A Logistical Multi-Agent System

Tamás Máhr	Jeroen Valk	Peet van Tooren
Almende BV	Delft University of Technology	Almende BV
The Netherlands	The Netherlands	The Netherlands

Logistical applications traditionally aim to decrease the cost of logistical activity. In accordance, researchers developed several techniques to compute economical routes for vehicles generating least costs. However, transportation companies experience a shift of emphasis in their planning nowadays. Due to the highly competitive nature of transportation, customer preferences play a more and more important role in planning. In our multi-agent model we chose to model customers as agents to emphasize their importance in planning. This is different from traditional models, where only vehicles are modeled. To have a package transported, a customer agent has to negotiate a contract with a transport agent. Any negotiation technique can be used to establish the contracts. To enable agents to adapt to changes, contracts can be broken and re-negotiated resulting in new agent plans.
This paper describes a multi-agent system where customer and vehicle agents dynamically change their contracts, hence change their planning. They use negotiation techniques like auctioning and decommitment to manipulate the contracts. Additionally, agents form coalitions to provide more sophisticated services, like multi-modal transportation.

Keywords: Logistics Multi-Agent Systems Decommitment Coalition Formation

1 Introduction

Transportation companies use several computer systems to support their planning and execution activities, but the planning actually is still made mostly by humans. Depending on the size of the problem, this may or may not lead to efficient plans. Since logistical problems have been in the center of interest for decades now, there are many different methods researchers have explored to solve the logistical planning problem (semi-)automatically. In this paper we describe a multi-agent system, where agents can use different negotiation strategies to perform distributed continual planning.

At first, we briefly introduce coordination and planning in multi-agent systems, and the necessity of distributed continual planning. We discuss the possible role of different negotiation and game theory mechanisms in multi-agent planning. Then we introduce our multi-agent model and discuss the implementation of it. Finally we propose experiments that will be carried out using a multi-agent system and in the end conclusions are drawn.

As agent modeling became popular to describe complex systems, more and more researchers applied the multi-agent paradigm to logistical planning. In a multi-agent system every agent has its own goals that it pursues. Every agent is only concerned with its own goals, so system-level behavior emerges from interactions of the different agents. One way to coordinate agents and ensure consistent system behavior is to find a joint plan. However, if agents also carry out their own planning, then the planning activities themselves need to be coordinated.

How we define cooperation between agents depends on what we model as an agent, especially w.r.t. granularity. Guo, Muller and Bauer [1] model different parts of a supply chain by agents to predict delays in delivery. Dorer and Calisti [2] cluster transportation vehicles by areas and make one agent responsible for one area. The finest grained model is provided by BrÃijckert, Fischer and Vierke [3] who employ a holonic multi-agent approach. Trucks, drivers, trailers are modeled by agents that form holons to perform transportation. The degree of modeling granularity corresponds to the degree of decentralization in the model. The finer the model, the more it is decentralized and the bigger the need for coordination is. Since we are particularly interested in coordination our agents represent low level entities like humans, trucks, etc.

Since demand for more and more realistic planning systems appeared, researchers have been looking for a way to distribute planning and to make plan creation and execution concurrent. The need for *distributed, continual planning* was first identified by desJardins et al., and a summary of what had been done up to that was published in [4]. Durfee and Lesser proposed *Partial Global Planning* [5] (generalized by Decker and Lesser [6]), while Clement and Barrett proposed *Shared Activity Coordination* [7] as a solution for the problem. All these techniques solve the coordination problem differently, using partial plans or shared activities. Our solution differs from these in using negotiation mechanisms for coordination.

In our research we are also seeking solutions to facilitate distributed continuous planning. We model customers as well as transportation vehicles as agents that make plans. Agents make their plans according to their preferences, and they cab change their plans to react to changes in their environment (in other agents' plans). This holds not only for vehicle agents, but also for customer agents. To coordinate the changes of agents' plans we use contracts. Customer agents and vehicle agents make contracts to commit to a certain transportation task. Once the contracts are given, agents can safely perform planning within the terms of the contracts, and the resulting plans will be consistent.

There is a wide literature on negotiation, on how to deal with contracts. Contracts are traditionally binding [8][9], which leads to suboptimal task allocation due to unexpected future events. To enable incorporation of future events leveled commitment contracts were introduced by Sandholm and Lesser [10]. They showed that the decommitment possibility increases each agent's expected payoff under very general assumptions. Anderson and Sandholm [11] analyzed different agent behaviors at different penalty levels and found that with proper penalty levels individual myopic agents perform nearly as good as social welfare maximizing agents. Pieter Jan 't Hoen et al. studied decommitment in a logistical setting [12], and they showed that there is a range in transportation capacity (compared to the amount of load to transport) where decommitment raises company profit. These results in the field of negotiations applied to coordination of the planning activities promise solutions for the continual planning problem.

Besides contracts, another way to coordinate multi-agent planning is the formation of coalitions. When coalitions are formed before the planning activities, then it imposes such constraints on the planning activities, that the produced plans will be coordinated without any further effort. The field of forming *coalitions* of cooperating agents originates from *game theory* but it is also studied by multi-agent systems researchers [13, 14, 15, 16]. How the agents form the coalitions or how the payoff is distributed are the key issues in this field. Coalition formation has some particular application in logistical planning. Multi-modal or multi-hop transportation can be realized by coalitions of agents.

The exploitation of negotiation and coalition formation techniques naturally follows from the choice of multi-agent modeling. Agents are expected to behave autonomously and the only way to influence their behavior (their planning) is to communicate (negotiate) with them. This way heterogeneous agents provided (designed, implemented) by different companies can also participate in the planning without sharing (too much) sensitive information.

Thus a multi-agent model of logistics enables us to choose a new approach to the logistical problem. Involving customers into the planning process and giving them a way to directly influence the planning puts the traditional problem into new light. Customers are not modeled anymore by a static set of preferences, but they actively participate in the planning process as agents along with the vehicle agents. They can change their behavior according to vehicle agents' behavior and influence them by their plans.

2 The Model

In this section we introduce our logistical multi-agent model, and discuss how it differs from the usual formal models.

In a multi-agent system customers are represented by agents that are responsible for enforcing solutions that satisfy the customer. In our formal multi-agent model packages or sets of packages associated with one customer order are represented by an agent (*order agent*). These agents negotiate transportation with *vehicle agents*, which results in contracts between order and vehicle agents (see figure 1).

A contract between a vehicle and an order agent defines the commitment of the two agents for the transportation task. High commitment level expresses strong intention to undertake the transportation task, while low commitment level of either party means that this contract is not so important for him. Once contracts are given every vehicle agent has to solve a traditional logistical problem where the set of orders (O) consists only of those orders the given vehicle has contract with. The value of such a set of contract is the value of

$$
\begin{array}{l}
\mathcal{V} : \{v_0, v_1, \ldots, v_n\} \\
\mathcal{O} : \{o_0, o_1, \ldots, o_m\} \\
\mathcal{L} : \{loc_0, \ldots, loc_k\} \\
o_i :< loc_p, loc_d > \\
contract_{ij} :< v_i, o_j, comm_i, comm_j > \\
\hline \\
C : \{contract_{ij} | j = 0..m\} \\
P_{v_i} : L_{v_i} \subset \mathcal{L}, \\
\forall o_j(loc_p, loc_d) : \exists contract_{ij}, \\
loc_p \in L_{v_i} \rightarrow loc_d \in L_{v_i}, \\
min(goal(P))
\end{array}
$$

Figure 1: Logistical multi-agent model

the final solution vehicle agents compute based on the contracts. This value is only comparable to values obtained using the same method to compute the subproblems per vehicle agent. The global goal is to find the set of contracts that defines feasible subproblems that can be solved by the individual vehicle agents.

The formal model does not go further than defining contracts and how they partition the logistical problem into subproblems that are solved by the vehicle agents. Our multi-agent system, however, considers contracts with commitment levels that can change due to changes in the environment of the agents. As an extreme case, agents may decide to break the agreement. In these situations contracts serve as coordination devices through which agents can handle the change. E.g. a truck agent may choose another package instead of a current one because it is more beneficial. Then breaking the contract with the current order agent will trigger a search for another truck and will result in a new contract. Additionally, the contract-breaking agent has to pay some penalty to its partner, which penalty is proportional to the other's commitment level.

Every agent is free to choose its commitment levels for different contracts. If agents experience a shortage in possible contracts, they can choose to raise their commitment levels. This means that their partners will have to pay more penalty in case they want to break the agreement. Thus, raising commitment levels yield in less broken contracts. On the other hand, if agents see lots of opportunities, they can lower their commitment level, allowing more changes in contracts, and generating even more opportunities for other agents.

In our model commitment levels are modeled by penalties. Penalties attached to contracts defines the payment the agents should receive in case the other party breaks the agreement. High penalty expresses strong commitment – the contract is important for the agent –, low penalty implies weak commitment. There are some basic penalty models described in the literature (fixed penalty, penalty as a function of contract value and/or time [11]) that can serve as comparisons.

In our first penalty model agents use one of the existing models, but they learn the proper value of the parameters of the model by consecutive turns. In every turn, they receive a feedback signal and modify their parameter values accordingly. The feedback signal can be based on individual agent performance or on group performance.

In another model agents actively monitor their environment to estimate how 'risky' their environment is. If they experience lots of opportunities, they can easily find new contracts in case a partner agents break their contract. In this case they can lower their penalties, since probably they can compensate their losses with new contracts. On the other hand, when there are few opportunities for new contracts, agents should insist on their current agreements, since it could be difficult to replace a lost contract. In this case agents should raise their penalty levels.

3 Evaluation and Experiments

The formal and actual multi-agent models are developed simultaneously. This section evaluates the two models so far and outlines future experiments for further evaluation.

The primarily goal of the formal model is to provide a clear definition of the problem. The current definition (in section 2) is turned into a Prolog program for verification. The program executes a full search of the problem space on two levels. On the top level, it searches the space of possible contracts, and for every set of contracts it computes the routes and schedules. This algorithm finds the optimal solution by enumerating all solutions. This is, of course, not feasible for any reasonably sized problem, since the number of contracts to check is exponential in the number of customer orders. In the future we will use this model and implementation to study the approximation algorithms that will be developed in connection with the actual multi-agent system.

In the multi-agent model we emphasize the ability of agents to choose their commitment levels. We study adaptive algorithms to set and change the commitment levels to achieve better results in finding the right contracts. In the following we describe the multi-agent platform that is used to implement the agents and some experiments that are considered future work.

The order and vehicle agents are implemented in the Common Hybrid Agent Platform (CHAP) [17] that is developed in the Distributed Engine for Advanced Logistics (DEAL) project. This platform is based on a thread pooling library that detaches agents from threads. This enables the existence of more agents than the number of allowed threads in the operating system, and also makes the execution more efficient. Agents are initialized from XML descriptions, thus they can be initialized from any data source that can be converted o XML (e.g. a database). This flexibility helps to design experiments based on different problem instances like standardized benchmark sets, or the operational database of a commercial (logistic) company.

Once instantiated, truck agents are ready to plan transportation of packages. Order agents are instantiated sequentially in given times. Every order agent organizes an auction to choose the cheapest truck. Once it is chosen, a contract is bound and the chosen vehicle agent modifies its plan to include the new order. This way all the orders are assigned to vehicles one-by-one. Decommitment of an order can occur in various situations. Should any unforeseen incident happen (traffic jam, truck breakdown, etc.), order agents might be forced to chose another vehicle. Upon bidding for an order, a vehicle agent could decide to decommit one of its existing contracts in favor of the new one. Order agent, monitoring the vehicle agents, could find a cheaper truck and decommit their current contract in favor of the other one. Given such agent behaviors, we study how can agents dynamically adapt to new situations, and what effect of this adaptability has on the resulting plans.

In the first two sets of experiments we will study the penalty setting models described in section 2. We expect to see the penalty settings for all agents to follow the changes in their environment. Our goal is to see if adaptability results in better overall performance.

Being able to set penalty (commitment) levels independently for every contract gives rise to an interesting opportunity to model companies as a group of agents. Suppose that a couple of vehicle agents form a group (or coalition) – a company –, and apply special rules inside the group. If an agent tries to break a contract with one of these agents, the penalty it has to pay depends on which other agent it will choose to contract with. If the new contract will involve another agent from the same group (company), a lower (possibly zero) penalty is due. If the new contractor will be an agent not in the company, a higher penalty may apply. This setting enable us to run simulations closer to real life.

The last set of experiments will aim to study how multi-modal transportation can be realized by coalition formation. Transport vehicles of different modes can form coalitions to transport orders which cannot be handled by any individual vehicle. These coalitions are very dynamic by nature and may exist only for the time of the transportation of one order. On the other hand long term agreements are also possible in which case positive synergy with the group-based decommitment strategy is expected.

4 Conclusions

The tough competition in logistics necessitates that planners take customer preferences into account when planning the routes of vehicles. In multi-agent systems this can be modeled by customer agents whose goal is to have the packages transported within conditions that are acceptable to the customer. This partly turns planning into a negotiation process where customer agents and vehicle agents set up transportation contracts.

Introducing contracts in the model enables the use of negotiation techniques, like decommitment, that enables on-line planning. Any event occurring in the system is handled by re-negotiating the contracts, making the system robust with respect to common errors (truck breakdown, traffic jam, etc.). On the other hand the system always maintains feasible plans that can be more or less optimal, but are always executable.

A formal model is developed and implemented to provide a quality measure for the multi-agent system solutions. The centralized Prolog implementation searches through all solutions generating first the contracts, then the routes and schedules for the agents. Since it finds all solutions it is possible to measure the quality of a solution provided by the multi-agent system by comparing it to the solutions provided by the model.

Once the reference model matured, we plan to conduct several experiments to study different negotiation strategies and coalition formation algorithms that can provide robust, good quality solutions for the logistical problem. Strategies to set the decommitment penalty will be studied such as fixed penalty, penalty as a function of contract value and/or time, and especially adaptive penalty setting. Experiments with coalition formation add the interesting perspective of inter company and intra company relations and also provides a way to implement multi-modal transportation.

References

[1] Yutao Guo, Jorg P. Muller, and Bernhard Bauer. A multiagent approach for logistics performance prediction using historical and context information. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1164–1171. IEEE Computer Society, 2004.

[2] K. Dorer and M. Calisti. Agent-based dynamic transport optimization. Technical Report WT-2004-05, Whitestein Technologies, 2004.

[3] Hans-JÃijrgen BÃijrckert, Klaus Fischer, and Gero Vierke. Transportation scheduling with holonic mas - the teletruck approach.

[4] Marie E. DesJardins, Edmund H. Durfee, Charles L. Ortiz, and Michael J. Wolverton. A survey of research in distributed, continual planning. *AI Magazine*, 4:13–22, 2000.

[5] E.H. Durfee and V.R. Lesser. Partial Global Planning: A Coordination Framework for Distributed Hypothesis Formation. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(5):1167–1183, September 1991.

[6] K. Decker and V. Lesser. Designing a Family of Coordination Algorithms. *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pages 73–80, January 1995.

[7] Bradley J. Clement and Anthony C. Barrett. Continual coordination through shared activities. In *AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 57–64, New York, NY, USA, 2003. ACM Press.

[8] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI-93)*, pages 256–262, Wahington DC, 1993.

[9] Sarit Kraus. Contracting tasks in multi-agent environments. Technical Report CS-TR-3254, 1994.

[10] T. Sandholm and V. Lesser. Leveled commitment contracts and strategic breach. *Games and Economic Behaviour*, 35:212–270, 2001.

[11] M. Andersson and T. Sandholm. Leveled commitment contracts with myopic and strategic agents. In *Proceedings of Fifteenth National Conference on Artificial Intelligence*, pages 38–44, 1998.

[12] P.J. 't Hoen and J.A. La Poutré. A decommitment stragtegy in a competitive multi-agent transportation setting. In *Proceedings of AAMAS*, 2003.

[13] V. D. Dang and N. R. Jennings. Generating coalition structures with finite bound from the optimal guarantees. In *3rd International Conference on Autonomous Agents and Multi-Agent Systems*, pages 564–571, New York, USA, 2004.

[14] Matthias Klusch and Andreas Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, 17(3):42–47, 2002.

[15] Onn Shehory and S. Kraus. Coalition formation among autonomous agents: Strategies and complexity. In *From Reaction to Cognition*, number 957, pages 57–72, 1995.

[16] J.P. Kahan and A. Rapoport. *Theories of Coalition Formation*. Lawrence Erlbaum Associates, London, 1984.

[17] Tamas Mahr and Mathijs de Weerdt. Distributed agent platform for advanced logistics. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 155–156, New York, NY, USA, 2005. ACM Press.