Planning under Uncertainty for Coordinating Infrastructural Maintenance [1]

Joris Scharpff Matthijs T.J. Spaan Leentje Volker
Mathijs M. de Weerdt
{j.c.d.scharpff, m.t.j.spaan, l.volker, m.m.deweerdt}@tudelft.nl

Delft University of Technology, The Netherlands

Abstract

Scheduling of infrastructural maintenance poses a complex multi-agent problem. Commonly a central authority is responsible for the quality and throughput of the infrastructure, while the actual maintenance is performed by multiple self-interested contractors. Not only does the central authority have to (economically) incentivise agents to consider quality and throughput, it is also burdened with the coordination of agents' activities on the network with contingent activity durations.

We introduce a coordination method that combines planning under uncertainty and dynamic mechanism design to coordinate agents on a network level. We apply this method on maintenance planning scenarios obtained through accurate modelling of the problem domain. To the best of our knowledge, this is the first application of dynamic mechanism design on a real-world problem. Finally, we validate the feasibility of our method through experimental evaluation and identify current open challenges for both the planning and scheduling as well as the mechanism design communities.

1 Introduction

Planning maintenance of (public) infrastructures, for example a highway network, is a challenging real-world problem. Although maintenance improves the infrastructure quality, it is accompanied by throughput reductions. For instance in the case of a highway network, performing maintenance requires at least a partial closure of the particular highway that is being serviced. Given the enormous economic impact of the resulting delays, it is essential to plan maintenance such that it only minimally disrupts the traffic flow. Especially when multiple agents are working concurrently on the network, they cause the throughput to degrade rapidly and hence coordination between agents is a necessity.

There are several major obstacles that make the coordination problem complex. First of all maintenance activities commonly face delays. Secondly, as the agents are self-interested, they focus only on profit and have to be explicitly rewarded (or fined) for their contribution to other objectives. Finally, we also need to consider the cost that agents have for doing the maintenance, otherwise they will not be interested in participating in the coordination mechanism at all. Difficult here is that maintenance costs are private and thus need to be elicited truthfully in order to be able to optimise the overall utility. If not, agents may increase their utility by lying about their costs, leading to inefficient plans.

2 Coordination Mechanism

We introduce a coordination method that combines the dynamic Vickrey-Clarke-Groves (VCG) [1, 2] mechanism with stochastic planning using Markov Decision Processes (MDP's) [4]. Informally, the mechanism first asks agents to reveal their maintenance costs after which an optimal joint policy is computed using these costs and penalties/rewards for quality and throughput. The policy dictates the

[1]This is an extended abstract, the full paper can be found in Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS), pages 425 – 433, AAAI Press, 2013

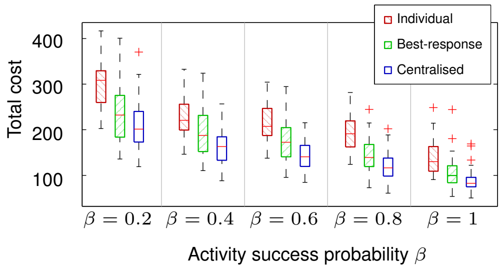

(a) Total cost using different planning approaches for the activities with failure problems (lower is better).

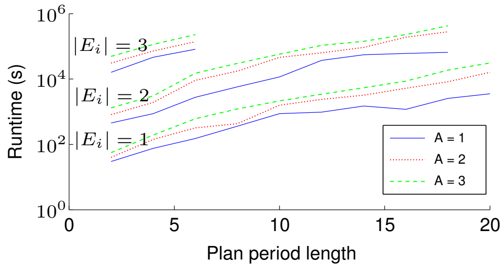

(b) Runtimes of best-response planning for portfolio management for various road set sizes $|E_i|$, activities per road $|A|$ and plan length $|T|$ (log scale).

activities to perform each state, therefore enabling us to account for all assessed maintenance risks. Truthfulness is achieved through the payments defined by the VCG mechanism.

Finding optimal solutions is achieved through a two-stage MDP encoding of the joint planning problem which we solve using the SPUDD solver [3]. This stage separation, enforced through tokens and penalties of infinite cost, enables us to first determine for each agent the activity to start or continue and subsequently ‘execute’ the combined set of planning choices. Effectively this allows us to reduce the size of the (exponential) joint MDP action set to a linear set size, at the cost of having a larger state space. However, as MDP solvers are typically geared toward exploitation of exponential state spaces, this results in a significantly better performance in terms of run time.

3 Evaluation

We have performed experiments on two scenarios that are both real-world problems obtained through interviews and discussions with domain experts [5]. We compared three planning approaches for both scenarios: one where agents plan individually, one where agents iteratively improve the joint plan (best response) and a centralised approach. The first scenario, dubbed *activities with failures*, concern unit-time activities that can fail and has applications in for example supply chains or factory scheduling (Figure 1a). The second scenario, *portfolio management*, serves to illustrate the complexity of scheduling a pre-allocated set of activities with possible delays (Figure 1b).

In conclusion, we present a first solution for coordination of multiple self-interested agents that are related through their planning choices. We argue that the theory of dynamic mechanism design and stochastic planning can be used for coordination but our experiments show that further research is required in order to apply it on real-world instances. Applications exist in several other domains such as bandwidth allocation or smart power grids, highlighting the need for a practical solutions.

Acknowledgements This research is part of the Dynamic Contracting in Infrastructures project supported by Next Generation Infrastructures and Almende BV. Matthijs Spaan is funded by the FP7 Marie Curie Actions Individual Fellowship #275217 (FP7-PEOPLE-2010-IEF).

References

[1] D. Bergemann and J. Valimaki. Efficient dynamic auctions. *Cowles Foundation Discussion Papers*, 2006.

[2] R. Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proc. of 9th ACM conference on Electronic commerce*, pages 220–229. ACM, 2008.

[3] J. Hoey, R. St-Aubin, A. Hu, and C. Boutilier. Spudd: Stochastic planning using decision diagrams. In *Proc. of Conf. on Uncertainty in Artificial Intelligence*, pages 279–288, 1999.

[4] M. L. Puterman. *Markov Decision Processes—Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, 1994.

[5] Leentje Volker, Joris Scharpff, Mathijs de Weerdt, Paulien Herder, and S Smith. Designing a dynamic network based approach for asset management activities. In *Proc. of the 28th Annual Conf. of the Association of Researchers in Construction Management (ARCOM)*, pages 655–664, 2012.