COORDINATION THROUGH PLAN REPAIR

Roman van der Krogt Mathijs de Weerdt*

Delft University of Technology
{R.P.J.vanderKrogt,M.M.deWeerdt}@ewi.tudelft.nl

This is an extended abstract. The complete paper has been published in *MICAI 2005: Advances in Artificial Intelligence, LNAI 3789, Springer [6].*

Introduction

In most application domains for distributed artificial intelligence, several autonomous agents (e.g., companies or personal assistants) each have their own goals. They need to plan for these goals, coordinate their actions, deal with uncertainty, and interleave all this with plan execution [2, 4]. To complicate matters even more, many agents are self-interested and require some privacy concerning their plans and the dependencies of actions in their plans on other agents' actions. Our goal is a system in which *self-interested* agents can *(i)* construct and *repair plans*, *(ii) coordinate* their actions, and do so while *(iii)* maintaining their *privacy*.

An auction for dynamic planners

Our idea is to combine a dynamic planning method for each agent with an auction for delegating (sub)tasks. This approach raises a number of questions. If agents do not want to construct a set of global constraints, how can we then guarantee that there will be no deadlocks upon execution of their plans? How can agents integrate the process of auctioning or bidding with the process of plan construction or adaptation? How to decompose a task into subtasks to auction to other agents? And what happens if agents discover that their assumptions about the initial state of their own actions turn out to be invalid? These questions are addressed in the paper, resulting in a multiagent plan repair system for self-interested agents. More specifically, we let agents schedule tasks on which others depend early in their plans to prevent cyclic dependencies, we let them alternatingly plan for one goal and bid on a task, we give them some high-level information about the capabilities of others, and we allow agents to come back on their commitment of accepting a task.

Such a system can be used, for example, to plan for logistics problems. In many logistic domains, some of the transport orders are only known in the nick of time, and, as we all know, traffic is often unpredictable. Consequently, plans need to be revised all the time. Furthermore, often a significant cost reduction can arise when transportation companies coordinate their actions well. For example, a company (agent) may assign a subtask to another either because the other can do it much more efficiently (because it is already in the neighbourhood), or because the first one cannot perform the task at all. Clearly, these companies are also each other's competitors, so they simply refuse to exchange more information than just a request for or an accept of a subtask.

Experimental results

In the paper we give experimental evidence that, in a simplified setup, self-interested agents that have a distinct set of resources can coordinate and repair their plans while only exchanging a very small amount of information. We use an existing plan repair algorithm in a goal-by-goal setting and a basic Vickrey auction, and we show how to prevent cyclic inter-agent dependencies, and how to deal with decommitment.

*This research is supported by the Technology Foundation STW, applied science division of NWO, and the technology program of the Ministry of Economic Affairs.

We study the difference in both plan size and planning time between multiagent plan repair and multiagent replanning from scratch. It turns out that although in some occasions completely replanning leads to slightly shorter plans, it usually takes much longer to reach those. Another advantage of plan repair is that inter-agent dependencies do not change that much compared to replanning from scratch, so there is much less costs for decommitting, which is especially important in real-life applications. Our distributed approach produces bigger plans than central solutions. This can be mainly attributed to our cycle-prevention heuristic, which is often too restrictive. However, it allows us to create valid multiagent plans without exchanging details about the plans, which is very important for self-interested agents.

Discussion

This system for coordinating self-interested agents using propositional plan repair is unique in that we integrate planning and coordination without assuming that the agents are *collaborating*. Agents may even be each other's competitors. Previous work on multiagent planning, although often more advanced in modeling problems realistically (by involving time constraints, minimizing costs, and efficient use of resources), assumes that the agents are collaborative. For example, in the Cougaar system [3] cooperative agents are coordinated by exchanging more and more details of their hierarchical plans until conflicts can be resolved (similar to [7]). And the General Partial Global Planning (GPGP) method [1] describes a framework for distributedly constructing a (partial) global plan to be able to discover all kinds of potential conflicts.

Our future work includes studying the interaction between plan repair and plan *execution* in a multiagent system, as well as trying to improve our goal-by-goal heuristic. Also we would like to test our method in other domains and study the conditions that define when this approach is feasible. Furthermore, we would like to have a method to estimate the costs of external actions. Typically those actions are more expensive than your own actions. If all actions have costs, we can try to optimize costs instead of plan length. Finally, we intend to study the applicability of (de)committing mechanisms [5] in relation to the specific requirements of the problem domain of these self-interested planning agents.

References

[1] K. Decker and J. Li. Coordinating mutually exclusive resources using GPGP. *Autonomous Agents and Multi-Agent Systems*, 3(2):113–157, 2000.

[2] M. DesJardins, E. Durfee, C. Ortiz, and M. Wolverton. A survey of research in distributed, continual planning. *AI Magazine*, 20(4):13–22, 2000.

[3] K. Kleinmann, R. Lazarus, and R. Tomlinson. An infrastructure for adaptive control of multi-agent systems. In *IEEE Int. Conf. on Integration of Knowledge Intensive Multi-Agent Systems*, pages 230–236, 2003.

[4] M. Pollack and J. Horty. There's more to life than making plans: Plan management in dynamic, multi-agent environments. *AI Magazine*, 20(4):71–84, 1999.

[5] T. Sandholm and V. Lesser. Leveled-commitment contracting: a backtracking instrument for multiagent systems. *AI Magazine*, 23(3):89–100, 2002.

[6] Roman van der Krogt and Mathijs de Weerdt. Coordination through plan repair. In *MICAI 2005: Advances in Artificial Intelligence*, volume 3789 of *Lecture Notes on Artificial Intelligence*. Springer, 2005.

[7] F. von Martial. *Coordinating Plans of Autonomous Agents*, volume 610 of *Lecture Notes on AI*. 1992.