Multi-Agent Coordination under Uncertain Communication

Nikhil Bhargava
Massachusetts Institute of Technology
Cambridge, MA
nkb@mit.edu

ABSTRACT

Multi-agent coordination problems have many layers of complexity that make them interesting research problems. What uniquely separates multi-agent coordination problems from large single-agent planning and execution problems is that in multi-agent problems, different agents have different understandings of global state. The way that diverging beliefs are reconciled is through communication. My thesis builds a real-time executive that manages and reasons about communication between agents and provides the supplementing theory to prove that these algorithms are practical and efficient.

KEYWORDS

Single and multiagent planning and scheduling; Coordination and control models for multiagent systems

ACM Reference Format:
Nikhil Bhargava. 2019. Multi-Agent Coordination under Uncertain Communication. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), Montreal, Canada, May 13–17, 2019,* IFAAMAS, 3 pages.

1 INTRODUCTION

The overall goal of my thesis is to build a real-time temporal executive that is capable of dispatching and reactively adapting multi-agent plans. A architecture diagram for a typical real-time executive can be seen in Figure 1a. The input to an executive is a partially ungrounded temporal plan which the executive is responsible for grounding and dispatching. In situations where there is no environmental uncertainty, an executive can immediately ground the plan and dispatch actions in an open loop fashion. In reality, however, there is non-determinism in the real world, and in order to handle that variability, an executive needs to receive feedback from the real world describing when events happen and incorporate that information to refine its plan and adapt its subsequent dispatches.

In order to build a temporal executive capable of managing multiple agents, my thesis focuses on, first, considering multi-agent plans, second, adding communication requirements to dispatches to enforce the eventual coalescing of diverging global state beliefs, and, third, examining how to move from a world with immediate state updates to ones in which state updates are delayed and noisy. The proposed architecture diagram for an improved multi-agent executive can be found in Figure 1b.

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 13–17, 2019, Montreal, Canada.

2 BACKGROUND

In the world of temporal planning, Simple Temporal Networks (STNs) provide a straightforward way to express constraints [8], but they assume that it is possible to schedule every last event, which is too strong of an assumption in practice. Simple Temporal Networks with Uncertainty (STNUs) [10] augment STNs to provide a way to model this kind of uncertainty by adding the ability to model actions whose durations are uncontrollable and, as such, are a useful starting point when considering how to model multi-agent coordination. From the perspective of one particular agent, an activity with uncertain duration can be used to represent the fact that the time window for another agent's activity is known a priori while the amount of time spent on that activity by that agent is not.

Definition 1. STNU [10]
An STNU is a 4-tuple $\langle X_b, X_e, R_c, R_g \rangle$ where:

- X_b is the set of activated timepoints
- X_e is the set of received timepoints
- R_g is the set of contingent constraints of the form $l_g \leq e_i - b_j \leq u_g$, where $e_i \in X_e$, $b_j \in X_b$
- R_c is the set of requirement constraints of the form $l_c \leq x_i - x_j \leq u_c$, where $x_i, x_j \in X_b \cup X_e$

When discussing the scheduling in the context of an STNU, we usually refer to its *controllability.* To characterize controllability in a multi-agent context, my prior work defines the notion of *delay controllability* [1]. Delay controllability generalizes strong and dynamic controllability in STNUs [10] and uses a delay function to parameterize what information the scheduling actor has when making decisions. The relevant definitions are reproduced below.

Definition 2. Delay Function
A delay function, $\gamma : X_e \to \mathbb{R}^+ \cup \{\infty\}$, takes a received timepoint and outputs the maximum amount of time that may pass after its assignment before its value is observed and the underlying uncertainty is resolved.

Definition 3. Delay Controllability An STNU S is delay controllable with respect to a delay function γ if it is possible to dynamically construct a schedule when learning about each received event x_e only after $\gamma(x_e)$ time has passed.

This definition allows us to model some of the most basic forms of communication in a temporal planning setting by modeling communication as a single event displaced in time.

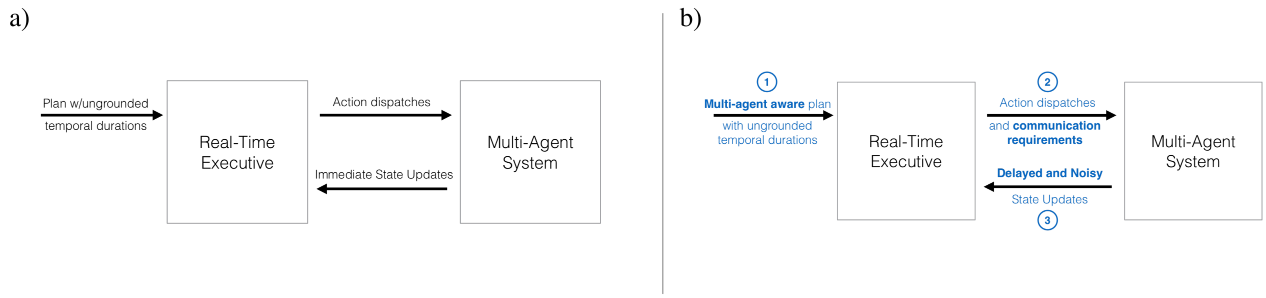

Figure 1: a) A typical architecture diagram for a temporal-based executive. b) An architecture for a temporal executive enabled by my research which incorporates asynchronous and uncertain communication.

3 RESEARCH DIRECTION

3.1 Alternative Models

My first area of focus considers whether other temporal models beyond the STNU are worth considering. STNUs are desirable systems for modeling because unlike many other types of networks, controllability can be determined in polynomial time. However, when used to model multiagent scenarios, they have quite low fidelity; it is either assumed that other agents can be controlled completely or that they act completely randomly. No uncontrolled but coordinated execution is permitted, unlike those considered in models like the Multi-agent STN [6] and Multi-agent STNU [7].

Recent works, including one awaiting publication at AAMAS, considers alternative temporal network models and considers how costly it is to model different forms of controllability across them [4]. By zeroing in on the subtle differences between the networks, my aim is to express a set of systematic rules for modelers to choose between different models, and by examining the jump in algorithmic complexity across these models, we get a strong justification for augmenting executives to use delay controllability.

3.2 Managing Communication Costs

Delay controllability is useful in that it lets us determine whether or not a particular fixed multi-agent strategy for communicating is sufficient to guarantee that all goals can be satisfied. In practice, however, the communication strategy used by a group of agents is flexible. Obviously, communicating about each event as soon as possible will increase the likelihood that the resulting temporal network is controllable, but agents may have preferences about when to stage their communication. For example, I might not want to update my friends about my evening plans in the middle of a meeting with my advisor, but if I wait until afterwards, I can still respect everyone's scheduling constraints. If we model this preference as a cost function over delay functions, we can begin to ask not whether a community strategy exists but rather what the optimal one might be for a given temporal network.

My previous work on managing communication costs described algorithms for deriving an optimal communication strategy that were optimal as well as ones that were suboptimal but fast in practice [2]. While the suboptimal approaches can be polynomially bad in theory, in practice they provide solutions that are quite close to optimal at a dramatic increase in speed.

Future work in this area will pay closer attention to the real world dynamics of communication. During execution, communication events may be missed entirely or may come in sooner than expected. In the interest of minimizing unnecessary communication overhead and maximizing user preference, it becomes important to update the upcoming communication strategy in order to avoid being overly conservative. While this problem can be solved by recalculating a communication strategy from scratch each time, smarter strategies are likely to more efficiently achieve the desired end.

3.3 Variable Delays

The models that have been considered so far are powerful in that they describe communication, but they assume a perfect transmission of information whenever it occurs. In practice, there may be noise in the signal itself that may make it difficult to know with certainty when the original event happened.

My previous work on *variable-delay controllability* characterizes these types of communication events and remarkably shows that controllability can be completely determined in polynomial time [3]. This is in contrast to networks that are similar but slightly more expressive, like the POSTNU [5], for which sound and efficient algorithms for determining controllability exist but for which no efficient sound and complete algorithms are known to exist.

In the next few months, I intend to augment this approach by applying a risk-bounding approach to this theory. In reality, the probability distributions associated with a communication event can have arbitrarily long tails. But the act of checking controllability ascribes undue weight to highly unlikely events. By applying a risk-bounding approach as seen in use by others in temporal networks [9, 11], we can expect to provide strategies that are highly likely to succeed and more likely to be used in practice.

REFERENCES

[1] Nikhil Bhargava, Christian Muise, Tiago Vaquero, and Brian Williams. 2018. Delay Controllability: Multi-Agent Coordination under Communication Delay. In *DSpace@MIT*. http://hdl.handle.net/1721.1/113340

[2] Nikhil Bhargava, Christian Muise, Tiago Vaquero, and Brian Williams. 2018. Managing Communication Costs under Temporal Uncertainty. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 84–90.

https://doi.org/10.24963/ijcai.2018/12
[3] Nikhil Bhargava, Christian Muise, and Brian Williams. 2018. Variable-Delay Controllability. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*. International Joint Conferences on Artificial Intelligence Organization, 4660–4666. https://doi.org/10.24963/ijcai.2018/648
[4] Nikhil Bhargava and Brian Williams. 2019. Complexity bounds for the controllability of temporal networks with conditions, disjunctions, and uncertainty. *Artificial Intelligence* (2019).
[5] Arthur Bit-Monnot, Malik Ghallab, and Félix Ingrand. 2016. Which contingent events to observe for the dynamic controllability of a plan. In *International Joint Conference on Artificial Intelligence (IJCAI-16)*.
[6] James Boerkoel and Edmund H Durfee. 2013. Distributed reasoning for multiagent simple temporal problems. (2013).
[7] Guillaume Casanova, Cédric Pralet, Charles Lesire, and Thierry Vidal. 2016. Solving Dynamic Controllability Problem of Multi-Agent Plans with Uncertainty Using Mixed Integer Linear Programming.. In *ECAI*. 930–938.
[8] Rina Dechter, Itay Meiri, and Judea Pearl. 1991. Temporal constraint networks. *Artificial intelligence* 49, 1-3 (1991), 61–95.
[9] Cheng Fang, Peng Yu, and Brian C Williams. 2014. Chance-constrained probabilistic simple temporal problems. (2014).
[10] Thierry Vidal and Helene Fargier. 1999. Handling contingency in temporal constraint networks: from consistency to controllabilities. *Journal of Experimental & Theoretical Artificial Intelligence* 11, 1 (1999), 23–45.
[11] Peng Yu, Brian C. Williams, Cheng Fang, Jing Cui, and Haslum. 2017. Resolving Over-Constrained Temporal Problems with Uncertainty through Conflict-Directed Relaxation.