APPROPRIATE COMMITMENT PLANNING FOR AUV CONTROL*

Erik Albert, Elise H. Turner, and Roy M. Turner
Department of Computer Science
University of Maine
Orono, ME 04469–5752
albert17@umcs.maine.edu, elise_turner@umit.maine.edu, rmt@umcs.maine.edu

Abstract

We propose a method for reactive mission planning in dynamic environments that also allows for *appropriate* commitments to future actions and goals. We consider appropriate commitments to be those based on predictable features of the situation and environment that are unlikely to change. By organizing actions based on appropriateness, we hope to create a planner that is not only able to react quickly to unanticipated events, but also be able to coordinate its actions with other agents, handle newly-arising goals, and perform missions efficiently and effectively.

This approach is currently being implemented in the Orca 3 AUV Mission Planner.

Introduction

For some AUV missions, it is sufficient that the AUVs be able to follow pre-compiled mission plans or to take action based on their own suite of simple behaviors. However, for many current and future missions of interest to AUV users, it is important that the AUVs themselves be able to take the initiative in planning the mission and in replanning when things go wrong. This is the case when the AUV must operate with little prior knowledge of its environment, and when it must be out of contact with humans for a significant time.

Of special interest is the ability to plan autonomously during complex, long-duration missions and missions involving multiple agents working cooperatively, such as autonomous oceanographic sampling networks (AOSNs) [8]. AOSNs are groups of AUVs and other instrument platforms that work cooperatively to return data about an area of interest, possibly over a long time period. For example, several AUVs and moorings might be fielded to characterize seasonal convective overturn events in the boreal ocean, hypothesized to play an important role in carbon sequestration in the deep ocean. AOSNs might also be used for long-term surveillance (e.g., for military or drug interdiction purposes), environmental monitoring, mine detection and clearing, or oceanographic data collection. An AUV controller participating in an AOSN, in addition to controlling its AUV, must be able to communicate and cooperate with other AUVs as part of the multi-agent system.

The task of mission planning in complex, dynamic environments is a difficult problem. An AUV should be able to react quickly to unanticipated events such as an effector failure or an unforeseen obstacle, but still be able to prepare for expected future actions such as a rendezvous with another agent. Further complicating the matter is that agents in these types of missions often do not have complete information about all of the goals that they will need to complete. New goals can be introduced by an operator, by the planner itself (e.g., for acquiring resources or repairing failures), or by another agent. In general, collaboration between cooperating agents that do not share a planner will involve performing tasks and solving goals that were not known (at least to one of the agents) before the start of the mission. An extreme, yet interesting, example of this is a mission that requires the agent to "join and take part in an existing AOSN." In this case, none of the actual goals of the mission will be known until the AUV is situated within the AOSN organization, and has been allocated tasks to perform or goals to achieve. Regardless of the source, an agent should be able to incorporate newly-arising goals into its

*This work was funded in part by grants from the United States Office of Naval Research (N00014-01-1-0818, N0001-14-96-1-5009, and N0001-14-98-1-0648). The content of the paper does not necessarily reflect the position or the policy of the U.S. government, and no official endorsement should be inferred.

overall mission plan.

Most previous approaches to flexible on-board mission-level control of AUVs can be thought of as variants of *reactive planning* [1, 9, 10]. This includes both the more extreme, behavior-based variants of reactive planning (e.g., [5, 11, 21, 14, 3, 4]) as well as more moderate forms of reactive planning, including the authors' own prior work on the Orca AUV controller (e.g., [18, 19]). The problem with the former approach is that there is no commitment to future actions at all: the AUV is completely reactive, and hence, it is difficult to reliably predict or bound its behavior. The problem with the second approach is that although there is a plan, typically there is minimal commitment to future specific actions, even those that it can be predicted with a high degree of confidence will be needed. This was a recognized shortcoming in Orca, for example, that was left for later work while attention was paid to context-sensitive behavior [20] and other aspects of reasoning.

What is needed is an approach that extends reactive planning to allow for *appropriate* commitment to future actions and goals. Such an approach would ensure that the AUV accomplishes necessary goals and perform missions efficiently by better organizing its other actions and plans around these commitments.

We consider appropriate commitments to be those which are based on predictable features of the situation and which are important enough to justify the risk of those predictions not being realized. In order to successfully complete a mission, it is important that actions be performed effectively and efficiently. By committing to and organizing around actions that are appropriate, we hope to create a reactive planner that can be more successful at completing complex missions in dynamic environments.

This paper is organized as follows. In the first section, we introduce the problems associated with autonomous mission planning in dynamic environments and the need to extend reactive planning to allow for appropriate commitments to future actions. We then introduce three existing planning systems and highlight their contributions towards the concept of appropriate commitment planning. An overview of the new planning technique is then given, detailing how appropriate commitments are made and the method that the planner uses to determine what to do next. Next, we describe the testbed planning system that is being developed to support this research. Finally, we give our conclusions and describe areas of future research.

Previous Work

Our work builds on our previous work on the Orca reactive planner, on JUDIS [15], a discourse control front end for a distributed system, and on the NBA-Planner [16, 6], a hierarchical path-planner for AUVs.

Orca is a context-sensitive, reactive mission planner for the AUV domain. Orca uses a knowledge base of schemas for both procedural and contextual reasoning. Procedural schemas (p-schemas) [18] form a generalization/specialization hierarchy and each specify the steps that must be taken to achieve a goal. These steps can be primitive actions, other p-schemas, or subgoals. Multiple p-schemas can exist in the knowledge base to achieve a single goal.

Orca uses an agenda to list the goals that comprise its intentions. At any time, Orca focuses on the one best goal from the agenda for the current situation. A p-schema is found to achieve the goal and Orca expands the partial plan, executing primitive actions as they are found. If the situation changes, Orca will look for a more-appropriate specialization of the existing p-schema, or will begin work on a different goal from the agenda.

The shortcoming of this planning system is that the focus of attention is always on a single goal and there is no overarching view of the mission as a whole. Orca is unable to predict or to summarize its intended actions over the course of its mission, other than just listing the goals and partially expanded p-schemas in its agenda. Also missing is the ability to intelligently interleave the execution of related actions from multiple p-schemas.

JUDIS' domain of dialogue planning required the system to follow conversation conventions to create a coherent discourse. It is important enough to follow these conventions, and the conventions provide enough predictions, so they can be used to organize new conversation goals as they arise. The NBA-Planner was an attempt to bring JUDIS' mechanism for organizing newly-arising goals to the task of path-planning for AUVs. Natural boundaries are permanent, or extremely long-duration, features of the environment that are costly for an AUV to cross. As goals arise, they are associated with naturally-bounded areas (NBAs) so that crossings of natural boundaries can be minimized. The planned ordering of visits to NBAs provides the organizational structure in which newly-arising goals are placed.

FOCUS OF ATTENTION – SEE ELISE'S NOTES

JUDIS and the NBA-Planner use domain- and task-dependent knowledge to select predictions, and both commit to placing all newly-arising goals within an inflexible organizational structure. Orca, on the other hand, is meant to be a general purpose problem solver able to handle a wide array of real-world tasks. They have very different levels of reactivity and commitment. Where JUDIS and the NBA-Planner may overcommit to organizational structure, Orca only commits to actions that are ready to be executed and not to any predictions that can be used to organize newly-arising goals. These disparities also lead to differences in mechanisms for focusing attention. Where Orca chooses the highest-priority action that can be executed at the time, JUDIS makes following its organizational structure an important factor in choosing the next goal.

Making Appropriate Commitments

Our current work on appropriate commitment planning seeks to combine the previous work described above to create a general purpose mission planner which can make appropriate commitments to predictions about real-world domains, and use those predictions to choose its course of action and handle newly-arising goals.

The properties that we consider tantamount to appropriate commitment planning include

Being able to dynamically order and reorder actions and goals, including interleaving actions from multiple goals and identifying when multiple subplans contain the same parts (and thus can be combined into a single unit). By organizing around features of the plan (and environment) that are difficult or costly to achieve, the planner should be able to accomplish tasks opportunistically and efficiently.

Being able to perform mutually-exclusive actions from disjoint goals in parallel. For example, since the AUV body is able to use sonar and GPS devices at the same time, these actions shouldn't have to be performed sequentially due to restrictions of the planner.

Remaining able to maintain conventions to remain predictable and able to cooperate with other agents that follow conventions. Conventions are especially important for preserving the coherency of conversations.

Remain responsive to changes in the environment and mission and be able to quickly adapt without the need for replanning from scratch. This includes being able to incorporate newly-arising goals into the current plan. And[1]

Having an overarching view of the plan as a whole that can be used to coordinate activies, plan future events (such as scheduling a time to surface for recharging via solar power), and summarize the intended actions over the course of a mission.

In order to achieve these properties in a planner, we have developed an organizational structure that allows us to explicitly represent and reason about plan components (including alternatives) and features that can be used to organize these components. In the following sections we describe this structure, introduce the rationale and methods for organizing around predictive features, and explain how the planner decides "what to do next" (its focus of attention") at any time during the mission.

The Reactive Plan Network

Information about the current plan will be encoded explicitly in a *reactive plan network*. A reactive plan network is a combination of plan components (goals, partially expanded p-schemas and information about alternatives being considered, and primitive actions) and *organizational nodes* (inspired by the template of JUDIS). The reactive plan network allows the planner to dynamically order and reorder actions and goals without breaking convention, group actions that can be executed together or at the same location, and consider alternative courses of action.

Organizational nodes (also called organization nodes) are nodes in the network that are used to group goals and activities around the predictable features of the plan and environment. These components will allow the planner to intelligently group and interleave actions from disjoint p-schemas, prioritize actions that require a limited resource, choose between alternatives, and identify actions that cannot be accomplished (e.g., because of resource exhaustion or a missed deadline).

Figure 1 shows a simple example of a reactive plan network. In this diagram, plan components are drawn using circles (goals), triangles (p-schemas), and squares (executable actions). There is also an organizational node, drawn as a double-circle, connected to several actions in the plan.

An important benefit of the reactive plan network over a simple agenda is that it allows for an overarching view of what specific tasks the agent is currently working on, as well as a general plan (or summary) of its future course of action. This view of

[1] where to put this and?

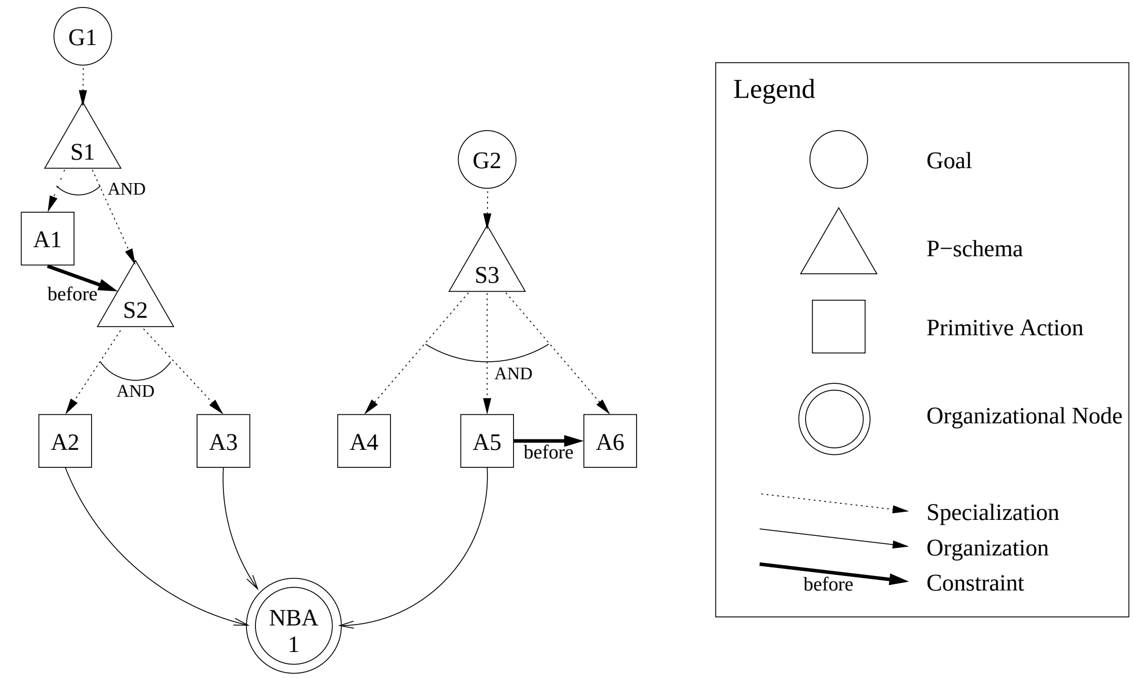

Figure 1: A Simple Reactive Plan Network

the network is important for coordinating activities when working cooperatively with other agents.

Predictive Features

Definition and rationale for predictive features. For the initial version of the planner, we chose to focus on location, time, and resources as the predictive features important enough for incluson in the reactive plan network.

Location is an important predictive feature when it comes to maximizing the efficiency of a mission plan. It makes sense for an AUV to perform as many tasks as it can while at a particular location, even if these tasks advance work on separate goals. For example, consider an AUV in a long-term oceanographic mission. Needing to recharge its batteries, the AUV plans to dock with a solar powered mooring. If the AUV has another goal that requires sending a long-range communication (another capability of the mooring), it would make sense for the AUV to also commit to performing this action while docked. In fact, the AUV's planner may opt to do some exploration of its entire mission to see which actions (that work to accomplish any of its goals) it can perform while at the mooring.

The AUV can also organize actions based on the natural boundaries of the environment, such as strong currents or areas with thick vegetation, that are costly to cross. The NBA-Planner showed that the utility of plans can be improved when an AUV organizes actions based on these naturally-bounded areas before using straight-line distances. It is beneficial to perform as many of the actions that need to be performed in an NBA before crossing the boundary, in order to minimize the total number of crossings.

In addition to location, time is also an important predictive feature that can be used to organize plan components. Many goals will have a deadline for completion, or other time constraints. Reactive planners have traditionally not been able to handle actions with hard deadlines, even though it is very important for an agent to be able to accomplish time-sensitive tasks as well as to be able to coordinate actions with other agents. Many types of time constraints, such as a scheduled rendezvous, offer important information that can be used to make commitments and structure the plan.

The final type of predictive feature that we will address in this first version of the planner is the use of various types of resources by the agent. There are two types of material resources: replaceable[2] and non-replaceable [12, 13] Replaceable resources can generally be considered as "tools", meaning that as they are used they are not consumed (such as a hammer or a resource 2). Many of an AUVs replaceable resources will be equipped before a mission, such as

[2] perishable not talked about in this section...

sonar or GPS devices, but others may have to be acquired during the course of a mission. For example, if the AUV is deployed for a long-duration mission, it may need to rendezvous with a ship at some point to be outfitted with specialized equipment not known to be needed or not available when deployed.

Non-replaceable resources are resources that are consumed when used. Examples of these resources include power, fuel, and money. These resources may be sharable (e.g., power is shared between all of the electrical components in an AUV), producible (e.g., power from a solar panel), or exchangeable (e.g., money for fuel). All of these properties of resources must be taken into consideration when they are used to organize, or choose between, actions in a plan. For example, a solar AUV may be more inclined to use high-cost electrical equipment on a clear day, but may reserve its power for locomotion when the skies are overcast.

Replaceable resources – i.e., tools – often require non-replaceable resources. For example, a sonar uses power. Information about such interactions is often included in the description of the action that uses the replaceable resource. For example, an executable action that uses a sonar would have as part of its description that it uses power. We choose, however, to explicitly treat one resource that uses another (such as the sonar using power) as a second order resource. By keeping their relationship explicit, we can then create a more robust network of organizational nodes that communicate and share activation. This will also allow us to have one generic action that can use any potential resource, rather than having to enumerate all of the actions with all of the resources that have different costs.

Organization Nodes and Links

The planner will explicitly use the predictive features detailed above in order to organize and select actions and activities throughout the duration of the mission. For each predictive feature, the planner will create a new organizational node in the reactive plan network. In order to keep the number of nodes small, the planner will initially create nodes only for the resources that the AUV already possesses. Links to other resources will be connected to a special network component, called the *potential resources node*, that keeps track of all of the resources that may need to be acquired.

There are two types of organization links in a reactive plan network. The first type of link, the *examine link*, connects a goal or partially expanded schema to the predictive features that might be used in order for the goal to be achieved. Information about these predictive features can be compiled into the description of each goal or schema based on the possible alternatives that exist in the knowledge base that achieve the goal or specialize the schema.

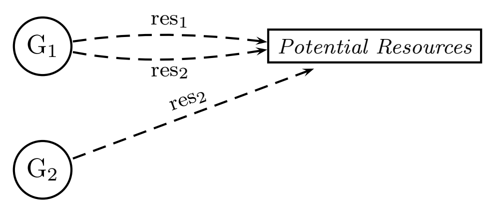

(a) A reactive plan network with two goals

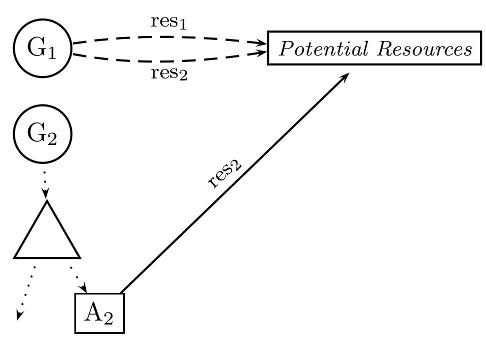

(b) Planner commits to action requiring res_2

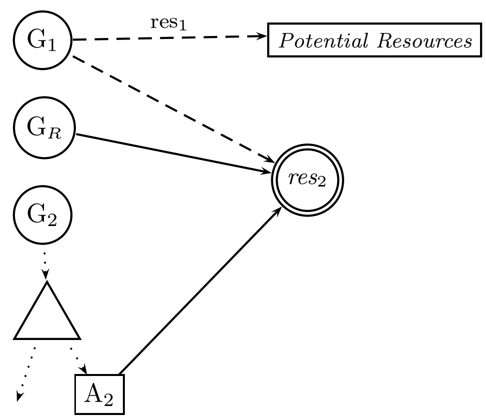

(c) Planner plans to acquire res_1

Figure 2: Instantiation of an organization node

Figure 2(a) shows a portion of a reactive plan network that consists of two goals, G_1 and G_2, which can potentially use two resources, res_1 and res_2. Since the agent has neither of these resources, the two goal nodes are connected to the potential resources node using examine organization links (shown using dashed lines).

At some point in the planning process, figure 2(b), the planner finds, expands, and commits to a p-schema that achieves the goal G_2 using the resource res_2 (the rest of the expansion is omitted for clarity). The examine link between G_2 and the potential resources node is replaced with the second type of organization link, the *use link* (shown using a solid line) that originates from the action that uses the resource (A_2).

Now that the agent knows that it will need to acquire the resource (a use link is connected to the potential resources node), it adds this goal (G_R) to the network and instantiates a resource node for res_2 (figure 2(c)). As the planner works to achieve G_R, G_1 will gain activation through its examine link, and the planner will begin to look for ways to achieve the goal. If res_1 and res_2 are used by two separate alternative p-schemas, the planner will choose the alternative that uses res_2.

Both types of organization links are used to transfer activation used for focusing attention, however, examine links transfer only a fraction of the activation passed through them. This ensures that the planner will only attempt to acquire high-cost resources (or other state[3] based on predictive features, such as traveling to a remote location) if the payoff is large.

Organizational Strategies

Organizational nodes lend activation to the goals, p-schemas, and executable actions to which they are connected based on the following organizational strategies:

Location Strategy All of the actions that need to be executed at a specific location are connected to the location organization node (figure 3). This organization node disperses activation to each of connected plan components based on the strategy used by the NBA planner. Activities located in the current naturally bounded area receive an amount of activation in proportion to their position in an ordering determined by a path planning algorithm.

***E*xample showing position of AUV and activities with related activations?**

Time Strategy ... *not ordering*

Replaceable Resource Strategy Because replaceable resources can be reused without being consumed, the organizational strategy is simply to give a constant level of activation to all of the connected plan components if the agent has acquired the resource, otherwise no activation is given. In the latter case, the node will instead collect activation from its connections and the planner may choose to plan to acquire the resource.

If a replaceable resource is perishable, i.e., the agent has access to the resource for a limited time, the amount of activation given will be inversely-proportional to the remaining time of availability.

When a replaceable resource is a second order resource, the amount of activation given is derived from the amount of activation given to that node by the resource(s) to which the node is connected.

[3]state's not the right word

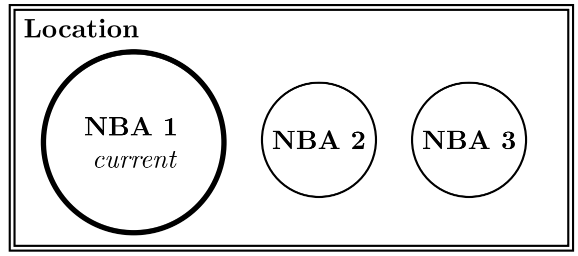

Figure 3: The location organization node

Non-replaceable Resource Strategy ...

Sub-strategies for different properties.

if obtainable or producible, but exhausted/nearly, add goal to acquired more

Of course, the organizational nodes can also influence the planner *not* to choose specific actions or courses of action or to postpone expanding certain areas of the plan. For example, when a non-replaceable resource is highly-used, about to be exhausted, or "reserved" for high-priority actions, it may lend low or negative activation to actions that would further consume it, causing alternative methods of achieving a goal without using the resource (if available) to be chosen instead.

The appropriate commitment planner will include generic classes and subclasses of organization nodes for each of these organizational strategies. The class of node chosen for resources will depend on the description of the resource in the domain knowledge. A domain engineer will also be able to create new, specific classes for predictable feaures of their domain[4].

The activation distributed by each organization node can be viewed as each node's preference for

[4]not sure that I like this sentence

what should be done next. The planning process can then be thought of as finding a consensus between the desires of the organization nodes and the intentions of the planner that doesn't violate convention (or does so only if necessary).

Focus of Attention

Focus of attention in the appropriate commitment planner, for both execution of primitive actions and plan refinement, is driven by activation [17]. Sources of activation include *intention* and *predictions*. Intentional activation is determined by the importance (priority) of a goal in the mission and causes the planner to focus on completing the primary tasks. In addition to importance, intention also lends activation based on urgency; as deadlines approach, time-critical goals and actions receive additional activation to ensure that they are completed in time.

The organizational nodes of the reactive plan network can lend predictive activation to related actions in the plan. Different classes of organizational nodes can have different rules about when and how activation is distributed to, or accumulated from, connected plan structures.

The final order of actions that the agent undertakes is determined by a combination of convention (the order of actions in a p-schema) and activation. The planner avoids breaking convention for two reasons: firstly, since the agent may not have information about the reason for the ordering within a p-schema, it cannot know whether a reordering will result in a successful plan. The second reason to maintain convention is to remain predictability of the agent. Predictability is an important feature of an agent that is working in cooperation with other agents *and/or that needs to appear rational to human observers*[5].

SOME SORT OF SUB-CONCLUSION HERE?

The Orca 3 Mission Planner

The next generation (version 3) of the Orca intelligent AUV mission controller is being constructed to embody our approach to appropriate commitment planning. Orca has been completely redesigned as an agent-based system where the main functionalities have been encapsulated as intelligent components. To avoid confusion, agents that comprise Orca are called modules and the term agent is reserved for entities which can take part in a mission (such as an AUV).

Figure 4 shows the proposed configuration for this new version of the planner. The basic Orca framework consists of a shared communications bus and support for accessing a shared working memory. In order to be able to be situated on different AUV systems (as well as other types of architectures), Orca communicates with its host (agent body) solely via a Low Level Architecture Interface Module. For each host, there will be a different interface module that can translate between Orca's and the architecture's representation of data.

The controller is an application that allows for the planner to be stopped, started, stepped, or completely reset. This component may be controlled by a user, such as when an AUV is being prepared for deployment, or by another piece of software, such as when Orca is being used to control an agent body in simulation.

Shared between all modules is Orca's Extendible Semantic Planning Memory (ESPmem). This memory includes all of the facts known about the state of the environment. Eventually, this module will be enveloped by the ConMan (the Context Manager, previously called ECHO) in order to supply context-sensitive values.

The Event Manager module filters all of the telemetry, sensor, and other information received from the low level architecture and updates Orca's working memory with the new data. A primary task for the Event Manager is detection and reporting of anticipated events (e.g., arriving at a waypoint) to other modules.

The Communications Manager module will handle communication from other agents with which Orca may be cooperating or otherwise interacting. Because of the natural language processing heritage of the appropriate commitment planner, it is expected that some or all of this module's functionality may be implemented by the planner itself.

Support modules are any modules that can assist in the overall operation of Orca such as debuggers, error handlers, and loggers; as well as specialized reasoners such as expert systems, case-based reasoners, constraint satisfaction reasoners, or path-planners, which can aid the planning processes. Orca can even include in its tool set one or more from-scratch planners that could be used when there lacks a schema specific enough to solve a goal in the system.

The Planner/Executive is the module in which the appropriate commitment planner will be implemented. This module has two main responsibili-

[5] hmm?

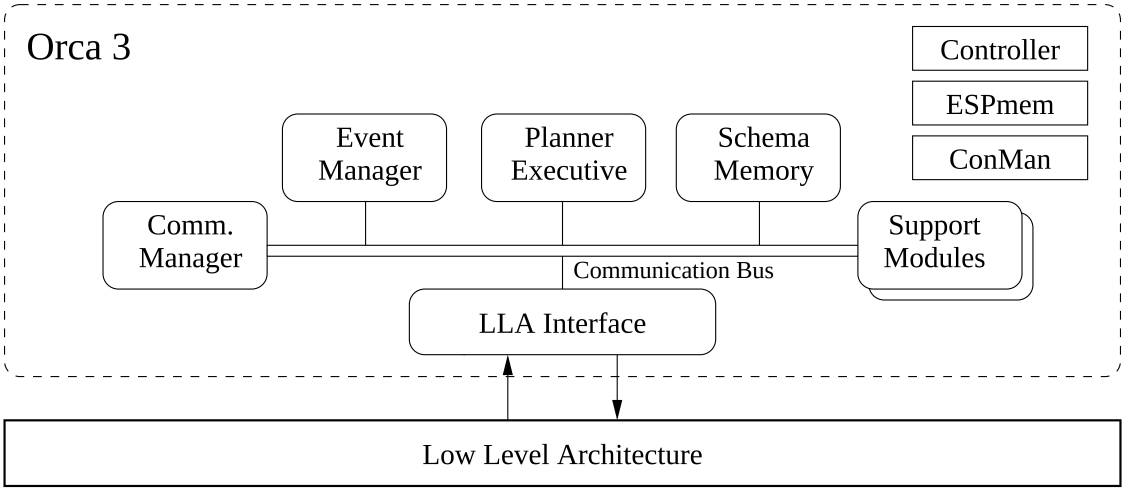

Figure 4: The Orca 3 Mission Planner

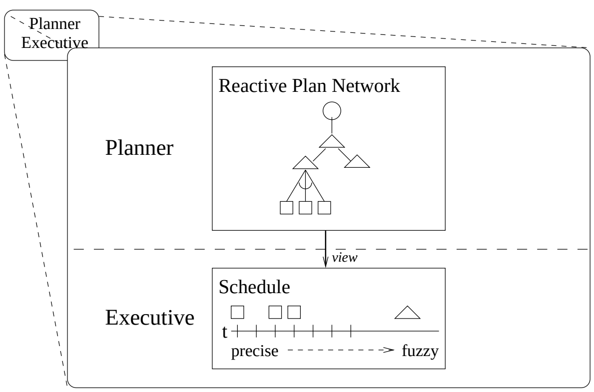

Figure 5: The Planner/Executive Module

ties[6]. The first is to plan a course of action by continuously updating and refining the reactive plan network. The second responsibility is to execute actions that are ready and to monitor the results of the execution. The executor uses a temporally-ordered view of the reactive plan network as its "schedule."

Conclusions and Future Work

CONCLUSIONS

Future areas of research for this project include

- Developing a strategy for determining the proper level of detail to which components of a plan should be expanded. JUDIS fully expands all of its plan components to the full detail, and Orca fully expands the current focus of attention, leaving all other components unexpanded. The appropriate commitment planner will need to expand components to various levels of details in order to predict and exploit organizations of actions based on the strategies defined in this paper without overexpanding (and thus overcommitting).

- Developing a method to determine what additional features of the plan or the world can be used in developing new organizational nodes for the reactive plan network and creating new organizational strategies for these features.

- If there are two alternatives to achieve a goal, one being costly or difficult and the other being inexpensive or simple (but currently infeasible), can the planner commit to the difficult alternative while keeping around the simple alternative in the event that it eventually becomes feasible? We will look into when it makes sense to *hedge our bets* in this manner, and how to decide when to remove the backup alternative from the reactive plan network because the cost of switching to the simple alternative outweighs the cost of finishing the difficult alternative.

- If all (or most) of the alternatives for a p-schema have the same primitive action(s), it is reasonable to assume that these actions will need to be executed, and therefor the planner can commit to them, before the agent commits to a particular schema alternative. In order to make use of this property, we will need to research a mechanism for overlaying p-schemas in

[6]The Planner/Executive may become two separate modules in the future.

order to identify similarities.

Other work is ongoing focused on context management for autonomous agents (e.g., [20]). This work will be incorporated into the Orca project as it becomes ready.

The CoDA multi-agent system project [22], which focuses on organization and reorganization of autonomous oceanographic sampling networks, will provide a rich testbed for Orca. After being developed and tested in simulation (using the CoDA and CADCON [7, 2] systems), Orca will be fielded on our mobile land robots and, ultimately, aboard AUVs.

References

[1] P. E. Agre and D. Chapman. Pengi: An implementation of a theory of activity. In *Proceedings of the National Conference on Artificial Intelligence*, pages 268–272, Los Altos, California, 1987. Morgan Kaufmann.

[2] E. Albert, J. Bilodeau, and R. M. Turner. Interfacing the CoDA and CADCON simulators: A multi-fidelity simulation testbed for autonomous oceanographic sampling networks. In *Proceedings of the International Symposium on Unmanned Untethered Submersible Technology (UUST)*, Durham, NH, 2003. The Autonomous Undersea Systems Institute.

[3] J. Bellingham, C. Goudey, T. Consi, J. Bales, and D. Atwood. A second-generation survey AUV. In *Proceedings of the 1994 IEEE Symposium on Autonomous Underwater Vehicle Technology (AUV'94)*, pages 148–155, Cambridge, MA, USA, July 1994.

[4] J. G. Bellingham, T. R. Consi, R. M. Beaton, and W. Hall. Keeping layered control simple. In *Proceedings of the Symposium on Autonomous Underwater Vehicle Technology (AUV '90)*, pages 3–9. IEEE, 1990.

[5] R. P. Bonasso and J. Barratt. A reactive robot system for find and visit tasks in a dynamic ocean environment. In *Proceedings of the Eigth International Symposium on Unmanned Untethered Submersible Technology (AUV'93)*, pages 69–80. MSEL, Northeastern University, 1993.

[6] T. L. Briggs. Exploiting natural boundaries in the autonomous underwater vehicle domain to limit resource utilization. Technical Report 91-15, University of New Hampshire, December 1991. M.S. thesis.

[7] S. G. Chappell, R. J. Komerska, L. Peng, and Y. Lu. Cooperative AUV Development Concept (CADCON) - an environment for high-level multiple AUV simulation. In *Proceedings of the 11th International Symposium on Unmanned Untethered Submersible Technology (UUST)*, Durham, NH, August 1999.

[8] T. Curtin, J. Bellingham, J. Catipovic, and D. Webb. Autonomous oceanographic sampling networks. *Oceanography*, 6(3), 1993.

[9] R. J. Firby. An investigation into reactive planning in complex domains. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 202–206, Seattle, Washington, 1987.

[10] M. P. Georgeff and A. L. Lansky. Reactive reasoning and planning. In *Proceedings of the National Conference on Artificial Intelligence*, pages 677–682, Los Altos, California, 1987. Morgan Kaufmann.

[11] R. Komerska, S. G. Chappell, L. Peng, and R. Blidberg. Generic behaviors as an interface for communication, command and monitoring between AUVs. Technical Report 9904-01, Autonomous Undersea Systems Institute, 86 Old Concord Turnpike, Lee, NH, 1999.

[12] D. Long, M. Fox, L. Sebastia, and A. Coddington. An examination of resources in planning, 2000.

[13] D. N. Morley, K. L. Myers, and N. Yorke-Smith. Continuous refinement of agent resource estimates. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 858–865, New York, NY, USA, 2006. ACM Press.

[14] S. Smith, K. Ganesan, S. Dunn, and P. An. Strategies for simultaneous multiple AUV operation and control. In *IARP'96*, France, 1996.

[15] E. H. Turner. Organizing dialogue from an incoherent stream of goals. In *Proceedings of the Fifteenth International Conference on Computational Linguistics (COLING–92)*, pages 338–344, Nantes, France, 1993.

[16] E. H. Turner and T. L. Briggs. Responding to unanticipated goals when planning travel for autonomous underwater vehicles. In *Proceedings of of the Tenth IEEE Conference on Artificial Intelligence Applications (CAIA-94)*, pages 493–494, 1994.

[17] E. H. Turner, R. M. Turner, and E. Albert. Placing newly-arising goals in the proper context. In *Workshop notes for Cooperative Systems and Context Workshop of the Fifth International and Interdisciplinary Conference on Modeling and Using Context (CONTEXT-05)*, Paris, France, July 5–8 2005.

[18] R. M. Turner. *Adaptive Reasoning for Real-World Problems: A Schema-Based Approach*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1994.

[19] R. M. Turner. Intelligent control of autonomous underwater vehicles: The Orca project. In *Proceedings of the 1995 IEEE International Conference on Systems, Man, and Cybernetics*. Vancouver, Canada, 1995.

[20] R. M. Turner. Context-mediated behavior for intelligent agents. *International Journal of Human–Computer Studies*, 48(3):307–330, March 1998.

[21] R. M. Turner, D. R. Blidberg, S. G. Chappell, and J. C. Jalbert. Generic behaviors: An approach to modularity in intelligent systems control. In *Proceedings of the 8th International Symposium on Unmanned Untethered Submersible Technology (AUV'93)*, Durham, New Hampshire, 1993.

[22] R. M. Turner and E. H. Turner. A two-level, protocol-based approach to controlling autonomous oceanographic sampling networks. *IEEE Journal of Oceanic Engineering*, 26(4), October 2001.