

Towards Planning Uncertain Commitment Protocols

(Extended Abstract)

Felipe Meneguzzi
Pontifical Catholic University
of Rio Grande do Sul
Porto Alegre, Brazil
felipe.meneguzzi@pucrs.br

Pankaj R. Telang
North Carolina State
University
Raleigh, NC, U.S.A.
ptelang@gmail.com

Neil Yorke-Smith
Olayan School of Business
American University of Beirut
Beirut, Lebanon
nysmith@aub.edu.lb

ABSTRACT

In the context of a business process modeled by commitments, agents enact a protocol by carrying out goals that service their part of commitments. In a competitive or even in a cooperative setting, an agent does not know for sure that its partners will successfully act on their part of the commitments. We introduce uncertainty into a successful recent approach of planning first-order commitment protocols. Probabilities reflect a semantics of the belief of an agent about the successful completion of tasks by other agents within the protocol, capturing notions of trust. We take a deterministic Hierarchical Task Network (HTN) planner, introduce probabilities into the task networks, and derive a protocol enactment which maximizes expected utility from the point of view of one agent. We illustrate our approach on a business scenario in e-commerce.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligent—*Multi-agent systems*; I.2.8 [Artificial Intelligence]: Plan execution, formation, and generation—*Planning under uncertainty*

Keywords

Multi-agent planning; Uncertainty; Commitments; Hierarchical Task Networks

1. MOTIVATION AND CONTEXT

Commitments help model interactions in multiagent systems in a computationally realizable yet high-level manner—without compromising the autonomy and heterogeneity of the member agents [4]. The problem of multiagent commitment protocol planning is important because agents can enact business protocols through commitments [1]. Recent work shows how to combine commitments with goals and apply planning methods to enable agents to not only determine their actions [3], but also synthesize business protocols through commitments in a deterministic setting.

Consider a purchase scenario within an e-commerce marketplace. The buyer has a goal of buying a book and the

seller a goal of getting payment. At its simplest, the buyer commits to the seller to paying if the seller provides the book, and the seller commits to the buyer to provide the book if the buyer pays. Possible plans that satisfy the goals and commitments of both agents can be determined: (1) the seller provides the book, then the buyer pays, and (2) the buyer pays the seller, then the seller provides the book. If prudent and rational agents, the buyer and seller will plan their actions bearing in mind that the other agent in the protocol might fail to satisfy some of its commitments.

However, when planning about commitments in real scenarios, simply discovering that there are possible realizations for commitment protocols might not be sufficient to ensure that these plans are feasible, nor (if feasible) optimal, e.g., in cost. In our example, the payment may fail due to insufficient funds or due to automated rejection on security grounds by the credit card company. Prior work in this context does not allow agents to reason about the likelihood of other agents' actions in the protocols. In this context we assume a mediating planning agent to whom the participants agree to provide information, such as (some of) their goals, commitments, actions, and utilities. This agent will produce a plan—in our approach, a contingent plan—that maximizes a notion of global utility for the set of agents.

In order to address this challenge, our starting point is the planning formalism based on Hierarchical Task Networks (HTNs) developed by Telang, Meneguzzi and Singh [5, 3] that implements a semantics of commitments and goals [6]. This formalization allows for the encoding of complex, parameterized commitments allowing agents to reason about the realizations of their commitments including features such as numeric values in commitment parameters (e.g., to reason about monetary costs, payments) and realization patterns. It does not, however, reason about stochastic action outcomes and overall utility maximization.

We take an established algorithm for non-deterministic HTN planning [2] and modify it to reason about probabilities and utilities, in order to derive the plans which maximize expected utility from the point of view of one agent. The contingent plan produced contains all contingencies, even those of low probability. However, we can choose to eliminate some contingencies from the plan to reduce its size. Note that our approach accommodates a first-order representation of the actions as well as explicit modelling of time, which is problematic for MDP-based approaches.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015)*, Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.

Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

2. ILLUSTRATIVE SCENARIO

In an e-commerce marketplace akin to eBay, buyer and seller agents interact and make transactions, mediated by the marketplace (mp) agent. We describe a scenario from mp’s point of view. Buyer wants a copy of a classic Dicken’s book; seller is a book dealer who has a copy. The goals of the agents are: $G(\text{buyer}, \text{book})$: buyer’s goal for the book, $G(\text{mp}, \text{fee})$: mp’s goal to receive transaction fees, and $G(\text{seller}, \text{paid})$: seller’s goal to be paid for the book. The commitments the agents make according to the norms of the marketplace are: $C(\text{seller}, \text{mp}, \text{paid}, \text{fee})$: seller commits to paying a fee to mp if the seller is paid by the buyer, $C(\text{mp}, \text{seller}, \text{fee}, \text{paid})$: mp commits to ensuring that the buyer pays the seller if the seller pays a fee to mp, $C(\text{buyer}, \text{seller}, \text{book}, \text{paid})$: buyer commits to the seller to paying if the seller provides the book, and $C(\text{seller}, \text{buyer}, \text{paid}, \text{book})$: seller commits to the buyer to provide the book if the buyer pays. Agent mp believes the following estimates of the chance of success of key actions (by the agent indicated): $\text{sendBook}(\text{seller})$: 70%, $\text{sendPay}(\text{buyer})$: 60%, $\text{payFee}(\text{seller})$: 80%. Further, agent mp’s beliefs about the utility of outcomes are as shown in Table 1. In the Table, $-\text{action}$ denotes that action fails.

The valid plans and their utilities are: $\text{sendPay}, \text{sendBook}, \text{payFee}$ ($0.6 * 1 + 0.7 * 2 + 0.8 * 3 = 4.4$), $\text{sendPay}, \text{payFee}, \text{sendBook}$ ($0.6 * 1 + 0.8 * 2 + 0.7 * 3 = 4.3$), and $\text{sendBook}, \text{sendPay}, \text{payFee}$ ($0.7 * 1 + 0.6 * 2 + 0.8 * 3 = 4.3$). The protocol of: the buyer paying the seller, the seller sending the book, and the seller paying a fee to mp, thus maximizes expected utility according to the marketplace. Our primary objective is to develop an algorithm that computationally generates such a plan. Note that if the planning neglected the probabilities, then mp might propose an inferior protocol.

3. ALGORITHM SKETCH

In order to reason about the goals and commitments semantics of Meneguzzi et al. [3] within a stochastic environment, we require an HTN planner that takes into account the possibility that each operator has multiple possible outcomes, each of which has an associated probability. As no extant HTN planner available to us has this capability, we implemented a probabilistic HTN planner inspired by the ND-SHOP2 algorithm [2], which, however, considers non-uniform operator outcome probabilities. Our implementation is based on the deterministic HTN planner Pyhop. We modified Pyhop to reason about ND-SHOP2 domains, and further modified it to consider domains with specific operator probabilities as discussed in the previous section. We call the resulting planning algorithm ND-Pyhop.

ND-Pyhop takes as input the initial state and task net-

Plan State	Utility
$\text{sendPay}, \text{sendBook}, \text{payFee}$	3
$\text{sendPay}, \text{sendBook}, -\text{payFee}$	2
$\text{sendPay}, -\text{sendBook}, \text{payFee}$	1
$\text{sendPay}, -\text{sendBook}, -\text{payFee}$	1
$-\text{sendPay}, \text{sendBook}, \text{payFee}$	2
$-\text{sendPay}, \text{sendBook}, -\text{payFee}$	1
$-\text{sendPay}, -\text{sendBook}, \text{payFee}$	1
$-\text{sendPay}, -\text{sendBook}, -\text{payFee}$	0

Table 1: Utility of outcomes, according to agent mp

work as well as the planning domain. ND-Pyhop optimizes the utilities calculated for the plans generated by a FORWARDSEARCH function, which performs a top-down left-to right expansion of the HTN, yielding a plan whenever the task network has no more tasks to be expanded. Otherwise, the algorithm either expands a primitive task by executing the operator associated to it and considering all possible outcomes of the stochastic operator, calculating probabilities and utilities in the process, or adds the subtasks of relevant methods that can expand a non-primitive task.

4. SUMMARY AND FUTURE WORK

The contribution sketched in this extended abstract is a principled approach to protocol planning under stochastic action uncertainty, in a highly expressive social planning setting. We define a mapping between protocol planning domains in stochastic environments into an extension of the HTN planning formalism, and develop a planning algorithm by which a mediating planning agent can choose the highest utility protocol realizations in the form of contingency plans. By developing a new stochastic HTN planning algorithm, our contribution has potential applications in other settings, for example, optimizing commitment protocols using different criteria than utility, which we plan to investigate in future work. Our implementation, like HTN planning in general, has exponential complexity. Nevertheless, our algorithm allows for quick anytime generation of a commitment protocol enactment in case an agent needs to quickly generate a realizable MAS commitment protocol, while an agent waits for an acceptable MAS commitment protocol.

Acknowledgments. We gratefully thank those who shared their code with us. Felipe thanks CNPq for support within grant numbers 306864/2013-4 under the PQ fellowship and 482156/2013-9 under the Universal project programs. Neil acknowledges support of the AUB URB grant number 102853 and the OSB grant OFFER_C1_2013_2014.

REFERENCES

- [1] A. K. Chopra, F. Dalpiaz, P. Giorgini, and J. Mylopoulos. Reasoning about agents and protocols via goals and commitments. In *Proc. of AAMAS’10*, pages 457–464, 2010.
- [2] U. Kuter and D. S. Nau. Forward-chaining planning in nondeterministic domains. In *Proc. of AAAI’04*, pages 513–518, 2004.
- [3] F. Meneguzzi, P. R. Telang, and M. P. Singh. A first-order formalization of commitments and goals for planning. In *Proc. of AAAI’13*, pages 1–8, 2013.
- [4] M. P. Singh. Commitments in multiagent systems. In *The Goals of Cognition: Essays in Honor of Cristiano Castelfranchi*, pages 601–626. College Publications, London, UK, 2012.
- [5] P. R. Telang, F. Meneguzzi, and M. P. Singh. Hierarchical planning about goals and commitments. In *Proc. of AAMAS’13*, pages 761–768, 2013.
- [6] P. R. Telang, N. Yorke-Smith, and M. P. Singh. Relating goal and commitment semantics. In *Proc. of ProMAS’11*, volume 7217 of *LNCS*, pages 22–37. Springer, 2012.