Implementation of an Embedded Sensor Network for the Coordination of Slocum Gliders for Coastal Monitoring and Observation

Ryan N. Smith[*] [†]
Robotic Embedded Systems Laboratory, USC
ryannsmi@usc.edu

Jnaneshwar Das
Robotic Embedded Systems Laboratory, USC
jnaneshd@usc.edu

Hordur Heidarsson
Robotic Embedded Systems Laboratory, USC
heidarss@usc.edu

Arvind Pereira
Robotic Embedded Systems Laboratory, USC
menezesp@usc.edu

Gaurav S. Sukhatme
Robotic Embedded Systems Laboratory, USC
gaurav@usc.edu

ABSTRACT

Autonomous Underwater Vehicles (AUVs) are revolutionizing oceanography through their versatility, autonomy and endurance. However, they are still an underutilized technology. For coastal operations, the ability to track a certain feature is of interest to ocean scientists. Adaptive and predictive path planning requires frequent communication with significant data transfer. Currently, most AUVs rely on satellite phones as their primary communication. This communication protocol is expensive and slow. To reduce communication costs and provide adequate data transfer rates, we present a hardware modification along with a software system that provides an alternative robust disruption-tolerant communications framework enabling cost-effective glider operation in *coastal* regions. The framework is specifically designed to address multi-sensor deployments. We provide a system overview and present testing and coverage data for the network. Additionally, we include an application of ocean-model driven trajectory design, which can benefit from the use of this network and communication system. Simulation and implementation results are presented for single and multiple vehicle deployments. The presented combination of infrastructure, software development and deployment experience brings us closer to the goal of providing a reliable and cost-effective data transfer framework to enable real-time, optimal trajectory design, based on ocean model predictions, to gather *in situ* measurements of interesting and evolving ocean features and phenomena.

Categories and Subject Descriptors

C.2 [**Computer-Communications Networks**]: Network Protocols, Network Operations, Local and Wide-Area Networks; J.2 [**Physical Sciences and Engineering**]: Earth and Atmospheric Sciences; B.4.1 [**Data Communications**]: Transmitters and Receivers

General Terms

Design, Implementation, Algorithms

Keywords

Autonomous Underwater Vehicles, Embedded Sensor Networks, Freewave, SLOCUM Glider, Trajectory Design.

1. INTRODUCTION

Effective observation and continual monitoring of a dynamic system as complex as the ocean cannot be done with one instrument in a fixed location. A combination of static and mobile sensors must be deployed, and the information gleaned from the collected data must be distributed to all members of the fleet. Successfully orchestrating a multi-sensor, long-term deployment additionally requires a robust, rapid and cost-effective communication network. Only when all of these components are in synchronous operation can scientists begin to make new discoveries and improve our overall understanding of the complex ocean environment.

We present the beginnings of such an infrastructure to aide the investigation of *coastal* ecosystems in Southern California. The goal of this paper is to present a reliable and cost-effective data transfer framework to enable real-time, optimal trajectory design, based on ocean model predictions, to gather *in situ* measurements of interesting and evolving ocean features and phenomena.

Driven by the underlying science, our collaborative research group strives to implement the necessary components to facilitate coastal ocean observation by combining autonomous systems and current predictive capabilities. Our mission is to track and collect daily information about an ocean process or feature which has a lifespan on the order of a week.

---

[*] Corresponding author.

[†] Department of Computer Science, University of Southern California, 3710 S. McClintock Avenue, Los Angeles, CA 90089, USA.

An ocean model is used to predict the behavior of an interesting artifact, *e.g.,* a fresh water plume, over a short time period, *e.g.,* one day. This prediction is then used to generate a sampling plan for the deployed sensor(s). The sampling plan is received and executed by the sensor platforms. The collected data is assimilated into the ocean model and an updated prediction is computed. A new sampling plan is created and the process repeats until the artifact is out of range or is no longer of interest. This entire procedure hinges upon reliable and rapid data transfer between the sensor platform(s) and a remote base-station.

In Section 2, we present an actual oceanographic phenomena of interest in Southern California, which motivates the presented research efforts. We follow this with a short description of the selected ocean predictive model and test-bed mobile sensor platform(s). We follow this with a detailed description of the design and implementation of a radio network that spans the coastal regions offshore of Los Angeles. We include the infrastructure specifics and the modifications performed to the test-bed sensor platform. Section 3.2 presents data on the current coverage of this network. Section 4 outlines one proposed usage of the installed radio modem network and provides some initial results from field deployments. We conclude with a discussion of the results, and comment on ongoing and future work.

2. BACKGROUND AND MOTIVATION

The motivation for this work comes from the desire to understand the complex dynamics and processes that occur in a *coastal* ocean environment through the use of a network of sensor platforms. To this end, the Marine Biology and Computer Science Departments at the University of Southern California have formed the Center for Integrated Networked Aquatic PlatformS (CINAPS, pronounced [sin-aps]) in a collaborative effort to monitor and observe the coastal ecosystem in Southern California, see [12].

Of particular interest to the CINAPS team is the formation and evolution of Harmful Algal Blooms (HABs). Combining an ocean prediction model, an array of static and mobile sensor platforms and an embedded communication network, we can perform continuous observation of the Southern California Bight (SCB)[1], which will lead to better predictions on when and where HABs may occur.

2.1 Harmful Algal Blooms

Microscopic organisms are the base of the food chain; all aquatic life ultimately depends upon them for food. Of these organisms, there are a few dozen species of phytoplankton and cyanobacteria that can create potent toxins when provided with the right conditions. Harmful algal blooms can cause harm via toxin production, or by their accumulated biomass. Such blooms can cause severe illness and potential death to humans as well as to fish, birds and other mammals. The blooms generally occur near fresh water inlets, where large amounts of nutrient rich, fresh water is deposited into the ocean. This water provides the excess food to support higher productivity and a *bloom* of microorganisms. It is of interest to predict when and where HABs may form, and which coastal areas they may affect. Harmful algal blooms are an active area of research along the western coast of the United States and are of large concern for coastal communities in Southern California. Impacts of HABs in the SCB are presented in [8, 14].

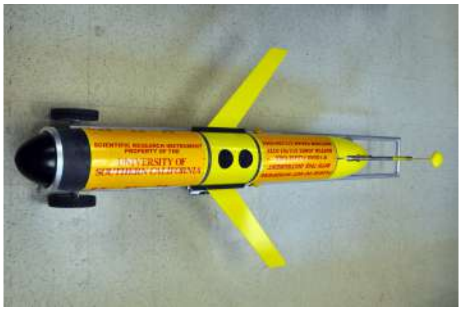

Figure 1: He Ha Pe, one of two USC *SLOCUM* gliders.

2.2 Regional Ocean Modeling System

The predictive tool utilized in this study is the Regional Ocean Model System (ROMS) [10] - a split-explicit, free-surface, topography-following-coordinate oceanic model. We use ROMS because it is an open-source, ocean model that is widely accepted and supported throughout the oceanographic and modeling communities. Additionally, the model was developed to study ocean processes along the western U.S. coast which is our primary area of study.

The version of ROMS used in this study is compiled and run by the Jet Propulsion Laboratory (JPL), California Institute of Technology. The JPL provides ROMS nowcasts and hourly forecasts (up to 36 hours) for Monterey Bay, the SCB and Prince William Sound, see [7] for more information. The JPL version of ROMS assimilates HF radar surface current measurements, data from moorings, satellite data and any data available from sensor platforms located or operating within the model boundary. The more *in situ* data assimilated into ROMS, the better the predictive skill of the model. Information regarding this version of ROMS and the data assimilation process can be found in [4].

2.3 Mobile Sensor Platform

The mobile sensor platform used in this study is a Webb *SLOCUM* autonomous underwater glider, as seen in Fig. 1 (http://www.webbresearch.com). The *SLOCUM* glider is a type of AUV designed for long-term ocean sampling and monitoring [9]. These gliders *fly* through the water by altering the position of their center of mass and changing their buoyancy. Due to this method of locomotion, gliders are not fast moving AUVs, and generally have operational velocities on the same order of magnitude as oceanic currents. The endurance and velocity characteristics of the glider make it a good candidate vehicle to track ocean features which have movements that are determined by currents, and have the capability to remain of scientific interest for weeks at a time.

We utilize autonomous gliders because our collaborative research group owns two of them, and hence field experiments

---

[1] The SCB is the oceanic region contained within 32° N to 34.5° N and $-117°$ E to $-121°$ E

can be readily performed. We have upgraded the communication capabilities of our gliders to make them a node in our local area wireless network. The details of this upgrade and network are topics contained in the sequel.

2.3.1 Standard Communication System

The standard means of communication used by a *SLOCUM* glider, and most other long-range AUVs, is via IRIDIUM satellite connection. The antennae for the IRIDIUM and GPS are integrated into the rudder assembly of the vehicle, and by use of an inflatable bladder, can be lifted to an effective communication height above the waters surface.

The global coverage of IRIDIUM provides the ability for a high-endurance vehicle, such as an autonomous glider, to operate in the open ocean for long-duration missions. However, satellite communication is plagued by very low data-rates ($\approx$ 2400 baud). Slow data rates imply longer times spent at the surface for data transfer. Long surface intervals are a safety concern in areas with high marine traffic, such as coastal regions, which is our primary area of interest.

In addition to the low data transfer rate, there are high costs for transmitted data and call time. Communication expenses alone can encompass a significant portion of the operating expenses for each mission. In our experiences, the communication cost for IRIDIUM usage during a three week glider mission is estimated at USD 2400; roughly half of the total expenses for the mission. Other groups, such as [9], report their IRIDIUM communications expenses to be USD 180/day for each glider. This translates to USD 3800 for a three week mission.

In an effort to reduce operational costs, deployed vehicles communicating via satellite phone are often configured to transmit subsets of the collected data rather than an entire dataset. This reduces the call time, as well as reduces the time the vehicle spends at the surface. However, scientists can then only get a snapshot of the information collected. For some applications, especially real-time, optimal trajectory design, a snapshot is insufficient information and waiting two weeks to retrieve the entire dataset is not practical. Thus, based upon the *coastal* research interests of the CINAPS research team, we propose an alternate, robust, disruption-tolerant communication framework enabling cost-effective glider operation in coastal regions by utilizing Freewave$^{\text{TM}}$ radio modems.

3. FREEWAVE NETWORK

Most AUVs, including *SLOCUM* gliders, come equipped with a wireless communication link (*e.g.,* Wi-Fi, Freewave$^{\text{TM}}$ radio modem) for short-range operator-vehicle communications. This typically occurs during deployment and recovery of the vehicle. Once deployed, communication is primarily via satellite phone.

To avoid the previously mentioned shortcomings of satellite phones, one may choose to implement acoustic modems, as presented in [2, 3, 15], or combined acoustic/optical strategies as seen in [13]. The obvious advantage of these techniques is that the vehicle is not required to surface to transmit data. Acoustic systems have low data rates, a high one-time implementation cost. and suffer from multi-path interference in coastal waters. Optical techniques are limited to short-range solutions.

Seeking a more flexible solution, we examined a Freewave$^{\text{TM}}$ radio modem. These radios are the typical radio modems used on AUVs, and are rated for a range of 60 miles line-of-sight. Thus, we propose to extend the use of these devices beyond just deployment, recovery and dockside operations to full-scale operations in near-coastal regions (*e.g.,* the SCB). With the appropriate infrastructure, an extended-duration, multi-AUV deployment, in a coastal region, could see a significant reduction in expenses if radio modems are the primary mode of communication.

Since line-of-sight is crucial to implement an effective radio network, we choose to install the on-shore base-stations at the several HF-Radar (CODAR) installations that are positioned at elevated locations throughout Southern California. These sites are always instrumented with an internet connection, and provide a cost-effective installation locale.

With minimal modifications to the gliders, and small additions to existing shore locations, it is possible to build a network of this type, which scales with multiple vehicles operating at a minimal cost. We continue this section by describing the installed infrastructure for our Freewave$^{\text{TM}}$ network and present a strategy to pave the way for reducing communication costs, and increasing data transmission capabilities for AUV deployments in coastal regions.

3.1 System Design

The proposed system consists of three main components. First, we have the communication module that handles the radio communications on the glider. Secondly are the on-shore, internet-connected base-stations that consist of the radio, computer and antenna to communicate with the deployed sensors and vehicles. Finally, we have the control server, a central, data-aggregation and command/control server. We illustrate the overall system design in Fig. 2. This design also implements the additional high-level control of autonomous retasking to glider operations, a feature that is unavailable with current systems.

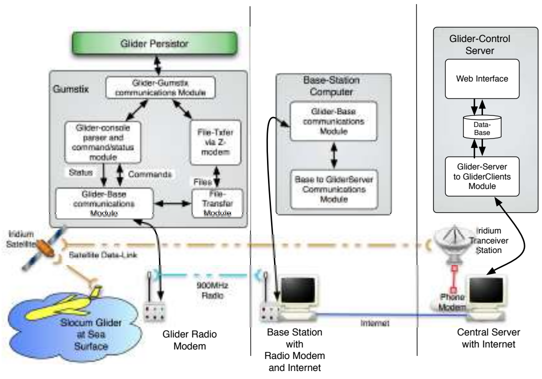

Figure 2: The System Block Diagram

3.1.1 Communication Module and Protocol

The communication module is a combination of hardware and software that handles the communication between the sensor platform and the on-shore base-stations. The hardware is specific to the sensor platform. The software incorporates general building blocks with a platform-specific interface, devoted to the interaction between the communication code and the control software on the vehicle. Based upon the test-bed platform presented in Section 2.3, the implementation and experiments presented are particular to this AUV, however the module can easily be implemented onto any other AUV or sensor platform.

The base-stations and sensor-platforms communicate via our own light-weight communication protocol. Each base-station can store and forward user-selected information between specified nodes (*i.e.,* gliders and other sensors).

We treat the Freewave$^{\text{TM}}$ modems as a serial link. There are incoming and outgoing packet queues which provide feedback to vary both the inter-packet delay, as well as the Freewave$^{\text{TM}}$ packet-sizes. Freewave$^{\text{TM}}$ modems make a "best-effort" delivery attempt on these packets, which can be fragmented into smaller pieces in the event of poor connections. This protocol (to be described in more detail in a future paper) also supports guaranteed delivery, as well as a non-guaranteed mode of transmission. The packet structure contains 14 bytes without payload data, and allows several applications to multiplex data, such as file transfers, status packets, data packets, terminal commands, etc. We use a *Selective-Acknowledgement* communication scheme with a fixed window, which works well in regions with *persistent line-of-sight* between antennae.

On our *SLOCUM* gliders, the communication module hardware consists of a *Gumstix* computer and the aforementioned software. Adding the Gumstix between the glider's control computer and the Freewave$^{\text{TM}}$ modem is a minimally intrusive way of adding new communication capabilities to the glider. By use of a separate computer to interface with the vehicle, we cleanly abstract the interface between high-level communication and low-level vehicle control.

We chose a *Gumstix* in this installation primarily for its small size and low power consumption (<120mA @5V). In addition, it is a fully-functional Linux-based computer that has excellent interfacing capabilities. Figure 3 displays the physical modifications made to the glider to incorporate the Gumstix. We need to make only five modifications to connectors on the glider to allow for communication with an external computer. Inevitably, adding the *Gumstix* increases the power consumption, which is a concern for low-power autonomous vehicles, like gliders. With this in mind, we designed the system so that the Gumstix is powered on along with the Freewave$^{\text{TM}}$ modem. This ensures that both devices get automatically turned on at the water surface, while remaining powered off when the glider is submerged.

The platform-specific software on the glider intercepts all messages between the glider persistor and the Freewave$^{\text{TM}}$ modem, as well as parses the ASCII strings from the glider. A flow chart for the basic control during each surfacing is displayed in Fig. 4. The software also reports necessary status information to the on-shore stations by use of our packet protocol. For the scientific data, the software gathers sensor data files from the persistor, compresses them and sends them to the on-shore base-station.

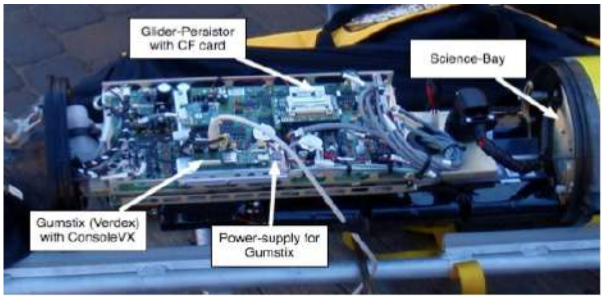

Figure 3: Hardware modifications to the USC *SLOCUM* glider.

If for any reason, communication with any base-station server via Freewave$^{\text{TM}}$ is unavailable, the glider automatically reverts to its IRIDIUM satellite phone for communication.

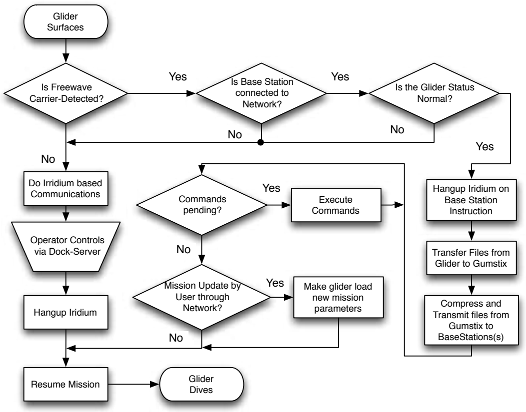

Figure 4: Control Flow during glider communications

3.1.2 *Base Stations*

The base-stations are the on-shore installations that directly communicate with the deployed sensors. The hardware consists of an internet-connected, Ubuntu Linux Server, an antenna and a Freewave$^{\text{TM}}$ radio modem. The computer at the base-station runs our communication software, which relays datagrams between the deployed sensors and the control server. Communication between these base-stations and the control server take place via TCP/IP. If a base-station loses connection with the control server, it disables its Freewave$^{\text{TM}}$ modem, to ensure that sensors connect to another base-station, or communicate via IRIDIUM.

3.1.3 *Control Server*

The control server is the orchestrator of the overall system. It is written in C++, runs on Ubuntu Linux Server and utilizes a MySQL database for data storage. The server maintains a connection with each of the base-stations, keeps track of the state of the system and issues any commands to connected sensors. Mission planners and algorithms can be run on any platform (*e.g.,* MatLab$^{\text{TM}}$, other programming languages and even other computers), and can issue commands to the control server via a database or a web service. This allows for great flexibility in the development of path planning algorithms. On top of the control server is a web-based user interface, written in Java$^{\text{TM}}$with Google Web Toolkit. This provides easy, accessible control and visualization to the end user of the system.

3.2 System Implementation and Coverage

One of the key goals in designing the communication network presented here, is to enable the coordination of multiple sensor platforms in the SCB. For coastal ocean observation, we need the ability to gather information from autonomous gliders, moored buoys, robotic or manned boats and fixed instrumentation on piers. The use of satellite communication to orchestrate all the data from these different platforms is not feasible. Additionally, utilizing multiple forms of communication protocols, unique to each platform, is also unreasonable. This provides further motivation to develop a single communication network to handle multiple platforms and data streams in one centralized location.

In addition to the flexibility of the system, we also provide safety with such a networked system. Upon surfacing, autonomous gliders are in a position of vulnerability. Since they do not protrude far from the waters surface, they are not easily seen and can be hit or run over by larger vessels. Hence, it is in our best interest to minimize their time on the surface. Through the implementation of a region-wide network, we can allow a glider to connect to the robotic or manned boats, other AUVs, or nearby buoys and moorings in the event that a base-station is out of range. Since these platforms are more visible to surface vessels, the gliders could use these boats as *data-mules* to transmit data to the on-shore base-stations. Coverage extensions of this network can be made, as off-shore buoys could act as relays in a multi-hop data transmission chain. Depending on the resources available, this *coastal* network has the capability to reach pelagic regions as well.

To date, we have only implemented on-shore base-stations. Our network consists of seven Freewave$^{\text{TM}}$ radio modem installations along the coast of Southern California from Malibu to Newport Beach, one installation near Two Harbors on Santa Catalina Island and one base-station located on the campus of the University of Southern California. The locations of these sites is depicted in Fig. 5 by the orange arcs. During multiple deployments in the SCB, we have been able to test the reliability and coverage of our network. Initial debugging experiments were conducted from a surface boat. With the network functioning as designed, we began intermittently communicating via our network, while maintaining the primary communication via IRIDIUM. Currently, all communication for one of our *SLOCUM* gliders is conducted primarily via the Freewave$^{\text{TM}}$ network. Our other glider is currently awaiting the installation of a Gumstix, so that it too will communicate primarily via our network. In Fig. 5, we display locations where communication with a glider was conducted by use of the Freewave$^{\text{TM}}$ network during our most recent deployments. We are currently gathering and analyzing data to provide communication statistics for our network. Preliminary results can be found in [6].

Figure 5: Freewave locations and locations of glider connections to the network.

4. DEPLOYMENT EXPERIENCE

Our intent is to utilize the presented Freewave$^{\text{TM}}$ network to rapidly and autonomously retask deployed mobile sensor platforms based on *in situ* data collected by one or more members of the entire network. To do this, we need the ability to transfer entire datasets from the deployed vehicle. Then, in near-real time, scientists can analyze and/or assimilate the data into predictive models. We continue here by providing examples of single and multi-vehicle deployments which will make use of this communication network.

4.1 Feature Tracking Based on Ocean Model Predictions

It is a goal of the CINAPS team to enable real-time, optimal trajectory design based on ocean model predictions for the use of tracking features of interest. The basic idea is to continually assimilate collected data into ROMS, which then predicts the evolution of a given feature. Since a complete prediction output from ROMS for the SCB takes roughly 12 hours, we plan the trajectory design and implementation in two steps. First, we identify a feature of interest in the SCB and get a ROMS prediction for a 12-16 hour period. The ROMS output is used as input to a waypoint generation algorithm that defines the trajectories that steer the AUV to regions of scientific interest based on the given feature. After the AUV executes the planned trajectory, the collected data is uploaded and assimilated into ROMS. A new prediction is generated and the process is repeated until the feature dissipates or is no longer of interest. This algorithm is presented in Algorithm 1.

At this point, the data transmission is handled via IRIDIUM, and we have successfully implemented one iteration of Algorithm 1. The ability to perform successive iterations of Algorithm 1, with retasking, depends heavily upon a rapid and reliable communication network; the Freewave$^{\text{TM}}$ network will facilitate such requests.

Considerable study has been reported on adaptive control of single and coordinated multi-glider systems, see for example

Algorithm 1 Ocean Feature Tracking Algorithm Based on Ocean Model Predictions

```
Require: A significant ocean feature is detected via direct observation or remotely sensed data such as satellite imagery.
  repeat
    A set of points ($\mathcal{D}$) is chosen which determine the current extent of the feature.
    Input $\mathcal{D}$ to ROMS.
    ROMS produces an hourly forecast for all points in $\mathcal{D}$.
    Input hourly forecast for $\mathcal{D}$ into the waypoint generation and trajectory design algorithm.
    Execute the trajectory design algorithm.
    Uploaded computed waypoints to the AUV.
    AUV executes mission.
    The AUV sends collected data to ROMS for assimilation into the model.
  until Feature dissipates, travels out of range or is no longer of interest.
```

[5] and the included references. In these papers, the trajectories given to the gliders were fixed patterns (rounded polygons) that were predetermined by a human operator. The difference between the method used in [5] and the approach described here, is that here the sampling trajectory is determined by use of the output of ROMS, and thus is, at first glance, a seemingly random and irregular sampling pattern. Such an approach is a benefit to the model and scientist alike. Scientists can identify sampling locations based upon ocean measurements they are interested in following, rather than setting a predetermined trajectory and hoping the feature enters the transect while the AUV is sampling. Model skill is increased by the continuous assimilation of the collected data; which by choice, is not a continuous measurement at the same location.

Based on the ocean feature tracking motivation presented in Section 2 and [1], we choose a freshwater plume as the proxy ocean feature. The low salinity and density imply that this feature will propagate through the ocean driven primarily by surface currents.

A waypoint generation and trajectory design algorithm based upon the ROMS output was first presented in [11]. This paper presented a single glider deployment with the mission to track the centroid of a freshwater plume. Here we briefly recall the deployment results from this trajectory design and retasking experiment, then continue with an extension to a multi-glider deployment.

4.1.1 Single Vehicle

A single Slocum glider is not optimal for rapidly surveying a large ocean feature, as it is built for endurance missions and traveling at low velocities. One primary location that is of interest in tracking ocean features is the centroid of the plume extent; analogous to the eye of the storm. Optimally, we would also like to gather samples on the boundary of the plume. However, the glider may not be able to reach the plume centroid and a point on the boundary while also keeping up with the movement of the plume.

A detailed algorithm to design a trajectory to track a plume centroid based on an output from ROMS is presented in [11]. This is the first known presentation of such a technology chain for adaptive ocean sampling. This trajectory design algorithm solves the 2-D motion planning problem by computing waypoints of interest for the vehicle to visit, which then defines the trajectory.

In summary, a set of points $\mathcal{D}$ (referred to as drifters) was input to ROMS as the delineation of initial plume location. The locations of the points in $\mathcal{D}$ were predicted for 15 hours. For safety reasons, we keep the glider from surfacing too often by setting the surfacing intervals to be ≥ 2.5 hours. Under this restriction as well as a very slow moving plume, the waypoint generation algorithm suggested visiting the plume centroid at hours $0, 5, 10$ and 15 (waypoints $1, 3, 5$ and 7, resp. in Fig. 6). Since these waypoints are 5 hours apart, three additional waypoints were computed as extra sampling locations (waypoints $2, 4$ and 6 in Fig. 6). This design strategy produced seven waypoints for the AUV to visit during the 15 hour mission. Note that we include the initial plume centroid as a waypoint, since the glider may not surface exactly at the predicted location.

Deployment Results

We deployed a Webb Slocum glider into the SCB on February 17, 2009 to conduct a two-month observation and sampling mission. For this deployment, the glider was programmed to execute a zig-zag pattern mission along the coastline by repeatidly navigating to six waypoints.

Unfortunately, weather and remote sensing devices did not cooperate to produce a rain event along with a detectable freshwater plume, so we were unable to retask the glider to track a real plume by use of our algorithm. Instead, we defined a pseudo-plume to demonstrate the proof of concept of this research.

In Fig. 6, we show the results of the retasking of a single glider to track the centroid of a dynamic ocean feature. Here, the initial delineation of the plume is denoted by the blue line and the 20 m and 30 m isobaths are denoted by the green and red lines, respectively. The waypoints computed by the trajectory design algorithm are given by the yellow diamonds, and the proposed glider trajectory is the black line. The red droplets represent the actual surfacing locations of the glider during the execution of this mission. We remark that the first waypoint was computed based on the location of a surfacing of the glider during its preset mission. We failed to get communication at this location, but were able to connect at the red droplet marked 1. From this location, the glider headed directly to waypoint two. Further details, and a discussion of the results of this deployment can be found in [11].

4.1.2 Multiple Vehicles

In addition to the centroid, tracking the boundary of an ocean feature is also of scientific interest. As previously mentioned, the gliders we use are slow, thus a single glider is practically not capable of gathering sufficient data for a large freshwater plume. To this end, we present an extension of the centroid tracking algorithm to include a boundary tracking component in the case that there are multiple gliders available.

Figure 6: Single glider centroid tracking. Image created by use of Google Earth.

The inputs to the boundary tracking portion of the algorithm are the same as for the centroid tracking algorithm. Since the boundary is not a specific point of the plume, but a continuous set of points of interest, the possible set of waypoints for the glider to visit is infinite. Additionally, we can not expect to survey the entire boundary of the plume continuously as it evolves. Thus, our algorithm delineates the boundary of the plume every four hours by connecting the predicted locations of the drifters given in the ROMS output. Let d be the distance that the glider can travel in 4 hours. For each of the drifter-delineated polygons, we compute the set of reachable points, *i.e.,* all points on the predicted boundary that are a distance d from the gliders current location. The next waypoint is a random point from this reachable set. If the set of reachable points is empty, the plume is moving faster than the glider. In this case, the next waypoint is computed by reckoning at an azimuth equal to the average azimuth of the plume's direction over the four hour period, for a distance d.

Simulation Results

We combined the centroid tracking algorithm with the boundary tracking component to simulate a two-glider, feature-tracking, trajectory design. In this case, we choose a location off the northern tip of Catalina Island as the initial location for the feature of interest. Although freshwater plumes may not exist in this general area, this region is known to contain eddys, which are equally as interesting.

For this experiment, we delineated a feature of interest with 12 drifters. We predicted the location of these drifters by use of ROMS for a period of 16 hours. Here, we assumed that two gliders would be deployed from the same location; one would track the centroid and the other would track the boundary. The computed trajectories are presented in Fig. 7. The yellow diamonds are the computed waypoints defining the boundary (B) and centroid (C) trajectories (green and magenta lines, respectively), the white stars depict the extra waypoints computed by the centroid tracker. The black line is the initial boundary of the feature and the blue (dashed) lines represent the predicted boundary of the feature near the computed waypoints for hours $4, 8, 12$ and 16. We are currently implementing a multi-glider feature tracking mission onto two gliders deployed in the SCB. It is our intent to utilize the aforementioned infrastructure of our Freewave$^{\text{TM}}$ network for uploading missions and downloading data to perform multiple iterations of Algorithm 1.

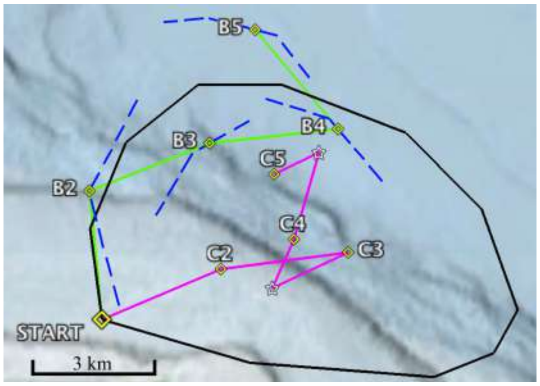

Figure 7: Trajectory design for two vehicles. The boundary tracker visits waypoints Bi, and a centroid tracker visits waypoints Ci and the additional two starred locations. Image created by use of Google Earth.

5. CONCLUSIONS AND FUTURE WORK

In the study of path planning for autonomous robots, planning the trajectory is usually less than half the battle, the real challenge comes in the implementation. This is exaggerated when dealing with underwater robots due to the complex environment and limited communications available. As one solution for coastal observation, we presented the design and implementation of a communication system using long-range RF-modems, to communicate with mobile and static sensor platforms. Field trials using *SLOCUM* gliders indicate promising results and provide us valuable insight for future improvement. Our communications protocol supports datagrams with priority, can maintain in-sequence transmissions, and has both reliable and un-acknowledged datagram modes. We utilize the Gumstix's processing power to compress data files before transmission, which provides us with a typical space saving of approximately 4x. The data rate that our system achieves in the field is 1.46KB/s (six times faster than IRIDIUM). The system can also simultaneously transmit multiplexed glider console information and status packets. This combined speed increase translates to a 24x improvement over conventional satellite phones.

We have successfully performed glider re-tasking via the Freewave$^{\text{TM}}$ network from a distance of 9.2km - a feature we will utilize in the future to enable multiple iterations of Algorithm 1. In the future, we hope to further automate the glider retasking by using our central server as the link between collected data and ROMS predictions.

Clearly, local conditions, such as the wave state of the oceanic region, affect communications and data transfer. In particular, the antenna of the *SLOCUM* glider cannot be raised more than a few tens of centimeters above the waters surface. Hence, local wave action can eliminate line-of-sight between the glider and the base-station radios. By the use of better queue management and introducing variations of re-transmit time and packet sizes based on the link state, significant improvement will be seen in the protocol-level link between a glider and base-station. Research is ongoing to map out the entire region of interest (SCB) for link quality. Equipped with such a map, we can then design planners which incorporate the knowledge of communication link availability to bias the communications of vehicles, such that they keep overall operation costs low.

Retasking autonomous gliders remotely involves patience, determination and many resources. Through the use of the Freewave^{TM} network, we hope to minimize retasking effort by reducing the monetary resources necessary for communication and providing the infrastructure to allow near-real time adaptive sampling based on collected *in situ* data.

Future work for the trajectory design involves the development of an optimization criterion to decide which vehicle is best suited for a certain task during a multi-vehicle deployment. Additionally, we plan to investigate a 3-D trajectory planner with an optimization criterion to decide the best vehicle or best path to visit a given waypoint.

6. ACKNOWLEDGMENTS

This work was supported in part by the NOAA MERHAB program under grant NA05NOS4781228 and by NSF as part of the Center for Embedded Network Sensing (CENS) under grant CCR-0120778. The ROMS ocean modeling research described in this publication was carried out by the Jet Propulsion Laboratory (JPL), California Institute of Technology, under a contract with the National Aeronautics and Space Administration (NASA).

The authors acknowledge Carl Oberg for his work with glider hardware making field implementations possible. We thank Ivona Cetinic and Matthew Ragan for their efforts in the glider deployment and communications during the missions.

7. ADDITIONAL AUTHORS

Additional authors: David A. Caron and Burton H. Jones (Department of Biological Sciences, University of Southern California, 3616 Trousdale Avenue, Los Angeles, CA 90089 USA, emails: `dcaron,bjones@usc.edu`).

8. REFERENCES

[1] I. Cetinic, B. Jones, M. Moline, and O. Schofield. Resolving urban plumes using autonomous gliders in the coastal ocean. *Journal of Geophysical Research/Americal Geophysical Union*, 2008.

[2] L. Freitag, M. Grund, S. Singh, J. Partan, P. Koski, and K. Ball. The whoi micro-modem: An acoustic communications and navigation system for multiple platforms. In *Proceedings of the IEEE Oceans Conference*, 2005.

[3] J. Heidemann, W. Ye, J. Wills, A. Syed, and Y. Li. Research challenges and applications for underwater sensor networking. In *Proceedings of the IEEE Wireless Communications and Networking Conference*, pages 228–235, Las Vegas, Nevada, USA, April 2006. IEEE.

[4] Z. Li, Y. Chao, J. C. McWilliams, and K. Ide. A three-dimensional variational data assimilation scheme for the regional ocean modeling system. *Journal of Atmospheric and Oceanic Technology*, 25:2074–2090, 2008.

[5] D. Paley, F. Zhang, and N. Leonard. Cooperative control for ocean sampling: The glider coordinated control system. *IEEE Transactions on Control Systems Technology*, 16(4):735–744, 2008.

[6] A. Pereira, H. Heidarsson, D. Caron, B. Jones, and G. Sukhatme. An implementation of a communication framework for the cost-effective operation of slocum gliders in coastal regions. In *Proceedings of The 7th International Conference on Field and Service Robotics*, Cambridge, MA, July, 2009. Submitted.

[7] Q. Vu. JPL OurOcean portal. `http://ourocean.jpl.nasa.gov/`, 2008. viewed February 2009.

[8] A. Schnetzer, P. E. Miller, R. A. Schaffner, B. Stauffer, B. Jones, S. B. Weisberg, P. M. DiGiacomo, W. M. Berelson, and D. A. Caron. Blooms of pseudo-nitzschia and domoic acid in the san pedro channela and los angeles harbor areas of the southern california bight. *Harmful Algae/Elsevier*, 6(3):372–387, 2003-2004.

[9] O. Schofield, J. Kohut, D. Aragon, E. Creed, J. Graver, C. Haldman, J. Kerfoot, H. Roarty, C. Jones, D. Webb, and S. Glenn. Slocum gliders: Robust and ready. *Journal of Field Robotics*, 24(6):473–485, 2007.

[10] A. Shchepetkin and J. McWilliams. The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Modeling*, 9:347–404, 2005. doi:10.1016/j.ocemod.2004.08.002.

[11] R. N. Smith, Y. Chao, B. H. Jones, D. A. Caron, P. P. Li, and G. S. Sukhatme. Trajectory design for autonomous underwater vehicles based on ocean model predictions for feature tracking. In *Proceedings of The 7th International Conference on Field and Service Robotics*, Cambridge, MA, July 2009. Submitted.

[12] R. N. Smith, J. Das, H. Heidarsson, A. Pereira, I. Cetinić, L. Darjany, M. Ève Garneau, M. D. Howard, C. Oberg, M. Ragan, A. Schnetzer, E. Seubert, E. C. Smith, B. A. Stauffer, G. Toro-Farmer, D. A. Caron, B. H. Jones, and G. S. Sukhatme. The Center for Integrated Networked Aquatic PlatformS (CINAPS): Observing and monitoring the southern california bight. *IEEE Robotics and Automation Magazine, Special Issue on Marine Robotics Systems*, 2009. In Preparation.

[13] L. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke. Data collection, storage, and retrieval with an underwater sensor network. In *Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 154–165. ACM, November 2005.

[14] E. Wood, A. Schnetzer, C. Benitez-Nelson, C. Anderson, W. Berelson, J. Burns, D. Caron, J. Ferry, E. Fitzpatrick, B. Jones, P. Miller, S. Morton, and R. Thunell. The toxin express: Rapid downward transport of the neurotoxin domoic acid in coastal waters. *Nature Geosciences*, 2009.

[15] X. Yu. Wireline quality underwater wireless communication using high speed acoustic modems. In *Proceedings of the IEEE Oceans Conference*, 2000.