A DIALOGUE-BASED APPROACH TO MULTI-ROBOT TEAM CONTROL

Nathanael Chambers, James Allen, Lucian Galescu and Hyuckchul Jung
Institute for Human and Machine Cognition
40 S. Alcaniz Street
Pensacola, FL 32502
{nchambers, jallen, lgalescu, hjung}@ihmc.us

Abstract This paper describes an approach to integrating a human into a team of robots through a dialogue-based planning assistant. We describe how the Rochester Interactive Planning System (TRIPS) is able to recognize and translate a user's intentions into collaborative problem solving acts. The user interacts naturally through language and avoids complex machine interfaces as TRIPS manages the appropriate lower level robotic commands and semantics. The system acts as a mediator to the user, managing all coordination and agent communication for a vast number of robots that would normally overwhelm a single user.

Keywords: agents, dialogue-based interaction, human-agent interface, mine search, multi-robot control, teamwork

1. Introduction

This paper describes recent work attempting to naturally integrate a human into a team of robotic agents. We describe the motivations for using a dialogue-based planning system to assist in robot communication and management. The system allows the user to interact with multiple robots through a single entity, thus reducing the burden of communicating simultaneously with multiple robots.

Although we describe a scenario of a single human and multiple robotic agents, much of this work may be extended to teams with varying numbers of humans and agents. Controlling a team of robots is inherently more difficult than single robot control: the human tries to manage multiple robots and to efficiently communicate with them, and the robots try to determine how to accomplish a global objective. As two or more robots attempt to communicate with the human, there is often cognitive overload for the human. A system that understands the human's goals and assists in building plans is essential in a framework that frees the user to work on higher level problems.

TRIPS acts as a mediator who collaborates with the human user to construct plans that can achieve the user's goals, communicating with the robots only when necessary. Alerts, notifications, and queries coming from the robots must go through the mediator who then presents them to the user through dialogue and/or an appropriate graphical interface. The user is freed from monitoring each robot, performing low-level system commands, and grasping difficult semantics and communication languages that robots use. One of the most general models of this interaction is an intent-driven, human-centered model based on natural human-human dialogue.

An intent-driven approach to human-robot interaction is relatively new, having been applied mainly to single software agent planners. We are abstracting the general properties of such an approach and applying it to the new domain of controlling and interacting with multiple agents (who may have plans of their own).

2. Background

TRIPS is a mixed-initiative planning assistant that interacts with a single user through dialogue and supplemental graphical interfaces. Its central components are based on a domain independent representation of linguistic semantics, illocutionary acts, and a collaborative problem-solving model. The system is tailored to individual domains through a domain specific ontology mapper that maps the independent into the specific domain's semantics. The interaction between the core reasoning components drives the system in assisting the user in *planning* and *execution*. Planning with the user involves creating collaborative problem solving acts and maintaining the user's goals, working with him/her to accomplish their needs. Communication with robotic agents is accomplished through a backend module that translates the user's high level intentions into low-level commands that robots can understand. Figure 1 gives a brief description of the main modules. For a complete description of the TRIPS architecture, see (Allen et al., 2001; Allen et al., 2002).

Our approach uses the KAoS (Bradshaw et al., 2004) environment for robot communication and control. KAoS provides the physical communication framework and a set of policy services that allow the user to adjust the capabilities of agents by restricting their behavior, access to resources, and task initiative.

Interpretation Manager (IM)	Coordinates the contextual interpretation of utterances.
Behavioral Agent (BA)	Maintains the collaborative problem solving state between the user, system, and robots.
Generation Manager (GM)	Realizes planned speech acts from the BA into spoken utterances and/or GUI events.
Discourse Context	Maintains discourse history to support reference and ellipsis. Identifies discourse acts, obligations, and expectations.
Ontology Manager (OM)	Maintains the domain independent logical form (LF) and the domain specific languages and provides a mapping between them.
SLIK	Simple Logical Interface to KAoS. Provides mappings between collaborative problem solving acts and external agent communications.

Figure 1. Key modules in TRIPS for interacting with a team of robots

3. Mine Search Scenario

We have applied the TRIPS architecture to a mockup mine detection scenario (using unmanned ground vehicles), an area of great concern to the Navy. One solution to this difficult task is to deploy a (potentially large) number of robots to quickly search and identify mines, allowing humans to then neutralize them. The robots work with humans through mounted cameras to identify unknown objects. This is an excellent example of a mixed-initiative task and requires coordination between robots and humans. The task also requires surface robots to facilitate in communication from the ship to underwater robots because of their extremely limited bandwidth. The robots need a significant level of autonomy to act in the absence of human intervention. A sample of our system interacting with the user can be seen in figure 2.

User:	Find a clear lane through this area <selects rectangular region on map>
TRIPS:	ok
TRIPS:	Alex Richard and Robert are available. <highlights robots on the map> Which do you want to use?
User:	all three

Figure 2. A user's interaction with TRIPS, initiating a team of robots to search an area for mines.

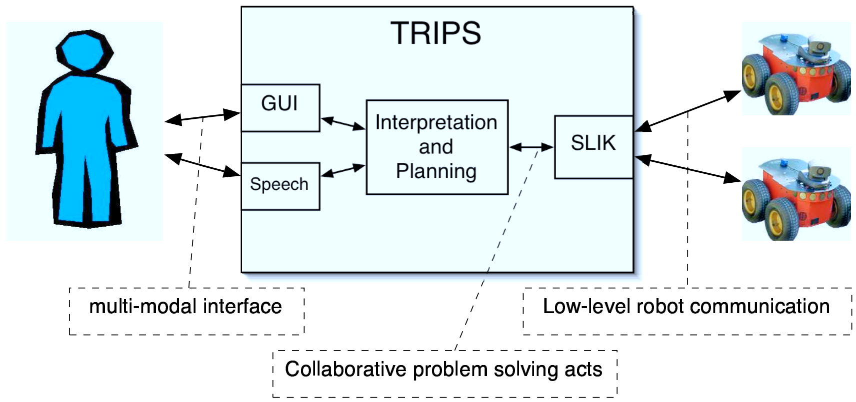

Figure 3. Control flow between the user, TRIPS, and a team of robots.

4. Advantages of a Dialogue-Based Approach

We believe a model for an effective human-robot collaboration must include a planning mediator that is able to interpret a dialogue-based interaction while hiding the user from lower-level robotic requirements. Figure 3 provides a description of our model.

4.1 Dialogue-Based Multi-Modal

Experience with this system has shown that the user greatly benefits from working with a single dialogue-based collaborative agent. As teams of agents grow larger, the user has difficulty understanding who he is/should be talking to. An assisting and planning collaborator reduces the complexity and allows the user to place the burden on TRIPS. Direct human to agent communication can be done through speaking agent names (we used human names in this scenario) or by selecting them on a GUI. In addition, the user can direct large numbers of robots at once and speak of them as a group or a single entity. The multi-modal aspect of TRIPS allows the user to speak in terms that are natural to humans, and use GUIs for deictic reference.

Speak	**GUI action**
'have this robot search a new leg'	<click robot icon on map>
'use these robots'	<highlight multiple robots on map>

4.2 A Planning Mediator

A significant problem in dealing with large teams of agents (or even mixed humans and agents) is that the agents don't necessarily act as helpful collaborators. They may perform complementary tasks with the human to solve a certain problem, but they don't help the user formulate his own plan and method of performing a task. The dialogue-based planning assistant serves as a helpful interface to robots, but it can also build and maintain plans collaboratively with the user. This kind of assistant is not usually found in the other agents of a team.

Secondly, the mine detection scenario makes it clear that the role of TRIPS as a mediator in limiting how information is conveyed to the user is essential to a manageable interface. While many research groups have implemented proxies, it is important to create a mediator that not only has a representation of the user's desires, but also of the user's problem-solving state in order to effectively present and hide varying levels of information coming from the external world.

4.3 Robot Communication

Robotic agents are often relatively primitive in their communication abilities. They may be very sophisticated in regard to the task at hand, but the communication barrier is a difficult one to overcome and is usually less of a priority to robot developers. For these reasons, we use the **Simple Logical Interface to KAoS** (SLIK) to act as a gateway to KAoS, keeping the undesired low level communication hidden from both TRIPS and the human user. SLIK receives a set of collaborative problem-solving acts from the Behavioral Agent in TRIPS and typically translates a *query* into an information *command* that is sent to a robot (our robots, like most robots, are command driven and currently understand OWL). The BA can send queries and receive answers from SLIK as if the agents themselves understood them. Likewise, TRIPS receives input from the robots as if they also communicate in higher level acts. This approach of a single gateway is beneficial in that it can handle requests and commands to teams of robots, not just a single one, in a high level semantic language.

5. Related Work

Most work involving one human and multiple robots use the human in a supervisory role with an emphasis on GUIs (Blake et al., 2000; Jones et al., 2002; Payne et al., 2000). Our approach differs in that the GUI modality is used to supplement the interaction, not dominate it.

Several groups have developed *proxies* (Kortenkamp et al., 2002; Scerri et al., 2003; Martin et al., 2003) for the human that help filter information for the user by maintaining some sort of user model. The work by Martin et al. is perhaps the most similar in that they use a global system planner with pre-defined plans for all the *agents*. Our approach contains a personal planner driven by a model of the *user's* intentions.

6. Conclusion

We have described a general approach for a dialogue-based mediator to assist humans in communicating with non-human-centric robots. By understanding the user's intentions and the current problem solving state, this assistant can reason about the status of multiple agents and lessen the cognitive load on the user. We believe multi-modal mixed-initiative planning assistants like ours are needed to achieve effective collaboration between a mixed team of humans and agents.

Acknowledgments

We thank the KAoS group at IHMC for their support and helpful insight into human-agent collaboration. This work was supported in part by ONR grant N000140310780.

References

Allen, J., Blaylock, N., and Ferguson, G. (2002). A problem solving model for collaborative agents. In *Proceedings of AAMAS-02*, Bologna, Italy.

Allen, J., Ferguson, G., and Stent, A. (2001). An architecture for more realistic conversational systems. In *Intelligent User Interfaces*, pages 1–8.

Blake, M. A., Sorensen, G. A., Archibald, J. K., and Beard, R. W. (2000). Human-assisted capture-the-flag in an urban environment. In *IEEE Int. Conf. on Robotics and Automation*.

Bradshaw, J. M., Jung, H., Kulkarni, S., Allen, J., Bunch, L., Chambers, N., Feltovich, P., Galescu, L., Jeffers, R., Johnson, M., Taysom, W., and Uszok, A. (2004). Toward trustworthy adjustable autonomy and mixed-initiative interaction in kaos. In *Proceedings of the AAMAS Workshop on Trust in Agent Societies*.

Jones, H. L., Rock, S. M., Burns, D., and Morris, S. (2002). Autonomous robots in swat applications: Research, design, and operations challenges. In *Proceedings of the 2002 Symposium for AUVSI*, Orlando, USA.

Kortenkamp, D., Schreckenghost, D., and Martin, C. (2002). User interaction with multi-robot systems. In *Workshop on Multi-Robot Systems*.

Martin, C., Schreckenghost, D., Bonasso, P., Kortenkamp, D., Milam, T., and Thronesbery, C. (2003). Aiding collaboration among humans and complex software agents. In *AAAI Spring Symposium on Human Interaction with Autonomous Systems in Complex Environments*.

Payne, T., Sycara, K., Lewis, M., Lenox, T., and Hahn, S. (2000). Varying the user interaction within multi-agent systems. *Autonomous Agents*.

Scerri, P., Pynadath, D. V., Johnson, L., Rosenbloom, P., Schurr, N., Si, M., and Tambe, M. (2003). A prototype infrastructure for distributed robot-agent-person teams. In *Proceedings of AAMAS-03*.