

A Drift-Tolerant Model for Data Management in Ocean Sensor Networks

Silvia Nittel^{*}
SIE and NCGIA,
University of Maine, USA
nittel@spatial.maine.edu

Francois Neville
School of Marine Science,
University of Maine, USA
fneville@umit.maine.edu

Niki Trigoni
Computing Laboratory,
University of Oxford, UK
niki.trigoni@comlab.ox.ac.uk

Arda Nural
SIE and NCGIA
University of Maine, USA
arda@spatial.maine.edu

Konstantinos Ferentinos
Agricultural Univ. of Athens,
Athens, Greece
kpf3@cornell.edu

Neal Pettigrew
School of Marine Science,
University of Maine, USA
nealp@maine.edu

ABSTRACT

Traditional means of observing the ocean, like fixed mooring stations and radar systems, are difficult and expensive to deploy and provide coarse-grained and data measurements of currents and waves. In this paper, we explore the use of inexpensive wireless drifters as an alternative flexible infrastructure for fine-grained ocean monitoring. Surface drifters are designed specifically to move passively with the flow of water on the ocean surface and they are able to acquire sensor readings and GPS-generated positions at regular intervals. We view the fleet of drifters as a wireless ad-hoc sensor network with two types of nodes: i) a few powerful drifters with satellite connectivity, acting as mobile base-stations, and ii) a large number of low-power drifters with short-range acoustic or radio connectivity. Using real datasets from the Gulf of Maine (US) and the Liverpool Bay (UK), we study connectivity and uniformity properties of the ad-hoc mobile sensor network. We investigate the effect of deployment strategy, weather conditions as well as seasonal changes on the ability of drifters to relay readings to the end-users, and to provide sufficient sensing coverage of the monitored area. Our empirical study provides useful insights on how to design distributed routing and in-network processing algorithms tailored for ocean-monitoring sensor networks.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous

General Terms

Design

^{*}This research was funded under NSF CAREER award No. 0448183.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiDE'07, June 10, 2007, Beijing, China.

Copyright 2007 ACM 978-1-59593-765-0/07/0006 ...\$5.00.

Keywords

Ad-hoc sensor networks, mobility, drifters, oceanography, geosensor networks

1. INTRODUCTION

Establishing a fine-grained model of local ocean currents is important since currents carry nutrients and other substances, which affect ecosystems in coastal regions. For example, researchers are interested in establishing current models for the Gulf of Maine (US) since they distribute a specific type of algae to shellfish off the coast of Maine during the warm summer months; the shellfish consuming the algae turn toxic for humans ('red tide' phenomenon)[1]. Today, major ocean currents are established using coastal radar; however, the information is spatially and temporally too coarse. We investigate the alternative deployment of a fleet of inexpensive ocean drifters which are passively propelled by the current and report their GPS-based location and trajectories to the end user.

Today, large-scale sensing platforms such as stationary moorings, research vessels, coastal radar (CODAR) or large gliders are state of the art in ocean monitoring. In the future, they will be combined with new technology developments such as small-scale, inexpensive drifters and/or autonomous underwater vehicles (AUVs) such as gliders. A drifter is a small computing platform the size of a basketball which floats on the ocean surface, usually equipped with a long underwater peddle which moves the sensor as the ocean current moves (instead of the wind). Sensor boards can be attached to detect phenomena such as oil spills or marine microorganisms [2]. In contrast to drifters, AUVs are self-propelled and determine their movement direction and travel speed in an autonomous way. Both types of platforms contain a battery supply, a GPS device, and wireless communication. Wireless communication media for shallow areas use acoustic signals with a communication distance of 5-10 Miles (shallow water areas are regions of the ocean that do not exceed depths of 100m).

Today, drifters are often deployed in a singular fashion, and use satellite communication to upload data to a centralized computer. In this paper, we explore the use of a fleet of inexpensive wireless drifters as an alternative flexible infrastructure for fine-grained ocean monitoring. We view the fleet of drifters as a wireless ad-hoc sensor network with

two types of nodes: i) a few powerful drifters with satellite connectivity, acting as mobile base-stations, and ii) a large number of low-power drifters with short-range acoustic or radio connectivity. Our objective is twofold: using a fleet of small-scale sensor nodes that communicate with each other using lower-energy acoustic signals instead of a satellite uplink saves large amounts of energy. Additionally, the fleet provides more detailed information by covering an ocean region in high density. The passive movement of drifters can be used to derive actual ocean current data on a detailed scale. Deploying a fleet of mobile ad-hoc sensor nodes on the ocean surface to track and monitor ocean currents in a fine-grained, near real-time scale is a novel research problem, both from the perspective of computer science and oceanography.

In this paper, we explore communication connectivity and sensing uniformity of a fleet of a mobile ad-hoc sensor network using real datasets from the Gulf of Maine (US) and the Liverpool Bay (UK). The challenge is to design, build and deploy drifter platforms that despite involuntary, passive movement over long time periods (up to 3 months) preserve energy power, long-term network connectivity, and sensing uniformity. Using simulation and real datasets, we investigate the effect of deployment strategy, weather conditions, and seasonal current changes on the ability of drifters to relay readings to the end-users, and to provide sufficient sensing coverage of the monitored area. Our empirical study provides useful insights on how to design distributed routing and in-network processing algorithms tailored for ocean-monitoring sensor networks.

The remainder of the paper is organized as following: Section 2 provides relevant technical background on the current state of the art of ocean sensor networks, drifter platforms and wireless communication technology for water environments. Section 3 explores the research questions and the approach of this paper in more detail. Section 4 contains our experimental results and we conclude with Section 5.

2. BACKGROUND

In this section, we review the current state of the art in ocean observation research. The research can be roughly divided into deep sea exploration using submarines and ocean bottom sensor platforms and robots connected by optical fiber cable (e.g. NEPTUNE [3]). Another large research area is in near-coastal observations using fixed large, sensor-equipped buoys like the moorings in the Gulf of Maine or Liverpool Bay, and extend the environments with coastal radar, gliders, and research vessels. Our interest in ocean surface drifters is with regard to near-coastal deployments in order to investigate these currents in greater detail.

2.1 Ocean drifters

Today, several projects and platforms for shallow water drifters exist. The first deployments of drifters were in the Gulf of Mexico and the Southwestern Caribbean Sea designed to explore the Gulf Stream in more detail. In 1998, so-called YOTO Drifters were deployed to collect information about this North Atlantic current flow in more detail.

Today, the international ARGO project [2] is one of the largest deployments of drifters in the world oceans. ARGO is an international program that began in 2000, and by 2007 the deployment of 3000 profiling drifters will be about 100% complete. The purpose of ARGO is to examine the global

currents, circulation and air-sea interaction, with the goal of improving climate models and predictions. Partners in the National Ocean Partnership Program (NOPP)/ARGO program include the University of Washington, the Scripps Institution of Oceanography in San Diego, the Woods Hole Oceanographic Institution, and others.

The Argo Drifter (also called "Davis Drifter") was designed to be a surface level (1 meter below surface) Lagrangian drifter which can report position via the Argos satellite-based data collection system. Location determination by GPS is also available. The unit consists of a central sealed tube which contains the electronics and power pack with a nominal operating life of 9 months. Argo augments existing upper-ocean observing networks, and extends their coverage in space and time, their depth range and accuracy, and enhances them through the addition of velocity measurements. The global array of 3,000 floats is distributed roughly every 3 degrees (300km).

Since 1993, the Minerals Management Service (MMS) has deployed over 800 satellite-tracked Davis drifters to measure the surface ocean currents in areas of active or prospective oil and gas leasing, primarily in the coastal waters in the Gulf of Mexico [4].

Currently, drifters are deployed in a singular fashion, and each drifter reports data via expensive satellite uplink instead of to other drifters or data mules (such as gliders, buoys or ships). The topic of fleets of surface level drifters using inexpensive acoustic, radio or optical communication is today a interesting research topic.

Networks of mobile wireless sensor nodes, however, are currently being investigated in shallow and deep sea applications such as the NEPTUNE project. For example, the Starbug AquaFlecks and Amour AUV, developed by MIT, are an underwater sensor network platform based on Fleck motes developed jointly by the Australian Commonwealth Scientific and Research Organization (CSIRO) and MIT CSAIL [5]. The 4m long AquaFlecks are combined with a mobile Amour AUV which acts as a data mule to retrieve data from the different sensor nodes. The Amour AUV uses 2 types of communication medium, i.e. an acoustic modem for long range communication and optical modem for short range. The WHOI acoustic modem has a data rate of 220 bits/s over 5000m, while the Aquacommod acoustic modem has a throughput of 480bit/s with range of over 200m consuming 4.5mJ/bit.

2.2 Wireless communication networks for ocean environments

Typically, underwater sensor nodes are connected to a network's surface station which connects to the Internet backbone through satellite communication or an RF link. The sensor nodes located in shallow or surface waters use diverse wireless communication technologies such as radio, acoustic, optical or electromagnetic signals. The different technologies vary with regard to communication *range* of the sender, *data rate* per second (data propagation speed), *energy consumption* and *robustness* with regard to noise or interference (such as Doppler effects) [6].

Radio signals are used in shallow water sensor networks, however, the travel speed of radio signals through conductive sea water is very low, i.e. about at a frequency of 30-300Hz. Experiments performed at the University of Southern California using Berkeley Mica2 Motes have reported to have a

transmission range of 120 cm in underwater at a 433MHz radio transmitter [7]. Optical waves do not suffer from such high attenuation but are affected by scattering. Also, transmission of optical signals requires high precision in pointing the narrow laser beams, which is less practical in water.

Basic underwater acoustic networks (UWA) are the most commonly used communication media for water-based sensor networks [8, 9]. Acoustic communication is formed by establishing two-way acoustic links between various sensor nodes. UWA channels, however, differ from radio channels in many respects. The available bandwidth of the UWA channel is limited, and depends on both range and frequency; the propagation speed in the UWA channel is five orders of magnitude lower than that of the radio channel. UWA networks can be distinguished into very long range, long, medium, short and very short communication range. As a rule, the shorter the communication range, the higher the bit rate. Typical ranges of acoustic modems vary between 10km to 90km in water. Furthermore, acoustic networks can be classified as horizontal or vertical, according to the direction of the sound wave. There are also differences in propagation characteristics depending on direction. Furthermore, acoustic signals are subject to multipath effects [10], large Doppler shifts and spreads, and other nonlinear effects.

Acoustic operation is affected by sound speed. Overall, the bit rate in water is about five orders of magnitude lower than in-air transmission. Sound speed is slower in fresh water than in sea water. In all types of water, sound velocity is affected by density (or the mass per unit of volume), which in turn is affected by temperature, dissolved molecules (usually salinity), and pressure. Today, the desired information transmission rate in the network is 100bit/s from each node. The available (acoustic) frequency band is 8-15 kHz. Uncertainty about propagation delays is typical of acoustic communication. Information is transmitted in packets of 256 bits, and nodes transmit at most 5 packets/h. Typical deployment of nodes can be as drifters or mounted on the ocean bottom, and separated by distances of up to 10km [11].

2.3 Data management for ocean sensor networks

Drifters are deployed to continuously collect data. At minimum, the end user is interested in the trajectory of the drifter itself since it contains relevant information about the ocean dynamics in the area covered. Furthermore, drifter platforms can be equipped with diverse sensors to sample the water. Today, salinity and temperature sensors are the most commonly used sensors. Drifter platforms can also carry accelerometers to measure wave speed or water acceleration for tsunami detection. Biological sensors detect marine microorganisms such as algae species and distribution.

Currently, drifters sense, store, and aggregate data locally until it is uploaded once a day via satellite connection to a centralized computer. Today, point sampling is common; region sampling via several collocated drifters during the same time period is rare. Typically, local data logger applications are run that contain limited processing and computing intelligence. Data collection is file-based, and reported in batch mode.

3. PROBLEM DEFINITION

3.1 Research Statement

Consider a set of n drifters $D = \{d_1, \dots, d_n\}$ deployed in the ocean to monitor a coastal area of interest. Let (t_i, x_i, y_i) be the time and location of initial deployment of drifter d_i . Drifters are designed to be passively propelled by local currents, are location-aware (using GPS), and are equipped with a variety of sensor devices to monitor different properties of the ocean surface. All drifters have local wireless communication capabilities that allow them to exchange messages with other drifters within range R ¹. A subset of the drifters ($B \subseteq D$) also has satellite connectivity, which allows them to propagate sensor data to oceanographers and other interested users around the globe. We refer to these special-purpose drifters as *mobile base-stations*, or simply *base-stations*. We thus view the set of drifters as a hierarchical mobile ad hoc network, wherein simple drifters forward their readings hop-by-hop to one of the mobile base-stations.

In order to predict drifter movement, we use a dataset *CUR* of coarse-grained radar measurements of current speed and current direction. Radar measurements are taken at regular intervals (e.g. every 1 hour) at various junction points of a grid spanning the area of interest (e.g. one pair of (speed, direction) measurements per $4km \times 4km$ grid cell). Current speed and direction conditions at all other locations are estimated using spline two-dimensional interpolation. Based on these current speed and direction measurements, we evaluate drifter locations over time, and we use the resulting trajectories as input to our simulations. In a real setting, drifter trajectories would be derived directly via GPS.

In this paper, we focus on empirically quantifying two aspects of drifter behavior: communication connectivity and sensing coverage.

Communication connectivity: We use two metrics of communication connectivity: i) *one-hop connectivity* of a drifter, which is the percentage of drifters within communication range of that drifter, and ii) *multi-hop connectivity* of the network, which is the percentage of drifters that can reach at least one of the base-stations on a multi-hop path. One-hop connectivity provides useful insight into whether drifters travel in clusters, or whether they disperse quickly out of each other's range soon after they are deployed in the ocean. Multi-hop connectivity is useful for quantifying the ability of drifters to relay their readings hop-by-hop to one of the base-stations, and eventually to the end-users.

Sensing coverage: We use two metrics of sensing coverage: i) *sensing density*, which is the number of connected drifters with multi-hop connectivity within the area of interest and ii) *sensing uniformity*, which denotes whether drifters are uniformly dispersed in the area of interest or congested in a small part of it. To quantify sensing uniformity, we adopt the definition of MRD (Mean Relative Deviation) proposed by Ferentinos and Tsiligiridis [12]:

¹In reality, the communication range is not a perfect circle, and the delivery ratio depends not only on the distance, but on a variety of environmental conditions. We leave the study of realistic communication models in ocean environments for future work.




Figure 1: One-hop connectivity: Percentage of drifters within communication range (1 mile) of each drifter (Liverpool Bay)

$$MRD = \frac{\sum_{i=1}^N |\rho_{S_i} - \rho_S|}{N\rho_S}$$

where N is the number of equally-sized overlapping sub-areas that the entire area of interest is divided into. Sub-areas are defined by four factors: two that define their size (length and width) and two that define their overlapping ratio (in the two dimensions). In the formula above, ρ_{S_i} is the spatial density of connected drifters within sub-area i and ρ_S is the spatial density of connected drifters in the entire area of interest. Thus, MRD is defined as the relative measure of the deviation of the spatial density of drifters in each sub-area to the spatial density of drifters in the entire area of interest. Perfect uniformity (MRD=0) is achieved when each sub-area has the same spatial density as that of the entire area of interest, while higher MRD values correspond to lower uniformity levels of drifters.

Given the set of drifters D deployed at specific times and locations, the subset of base-stations B and a real dataset of current information CUR that determines drifter trajectories, we would like to address the following questions:

- How is the one-hop connectivity of a drifter affected by the communication range, deployment location and deployment period?
- How is the multi-hop network connectivity affected by the number of base-stations, deployment location and deployment period?
- How is sensing density affected by the deployment location and deployment period?
- How is sensing uniformity affected by the deployment location and deployment period?

3.2 Application Background

Marine microorganism such as phytoplankton are exceedingly small (2-3 μ m), and are distributed at varying density in the ocean water. The 2005 bloom of *Alexandrium fundyense* at the New England coast was the most widespread outbreak of 'red tide' since a hurricane in 1972 spread the toxic algae throughout southern New England; the phenomenon received its name for the rust color that intense concentrations of algae sometimes paint ocean water. The type of red tide algae in New England contaminates shellfish, and can make people who eat the shellfish sick.

In most years, *Alexandrium fundyense* grows to toxic levels in Penobscot Bay and Casco Bay in Maine and in Candas Bay of Fundy. The more intense blooms can lead to the shut down of clam, oyster, and mussel beds to avoid paralytic shellfish poisoning of humans. The potent neurotoxin from *Alexandrium* accumulates in the meat of filter-feeding bivalves. While it does not harm them, it can cause paralysis and respiratory problems in humans and other animals that eat the shellfish.

In 2005, concentrations of toxic algae reached levels 40 times the norm, and the plants spread southward to regions of Cape Cod Bay, Massachusetts Bay, Nantucket Sound, and Buzzards Bay that are usually not affected by this species. Shellfish beds in Massachusetts, Maine, and New Hampshire, as well as 15,000 square miles of federal waters, were closed for more than a month at the peak of the seafood harvesting season.

Since the distribution of the algae is mainly influenced by ocean currents, our objective is to find out more information about ocean surface current dynamics by using current-propelled drifters. The floating drifters can have sensors attached, which measure algae occurrence. Overall, this information and ocean drifter network can be used as a monitoring, and early warning for red tide dangers.

4. SYSTEM EVALUATION

4.1 Experimental setup

In order to empirically address the questions posed in Section 3, we considered two realistic scenarios of deploying drifters in Liverpool Bay (UK) and the Gulf of Maine (USA). We used real datasets of surface current measurements monitored in the two coastal areas, to infer how drifters would move under the influence of these currents.

Liverpool Bay dataset: This data has been provided by the Proudman Oceanographic Laboratory Coastal Observatory Project, and it was measured by a 12-16MHz WERA HF radar system, which has been deployed to observe sea surface currents and waves in Liverpool Bay. In our simulations, we use current direction and current speed data measured hourly at the junction points of a 8×11 grid. The size of each grid cell is $4km \times 4km$, and thus the size of the monitored area is $28km \times 40km$. Current speed and direction conditions at locations inside the grid (other than the grid junctions) are estimated using two-dimensional spline interpolation. Drifter locations are estimated every 5 min-




Figure 2: Effect of communication range on average one-hop connectivity (Liverpool Bay)

utes. We simulate the deployment of 10 drifters from a single point, one drifter at a time at 1-hour intervals. Our simulations typically last for 1.5 days, which corresponds to 432 5-min time-steps. The default communication range is set to 1 mile and the default initial deployment location to 10km east and 4km north from the bottom left point of the grid.

Gulf of Maine dataset: This data was provided by the University of Maine’s Physical Oceanography Group, covering four months (March, June, September and December) of 2005. It was measured by a 4.3-5.4 MHz SeaSonde HF radar system, which is deployed to observe sea surface currents in the Gulf of Maine. In our simulations, we use current direction and current speed data measured hourly at the center of cells in a 36x24 grid. The size of each grid cell is 16km x 16km. Current speed and direction conditions at locations throughout the grid (other than at cell centers) are estimated using two-dimensional Gaussian interpolation. We simulate the deployment of 10 or more drifters from a single point, one drifter at a time at 1-hour intervals. Our simulations generally last for 30 days, which corresponds to 720 1-hr time-steps. The default communication range is set to 1 mile and the default initial deployment location to: Lat. 43.5° N, Long. -67° W.

4.2 One-hop connectivity

A drifter evaluates its one-hop connectivity by dividing the number of drifters within its communication range by the total number of drifters excluding itself (neighbors / 9). Figure 1 shows the one-hop connectivity of each one of the 10 drifters as time elapses, assuming a communication range of 1 mile. Since drifters are deployed every hour and time-steps last 5 minutes, the 1st drifter is deployed at time-step 0, the 2nd at time-step 12, and the 10th at time-step 108. From time-step 108 to 432, observe that most drifters have 2 to 3 immediate neighbors, with the exception of the 7th drifter, which has zero connectivity most of the time. For the particular 1.5-day deployment period, Figure 1 shows no consistent change in one-hop connectivity for all drifters as time passes, hence no clear trends of drifter dispersion or clustering.

Effect of communication range: To further investigate this, we measure average one-hop connectivity of the 10 drifters, and examine the impact of communication range on connectivity. Figure 2 confirms what the average one-hop connectivity does not deteriorate with time, except when




Figure 3: Effect of communication range on average one-hop connectivity (Gulf of Maine)




Figure 4: Effect of deployment location on average one-hop connectivity (Liverpool Bay)

the communication range is relatively low ($R=0.5$ mile). In addition, one-hop connectivity does not increase quadratically with the communication range as one would expect if drifters were uniformly distributed in the coastal area. In fact, increasing the communication range from 2 to 4 miles has only a small impact on the size of drifter neighborhoods. One-hop connectivity results in the Gulf of Maine (Figure 3) are also largely similar, as one might expect. In nearly all cases, however, one does observe a sharp decrease in connectivity begin within the first 48 hours.

Effect of deployment location: The previous two simulations concerned the same initial deployment location and deployment period. In order to draw more general conclusions about drifter behavior, we proceeded to evaluate the spatial and temporal variations of one-hop connectivity. Figure 4 shows one-hop connectivity values at different initial deployment locations; these values are averaged over all drifters during a particular deployment period. First, observe that the deployment location plays an important role in predicting the sizes of drifter neighborhoods. For example, if drifters are deployed at $[23km, 4km]$ they are likely to quickly cluster together and form a fully connected graph, whereas if they are deployed at $[20km, 12km]$ they have at most 1 to 2 neighbor drifters on average. Designers of ocean sensor networks should take this variability into account in order to select a deployment location that will yield sufficient connectivity for their purposes.




Figure 5: Effect of deployment period on average one-hop connectivity (Liverpool Bay)




Figure 6: Effect of deployment period on average one-hop connectivity (Gulf of Maine, comm. range = 1 mile)

Effect of deployment period: Surface currents, and thus drifter movements, are influenced by weather conditions, and thus vary across different time periods. Figure 5 measures one-hop connectivity at different 1.5-day deployment periods, averaged over 10 drifters deployed from a particular location. Depending on the deployment period, a drifter is shown to have from 2 to 9 drifters within a 1-mile range. In certain deployment periods (e.g. on days [1 – 2.5]), the average drifter connectivity remains relatively stable, whereas in others (e.g. on days [3 – 4.5]), it increases with time. In most periods, we observe signs of drifter clustering, rather than drifter dispersion. Temporal variability is also seen in the Maine simulations (Figure 6), where drifters deployed in the winter maintain their clusters longer than those deployed in the summer, indicative of the comparatively slower ocean current velocities occurring in the Gulf of Maine at that time of year. In future work, it would be interesting to associate one-hop connectivity values with weather conditions (e.g. wind and temperature) in order to be able to schedule drifter deployment in suitable time periods.

Effect of cluster deployment: We simulated deploying all drifters at the same time at a single point in the Gulf of Maine to test its effect on connectivity as compared to interval-based deployment. Figure 7 shows that single-hop connectivity declines precipitously in the first fifty hours, as we have come to expect, and all connectivity is lost within

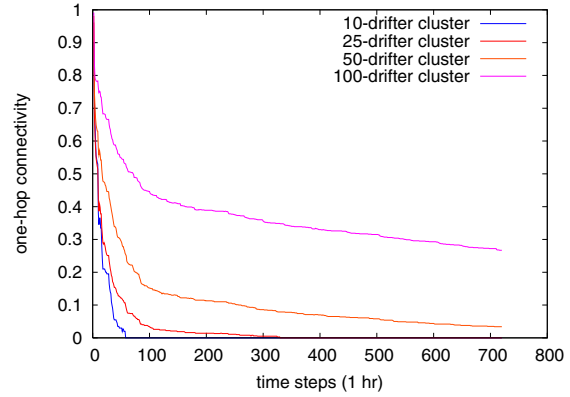


Figure 7: Effect of number of drifters deployed hourly at the same location (Gulf of Maine)




Figure 8: Multi-hop connectivity per base-station drifter (Liverpool Bay)

the first few days. However if the drifters are deployed in larger groups (e.g. 50-100) then some connectivity can still be observed up to a month later.

4.3 Multi-hop connectivity

The experiments on one-hop connectivity showed that drifters rarely form a fully-connected network. In order to send their sensor readings to the end-users they must either be base-station drifters with satellite connectivity, or relay their readings hop-by-hop to one of the *base-station* drifters. Out of the 10 drifters, we consider that up to 3 drifters (1st, 4th and 7th) act as base-stations. Figure 8 shows the percentage of drifters attached directly (through one hop) or indirectly (through multiple hops) to each one of the three base-stations. All drifters are deployed from the same default location [10km, 4km] during days [1 – 2.5] and the default communication range is 1 mile. Between time-steps 110 and 250, the three base-stations form three disjoint clusters; after time-step 250, the two clusters led by base-stations 1 and 3, merge into one cluster with multi-hop connectivity oscillating between 30% and 60%.

Effect of number of base-stations: The question that arises is how many base-stations we need to ensure multi-hop connectivity close to 100%. By summing up multi-hop connectivities for all base-stations in Figure 8, we observe that the first base-station alone is able to keep up to 50% of the drifters connected. The second base-station increased




Figure 9: Effect of number of drifter base-stations on the total multi-hop connectivity of the entire network (Liverpool Bay)




Figure 11: Effect of deployment period on multi-hop connectivity (Liverpool Bay)




Figure 10: Effect of deployment location on multi-hop connectivity (Liverpool Bay)

multi-hop connectivity significantly (up to 100%), whereas the third base-station does not yield any added benefits. This was observed because, in the particular deployment, two of the three clusters of drifters (around two of the base-stations) quickly merged into one cluster. To study the effect of base-stations in the more general case, we evaluated multi-hop connectivity averaged over multiple initial deployment locations and deployment periods. Figure 9 shows that 1, 2 and 3 base-stations are able to keep 60%, 80% and 90% of the drifters connected respectively.

Effect of deployment location: The next question that we addressed is whether multi-hop connectivity is sensitive to the drifter deployment location. We set the number of base-stations to 3 and for each deployment location, we evaluate multi-hop connectivity over a large number of deployment periods. Figure 10 shows that, depending on the deployment location, 60% to 100% of drifters are able to relay their readings to the end-users through the base-stations. Multi-hop connectivity is less sensitive to changes in the deployment location than one-hop connectivity. In contrast to one-hop connectivity that increases with time in the first 1.5 days for most deployment locations, multi-hop connectivity remains relatively stable for all deployment locations.

Effect of deployment period: The effect of deployment period on multi-hop connectivity is very similar to the effect of deployment location. Figure 11 shows multi-hop net-




Figure 12: Effect of deployment location on sensing density (Liverpool Bay)

work connectivity for different deployment periods. Observe that, depending on the weather at the time that drifters are deployed, their ability to remain connected varies between 60% and 100%. Multi-hop connectivity remains relatively constant throughout any of the 1.5-day deployment periods.

4.4 Sensing density

Sensing density is defined as the number of connected drifters within an *area of interest*; drifters counted must be connected directly or indirectly to at least one of the base-stations. The main distinction between multi-hop connectivity and density is that the former counts the percentage of connected nodes overall, whereas the latter counts the number of connected nodes within an area of interest. In our experiments, the area of interest is of constant size, but not centered around a fixed location. It is defined as follows: we deploy 10 drifters and let them be propelled by currents for 1.5 days. We evaluate the bounding box of their trajectories, and we select as *area of interest* the square area of size $15km \times 15km = 225km^2$ centered at the center of the bounding box. We use the same area of interest to evaluate sensing uniformity in the next subsection.

Effect of deployment location: Figure 12 shows changes in sensing density with time, for 5 different deployment locations. Sensing density is shown to be drastically affected by the deployment location. For example, by deploying drifters at $[13km, 9km]$ we manage to keep most of them within the




Figure 13: Effect of deployment period on sensing density (Liverpool Bay)

area of interest for the first day of the experiment, whereas by deploying them at $[23km, 4km]$ we observe that only 2-4 drifters remain in the area of interest. Sensing density may also fluctuate from 1 to 10 drifters within a day for the same initial drifter deployment (e.g. $[20km, 12km]$). By comparing Figures 10 and 12, we see that sensing density is more sensitive than multi-hop connectivity to deployment location.

Effect of deployment period: Figure 13 shows changes in sensing density with time, for 7 different deployment periods of 1.5 day each. The effect of deployment period on sensing density is very similar to the effect of deployment location discussed above. That is, the number of drifters within the area of interest fluctuates wildly across different deployment periods, as well as within a certain deployment period. If oceanographers are interested in monitoring a particular coastal area, they must select very carefully both the deployment location and the deployment period, in order to obtain a large number of sensor readings from that area. In the future, it would be interesting to associate weather conditions on different days, or known current patterns at certain locations, with the observed sensing density values.

4.5 Sensing uniformity

To quantify sensing uniformity, we adopt the definition of MRD (Mean Relative Deviation) as discussed in Section 3 and in [12]. In our experiments, we set the side lengths of each sub-area to be half the side length of the entire area of interest; we also set the overlapping ratio to half the side length of a sub-area. This means that there are 9 sub-areas within the area of interest. Figure 14 shows how MRD varies with time in a particular instance of deploying 10 drifters for 1.5 days, with 3 of the 10 drifters acting as base-stations. MRD values vary between 0.8 and 1.9 during the 1.5 day period of the experiment.

Figure 15 shows the positions of drifters within the area of interest in three different time-steps when MRD values are evaluated to 0.8 (see star positions), 1.3 (see circle positions) and 1.8 (see diamond positions) respectively. Notice that the distribution of drifters within the area of interest is not significantly different for the three MRD values.

Effect of deployment location and deployment period: Sensing uniformity fluctuates mildly (between 0.8 and 1.9) with time for several different deployment locations as shown in Figure 16. However, there is no deployment loca-




Figure 14: Sensing uniformity (MRD) of drifters deployed one-by-one at a particular location in the Liverpool Bay and left to float for 1.5 days. MRD is a measure of how drifters are dispersed in the area of interest. Higher MRD values correspond to lower uniformity.




Figure 15: Drifter locations in three time-steps with MRD values 0.8 (star positions), 1.3 (circle positions) and 1.8 (diamond positions) respectively. The inner square represents the area of interest and the outer rectangle represents the minimum bounding box of drifter trajectories. (Liverpool Bay)

tion that consistently yields high or low MRD values during the 1.5-day period of the experiment. Similarly, Figure 17 shows that MRD values fluctuate mildly (between 0.8 and 1.9) for most deployment periods, and no particular deployment period yields consistently higher MRD values than others.

5. CONCLUSIONS AND FUTURE WORK

In this paper, we proposed a new model of ocean monitoring using a fleet of drifters passively propelled by surface currents. These drifters form an adhoc sensor network, they task their local sensor devices to monitor coastal waters, and propagate their readings hop-by-hop to a few mobile base-stations. Our study on network connectivity and sensing coverage of the drifter network led us to the following conclusions.

For most drifters, their one-hop connectivity increases sub-linearly with the communication range (instead of quadrat-




Figure 16: Effect of deployment location on uniformity (Liverpool Bay)




Figure 17: Effect of deployment period on uniformity (Liverpool Bay)

ically) revealing non-uniform drifter distribution. Deployment locations and deployment periods seem to play an important role in one-hop connectivity.

Multi-hop connectivity is clearly improved by the increase in the number of available base-station drifters; an average connectivity of around 60% with one base-station becomes around 80% with two base-stations and almost 90% with three base-stations based on our Liverpool Bay measurements. Given a fixed number of base-stations, we observed that different deployment locations and deployment periods play an important role in the network’s multi-hop connectivity, with connectivity varying between 60% and 100%.

Sensing density is drastically affected by the deployment period and by the initial drifter location. Moreover, generally, the difference between sensing density and network connectivity increases with time, indicating that drifters may stay highly connected in several cases after a long period of time (over one day) but this does not necessarily mean that they are kept inside the corresponding area of interest. Actually, in most cases, sensing density decreases significantly after the first 24 hours.

Finally we observed no clear trends showing how deployment locations and deployment periods influence the values of sensing uniformity of connected drifters. In fact, the observed differences in uniformity values across different deployment scenarios was shown to be negligible, at least during the relatively short simulation periods (1.5 days).

In the future, we plan to verify our findings by deploying a real network of drifters and investigating the communication connectivity and sensing coverage in a real scenario. After verifying our findings, we would like to take them into account in designing routing and in-network processing algorithms that are particularly tailored for ocean monitoring applications. This will lead us to investigate important tradeoffs between delay, communication cost and reliability of disseminating sensor data from drifters to end-users.

6. REFERENCES

- [1] N. Pettigrew, J. H. Churchill, C. Janzen, L. Mangum, R. Signell, A. Thomas, D. Townsend, J. Wallinga, and H. Xue, “The kinematic and hydrographic structure of the gulf of maine coastal current,” *Deep Sea Research II*, vol. 52, pp. 2369–2391, 2005.
- [2] e. a. J. Gould, “Argo profiling floats bring new era of in situ ocean observations,” *EOS*, vol. 85, no. 19, pp. 179–184, 2004.
- [3] U. of Washington.
- [4] R. P. LaBelle, J. Price, W. Johnson, and C. Marshall, “Surface drifter deployment in experimental oil spill.”
- [5] I. Vasilescu, K. Kotay, D. Rus, M. Dumbabin, and P. Corke, “Data collection, storage, and retrieval with an underwater sensor network,” in *SenSys ’05: Proceedings of the 3rd international conference on Embedded networked sensor systems*. New York, NY, USA: ACM Press, 2005, pp. 154–165.
- [6] I. F. Akyildiz, D. Pompili, and T. Melodia, “Challenges for efficient communication in underwater acoustic sensor networks,” *ACM SIGBED Review*, vol. 1, no. 2, pp. 3 – 8, 2004.
- [7] B. Zhang, G. Sukhatme, and A. Requicha, “Adaptive sampling for marine microorganism monitoring,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2004.
- [8] I. F. Akyildiz, D. Pompili, and T. Melodia, “Underwater acoustic sensor networks: research challenges,” *Ad Hoc Networks*, vol. 3, no. 1, pp. 257–279, 2007.
- [9] J. Proakis, E. Sozer, J. Rice, and M. Stojanovic, “Shallow water acoustic networks,” *IEEE Communications Magazine*, vol. 39, no. 11, pp. 114–119, 2001.
- [10] M. Stojanovic, “Recent advances in high-speed underwater acoustic communication,” *IEEE Journal of Oceanographic Engineering*, vol. 21, pp. 125–136, 1996.
- [11] E. Sozer, M. Stojanovic, and J. Proakis, “Underwater acoustic networks,” *IEEE Journal of Oceanic Engineering*, vol. 25, no. 1, pp. 72–83, 2007.
- [12] K. P. Ferentinos and T. A. Tsiligiridis, “Adaptive design optimization of wireless sensor networks using genetic algorithms,” *Elsevier Computer Networks*, vol. 51, no. 4, 2006.