Optimal Temporal Decoupling in Multiagent Systems[1]

Léon R. Planken Mathijs M. de Weerdt Cees Witteveen

Faculty EEMCS, Delft University of Technology, The Netherlands

Abstract

When agents need to interact in order to solve some (possibly common) problem, resolving potential conflicts beforehand is often preferred to coordination during execution. Agents may lose some flexibility, but their course of action will be more predictable and often also more efficient. One way to resolve conflicts beforehand is to give extra constraints to each of the agents such that when they all meet these constraints, the resulting execution is conflict-free. A set of constraints that meets this requirement is called a *decoupling* of the original problem; if it also maximizes the social welfare (i.e. the sum of the valuations of all the agents), it is called optimal. Representing interesting multiagent problems as a constraint problem, we show that finding an optimal decoupling is at least as hard as finding a solution for the constraint problem. We therefore focus on a constraint problem that is efficiently solvable, but still very relevant and interesting in the context of multiple agents executing their actions, i.e. the Simple Temporal Problem (STP). Two more technical results, then, are that we resolve the open question whether finding an optimal decoupling of the STP is NP-hard (it is), and if all agents have linear valuation functions, this decoupling problem can be solved efficiently.

1 Introduction

In multiagent systems, a number of autonomous agents perform activities to solve some (possibly common) problem. There are usually many interdependencies between these activities, for example caused by shared goals, scarce resources, or consumer–producer relations. Consequently, these activities have to be *coordinated* in order to ensure a correct solution. In principle, such coordination can be achieved by solving all agents' problems together as one multiagent problem, completely coordinating all their activities in the process. We feel that this is counter to the idea of autonomy and privacy of the agents, thereby removing all flexibility, and is thus a useful approach only in rare cases. The other extreme is to coordinate during execution. For example, semaphores (e.g. traffic lights) can be used to prevent agents from using the same resource at the same time; alternatively, agents can wait for other agents to complete activities they depend upon. Such real-time coordination allows for last-moment changes and thereby offers the agents a significant amount of flexibility. Often, however, when coordination is performed before execution, the process is more efficient (obtaining a social optimum) and results are more predictable; moreover, communication during execution may be prohibitively expensive or simply impossible.

In this paper, we therefore propose to model the problem of coordinating the activities of agents beforehand by finding a *decoupling* [2, 3]. The idea behind such a decoupling is rather simple: additional constraints ensure that individual parts can be solved independently, and all combinations of (local) solutions to these parts can be merged to constitute a (global) solution of the original multiagent problem.

Since a decoupling introduces new constraints, the set of solutions of the original system might be reduced. One commonly used objective is then to choose among alternative decouplings the one that minimizes the loss of flexibility introduced by the additional constraints. For example, Fitoussi and Tennenholtz [1] study how to obtain minimal social laws, which define a set of forbidden strategies such that no strategy can be removed without losing usefulness (i.e. any combination of the allowed strategies is a global

[1]Published in: *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, Van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 789-796.

solution). In this paper we generalize this notion of optimality by optimizing a valuation, specified by agents, for the chosen decoupling. Beside flexibility, such a valuation can express other preferences as well. We then study the problem of finding a decoupling that is as good as possible for all agents. To be precise, we optimize the sum of the valuations of the agents for the chosen decoupling, i.e. the utilitarian social welfare.

In this paper, we use general constraint systems to represent many interesting multiagent problems, be they distributed constraint satisfaction, task/resource allocation, or scheduling problems. We show that finding an arbitrary decoupling is as hard as finding a solution for a constraint system. This implies that finding a decoupling in general is an NP-hard problem. Therefore, finding an optimal decoupling must be at least as hard, making it very unlikely to have an algorithm that always finds an optimal solution efficiently. For this reason, we concentrate on constraint systems where optimal solutions can be found in polynomial time. In addition, in the context of self-interested agents, optimal solutions are important for a second reason; if optimality cannot be guaranteed, agents may want to lie about their valuations, thus changing the objective [4]. Therefore, concentrating on constraint systems where finding a solution is tractable, we study the Simple Temporal Problem (STP). STPs occur as a subproblem of many interesting multiagent problems and some work has already been done on decoupling of STP by Hunsberger [3]. However, whether an optimal decoupling of STP can be found efficiently was still an open problem. We show that for general valuation functions, finding an optimal decoupling is NP-hard, while for linear functions this can be done in polynomial time. Then, it becomes possible to construct a Groves mechanism that ensures that the dominant strategy for an agent is to truthfully report its private information, viz. its valuation function.

2 Discussion and related work

In our paper, we presented the following contributions. First, we have proved that finding decouplings for arbitrary distributed constraint systems is as hard as solving the constraint system centrally. From this, it immediately follows that *optimal* decoupling is at least as hard, and that if we aim for efficient algorithms for optimal decoupling, we should concentrate our attention on problems that are efficiently solvable. We then solved the open problem put forward by Hunsberger regarding the complexity of finding an Optimal Temporal Decoupling (OTD) for the STP [3]. Surprisingly, OTD was shown to be intractable even for some quadratic objectives, but, fortunately, we were able to show that OTD is efficiently solvable when the objectives are linear, obtaining a socially optimal decoupling instead of a local optimum such as found by Hunsberger's algorithm. Finally, we have illustrated how a tractable optimal algorithm for decoupling can be combined with known results from mechanism design, obtaining a Groves mechanism that resolves coordination conflicts for rational, self-interested agents.

We believe that the mechanism design view on solving coordination problems between self-interested agents put forward in this paper can be an inspiration to continue work on social laws and decoupling. In particular, we expect significant contributions by studying other special cases like OTD that are polynomially solvable. Additionally, we leave for future work the study of decoupling for multiagent problems that are even harder than constraint satisfaction, such as e.g. multiagent planning, including nontemporal or mixed problem settings. In general, we cannot expect to be able to optimally decouple such intractable problems. This thus encourages us to consider recent work on for example second-chance or maximal-in-its-range mechanisms [4] to be able to use approximation algorithms in the context of self-interested agents.

References

[1] D. Fitoussi and M. Tennenholtz. Choosing social laws for multi-agent systems: Minimality and simplicity. *Artificial Intelligence*, 119(1-2):61–102, 2000.

[2] L. Hunsberger. Algorithms for a temporal decoupling problem in multi-agent planning. In *Proc. of the National Conference on Artificial Intelligence*, pages 468–475. AAAI Press, 2002.

[3] L. Hunsberger. *Group decision making and temporal reasoning*. PhD thesis, Harvard University Cambridge, Massachusetts, 2002.

[4] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. *Journal of AI Research*, 29:19–47, 2007.