

Adaptive Sampling of Thermoclines with Autonomous Underwater Vehicles

Nuno A. Cruz and Anibal C. Matos

INESC Porto and FEUP-DEEC

Rua Dr. Roberto Frias, 4200-465 Porto, Portugal

Email: {nacruz, anibal}@fe.up.pt

Abstract—The use of Autonomous Underwater Vehicles (AUVs) is providing the oceanographic community with very detailed ocean data at reasonable costs. Typical missions exploit their mobility and physical autonomy to achieve a wide coverage with minimal setup and operational costs. Most of this efficiency

In some scenarios, the ability to react to acquired data in real time is fundamental to acquire data in specific locations, without a-priori knowledge.

Demand to react to the environment

In this paper, we describe the

one of the examples is the ability to track the thermocline, specially designed for small scale variations of the thermocline, such as very shallow waters.

Data from field experiments show that the thermocline tracking behavior was able to maintain the vehicle in the vicinity of the thermocline, even in very shallow waters with badly pronounced features.

I. INTRODUCTION

Autonomous Underwater Vehicles (AUVs) are becoming ordinary tools for ocean sampling, with their mobility and physical autonomy being exploited to achieve wide coverage with minimal setup and operational costs. This yields improvements in efficiency as compared to traditional ocean sampling techniques, particularly when the main objective is to obtain a comprehensive 3-D view of a scalar field. On the contrary, if the objective is to search and map a specific underwater feature, then there is a complex trade-off between the dimension of the search grid, the velocity of the AUV and the total mission duration. This may result in a small percentage of useful data about the specific feature and moderate overall efficiency, which may be further stressed if the feature is dynamic. The concept of *adaptive sampling* has been suggested as a way to address this problem, by processing the environmental data in real time and commanding the AUV to sample the region of interest more densely. Although the concept is very appealing, it was only recently, with the exponential increase in computational power available on-board, associated with more mature software architectures, that there have been a few practical implementations.

One example of the utility of the *adaptive sampling* concept is the characterization of the thermocline – a vertical transition layer on the water column, separating the warm surface layer from the cold deep water below it. In a standard approach, the AUV would be programmed to follow a *yo-yo* pattern, resulting in very few data points in the vicinity of the thermocline. Alternatively, the ability to stay within the thermocline

region provides more *useful* data, but requires the AUV to process the payload data in real time and adjust the vertical controllers accordingly. For such an online implementation, it is important to adopt simple models for the feature and to compensate for model inaccuracy by using robust tracking algorithms and overall safety mechanisms. In a previous work [1], we followed this approach to develop an algorithm for tracking a thermocline, and we have demonstrated it on board the MARES AUV [2], a small sized vehicle developed at the University of Porto, in Portugal (Fig. 1). Using data from a CTD sensor, the AUV was able to detect the thermocline in real time, while moving in a standard *yo-yo* pattern. In this paper, we build on that previous work to increase the tracking performance and the utility of the manoeuvre and show how our *adaptive sampling* approach can be applied to maintain an AUV in the vicinity of a thermocline. We describe the algorithms implemented on board the MARES AUV and, finally, we demonstrate the tracking performance in the field, in a mission where the MARES AUV autonomously adapts the motion pattern to track a thermocline with space and time varying characteristics. The importance of detecting small scale variations on the thermocline parameters.




Fig. 1. The MARES AUV with an externally mounted CTD.

II. ON-BOARD ADAPTIVE SAMPLING

The concept of adaptive sampling in the ocean has emerged with the development of low cost robotic tools with increasing computational power and the ability to process large amount of information off-line in a short time. In fact, it is frequent to see *adaptive sampling* identified as the ability to *plan* new missions according to data arriving in real-time from a variety of sources (such as satellite data, CODAR seasondes or meteorological sensors, for example). In [3], for example, both AUVs and research vessels have been programmed to start missions as soon as a given triggering event has occurred. In this work, we consider adaptive sampling as the ability of the vehicle to process payload sensor data in realtime and decide *autonomously* the best trajectory to follow, in order to capture more information about a given oceanographic feature. This concept is sometimes referred to as *real-time adaptive sampling*.

There are many ocean processes that may be efficiently mapped using the adaptive sampling paradigm, particularly if they can be identified by sharp transitions or boundaries. Such efficiency will be accentuated if the processes show high spatial and temporal variations, since the vehicle trajectory can be set to concentrate measurements in the boundary region. Examples of these processes include thermoclines, fronts, eddies, and all sorts of plumes, such as emanating from hydrothermal vents or sewage outfalls.

Even though this is not a new concept, there have been very few practical implementations of real-time adaptive sampling. Some successful examples include the use of ROVs to track boundaries on the sea bottom ([4]), and, also, the use of AUVs for searching for the sources of chemical plumes, trying to mimic the real behavior of lobsters or bacterium in odor source localization ([5]–[7]). More recently, there has been some other thermocline tracking experiments, carried out on gliders ([8]) and other AUVs ([9], apart from the authors own implementations ([1]).

One of the main reasons for the shortage of implementations is vehicle safety. In standard AUV operations, the vehicles are programmed to move along geo-referenced trajectories that cover the area of interest. These trajectories are decomposed in sequences of elementary maneuvers, resulting in a *mission plan* to be executed by the on-board software. Although communications perform poorly underwater, there are many solutions for tracking the position of an underwater asset from a remote location and, therefore, it is relatively simple to detect any malfunction just by tracking the position of the AUV during a geo-referenced mission plan. On the contrary, during an adaptive sampling mission, the AUV *decides* the trajectory in real time and the vehicle motion directly depends on mapped features that cannot be easily related with the vehicle dynamics. It is, therefore, difficult to predict the vehicle motion and also to devise simple mechanisms to detect malfunctions and ensure the safe operation of the AUV.

In our case, we have followed a conservative approach to adapt the onboard software architecture of the MARES AUV

[2] to include the capability of real time adaptive sampling. The adaptive sampling maneuvers simply define reference inputs for the already existing real time controllers. This way, when an adaptive sampling maneuver is activated, the reference inputs for the control system are computed in real time by dedicated modules of software (that implement the adaptive sampling behaviors) and are fed to the controllers. In the particular case of thermocline tracking, the algorithm provides in real time the depth reference for the vertical controller, based on the acquired CTD data, as explained later. In any case, a protection layer of the control system sets limits for the reference values provided by the external modules, preventing the vehicle to enter unwanted regions and ensuring safety of operation.

III. APPLICATION SCENARIO – THE THERMOCLINE

In this work, we consider the scenario of detection and characterization of a thermocline in a given region. The thermocline is the transition layer in the water column that separates the warmer surface water from the cold deep water. The so-called *mixed layer*, closer to the surface, is the most easily influenced by atmospheric conditions (wind, rain and solar heat), resulting in a wide range of values throughout the year, while the deep-water layer is the less dynamic zone, with a slow variation with depth. In the thermocline, the water temperature drops as the depth increases, with a significant gradient as compared to the rest of the profile. Figure 2 shows an example of a temperature profile measured 2km off the Portuguese coast, in about 40m of water.

The thermocline can be permanent or seasonal and may develop both at sea [10] or in inner waters such as lakes and dam reservoirs [11]. The characterization of the thermocline is particularly relevant to marine biology, since the location and variability of the characteristics can provide valuable information about phytoplankton concentration and primary production [12]. The thermocline is also associated to water stratification, with strong impact on underwater acoustics and, therefore, its characterization can play an important role in underwater communications and military strategic planning. By *characterization*, we mean to gather enough data to provide a quasi-synoptic 4-D view of the thermocline.

Temperature is probably the most measured parameter in all water masses and given the huge amount of data that exists about the world oceans, there have been several attempts to represent these temperature profiles by simple sets of coefficients, in order to reduce the data warehousing requirements. This effort has been helped by the relatively regular *shape* of the profiles, that suggests a parametric representation. In any case, the proposed parametric models (for example [13]–[16]) require the availability of the full temperature profile to extract the parameters. Note that our idea is not to provide a very accurate model that approximates the true vertical profiles, minimizing some optimization criteria, but, instead, our main concern is that the model is simple enough to be iteratively estimated in real time. Even with only a partial vertical profile, the AUV has to be able to determine its main characteristics




Fig. 2. Temperature profile measured 2 km off the Portuguese coast, June 2009. The thermocline can be easily identified at about 11-14 meters of depth.




Fig. 3. Simplified temperature profile.

and maintain a trajectory in the vicinity of the thermocline. In order to achieve this, we assume a 3-layer piecewise linear model, similar to the one proposed by Haeger [13], with a high gradient at the thermocline, and low gradients both above and below. Figure 3 shows such a representation, where the three vertical zones can be easily identified.

IV. THE THERMOCLINE TRACKING MANEUVER

Overview

The thermocline tracking maneuver acts directly on the vertical controllers of the AUV, requiring a set of values for depth limits, which can be defined with respect to the surface or the bottom. In case the thermocline is not detected, the vertical motion of the AUV reverts to a standard *yo-yo* between those values. A specific process runs in the onboard computer, identifying the thermocline, filtering temperature data and fitting a gradient search algorithm implemented in real time. The thermocline is detected when the vertical temperature gradient exceeds a given threshold. The vehicle will proceed the vertical motion until the algorithm detects a *significant decrease* in the gradient, as compared to the current *maximum*. At this position, the vertical profile is reversed and the algorithm restarts on the other direction. In the absence of a positive identification of the thermocline in either direction,

then the vehicle will continue the vertical motion until the pre-specified limit is reached (either the surface or the maximum depth), as in a standard *yo-yo* motion. All thresholds are dynamic and depend on the gradient information acquired on previous profiles, so that the AUV can be used to track a time- or space-varying thermocline while moving along a pre-defined *lat-lon* trajectory.

During the thermocline tracking maneuver, the model parameters can be passed on to other AUV processes, if necessary. This allows, for example, the AUV to switch on any special sensor or to trigger an underwater sampler at the right instant, in order to capture a relevant sample of water for later laboratory analysis, such as suggested in [17].

The MARES AUV has 4 independent controllers for 4 degrees of freedom, and the vehicle is able to control the vertical velocity, from zero (*i.e. hovering*) to a maximum value around 40cm/s. This means that the vehicle can implement this algorithm to track the vertical thermocline at a single *lat-lon* location, or while following a completely independent *horizontal* trajectory. In most other AUVs, however, the vertical motion is obtained using horizontal fins or deflectors, in which case the vertical motion requires a minimum value for the horizontal velocity. In any case, the same principles described here can be implemented as a particular case of the *yo-yo* maneuver.

The current implementation of the thermocline tracking maneuver is the evolution of the first implementation, described in [1]. The main new features can be summarized as:

- all parameters are computed on line, requiring no *a-priori* estimation of the thermocline;
- the thermocline tracking manoeuvre has been extended to include additional parameters adjusted to the objectives for the mission (for example, extend the vertical span downwards to map possible chlorophyll or phytoplankton patches below the thermocline);
- An auto-diagnosis performance index is computed in real time and transmitted to a mission supervisor.

Tracking the thermocline

The algorithm developed for thermocline tracking can be described by the state machine represented in fig. 4. Note that the darker arrows represent the transitions that are expected during a normal tracking maneuver. These transitions will cycle the state machine through the most relevant states:

- TOP - The vehicle is located above the thermocline.
- TC2BOT - The AUV is within the region of the thermocline, on a downward vertical motion.
- BOTTOM - The vehicle is located below the thermocline.
- TC2TOP - The AUV is within the region of the thermocline, on an upward vertical motion.

The process usually starts when the vehicle is at the surface, so the first state is set to TOP. On entering this state, the depth reference is set to Z_{max} , which means the AUV will try to dive in the water column. The vehicle will then evaluate the vertical temperature gradient and compare it with a given threshold,

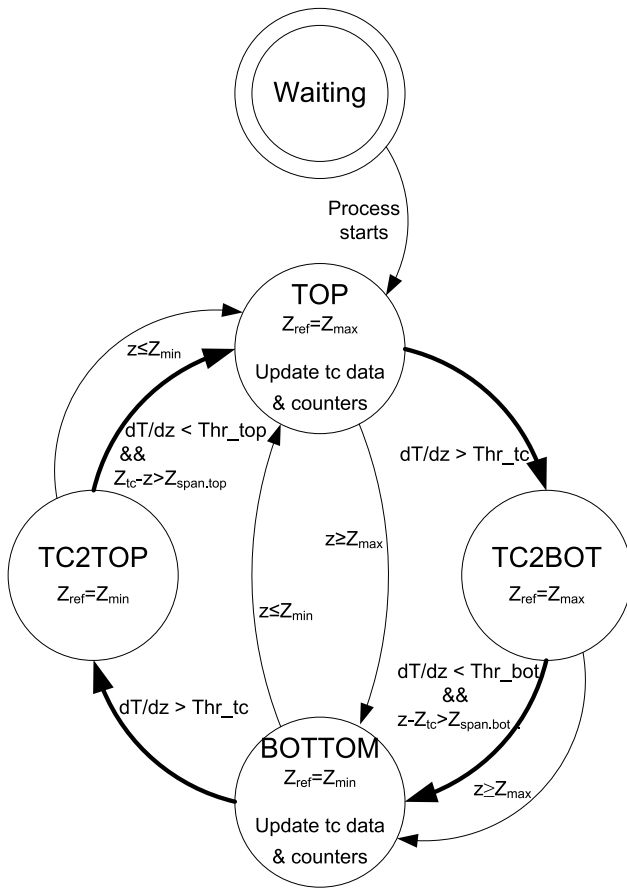


Fig. 4. State machine and transitions representing the thermocline tracking maneuver. The bold arrows represent the expected cyclic transitions during normal tracking.

Thr_{tc} . When this threshold is exceeded, the vehicle will assume the thermocline has been detected on the downward motion, entering the TC2BOT state. This can be seen as the *upper limit* of the thermocline region and then the vehicle will try to detect the *lower limit* of this region, by diving deeper. Therefore, the vertical direction will remain the same as before, with $Z_{ref} = Z_{max}$. At this state (TC2BOT), the gradient search algorithm will try to find if the level decreases below another threshold, Thr_{bot} . Note that this low gradient level has to be searched only for depths greater than the depth of the maximum gradient. In order to confirm the *lower limit* of the thermocline and avoid (early) false detections, an additional test is performed, verifying that the vertical span is large enough, *i.e.* if $z - Z_{tc} > Z_{span}$, where z is the current depth, Z_{tc} is the depth of the maximum gradient and Z_{span} his an optional parameter set by the user. When both these conditions are met, the vehicle enters the BOTTOM state.

Note from the state machine of fig. 4 that the BOTTOM state is also reached if $z \geq Z_{max}$. This is a safety mechanism to ensure that the maximum depth the AUV will be limited to Z_{max} , even if the algorithm is not able to positively *find* the thermocline. More, this condition contributes to a performance

index, since it signals a failure in the algorithm.

When the vehicle enters the BOTTOM state, the thermocline characteristics are extracted from the previous vertical profile (in particular, the maximum gradient and the thermocline limits) and this information is used to adapt the thresholds for the thermocline detection during the next vertical profile. At this state, the depth reference changes to Z_{min} , which means the AUV will now move towards the surface. The temperature and depth values arriving from the CTD sensor will be used to determine when the vertical gradient exceeds the new thermocline threshold, Thr_{tc} . This will change de state machine into the TC2TOP state, signaling the *lower limit* of the thermocline. In order to find the *upper limit*, the algorithm will proceed in much the same way as in the downward motion, but with a natural symmetry: the reference for the vertical controllers will be maintained at Z_{min} and the gradient search will be done for depths lower than the depth of the maximum gradient. Once again, the TOP state is only confirmed if a minimum vertical span has been covered.

As long as this process is active, the above cycle will be maintained, resulting in a *yo-yo* pattern around the thermocline. In case the thermocline is not detected in one of the vertical profiles, the vehicle will extend the vertical span up to Z_{min} or Z_{max} . It should be noted than on the very first time the algorithm runs (usually during the first descend), it is possible either to use *a priori* data to define the detection thresholds, or to use no information at all and let the AUV acquire a full vertical profile to determine those values.

V. PRACTICAL IMPLEMENTATION ISSUES

In order to estimate the temperature gradients, we cluster data points into bins and take the differences of the averaged values. The size of the bins is adjusted dynamically according to the thickness of the thermocline, the vertical velocity of the AUV and the sampling rate of the CTD. At an initial stage, the vehicle makes a full *yo-yo* profile between depth limits, in order to compute a first estimate of the relevant parameters for the algorithm: maximum gradient, depth bins, etc.

During normal tracking, the transitions between states are triggered by the detection of temperature gradients above or below specific thresholds. These thresholds are updated for each new profile, to allow for temporal and spatial variations of the thermocline parameters. All data from the thermocline model can be passed on to other AUV processes, if necessary. This allows, for example, the AUV to switch on any special sensor or to trigger an underwater sampler at the right instant to capture a relevant sample of water for lab analysis, such as suggested in [17]. We can also use this information to extend the vertical guidance of the AUV, towards the surface or towards the bottom. This is useful to correlate thermocline information with other ocean data, such as Chlorophyll concentration, for example. In terms of the state machine, the only difference is that the AUV needs to verify, not only the temperature gradient, but also if the vertical span of the motion has been exceeded in either direction.

The manoeuvre performance index is computed in real time, simply by counting the number of times the AUV has successfully detected the thermocline (by means of gradient thresholds) and the number of times the vertical limits have been reached without thermocline detection. If the AUV is not able to maintain tracking of the thermocline, this index gets too low and a proper action is triggered, according to instructions given by the user when preparing the mission (abort the mission or revert to a broader *yo-yo* survey).

there are 3 counters used to evaluate the performance of the algorithm, depending on the state transitions.

Iterative gradient estimation

When looking at a full temperature profile such as the example of fig. 2, one can clearly visualize a thermocline and imagine an online algorithm to extract 3 regions from that profile. However, this seemingly simple task is visually facilitated by the *long* low gradient regions, both above and below the thermocline. In a practical implementation, the challenge is to maintain the AUV in the region of the thermocline, detecting as early as possible a significant decrease in the temperature gradient and avoiding as much as possible to navigate within those *flat* regions. A difficult problem in detecting the thermocline is then to estimate the derivatives $\frac{dT}{dz}$ for a limited number of previous values of depth and decide if those derivatives are sufficient to conclude that the thermocline has already been passed and there is no need to proceed further. More, this data is updated several times per second (16, in our case), it may show small scale variations, it is not uniformly distributed and, surely, may have errors.

In order to estimate the gradient, we start by clustering the depth and temperature values into bins, as soon as they are available. For each new data point, only the derivatives (or differences, to be more accurate) affected by this data are updated. The derivatives are then estimated by the differences of the averaged bin values. In order to prevent biases in the average values of depth, we consider the average of both temperature values and depth values within a bin, as compared to considering an average temperature in the middle of the depth bin. This is particularly relevant if we increase the size of the bins, because as the vehicle ascends/descends in the water column, the first few samples within a depth bin are biased towards the bin boundary.

Robustness and performance

Even with the clustering of data into bins, there may exist some false detections of gradients if we only consider a single bin at a time. In order to avoid these false detections, we need a consecutive number of differences, above or below a threshold, to confirm the change of state. Note that if we increase the number of confirmations required to validate a low gradient, the main consequence will be for the AUV to dive further into the *flat* regions of the profile, with a minor impact on performance. On the contrary, if we increase the number of confirmations for the high gradients, we may not be able to detect a very thin thermocline.

The described algorithm provides information to monitor the performance in real time, by counting the times the thermocline boundaries are detected (the state machine of fig. 4 passes in the dark transitions) and the times the limits of depth are reached (both Z_{\max} and Z_{\min}). The percentage of detections is then a good performance metric and if it gets too low, a proper action may be triggered, according to instructions given by the user when preparing the mission. For example, the user can define that if the AUV doesn't detect a proper temperature variation, then the mission is aborted or, alternatively, the vehicle just reverts to a broader survey to try to reacquire the feature – a standard *yo-yo* pattern with user-defined vertical limits.

Magic numbers

The size of the depth bins is an important parameter for the algorithm, since it acts as a low pass filter which may affect the ability to detect gradients. If the bins are too small, only few data points will fall into that bin, resulting in large errors in gradient estimation. If the bins are too large, then they will smooth the temperature variations and the algorithm will have difficulty in separating significant variations. After several tests with both real and simulated vertical temperature profiles, we've found that we should have at least 5 to 8 bins spanning the depths of the thermocline to have significant changes in the gradients. As a rule of thumb, we usually set this value so that the first estimated thermocline spans about 10 bins.

The number of consecutive detections required to confirm a given threshold also depends on the size of the bins. For the size of our bins, we've considered a minimum number of 3 consecutive detections to confirm the gradient thresholds. We've also implemented a mechanism of validating a bin, depending on the quantity of data points used. For our experiments, we've considered a minimum number of 5 data points in a bin for its data to be used for gradient detection.

VI. FIELD TESTS AND RESULTS

A. Operational setup

The thermocline tracking maneuver was validated by a set of field trials carried out with the MARES AUV in a dam reservoir in the Douro river, close to Porto, in the north of Portugal. This location is only a 30 minute drive from our lab and it is known to develop a seasonal thermocline from late spring to early autumn. Just before the mission, we used a small support boat to deploy 2 buoys with acoustic beacons, located 350m apart, which provided a baseline for acoustic navigation.

difficult thermocline.

We started by running a vertical profile in the operation area, in order to confirm the thermocline and also to extract some initial data for the algorithm. This profile can be seen in fig. 5. Comparing this profile with a typical ocean profile (see for example, fig. 2), it is clear that the reservoir thermocline has much more small-scale variations, particularly in the mixed layer, above the thermocline. For the thermocline tracking




Fig. 5. Temperature profiles measured by the MARES AUV at approximately the same location of the Crestuma reservoir. The thermocline is visible in the summer profiles, reaching the surface in the last plot.

algorithm, this provides a much harder scenario than a smooth variation and, therefore, it is a good way of demonstrating its effectiveness. From the vertical profile, we estimate a maximum temperature gradient of around $1^{\circ}\text{C}/\text{m}$, with a very thin thermocline between about 6.1 and 7.3 meters of depth. As described earlier, taking these numbers, we chose a bin size of 12cm so that we would have a thermocline spanning around 10 bins. For a vertical velocity of about 10cm/s and a 16Hz sampling rate, this results in about 18 samples per bin.

The MARES AUV was then programmed to perform a *yo-yo* pattern from a minimum depth of 2 meters up to a depth of 12 meters, while the thermocline tracking algorithm was running in parallel. Fig. 6 shows the details of the depth during the *yo-yo* motion, the temperature readings and the corresponding state of the algorithm state machine, as described earlier: TOP (3), TC2BOT (2), BOTTOM (0), and TC2TOP (1). Note in the first graph the solid red line indicating the depth reference suggested by the thermocline tracking algorithm. It can be seen that the algorithm suggests an inversion in the direction of the vertical motion around $t = 95$ seconds, when the vehicle was at 7.9 meters of depth. On the upward motion, the algorithm also indicates the *upper limit* of the thermocline to be around 5.5 meters of depth.

B. Results

Although the AUV was not directly controlled by the thermocline tracking algorithm, it is possible to simulate the behavior of the vehicle if the depth reference suggested by the algorithm were transmitted to the vehicle controllers. We assume a simple kinematic model for the vehicle, which is a reasonable assumption since the vertical velocity is very slow. We also assume that the temperature profile is constant, which is also reasonable for such a short duration. Figure ?? shows the resulting profile if the vehicle were to use the information




Fig. 6. MARES AUV tracking a thermocline in the Crestuma reservoir, August 2010.

provided by the algorithm, superimposed to the original *yo-yo*. With the new algorithm, the initial 10 meters of vertical span are reduced to about 2.5 meters, resulting in a corresponding 4-fold increase in the sampling frequency of the thermocline (from the original 200 seconds period of the *yo-yo* down to 50 seconds). It is obvious that this increase can be much more significant when the ratio between the full water depth and the thickness of the thermocline layer increases, such as the case of deeper water scenarios.

refer to state machine to explain transitions

VII. CONCLUSIONS AND FUTURE WORK

The MARES onboard software architecture was adapted to incorporate the capability of real time adaptive sampling, providing a framework to implement innovative guidance algorithms, not only to track the thermocline, but also to efficiently characterize other ocean processes under the same paradigm.

In this work, we have implemented a new approach to AUV surveys that is to replace the standard predefined *mission* by an adaptive sampling strategy. Such an approach exploits the on-board computational power to infer some characteristics of the oceanographic environment and react to this environment by making decisions about the best sampling strategy to use. By continuously interpreting collected data, this decision can be made in real time so that the vehicle can use most of the available resources (mainly power) in sampling the ocean in regions of interest. Naturally, this will contribute to faster and more efficient surveys.

In our implementation, we assume that the thermocline can be coarsely approximated by a very simple model with time varying parameters. We've described how a small size AUV can change the sampling pattern in order to maintain tracking of a thermocline and therefore contribute to a more efficient characterization of a space- and time-varying thermocline.

We've developed a very low complexity algorithm that may be implemented in real time, with minor impact on CPU load and memory usage. Using an illustrative example, we've shown that the new maneuver allows the AUV to gather much more information about the thermocline than a standard *yo-yo*. This improvement may be accentuated in scenarios with a thin thermocline layer, particularly in deeper waters. We estimate that a similar approach can be followed to increase the efficiency in the sampling process of other ocean processes.

Even though the acquired data affects the vehicle motion pattern, the algorithm was implemented on the MARES AUV in such a way as to ensure safety of operation – in case the algorithm is not able to positively detect the thermocline, then the AUV will revert to a standard *yo-yo* pattern. This has been extensively tested with both real and simulated vertical temperature profiles, in order to make sure that the inclusion in the library of maneuvers will result in a safe operation. During the summer of 2010, the MARES AUV will be using this new maneuver to follow *lat-lon* transects while tracking the thermocline. At the same time, the vehicle will gather data from the optical sensors, in order to evaluate the correlation between thermocline characteristics and chlorophyll concentration. If required, the algorithm may be adapted for different sampling requirements, for example, the AUV may extend the bottommost part of the *yo-yo* to evaluate a phytoplankton layer trapped below the thermocline.

As far as the thermocline tracking is concerned, we expect to improve on the mechanism of initializing the algorithm, by calculating the size of the depth bins automatically. This will be done using information gathered on the first vertical profile and adjusting the bin size to maximize detectability. We also plan to improve the real time filtering of the parameters of the algorithm, so that the tracking mechanism may benefit more from the results of previous profiles, such as thermocline depth limits, detection thresholds, and temperature levels both above and below the thermocline.

the algorithm incorporates additional mechanisms of redundancy in the detection of the thermocline;

REFERENCES

- [1] N. Cruz and A. Matos, "Reactive AUV motion for thermocline tracking," in *Proc. IEEE Int. Conf. Oceans'10*, Sydney, Australia, May 2010.
- [2] —, "The MARES AUV, a modular autonomous robot for environment sampling," in *Proc. MTS/IEEE Int. Conf. Oceans'08*, Quebec, Canada, Sept. 2008.
- [3] O. Schofield, T. Bergman, P. Bisset, J. F. Grassle, D. B. Haidvogel, J. Kohut, M. Molin, and S. M. Glenn, "The long-term ecosystem observatory: An integrated coastal observatory," *IEEE J. Oceanic. Eng.*, vol. 27, no. 2, pp. 146–154, Apr. 2002.
- [4] C. Barat and M. J. Rendas, "Benthic boundary tracking using a profiler sonar: A mixture model approach," in *Proc. MTS/IEEE Int. Conf. Oceans'03*, San Diego, CA, USA, Sep. 2003, pp. 1409–1416.
- [5] J. A. Farrell, W. Li, S. Pang, and R. Arrieta, "Chemical plume tracing experimental results with a REMUS AUV," in *Proc. MTS/IEEE Int. Conf. Oceans'03*, San Diego, CA, USA, Sep. 2003, pp. 962–968.
- [6] S. Pang and J. A. Farrell, "Chemical plume source localization," *IEEE Trans. Syst., Man, Cybern. B*, vol. 36, no. 5, pp. 1068–1080, Oct. 2006.
- [7] W. Naeem, R. Sutton, and J. Chudley, "Chemical plume tracing and odour source localization by autonomous vehicles," *The Journal of Navigation*, vol. 60, no. 2, pp. 173–190, May 2007.
- [8] H. C. Woithe and U. Kremer, "A programming architecture for smart autonomous underwater vehicles," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems IROS'09*, St. Louis, MO, USA, Oct. 2009.
- [9] S. Petillo, A. Balasuriya, and H. Schmidt, "Autonomous adaptive environmental assessment and feature tracking via autonomous underwater vehicles," in *Proc. IEEE Int. Conf. Oceans'10*, Sydney, Australia, May 2010.
- [10] G. L. Pickard and W. J. Emery, *Descriptive Physical Oceanography – An Introduction*, 5th ed. Pergamon Press, 1990.
- [11] B. M. Johnson, L. Saito, M. A. Anderson, P. Weiss, M. Andre, and D. G. Fontane, "Effects of climate and dam operations on reservoir thermal structure," *J. Water Resour. Plann. Manage.*, vol. 130, no. 2, pp. 112–122, Mar./Apr. 2004.
- [12] J. Sharples, "Investigating the seasonal vertical structure of phytoplankton in shelf seas," *Mar. Model.*, vol. 1, no. 1–4, pp. 3–38, 1999.
- [13] S. D. Haeger, "Vertical representation of ocean temperature profiles with a gradient feature model," in *Proc. MTS/IEEE Int. Conf. Oceans'95*, San Diego, CA, USA, Oct. 1995, pp. 579–585.
- [14] P. C. Chu, Q. Wang, and R. H. Bourke, "A geometric model for the Beaufort/Chukchi Sea thermohaline structure," *J. Atmos. Oceanic Technol.*, vol. 16, no. 6, pp. 613–632, Jun. 1999.
- [15] P. C. Chu, C. Fan, and W. T. Liu, "Determination of vertical thermal structure from sea surface temperature," *J. Atmos. Oceanic Technol.*, vol. 17, no. 7, pp. 971–979, Jul. 2000.
- [16] V. Babovic and H. Zhang, "Modeling of vertical thermal structure using genetic programming," in *Proc. 7th OMISAR Workshop on Ocean Models*, Singapore, Sep. 2002. [Online]. Available: <http://sol.oc.ntu.edu.tw/OMISAR/wksp.mtg/WOM7/1.pdf>
- [17] Y. Zhang, R. S. McEwen, J. P. Ryan, and J. G. Bellingham, "An adaptive triggering method for capturing peak samples in a thin phytoplankton layer by an autonomous underwater vehicle," in *Proc. MTS/IEEE Int. Conf. Oceans'09*, Biloxi, MI, USA, Oct. 2009.