

Active Control of Marine Vehicles in the Presence of Strong, Dynamic, Uncertain Currents

Martina Troesch, Steve Chien

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
{firstname.lastname}@jpl.nasa.gov

Yi Chao, John Farrara

Remote Sensing Solutions
2824 East Foothill Blvd
Pasadena, CA 91107
ychao@remotesensingsolutions.com
farraraj@yahoo.com

Abstract

We address the control of a vertically profiling float using ocean-model-based predictions of future currents. While these problems are in reality continuous control problems, we solve them by searching a discrete space of future actions. Additionally, while the environment is a continuous space, the ocean model we use is a discrete cell-based model. We show that even with an imperfect model of ocean currents, planning in the ocean current model can significantly improve results for a specific problem of controlling a vertically profiling float when trading off remaining at the same location as a virtual mooring and collecting more data with more profiles.

Introduction

The study of ocean dynamics is an important problem with many ramifications including environment and climate, food production, defense, and leisure. There are a variety of different techniques currently being used to monitor and measure ocean conditions such as currents, salinity, and temperature. Recent advances in robotic marine vehicles such as floats, gliders, and autonomous underwater vehicles (AUV) provide a plethora of new tools and techniques for ocean measurement. Another important way that scientists are able to collect measurements is by using moored buoys. This allows scientists to collect data at a fixed location over time. However, physically mooring a buoy is a significant financial investment and the mooring is only able to collect data at a fixed location.

As an alternative, a *virtual mooring* has been proposed in which a dynamically controlled vehicle uses a control policy in order to maintain its position. Specifically, one proposal is to deploy a vertical profiling float to the location of desired data collection and to use predictive ocean models to plan a control sequence for changing depths that best keeps the float near the same location using the ocean currents. This method has multiple benefits over using a float at a physical mooring. First, the float could be retrieved and re-deployed when desired. Second, there is more flexibility in that the float could be programmed to track a moving target or to drift to facilitate deployment or retrieval. Third, the de-

ployment would be less expensive than building a physical anchor location.

Although using an AUV would provide better control to remain at a fixed location, more capable vehicles are more expensive - the following table shows approximate costs for families of marine vehicles.

Asset	Control	Speed	Endurance	Cost
Floater	None	None	Weeks	\$100's
Vertical Profiling Float	Vertical	None	Years	\$10K's
Seaglider	Horizontal	1 kts	Months	\$50's - \$100'sK
AUV	Horizontal & Vertical	5 kts	1 Day - Weeks	\$100K - \$M

We believe that deploying a vertical profiling float is a cost effective alternative to using physical moorings or AUVs by using predictive ocean models to generate a control sequence that allows the float to act as close as possible to a virtual mooring considering the ocean currents and that when flexibility is available, a trade-off can be made to collect more data profiles and float away from the mooring location.

Problem Definition

This experiment generates a control sequence for a vertical profiling float to make it act as a virtual mooring. Given a start location, a goal location, a mission duration, and a planning model of ocean currents, the goal is to produce an optimal sequence for a float with a particular diving and surfacing speed in the planning model and to verify that it does well in the nature model, based on a desired objective function. To verify that the plan does well in the nature model, the goal is to show that planning a control sequence performs better than simply performing even profiling.

The start location represents the drop-off location of the float and the goal location represents the location of the virtual mooring, which is equivalent to the desired data collection site. The mission duration is the length of time for which a control sequence will be generated. In an ideal situation, a data measurement would be collected at all times and at all depths at the same latitudinal and longitudinal location. However, given a single float, a trade-off has to be made between profiling across all depths and drifting away from the

mooring location. Therefore, the objective function should take this into account and make a compromise between collecting data at different depths or staying close to the desired location. The ocean models contain current information at varying depths, which is used to determine where the float will move during the mission duration. Since predictive ocean models cannot perfectly forecast the ocean currents, the planning model represents this imperfect knowledge and the nature model represents the actual ocean currents. In this way, it is possible to approximate a scenario of generating a control sequence in an ocean model and executing it in the ocean.

Experimental Inputs


For this experiment, the vertical profiling float was given a vertical speed of 0.12 m/s taken as the speed of the operational Electromagnetic Autonomous Profiling Explorer (EM-APEX) (Sanford et al. 2005) floats used in prior deployments. The maximum profiling depth was chosen to be 500 m. The mission duration was set to be one day. The start location and the goal location were (latitude, longitude) locations at the surface of the ocean. They were set to be the same location so that it represents a float that is placed at the location of desired data collection. In particular, nine different start/goal locations in the Monterey Bay Area were chosen for this experiment. These locations, as well as their latitude and longitude values, are shown in the following image and table.



latitude:	36.872217	36.072217	35.272217
longitude:	-123.734853	-123.134853	-122.534853

The ocean model used for this experiment was the Regional Ocean Modeling System (ROMS) (Chao et al. 2009; Li et al. 2006; Farrara et al. 2015), a discrete, cell-based model. For our experiments we used a ROMS California

coast configuration with a grid size of approximately 3 km in the latitudinal and longitudinal directions, 1 hour in the time dimension over 72 hours, and fourteen depths ranging from 0 to 1000 m in non-uniform intervals. In order to approximate a scenario of planning with the predicted ocean model and executing the control sequence in the ocean, two different models were used. Using the standard practice in ocean modeling, the planning model uses information from ten days before the desired day of execution without HF radar assimilation. The nature model incorporates the HF radar assimilation and represents the conditions of the ocean on the desired day of execution and is the best model of the ocean on that day. The models for this experiment were generated for April 17, 2015. The following figure demonstrates how the zonal and meridional currents vary between the two different models. The zonal currents are in the west-east direction and the meridional currents are in the north-south direction. The currents are shown for location (34.66°lat, 235.68°lon) at the surface over 72 hours.



The objective function was designed to trade-off profiling and remaining close to the mooring location and to favor one or the other by changing the relative weights between the terms. The equation for the objective function is

$$\sum_n \sum_d (w_T T_d + w_D D_d)$$

where n are the nodes in the path, d are the depth choices, and w_T and w_D are weighting terms. T_d is the number of seconds since the last time the path was at depth d divided by 1000. D_d is the distance in kilometers that the path was from the goal location the last time the path was at depth d .

In other words, at each time step the most recent node at the surface and at the profiling depth are found and $(w_T T_d + w_D D_d)$ is added to the objective function. Since T_d can be used as a proxy for determining the time since the last profile, a smaller w_T to w_D ratio favors control sequences

that keep the float closer to the desired data collection site, whereas a higher w_T to w_D ratio favors control sequences that have more profiles.

Algorithm

Even though, in reality, the problem space is continuous, in order to make the search tractable the future actions of the float are determined in a discrete manner. At each time step, the float can be at either the surface depth or the deepest profiling depth. In other words, if the float is at the surface it can either stay at the surface or move to the profiling depth and if the float is at the profiling depth it can remain there or return to the surface. If the float stays at a depth, the duration is equivalent to the amount of time required to move between the surface and the deepest depth. This means that each time step is equal in length. One profile is defined to be one trip down and up again.

The current information that was used to calculate the motion of the float was determined based on the float's latitudinal and longitudinal location, its depth, and the time. If remaining at the same depth, the data from the model cell at the closest latitude, longitude location and previous hour index of the beginning of the time step was used. If moving through depths, the time to move through the discrete depths that are available in the ROMS model is calculated such that a new latitude, longitude location and time is calculated for the arrival at each depth. It is that data that is used to choose the cell from which to extract the current information in the same manner as when staying at the same depth: the closest latitude, longitude cell is chosen and the time index is for the previous hour. Currently, no smoothing is being applied between the grid points, neither in location nor in time.

The algorithm used to search for the control sequence is an A* algorithm with a zero heuristic estimator. A path is considered finished once the mission duration has been reached. Once a path has reached the length of the mission duration, the score is saved as the best path score and the queue is flushed to finish all remaining paths while they have a score that is less than the best finished path so far. If a finished path has a score that is better than the best path, that score becomes the best path score. Since the queue is completed to expand all paths with a lower score after finding the first complete path, this algorithm is optimal.

The following pseudocode summarizes the execution of the algorithm.

Experimental Procedure

At each location, the algorithm described above was run using the planning model with a ratio of w_T to w_D ranging from 0.4 to 20.0 in steps of 0.4. Two more control sequences were also run on the planning model. The first required the float to remain at the surface and the second required the float to continuously profile. Once the optimal control sequences were determined in the planning model, these sequences are then executed in the nature model. The surface and continuous profiling sequences were also run in the nature model. In order to analyze whether planning for the control sequence in the planning model improves the goal of

```

Q ← start location
while Q not empty do
  curPath ← lowest objective score path in Q
  curPath add node at surface and profiling depth
  if curPath duration > mission duration then
    bestPath ← curPath
    bestScore ← curPath score
    break
  end if
  Q.push(curPath)
end while
while Q not empty do
  remainingPath ← lowest objective score path in Q
  if remainingPath score > bestScore then
    break
  end if
  curPath add node at surface and profiling depth
  if curPath duration > mission duration then
    bestPath ← curPath
    bestScore ← curPath score
  else Q.push(curPath)
  end if
end while


```

remaining close to the virtual mooring location, a baseline control sequence needed to be chosen as a comparison. As our baseline we chose evenly spaced profiles at each available number of profile choice. Specifically, since 10 full profiles can fit within the mission length, the baseline sequences were 1 through 9 evenly spaced profiles. Remaining at the surface (0 profiles) and full profiling (10 profiles) were calculated separately, as described previously.


Empirical Evaluation in Simulation

Each location and weight combination results in a control sequence for the float. This means that for each location, there are fifty different control sequences that make different trade-offs between profiling and remaining near the goal location. However, due to the nature of the values of the ratio of the weights that are applied to the objective function, many control sequences are identical among similar weights. Furthermore, due to the discrete selection of weights, not all numbers of possible profiles are generated at each location.


The following image shows, as an example, the path that was generated using the optimal path at location (35.272217°lat, -123.734853°lon) and weight ratio 2.4 in the planning model and the nature model. The blue path is the path in the planning model and the red path is in the nature model.



The next figure shows the control sequence for that path as well as the distance and sum of distances at each time step. In other words, the middle graph shows the instantaneous distance from the goal and the bottom graph shows the instantaneous sum of the distances so far from the goal. This particular example shows that in the planning model, the path results in 5 full profiles, ends at a distance of approximately 2 km, and the sum of the distances from the goal location at each step in the path ends at almost 17 km.



The next figure shows this same control sequence starting at the same goal location executed in the nature model. The depth profile verifies that the same control sequence was used. The distance from the goal at the end of the path is now approximately 15 km and the sum of the distances is approximately 160 km.



In this particular example, the execution in the nature model resulted in an end distance and sum of distances greater than those in the planning model. However, this may not always be the case. It is expected that sometimes the currents in the nature run will be more favorable and sometimes they will be less favorable to keeping the vertical profiling float near its goal location, but that on average the difference will be close to zero. What this would indicate is that the planning model is a good predictor for the nature model and thus provides effective planning information.

Let D_{End} be the difference between the distance from the goal at the end of the mission in the planning run and the nature run in kilometers. A positive difference means that the distance in the nature run was larger than that of the planning run.

Let D_{Sum} be the difference between the sum of the distance from the goal at each node in the path in the planning run and the nature run in kilometers. A positive difference means that the sum of the distances in the nature run was larger than that of the planning run.


Avg D_{End}	Stdev D_{End}	Avg D_{Sum}	Stdev D_{Sum}
5.63	6.63	54.60	64.47

These results show that the previously held expectation does not hold for our sample size. This indicates that the planning model lacks information and is not a perfect predictor, even on average. This is further confirmed by observing the correlation coefficient between the zonal currents in the planning and nature models as well as the meridional currents in the two models. The correlation coefficient for the two currents in the planning and nature models at the surface depth and first time index was calculated using the cells in the search region around the example starting location. These coefficients are shown in the table below.

Zonal Correlation Coeff	Meridional Correlation Coeff
0.1881	.4390

Those small, positive correlation coefficients substantiate the claim that the planning model provides some predictive knowledge, but it far from a perfect predictor for the nature model.

The following graph shows the results for all of the weights at the same location used in the previous example: (35.272217°lat, -123.734853°lon). To properly convey the results in terms of the objective function, it would be best to show the distance from the goal and use the time since the float was at a particular depth. However, this time value is difficult to conceptualize in a concrete manner, and since it is a proxy for the number of profiles that are achieved, it was decided to display the results in terms of the number of profiles instead of the time term in the the objective function. The open circles are the results in the planning model and are labeled by the weighting ratio w_T/w_D that resulted in that result. The open squares show the results of remaining at the surface or purely profiling in the planning model. The closed circles are the results of using the control sequences generated in the planning model in the nature model and are also labeled with the weighting ratio that was used in planning the sequence. The closed squares show the results of the reference control sequence of evenly spaced profiles in the nature model.



This particular example demonstrates how the span and discrete selection of the weighting ratio does not result in all possible number of profiles and that multiple weighting ratios result in the same control sequence. This graph also shows that for this location, for all weights, the distance in the nature model was greater than that in the planning model and that the baseline control sequence always did worse than the planned sequence in the nature model.

Let D_{RefEnd} be the difference between the distance from the goal at the end of the mission in the nature model with the reference control sequence and the planned control se-

quence. A positive difference means that the distance using the reference control sequence was larger than that of the planned sequence.

Let D_{RefSum} be the difference between the sum of the distance from the goal at each node in the path in the nature model with the reference control sequence and the planned control sequence. A positive difference means that the distance using the reference control sequence was larger than that of the planned sequence.

The following tables summarize the results in terms of D_{RefEnd} and D_{RefSum} over all locations and weighting ratios. Paths from control sequences that resulted in the same number of profiles were compared to each other. Since not all weighting ratios achieve all possible number of profiles, there are a different number of paths that are compared for the different profile numbers. The column headed "Num" shows this number. What these metrics demonstrate are whether, when there is flexibility as to when the float will profile, planning for the timing of the profiles in an ocean model provides for better virtual mooring than evenly spacing out the profiles, which is like doing a random plan.

Profiles	Num	Avg D_{RefEnd}	Stdev D_{RefEnd}
1	7	0.62	8.70
2	4	2.70	3.62
3	7	-5.87	8.62
4	7	2.72	4.63
5	12	3.35	2.44
6	14	3.14	2.51
7	16	1.21	2.13
8	19	0.56	1.60
9	14	-0.41	1.51

Profiles	Num	Avg D_{RefSum}	Stdev D_{RefSum}
1	7	13.38	87.88
2	4	31.57	53.71
3	7	-47.83	81.84
4	7	23.50	40.12
5	12	24.27	20.11
6	14	20.12	23.28
7	16	1.67	21.19
8	19	1.57	14.12
9	14	-0.60	8.96

The values in these tables indicate that, in general, planning is beneficial for remaining near the goal location instead of using evenly spaced profiles, as was expected from the previous analysis of the correlation coefficients. Looking at the results for 3 and 9 profiles, the distances are negative, meaning that planning did worse than the baseline for these two cases. One possibility is that more data needs to be collected, possibly over different locations and different days, in order to smooth out the noise and get a better representation of the problem, and that on this particular day, the nature model had particular characteristics that were not captured in the planning model for those paths. More data collection would need to be done to analyze this anomaly further.

Related Work


Path planning for underwater vehicles has been widely studied; however, there has been relatively little work in float

planning. A notable exception is (Dahl et al. 2011) which examined the problem of optimizing coverage across the oceans for a large number of floats, but only considered a constant depth and a greedy algorithm. Much more research has been done on glider planning, where there is some control for choosing a direction of motion, but it is less than the current velocity. (Thompson et al. 2010) also uses the ROMS model, but calculates reachability envelopes using wavefront propagation for glider path planning. The work in (Eriksen et al. 2001) describes Seaglider, a glider that is manually controlled from the shore, and is sometimes controlled to maintain position. No ocean model similar to ROMS was used. (Alvarez, Garau, and Caiti 2007) also does not use an ocean model, but instead uses synthetic data with general algorithms to control a set of floats and gliders. Like the work in this paper, (Rao and Williams 2009) uses an A* graph search algorithm; however, that work assumes that currents change slowly with time and compute the path across many nodes in a single time step, whereas we have many time steps within a single cell. Instead of trying to remain near a specific location, (Pereira et al. 2013) focuses on gliders that are attempting to avoid surfacing in dangerous areas, such as shipping lanes. (Grasso et al. 2010) focuses on the prediction of the glider location, analyzes the accuracy of the predictive model, and uses a physics based control model. Using an asset with more control, (Cashmore et al. 2014) explores the problem of autonomously maneuvering near a site for inspection using an AUV with probabilistic modeling for uncertainty.

Future Work

There are future extensions to further validate the model-based planning approach to sea asset control. Smoothing, or interpolation, of the grid points in the model could be performed when determining the current velocity at a location and time. Different types of goal locations could be used, such as a line, a grid, or even a moving point. Similarly, the start location could be different from the goal location. The deployment cost and retrieval cost could be incorporated into the overall objective function. The search could include asset selection with an asset cost metric, e.g. the asset could be changed from a float to a glider or an autonomous underwater vehicle to add more flexibility to the control sequence when such an asset is available. Multiple assets of different types could be planned for at once with different goal locations. The algorithm could be changed to be an incomplete or suboptimal algorithm in order to decrease the run time. Another scenario could be to re-plan the control sequence with every surfacing of the asset. This would simulate a scenario in which the most recent, real information about the location of the asset could be used to create a better control sequence during the execution. Finally, we are exploring use of these algorithms in experiment design - in which a range of locations within a specific region of interest are considered and a virtual mooring location is chosen that is expected to be location maintainable. The figure below shows the variability of the expected float drift based on the deployment location from a prior deployment to the Monterey Bay region. The blue paddle indicates the start loca-

tion and the green path shows where the float was projected to drift over time.



Conclusions

Based on these results, it is clear that modeling the ocean for the purpose of planning trajectories is a very difficult problem. We saw that the currents in the planning model and the nature model are weakly, positively correlated, so although the planning model does not perfectly capture the ocean movement in the nature model, it is a useful tool that leads to improved control sequences, as was demonstrated by the results of the planned control sequences compared to the baseline sequences. This method could also be used to help in the selection of a virtual mooring site in order to find the location where it is least likely that the float will be carried away from the goal location.

Acknowledgments

Portions of this work were performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

- Alvarez, A.; Garau, B.; and Caiti, A. 2007. Combining networks of drifting profiling floats and gliders for adaptive sampling of the ocean. In *Robotics and Automation, 2007 IEEE International Conference on*, 157–162. IEEE.
- Cashmore, M.; Fox, M.; Larkworthy, T.; Long, D.; and Magazzeni, D. 2014. Auv mission control via temporal planning. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, 6535–6541. IEEE.

Chao, Y.; Li, Z.; Farrara, J.; McWilliams, J. C.; Bellingham, J.; Capet, X.; Chavez, F.; Choi, J.-K.; Davis, R.; Doyle, J.; et al. 2009. Development, implementation and evaluation of a data-assimilative ocean forecasting system off the central california coast. *Deep Sea Research Part II: Topical Studies in Oceanography* 56(3):100–126.

Dahl, K. P.; Thompson, D. R.; McLaren, D.; Chao, Y.; and Chien, S. 2011. Current-sensitive path planning for an underactuated free-floating ocean sensorweb. In *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, 3140–3146. IEEE.

Eriksen, C. C.; Osse, T. J.; Light, R. D.; Wen, T.; Lehman, T. W.; Sabin, P. L.; Ballard, J. W.; and Chiodi, A. M. 2001. Seaglider: A long-range autonomous underwater vehicle for oceanographic research. *Oceanic Engineering, IEEE Journal of* 26(4):424–436.

Farrara, J. D.; Chao, Y.; Zhang, H.; Seegers, B. N.; Teel, E. N.; Caron, D. A.; Howard, M.; Jones, B. H.; Robertson, G.; Rogowski, P.; and Terrill, E. 2015. Oceanographic conditions during the orange county sanitation district diversion experiment as revealed by observations and model simulations. Submitted to *Estuarine, Coastal and Shelf Science*.

Grasso, R.; Cecchi, D.; Cococcioni, M.; Trees, C.; Rixen, M.; Alvarez, A.; and Strode, C. 2010. Model based decision support for underwater glider operation monitoring. In *OCEANS 2010*, 1–8. IEEE.

Li, P.; Chao, Y.; Vu, Q.; Li, Z.; Farrara, J.; Zhang, H.; and Wang, X. 2006. Our ocean—an integrated solution to ocean monitoring and forecasting. In *OCEANS 2006*, 1–6. IEEE.

Pereira, A. A.; Binney, J.; Hollinger, G. A.; and Sukhatme, G. S. 2013. Risk-aware path planning for autonomous underwater vehicles using predictive ocean models. *Journal of Field Robotics* 30(5):741–762.

Rao, D., and Williams, S. B. 2009. Large-scale path planning for underwater gliders in ocean currents. In *Australasian Conference on Robotics and Automation (ACRA), Sydney*. Citeseer.

Sanford, T. B.; Dunlap, J. H.; Carlson, J.; Webb, D. C.; Girtton, J. B.; et al. 2005. Autonomous velocity and density profiler: Em-apex. In *Current Measurement Technology, 2005. Proceedings of the IEEE/OES Eighth Working Conference on*, 152–156. IEEE.

Thompson, D. R.; Chien, S.; Chao, Y.; Li, P.; Cahill, B.; Levin, J.; Schofield, O.; Balasuriya, A.; Petillo, S.; Arrott, M.; et al. 2010. Spatiotemporal path planning in strong, dynamic, uncertain currents. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 4778–4783. IEEE.