

Slocum Glider Energy Measurement and Simulation Infrastructure

Hans C. Woithe*, Ilya Chigirev[†], David Aragon[‡], Murium Iqbal[†], Yuriy Shames[†]
Scott Glenn[‡], Oscar Schofield[‡], Ivan Seskar[†], Ulrich Kremer*

*Department of Computer Science, [†]WINLAB, [‡]Institute for Marine and Coastal Sciences,
Rutgers University

Abstract—Autonomous underwater vehicles (AUVs) are indispensable tools for marine scientists to study the world's oceans. Depending on their missions, AUVs are equipped with advanced sensors (sonar, cameras, acoustic communication, bio-sensors), have on-board computers for data analysis (image analysis, data compression), and are capable of on-board decision making (resource planning, swarming). Since AUVs operate solely on battery power, power and energy management is a crucial issue. Mission-critical tradeoff decisions have to be made between energy consumption and sensing, data processing, and communication activities. Mission planning has to consider these tradeoffs when provisioning resources for expected future events, or when dealing with changing environmental conditions such as weather, water currents, and seafloor profiles. Effective power and energy management requires knowledge about the actual energy consumption of each active component within the AUV. Effective planning requires simulators that can predict energy consumptions based on expected future events and environmental conditions.

In this paper, we discuss the design and implementation of a power measurement infrastructure for the Teledyne Webb research Slocum glider. This infrastructure can be used for online power/energy management or to better understand the time-dependent energy consumption profile of the active glider components during a particular mission. We also discuss the design of a new simulation environment for the Slocum glider which uses the power/energy data obtained by our measurement infrastructure, in addition to seafloor and coastal radar information. We illustrate the effectiveness of the new tools in the context of planning a glider flight across the continental shelf off the coast of New Jersey.

I. INTRODUCTION

The mission endurance of today's Autonomous Underwater Vehicles (AUVs) depends highly on the capacity and usage of the vehicle's batteries. Typically, missions for the Slocum Electric Glider last about 30 days [8]. Longer missions, such as the 221 day mission to cross the Atlantic by RU27 from Rutgers University [10] are possible through an increase in the number of batteries and through the careful planning of the usage of the vehicle's devices. Such planning is also crucial

for shorter missions when gliders are equipped with advanced sensors such as an Acoustic Doppler Current Profiler (ADCP) or acoustic underwater communication.

With the recent integration of the coulomb meter into the glider, measuring the discharge of the battery has become more accurate. Knowing the rate at which energy is used and how much remains is vital to mission planning. However, the glider's coulomb meter only measures whole vehicle current. To perform more precise mission planning, being conscious of the energy consumption of individual components is necessary. We have developed a measurement infrastructure which captures the currents drawn from distinct components of the Slocum Glider. The infrastructure has been deployed in test missions off of the coast of New Jersey, and the data collected have been integrated into a Slocum Glider simulator. Our measurement board and simulation framework can be used to assist in the planning and decision making of missions and shows possible tradeoffs, for instance, between mission duration, speed, and energy consumption.

II. MEASUREMENT INFRASTRUCTURE

We have created a measurement infrastructure to measure and record the electric current drawn by individual devices of the Slocum glider. The infrastructure consists of a measurement board and a data logger. The design philosophy in creating the infrastructure was to not compromise the safety of the vehicle, even if quality of the resulting measurements are affected. The




Fig. 1. Measurement board mounted on a weight bar used for ballasting the Slocum glider.

glider components measured are: the main, external, and emergency power, the buoyancy pump and brake, and the pitch and fin servos.

The measurement board, shown in Fig 1, was intended to be housed above the glider's mainboard in the aft section of the vehicle. However, due to the different space constraints between different generations of gliders, the board was moved to the center payload bay. This allows the board to be quickly uninstalled and re-equipped onto another glider.

The board makes use of eight Hall Effect sensors which do not interfere with the vehicle's current flow. This ensures that in the event of sensor failure, the glider will continue to operate normally. Three 20A sensors are used for the main, external, and emergency power, while two 5A sensors are used for the buoyancy pump. Three 3A sensors are used for the buoyancy pump brake, pitch servo and fin servo. The sensors were over-provisioned for safety, but still allow the capture of large spikes in the current.

The microprocessor used in our design was the PIC16F767. The processor typically operates at less than 2mA at 8MHz. It contains eleven 10-bit analog-to-digital (A/D) channels of which eight are in use to measure the currents drawn by the glider using the Hall Effect sensors. The microprocessor has been programmed to use interrupts to generate constant samples at 32ms intervals. These samples are transmitted to the glider's science bay processor.

The measurement board communicates its samples via a 9,600 baud serial connection to the science bay processor. The stock 6.38 software version of the glider's science computer software has been retrofitted to record the samples produced and transmitted by our measurement board. The science processor, a CF1 from Persistor Instruments Inc. [6], is typically clocked to run at 3.68MHz using the stock software, but is usually run at higher clock speeds when collecting data as part of our infrastructure.

The power consumption of the science processor is shown in Table I. The CF1 as programmed during our previous deployments, was clocked at 14.72MHz which consumed approximately 520mW more power than the stock software release. We have since developed optimizations in the logging process to reduce energy consumption. These improvements allow the clock speed to be lowered down to 3.68MHz provided that the mission specifications allow for the trade off of four second, instead of two second, sampling from the conductivity, temperature and density (CTD) sensor. The CTD is a standard sensor on a Slocum glider.

The power consumption of the measurement board

TABLE I
CF1 POWER CONSUMPTION

Software	Clock Rate (kHz)	Power (Watts)
Stock 6.38	3680	0.19
Deployed 6.38	14720	0.71
Development 6.38	3680	0.35
Development 6.38	7360	0.49

TABLE II
MEASUREMENT BOARD POWER CONSUMPTION

Description	Channels	Power (Watts)
Deployed	8	0.76
Development	6	0.58




Fig. 2. Assessment of the measurement board's accuracy using a Tektronix MSO4034 oscilloscope.

itself is shown in Table II. After deploying gliders equipped with our infrastructure we have found little use in measuring the external and emergency current. The external power supply is only active on bench, so it is unnecessary for a board which will be deployed at sea. In the event of a emergency, the safe recovery of the vehicle is of higher priority than collecting good data. The presence of the emergency sensor could not be justified for the additional power it consumes. By removing the two Hall sensors, 180mW of power was saved. An additional benefit comes from the fact that the measurement board and the science processor now measure, transmit, and log less samples allowing for more data to be collected.

The measurement infrastructure has been extensively tested to ensure that recorded current samples are representative of the actual events. Fig 2 shows the results of a test where a current of 1A was applied to one of the sensors for approximately six seconds. The event was measured and logged by a Tektronix MSO4034 oscilloscope as well as a PC connected to our measurement board. The results of these experiments indicate that the samples collected are within the expected error of the sampling rate, A/D conversion and the sensors




Fig. 3. Current draw of the fin servo during a “wiggle.”




Fig. 4. Current draw of the pitch servo during a “wiggle.”

themselves.

The infrastructure as installed on the gliders records all measured samples. Without compression, data can be recorded for mission lasting up to 26 days. However, multi-week missions using this revision of the board with alkaline batteries are not feasible due to the significant energy overhead. It may be possible if lithium batteries are used instead. In future work we hope to significantly reduce the power dissipation of the system so that full length deployments are possible.

III. DEPLOYMENTS AND MEASUREMENTS

The measurement infrastructure has been installed and deployed on two Slocum gliders. It has been used to collect current readings of the vehicles on the bench top as well as at sea. The sea trials took place off of the coast of New Jersey in September of 2009 and in February of 2010.

Before the measurement infrastructure was trusted to be deployed, it was installed in the glider and extensively tested on the bench top. To ensure the vehicle components still performed up to par in the presence of the measurement board, the vehicle’s motors were subjected to “wiggle” tests. This entails moving its motors through their full range of motion. Sample results of such a wiggle of the fin and pitch servos are depicted in Fig. 3




Fig. 5. Current drawn from the buoyancy engine during our deployment. The flight profile is shown together with the current draw of the buoyancy pump. It can clearly be seen that the buoyancy engine activities align with inflection points, and that the power consumption at depth is significantly higher than near the surface.

and Fig. 4, respectively. The Hall Effect sensors used for these devices are bipolar so the reported currents show the current flow in both directions as the servos move the opposite direction. The fin is used to steer the vehicle, and the pitch motor is used to fine tune the vehicle to the commanded pitch by moving an internal battery pack. The power draw of these two motors is generally very low, and during a mission motor activities typically occur in brief bursts. Through wiggle, overnight, and weekend tests the system was deemed stable and reliable for sea tests.

The first sea trial involved two short mission segments of approximately thirty minutes in length each. The glider was instructed to perform yos, sequences of dives and climbs, between 1 and 20 meters. The glider depth profile along with the current draw of the buoyancy pump of one mission segment are illustrated in Fig 5. The glider never reached a depth of 20 meters because the seafloor was not sufficiently deep enough at the deployment location. The experiences gained in the sea trial were used to prepare the infrastructure for a longer term mission.

The second deployment was a 6.5 day mission in early February of 2010. A map of the glider’s path is shown in Fig 6. The mission’s goal was to fly to the continental shelf to gather buoyancy engine readings at depths of up to 100 meters. The mission was cut short due the combination of inclement weather and the high power consumption of our measurement infrastructure. After heading east toward the shelf for two days, the vehicle was commanded to head north because a Nor’easter storm was expected to push the vehicle south. After being forced south for two days, it was again commanded to head east towards the shelf to gather readings at deeper




Fig. 6. Flight path of the mission deployed with the measurement infrastructure in February of 2010.




Fig. 7. Current draw of the buoyancy engine during an inflection at approximately 12 meters.

depths for a short time. Unfortunately, another Nor'easter was imminent so the mission was aborted and the glider spent the remaining time flying back to shore to be retrieved.

The buoyancy engine of the electric Slocum glider consists of a buoyancy pump and a brake mechanism. The pump moves a piston to change the vehicle's buoyancy by altering its displacement of water. The brake locks the pump's position in place which would otherwise be forced to retract due to water pressure. The current draw of the buoyancy engine is shown in Fig. 7. When commanded to inflect from a dive to a climb, or from a climb to a dive, the brake first unlocks the pump. The pump follows by moving the piston to the commanded position. When the position is reached, the brake again locks the pump's position in place.

The energy used by the buoyancy pump increases with depth because the pump must work harder to battle the additional pressure. This was confirmed by our first sea trials, Fig. 5, where inflections from a dive to a climb state used more energy when the inflections occurred at three, six and twelve meters. Fig. 8 depicts the measured




Fig. 8. Energy use of the buoyancy pump at depth.

energy used by the pump during the deployment in February of 2010. The energy used for similar depths in the two seal trials were comparable considering different gliders were used. In both missions, however, the energy necessary for the pump to perform inflections from a climb to a dive at shallow depths is at times less expensive than the cost associated with the brake. Having detailed knowledge of the cost of components is important when trying to optimize vehicle flights.

IV. SIMULATOR

To assist in planning of future missions, we have created a simulator for the Slocum glider. The simulation environment incorporates energy, speed, seafloor and sea surface current models, and is used to predict the flight path, longevity and energy usage of a mission. The simulation environment has been validated against Teledyne Webb's Shoebox simulator and compared to a deployment to the continental shelf off of the coast of New Jersey.

A. Implementation

The longevity of missions performed by AUVs rely on the limited energy resources the vehicle carries on board in its batteries. This resource limit can effect the quality of missions. Missions which require the vehicle to maintain a constant presence at a location or require traveling to an area of interest are constrained in the amount of information they can collect. For the aforementioned reasons, we have developed and implemented new energy models in our simulation environment.

The energy models were formulated from the samples recorded by our infrastructure along with the voltages reported by the glider. Our simulator uses models for the buoyancy pump, brake and steady state load, where no motor and most devices are not in use. The average observed cost associated with the brake is applied at




Fig. 9. Speed distribution derived from over four years of glider flights.

every inflection point. The expense of inflections near the surface where the vehicle state changes from a climb to a diving state is modeled as a constant. Inflection performed at depth from a diving to a climbing state use a linear cost function. The function has been fitted to the data points from the deployment in February 2010. The function is shown in Fig. 8 labeled "Predicted Energy Use." The energy used in simulated missions is dependent on the vehicle's pitch angle and speed. The pitch angle of the flight impacts the number of inflection points, and thus the use of the buoyancy engine. The speed determines the amount of time required to complete the mission.

The simulation environment makes use of two types of speed models. The first is a model similar to that of the Slocum glider's shoebox simulator. The Shoebox, named after its physical similarities to a shoe box, contains the essential glider electronics to perform simulations in real time. The software running in the Shoebox is the same software used during deployments but makes use of simulated device drivers. The speeds and missions generated by this model when used in our simulator should be similar to that of the commercial Shoebox. However, unlike the Shoebox, our framework is able to simulate missions significantly faster than real time.

The second speed model integrated into the simulator is based on speed distributions which were empirically derived from over four years worth of glider flight data. The flights took place off of the New Jersey coast between the years of 2003 and 2009. The resulting distributions are shown in Figure 9 and were constructed by measuring the distance covered in each dive segment and the time necessary to travel the segment. A dive segment starts when a glider submerges and ends when it resurfaces. The 25° distribution was comprised of 2,539 segments, covering 6,263km over 293 days, while the 26° distribution span 16,411 segments, 32,527km and

3.48 years. Sufficient data to build speed distributions were available only for 25° and 26°, which are the most common angles used by the Slocum Glider. These speeds are sampled by the simulation environment to produce realistic over-the-ground speeds. Although very similar, the 26° distribution is slightly faster than the 25° distribution. Along with the dive and climb pitch angles specified by the mission, the depth rate is calculated and used to position the glider in space. The depth rate and the seafloor determines the number of inflections that occur during flight.

The simulation environment also supports the use of a seafloor terrain. The seafloor model used may be artificial, come from prior deployments as measured by a glider, or can be interpolated from NOAA's National Geophysical Data Center's (NGDC) bathymetric data set [7]. The current data set (from the NGDC) used by the simulation environment is of the coast of New Jersey at a resolution of one arcminute. The addition of a seafloor model improves the quality of the vehicle's predicted energy usage especially in shallow waters. Simulated open ocean deployments, or deployments where it is known that the glider will never reach the seafloor will not benefit from a seafloor model, and could therefore be removed for such missions.

Time dependent sea currents can significantly impact the flight profile of a glider, and are therefore modeled within our simulation framework. The currents may be artificially and dynamically generated, or can be interpolated much like the seafloor. The use of Coastal Ocean Radar (CODAR) [5] data from Rutgers University has been integrated into the framework. This data describes the sea surface currents of the New Jersey area at a spatial resolution of six kilometers and a temporal resolution of one hour. The addition of sea currents add another degree of realism which should improve the prediction quality.

Our simulation framework can be used to analyze past glider flights, support active deployments, or help to plan future missions. CODAR information is valuable when simulating past flights and can be used in the decision making of active missions. For example, if recent sea surface current data is available, it can be used to predict the location of where the glider may resurface next. With the utilization of weather trend or prediction models, such as the Regional Oceanic Modeling System (ROMS) [9], the simulator could also forecast the general outlook of missions.

B. Validation

To validate our simulation infrastructure we have compared its predictions to that of Teledyne Webb's




Fig. 10. Validation of the simulation environment with respect to the Shoebox simulator.

Shoebox simulator. The mission executed on both the Shoebox and our simulation framework entailed three yos (a sequence of climbs and dives) between 2 and 25 meters.

The depth profile of the simulations are shown in Fig. 10. The Shoebox profile describes the flight as performed by the Shoebox simulator. SimShoebox and SimDist are the flight profiles generated by our simulation environment. SimShoebox generates speeds similar to that of the Shoebox, while SimDist sampled speeds from the distribution in Fig. 9.

The time necessary to complete the missions for Shoebox and SimShoebox are very similar. Like the Shoebox, our simulated vehicle also slightly overshoots the commanded depth limits. On average, the SimShoebox is slightly slower, taking several seconds longer to complete the mission. The results produced are however a reasonable representation of what may be generated by the manufacturer's simulator. The advantage of our simulation framework lies in the runtime necessary to produce the simulated mission. The Shoebox took approximately 15 minutes to simulate the sample mission, while SimShoebox required only 0.35 seconds on a 2.2GHz dual core processor.

The flight simulation which applies the speed distribution model, SimDist, requires an additional 380 seconds longer in mission time than both the Shoebox and SimShoebox. The simulation took 0.5 seconds. This suggests that the vehicles speed is on average slower using this model than that of the Shoebox. We believe the speed model based on speed distributions is more accurate than the Shoebox model since it is derived from over four years worth of vehicle flight time. The speed distribution model should then not be compared to that of the Shoebox simulator but against an actual deployment.

V. DEPLOYMENT SIMULATIONS

To validate the simulator and its speed distribution model, a deployment from September 2009 is compared to similar flights in our simulation environment. The goal of the original mission was to fly to the continental shelf from the coast of New Jersey and back at 26° . The flight path of the mission is illustrated in Fig. 11(a). Due to strong currents for portions of the mission, the glider was pushed south preventing it from making steady forward progress towards the target waypoint. An operator interfered with the flight and changed the target waypoint due west back to shore before the vehicle reached the commanded waypoint near the continental shelf. Waypoints were changed further throughout the mission, causing the vehicle to reach none of the target waypoints except the last which was used to collect the vehicle. The total length of the deployment was 14.84 days. To validate the simulation framework a similar deployment length should be achieved.

A. Baseline

The baseline simulation assumes that no seafloor or currents exist in the environment. Consequently, the runtime needed to simulate the mission is small, but the predicted mission will also not be very accurate. The simulated mission flown in the remainder of this section will be that of Fig. 11(a) except that the vehicle will be commanded to keep flying until it has reached all its waypoints. It is difficult to reenact the intentions or reasoning behind the operator's actions so they are ignored.

The SimShoebox simulation of the mission predicts a mission length of 7.9 days, with the energy usage of 707kJ and a flight path as depicted in Fig. 11(b). A runtime of 20 seconds was needed to simulate the mission. SimShoebox, which has been shown in the previous section to be fairly representative of the Shoebox simulator, would suggest a real time simulation of 7.9 days. If the speed distribution is used instead in the simulation (SimDist), the mission length increases to 11.5 days, 785kJ and a runtime of 86 seconds. SimShoebox in this scenario has erroneously estimated the mission length by 6.94 days while SimDist by 3.34 days. Unlike the previous validation experiment, the speed distribution produces a better estimate when compared to a real deployment.

B. Seafloor Model

To add a layer of realism, the simulation environment can use a seafloor as previously described. Instead of flying to the full commanded depth, the vehicle must inflict several meters above the seafloor to avoid impact.




Fig. 11. (a) The flight path of the mission being simulated. (b) The flight path of the baseline and seafloor simulations. (c) The simulated mission using both seafloor and CODAR data.

TABLE III
SPEED DISTRIBUTION SIMULATION RESULTS

Mission	Seafloor	Currents	Time (days)	Energy (kJ)	Runtime (min)
Actual	N/A	N/A	14.84	N/A	N/A
Baseline	No	No	11.5	785	1.4
Seafloor	Yes	No	11.5	984	5.7
CODAR	Yes	Yes	14.89	1,235	20

This will increase the total number of inflections points in the mission which directly translates into more energy use because the buoyancy engine is activated at each inflection. The mission length and flight path for both SimShoobox and SimDist remain nearly identical to the baseline but the energy usage increase to 892kJ and 984kJ, respectively. The modeling of the seafloor is paramount so that missions may be more accurately predicted and planned for.

C. Seafloor And CODAR Models

The final model supported by our framework is that of the sea currents. The CODAR sea surface currents of the days surrounding the deployment of Fig. 11(a) were integrated and applied to the simulated mission. The flight map of SimDist is shown in Fig. 11(c). The SimDist mission flew for 14.89 days using 1,235kJ of energy and required 12 minutes to simulate. SimShoobox's mission flew for only 8.88 days, used 986kJ, and had a runtime of 5 minutes.

The presented simulation results indicate that the speed distribution model was more representative of the deployment in Fig. 11(a) than that of the speed model which is similar to the Shoobox. A summary of the simulations for SimDist is listed in Table III. The final SimDist mission using both the CODAR and seafloor resulted in a mission time slightly longer than that of

the real deployment. This is however expected since the simulated mission flew a slightly different mission where the vehicle actually reached the waypoints and was not interrupted by an operator. Modeling the supervision as part of the mission is a difficult task because the intentions of the operator at the time are not known. Errors associated with the spatial and temporal resolution of the seafloor and CODAR data also add to the difficulty of recreating the original mission.

VI. RELATED WORK

The modeling of underwater gliders has been extensively studied in the previous work [1][2][3][4][11]. The primary focus lies in the formulation of hydrodynamic models that try to closely emulate the vehicle as it flies through the water. Our work differs in approach in that we use simpler mathematical models and make use of years of glider flight information. Our simulation environment also incorporates empirically derived energy costs of a subset of the vehicle's devices to assist in the planning and prediction of deployments.

VII. FUTURE WORK

We have begun the design of the second revision of the measurement infrastructure. Using the knowledge gathered from the deployments using the first revision,

we are choosing more appropriate sensors for each measured component of the glider. Unlike the first revision, we plan to have the ability to customize the rate at which current readings are sampled and logged. This would also entail adding triggers so that only samples of interest are recorded, saving precious storage which would otherwise go to waste recording silent or noisy data. We will also add the ability to log data locally on the board while still maintaining the capability to send a subset of data to the science computer for transmission to shore via a satellite modem. Running the measurement board on a separate power source is also desirable so that longevity of the mission and the samples themselves are not influenced by the presence of the board. Finally, we would like to expand the number of glider components we monitor. Other devices such as the air pump (used to breach and keep the glider at the surface), and the iridium satellite modem use a great deal of energy over time due to the length and frequency that the vehicle surfaces. New sensor payloads such as an ADCP or an acoustic modem will require careful power and energy management as well.

In the future we will continue to improve the simulation environment by expanding and implementing more complex models. Integration of ROMS [9] may become an essential component to aid in the prediction and planning of future flights. Additionally, although energy models from the components we have measured have already been incorporated into the simulator, the focal point thus far has been the energy usage of the buoyancy pump. Refining the energy characteristics of the other devices in the simulator would lead to more accurate mission predictions. Once complete, we would like to simulate the deployment from February 2010 and compare the simulated energy usage against the actual energy usage measured by our measurement board.

During the development of the next revision of the measurement infrastructure we will continue to improve the current system to prepare it for additional deployments. Specifically, we are planning to fly to the edge of the continental shelf to gather more buoyancy pump samples at depths of up 100 meters, which is the operating depth of most of our gliders. The additional samples could help determining the accuracy of the buoyancy pump's energy cost function.

Finally, integration of the presented work with our previous work [12] is underway. The simulation framework will be used to execute and analyze missions specified in our programming language. The complete system will allow for the development and testing of complex missions. The Linux single board computers we have

integrated into the Slocum glider allow for simulations to be run online and may assist in the steering of the vehicle so that dead reckoning error may be reduced.

VIII. CONCLUSION

We have described the implementation of a measurement and simulation infrastructure for the Slocum Glider. Energy cost models derived from two sea trials have been incorporated into the simulation environment. Using over four years of previous glider flight data we have constructed distributions to define the vehicle's speed over ground. Along with the use of a sea floor data set from NOAA, and sea surface current data from Rutgers University, we were able to simulate a vehicle's flight path, mission time, and energy usage. The framework has been evaluated against a simulator produced by the Slocum glider's manufacturer as well as a deployment off the coast of New Jersey. Simulation results indicate that the framework can produce sensible mission estimations with low computation costs.

Acknowledgement: This research has been partially funded by NSF awards CSR-CSI #0720836 and MRI #0821607.

REFERENCES

- [1] M. Arima, N. Ichihashi, and Y. Miwa. Modelling and motion simulation of an underwater glider with independently controllable main wings. pages 1–6, May 2009.
- [2] A. Bender, D.M. Steinberg, A.L. Friedman, and Williams S.B. Analysis of an autonomous underwater glider. 2008.
- [3] P. Bhatta and N.E. Leonard. Stabilization and coordination of underwater gliders. volume 2, pages 2081–2086 vol.2, Dec. 2002.
- [4] Joshua G. Graver, Ralf Bachmayer, and Naomi Ehrich Leonard. Underwater glider model parameter identification. In *Proc. 13th Int. Symposium on Unmanned Untethered Submersible Tech.*, 2003.
- [5] CODAR Corporate Headquarters. Codar. Mountain View, CA. <http://www.codar.com>.
- [6] Persistor Instruments Inc. Cfl computer system. Marstons Mills, MA. <http://www.persistor.com>.
- [7] National Oceanic and Atmospheric Administration. National geophysical data center. <http://www.ngdc.noaa.gov/>.
- [8] Teledyne Webb Research. Slocum glider. Falmouth, MA. <http://www.webbresearch.com/slocum.htm>.
- [9] A.F. Shchepetkin and J.C. McWilliams. The regional oceanic modeling system (roms): a split-explicit, free-surface, topography-following-coordinate oceanic model. volume 9, pages 347–404, 2005.
- [10] Rutgers University. The scarlet knight's trans-atlantic challenge. <http://rucool.marine.rutgers.edu/atlantic/>.
- [11] C.D. Williams, R. Bachmayer, and B. deYoung. Progress in predicting the performance of ocean gliders from at-sea measurements. pages 1–8, Sept. 2008.
- [12] H.C. Woithe and U. Kremer. A programming architecture for smart autonomous underwater vehicles. pages 4433–4438, Oct. 2009.