DEMO: MA−RADAR – A Mixed-Reality Interface for Collaborative Decision Making

Sailik Sengupta* and **Tathagata Chakraborti*** and **Subbarao Kambhampati**

Arizona State University, Tempe, AZ 85281 USA

{ sailiks, tchakra2, rao } @ asu.edu

Abstract

There has been a lot of recent interest towards adapting automated planning techniques for the role of decision support for human decision makers in the loop. A unique challenge in such settings is the presence of multiple humans collaborating during the planning process which not only requires algorithmic advances to handle issues such as diverging mental models and the establishment of common ground, but also the development of user interfaces that can facilitate the distributed decision making process among the human planners. We posit that recent advances in augmented reality (AR) technology is uniquely positioned to serve this need. For example, a mixed-reality workspace can be ideal for curating information towards the particular needs (e.g. explanations) of the individual decision makers. In this demonstration we present MA−RADAR – a mixed-reality interface for collaborative decision making with multiple humans in the loop.

In this demonstration, we extend the decision support system RADAR (Sengupta et al. 2017) to handle multiple human decision makers (Kim and Shah 2017; Chakraborti et al. 2017b) in the loop –

- We show how AR provides an effective medium of augmenting the shared GUI with private information (as studied in planning literature (Brafman and Domshlak 2008)) – thus the same plan will appear differently on the shared GUI than in the mixed-reality view where the private actions will be coupled with the public plan; and
- We show how AR can reduce irrelevant information on the screen by porting them into the mixed-reality view. Such situations can occur, for example, when one user asks for an explanation, which the others may not require and thus should not appear on the shared GUI.

The Fire-Fighting Domain The fire fighting domain involves extinguishing fire at a particular location. It requires two commanders (henceforth referred to as Comm-I and Comm-II) to come up with a plan or course of action (CoA) which involves coordination with the police, medical and transport authorities. Each commander might have a personalized model of this domain, which (1) may have certain actions that are *private* to them, i.e. unknown to the other commanders; and (2) incorrect ideas about the actual domain, for example, an incorrect action definition (according to the model of the decision support agent). A detailed description of the domain is available in (Sengupta et al. 2017).

*Authors marked with * contributed equally.

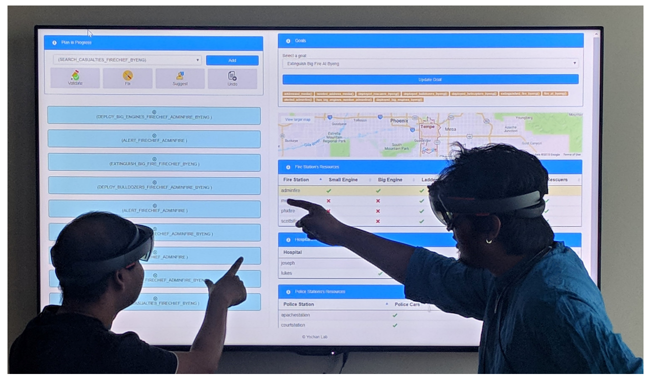

Figure 1: Multiple commanders involved in the collaborative planning process on the MA−RADAR interface.

Privacy Preserving Planning In (Brafman and Domshlak 2008) authors explored multi-agent planning scenarios where each agent has a different domain model with individual actions that can have private preconditions and effects which are not accessible to other agents. Planning in such scenarios becomes more complex because state-space search techniques have to ensure that private state variables of an agent are not exposed to other agents (Brafman 2015).

Here, we assume that apart from the main task of extinguishing the fire, each of the commanders have specific tasks they need to achieve. Furthermore, only the commander (in charge of a specific task) and MA−RADAR have the knowledge of these private tasks. Comm-I is in charge of handling the communication with the media, which is an important aspect in the case of disaster response scenario, Comm-II needs to take care of all communication and deployment of medical help for rescued victims. The private actions of the two commanders follow. When the commanders ask MA−RADAR to suggest missing actions or complete the plan in order to achieve the goal of extinguishing the (big) fire, it communicates the private actions, but only to the specific commander in charge of the private task.

Multi-Model Explanations The second demonstration looks at *plan explanations for model reconciliation* introduced in (Chakraborti et al. 2017a; Sreedharan, Chakraborti, and Kambhampati 2018) – the aim of explanations of this form is to provide updates to the user's possibly faulty understanding of the planning problem to make sure that the optimal plans in the planner's model are also optimal in the

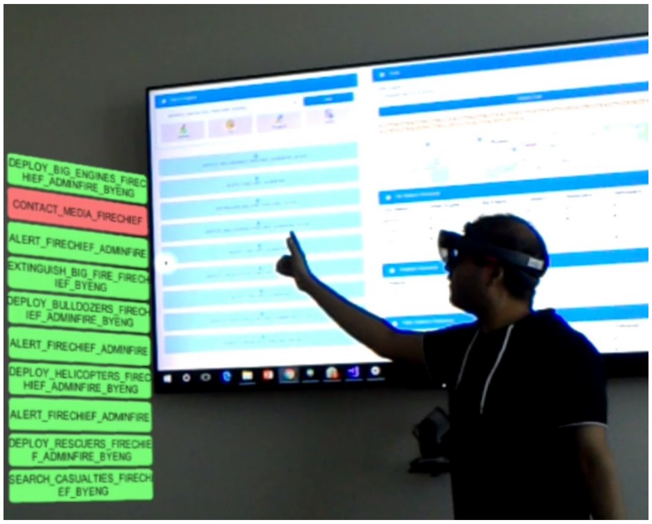

(a) Comm-I, who is responsible for communication with the media, has a private action to contact media as visible only in his POV.

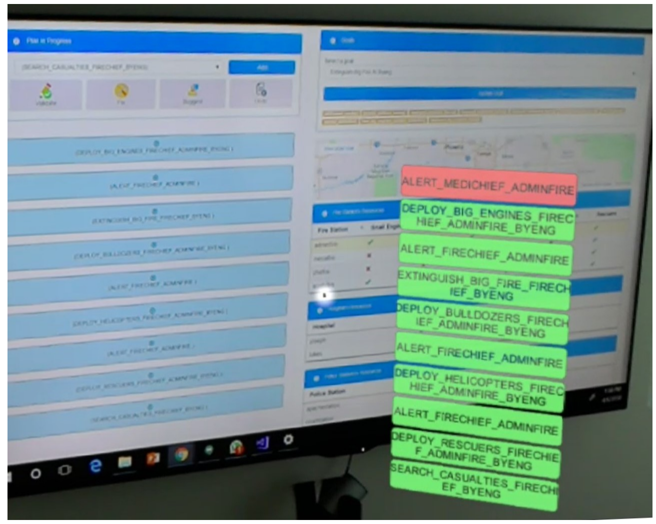

(b) Comm-II, who is in charge of communicating with the medical units, has a private action to alert the medical chief in the area.

Figure 2: Mixed-reality capture illustrating how the public plan in the shared GUI can be overlayed with information on private actions (private actions are in red; public actions are in green) (Brafman and Domshlak 2008) of individual decision makers.

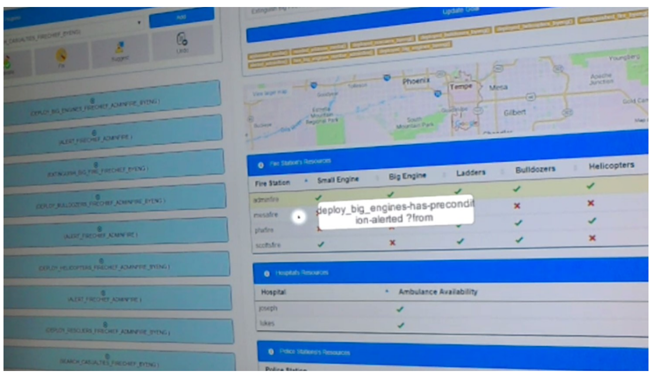

(a) Comm-I, who is unaware of the procedure that a fire-chief needs to be alerted first before deploying the fire engines, is provided this explanation to justify the suggested (public part of the) plan.

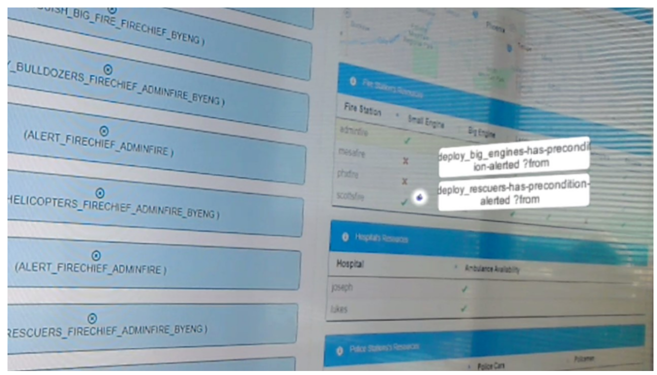

(b) Comm-II, unaware that a fire-chief needs to be alerted before deploying any kind of resources (fire engines or rescuers) from a fire station, is provided both of these model updates as explanations.

Figure 3: Mixed-reality capture illustrating how the multi-model explanation generation algorithm (Sreedharan, Chakraborti, and Kambhampati 2018) can be used to provide targeted explanations to each commander based on their models.

human's. Thus the process of model reconciliation is crucial in maintaining that the decision support agent is on the same page as the human in the loop and thus the establishment of common ground. From the point of view of the interface, there still remains the matter of filtering out superfluous information (due to the single explanation or model update being computed that suffice for all the models) as they are being presented to the individual users.

Here, Comm-I is unaware that the precondition for alerting the authority at a fire-station is missing for deploying (big) fire engines while Comm-II is unaware that fire-stations need to be alerted in order to deploy fire engines or rescuers. When the commanders ask `MA-RADAR` to suggest a plan (or complete a plan) in order to achieve the goal of extinguishing big fire, it will suggest a plan that has both the actions of deploying big engines and rescuers. Since both of these actions need to alert the authority at the fire station, there will be two `alert_firechief` actions which makes the `alerted_firechief` proposition (which is a precondition of these two actions in the original domain) true. Although both the commanders might ask for an explanation, Comm-I just needs to be told about the missing precondition of the `deploy_big_fire_engine` action, whereas Comm-II also needs to be informed about the missing precondition of the action `deploy_rescuers`.

A full description of the system (Sengupta, Chakraborti, and Kambhampati 2018) (presented at the ICAPS-2018 Workshop on User Interfaces and Scheduling and Planning) is available online at http://rakaposhi.eas.asu.edu/ma_radar.pdf.

Acknowledgements. This research is supported in part by the AFOSR grant FA9550-18-1-0067, the ONR grants N00014-16-1-2892, N00014-13-1-0176, N00014-13-1-0519, N00014-15-1-2027, N00014-18-1-2442 and the NASA grant NNX17AD06G. Chakraborti is also supported by the IBM Ph.D. Fellowship for the years 2016-18.

References

Brafman, R. I., and Domshlak, C. 2008. From one to many: Planning for loosely coupled multi-agent systems. In *ICAPS*, 28–35.

Brafman, R. I. 2015. A privacy preserving algorithm for multi-agent planning and search. In *IJCAI*, 1530–1536.

Chakraborti, T.; Sreedharan, S.; Kambhampati, S.; and Zhang, Y. 2017a. Explanation generation as model reconciliation in multi-model planning. Technical report.

Chakraborti, T.; Talamadupula, K.; Dholakia, M.; Srivastava, B.; Kephart, J. O.; and Bellamy, R. K. 2017b. Mr. Jones – Towards a Proactive Smart Room Orchestrator. *AAAI Fall Symposium on Human-Agent Groups*.

Kim, J., and Shah, J. 2017. Towards intelligent decision support in human team planning. In *AAAI Fall Symposium on Human-Agent Groups*.

Sengupta, S.; Chakraborti, T.; Sreedharan; and Kambhampati, S. 2017. RADAR - A Proactive Decision Support System for Human-in-the-Loop Planning. In *ICAPS Workshop on User Interfaces for Scheduling and Planning (UISP)*.

Sengupta, S.; Chakraborti, T.; and Kambhampati, S. 2018. MA-RADAR – A Mixed-Reality Interface for Collaborative Decision Making. *ICAPS UISP*.

Sreedharan, S.; Chakraborti, T.; and Kambhampati, S. 2018. Handling Model Uncertainty and Multiplicity in Explanations as Model Reconciliation. In *ICAPS*.