MIP-Nets: Enabling Information Sharing in Loosely-Coupled Teamwork

Ofra Amir, Barbara J. Grosz and **Krzysztof Z. Gajos**
Harvard School of Engineering and Applied Sciences
{oamir,grosz,kgajos}@seas.harvard.edu

Abstract

People collaborate in carrying out such complex activities as treating patients, co-authoring documents and developing software. While technologies such as Dropbox and Github enable groups to work in a distributed manner, coordinating team members' individual activities poses significant challenges. In this paper, we formalize the problem of "information sharing in loosely-coupled extended-duration teamwork". We develop a new representation, Mutual Influence Potential Networks (MIP-Nets), to model collaboration patterns and dependencies among activities, and an algorithm, MIP-DOI, that uses this representation to reason about information sharing.

Introduction

Distributed teamwork is becoming increasingly prevalent as technology enables groups of people distributed over vast distances, with few opportunities for synchronous interaction, to work together on complex activities extended in time. Technologies such as Google Drive, Dropbox and Github enable teams to share work artifacts remotely and asynchronously. The coordination of their activities remains a challenge, however, because these technologies do not have capabilities for focusing people's attention on the actions taken by others that matter most to their own activities.

Coordination is especially challenging in teamwork that is extended over a long time period and that is loosely-coupled in nature (Amir et al. 2015; Pinelle and Gutwin 2006). While loosely-coupled teamwork allows collaborators to focus on their individual tasks and reduces coordination needs, it also makes it harder to identify the dependencies and conflicts between collaborators' activities (Hutchins 1995).

This paper defines the problem of Information Sharing in Loosely-coupled Extended-duration Teamwork (ISLET) and presents new methods for addressing it. Successful solutions to the ISLET problem need to identify and share with team members information that is *relevant* to their activities and to do so within a *limited communication budget* so as to not overwhelm them with too much information.

We extend prior research in multi-agent systems, which has developed a variety of methods for reasoning about information sharing (Kamar, Gal, and Grosz 2009; Wu, Zilberstein, and Chen 2011). These approaches rely on a *complete plan knowledge assumption*. They use a model and knowledge of the team's plans or policies to compute the value of information. Although some approaches assume only incomplete knowledge of agents' plans (Barrett and Stone 2015), these approaches still assume a known planning domain. In real-world distributed human teamwork settings, such plan models are not available for modeling (Amir et al. 2013).

Our approach avoids the complete plan knowledge assumption by utilizing the extended duration of the teamwork to learn team members' role allocation and dependencies between activities. This information is implicitly represented in a new representation, MIP-Net. We define a MIP-DOI algorithm, which uses the MIP-Net to reason about information sharing decisions. We evaluate our approach in simulation, showing that it is capable of learning collaboration patterns and sharing relevant information with team members.

The ISLET Problem

An *ISLET problem setting* comprises the following: (1) P: a set of collaborating partners. The set can change over time with partners joining or leaving the team; (2) O: a set of objects that partners interact with. The set can change over time as a result of partners' actions; (3) A: the set of act-types $\{ADD, MOD, DEL\}$ for adding, modifying or deleting objects, and (4) S: A set of interaction sessions of partners. A session $s(p, t, (\langle a_1, o_1\rangle, ..., \langle a_{|s|}, o_{|s|}\rangle))$ is defined by a triple: the partner acting , the time of the session, and a set of pairs of act-types and the objects they operate on ($\langle a_i, o_i\rangle$). For brevity, we denote a session recorded at time t as s_t. Figure 1(a) shows an example session.

The *ISLET problem* is to determine a set of l objects $O_{share} \subset O$ to inform a partner $p \in P$ about when p starts session s_t, given sessions s_1 to s_{t-1} and the identity of p. The communication budget l restricts the amount of information that can be shared, reflecting the need not to overwhelm partners with too much information. The set O_{share} should include objects that are *relevant* to the partner.

Mutual Influence Potential Networks

MIP-Nets represent interactions between partners and objects and dependencies between different objects. Partners and objects are represented by nodes. A particular partner $\mathbf{p}$ and a particular object $\mathbf{o}$ are represented by nodes n_p and n_o, respectively. The nodes n_p and n_o are connected by an edge if $\mathbf{p}$ performed an action on $\mathbf{o}$. The edge weight corresponds to the extent of the interaction. Similarly, n_o and

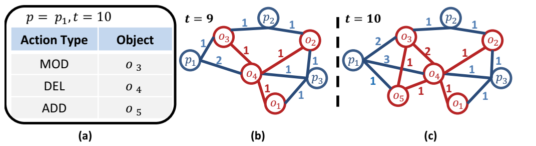

Figure 1: (a) An interaction session s_{10}; (b) The MIP-Net after sessions $s_1 - s_9$. Partner nodes and edges connecting partners and objects are shown in blue. Object nodes and edges connecting them are shown in red. Numbers on edges represent the edge weights; (c) The updated MIP-Net after session s_{10}.

$n_{o'}$ are connected by weighted edges based on the frequency of the objects they represent being modified in the same sessions. Figure 1(b) shows a sample MIP-Net. MIP-Nets are constructed and revised over time based on partners' sessions. At the end of each session s_t, the MIP-Net update procedure increments the weights of edges connecting n_p with nodes representing objects that were modified in the session and the weights of edges connecting object nodes representing objects that the partner interacted with in the same session.

To reason about information sharing, the MIP-DOI algorithm uses the MIP-Net to quantify the relevance to $\mathbf{p}$ of modifications to object $\mathbf{o}$. We use the concept of *Degree-Of-Interest* (DOI) (Furnas 1986), which into consideration the a priori importance of an item and its proximity to the users' focus of attention. In our formulation of DOI, we consider two different nodes as representing $\mathbf{p}$'s focus of attention: the node representing the partner (n_p), and the node representing the object that the partner acts on at the beginning of a session, denoted o_f for "focus object". DOI is computed by:

$$\mathrm{DOI}(o \mid p, o_f) = \alpha \cdot API(n_o) + \beta_1 \cdot D(n_o, n_p) + \beta_2 \cdot D(n_o, n_{o_f})$$

The distance values $D(n_o, n_p)$ and $D(n_o, n_{o_f})$ take into account the weight on the edge connecting the two nodes and their shared neighbors. The *a priori* importance ($API(n)$) is the degree of n. MIP-DOI computes $DOI(o \mid p, o_f)$ for each $o \in O$ and chooses the l objects with the highest DOI to share with $\mathbf{p}$.

Empirical Evaluation

We evaluated the MIP-DOI algorithm in a simulation environment which uses a collaborative graph coloring problem: the partners (P) need to color a graph such that no two neighboring vertices are assigned the same color. The graph vertices are the set of objects (O). In each turn, a partner $\mathbf{p}$ colors a set of k vertices, denoted O_{modify}, which includes a focus object o_f. Before choosing their colors, $\mathbf{p}$ receives information about the color of a set of l vertices (O_{share}) from an information sharing agent. $\mathbf{p}$ knows the graph structure (vertices and edges), but only knows the colors of vertices that were shared with it, and assume colors have not changed until being informed otherwise. The information sharing agent only knows about the existence of vertices that partners interacted with, but does not have information about edges (reflecting the lack of knowledge about the task structure in ISLET settings). Its aim is to share relevant information: an object $\mathbf{o}$ is relevant, if there is an edge connecting $\mathbf{o}$ to an object in O_{modify}, as such information can directly affect $\mathbf{p}$'s actions.

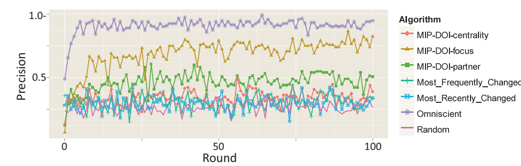

Figure 2: Average precision by round.

We evaluated agents using 3 configurations of MIP-DOI: MIP-DOI-centrality, which only considers objects' centrality ($\alpha = 1$); MIP-DOI-partner, which only considers objects' proximity to n_p ($\beta_1 = 1$), and MIP-DOI-focus, which only considers objects' proximity to o_f ($\beta_2 = 1$). We compared these MIP-DOI variations with 4 baselines: an Omniscient agent which has access to the graph structure and chooses objects in proportion to their distance from o_f, an agent that shares the most frequently changed objects; an agent that shares the most recently changed objects, and an agent that chooses objects randomly. We note that the algorithms' goal is not to solve the graph coloring problem, but rather to share relevant information with partners. Thus, they are incomparable to distributed CSP algorithms.

Figure 2 shows the precision obtained by each of the algorithms with $l = 3$, averaged over 10 different graph instances with 5 runs for each graph instance. As can be seen in the figure, MIP-DOI-focus significantly outperforms all other uninformed algorithms, i.e., the baseline algorithms except "omniscient" which has access to the true graph structure. Over time, its performance becomes close to that of the omniscient algorithm as more information about role allocation and task structure is accumulated in the MIP-Net. The other MIP-DOI configurations also outperform all uninformed baselines.

Acknowledgments. The research was supported in part by a grant from the Nuance Foundation.

References

Amir, O.; Grosz, B. J.; Law, E.; and Stern, R. 2013. Collaborative health care plan support. In *AAMAS*.

Amir, O.; Grosz, B. J.; Gajos, K. Z.; Swenson, S. M.; and Sanders, L. M. 2015. From care plans to care coordination: Opportunities for computer support of teamwork in complex healthcare. In *CHI'15*.

Barrett, S., and Stone, P. 2015. Cooperating with unknown teammates in complex domains: A robot soccer case study of ad hoc teamwork. In *AAAI*.

Furnas, G. W. 1986. *Generalized fisheye views*, volume 17. ACM.

Hutchins, E. 1995. *Cognition in the Wild*. MIT press.

Kamar, E.; Gal, Y.; and Grosz, B. 2009. Incorporating helpful behavior into collaborative planning. In *AAMAS*, 875–882.

Pinelle, D., and Gutwin, C. 2006. Loose coupling and healthcare organizations: deployment strategies for groupware. *Computer Supported Cooperative Work (CSCW)* 15(5-6):537–572.

Wu, F.; Zilberstein, S.; and Chen, X. 2011. Online planning for multi-agent systems with bounded communication. *Artificial Intelligence* 175(2):487–511.