

Adaptive Control of AUVs Using Onboard Planning and Execution

The Role of Artificial Intelligence Techniques In AUVs for Ocean Exploration

By Kanna Rajan

Principal Researcher for Autonomy

Frederic Py

Software Engineer

Monterey Bay Aquarium Research

Institute

Moss Landing, California

and

Conor McGann

Senior Engineer

Willow Garage

Menlo Park, California

Many of the complex multidisciplinary phenomena in the ocean that scientists seek to understand, such as blooms, riverine and estuarine plumes, and sediment transport processes, have unpredictable spatial and temporal expressions. Conventional methods of sampling and observing the ocean have relied on ship-based observations, which often require on-station presence. Such methods have given way to more cost-effective and capable exploration by autonomous underwater vehicles (AUVs). These can support a diverse array of sensors to resolve interacting physical, chemical and biological phenomena.

Current AUV control systems are a variant of reactive approaches that rely on manually scripted plans generated *a priori*. The controller is responsive to its immediate environment (e.g., passing through a front with a temperature gradient, changing the vehicle pitch to avoid an obstacle), generating actions while disregarding consequences. While such an approach is adequate, it prevents *in-situ* adaptation of mission plans for pursuing unanticipated

opportunities. Safe and effective adaptation requires a balanced consideration of mission objectives and the effects of adaptive actions, environmental conditions and available resources on the vehicle. In addition, as AUV mission durations increase, sustained presence in the ocean requires the ability to deal with atypical conditions. This calls for deliberation on alternative future courses of action while balancing imminent needs.

Probabilistic State Estimation

The autonomous systems group at the Monterey Bay Aquarium Research Institute (MBARI) has designed, developed and deployed an onboard adaptive control software system that integrates automated planning and probabilistic state estimation within a hybrid executive. Probabilistic state estimation integrates a number of scientific observations to determine the likelihood that the vehicle sensors perceive a feature of interest. Interleaved onboard planning and execution enable goal-directed commands within the context of projected mission state and allow the vehicle to replan for opportunistic events. In the process, human operators can describe their goals rather than encoding detailed control actions.

The software, called the Teleo-Reactive Executive (T-REX), is built around the paradigm of sense-plan-act and is deployed on MBARI's Dorado AUV. The vehicle uses its instruments to sense the environment, generates a plan or course of action for the future by deliberating about the best way to

accomplish desired goals and then acts on (or executes) this plan. T-REX's architecture handles short reaction times and extended deliberation with modules called reactors, each of which has a planning and execution component.

Deliberation of mission goals allows the software to reason how best to carry out the specified goals, taking into account the dynamic environment and the potentially changing state of the vehicle itself (e.g., diminishing battery power). Since plans rarely reflect the state of the environment at execution time, T-REX replans automatically when the perceived state of the environment does not match expectations encoded in a synthesized plan.

Reactors are differentiated on the basis of functional scope (indicating appropriate variables of concern for reasoning) and temporal scope (indicating the look-ahead window over which to plan and timing requirements to deal with how fast a reactor should reason).

Synchronizing the Reactors

Each control loop is embodied within a reactor that encapsulates all details of how to accomplish its control objectives, with the overall T-REX agent coordinating all of the concurrent control loops. For example, the fielded agent on the Dorado AUV has three reactors: a mission manager (MM) that plans and reasons about mission requirements, a navigator that manages the MM subgoals and determines navigation and localization for the vehicle, and an executive that dispatches commands to the low-level

functional layer of the vehicle to actuate control. The MM can reason about alternatives to the plan in the order of minutes, the executive has to resolve its goals with little or no delay and the temporal scope of the navigator lies in between.

All reactors use a single and explicit notion of time synchronized by an adjustable internal clock. This unit of time (a tick) signifies when synchronization of all reactors with the external state must occur, while between ticks reactors may deliberate.

Reasoning Response to System State

To reason about system state, T-REX uses a single domain model across all reactors that describes the vehicle's operational constraints. These constraints are written in a rich procedural language that can represent the notion of metric time and resources. Examples of constraints include: the AUV must surface to localize every 30 minutes; any two water samples must be separated by at least 500 meters; and there can be at most two water samples per transect of less than 3,000 meters.

State variables of the agent are represented by timelines that describe the evolution of the variable over time. Timelines, in turn, are a sequence of temporally qualified assertions, called tokens, that describe vehicle state with start/end times as intervals of time expressing temporal flexibility. This allows the model to capture the uncertainty, for example, in completing an ascent with a token $[e_s, l_s]\text{Ascend}[e_e, l_e]$ having variable early/late start and end time points. During deliberation, the token duration can be squeezed or expanded depending on the prevailing view of the plan; during replanning, the navigator could decide to remove a token to opportunistically retarget a feature of interest.

Timelines are distributed across reactors depending on their functional scope. One timeline is owned by one and only one reactor, is internal to it and controls the evolution of the timeline. Conversely, a different reactor can declare the same timeline as external, providing the possibility to observe its evolution and suggest new goals for its future value. The new goals are then transmitted to the owner of the corresponding internal timeline, which will deliberate to find a solution that satisfies these objectives. Such a plan can

impact other timelines owned by alternate reactors in the same way, implying a natural refinement from reactors with very abstract objectives toward low-level commands directly executable by the vehicle.

As time advances, each state variable encapsulates observations as tokens indicating the current state of the environment and the vehicle. This reflection of reality, in turn, might impact one or more planned actions. For example, a new state update might indicate a sudden change in battery power, or a conductivity, temperature, depth (CTD) instrument might indicate a sudden temperature gradient. An observation that is not consistent with expectations triggers replanning within the reactor in which divergence occurs. The reactor attempts to find a solution, and if a solution is not possible, plan failure is passed up the hierarchy of reactors to be handled by more abstract reactors until a new plan is generated or the failure results in a graceful degradation of the entire system.

Such an approach allows system degradation to be controlled, localizes plan formulation and problem-solving by partitioning the problem, and allows for modular software engineering. The architecture is open, permitting specialized implementations of reactors to be integrated.

Database Planning

Model rules are applied automatically through a combination of propositional inference and constraint propagation to check consistency and prune infeasible elaborations of the plan maintained in a temporal database. The database propagates the results of the synchronization of sensor data to the future. The dispatcher publishes goals on external timelines according to well-defined semantics. A temporal, constraint-based artificial intelligence planner with a rich legacy is used to deliberate over the specified temporal and functional scope of the reactor.

Ocean features in the water column are dynamic, unpredictable and generally difficult to track. To classify and enable online estimation, off-line machine learning techniques for feature detection are used. Statistical inference from previous AUV surveys are used for generating clusters of data

that are identifiably significant for an environmental feature set. These clusters are then integrated on board with incoming sensor data. Online classification based on real-time data then allows suitably thresholded sensor data to be marked as significant. Multiple and monotonic hits then activate the planner to dynamically replan to trigger instruments or alter the vehicle's navigation.

Results

T-REX has been used in a number of science missions for upper water column studies. The hardware setup on the AUV uses two onboard computers: a main vehicle computer, which is a 244 megahertz PC/104 stack running Winsystems Inc.'s (Arlington, Texas) QNX real-time operating system, and a 367 megahertz Winsystems EPIC EPX-GX500 AMD Geode stack running Red Hat Linux and T-REX. A socket-based communication protocol allows the exchange of command requests (i.e., goals) and state updates (i.e., observations) between T-REX and the functional layer.

A pioneering application of T-REX has been finding, mapping and sampling episodic intermediate nepheloid layers (INLs). These are thin sheets of suspended sedimentary particles that form on subsurface density interfaces and are transported from the seafloor to the surface. Scientists have used T-REX to precisely map and bring back water samples from INL hot spots for analysis using a water sampler integrated into the AUV. In doing so, T-REX has taken a high-level mission goal, broken it into component goals, detected the presence *in situ* of INL features and altered the vehicle's navigation to ensure high-density sampling is done within the INL. To cover more volume, wider spacings between transects are used in the absence of a feature signal.

In another application, T-REX was tasked with detecting a sediment plume laden with agricultural runoff that was flowing into Monterey Bay with every ebb tide. The constituents and strength of this plume are characterized by low salinity and high nitrate signatures. T-REX was used to map high-density transects within the plume and activate the water sampler with geometric constraints to appropriately separate sample acquisition.

More recently, a mapping transect

on the northern Monterey Bay shelf generated condensed CTD data which T-REX sent to shore via the Iridium (Bethesda, Maryland) satellite network. There it was analyzed by a scientist for temperature gradients indicative of ocean fronts. A specific front was targeted for higher resolution mapping, and high-level goals were dispatched from the scientist's desktop to direct autonomous front tracking and high-resolution mapping. T-REX received these goals and replanned on the fly by turning around to lock onto the front's hot spot using the shore-specified goals. This is the first known shoreside command of an AUV to date.

Conclusions

T-REX is an open-source, novel robotic controller that enables a new level of scientific observation in the ocean—to sample precisely in the right place and time. While T-REX was built for a specific underwater robotics application, it is also being used for terrestrial robot planning by Willow Garage, and in the near future, it will be used for the planning and execution of missions for a European Space Agency rover test bed.

The long-term goal is increased autonomy of AUVs for persistent ocean presence that would allow reconfiguration from hardware as well as software anomalies. This would allow for sustained ocean exploration for water column and seafloor studies. One can imagine the scientific promise of AUVs capable of adapting to opportunistic discoveries in the course of routine observations: These AUVs could alter mission plans dynamically and return with highly targeted data for analysis.

Acknowledgments

This work is funded by the David and Lucile Packard Foundation. Thanks to MBARI collaborators Tom O'Reilly, Hans Thomas, John Ryan, Thom Maughan, Brent Roman, Rob McEwen, Rich Henthorn, Chris Scholin, Bob Vrijenhoek, Larry Bird and Alana Sherman. Thanks to Yvan Petillot of Heriot-Watt University for comments on an early draft and to the crew of the R/V *Zephyr* for help with deployments.



Visit our Web site at www.sea-techology.com and click on the title of this article in the Table of Contents to be linked to the respective company's Web

site.

Kanna Rajan joined the Monterey Bay Aquarium Research Institute (MBARI) as the principal researcher for autonomy after a decade at NASA's Ames Research Center. He was the principal investigator for activity planning for NASA's Mars Expedition Rover and a principle of the Remote Agent Experiment, which flew the first artificial intelligence-based onboard planner in space.

Frederic Py is a software engineer at MBARI, which he joined after earning a Ph.D. from Laboratoire d'Analyse et d'Architecture des Systèmes. His area of interest is automated planning and verification of complex systems.

Conor McGann is a senior engineer at Willow Garage. He has previously worked at MBARI and NASA. His research interests are dominated by enabling robust autonomous robotic exploration of the physical world.

[Figure 1]

T-REX is made up of reactors with different functional and temporal scopes. More abstract reactors are deliberative, while lower-level reactors are reactive.

[Figure 2]

A T-REX-controlled AUV transects within an INL. Red indicates high probability of INL presence, and S1-S5 shows locations of 10 samplers triggered.

[Figure 3]

Tracking fronts with T-REX via science-driven commands from shore. The transects mix mapping and search phases driven by the scientist's intent to track dynamic water-column features.

[Figure 4]

The MBARI Dorado AUV and its support vessel, the R/V Zephyr.