Estimation of Wave Parameters from Accelerometry to Aid AUV-shore Communication

Arvind A. Pereira
Robotic Embedded Systems Laboratory
University of Southern California
menezesp@usc.edu
http://robotics.usc.edu/~ampereir

Gaurav S. Sukhatme
Robotic Embedded Systems Laboratory
University of Southern California
gaurav@usc.edu
http://robotics.usc.edu/~gaurav

Abstract*—Commercial off-the-shelf autonomous underwater vehicles (AUVs) are commonly equipped with short range radios (for high bandwidth, low-cost close range operation) and satellite modems (for low bandwidth, global operation). In a coastal observatory it should thus be possible to exploit an AUVs' proximity to the coast to use the short range radios when vehicles are within *useful communication range* of coastal base stations, and to use satellite communications otherwise. Measurements in the coastal ocean using a Slocum glider and a network of elevated, coastal base stations show that the useful communication range depends strongly on two parameters 1) the distance between the vehicle and the base station, and 2) the local sea state experienced by the vehicle. Here we describe measurements of communication quality in a coastal setting over 206 glider surfacing events at ranges between 2.4 and 16 km from the base station. The glider was instrumented with an accelerometer. At each surfacing, the accelerations experienced by the glider due to the local sea state were recorded concurrently with communication quality. Analysis of these measurements provides the first evidence that it is possible to predict communication quality from the sea state at a vehicle's location and the distance between the vehicle and the nearest base station. This enables a simple, yet useful, piece of logic that allows onboard decision making on whether the vehicle should use its short range radio or its satellite modem. We discuss how this finding enables online modification of low-level communication protocol parameters.

I. Introduction

Autonomous Underwater Vehicles (AUVs) are helping oceanographers by improving sampling capabilities and acting as an extension to traditional ship-based methods by vastly increasing spatiotemporal sampling resolution. Gliders are a type of low-speed, long-range AUV that use gravity and buoyancy for energy-efficient propulsion [1], [5], [6], [7], [11], [12]. These vehicles can be deployed unattended at sea for long periods of time (*e.g.,* several weeks) during which they intermittently communicate with their operators, typically via satellite phone (*e.g.,* Iridium).

Networked coastal observatories are coming online [10], [14]. When operating these vehicles in coastal regions it is possible to utilize the radio modems on these vehicles [3], [13]. This reduces operational costs while improving communication speeds (compared to slower satellite-based communication methods). However such a setup requires startup costs related to setting up the coast base stations. A coastal communication framework in the Southern California Bight (SCB)[1] is discussed in our prior work [9]. In this work we reported the development of a lightweight TCP-like communication protocol that performs communication between a glider at the water surface and a network of coastal base stations (each station is equipped with a Freewave radio modem). The primary motivation for doing this was to lower operational cost by reducing Iridium usage which constitutes a significant portion of operational costs for the vehicle. Our protocol allows the transfer of large files at high speed, providing an improvement of ∼24x over Iridium-based communications. This allows the gliders to spend less time at the surface thus improving operational safety; an important consideration in a busy coastal area such as the SCB. The communication protocol was designed from the ground up to handle multiple gliders communicating with multiple Freewave sites - a feature which is not easily possible with the commercial off-the-shelf Slocum gliders.

Figure 1 (based on an earlier version in [9]) shows how a glider at the surface can use either its Iridium satellite phone to communicate directly with a modem connected to the shore-side glider control server or it can communicate with the server via shore-based base stations using its Freewave radio modem. The figure also gives a logical overview of the typical data flow between the glider's main computer and the communication logic running on the glider's secondary computer which we have added to the commercial off-the-shelf system.

In a typical deployment where the glider's antenna is approximately 30 cm above the sea surface during a surfacing, waves near the glider can cut off visibility of the shore based antenna. However there are short periods of time when the glider may be riding the crest of a wave which is locally higher than others, thus giving it line-of-sight (LOS) to the base station. Since sea waves are periodic, we expect to periodically obtain communication between the glider and the base station when the glider is at such locally elevated positions. A long term objective of this research is to exploit the periodicity of these drops in LOS (and hence the link itself) to adapt packet size to improve the link-state protocol between the glider and the base stations.

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

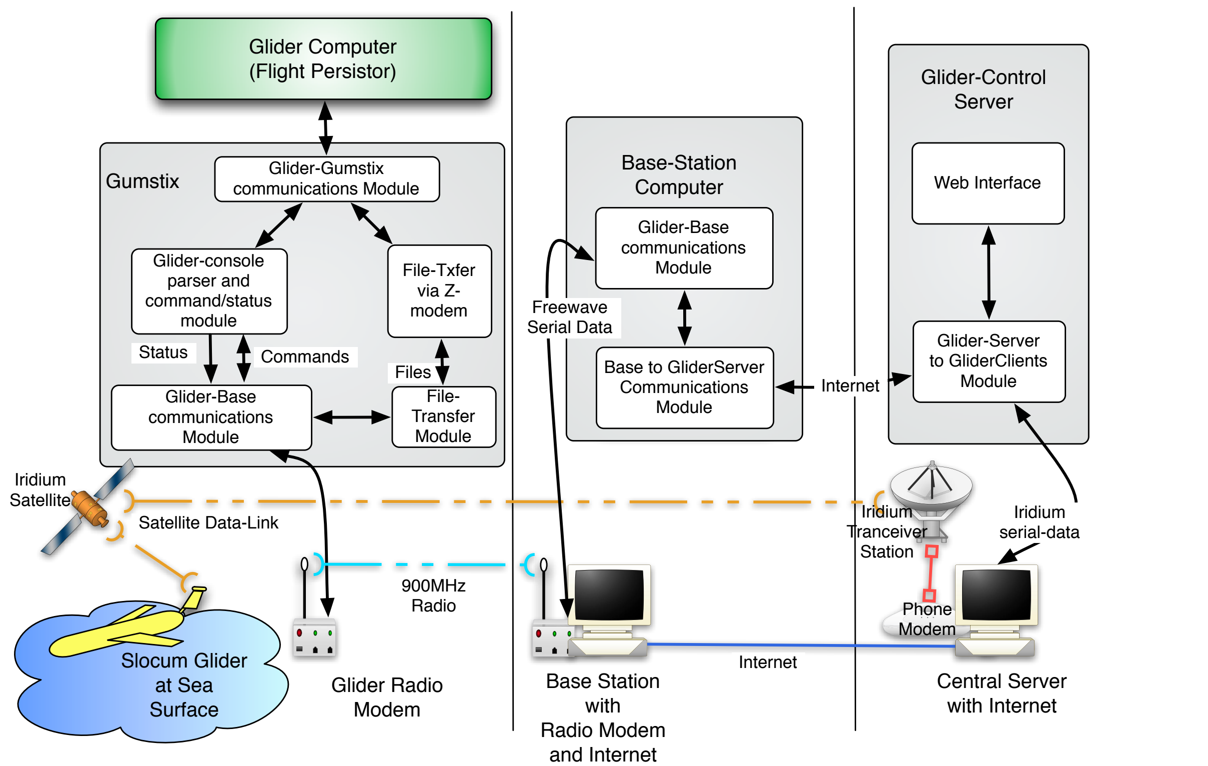

Fig. 1. A Block diagram depicting the communication system and the shore-based Glider Mission control server which allows users to plan missions and view the status of the glider through a web-based interface.

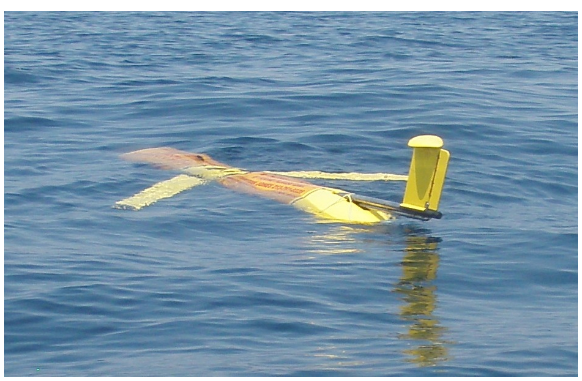

Fig. 2. A deployed Slocum glider. Note the *nose-down* posture with the tail elevated to facilitate communication. In this picture, the antenna is located approximately 30 cm above the sea surface.

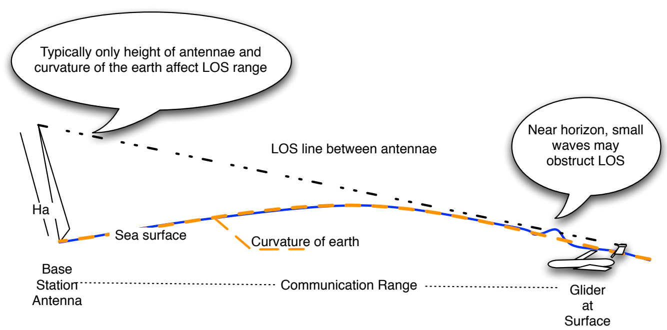

Fig. 3. Schematic depiction of glider to shore line-of-sight (LOS) obstruction due to waves. (not to scale)

The Freewave modem [15] we use in the experiments described here has in-built buffering and re-transmissions to ensure a reliable modem to modem link. If the link between the modems drops long enough, the modems lose carrier. In this case it takes longer for successful communication to take place.

While our ultimate goal is to determine appropriate communication parameters from test packets and accelerometry, here we discuss progress in the first step towards this goal. We present results from recent field tests to measure the impact of sea state on communications. These measurements establish that the useful terrestrial radio communication range depends strongly on the vehicle to base station distance as well as the local sea state experienced by the vehicle. Our analysis suggests that it is possible to predict communication quality from the sea state at the vehicle's location and the distance between the vehicle and the nearest base station. This, in turn, enables the vehicle to make a decision on whether it should use its short range radio or its satellite modem. We discuss how this finding enables online modification of low-level communication protocol parameters (*e.g.,* packet size, in analogy with GSM networks [8])

II. Hardware Modifications and Communication Protocol Design

In this section we describe, in brief, modifications performed on a commercial off-the-shelf glider to enable them to use our communication protocol as well as to collect accelerometer data to assess the effect of sea surface conditions on our communication system, in a minimally intrusive manner. Next, we describe the communication protocol used to communicate between a glider and a base-station it is communicating with.

A. Hardware Modifications

Our communication experiments required making minor modifications to the hardware on the glider. We added a low-power Linux-based computer to the glider called the *Gumstix Verdex* which sits between the glider's main flight computer (the Persistor) and the radio modem. The Gumstix appears as a Freewave modem to the Persistor. It takes a serial bit stream from the Persistor and packetizes it before sending on to the Freewave radio using the protocol described in [9]. We developed this protocol to enable the operation of multiple gliders in a region where the same base station might end up servicing several gliders at the surface as and when they make (break) connections with it.

We installed a *Phidget* 3-axis accelerometer on the glider, which is sampled at 50Hz. The accelerometer is interfaced with the Gumstix via USB, and is mounted at the rear of the vehicle (Fig. 4). The Gumstix shares its power supply with the Freewave modem, and is powered up only when the glider is at the surface and transmitting data. This helps to keep power consumption to a minimum, since the glider does not typically surface very often during normal mission operation.

When the glider detects that it is about to surface, (based on readings from its pressure transducer), it switches on the Freewave modem (and consequently the Gumstix and accelerometer). The Gumstix boots up and executes our communication code, which attempts to look for a carrier on the Freewave modem (indicative that the glider can connect to a shore-based station). Figure 5 shows the typical procedure that the Gumstix performs during a glider surfacing.

We programmed the Gumstix to perform communication-related tests over two field deployments conducted in May 2009 and September 2009. During the trials in May 2009, the goal was to test the feasibility of using the shore station-based communication network and to determine file transfer speeds. In September 2009, the goal was to concurrently record accelometry and communication quality (using test packet sequences) and develop some insight into the dependence of communication quality on location and local sea state. In time, we expect to use these data to develop algorithms for modifying communication parameters (*e.g.,* packet size) in real time based on local sea state at each surfacing to improve glider to shore communication.

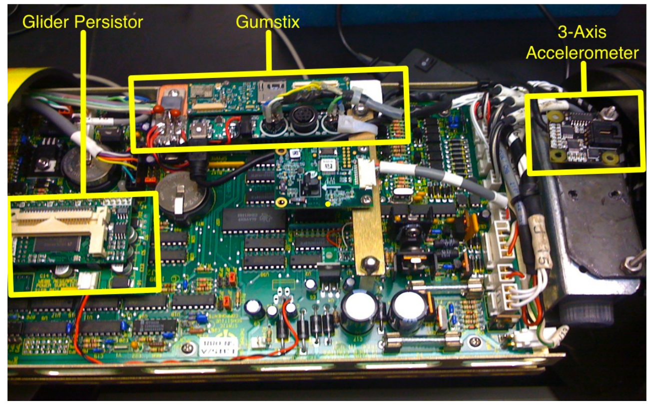

Fig. 4. The glider Persistor, additional Gumstix computer which runs our communication protocol stack, and accelerometer as installed on the USC Slocum glider.

B. The Communication Protocol Used on the Vehicle

The communication protocol we have implemented on the glider is a simplified lightweight version of TCP with selective acknowledgement. Outgoing packets are queued for transmission while incoming packets are buffered and analyzed for errors before being processed. Each valid packet contains a header which contains a packet identifier field used to distinguish between data and (various) command packets. In our implementation these could be file transfer packets, console packets, glider status packets, acknowledgement packets or test packets. In the case of some packets we want to ensure reliable transmission. These are placed in a re-transmission packet queue to be re-transmitted in the event of them not being acknowledged before the retransmission timeout (RTO) occurs. The acknowledgement is based on a selective acknowledgement scheme. Packets that need not be reliably transmitted are not copied to the re-transmission queue.

When the modems have difficulty communicating, we reduce the packet sizes to improve the chances of a successful transmission. We refer to this as varying the Maximum Transmission Unit (MTU). Acknowledgement packets provide feedback to vary the MTU and the inter-packet transmission time. Successful transmissions (indicated by acknowledgement packets) are used to increase the MTU while reducing the inter-packet transmission time. On the contrary if we timeout on the acknowledgement of packets, we reduce the MTU and increase the inter-packet transmission time. An increase in the inter-packet transmission time provides the radio modems more time to attempt re-transmission of data using their own internal protocol. The radio modem uses small packet sizes, hence we need to develop our own protocol for communication.

As long as the communicating modems have a persistent LOS this model for communication tends to work fairly well. We have been able to utilize the communications computer (Gumstix) to perform additional processing to zip data files from the glider to reduce file sizes drastically. We can also multiplex file transfers between the glider while simultaneously being able to command the Gumstix to perform other tasks such as compressing data files or retrieving new files from the Glider's computer (Persistor). We point out that since the Gumstix is powered up along with the radio modem, (which happens only at the surface), we still require to transmit data between the glider persistor and the Gumstix. This transmission is relatively fast since there is a reliable wired data link between these two computers. After a single file begins transferring between the glider and a base station, subsequent files can be transferred between the glider and the Gumstix in parallel with transmissions from the Gumstix to the base stations. We have also implemented capabilities to differentially build data files to prevent the need to initiate a new file transfer if the file was partially transmitted during a previous data transfer attempt. This feature (which was already present in the glider Persistor's Z-modem protocol) is extremely useful during file transfers.

File transfers are typically high bandwidth applications and tend to take up most of the bandwidth available on the modems. To ensure fairness as well as the ability to provide higher reliability for important packets (*e.g.,* control packets) we also provide a priority field in the packet header. The packet transmitter sorts the outgoing packets based on priority to attempt to transmit the higher priority packets before the lower-priority packets. File transfer packets are usually assigned a lower priority than command packets. This also provides fairness and improves the interaction with the glider during file transfers by improving the speed of the relay of these packets. As the radio modems move further apart, there is a natural degradation in the signal quality and consequently a reduction in the available bandwidth between the modems. Local surface waves intermittently obscure the LOS between the glider and base station radio antennae. This is primarily due to the glider antenna being too close to the surface of the sea. Finally, we note that the curvature of the earth also affects the LOS and limits or reduces the bandwidth for communication between the modems.

III. Data Collection and Processing

We tested the modified glider in the SCB during May 2009 off the Santa Catalina island in California, USA. The goal was to measure coverage and quality, *i.e.,* to build a spatial map of communication quality. We conducted a second set of tests during September 2009 where the glider was modified with the addition of the accelerometer to collect information about the sea state while concurrently measuring communication quality. This section describes the data collection and processing involved to compute features used to classify communication quality based on accelerometer data.

Unlike National Data Buoy Center (NDBC) buoys which generally rely on 20 minute time series at 2 Hz, we collect shorter accelerometer time series during glider surfacings. The surfacing event durations range between 3 and 10 minutes during which we sample the accelerometer at 50 Hz. The glider experiences roll, pitch and yaw at the surface due to its hydrodynamic interactions with the water. In Section III we describe an algorithm for processing accelerometer timeseries to obtain acceleration in the heave axis, from which wave parameters are computed using standard techniques. As reference we compare these estimates with those from the nearest NDBC buoy (approximately 25 km away).

A. Data Gathering at Sea

In the course of the May 2009 tests for communication quality and coverage (primarily file transfers at different surfacing locations) we noticed significant variations in performance at comparable glider to base station distances. This variation led to the hypothesis that sea state was a significant factor that affected communication quality [9].

A second set of tests were conducted in September 2009 during which the glider was programmed, upon surfacing, to initiate transmission of a special sequence of bytes (described next), before switching back into a regular communication with the base stations using data packets following the communication protocol described in Section II. During these tests accelerometry in all 3 axes was logged onboard the glider.

During the field trials in May 2009, we collected data to build a coverage map of a region around a radio modem base station installed at a height of approximately 70 m on Catalina island. We collected three basic types of communication metrics which are shown in Figs. 6, 7, 8 and 9. These data are interpreted in Section IV.

B. Accelerometer Data Processing Pipe-line

During the September 2009 tests, the glider collected accelerometer data while simultaneously transmitting test sequences on the Freewave radio. The test sequences were designed to allow us to understand the quality of communication between the glider and the base station at various surfacing locations and sea states. Here we describe the steps involved in processing the accelerometer data to arrive at the wave parameters in which we are interested. While this is a standard procedure and is well documented in standard wave mechanics texts for wave gauges and buoys (see [2]), we include it here to facilitate the reader.

Our accelerometer is programmed to sample each axis if it detects a minimum acceleration of 0.025 g. Since the accelerometer data is sampled asynchronously, we use a zero-order-hold while resampling the data at 50 Hz to obtain the measured time-series. We smooth the data using a low-pass filter. If we call the smoothed accelerometer time-series data as $a(n\Delta t)$ where, n is the sample index and Δt is the time-step between subsequent samples, then we segment this data into J segments of length L. Each segment is $a(j, n\Delta t)$ where a represents the values in each segment and j is the segment number. The segments are divided as shown below based on the index j.

$$
\begin{array}{ll}
1 & a(1,n\Delta t) = a(n\Delta t), n = 0..., L-1 \\
2 & a(2,n\Delta t) = a(n\Delta t), n = \frac{L}{2}, ... \frac{3L}{2} - 1 \\
\vdots & \vdots \\
J & a(J,n\Delta t) = a(n\Delta t), n = (J-1)\frac{L}{2}, ..., (J+1)\frac{L}{2} - 1,
\end{array} \tag{1}
$$

where

$$J = 2(\frac{N - \frac{N}{2}}{L}). \tag{2}$$

We remove the mean from the accelerometer data for each segment. After the removal of the means, we obtain wave spectra by correcting for roll and pitch using the glider's tilt sensor. Next we perform spectral leakage reduction using a Hanning window, which is a cosine bell taper over each data segment. The bell window is given by

$$W(n\Delta t) = \frac{1}{2}(1 - \cos(\frac{2\pi n}{L}))), 0 \leq n \leq L-1, \tag{3}$$

where L is the record or data segment length.

The data is then multiplied with the Hanning window

$$a_w(j, n\Delta t) = a(j, n\Delta t) * W(n\Delta t) \tag{4}$$

Next, we calculate the Fast Fourier Transform for the data. The Fast Fourier Transform provides the frequency domain representation, A of the measured time-series, a (or a_w with the Hanning window).

$$A(j, m\Delta f) = \Delta t \sum_{n=0}^{L-1} a_w(j, n\Delta t) e^{-i\frac{2\pi mn}{L}}, \tag{5}$$

where

$$
\begin{array}{lll}
m = & 0, 1, 2, \cdots, \frac{L}{2} & \text{L even,} \\
m = & 0, 1, 2, \cdots, \frac{L-1}{2} & \text{L odd.}
\end{array}
$$

The real and imaginary parts of A are given by

$$
\begin{aligned}
Re[A(j, m\Delta f)] &= \Delta t \sum_{n=0}^{L-1} a_w(j, n\Delta t) \cos(\frac{2\pi mn}{L}), \text{ and} \\
Im[A(j, m\Delta f)] &= -\Delta t \sum_{n=0}^{L-1} a_w(j, n\Delta t) \sin(\frac{2\pi mn}{L}).
\end{aligned}
$$

We obtain spectral estimates at Fourier frequencies, $m\Delta f$, where the interval between the frequencies is given by

$$\Delta f = \frac{1}{L\Delta t}. \tag{6}$$

Power Spectral Density (PSD) estimates for the j^{th} segment are obtained from

$$S_{aa}(j, m\Delta f) = \frac{A^*(j, m\Delta f) A(j, m\Delta f)}{L\Delta t} \tag{7}$$

where A^* is the complex conjugate of A. In particular we are interested in the Co-Spectral Density given by

$$C_{aa}(j, m\Delta f) = \frac{Re[A]Re[A] + Im[A]Im[A]}{L\Delta t} \tag{8}$$

The Final co-spectral estimate is obtained by

$$C_{aa}(m\Delta f) = \frac{1}{J} \sum_{j=1}^{J} C_{aa}(j, m\Delta f) \tag{9}$$

Now, we compute some of the standard wave parameters which are usually measured by wave-gauge buoys, which are the *significant wave-height, peak time period and average time period*. The first of these parameters is the significant wave height (H_{mo} or H_s), which is calculated from the wave elevation variance (also the zero moment, m_0 of a non-directional wave spectrum) using

$$H_s = 4.0\sqrt{m_0} \tag{10}$$

where

$$m_0 = \sum_{N}^{n=1} C_{11}(f_n) df_n. \tag{11}$$

The peak time period is given by:

$$T_p = \frac{1}{f_p} \tag{12}$$

where f_p (the peak frequency) is the frequency at which the spectral wave energy density is a maximum. The average period T_{av} is calculated by the following

$$T_{avg} = \frac{m_0}{m_1}. \tag{13}$$

The spectral moments are given by:

$$m_i = \sum_{n_b}^{n=1} f_n^i C_{11}(f_n) df_n, \qquad \text{i=0,1.} \tag{14}$$

C. Communication parameters processing

During the September 2009 trials, the glider was programmed to begin transmitting a sequence of bytes at the maximum baud-rate possible. It makes five attempts to transmit a sequence of bytes from 0-65535, and 2000 padding bytes around the sequence. We chose this sequence because we were able to detect dropped bytes and by cycling through these bytes, we are able to detect successful and unsuccessful parts of the transmission. While transmitting this data, the glider is simultaneously collecting accelerometry which is used to compute the wave parameters described in Section III-B.

The base station is programmed to continuously monitor its Freewave modem for carrier or incoming data. Any new activity at the base station is assumed to be from the glider, since it has the only Freewave modem programmed to be able to communicate with the base station during the course of the experiment. The base station stores any change to its carrier detect state, as well as logs incoming byte streams with timestamps. We sample bytes collected from the incoming byte stream every 10 ms. The data rate between the radio modem and the base station computer is 115200, which translates into a maximum of 115 bytes being received between any two samples.

If we designate a glider surfacing as the time between obtaining a carrier detect on the modem, and losing carrier detect for the final time (during a particular glider surfacing episode), we can compute the total time carrier detect was high during this period of time. We use this ratio (expressed as a percentage of surfacing time) as one of the indicators for communication quality. Carrier detect is a necessary but not sufficient indicator of successful delivery of the data payload.

The most useful communication parameter we can measure is the link state ratio. This is the ratio of the time we had a persistent communication link between the glider and the base station to the total surfacing time. Since the only way to ensure there is a link between the two is through the exchange of data packets, the granularity of this measure is governed by the maximum time between transmission of any two packets from the glider. If the glider has nothing to transmit during a surfacing (*e.g.,* files or console packets), it tries to transmit a status packet containing health information (*e.g.,* position, battery voltage, mission name *etc.*) to the base station every 5 seconds. The link is assumed to be lost if no packet is received from the glider at the base station for more than 10 seconds. For many of the glider surfacings at distances greater than 12 km, despite good carrier detect ratios, the link state percentage is low.

Another indicator of the communication quality is a count of the number of times the modems had a drop on their carrier. Similarly we can also use the number of times the link was lost as an indicator of how poor the communication is. These can be used to indicate poor communication conditions.

From the experiments we performed in September, we perform a correlation of the byte sequences received at the base station, to determine the total correlation, number of bytes dropped or skipped during transmission, as well as the average correlated byte length. This gives us substantial information on how to vary the communication parameters at each transmission time. The average correlated byte length gives us an indicator for the MTU, while we can compute the length of time occupied by the average skipped count to adjust the retransmission timeout (RTO) to a better value for the given conditions.

The wave parameters computed from the acceleration spectra are significant wave height (H_s), peak time period (T_p) and average time period (T_{avg}). We find that the wave parameters obtained using the accelerometer on the glider show a similar trend to those gathered from the nearest data buoy which is NDBC buoy 96222 which is located approximately 25 km from the region the glider was operated within, (see Fig. 11). We trained a linear perceptron classifier (described in [4]) on a subset of the data collected at surfacing events to assess how well it predicts the communication quality, specifically link percentage (the ratio of the time a persistent communication link was present to the total time at the surface). We notice from our experiments that values greater than 90% are good for file transfers.

IV. Experimental Results and Analysis

We present the coverage maps from the May 2009 tests for the Carrier-detect percentage, Link-state percentage and

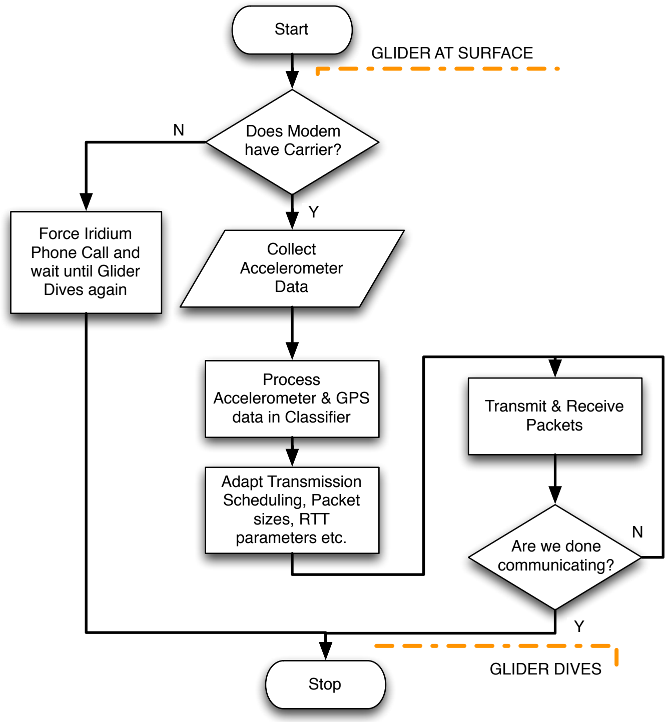

Fig. 5. Flowchart for a typical glider surfacing with the Gumstix and accelerometer additions.

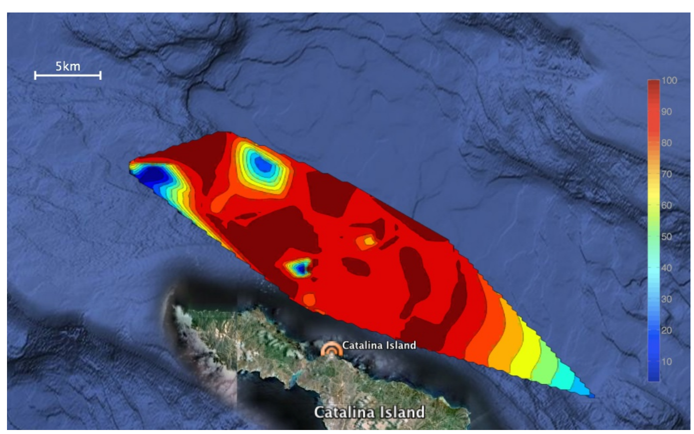

Fig. 6. Carrier detect expressed as a percentage of time with modem-modem carrier detect on from the May 2009 measurements. Some glider surfacing events had poor carrier detect in spite of being closer to the base station than other locations. This motivated us to use an accelerometer on the glider to test if the sea state was responsible for the poor communication performance.

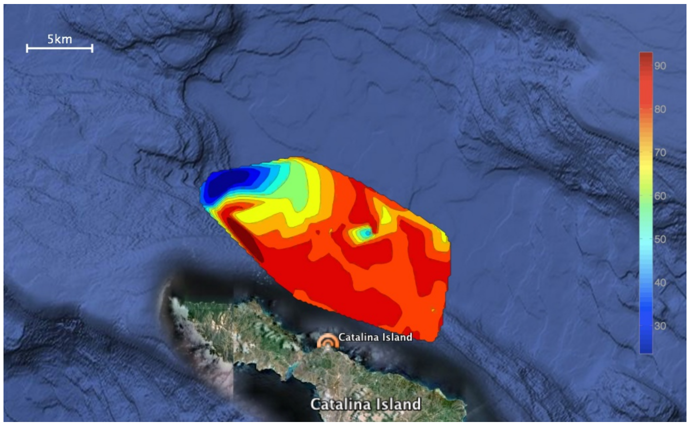

Fig. 7. Protocol link state expressed as a percentage of time with protocol link on during the May 2009 trials. The link state is a ratio of the total time the glider had communication to the total time spent by the glider at the surface. Links closer to the base station are generally better although there are exceptions. Having carrier detect between modems does not imply a good link state.

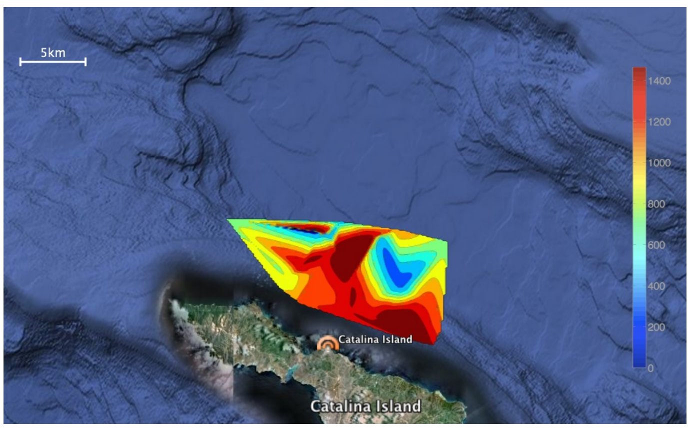

Fig. 8. File transfer speeds in bytes/s from the May 2009 trials. In these trials, we used a standard set of files as a payload to measure file transfer performance after the glider surfaces. No data compression was performed before transmission. For these trials, we used conservative estimates on retransmit timeouts.

file transfer percentages in Figs. 6, 7, 8 respectively. We also present a scatter plot for the link state percentage vs. distance for all the observations made during the May 2009 tests. Figure 10 shows a link state coverage map from the September 2009 tests. Just as we observed for the tests in May, we observe that there are regions of lower performance closer to the base station compared to some regions further away from the base station.

Figure 11 compares the significant wave height (H_s) obtained by processing our accelerometer data, with that obtained from NDBC buoy 96222 located 25 km from most of the glider surfacings. This buoy is the nearest reference providing wave data to the region where our glider experiments were performed. At the time of writing this paper, preliminary analysis compares significant wave height at glider surfacing location directly with that at the buoy during the same time

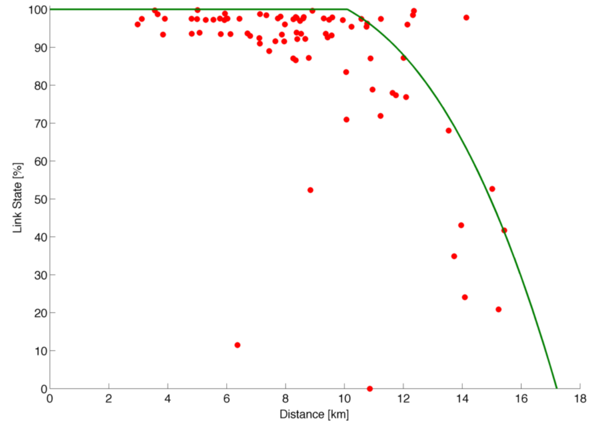

Fig. 9. Protocol link-state expressed as a percentage vs. distance (km) between the glider and the nearest base station. Our communication protocol gives a link-state of >90% up to 12 km with a base station located at a height of 70 m above the sea surface. We notice that there are instances of poor communication at ranges under 10 km. We hypothesize that these are due to surface conditions that affect communication between the glider and the base station.

interval. An analysis is ongoing to provide better comparisons by incorporating the direction and velocity of the waves. The general trend in the significant wave-height measured at the buoy and the glider is similar, although due to differences in the direction of the sea-waves, significant distance between the observations at the buoy and the glider, as well as the smaller time-window of accelerometer data-collection on the glider, variations may be expected.

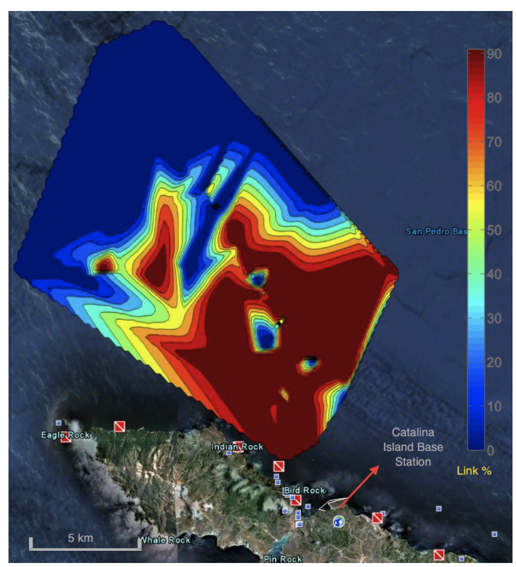

Fig. 10. Link state obtained from the September 2009 glider deployment. While the link state is generally very good within 8 km, we find several surfacing locations where the link state was poor nearer the base station on Catalina island.

Figure 12 depicts the scatter plot of the correlated byte length and the distance between glider surfacing locations. We observe that there is a significant drop in the average correlated

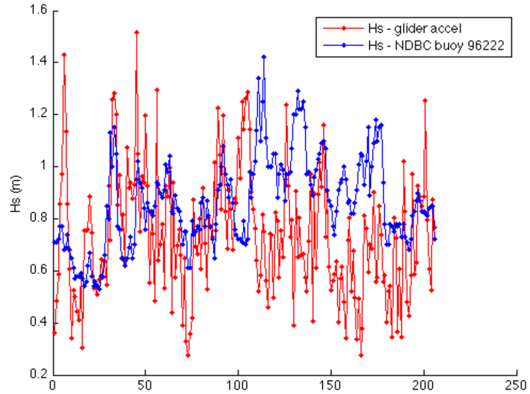

Fig. 11. Comparison of H_s from the glider accelerometer and NDBC Buoy 96222 located approximately 25 km away. The general trend in the significant wave-height measured at the buoy and the glider is similar, although due to differences in the direction of the sea-waves, significant distance between the observations at the buoy and the glider, as well as the small time-window of accelerometer data-collection on the glider, variations may be expected.

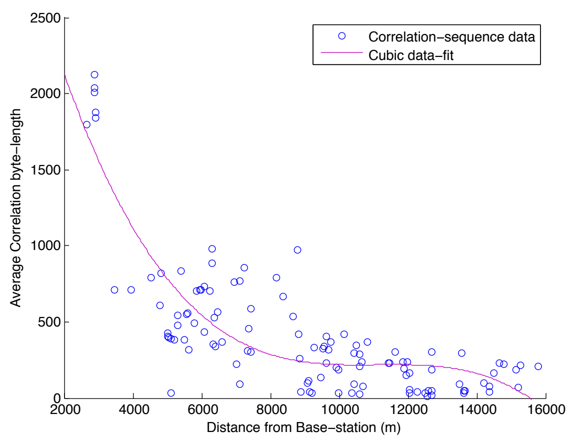

Fig. 12. Scatter plot showing the average correlated byte length and the distance (m) between glider surfacing location and the base station. A cubic data fit is shown to highlight the trend followed by the data as distance increases. From the plot it is obvious that there is a sharp drop in correlated byte lengths beyond 3 km. There is considerable spread in the correlation scores between 5 km and 11 km, which we will investigate further. Beyond 12 km, the average correlated byte length is under 300 bytes.

byte length from 2200 bytes at a range of 2.8 km to under 1000 bytes beyond 4 km. The trend from this point onwards is a slow decline beyond 4 km. We notice that there is significant spread in the average correlated byte sequence length between 4 km and 10 km. We speculate that this is related to the conditions at the sea surface affecting the performance of the radio modem on the glider and its antenna. Beyond 8 km, there is a continued drop in the average byte length which falls to approximately 250 bytes beyond 12 km. We see several cases where the average correlated byte length is almost zero beyond 8 km.

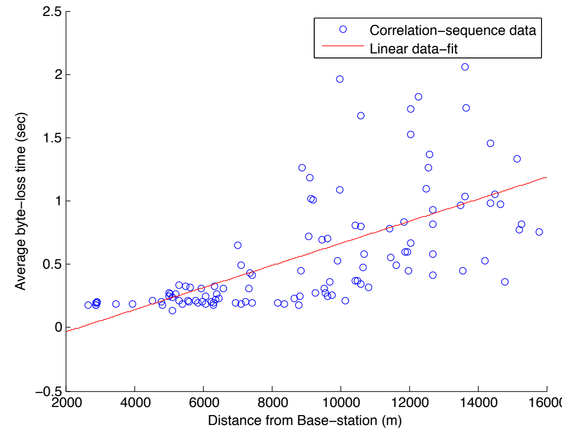

Fig. 13. Scatter plot showing the average time occupied by dropped bytes in the received correlation sequence. This time provides us with a figure that includes delays due to transmission errors, modem-modem retransmissions, intermittent losses in carrier detect, line-of-sight between modems and so on. We see an increasing trend (as expected), with a large spread as glider-surfacing locations beyond 8km become progressively more distant from the base station.

Figure 13 is a scatter plot of the average time occupied by dropped bytes during the transmission of the correlation sequence. This serves as a measure of how much time is wasted during a typical transmission due to a combination of various factors which may include loss in LOS between the modems, a drop of carrier between modems, errors in reception, delays due to retransmission between modems, bytes dropped due to buffer overruns in the modem and so on. We see an increasing trend (as expected), with a large spread as glider-surfacing locations beyond 8 km become progressively more distant from the base station. Our approach here is to treat the modems as a black box, and this is why we do not explicitly model the effects of each factor that might result in transmission delays.

By looking at Fig. 12 and Fig. 13, we see that the overall throughput of transmission must drop very quickly with distance. This is because there is a significant increase in the amount of time between successful transmissions of data bytes, while there is a sharp drop in the number of successfully transmitted bytes at increasing distances between the modems. This translates into a sharp drop in data throughput, because the modems begin taking longer to transmit smaller amounts of data. If the glider attempts to transmit large files at the surface, it might risk staying at the surface for an unduly long period of time, which is undesirable.

Figure 14 is a scatter plot of the maximum throughput achievable by the glider by measuring the total number of bytes from the transmitted sequence successfully obtained at the base station, without any retransmissions from the glider side. This transmission speed is therefore governed primarily by the maximum transmission rate which the modems can handle, and it goes down with increasing distance. The the-

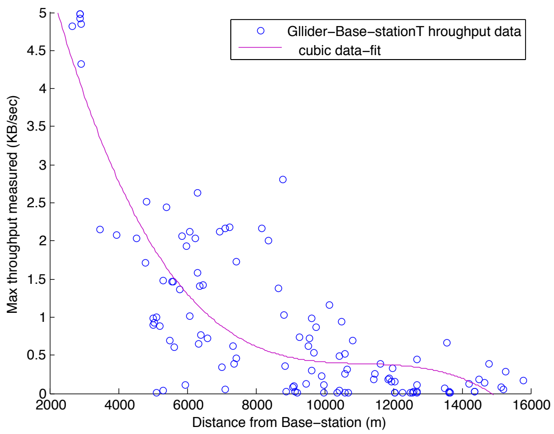

Fig. 14. Scatter plot showing the maximum throughput between the glider and the base station from the received correlation sequence. Each data point is based on the achieved throughput at the base station based on the ratio of the total data successfully transmitted in the given time. This is the maximum throughput we can expect to achieve, since the modems cannot be made to transmit any faster. In practice, we have lower throughput because we need to retransmit data packets that have errors in them, due to which the rate goes down.

oretical maximum transmission rate between two Freewave modems such as those on the glider and base station is 11.5 KB/s. In practice, with the glider we have been testing and our base station the maximum throughput we have achieved is approximately 5 KB/s for distances under 3 km. The throughput goes down to approximately 2 KB/s in good sea conditions between 4 km and 9 km. Beyond 9 km, there is a gradual drop in throughput.

TABLE I
COMPARISON OF CLASSIFIER PERFORMANCE ON LINK-QUALITY

Training data-set	Range >6 km Link <30%	Range >6 km Link >30%	Range >6 km Link >90%
A	91.14%	90.91%	91.19%
B	90.97%	91.33%	89.40%
C	94.39%	93.62%	93.79%

TABLE II
COMPARISON OF CLASSIFIER PERFORMANCE ON LINK-QUALITY

Training data-set	Range <8 km Link <30%	Range <8 km Link >30%	Range <8 km Link >90%
A	97.25%	97.83%	96.35%
B	98.62%	97.14%	98.36%
C	98.20%	97.35%	97.72%

Tables I, II, III show results of averaging the comparison of 30 trials using a classifier with data from 206 glider surfacings for A - classification with range only, B - classification with T_p, T_{avg}, H_s only, and C - classification with T_p, T_{avg}, H_s and range (distance between the base station and the glider). We used 80% of the permuted feature vectors for training the classifier and used the remaining feature vectors to assess its accuracy from the data we already have. We find that given processed wave parameters and a glider's surfacing location, we can determine if we should use Iridium (Link $< 30\%$), send out only a status report, or do transmission of files (Link $> 90\%$). Range plays a major role as is also visible in Fig. 10.

TABLE III
COMPARISON OF CLASSIFIER PERFORMANCE ON LINK-QUALITY

Training data-set	Range >8 km Link <30%	Range >8 km Link >30%	Range >8 km Link >90%
A	90.62%	91.99%	90.43%
B	91.99%	91.34%	89.67%
C	94.56%	96.05%	91.67%

V. DISCUSSION

Throughput is a good indicator of communication quality. It also serves as a good benchmark to decide where the cut-off between the use of Iridium or Freewave should be, since there is a direct comparison which is possible. The maximum transmission speed on Iridium is 240 byte/s. As we can see from the data, there are several surfacings when the throughput is far lower than this - even at distances of approximately 5 km between the glider and the base station. Hence, although it appears that in general, we can perform file transfers via Freewave all the way up to 14 km, there are times when terrestrial communication is poor enough that Iridium even at a low data rate performs better.

Some other key observations from the trials conducted are that the classifier has a very high accuracy rate in predicting how poor the communication quality is (98.62%), without the use of range in the feature vector for training. Similarly, there is a very high accuracy in predicting a good link quality between the two modems (98.36%) using accelerometer data only. Accelerometer data appears to capture the effect of some of the external conditions that might affect the link quality experienced by the glider during communication. In fact, the use of accelerometer data has a lower predictive error rate than the use of a classifier that uses range alone (as we can see from Table II).

The use of the communication sequence yielded insights into the trend followed by communication throughput and even provided us with a means of choosing appropriate values of MTU and RTO for use between the glider and the particular base station at Catalina island. Since the range achievable by an antenna is dependent on its elevation (among many other factors), coming up with a generalizable version of this will need further tests with antennas at multiple elevations.

For our antenna which is situated at an elevation of 70 m, the MTU can take the maximum value of 2000 bytes if the glider is closer to the base station than 3 km. Beyond this distance, a safe MTU is below 200 bytes (see Fig. 12). Similarly, we find that the RTO can be very high at distances beyond 6 km. Up until 6 km, a reasonable value for the RTO is 1 second. Beyond 6 km, there are occasions when the RTO can be higher than 4 seconds. Such a high RTO is not practical for communication that relies on acknowledgements, while having low bandwidth. Hence, we would likely cap the RTO to a maximum of 2 seconds for acknowledged data transfers.

VI. Conclusion

We find that wave parameters are useful in estimating expected quality of communication. The sea state was generally calm during the time our data was collected, with significant wave heights between 0.4 m and 1.7 m, and peak time periods ranging between 2.86 sec and 22.22 sec. By knowing how well we can expect to communicate over the Freewave radio based on the sea state, the time at the surface can be better utilized. If the sea state indicates that the expected communication quality will be poor, file transmissions can be postponed to a future surfacing. Accelerometer time series from the glider at the surface can be utilized to predict the expected quality of communication with shore stations. While range plays a major role in predicting the expected communication quality, Table II shows that for ranges under 8 km, the classifier that relies on acceleration information is more accurate in predicting when the communication quality is expected to be poor. This provides evidence supporting the hypothesis that sea state plays a role in determining the quality of communication and can be used to enhance prediction of communication quality.

VII. Future Work

We have been able to find evidence supporting the hypothesis that sea state plays a role in determining the quality of communication. We have also been able to obtain some parameters describing sea state from accelerometer data gathered on the vehicle. The next step is to test if we obtain gains from our approach in the course of a real deployment. The goal in these future experiments will be to analyze the accelerometer data in real time on the glider, to determine if the glider should transmit data files or stick to transmitting only status information packets or to give up on attempting a Freewave call and settling to use the Iridium satellite phone instead.

The ability to predict expected communication regimes on the glider will in the future allow us to do better path planning for the glider where we could include a waypoint which might provide the glider with the ability to transmit data files more easily. The glider may also be able to analyze science data and look for values that indicate that there is merit in transmitting data over postponing transmission. While our network is designed to work with multiple gliders, we have tested it with a single glider at a time so far. A natural extension is to ascertain that it is capable of handling multiple gliders simultaneously.

VIII. ACKNOWLEDGMENTS

The authors gratefully acknowledge the help, valuable advice and support provided by Ryan Smith, Hordur Heidarsson, Burton Jones, David Caron, Carl Oberg, Gordon Boivin, Ivona Cetinic, Matthew Ragan, Jonathan Binney, Gerry Smith, Trevor and other members of the CINAPS team, in the course of the experimental trials. Hordur Heidarsson was also involved in developing the glider network network. 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, by NSF grants CNS-0520305 and CNS-0540420, by the ONR MURI program (grants N00014-09-1-1031 and N00014-08-1-0693) by the ONR SoA program and a gift from the Okawa Foundation.

References

[1] Ralf Bachmayer, Naomi Ehrich Leonard, Joshua G. Graver, E. Fiorelli, Pradeep Bhatta, and D. Paley. Underwater gliders: recent developments and future applications. In *International Symposium on Underwater Technology*, number 0-7803-8541-1, pages 195–200, 2004.

[2] Ronald Brown, James Baker, and Jerry McCall. Nondirectional and directional wave data analysis procedures. Technical Report 01, National Data Buoy Center, NOAA, Stennis Space Center, January 1996.

[3] Jnaneshwar Das, Kanna Rajan, Sergey Frolov, John Ryan, Frederic Py, David A. Caron, and Gaurav S. Sukhatme. Towards marine bloom trajectory prediction for AUV mission planning. In *IEEE International Conference on Robotics and Automation*, 2010.

[4] Richard O. Duda, Peter E. Hart, and David G. Stork. *Pattern Classification*. A Wiley-Interscience Publication, 2nd edition edition, 2001.

[5] Charles C. Eriksen, T. James Osse, Russell D. Light, Timothy Wen, Thomas W. Lehman, Peter L. Sabin, John W. Ballard, and Andrew M. Chiodi. Seaglider: A long-range autonomous underwater vehicle for oceanographic research. *IEEE Journal of Oceanic Engineering*, 26(4):424–436, October 2001.

[6] Gwynn Griffiths, Clayton Jones, James Ferguson, and Neil Bose. Undersea gliders. *Feeding and Healing Humans*, 2(2):64–75, 2007.

[7] Clayton Jones, Elizabeth L. Creed, Scott Glenn, John Kerfoot, Josh Kohut, Chhaya Mudgal, and Oscar Schofield. Slocum gliders - a component of operational oceanography. In *Autonomous Undersea Systems Institute Symposium Proceedings*, 2005.

[8] Reiner Ludwig, Almudena Konrad, Anthony D. Joseph, and Randy H. Katz. Optimizing the end-to-end performance of reliable flows over wireless links. *Wireless Networks*, 8(2-3):289–299, November 2002.

[9] Arvind A. Pereira, Hordur Heidarsson, Carl Oberg, David A. Caron, Burton H. Jones, and Gaurav S. Sukhatme. A communication framework for cost-effective operation of AUVs in coastal regions. In *The 7th International Conference on Field and Service Robots*, Cambridge, Massachusetts, Jul 2009.

[10] Oscar Schofield, Scott Glenn, Robert Chant, Mark A. Moline, Paul Bissett, Dale Haidvogel, and John Wilkins. The evolution of a nearshore coastal observatory and the establishment of the new jersey shelf observing system. *Sea Technology*, 44(11):52–58, 2003.

[11] Oscar Schofield, Josh Kohut, David Aragon, Elizabeth L. Creed, Josh Graver, Chip Haldeman, John Kerfoot, Hugh Roarty, Clayton Jones, Douglas C. Webb, and Scott Glenn. Slocum gliders: Robust and ready. *Journal of Field Robotics*, 24(6):473–485, 2007.

[12] Jeff Sherman, Russ E. Davis, W. B. Owens, and J. Valdes. The autonomous underwater glider "spray". *IEEE Journal of Oceanic Engineering*, 26(4):437–446, October 2001.

[13] Ryan N. Smith, Yi Chao, Peggy P. Li, David A. Caron, Burton H. Jones, and Gaurav S. Sukhatme. Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model. *International Journal of Robotics Research*, 2010. Submitted To.

[14] Ryan N. Smith, Jnaneshwar Das, Hordur Heidarsson, Arvind A. Pereira, Filippo Arrichiello, Ivona Cetinic, Lindsay Darjany, Marie-Eve Garneau, Meredith D. Howard, Carl Oberg, Matthew Ragan, Erica Seubert, Ellen C. Smith, Beth Stauffer, Astrid Schnetzer, Gerardo Toro-Farmer, David A. Caron, Burton H. Jones, and Gaurav S. Sukhatme. USC CINAPS builds bridges: Observing and monitoring the Southern California Bight. *IEEE Robotics and Automation Magazine*, pages 20–30, Mar 2010. Special Issue on Marine Robotic Systems.

[15] FreeWave Technologies. Spread spectrum wireless data traceiver user manual. Technical Report 6.3, FreeWave Technologies Inc., 2005.