

# Research Summary: Communication-Sensitive Decision Making in Multi-Agent, Real-Time Environments

Marie desJardins, Karen Myers, David Morley, and Michael Wolverton

SRI International  
333 Ravenswood Ave.  
Menlo Park, CA 94025  
{marie,myers,morley,mjw}@ai.sri.com

## Overview

In a recently started project, we are developing techniques for intelligent agent control and coordination in a dynamic, real-time, multi-agent setting. The application domain, consisting of teams of autonomous air vehicles (AAVs), is characterized by dynamic environments, real-time response requirements, limited information, and unreliable, low-bandwidth communications.

We have developed an initial framework that models decision making, reasoning about constraints, and learning at multiple levels of abstraction and multiple time scales, within and across agents. Decision making in this framework is sensitive to communication availability and costs, tradeoffs among multiple objectives, and reliability of information about other agents (friendly and hostile) in the environment.

We plan to focus our research on three key areas:

- Reasoning about communication constraints to inform action selection during planning and execution.
- Integrating deliberative planning, reactive planning, and continuous real-time control.
- Learning methods to improve performance within the AAV environment.

The first of these, communication-sensitive decision making, is the focus of this research summary. In the next section, we describe the AAV application domain. We then outline our approach to communication-sensitive decision making and discuss related work.

## Autonomous Air Vehicles

The Office of Naval Research, which sponsors this research, is interested in developing autonomous air vehicles (AAVs) that can be deployed to accomplish predefined missions without continual ground-based guidance from a human operator.

These vehicles must operate in hostile, highly dynamic environments. Their behavior must be goal-directed, yet sensitive to the current situation, and they must be able to respond rapidly to unexpected events. They must be able to

operate as a team, with multiple AAVs coordinating their behavior to accomplish joint missions, in an environment with limited, unreliable, and low-bandwidth communication.

AAV tasks range from the very high level (tasks to provide surveillance for a given region) down to the very low level (hover in a fixed location, execute a high-speed turn with a given radius). The environment can be characterized, and must be understood by the agents, at multiple time scales (hours down to milliseconds) and multiple spatial scales (kilometers down to centimeters). Conceptually, then, the agents may need to model actions and objects in the world at many different levels of abstraction simultaneously.

These AAVs are expected to be able to handle a wide range of missions. For our research, we are focusing on two mission categories: and surveillance and reconnaissance (S&R) and pursuit and evasion (P&E). S&R scenarios require coordinated response from a team of AAVs in order to gather information about a particular region of interest, about which the agents may initially have very limited information. P&E scenarios involve teams of agents locating and tracking one or more enemy units.

The characteristics of the tasks in these scenarios vary along several dimensions. In particular, tasks may be fixed in duration and scope (e.g., fly to a particular target, collect imagery data, and return to base), or they may be open-ended and continuous (e.g., patrol a given region for an extended period of time). Teams of agents may be assigned tasks that are relatively independent of each other (e.g., each patrol a pre-assigned region) or, conversely, that require substantial coordination (e.g., jointly harry a group of ground-based hostile agents towards a particular region). Finally, tasks may be primarily reactive (e.g., fly to a series of pre-determined waypoints, avoiding detection by known enemy ground-based installations) or require substantial strategic planning (e.g., gather intelligence on enemy movements in a previously unmapped region).

For each of these task characteristics, the former classes of tasks (fixed scope, independent, reactive), while by no means trivial to achieve, are relatively simpler than the latter categories (open-ended, coordinated, strategic). We are primarily interested in the issues raised by these latter task categories.

## Communication-Sensitive Decision Making

We are interested in developing a framework for coordinated activity of AAV teams in the presence of unreliable, intermittent, and low-bandwidth communications channels. In particular, the agents should be able to reason about how their actions will influence the reliability of communication, and use this knowledge to select plans that are likely to succeed. The plans selected by the agents may include actions to improve the reliability of communication, such as positioning themselves to set up a series of relay points to enable communication between geographically distant agents. Conversely, the agents may recognize the fact that their planned actions will create a situation with limited communication availability, and accordingly select strategies that are likely to succeed despite these conditions. For example, they may select a plan that includes default behaviors—in case coordination messages are dropped—and a joint plan for regrouping at the end of the mission.

Imagine a P&E scenario in which a team of two AAVs is pursuing an enemy air vehicle along the ridgeline of a low mountain range. If the enemy flies to one side of the ridgeline, the agents may agree that one AAV will pursue the enemy along that side, while the other AAV parallels their track along the other side of the ridge. The latter AAV can then scout for additional enemy units, and can intercept the enemy if it crosses the range. These agents plan to communicate periodically in order to synchronize their location and share information about enemy positions.

However, suppose that the ridgeline blocks communication, so that any message sent is very unlikely to be received. The agents may instead choose an alternative plan that minimizes communication requirements, such as pursuing the agent jointly along the side of the ridge, establishing a default time and location to regroup. Alternatively, they may add actions designed to increase communication availability and reliability, such as agreeing to periodically “pop up” (increase their altitude) in order to synchronize and communicate.

The goal of our research is to explore methods by which a team of agents can predict the likely communication availability for a set of alternate plans and environments, and use these predictions to evaluate and select from among the plans. These techniques could be applied during execution time to select the best course of action given each option’s expected communication requirements, and the anticipated communication availability if that particular option is selected. The techniques could also be applied offline to analyze the robustness of a knowledge base of planning operators or reactive schemas with respect to possible communication failure scenarios.

We are currently constructing a model of communication quality as a function of the environment. The quality of communication is characterized by factors such as bandwidth, signal strength, and probability of packet loss. The communication model abstracts away from low-level network management issues such as routing and connectivity. The environment model includes factors that influence communication quality, such as distance, terrain, weather, antenna position and orientation, and communication equipment ca-

pabilities and failure modes.

We next plan to develop predictive techniques for assessing the likely communication situation given a plan (or set of reactive behaviors) and an environment. This will allow the agent to assess the likelihood of communication failures, which in turn leads to an assessment of the likelihood (or quality) of mission success. A plan can then be selected that results in optimal (or “good enough”) performance, given the expected environment.

Both the communication quality model and the predictive model will be probabilistic, since a fundamental characteristic of the domain is limited information about the situation and about the behavior of the environment.

## Related Work

To our knowledge, there has been very little work that specifically factors communication requirements and costs into plan selection or that evaluates the quality (robustness) of these plans in the event of communication failures. Stone and Veloso (1999) explored methods for robust autonomous team behavior in RoboCup, which is an environment with low-bandwidth, unreliable communication. They addressed the problem of communication failures by developing behaviors that are not completely dependent on receiving messages from other agents. For example, while waiting for a response, agents continue to perform their current roles. The methods developed are fairly specific to the RoboCup domain.

Simmons *et al.* (2000) addresses the related problem of task selection in multi-robot exploration and mapping of an unknown environment. They present an algorithm that maximizes overall utility by trading off the individual preferences of robots of which regions to explore with the global utility of avoiding overlapping explorations. While communication is not a factor in this work, the approach of maximizing overall information gain, taking into account individual costs of exploration for alternative tasks, is analogous to our goal of maximizing overall goal achievement, taking into account costs of communication for alternative tasks.

Several researchers have studied the value of shared information in multi-agent environments. Balch and Arkin (1994) presented empirical results in simulated robot environments on the value of different types of communication for a particular suite of tasks. In our previous work, we have explored the use of irrelevance reasoning to determine which parts of a plan are valuable to share with other agents in distributed hierarchical task network planning (Wolverton & desJardins 1998). In several other approaches (Gmytrasiewicz & Durfee 2000; Xuan, Lesser, & Zilberstein 2000; Tambe 1997), a utility analysis is used to decide whether to send a particular piece of information to another agent. We hope to use these existing approaches as the basis for modeling the *value* of communication. Our new contribution will be in modeling the *quality and availability* of communication, and using these models to perform action selection, thus ensuring robustness in the face of unreliable communication.

## Conclusions

We have described the autonomous air vehicle (AAV) domain, a dynamic, real-time, multi-agent environment. We outlined some of the key domain issues and research areas for AAV, and discussed our ongoing work in developing communication-sensitive decision making methods for the AAV domain. This work is in the preliminary stages.

## References

- Balch, T., and Arkin, R. C. 1994. Communication in reactive multiagent robotic systems. *Autonomous Robots* 1:1–25.
- Gmytrasiewicz, P. J., and Durfee, E. H. 2000. Rational communication in multi-agent environments. *Autonomous Agents and Multi-Agent Systems* 3(4):319–350.
- Simmons, R.; Apfelbaum, D.; Burgard, W.; Fox, D.; Moors, M.; Thrun, S.; and Younes, H. 2000. Coordination for multi-robot exploration and mapping. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*, 852–858.
- Stone, P., and Veloso, M. 1999. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence* 110(2):241–273.
- Tambe, M. 1997. Towards flexible teamwork. *Journal of Artificial Intelligence Research* 7:83–124.
- Wolverton, M., and desJardins, M. 1998. Controlling communication in distributed planning using irrelevance reasoning. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence*, 868–874. AAAI Press.
- Xuan, P.; Lesser, V.; and Zilberstein, S. 2000. Formal modeling of communication decisions in cooperative multi-agent systems. In *Proceedings of the Second Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT 2000)*.