The Twenty-Fourth AAAI/SIGAI Doctoral Consortium (AAAI-19)

Multi-Agent Coordination under Uncertain Communication

Nikhil Bhargava
Massachusetts Institute of Technology
nkb@mit.edu

Abstract

Multi-agent coordination is not a simple problem. While significant research has gone into computing plans efficiently and managing competing preferences, the execution of multi-agent plans can still fail even when the plan space is small and agent goals are universally aligned. The reason for this difficulty is that in order to guarantee successful execution of a plan, effective multi-agent coordination requires communication to ensure that all actors have accurate beliefs about the state of the world. My thesis will focus on the problem of characterizing, modeling, and providing efficient algorithms for addressing planning and execution when there agents cannot maintain perfect communication.

Introduction

Multi-agent coordination is not a simple problem. While significant research has gone into computing plans efficiently and managing competing preferences, the execution of multi-agent plans can still fail even when the plan space is small and agent goals are universally aligned. The reason for this difficulty is that in order to guarantee successful execution of a plan, effective multi-agent coordination requires communication to ensure that all actors have accurate beliefs about the state of the world. The communication of information in and of itself can be quite complex and nuanced, but despite the importance of understanding communication during execution, little work has been done to understand how different patterns of communication affect the robustness of generated plans.

The aim of my thesis is to characterize the effects of communication in temporal plans and to provide algorithms that can efficiently construct schedules for temporal plans and provide a robust execution strategy for those plans subject to communication constraints. By incorporating communication into temporal planning, the resulting plans will both have higher utility and lower risks of failure.

Background

In temporal planning, Simple Temporal Networks (STNs) provide a straightforward way to express constraints (Dechter, Meiri, and Pearl 1991) but assume it is possible to schedule every last event, which is too strong of an assumption in practice. Simple Temporal Networks with Uncertainty (STNUs) (Vidal and Fargier 1999) 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 (Vidal and Fargier 1999)
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* (Bhargava et al. 2018a). Delay controllability generalizes strong and dynamic controllability in STNUs (Vidal and Fargier 1999) 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, $\gamma : X_e \rightarrow \mathbb{R}^+ \cup \{\infty\}$, 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.

Research Direction

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 communication strategy exists but rather what the optimal one might be.

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 (Bhargava et al. 2018b). While the suboptimal approaches can be polynomially bad in theory, in practice we see that Least Cost Resolution Search provides solutions that are quite close to optimal at a dramatic increase in speed.

Future work in this area will focus on the real world dynamics of communication. During execution, communication events may be missed entirely, or they may come in sooner than expected or may arrive for events for which we expected no information. In the interest of minimizing unnecessary communication overhead and maximizing user preference, it becomes important in these cases to also update the upcoming communication strategy in order to avoid committing to an overly conservative approach. While this problem can be solved by recalculating a communication strategy from scratch each time, smarter strategies are likely to be much more efficient at achieving the desired end.

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, 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 (Bhargava, Muise, and Williams 2018). This is in contrast to networks that are similar but slightly more expressive, like the POSTNU (Bit-Monnot, Ghallab, and Ingrand 2016), 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 (Yu et al. 2017), we can expect to provide strategies that are highly likely to succeed and more likely to be used in practice.

Alternative Models

My final body of work considers whether temporal models beyond the STNU are worth considering. STNUs are desirable 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 (Boerkoel and Durfee 2013) and Multi-agent STNU (Casanova et al. 2016).

Work that is currently accepted at AIJ considers alternative temporal network models and considers how costly it is to model different forms of controllability across them. 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.

References

Bhargava, N.; Muise, C.; Vaquero, T.; and Williams, B. 2018a. Delay controllability: Multi-agent coordination under communication delay. In *DSpace@MIT*.

Bhargava, N.; Muise, C.; Vaquero, T.; and Williams, B. 2018b. Managing communication costs under temporal uncertainty. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, 84–90. International Joint Conferences on Artificial Intelligence Organization.

Bhargava, N.; Muise, C.; and Williams, B. 2018. Variable-delay controllability. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, 4660–4666. International Joint Conferences on Artificial Intelligence Organization.

Bit-Monnot, A.; Ghallab, M.; and Ingrand, F. 2016. Which contingent events to observe for the dynamic controllability of a plan. In *International Joint Conference on Artificial Intelligence (IJCAI-16)*.

Boerkoel, J., and Durfee, E. H. 2013. Distributed reasoning for multiagent simple temporal problems.

Casanova, G.; Pralet, C.; Lesire, C.; and Vidal, T. 2016. Solving dynamic controllability problem of multi-agent plans with uncertainty using mixed integer linear programming. In *ECAI*, 930–938.

Dechter, R.; Meiri, I.; and Pearl, J. 1991. Temporal constraint networks. *Artificial intelligence* 49(1-3):61–95.

Vidal, T., and Fargier, H. 1999. Handling contingency in temporal constraint networks: from consistency to controllabilities. *Journal of Experimental & Theoretical Artificial Intelligence* 11(1):23–45.

Yu, P.; Williams, B. C.; Fang, C.; Cui, J.; and Haslum. 2017. Resolving over-constrained temporal problems with uncertainty through conflict-directed relaxation.