Impact of Drifter Deployment on the Quality of Ocean Sensing

Konstantinos P. Ferentinos[1], Niki Trigoni[2], and Silvia Nittel[3]

[1] Agricultural University of Athens, Athens, Greece
kpf3@cornell.edu
[2] University of Oxford, Oxford, UK
niki.trigoni@comlab.ox.ac.uk
[3] University of Maine, Orono, USA
nittel@spatial.maine.edu

Abstract. Traditional means of observing the ocean, like fixed moorings and radar systems, are expensive to deploy and provide coarse-grained data measurements of ocean currents and waves. In this paper, we explore the use of an inexpensive wireless mobile ocean sensor network 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 basestations, and ii) a large number of low-power drifters with short-range acoustic or radio connectivity. We study connectivity and uniformity properties of the ad-hoc mobile sensor network. We investigate the effect of deployment strategy. The objective of this paper is to address the following challenge: how can we trade the usage of resources (e.g. number of drifters, and number of basestations vs. communication range) and which deployment strategy should be chosen (e.g. grid-like, star-like, etc.) to minimize energy costs, whilst satisfying application requirements for network connectivity and sensing density. Using simulation and real dataset, we investigate the effects of deploying drifters with regard to the following questions: i) where/when should drifters be placed initially? ii) how many drifters should initially be deployed?, iii) the effect of the number of basestations (drifters with satellite connectivity) on the overall network connectivity, and iv) the optimal communication range of the basic drifters. Our empirical study provides useful insights on how to design distributed routing and in-network processing algorithms tailored for ocean-monitoring sensor networks.

1 Introduction

Marine microorganisms such as phytoplankton are exceedingly small (2-3 mm), and coastal ocean currents are a significant factor influencing the transport and circulation of these marine microorganisms. Establishing a detailed model of

S. Nittel, A. Labrinidis, and A. Stefanidis (Eds.): GSN 2006, LNCS 4540, pp. 9–24, 2008.

ocean currents in a relatively small area such as a bay or a specific coastline is important since currents carry nutrients and potential toxins, which affect ecosystems and humans in coastal regions. For example, currents can distribute the Alexandrium fundyense algae during the warm summer months. This type of algae contaminates shellfish, and can make people who eat the shellfish seriously ill [7]. Today, the knowledge of the behavior of ocean and coastal currents is limited, and predicted trajectories of armful algae imbedded in the flow are inadequate from the point of view of public health.

Today, major ocean currents are established using coastal radar; however, the information is spatially and temporally coarse, and the gaps need to be filled in by interpolation techniques. In this paper, we investigate the alternative deployment of a fleet of inexpensive, networked ocean drifters, which are passively propelled by the currents and report their GPS-based location and trajectories to the end user.

Today's available drifter platforms such as [4] are deployed in a singular fashion, and use expensive satellite communication to upload data to a centralized computer. Satellite communication is high cost, from the standpoint of the equipment price, with regard to energy consumption and also the service fee for the satellite link service. In this paper, we explore the use of a fleet of inexpensive drifters as an alternative flexible infrastructure for fine-grained ocean monitoring. We view the fleet of drifters as wireless ad-hoc sensor network with two types of nodes: i) only a few powerful drifters with satellite connectivity, acting as mobile basestations, and ii) a large number of low-power, inexpensive 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 currents on a detailed scale.

Deploying such a fleet of mobile ad-hoc sensor nodes on the ocean surface is a novel research problem, both from the perspective of computer science and oceanography. We have explored communication connectivity and sensing uniformity of a fleet of a mobile ad-hoc sensor network using real datasets from the Liverpool Bay. 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 [5]. The objective of this paper is to address the following challenge: how can we trade the usage of resources (e.g. number of drifters, and number of basestations vs. communication range) and which deployment strategy should be chosen (e.g. grid-like, star-like, etc.) to minimize energy costs, whilst satisfying application requirements for network connectivity and sensing density. Using simulation and real dataset, we investigate the effects of deploying drifters with regard to the following questions: i) where/when should drifters be placed initially? ii) how many drifters should initially be deployed?, iii) the effect of the number of basestations (drifters with satellite connectivity) on the overall network connectivity, and iv) the optimal communication range of the basic drifters. 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: in Section 2, we provide relevant technical background on the current state of the art of ocean sensor networks, drifter platforms and wireless communication technology for water environments. In Section 3, we explore the research questions and the approach of this paper in more detail. Section 4 and Section 5 contain our experimental results and we conclude with Section 6.

2 Background

In this section, we provide background information about 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 [6]). Another main research domain is in near-coastal ocean observations using fixed, large moorings equipped with sensor and satellite connection for data upload. These sensor environments are extended with coastal radar stations, autonomous gliders, and research vessels. Our interest in ocean surface drifters is with regard to near-coastal deployments.

2.1 Ocean Drifters

Today, several projects and platforms for shallow water drifters exist. Initial larger-scale deployments of drifters were in the Gulf of Mexico and the South-western Caribbean Sea designed to explore the Gulf Stream in more detail.

Today, the international ARGO project [4] is one of the largest deployments of drifters in the world oceans. ARGO is an program that began in 2000, and now the deployment of 3000 profiling drifters is 100% complete. The purpose of ARGO is to examine the global ocean currents, circulation and air-sea interaction, with the goal of improving climate models and predictions.

The Argo drifter (also called "Davis Drifter") was designed to be a surface level (1 meter below surface) drifter which can report position via the Argos satellite link-based data collection system. Location determination by GPS is also available. The unit has a nominal operating life of 9 months. The global array of 3,000 floats is distributed roughly every 3 degrees (300km).

Currently, drifters are deployed in a singular fashion, and each drifter reports data via expensive satellite uplink instead 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, novel research topic.

Networks of mobile wireless sensor nodes are currently investigated in deep sea applications such as the NEPTUNE project or in small-scale surface deployments. 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 [11]. The 4 inch 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 two types of communication, i.e. an acoustic modem for long-range communication and optical-modem for short range. The acoustic modem has a data rate of 220 bits/s over 5 km, while the optical modem has a throughput of 480bit/s with range of over 200m consuming 4.5mJ/bit.

2.2 Wireless Communication 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) [1].

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 [12]. 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 [2,8]. 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

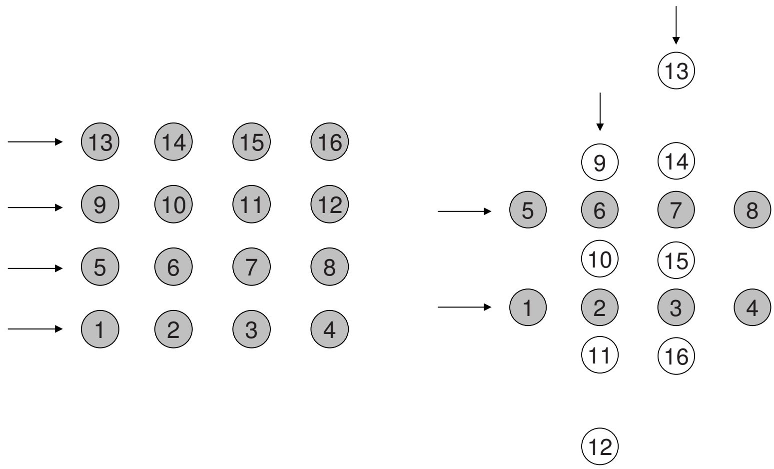

Fig. 1. Grid and hash deployments of drifters

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 [9].

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.

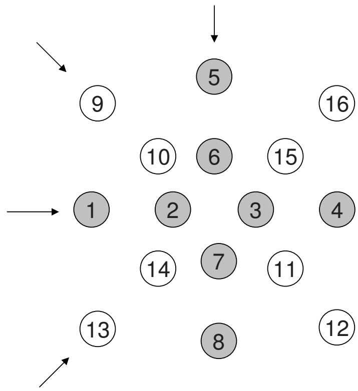

Fig. 2. Star deployment of drifters

3 Problem Description

Model assumptions: Consider a set of n drifters $D = \{d_1, \ldots, d_n\}$ 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.

[1] In 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.

Basestations	Grid Deployment	Hash Deployment	Star Deployment
1			
2	1,11	1,15	1,11
4	1,6,11,16	1,7,12,14	1,6,11,16
8	1,3,6,8,9,11,14,16	1,3,6,8,10,12,13,15	1,3,6,8,9,11,14,16

Fig. 3. Identifiers of drifters selected as basestations in the grid, hash and star deployments

Metrics: In this paper, we focus on empirically quantifying two aspects of drifter behavior: multi-hop connectivity and sensing coverage.
Multi-hop connectivity: This is defined as the percentage of drifters that can reach at least one of the base-stations on a multi-hop path. Multi-hop connectivity is useful for quantifying the ability of drifters to relay their readings hop-by-hop to the end-users through one or more base-stations.
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 [3]:

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

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 from 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.

Main objectives: 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 two questions:

– How does the deployment strategy, which is defined by the positions and timestamps of drifters when they are first deployed, impact the communication connectivity and sensing coverage of the fleet of drifters? In particular, our goal is to determine whether grid-like drifter formations are preferred to star-like drifter formations. We study this problem in Section 4.

– In a specific application framework, engineers will be faced with the question of how to balance network resources to achieve a certain density of connected drifters in an area of interest. They will have the option of varying the number of drifters, the number of basestations or the communication range between drifters. In this paper, we present a study of how these three factors impact density of connected drifters (Section 5). Based on this study, engineers can make informed decisions about how to build a network of drifters, taking into account the prices of drifters, basestations, or communication devices with varying ranges.

Real Datasets: In order to empirically address the questions posed before, we considered a realistic scenario of deploying drifters in the Liverpool Bay (UK). We used real datasets of surface current measurements monitored in the coastal area, to infer how drifters would move under the influence of these currents.

The Liverpool Bay 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 center of cells of a 7×10 grid. The size of each grid cell is $4km \times 4km$, and thus the size of the monitored area is $28km \times 40km$. A smaller 5×5 grid is considered to be the *area of interest* (of size $20km \times 20km$); recall that among connected drifters we only count those in the area of interest to measure sensing density, as defined in Section 3. This is in contrast with multi-hop connectivity, which is the percentage of all connected drifters, in and out of the area of interest. 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 minutes. Our simulations typically last for 1.5 days, which corresponds to 432 5-min time-steps.

4 Impact of Deployment Strategy

In our study, we compare the performance of three different deployment strategies. In all of them, we assume that four boats start deploying drifters simultaneously, at time t=0, and in 1-hour intervals. Four drifters are deployed per boat (at times t, t+1hr, t+2hr and t+3hr). The initial positions and routes of the four boats vary across different strategies:

– *Grid deployment:* Four horizontal lines (left deployment in Figure 1)
– *Hash deployment:* Two horizontal parallel lines and two vertical parallel lines (right deployment in Figure 1)
– *Star deployment*: Two pairs of perpendicular lines in a 45-degree rotation (four lines in star formation in Figure 2)

In all three deployments, boats scan the same square area of interest of size $20km \times 20km$. Each boat starts deploying drifters along a line as shown in

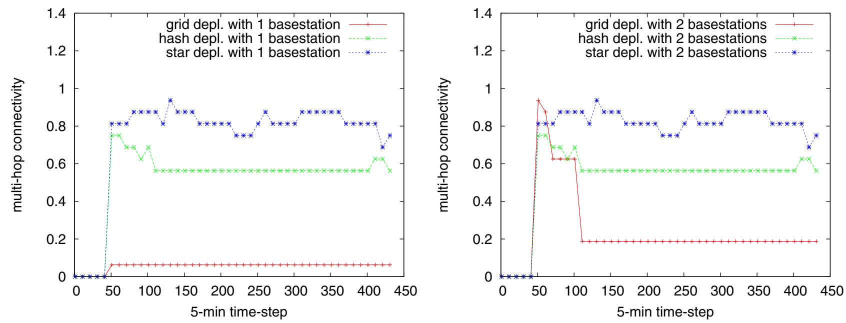

Fig. 4. Multi-hop connectivity with 1 and 2 basestations

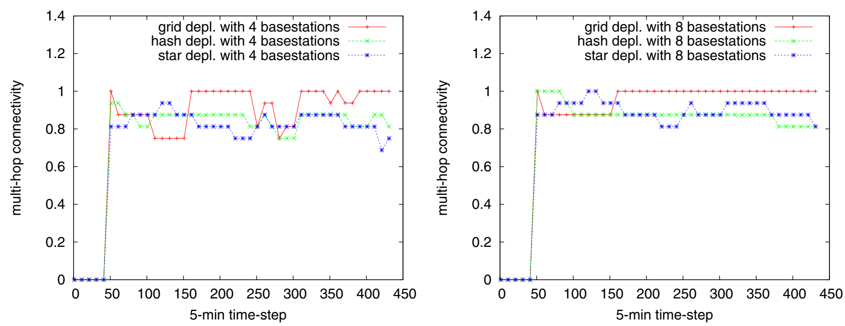

Fig. 5. Multi-hop connectivity with 4 and 8 basestations

Figures 1 and 2. The initial positions of the four boats, where the first four drifters are deployed, are illustrated by an arrow. It takes 1 hour for each boat to cross a distance of $4km$ where the next drifter is deployed.

Some of these drifters have satellite connectivity and act as mobile basestations. Table 3 shows which particular drifters are selected as basestations in scenarios with 1, 2, 4 and 8 basestations.

To compare the three deployment strategies we use the three metrics introduced in Section 3, namely multi-hop connectivity, sensing density and sensing uniformity (MRD: Mean Relative Deviation). We are interested in investigating which is the preferred deployment strategy in terms of connectivity and sensing, and how our decision may be affected by the number of basestations.

Multi-hop connectivity: Figures 4 and 5 show the behavior of the three deployment strategies as time elapses and as we increase the number of drifters acting as basestations. Notice that in scenarios with very few basestations, the grid deployment exhibits very low connectivity. In fact, only basestations manage to get readings through to the users using their satellite links, whereas all other drifters remain disconnected. The low connectivity is due to the fact that

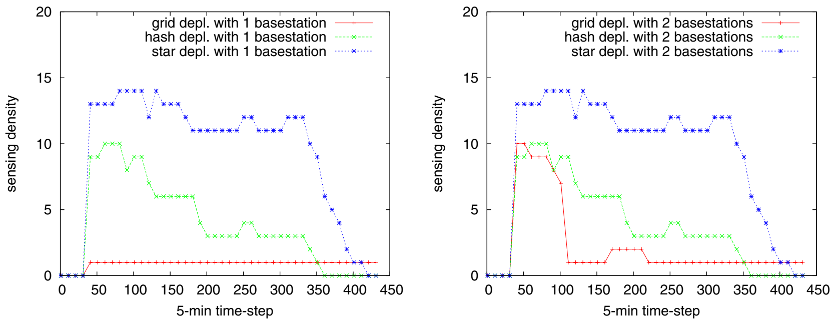

Fig. 6. Number of connected drifters in the area of interest (sensing density) with 1 and 2 basestations

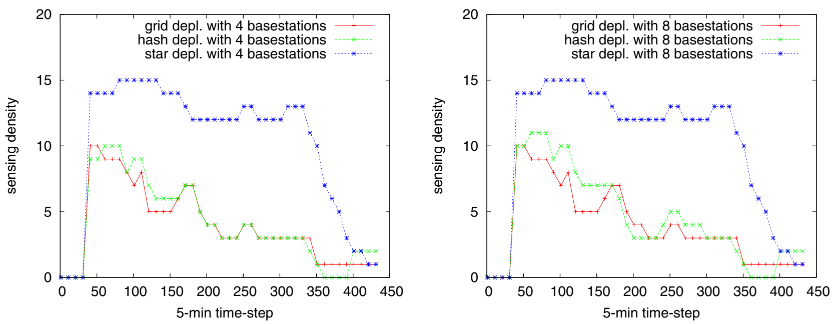

Fig. 7. Number of connected drifters in the area of interest (sensing density) with 4 and 8 basestations

the communication range (4 km) is equal to the distance between the locations of two consecutive drifters as they are initially deployed by a boat. Surprisingly the grid deployment exhibits very high multi-hop connectivity when 4 or 8 of the 16 drifters act as basestations.

The connectivity of the hash deployment is 10 times higher than that of the grid deployment in the case of 1 basestation, and 3 times higher in the case of 2 basestations. However, the relative difference between the two deployment strategies diminishes as we further increase the number of basestations (to 4 or more basestations).

The star deployment is shown to exhibit the highest multi-hop connectivity in most cases. Observe that the connectivity achieved with 1 basestation is very similar to that achieved with 8 basestations. This consistent behavior of the star deployment makes it very desirable, since it enables oceanographers to obtain readings from most drifters without using many expensive drifters with satellite connectivity (basestations). Before selecting the star deployment as the preferred strategy, we should first investigate whether it is equally efficient in terms of sensing density and sensing uniformity.

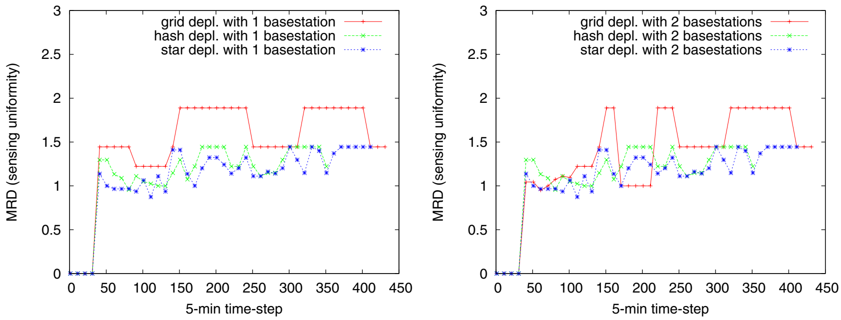

Fig. 8. Drifter uniformity in the area of interest (MRD) with 1 and 2 basestations

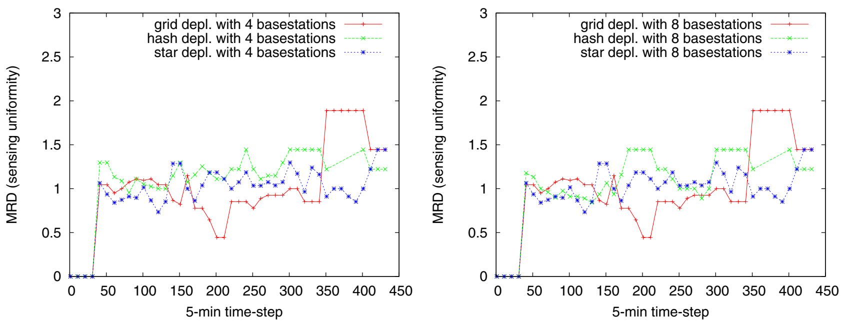

Fig. 9. Drifter uniformity in the area of interest (MRD) with 4 and 8 basestations

Sensing coverage: We examine two distinct metrics of sensing coverage, namely sensing density in Figures 6 and 7 and sensing uniformity in Figures 8 and 9. Recall that sensing density is the number of drifters within a region of interest that are connected through a multi-hop path to at least one of the basestations. In our experiments, the area of interest has size $20km \times 20km$, and each deployment has the same center as the area of interest.

Figures 6 and 7 show that for all deployment strategies, the number of connected drifters that remain within the area of interest decreases over time. It takes almost 400 time-steps ($400 \times 5 = 2000$ mins) for most drifters to be passively propelled by currents out of the area of interest.

In scenarios with few (1 or 2) basestations, the hash deployment yields higher density than the grid deployment. The gap, however, decreases significantly as we increase the number of basestations. The star deployment significantly outperforms the other two strategies in all cases, irrespective of the number of basestations. The density of drifters in the star deployment decreases more slowly than in the other two strategies up to time-step 300. Shortly after, we observe a very sharp descent of the density in the star deployment, until all drifters move outside the area of interest.

Observe that in scenarios with many basestations ($>= 4$), our results on multi-hop connectivity are very different from our results on sensing density. More specifically, in terms of multi-hop connectivity, the three deployment strategies have similar performance, whereas in terms of sensing density, the star deployment outperforms the others for prolonged time periods. This means that the star deployment places drifters so that most of them remain within the area of interest longer than if we had used grid or hash deployments.

The next question that arises is whether connected drifters remain uniformly dispersed in the area of interest, or whether they are clustered in small parts of the area. Figures 8 and 9 illustrate the performance of the three deployment strategies in terms of MRD (Mean Relative Deviation), which is a metric of sensing uniformity. The lower the MRD value the more uniform the spatial distribution of connected drifters. Recall that in scenarios with 1 or 2 basestations, the grid and hash deployments perform very poorly in terms of sensing density (Figures 6 and 7). Very few connected drifters are in the area of interest, and it is of little interest to compare the algorithms in terms of sensing uniformity. In scenarios with 4 and 8 basestations, the MRD values of the three deployment strategies fluctuate between values 0.5 and 2, with the grid deployment exhibiting slightly higher uniformity (lower MRD) than the other two approaches.

Summary of results: The star deployment exhibits significantly higher multi-hop connectivity than the hash and grid deployments in scenarios with few basestations (1 or 2 basestations out of 16 drifters). The three strategies behave similarly in scenarios with 4 or more basestations. The star deployment yields higher sensing density than the other two strategies in all cases. The grid deployment is slightly better in terms of uniformity, however, this benefit is too small to counteract the superiority of the star deployment in terms of sensing density. Our results are based on a dataset of real radar readings of current direction and velocity in the Liverpool Bay. They show that 1) the strategy used to deploy drifters has a significant impact on the network connectivity and sensing coverage, and 2) based on that dataset, the star deployment is preferred to the grid and hash deployments.

5 Resource Management

In the previous section, we compared three deployment strategies and selected the star deployment as the preferred approach. In this section, we examine variants of the star deployment in detail. Our aim is to understand the tradeoffs involved in varying the use of three network resources - number of drifters, number of basestations and communication range of simple drifters- to achieve a desirable sensing density. In each experiment, we fix the value of one of the three resources and measure the average sensing density as we vary the other two. Sensing density is averaged over 200 time-steps of 5 mins each (from time-step 100 to time-step 300).

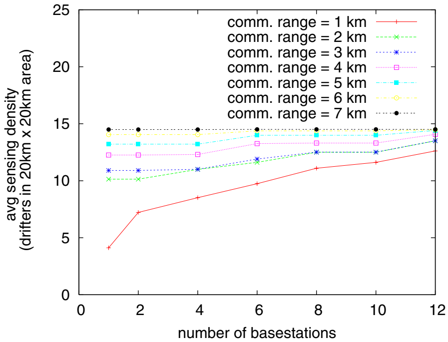

Fig. 10. Density of drifters in the area of interest as we vary the number of basestations and the communication range of simple drifters

We aim to address the following three questions:

1. How is sensing density affected by the number of basestations for various communication ranges? For example, if we want to increase the sensing density by δ, is it better to increase the communication range of drifters or convert some of the drifters into basestations by equipping them with satellite connectivity?
2. How is sensing density affected by the number of drifters for various communication ranges? If we want to increase network density by δ, is it better to increase the communication range of drifters of buy more drifters with the same communication range?
3. How is sensing density affected by the number of basestations for various numbers of drifters? If we want to increase network density by δ, is it better to increase the number of basestations or the number of drifters?

Figure 10 studies the tradeoff between the communication range of simple drifters and the number of basestations. The total number of drifters is set to 16. For small communication ranges, increasing the number of basestations is very beneficial in terms of sensing density. However, the benefit of adding a basestation is significantly decreased as we increase the communication range of simple drifters. Figure 10 shows that with 1 basestation and a 2 km communication range we observe an average of 10 drifters in the area of interest. To increase the sensing density from 10 to 14 drifters, we could either turn 10 drifters into basestations or we could increase the communication range of all drifters to 7 km. Depending on the price of each resource, network engineers can make an informed decision about how to reach their density goal.

Figure 11 studies the tradeoff between the total number of drifters and their communication range. The number of basestations is set to 2. Interestingly for small communication ranges (1-3km) the sensing density is initially sublinear in the number of drifters and then it becomes superlinear. Initially, when we insert a new drifter in the network, it might be disconnected from the other

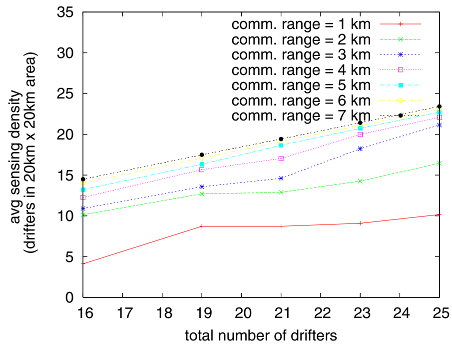

Fig. 11. Density of drifters in the area of interest as we vary the total number of drifters and their communication range

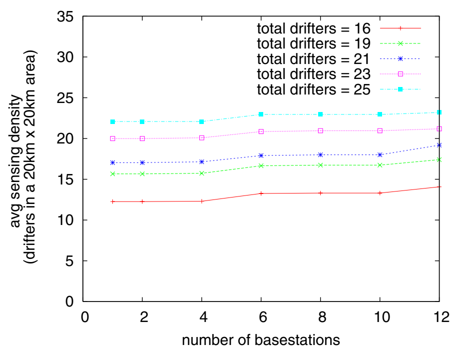

Fig. 12. Density of drifters in the area of interest as we vary the total number of drifters and the number of basestations

drifters because of the low communication range. However, as we keep inserting new drifters, density increases faster, because the new drifters connect not only themselves, but also previously disconnected drifters, to one of the basestations. Given a fixed number of basestations, if we want to increase network density, we have the option of increasing the number of drifters or their communication range. For example, Figure 11 shows that with 19 drifters and a 3 km communication range, we have on average 13 drifters in the area of interest. To increase the sensing density to 15 drifters we could either insert 3 more drifters or increase the communication range of all drifters to 4 km.

Figure 12 illustrates the tradeoff between the total number of basestations and the total number of drifters (sum of simple drifters and basestations). The communication range of simple drifters is set to 4 km. As we observed in Figures 6 and 7 of Section 4, the density of the star deployment is consistently high and it is not affected by the number of basestations. This is also confirmed in Figure 12, which shows that to increase network density, the only way is to increase the number of drifters.

6 Conclusions

In this paper, we investigated the impact of the drifter deployment strategy on the quality of ocean sensing performed using a fleet of drifters passively propelled by surface currents. These drifters form an adhoc sensor network; they monitor coastal waters with local sensor devices and propagate their readings hop-by-hop to a few mobile base-stations. Our comparison of the performance of three different deployment strategies, namely the grid, hash and star deployments, led us to the following conclusions.

The influence of deployment strategy over multi-hop connectivity depends on the number of basestations in the network. In cases with few basestations (1 or 2 basestations out of 16 drifters) the star deployment clearly exhibits significantly higher multi-hop connectivity than the hash and grid deployments. However, in cases with 4 or more basestations, all three strategies behave similarly. As far as sensing density is concerned, the star deployment outperforms the other two deployment strategies, irrespectively of the number of basestations in the network. In terms of uniformity of sensing points, the grid deployment is slightly better, however this benefit is too small to counteract the superiority of the star deployment in terms of sensing density. Thus, we conclude that 1) the strategy used to deploy drifters has a significant impact on the network connectivity and sensing coverage, and 2) the star deployment is preferred to the grid and hash deployments in the context of a network of drifters moving under real current conditions at Liverpool Bay.

Finally, we addressed three major tradeoff issues concerning ways to achieve specific increase in sensing density by an existing network of drifters. In the tradeoff between increasing the communication range of drifters or converting some of the simple drifters into basestations, we concluded that both ways are effective and the decision would be purely based on economical parameters. In the tradeoff between increasing the communication range of drifters or simply adding more drifters to the network, we concluded that for small communication ranges, increasing the communication range is more effective than adding more drifters, while for larger communication ranges both solutions are effective and again economical parameters come into play. Finally, in the tradeoff between adding more drifters to the network or converting some of the existing drifters into basestations, we concluded that the former approach is by far the most effective solution.

References

1. Akyildiz, I.F., Pompili, D., Melodia, T.: Challenges for efficient communication in underwater acoustic sensor networks. ACM SIGBED Review 1(2), 3–8 (2004)
2. Akyildiz, I.F., Pompili, D., Melodia, T.: Underwater acoustic sensor networks: research challenges. Ad Hoc Networks 3(1), 257–279 (2005)
3. Ferentinos, K.P., Tsiligiridis, T.A.: Adaptive design optimization of wireless sensor networks using genetic algorithms. Elsevier Computer Networks 51(4), 1031–1051 (2007)

4. Gould, J., et al.: Argo profiling floats bring new era of in situ ocean observations. EOS 85(19), 179–184 (2004)
5. Nittel, S., Trigoni, N., Ferentinos, K., Neville, F., Nural, A., Pettigrew, N.: A drift-tolerant model for data management in ocean sensor networks. In: MobiDE, pp. 49–58. ACM Press, New York (2007)
6. U. of Washington. The neptune project page
7. Pettigrew, N., Churchill, J.H., Janzen, C., Mangum, L., Signell, R., Thomas, A., Townsend, D., Wallinga, J., Xue, H.: The kinematic and hydrographic structure of the gulf of maine coastal current. Deep Sea Research II 52, 2369–2391 (2005)
8. Proakis, J., Sozer, E., Rice, J., Stojanovic, M.: Shallow water acoustic networks. IEEE Communications Magazine 39(11), 114–119 (2001)
9. Sozer, E., Stojanovic, M., Proakis, J.: Underwater acoustic networks. IEEE Journal of Oceanic Engineering 25(1), 72–83 (2000)
10. Stojanovic, M.: Recent advances in high-speed underwater acoustic communication. IEEE Journal of Oceanographic Engineering 21, 125–136 (1996)
11. Vasilescu, I., Kotay, K., Rus, D., Dunbabin, M., Corke, P.: Data collection, storage, and retrieval with an underwater sensor network. In: SenSys 2005: Proceedings of the 3rd international conference on Embedded networked sensor systems, pp. 154–165. ACM Press, New York (2005)
12. Zhang, B., Sukhatme, G., Requicha, A.: Adaptive sampling for marine microorganism monitoring. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (2004)