The Trucks

The trucks drive round trips in the course of a day. Each individual truck starts from the same initial location each day, to return to this location at the end of the day. Multiple round trips on the same day are allowed as long as sufficient time remains to complete each trip the same 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, a repeating pattern is common as population and business centers do not change dramatically overnight. In our simulations, the truck start their trips at 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).

⁵ See www.ups.com.

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 typical working day of eight hours. We also assume (for simplicity) that the trucks travel with a constant “average” speed.

3 The Market Mechanism

Each piece of cargo is sold in a separate auction. Auctions for loads are held in parallel and can continue over several rounds. The auctions continue until all cargo is sold or until no further bids are placed by the agents in a round. After a load is sold, it awaits pickup at its depot and is no longer available for bidding.

Agents are not allowed to bid for bundles of cargo. Such a combinatorial auction type is as yet beyond the scope of our research 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 for each of which traveling salesman problems have to be solved.).⁶ We also do not allow agents to participate simultaneously in multiple auctions with all implied complications [16, 3, 26]. An agent’s valuation for a load is typically strongly dependent on which other loads are won, and at what cost.⁷ For this reason, and for the sake of computational feasibility, we allow each agent to 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, 22]), although they repair (some of) their non-optimal local decisions through a decommitment strategy (see Section 4).

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.⁸ In our model, neither the number of participants nor the submitted bids are revealed by the auctioneer.⁹ 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, 8].¹⁰ Another attractive property of the Vickrey auction is that

⁶ 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].

⁷ Schillo *et al.* analyzed the risk of over-bidding when participating in simultaneous auctions and propose a strategy with a constrained number of decommitments (and associated penalties).

⁸ Ties are broken at random.

⁹ 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.

¹⁰ 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.

a limited amount of communication between the auctioneer and the bidders is required (as opposed to, for example, 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. The valuation of an added load is equal to added profit for this load (the amount of money which the truck receives when the load is delivered minus the additional costs associated with the new path). The application of more elaborate valuation functions can also be useful. 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 [28, 13], 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 however 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 that 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 case. This, in general, also 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, 23, 1], a more general protocol with continuous levels of commitment is proposed and analyzed. The key ingredient of this protocol is the option to break an agreement, in favor of, hopefully, a better deal, at the possible cost of a prenegotiated penalty. In our the experiments, an agent with a decommitment strategy can improve its immediate profits by bidding for a new load with the additional possibility to discard a load to which it committed earlier. The agent is hence more flexible in the choice of loads to choose to bid on, at the cost of discarding a previously won bid.

Trust and reputation are however of importance in the world of (electronic) contract negotiation [10, 5]. A bad track-record can, for example, lead to the shunning of a party in negotiations. How an auctioneer or a client will change its attitude towards a party which in the past has decommitted from a negotiated contract has to be quantified for specific areas of application. For example, for many bulk transports, a delayed delivery is not too detrimental as another

transporter can easily be found and the transport does not have a tight delivery schedule. This is however not the case for expensive, quickly perishable goods.

In our market mechanism, we circumvent the above quantification issue. We achieve this by delivering decommitted cargo by a truck of the same company as the truck that decommitted the load (with consideration of delivery constraints). We thus “hide” the process of rejecting deals from the customer who offered the load at auction: a truck only postpones the transport of decommitted cargo until another truck of the same company becomes available. A company that uses a decommitment strategy in this fashion retains its reputation and performs according to the contract. For more complex scenario’s (not considered here) where there is no “hiding” the decommitment and where a good cost function is available to quantify the impact of decommitment on trust, we however expect the benefits of a decommitment strategy to increase. The agents then have more options available to optimize their choice of loads.

The “hiding” of the decommitment strategy is achieved by internal reauctioning of loads. 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. Effectively, through a “hidden” decommitment strategy, tasks are redistributed between the agents of one company. Implicitly, the agents renegotiate their concurrent plans.

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 however 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 winning the decommitted load in the original auction are covered. As an alternative, a decommitted load could be offered in a public auction to other companies, i.e. outsourcing, a common practice in the transportation world. The internal resale auctions of decommitted loads are held in parallel with the public auctions as experiments showed that this as a good approach to maintain a sufficient degree of competition with the other companies on the auctions for publicly available loads.

For simplicity and from a computational viewpoint, we allow agents to discard only one load in each round of bidding. Furthermore, only 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.

Furthermore, we do not allow agents to decommit cargo which must be delivered today (see Section 2.1) to minimize the chance of a too-late delivery. Additionally, we have constrained the possible backlog of decommitted loads by only allowing a decommitment by an individual truck if the total number of currently unassigned, decommitted loads does not exceed the number of trucks in the company.¹¹ This approach leads to good results: In the computational experiments less than 0.2% of the decommitted loads were delivered too late.

¹¹ Alternative, more sophisticated heuristics are a topic of research.

Penalties for too-late delivery will hence have to be exorbitant in order to offset the benefits of decommitment presented in Section 6.

5 Conditions for Decommitment

We observe in the computational experiments that decommitment of a load occurs predominantly when trucks are close to filling 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 a load 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 case of a large selection of loads to choose from, a new load, which (closely) fills the remaining capacity of the truck is mostly available. Again, decommitment does not occur as adding a load which fits is more profitable than fine tuning profits at the cost of another load which is dropped.

Figure 1 illustrates the impact of decommitment for a range of offered loads for a single truck to bid on. We plot the number of transported loads as a function of the number of loads presented. On the far left, the number of available loads is low. As a consequence, the available loads are almost all picked up and transported. If the production rate increases, we move to the right in Figure 1. The (positive) effect of decommitment then increases, until the trucks reach their capacity limits. On the far right in Figure 1, the number of offered loads is very high. In this case (an excess of cargo), the added value of decommitment also decreases as the maximum number of tasks that the truck is able to handle can be achieved. Note that for specific scenario's a slightly higher performance can be reached than without the use of a decommitment strategy, but in the limit of available loads (tasks) the added benefit of decommitment will disappear.

Hence, we hypothesize a decommitment strategy is most beneficial when a truck is close to reaching its maximum capacity and has a limited number of extra tasks to choose from. We believe this is a general result for an agent capable of doing multiple tasks in parallel. This hypothesis must be kept in mind when evaluating whether to apply a decommitment strategy.

In our experiments, we observe for a company with multiple trucks, the use of a decommitment strategy only has a strongly positive effect when a significant fraction of its trucks actually decommit loads. When the supply of loads during one day approximately matches the carrying capacity of the trucks, the above condition is met. We note in real-life situations that there are often economic incentives which drive the market to such a balanced situation, if supply and demand do not match. Hence, a decommitment strategy can be expected to have an impact in real markets.

In our simulations, we keep the number of companies and trucks constant when observing the performance of the companies over a number of days. In case of a balanced market, this implies that the amount of cargo which is transported




Fig. 1. The added value of decommitment for a wide range of number of offered loads for one truck. The decommitment strategy only has a strong impact for a subset of the range of number of offered loads.

per day is relatively constant. To this end, we search for an equilibrium “production” of new cargo. In a sense, this is a reversion of the normal market operation. The addition or removal of a truck is however an operation with a large impact. It is not straightforward to formulate criteria in terms of profits which make the addition/removal of a truck an issue, especially over a short time period. Furthermore, differentiation between the various companies in composition makes evaluation of the experiments non trivial. We hence set the production level at a good initial estimate and adapt towards the equilibrium for the strategy used.

We in our experiments achieve this equilibrium in supply and demand by setting the production level of loads to match the approximate carrying capacity of the trucks, while as yet not using a decommitment strategy. An initial number of loads is generated and new loads are produced in the course of the day. The level of production is chosen so as to arrive at a constant number of loads available for transport the next day (within 5% of the initial number of loads available). When this constraint is met, the number of loads and the carrying capacity of the trucks on the grid are in equilibrium over the days of the simulation. With the derived production schedules, we rerun the experiments, but with the additional possibility of a truck to decommit an earlier won bid. The performance of the regular bidding strategy versus the decommitment strategy can then be calculated.

6 Results

In this section, we study the performance of a companies who uses a decommitment strategy relative to companies which do not. Section 6.1 contains results for a Sugarscape-like model. In this model, the edges of the transportation grid are connected (to suppress boundary effects). In Section 6.2 we consider a finite-

size model with a Gaussian distribution of the production. In Sections 6.1–6.4, 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 6.5 and 6.6. Similar results for the above two models were also found using benchmark data from www.opsresearch.com and www.sintef.no/static/am/opti/projects/top/vrp/ for location of depots and scheduling of loads. We feel that our results hence hold for a wide scheme of settings as long as the number of offered loads meets the requirements given in Section 5.

In the experiments, the performance of the bidding strategies is tested over a period of days (15) in order to measure not only immediate performance but also the effect of a bid (or decommitment) over a longer time period. All companies place an equal number of trucks at each depot for fair competition. Unless stated otherwise, we use one truck per depot per company. See also Appendix A for settings of the experiments.

6.1 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 grid lines (i.e., they cannot move diagonally). We place the depots with equal spacing 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 for one company without and with a decommitment strategy is summarized in Table 1 for respectively 4, 9, and 25 depots. We consider two companies in these experiments of which only one can use a decommitment strategy. In Table 1, we report the number of transported loads and the profit that is generated (in 1000 monetary units), with and without use of a decommitment strategy. Note that the grid is already filled densely in case of 25 depots (out of 100 possible locations). Competition between the two companies then becomes intense and profit margins drop as competition in the auctions increases.

6.2 A Gaussian Distribution Model

The Sugarscape transportation model of Section 6.1 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 grid lines. 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

Table 1. Results for a Sugarscape model.

depots	decommitment?	loads	profit
4	no	940	91
4	yes	987	99
increase		5%	8.7%
9	no	1826	420
9	yes	1920	446
increase		5.1%	10.6%
25	no	3704	585
25	yes	4197	627
increase		10.6%	7.1%

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 transportation¹². Such a model is representative of a large city or a major business center which is surrounded by smaller cities or businesses [11]. 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


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

the positive effect of decommitment on a company's profit. 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 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.

¹² Production is maximized by maximizing the standard deviation of the Gaussian.




Fig. 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.

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.

6.3 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 grid increases.

6.4 Multiple Decommitting Firms

The previous results show that the use of 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 counter measures can be expected if a firm uses a superior strategy [9, 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 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, as can be expected, the relative increase in performance with respect to the competing companies drops with a growing number of decommitters.




Fig. 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.

6.5 Depot-to-Depot Routing

An interesting modification of the transportation model is to restrict the destination 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 [2] and [25].




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

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

6.6 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 limiting the distance over which an agent can sample the grid for available loads.

In Figure 6, we show the strong relative increase in profits when a quadratic price function is used.¹³ A similar effect as visible in Figure 6 occurs if the price




Fig. 6. The effect of decommitment in case of linear and nonlinear (quadratic) price functions.

for delivery increases sharply as the deadline for 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 Figure 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.¹⁴ This figure shows that the absolute and relative impact of decommitment increases in

¹³ The price for a load l is $40 + \text{weight}(l)^2 + \text{distance}(l)$, instead of the linear price function given in Appendix A.

¹⁴ We use an operating range of one quarter of the size of the grid.




Fig. 7. The role of decommitment in case of “myopic” bidding agents.

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.

7 Conclusions and Discussion

We study the use of a decommitment strategy in case of on-line bidding for cargo by agents in a multi-company, multi-depot transportation setting. In our model, an agent bidding for a truck can decommit a load in lieu of a more favorable item of cargo. We observe significant increases in profit that scale with the size of operations and uncertainty of future prospects. The observed profit margins are significant in the competitive market of transport where a 4% profit is considered exceptional. For example, the average profit margin before taxes for the Dutch road transport sector (from 1989 to 1999) was only 1.6% [27]. Adoption of a decommitment strategy can thus give a company a significant edge.

For specific applications beyond that of our model and for novel areas, 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 hypothesize 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 sub optimal decision [21].

We also observe 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] 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.
- [3] 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.
- [4] S. Clearwater, editor. *Market based Control of Distributed Systems*. World Scientific Press, Singapore., 1995.
- [5] C. Dellarocas. Goodwill hunting: An economically efficient online feedback mechanism in environments with variable product quality. In W. Walsh, editor, *Proceedings of the 4th Workshop on Agent Mediated Electronic Commerce at AAMAS 2002. Lecture Notes in Artificial Intelligence from Springer-Verlag volume 2531*, pages 238–252. Springer, 2002.
- [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] P. Klemperer. Auction theory: a guide to the literature. *Journal of economic surveys*, pages 227–286, 1999.
- [9] E. N. Luttwak. *Strategy, The Logic of War and Peace*. The Belknap Press of Harvard University Press, 1987.
- [10] L. Mui, A. Halberstadt, and M. Mojdeh. Notions of reputation in multi-agents systems: A review. In *Autonomous Agents & Multiagent Systems*. ACM press, 2002.
- [11] 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.
- [12] 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.
- [13] P.J. 't Hoen and S. Bohte. Collective INtelligence with sequences of actions. In *14th European Conference on Machine Learning*, Lecture Notes in Artificial Intelligence, LNAI 2837. Springer, 2003.
- [14] 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. Lecture Notes in Artificial Intelligence from Springer-Verlag volume 2531*, pages 273–288. Springer, 2002.
- [15] 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.
- [16] 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.
- [17] Private communication with E. Tempelman, author of [26].

- [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. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1-2):99–137, 1997.
- [23] T. Sandholm and V. R. Lesser. Leveled commitment contracts and strategic breach. *Games and Economic Behavior*, 35:212–270, 2001. Special issue on AI and Economics. Early version: *Advantages of a Leveled Commitment Contracting Protocol* in the proceedings of the National Conference on Artificial Intelligence (AAAI), pp. 126–133, Portland, OR, 1996.
- [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] P. Stone, R. Schapire, J. Csirik, M. Littman, and D. McAllester. ATTac-2001: A learning, autonomous bidding agent. In W. Walsh, editor, *Proceedings of the 4th Workshop on Agent Mediated Electronic Commerce at AAMAS 2002. Lecture Notes in Artificial Intelligence from Springer-Verlag volume 2531*, pages 143–160. Springer, 2002.
- [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.

A Experimental Settings

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.

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 6.1, 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 to 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.