

SELECTING ORGANIZATIONAL STRUCTURES FOR ADVANCED MULTI-AUV SYSTEMS*

Roy M. Turner and Elise H. Turner
Department of Computer Science, University of Maine
5752 Neville Hall, Orono, ME 04469
(207) 581-3909 {rmt,eht}@umcs.maine.edu

Abstract

Current multi-AUV systems rely on a human to define their organizations or else organize themselves into very simple kinds of organizations. While this is sufficient for simple missions involving few AUVs, it will not be adequate for future multi-AUV systems involving large numbers of AUVs that perform complex missions, are fielded for long periods of time, and/or cannot communicate with a human. In the CoDA project, we are developing techniques to allow such advanced systems to self-organize and to reorganize as necessary to effectively carry out their missions. This paper discusses autonomous organization design for such systems. The approach draws on the organization design literature and will represent organization structures (e.g., hierarchies, teams, markets, etc.) appropriate for the AUV domain. These will be linked to representations of contexts in which the organizations is predicted to be useful. Once CoDA recognizes the current situation as an instance of a known context, then the organization structure can be retrieved and instantiated. The work is being developed in the domain of autonomous oceanographic sampling networks (AOSNs), in which groups of AUVs and other instrument platforms cooperate to return data over a long time period from an area of interest in the ocean.

The concept of autonomous multiagent systems for ocean applications dates back to the early MAUV (Albus, 1988) and MAVIS (Turner *et al.*, 1991) projects and was given significant substance and refinement by the idea of autonomous oceano-

graphic sampling networks (AOSNs; Curtin *et al.*, 1993). Recent years have seen the development and fielding of simple multiagent systems comprised of autonomous underwater vehicles (AUVs), many of which were reported at a recent workshop focused on multi-AUV systems.¹ These have taken on real-world missions, typically under human control, and usually with few vehicles. Much remains to be done, however, before large-scale systems can function autonomously for long periods of time doing complex, demanding missions, especially when the system will need to be re-tasked or reconfigured over time. Issues of autonomous organization and reorganization, task assignment, common command languages, communication languages, and problem-solving protocols must all be addressed before the vision of a fully-capable AOSN can become a reality.

The CoDA (Cooperative Distributed AOSN control) project (Turner & Turner, 1998, 2001) focuses on developing mechanisms to support the realization of heterogeneous, flexible multiagent systems of AUVs for ocean science. We are concerned with advanced AOSN-type systems that will be comprised of a large number of heterogeneous agents (AUVs, other instrument platforms, etc.). Furthermore, these systems will be deployed for long periods of time. They will be *open systems* (Hewitt, 1986), meaning that agents will come and go over time, for example, due to failure. Such systems must be able to self-organize to fit their initial circumstances and to reorganize as the situation—and the system itself—changes. Work in CoDA is in three areas: problem-solving protocols to support organization, reorganization, and work on mission goals; assigning mission tasks to agents based on their capabilities; and selecting the appropriate organization for the agents given the current situation. Previous pa-

*This work was funded in part by grants N0001-14-96-1-5009 and N0001-14-98-1-0648 from the Office of Naval Research. The content of this paper does not necessarily reflect the position or the policy of the U.S. government, and no official endorsement should be inferred. The authors would like to thank Kim McKeage, Wenyu Zhu, and Sheri St. Pierre for their contributions to this work.

¹IEEE/OES Autonomous Underwater Vehicles 2004 (AUV'04): A Workshop on Multiple AUV Operations, Sebasco Estates, Maine.

pers have reported on the first two of these (Turner & Turner, 1998, 2001). In this paper, we focus on work in progress on organization selection.

The task of selecting an organization for a group of agents is context-dependent: which organization will work best depends on the environment, the mission, and which agents are available. For example, the severe communication bandwidth limitations of the underwater environment precludes some communication-intensive organizations, such as those based on consensus, from being used in most situations. Missions that are readily decomposable into independent or nearly-independent tasks lend themselves to hierarchies, whereas others may require an organization more like a team or a committee. If the system contains no agents capable of managing others, then hierarchies are impossible. The task of selecting an organization, then, is one of matching the current situation's features to those for which a given organization is appropriate. Since the organization depends on these kinds of features of the situation, as the situation changes, the system will need to be reorganized.

Our approach treats organization selection as a diagnostic process. One or more agents in the multi-AUV system use features of the system's current situation to diagnose the situation as being an instance of one or more known classes of situations, called *contexts*. The agents will have been given knowledge about the contexts, in particular how to behave when in the contexts, as part of knowledge structures called *contextual schemas* (c-schemas; Turner, 1998). Part of the knowledge contained in a c-schema representing a kind of situation is what organizational structure (e.g., a hierarchy of some form, a team, a committee, etc.) is appropriate for the system to use in that context.

For this approach, we draw on our past and ongoing work on context-sensitive reasoning for autonomous underwater vehicle control. We also are making use of the extensive literature on organizational design for human organizations for information about the kinds of organizations that are possible. Obviously, some human organizations may not be appropriate for AOSNs. However, many are, at least in broad outline, and we are identifying the features of situations in the AOSN domain that would suggest which organizational structure to use.

In the remainder of the paper, we first discuss organizing multi-AUV systems, then turn to the particular approach being developed in CoDA, including how to represent organization structures, how to link situations to the organization structures appropriate for them, and how to assess the agents'

current situation.

Organizing Multi-AUV Systems

Not all multi-AUV systems will need the ability to create their own organizations. For example, current work on AOSNs either has a human design the agent's organization or the agents self-organize into extremely simple organizations, such as a swarm.

However, we are most interested here in those circumstances in which a multi-AUV system will need the capability to self-organize and reorganize as the situation demands. This need may arise, for instance, in missions in which there is uncertainty and little or no possibility of communication with human users once the system is deployed. An example would be a mine-hunting mission in which AUVs are deployed some distance away from the minefield (e.g., by submarine launch or air drop). It is likely in this case that: there is uncertainty about the conditions obtaining at the work site; some of the AUVs may not make it to the work site, and so the composition of the overall system is uncertain; vehicle failure is likely during the mission; and communication from the work site is likely to be extremely limited or nonexistent. Other examples include teams of AUVs operating under ice (AUV failures may occur, communication limited or impossible); and recovery of a downed aircraft's "black box" under conditions of severe sea state (uncertain work environment, difficult to communicate with base); and long-duration AOSN missions (vehicles entering or leaving, environmental change, desire for no dedicated human controller).

Once we decide to give a multi-AUV system the ability to self-organize (and reorganize), the questions naturally arise:

- Which organizations are possible to choose from?
- How to select the best organization for the current situation?

These questions have long been examined for human organizations in the field of organizational design. We can look there for guidance for multi-AUV systems, realizing that some organizations may not work in the non-human case and that some changes will need to be made to some organizations.

Organization Structures

An *organization structure*² is “the system of task, reporting, and authority relationships within which the work of the organization is done” (Moorhead & Griffin, 1998).

Many different organization structures have evolved for human organizations, some of which have direct analogs for multi-AUV systems. These can be grouped into several high-level categories.

Probably the most familiar type used in current multi-AUV systems (e.g., Phoha *et al.*, 2001) is the simple hierarchy, in which a single manager controls a group of subordinates. Other kinds of hierarchies exist as well. Multi-layer hierarchies, for example, introduce one or more layers of middle managers between the topmost manager and the workers.

Hierarchies can also be distinguished by how work is distributed, as well. Product hierarchies are those with departments associated with products (e.g., consumer product division, etc.), while functional hierarchies are those with departments associated with functions (e.g., engineering, marketing, sales, etc.). For AUV systems, a product hierarchy might correspond, for example, to organizing the system with one “department” working on search while another works on data sampling, or organizing “departments” by location. A functional hierarchy would correspond to having “departments” of AUVs with similar sensors, effectors, or problem-solving expertise.

There are many other organizations in addition to hierarchies. Small groups, for example, may run by consensus, in which there are no authority relationships and all decisions are made by mutual consent. Teams are another form of small-group organization that are distinct from hierarchies in the level of participation allowed the subordinates in the decision-making process. These have been used extensively in multiagent systems (MAS) (e.g., Tambe, 1997), including for such real-time MAS as robo-soccer teams. Organizations can also run by voting; examples exist from committees to municipalities to countries. Market-based systems are also common and have been extensively investigated in the MAS research community (e.g., Smith, 1980; Sandholm, 1993). Even more esoteric organizations are possible, such as that of the scientific community; this, too, has been modeled in multiagent systems research (Kornfeld & Hewitt, 1981). And there are, of course, heterogeneous organizations, such as hierarchies in which departments bid for contracts from each other; many companies are organized at

least in part in this manner.

The question naturally arises: Which of these human organization structures are good for multi-AUV systems? The answer to this depends on the answer to another question: What distinguishes one organization structure from another? Once this is known, then the features of multi-AUV systems can be matched to organizations that are most appropriate in general.

One differentiating feature is how good an organization structure is for environmental and task complexity. As complexity increases, the bounded rationality (Simon, 1957) of an individual agent can quickly become overwhelmed, leading to hierarchies with their concomitant division of labor. Increasing complexity, however, can overwhelm even a hierarchy, leading instead to market-based organizations (Fox, 1981). Markets, relying as they do on local evaluation of contracts (e.g., Smith, 1980), can in some circumstances reduce the cognitive load on individual agents. For complex multi-AUV missions, then, a market-based approach might be reasonable, especially if the cognitive capacity of the individual agents is low.

On the other hand, uncertainty tends to favor hierarchies over markets, as markets under conditions of uncertainty have to resort to complex contingent contracts (Fox, 1981), whereas a hierarchy can simply mandate new task–role assignments. Some multi-AUV missions will entail significant uncertainty, due to lack of knowledge of the environment, poor sensors, ill-specified missions, a highly-dynamic environment, or a combination of these factors. In these cases, hierarchies may have the edge over (e.g.) markets.

Communication bandwidth required is one very important differentiating feature between organization structures. Markets tend to require more communication, for example, than hierarchies (Malone, 1987). Consensus-based organizations, even those such as involved in *partial global planning* (Durfee & Lesser, 1987) in which only local consensus is reached, still impose a great communication bandwidth cost. At the other extreme, convention-based organizations (e.g., automobile traffic) and some of those based on game theory (e.g., Genesereth *et al.*, 1988) require little or no communication at all. Communication will almost always be limited in multi-AUV settings, sometimes extremely so. Consequently, the communication bandwidth available compared to that needed by an organization structure may be a strong determinant of which one is selected.

Organizations also differ in their reliance on au-

²Also called an *organizational structure*.

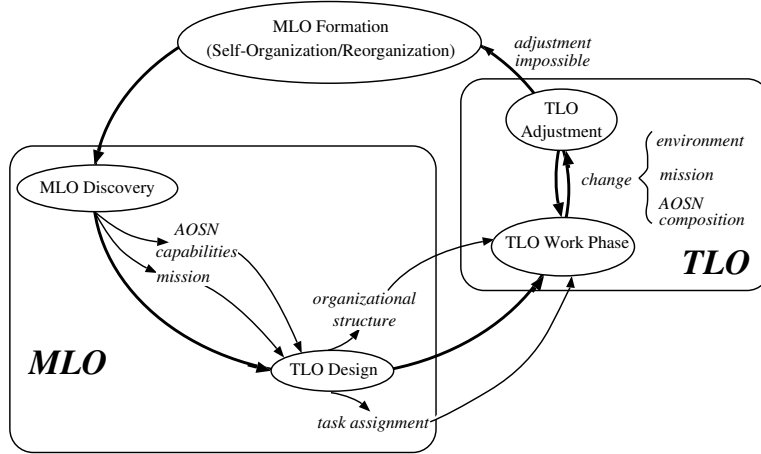


Figure 1: Overview of CoDA. (From Turner & Turner, 2001; copyright © 2001 IEEE.)

thority relationships. For example, a hierarchy requires either a means to enforce authority relationships (appeals to law, force, or loss of benefits, or, in the case of artificial agents, a common designer) or, in the case of artificial agents, a common designer) or, in the *acceptance of authority* perspective (Moorhead & Griffin, 1998), a willingness on the part of the subordinates to obey orders. Market-based organizations do not have this problem, exactly, but they do require some means of ensuring that contracts are enforced, either by coercion (e.g., laws), complete trustworthiness (e.g., common designer), or self-interest (e.g., actual monetary payments with agent memory of past contract fulfillment by others). In general, then, organizations differ in their assumptions about just how cooperative the agents are. Hierarchies assume complete cooperation, while market-based organizations can tolerate self-interested agents. For some multi-AUV systems, this will not be a problem, as the AUVs can simply be instructed to obey authority relationships. In other situations, however, AUVs may come from different labs or may vary in their trustworthiness, etc., and so it will become a factor.

Organization structures also differ in their assumptions about the ability to maintain a global perspective. Hierarchies are based on the idea that it is possible to have a global, albeit necessarily abstract, view of the world, and that it is possible to force global coherence of the solution. Markets and such structures as partial global plans take the often more realistic view that global coherence is impossible, and so content themselves with *functionally accurate* (Lesser & Erman, 1981) solutions or plans. In the underwater domain, where factors include poor sensors, uncertain knowledge, limited communication bandwidth, and limited cognitive capacity,

an assumption of a detailed global perspective is not realistic, although assuming that there can be an abstract global perspective may be, depending on the situation.

Organization Design

Organization design involves selecting or creating an appropriate organization structure for the current situation and instantiating it to form an organization. The parameters of an organization, including tasks, roles, and properties of the organizational structure itself, define a space of possible organizations (Carley & Gasser, 1999) which can be quite large. The organization design task can be viewed as searching this space for the organization that is appropriate in the current situation.

Heuristics used to search the space can make use of factors of the environment, the agents available, and so forth. This approach is similar to one that has received the most attention in recent years in the organization design literature, the *structural imperatives* approach (e.g., Moorhead & Griffin, 1998). This approach views the environment, the technology, and the size of the organization (the imperatives) as the primary factors affecting organizational design. This is in keeping with our approach, at least with respect to the environment and organization size—and, somewhat stretching a point, the “technology” could be construed to mean the sensors and effectors present on the vehicles.

As with any heuristic search, the search process can be drastically short-circuited if there is some a priori knowledge of the search space. In this case, we know many organization structures that have worked for human organizations. If we restrict the

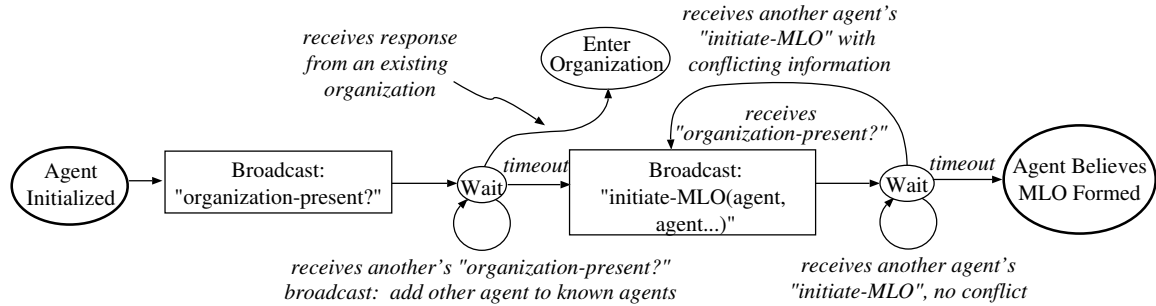


Figure 2: MLO formation protocol. (From Turner & Turner, 2001; copyright © 2001 IEEE.)

search to variants of these, then we reduce the size of the search space immensely from all possible organizations to those based on one of the known organization structures. In this view, organization design, whether for a human organization or a multiagent system, requires selecting an organization structure, possibly adapting it to the situation, and instantiating it by assigning agents to roles and, for some structures, assigning tasks to roles.

An Approach to Organization Design for Multi-AUV Systems

The CoDA project has developed an overall approach to self-organization and reorganization for multi-AUV systems. In this section, we briefly describe this approach and organization design’s place within it. In the following sections, we discuss current work on representing organizations, linking situations to organizations appropriate for them, and determining what the situation is, so that an organization can be selected.

CoDA takes a two-level approach to organizing a multi-AUV system such as an AOSN, as shown in Figure 1. Since the AUVs or other instrument platforms initially present may know very little about each other, including which others are present, no highly-specific organization can be given to them in advance. Consequently, the AUVs first self-organize into a very flexible, though likely inefficient, *meta-level organization*, or MLO. This organization is simple to create and is composed of the most intelligent of the AUVs present. The job of the MLO is to analyze the situation, including which AUVs are present, their capabilities, and the mission, and to then design an efficient *task-level organization* (TLO) to actually carry out the mission. When the situation changes beyond the TLO’s ability to cope, a new MLO is formed to design a TLO for the

changed situation.

To participate in the system, all agents (AUVs and instrument platforms) agree to abide by a set of *cooperation protocols* that governs their interactions with the other agents. The protocols are rule-like descriptions of what the agent should do under particular circumstances. There are protocols for the different phases of organization shown in the figure, and there are protocol variants for different kinds of agent. For example, more intelligent agents follow different protocols than those that do not have decision-making capabilities.

Figure 2 shows an example protocol, which governs an intelligent agent from the time it arrives at the work site until a new MLO is formed, or until the agent enters an existing organization. Other protocols are reported elsewhere (Turner & Turner, 2001).

In this paper, we are concerned more with the portion of CoDA having to do with organization design, the TLO design phase. In previous work, we have assumed that TLO design would be done by a single agent, chosen by some mutually agreed-upon convention (e.g., lexicographic order of agent names), and that the result would be a simple multi-layer hierarchy. In the current work, however, the protocol has been expanded, as shown in Figure 3.

For now, we still assume that a single agent will be selected by convention, and with no communication, to design the TLO; in future work, we will consider the distributed case. Once the this planning agent has been selected, it asks the other agents in the MLO for information about the capabilities they “control”—i.e., that they themselves have or that they know about from their interactions with other agents during a prior phase of problem solving (see Turner & Turner, 2001). We assume that by this time in the MLO, the agents have all come to mutual understanding of the mission. The planner then assesses the situation, using features of the

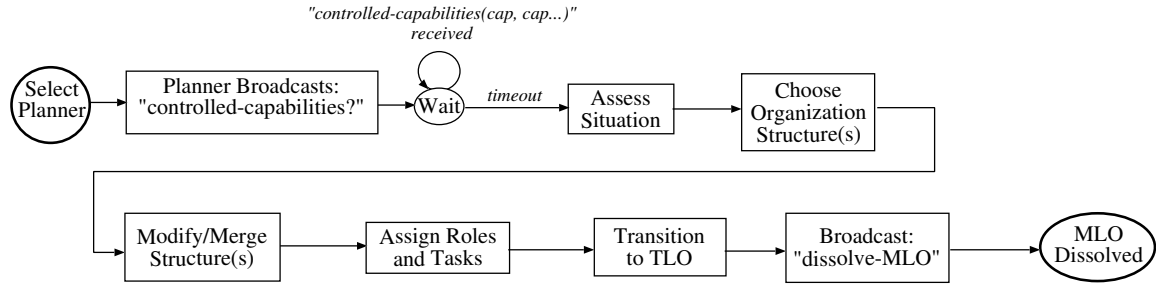


Figure 3: The new TLO Design protocol.

environment, the mission, and the agents, including their capabilities. It then uses that assessment to determine which organization structure or structures might be appropriate for the current situation, as described below. If there is more than one structure possible, then the agent must decide which one to use or whether to merge them into a hybrid structure. It then instantiates the structure to form the task-level organization by assigning agents to roles and, for many organization structures, tasks to the agents.³ At that point, the TLO is informed that it is in control—a process that will differ based on the kind of organizational structure has been selected—and the planner dissolves the MLO by sending a message to its peers.

Representing Organizations

In order to reason about organization structures, a way of explicitly representing organizations and organization structures is needed. From one point of view, it does not matter much which knowledge representation scheme is used; at one level, all knowledge representation schemes are more or less equivalent. For the present, in CoDA we intend to use a frame-based mechanism, since that in keeping with the rest of the system’s representation. We could just as easily use first-order logic, semantic nets, or other schemes.

What is important is to determine *what* to represent. In our view, any representation of an organization structure needs to include at least information about:

- **effectiveness for uncertainty:** How well do organizations based on this structure perform under conditions of uncertainty?

³For some organizations, such as markets, tasks would not be assigned here, but rather would be given to the organization as a whole to assign via mechanisms of the organization, for example, bidding for contracts.

- **effectiveness for complexity:** How well do organizations of this type perform when the task and environment are very complex?
- **size of organization:** How big are the organizations controllable using this structure?
- **intelligence needed:** How intelligent do agents need to be to participate in this type of organization? This might be expressed, for example, by listing the different cognitive capabilities (planning, scheduling, constraint satisfaction, etc.) needed in general in the organization.
- **communication bandwidth requirements:** How much message traffic is needed (possibly broken down by time periods or phases of operations)? What kinds of messages are needed?
- **controllability:** Where does the resulting organization lie on the spectrum of fully emergent properties to fully specifiable outcomes?
- **coordination mechanism:** How does coordination happen in this organization—voting, one role giving orders to another, contracting, etc.? The information of this sort may include protocols for the agents to use when participating in the organization.
- **roles:** What roles are possible in the organization, and which should be present?
- **links between roles:** This includes both authority links and communication links. Authority links specify, for applicable organizations, who controls whom, as well as the degree of absoluteness of that control. Information about communication links includes communication pathways (role-to-role, broadcast) as well as the kinds of things that can be said.

In addition, some organization structures may have additional types of information. Some, for example, may have information about span of control (how many subordinates a particular manager role can/should control), or about the kinds of bids acceptable, etc. It also may be advantageous for organizations to have information about under what

conditions they fail, how graceful that failure might be, etc.

Figure 4 shows an example of how an organization structure might be represented; the one shown is a simple hierarchy, with one manager and many workers. This is meant only as an illustrative sketch of an actual implementation. As such, it is far simpler than the actual representation will be.

An organization is an instantiation of an organization structure. Consequently, in CoDA the representation of an organization will look something like the representation for an organization structure. Of course, additional information will be present to record the linkages between agents and roles, etc.

Linking Organizations to Situations

In order to link an organization structure to a situation in which it is useful, there needs to be some way to represent situations themselves. CoDA will use the same mechanism as its sister project, Orca: contextual schemas (Turner, 1998).

In our approach, an agent's *situation* consists of all the features of its world and itself at a given time. A situation therefore both has indeterminate (and vast) extent and is unique. Situations, however, naturally fall into classes that have implications for how the agent should behave. We use the term *context* to refer to such a class of situations, since this is similar to at least one usual sense of the word, as in "the agent's context was being in a harbor."

A *contextual schema*, or *c-schema*, is an explicit representation of a context. A c-schema represents, then, one or more—usually many—situations, although many details of any given situation will necessarily be omitted or abstracted. C-schemas in our approach are frame-like knowledge structures with two parts: descriptive knowledge and prescriptive knowledge.

A c-schema's descriptive knowledge describes the situations that are instances of the represented context. For example, a c-schema representing "in a harbor" would contain information predicting that the water column will be fairly shallow, there will be surface traffic, there may be a main channel in which ships navigate, land will be nearby, and so forth. Such knowledge not only helps identify the situation as being an instance of the context, but also can be used as a source of context-dependent predictions to help interpret sensor data. Also included is knowledge about what concepts mean in the context. For example, different c-schemas would provide different

membership functions for the fuzzy value "nominal" of the fuzzy variable "depth" (Turner, 1997) or, in systems using neural networks, different weights for different contexts (Arritt & Turner, 2003a).

Prescriptive information tells the agent how to behave in the context. This has included in the past:

- context-specific ways to achieve goals;
- information about event-handling in the context, including how to recognize that an event has occurred, how to diagnose the event, how to assess its importance, and what action to take in response;
- attention-focusing information, such as information about the importance of various goals in the context; and
- standing orders, that is, parameter settings that should be automatically in effect when in the context.

Together, this information allows the agent to automatically behave appropriately for the context, once that context has been recognized.

In CoDA, we will add to this prescriptive knowledge information about what organization structures are appropriate to use in contexts where multiagent systems are being considered. This means that in addition to the usual contexts characterized by features of the environment ("in a harbor") or the mission ("search mission"), we will add c-schemas representing being in a multiagent context. The descriptive information of these c-schemas will contain knowledge about the other agents present, etc., and the prescriptive information will provide information about organization structures.

Situation Assessment for Organization Design

The process of finding an organization structure for a group of agents depends, in our approach, on finding the c-schema or c-schemas that contain the appropriate structure, that is, on situation assessment. Situation assessment in CoDA will mirror that in Orca. This process has been reported elsewhere (Arritt & Turner, 2003b), so here we will be brief.

Situation assessment can be fruitfully viewed as a diagnostic task in which observed features of the situation are used to first evoke, then verify hypotheses about the context. Our approach is based on work on medical diagnostic reasoning, in particular the abductive differential diagnosis process developed for the INTERNIST/CADUCEUS program (Miller *et al.*, 1982).

<p>uncertainty-effectiveness: high complexity-effectiveness: low size: low agent-intelligence: $\exists x \text{ agent}(x) \wedge \text{has-capability}(x, \text{manage})$ bandwidth: low controllability: moderate coordination-mechanism: direct control roles: types: manager, worker constraints: $\exists!x \text{ manager}(x)$ <i>;; there is only one manager</i> $\exists x, y \text{ worker}(x) \wedge \text{worker}(y) \wedge x \neq y$ <i>;; exists > 1 worker</i> $\forall x, y \text{ manager}(x) \wedge \text{worker}(y) \wedge x \neq y$ <i>;; managers don't work</i> authority-links: $\exists x \forall y \text{ manager}(x) \wedge \text{worker}(y) \Rightarrow \text{controls}(x, y)$ <i>;; the manager is in charge</i> $\forall x \forall y \text{ worker}(x) \wedge \text{worker}(y) \wedge x \neq y \Rightarrow \neg \text{controls}(x, y)$ <i>;; no worker controls another</i> communication-links: $\forall x \forall y \text{ worker}(x) \wedge \text{worker}(y) \Rightarrow \neg \text{communicates}(x, y)$ <i>;; workers do not communicate</i> $\exists x \forall y \text{ manager}(x) \wedge \text{worker}(y) \Rightarrow \text{communicates}(x, y)$ <i>;; manager and workers communicate</i></p>
--

Figure 4: An example organization structure: a simple hierarchy.

In this approach, c-schemas are stored in an associative memory that can retrieve items based on features of the current situation (e.g., Kolodner, 1984), along with a measure of how strongly the features evoked the item. The diagnostic process then takes all of the c-schemas thus evoked and groups them into competitor sets. Each c-schema is rated based on how strongly it was evoked as well as on what it explains successfully in the current situation, and penalized by what it does not explain, but should. Information to support this rating process comes from the c-schema’s descriptive information.

Once the hypotheses have been rated, the competitor set containing the topmost hypothesis (i.e., the topmost set) is selected for further work. The agent then compares the hypotheses in the set to determine what information would serve to separate the topmost hypothesis from the others sufficiently to consider it the “solution” to the set. Once this information has been gathered (e.g., from sensors or asking other agents), then new competitor sets are formed and the process continues.

When a set has been solved, then that hypothesis is considered confirmed, which means that the c-schema represents (at least partially) the agent’s context. When all sets have been solved, then the agent has a set of c-schemas that, when merged, will represent a coherent picture of its current context. From the merged c-schemas will come the organization structure the agent should use for the current situation.

Conclusion and Future Work

Work has been ongoing for some time in CoDA on mechanisms and protocols to allow groups of AUVs and other instrument platforms to self-organize into effective multiagent systems. Work has now begun to devise ways for such systems to choose an organization that is most appropriate for the situation at hand. The mechanism being developed, as discussed in this paper, relies on explicitly representing classes of situations that have implications for an agent’s behavior, then bringing information contained in these c-schemas to bear in similar situations. This mechanism already is in use a sister project. In CoDA, the repertoire of c-schemas will be expanded to include multiagent situations, and to the kinds of predictive information that can be represented will be added organization structures that are appropriate for the situations represented.

Many issues will have to be addressed as this work progresses. The representation of organization structures will need to be fleshed out more fully. The representation of instantiated organizations will also need to be worked out in more detail. Task assignment, for which a constraint-based mechanism has already been devised in CoDA, will need to be integrated into the process, at least for some kinds of organization structures. C-schema merger is already being worked on, but as part of this, we will have to determine how best to merge organization structures when more than one is suggested by the

c-schemas representing the context. This may include devising ways to create “hybrid” or blended organizations, with some parts organized differently than other parts. Distributed organization design will also need to be addressed, as the current scheme has the drawback of a single point of failure and a processing bottleneck. Mechanisms to learn from experience will also be needed to acquire additional c-schemas as well as to adapt those the agent already has over time to its problem-solving environment. Case-based reasoning (Kolodner, 1993) is an obvious and natural choice for a first step, as c-schemas evolved as generalized cases from work on case-based reasoning.

The end result of the work will be a general mechanism for selecting context-appropriate organization structures for multiagent systems. This mechanism will first be developed and tested in the domain of multi-AUV systems such as AOSNs, then broadened to include other multiagent systems.

References

- Albus, J. S. (1988). Multiple Autonomous Undersea Vehicles. Technical Report NIST Technical Note 1251, National Institute of Standards and Technology, Gaithersburg, MD 20899.
- Arritt, R. P. & Turner, R. M. (2003a). Context-specific weights for a neural network. In *Proceedings of the Fourth International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT'03)*, pages 29–39, Stanford, CA. Springer, New York.
- Arritt, R. P. & Turner, R. M. (2003b). Situation assessment for autonomous underwater vehicles using a priori contextual knowledge. In *Proceedings of the Thirteenth International Symposium on Unmanned Untethered Submersible Technology (UUST)*, Durham, NH. The Autonomous Undersea Systems Institute, Lee, NH.
- Carley, K. M. & Gasser, L. (1999). Computational organizational theory. In Weiss, G., editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press, Cambridge, MA.
- Curtin, T., Bellingham, J., Catipovic, J., & Webb, D. (1993). Autonomous oceanographic sampling networks. *Oceanography*, 6(3).
- Durfee, E. H. & Lesser, V. R. (1987). Using partial global plans to coordinate distributed problem solvers. In *Proceedings of the 1987 International Joint Conference on Artificial Intelligence*, pages 875–883.
- Fox, M. S. (1981). An organizational view of distributed systems. *IEEE Transactions on Systems, Man and Cybernetics*, 11:70–80.
- Genesereth, M. R., Ginsberg, M. L., & Rosen-schein, J. S. (1988). Cooperation without communication. In Bond, A. H. & Gasser, L., editors, *Distributed Artificial Intelligence*, pages 220–226. Morgan Kaufmann Publishers, Inc., San Francisco, CA.
- Hewitt, C. (1986). Offices are open systems. *Communications of the ACM*, 4(3):271–287.
- Kolodner, J. (1984). *Retrieval and Organizational Strategies in Conceptual Memory*. Lawrence Erlbaum Associates, Publishers, Hillsdale, New Jersey.
- Kolodner, J. L. (1993). *Case-Based Reasoning*. Morgan Kaufman, San Mateo, CA.
- Kornfeld, W. & Hewitt, C. (1981). The scientific community metaphor. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1).
- Lesser, V. R. & Erman, L. D. (1981). Functionally-accurate, cooperative distributed systems. *IEEE Transactions on Systems, Man and Cybernetics*, 11(4):81–96.
- Malone, T. W. (1987). Modeling coordination in organizations and markets. *Management Science*, 33(10):1317–1332.
- Miller, R. A., Pople, H. E., & Myers, J. D. (1982). INTERNIST-1, an experimental computer-based diagnostic consultant for general internal medicine. *New England Journal of Medicine*, 307:468–476.
- Moorhead, G. & Griffin, R. W. (1998). *Organizational Behavior: Managing People and Organizations*. Houghton Mifflin Company, New York, NY.
- Phoha, S., Peluso, E., & Culver, R. (2001). A high-fidelity ocean sampling mobile network (SAMON) simulator testbed for evaluating intelligent control of unmanned underwater vehicles. *IEEE Journal of Oceanic Engineering*, 26(4):646–653.
- Sandholm, T. (1993). An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of Eleventh National Conference on Artificial Intelligence (AAAI-93)*, pages 265–262.
- Simon, H. A. (1957). *Models of Man*. Wiley, New York.
- Smith, R. (1980). The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, C-29(12):1104–1113.
- Tambe, M. (1997). Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124.
- Turner, R. M. (1997). Determining the context-dependent meaning of fuzzy subsets. In *Proceedings of the 1997 International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-97)*, Rio de Janeiro.
- Turner, R. M. (1998). Context-mediated behavior for intelligent agents. *International Journal of Human-Computer Studies*, 48(3):307–330.

- Turner, R. M., Fox, J. S., Turner, E. H., & Blidberg, D. R. (1991). Multiple autonomous vehicle imaging system (MAVIS). In *Proceedings of the 7th International Symposium on Unmanned Untethered Submersible Technology (AUV '91)*, pages 526–536, Durham, NH.
- Turner, R. M. & Turner, E. H. (1998). Organization and reorganization of autonomous oceanographic sampling networks. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation (ICRA '98)*, pages 2060–2067, Leuven, Belgium.
- Turner, R. M. & Turner, E. H. (2001). A two-level, protocol-based approach to controlling autonomous oceanographic sampling networks. *IEEE Journal of Oceanic Engineering*, 26(4).