

GoCo: Planning Expressive Commitment Protocols (Article Abstract)*

Neil Yorke-Smith^{1,2}[0000-0002-1814-3515]

¹ Delft University of Technology, The Netherlands

² American University of Beirut, Lebanon
n.yorke-smith@tudelft.nl

Abstract. The full article published in *Autonomous Agents and Multi-Agent Systems* addresses the challenge of planning coordinated activities for a set of autonomous agents, who coordinate according to social commitments among themselves. The article develops a multi-agent plan in the form of a commitment protocol that allows the agents to coordinate in a flexible manner, retaining their autonomy in terms of the goals they adopt so long as their actions adhere to the commitments they have made. The setting considered allows parameterized predicates with probabilistic uncertainty over action outcomes. For this rich setting, the article contributes the first practical means to derive protocol enactments which maximise expected utility from the point of view of one agent. To do so, the article employs hierarchical planning techniques to check whether a commitment protocol can be enacted efficiently, and to generate protocol enactments under a variety of conditions. The approach is illustrated on a real-world healthcare scenario.

1 Coordination in Socio-Technical Systems

The full article by Meneguzzi et al [3] recognizes that agents in an open system, although autonomous, may be interdependent in subtle ways. The physical, social, or organisational environment in which they interact can be complex. We need ways to accommodate the environment while supporting decoupling of the agents' internals from their interaction, thus facilitating the composition of multi-agent systems. The notion of a socio-technical system (STS) provides a basis for representing such interactions between agents in the context of an organisation, and respecting technical artefacts required for the effective operation of the organisation. Activities in an STS are characterized by a combination of the goals sought by and the specified interactions between the agents.

An agent's goals relate to its social commitments at two levels: at design time, the collaborative STS design process produces a commitment protocol; and at run time, agents consider the protocol and their respective goals and make their

* This is an extended abstract of [3].

decisions accordingly. This view leads us to two main challenges: (1) Where do the protocols come from? and (2) How can an agent gain assurance that its goals are indeed achievable when the STS is instantiated?

2 Hierarchical Planning Approach

The contribution of the article is to address the foregoing challenges by moving from a purely qualitative generated specification to one that allows the verification of the properties of actual instantiations of the commitment protocols. While previous approaches have addressed the challenge of whether a protocol can be verified in the sense that participants can enforce it [1] (i.e., observe that agents comply with it), we address a more fundamental question of whether a protocol can be enacted in the environment for which it was designed. Specifically, the problem addressed by the article is the design of commitment protocols for agents acting in uncertain environments, and the validation of the feasibility of the commitment protocol from a centralized perspective.

The planning-based approach developed in the article provides a computational mechanism to reason about a number of properties of commitment protocols and their enactments. An enactment is considered to be an instantiation of a protocol defined in terms of goals and commitments that corresponds to the full hierarchical decomposition of an Hierarchical Task Network (HTN) task using a method library [2]. An enactment is optimal if it maximises the expected utility across all alternative enactments. First, the approach can identify whether or not a commitment protocol is compatible with the agent's goals, i.e., whether there is at least one enactment of the protocol that achieves one or more individual goals. Second, the approach can quickly generate a suboptimal enactment to prove a protocol is feasible as well as generate all possible enactments, if needed. Third, the approach can provide a quantitative assessment of the utilities of each possible enactment of the commitment protocol. Fourth, the approach can use the exhaustive generation of enactments to select among them regarding their utility to one or more of the participating agents.

Acknowledgements This work is supported in part by the Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) process numbers 306864/2013-4 and 482156/2013-9, the American University of Beirut University Research Board grant number 102853, and the Olayan School of Business grant number OFFER_C1.2013-2014.

References

1. Baldoni, M., Baroglio, C., Chopra, A.K., Singh, M.P.: Composing and verifying commitment-based multiagent protocols. In: Proc. of IJCAI'15. pp. 10–17 (2015)
2. Erol, K., Hendler, J., Nau, D.S.: HTN planning: Complexity and expressivity. In: Proc. of AAAI'94. vol. 2, pp. 1123–1128 (1994)
3. Meneguzzi, F., Magnaguagno, M.C., Singh, M.P., Telang, P.R., Yorke-Smith, N.: GoCo: Planning expressive commitment protocols. *Autonomous Agents and Multi-Agent Systems* **32**(4), 459–502 (2018)