Low-Cost, Intelligent Drifter Fleet for Large-Scale, Distributed Ocean Observation

Eric Cocker, Julie A. Bert, Francisco Torres, Matthew Shreve, Jamie Kalb, Joseph Lee, Michael Poimboeuf, Paloma Fautley, Samuel Adams, Joanne Lee, Jengping Lu, Chris Chua, Norine Chang
Palo Alto Research Center
Palo Alto, CA
ecocker@parc.com; jbert@parc.com; torres@parc.com; mshreve@parc.com; jkalb@parc.com; joelee@parc.com; mpoimboeuf@parc.com; pfautley@parc.com; sadams@parc.com; jolee@parc.com; jplu@parc.com; chua@parc.com; nchang@parc.com

Steven Neltner
Bright Ceramic Technologies
Palo Alto, CA
Steven.Neltner@Bright-CT.com

Michael Gray
Dept. of Marine, Earth and Atmospheric Sciences
North Carolina State University
Raleigh, NC
magray@ncsu.edu

***Abstract*— We have developed Persistent Environmental Awareness Reporting and Location (PEARL) ocean drifters. PEARL drifters are small, rugged, low-cost, autonomous, environmentally friendly ocean drifters that represent a significant opportunity for high-impact applications in both national security and environmental ecosystem monitoring. Drifters record and report data which is processed by advanced edge analytics before being compressed for satellite transmission to populate a large data repository with sensor data that is combined and analyzed to discover signals of interest in the ocean environment with the goal of increasing distributed maritime awareness. Each drifter is entirely self-contained, powered by solar panels and backup batteries, with a large array of sensors, compute module for onboard data processing, and satellite modem for data reporting. The drifter architecture is flexible and can be customized for a specific purpose. The complete data record is stored locally and processed by the onboard compute module, which runs anomaly detection algorithms that detect nearby activity. Anomalous events, as well as baseline environmental data, are reported to a cloud database using satellite short burst data transmission. Though each independent drifter is a powerful sensing tool, the low unit cost permits large scale deployment. To date thousands of drifters have been deployed over vast areas of the ocean and are reporting data to a remote database where cloud-based analytics algorithms develop global situational awareness and update local edge algorithms on the drifters based on learnings across the full network.**

Keywords*— lagrangian drifter, edge processing, anomaly detection, environmental data, cloud-edge algorithms, data reduction, ocean circulation, maritime domain awareness, vessel tracking

I. INTRODUCTION

The oceans of the world are vast, covering 71% of the total surface. At the same time, however, they remain largely inaccessible with marine vessels cutting narrow tracks mostly along well trafficked shipping paths and the waters near shore. Manned research vessels and autonomous underwater vehicles can purposely explore specific ocean area, but are high in cost. Satellites have advanced measurement tools to measure ocean properties, but they can only cover a limited area at any point in time and are necessarily limited by long-distance sampling. Global sensing programs like Argo[1] and the Global Drifter Program[2] have a current deployment of 1,000-4,000 instruments, but they are distributed globally with an average separation of 300-500 km. Despite all these sensing systems, wide areas of the ocean remain un-instrumented and therefore key environmental and activity measurements are sparse.

As a performer on DARPA's Ocean of Things program[3], we have developed a compact yet highly capable surface drifter to fill this gap. The low price point, made possible by leveraging advances in consumer-grade processing and sensors developed for the mobile phone and automotive industries, and high persistence (> 1 year at most latitudes) make it feasible to deploy large numbers of drifters in select regions to generate high spatio-temporal resolution observations documenting both environmental parameters and human activity across a region of the ocean. Each drifter has a collection of standard low-cost sensors and a single high-fidelity sensor that varies for different models. When these differently equipped drifters are deployed together, the fleet benefits from heterogeneous sensing modalities with cloud analysis aggregating data from the different sensors across the deployed fleet.

Keeping the cost of the drifter low while still generating high-fidelity data, presented several challenges. The chassis needed to be inexpensive to produce and assemble while remaining robust to the ocean surface environment for the intended deployment lifetime of greater than one year. Drifter size constraints, dictated by the desire for high volume deployments, limited the total area of the solar panel on each drifter. The resulting power limitation required highly efficient power and electrical system design as well as judicious regulation of onboard sensing and algorithm uptime to adapt to varying solar conditions. Finally, the difference between the large amount of data collected locally via the onboard sensors

This research was supported with funding from the Defense Advanced Research Project Agency – Strategic Technology Office (DARPA STO)

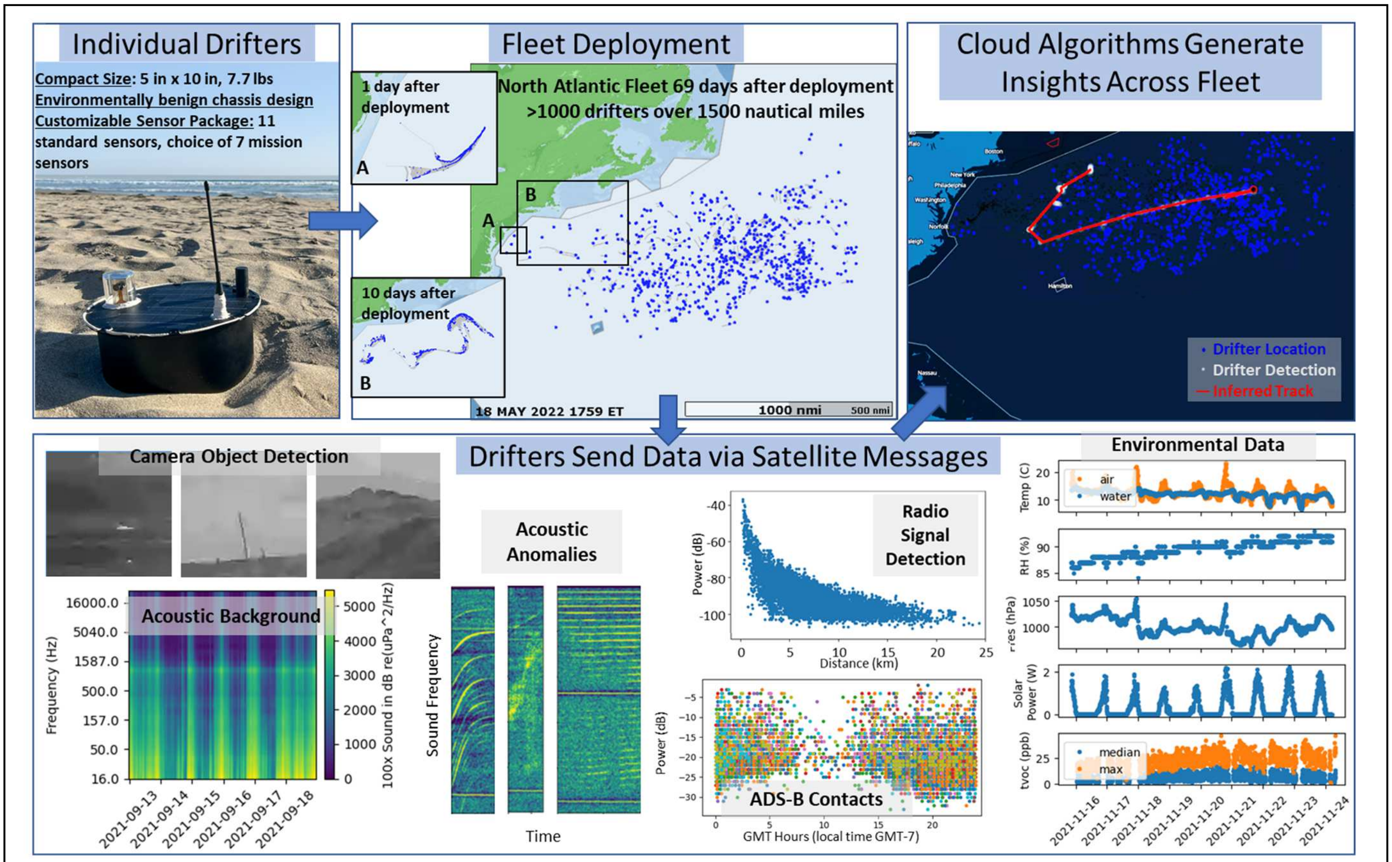

Fig. 1. Overview of operational concept and sample data produced by PEARL drifters.

compared with the small amount of bandwidth available via the satellite communication channel required development of specialized algorithms for ensuring efficient exfiltration of the most valuable information. How each of these challenges was solved is detailed in the following sections.

II. Comparable Technologies

Over the past half century instrumented freely drifting devices have been deployed in the ocean in order to increase scientific understanding, chiefly as part of the government funded Argo and Global Drifter Programs. Organizations including NOAA and Argo maintain databases and satellite communications for these collections of devices[2][4]. These devices are classified as either drifters (i.e., passively drift with ocean currents near the surface of the water) or floats (i.e., able to actuate, e.g., change depth or hydraulic profile). The size of the device ranges from ~1 meter to ~10 meters, depending on the design. Both drifters and floats use the underlying currents and wave action to drive motion. At a minimum, drifters measure their local position as a function of time and their movements map ocean circulation patterns. Additionally, drifters can be outfitted with sensors to measure physical parameters of the ocean, e.g., mixed layer currents, sea surface temperature, atmospheric pressure, winds and salinity[2]. Drifters often deploy drogues to use currents to drive the motion and sample ocean parameters without being affected by surface wind activity. The Global Drifter program currently has over 1200 active drifters, which cost $1,700 each, covering 84% of the 5°x5° lat/lon (556 km) bins across the world's oceans. Floats, on the other hand, can be much more complex: they are designed to change their buoyancy in order to profiles the ocean depths. The Argo program has just under 3800 floats currently deployed, which cost $20,000 each. The floats are approximately distributed in 3°x3° lat/lon bins (333 km) and collect position temperature and salinity measurements while they cycle to depths of 2,000 m every 10 days.

A few commercial entities have also entered ocean instrumentation space with deployment footprints starting to approach the government programs discussed above. Notably, Sofar Ocean Technologies has developed a drifter, Spotter, and has sold or deployed over 2000 of them across multiple oceans. These drifters retail for $4,900 and report ocean weather conditions including: wave height, wave direction, wave period, float position, solar voltage, humidity, surface water temperature and battery voltage. Sofar also sells access to a real time ocean weather online dashboard which aggregates the data reported from its float deployment[5].

Pacific Gyre sells a variety of drifters for different ocean current monitoring applications. Most of their drifters only measure position and users can choose various drogue shapes for deployments in different ocean environments (shallow reefs, shallow coastal waters, deep open ocean). One of their drifter products, the CARTHE drifter, is a commercialized version of the custom-made drifters deployed in the Grand LAgrangian Deployment (GLAD) experiment. GLAD was the largest-scale experiment of its kind where 300 drifters were deployed in August 2012 to probe multi-scale ocean flows in the Gulf of Mexico near the Deepwater Horizon spill site[6]. Unlike the global or ocean-scale programs described above, the GLAD experiment deployed the drifters much closer together and had them report position at a much higher frequency, creating a highly detailed dataset to study the complex surface ocean currents that transport pollutants[7].

In all cases of described above, the parameters measured have been either position only or physical (e.g., temperature, salinity, etc.), owing to the challenge of measuring other types of data in a similarly autonomous manner. For example, acoustic monitoring is much more intensive. Sonobuoys are energy intensive and have a lifetime of hours instead of hundreds of days. Autonomous, passive acoustic monitors are mostly deployed to the deep ocean, on the sea floor itself, to maximize their range and to monitor the soundscape without interference from surface noises. They produce large datasets and must be physically recovered. Hydrophone arrays can be towed behind ships.

Our drifters contain the basic position and physical sensors (temperature, pressure, relative humidity, and solar power generation) but also capabilities to measure more complex or advanced signals (directional wave spectra, RF communication, automated image collection, object identification and processing, magnetic signals from vessels, acoustic signals, volatile organic compounds from ship emissions, and local algae and dinoflagellates populations), with the goal of gathering and using more data for additional monitoring and analytics. Our drifters are adaptable to additional missions and are designed with this flexibility in mind. Moreover, the low cost of the drifter, high production volume and high persistence (> 1 year at most latitudes) make it feasible to deploy thousands of drifters to generate a fine-grained, both spatially and temporally, dataset detailing environmental parameters and human activity across a vast region of the ocean. This data can be used by oceanographers and other environmental scientists to improve ocean models and for tracking human generated marine activity for increased operational intelligence.

III. Drifter Design

A. Mechanical

The drifter is a cylinder, 21.8 cm in diameter, 11.5 cm tall, with a mass of 3.5 kg. The buoyancy and arrangement of internal components (including a small steel ballast) were designed to maintain a low freeboard of 2.1 cm and to keep the center of mass below the center of buoyancy, allowing the drifter to self-right when it gets flipped by waves. To minimize environmental impact, the drifter design maximizes the use of biologically inert materials; >90% of the mass is metal (aluminum and steel), and the use of plastics is minimized. The drifter chassis is made of aluminum with an anodization layer that helps reduce corrosion. The top surface is a solar panel with protrusion holes for antennas (Iridium Satellite, VHF radio, and an optional UHF radio antenna), a service window to view status LEDs, a gore vent to allow atmosphere exchange, and, in one float variation, a glass housing for back-to-back cameras. All protrusions are located near the perimeter to maximize the area available to solar cells. A small, 1.5 cm in width, rim extends around the top edge of the drifter to provide a sealing surface between the aluminum chassis and the solar panel as well as a mechanical support for the internal electronics. All holes in the chassis are either sealed with a compressed o-ring or with UV-resistant marine sealant. To ensure the sealing methods would survive in the open ocean, the complete drifter design was tested for IP68 compliance, submerged at depth of 1.5m for over 30 minutes without water incursion (test was halted after 12 hours).

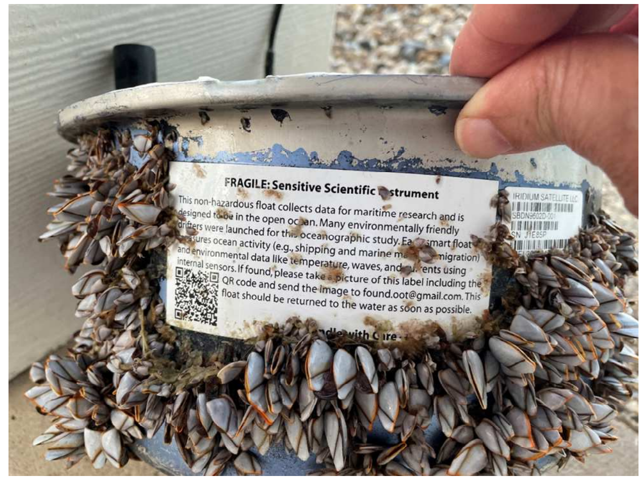

Fig. 2. Drifter recovered after several months of deployment in the Gulf of Mexico showing wear of anodized aluminum chassis, particularly above the water line, and accumulation of biological growth.

One unavoidable source of plastic in the drifter is the electronics components including printed circuit boards, integrated circuit packages and the insulation materials for wires and antennas. On average the plastic content in typical printed circuit board assemblies (PCBAs) is 36.2% by weight, determined by combustion analysis[8]. The remaining content consists of metals, ceramics and glass fibers. The total weight of all electronic components (PCBAs, solar panel, antennas, and cabling) is 487 g. Assuming 36.2% of that mass is plastic results in just 176 g of plastic per drifter. Consequently, plastics make up just 5% of our drifter by mass.

During deployment, the anodization of the aluminum chassis does appear to wear off over time, particularly above the waterline. Biological growth also tends to accumulate below the waterline but not sufficient to pull a drifter underwater (Fig. 2).

B. Drifter Operation

The basic drifter operation has three states: sleep state, measurement state, and analytics state. The system starts in the measurement state, collecting data from the onboard sensors for one minute, then the system runs in a low power sleep state for four minutes. Every 4th wake/sleep cycle the system enters the analytics state, utilizing a higher power single-board computer to run more complex data analytics, anomaly detection, and compression algorithms. At the end of the analytics cycles the systems packages the processed data and sends an Iridium message. The durations and duty cycles for each state are configurable and can vary based on sensor configuration.

C. Electronics

The onboard electronics system is centered around a pair of processing units, a low-power ARM microprocessor and a single-board computer. The microprocessor manages the basic operation of the drifter such as power management and wake/sleep cycle timing as well as data collection from the majority of the attached sensors. The single-board computer runs data processing algorithms and handles communication over the Iridium satellite connection. Data is stored on a SD card that is shared between the two processing units. A WiFi interface can be used to download raw sensor data from the SD card or provide software updates to either of the processing units.

D. Power

The main power supply for the drifters is a 20.8 Ah rechargeable Li-ion battery pack. The battery is charged by a 4.2 W solar panel that covers most of the top surface of the drifter and a 54 Ah alkaline back-up battery that can provide emergency power to the drifter during extended periods with low solar irradiance. We measured the efficiency of the solar panel and solar charging circuit to be 18% by comparing the power produced by the solar charging circuit and with the incident solar irradiance, which was measured separately with a pyranometer (Fig. 3 left, middle). The total active area of our solar panel is 0.02 m^2. The efficiency and area are used to estimate the total power generated by the solar charging system if the float were deployed at a variety of latitudes. Fig. 3 (right) shows the average daily power generated by the float using simulated solar data without cloud cover and historical actual solar data with realistic cloud cover represented by 1-day and 7-day rolling averages[9][10]. The simulated data and the 7-day rolling average are computed at different latitudes.

The typical power consumption of the system varies from 380 mW (VOC hardware configuration) to 630 mW (SDR hardware configuration); though this is highly tunable and depends on the drifter's parameter configuration (described below). When the battery state of the rechargeable battery drops below 30% state of charge, the float enters a low power mode. This mode reduces the sensor measurement uptime and lengthens the time between runs of the higher order data process

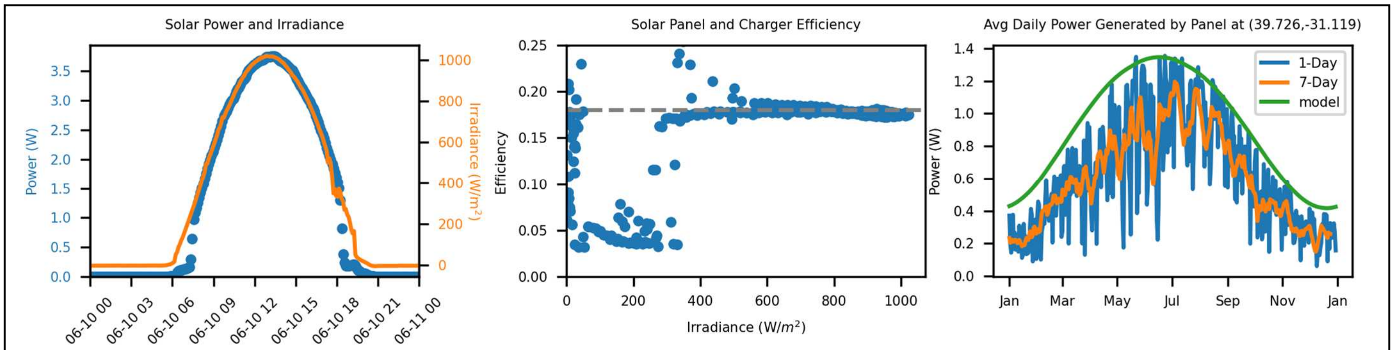

Fig. 3: [Left] (blue) Power generated by our solar charging circuit over a 24-hour period on June 10, 2021. (orange) Solar irradiance data measured by a commercial weather station with an integrated pyranometer. The drifter was shaded in early morning and late evening creating a deviation between the power generated by the panel and measured irradiance. [Middle] Efficiency of the solar panel. Dashed line at 18% is a guide to eye and represents the efficiency value used in future power calculations. [Right] Average daily power generated by the float using simulated solar data without cloud cover (green) and historical actual solar data with realistic cloud cover represented by 1-day (blue) and 7-day (orange) rolling averages.

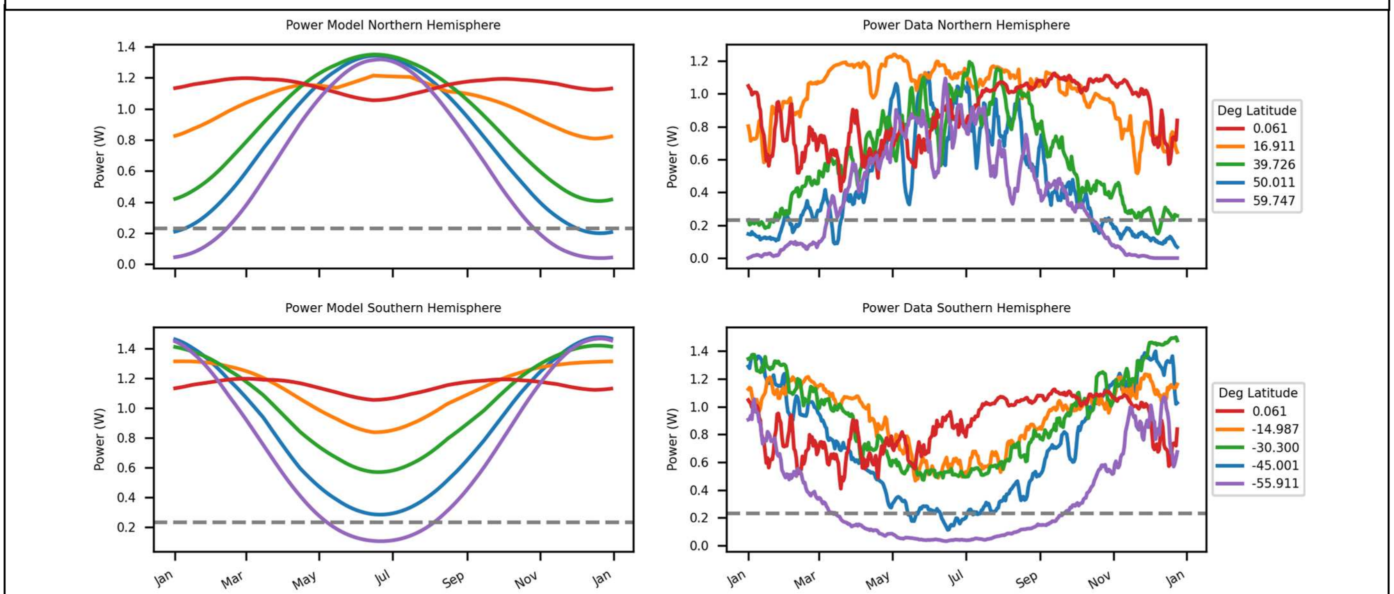

Fig. 4: [Left] Simulated daily solar power generated at different latitudes. [Right] 7-day rolling average historical data at different latitudes. The gray dashed line in all four plots represents the average power draw of the float in low-power mode.

algorithms and satellite data transmissions. In this mode, the average power consumption drops to 231 mW. This power draw is represented by the gray dashed line in Fig. 4. This shows that the drifters can survive indefinitely at lower latitudes. The primary back-up battery can provide 45 days of power in zero light conditions, which lengthens the total lifetime but would still not be sufficient to allow the drifter to survive the darkness of winter at latitudes of 60 deg or greater.

The power draw of the system depends on the deployed hardware configuration and a set of parameters that control reporting frequency, measurement configurations, sensor uptime and downstream processing algorithms. The parameter set is configurable over the air to control drifter operation and power consumption during deployment. During operation, the low power ARM microcontroller manages the power state of the various subsystems. During sampling periods, the low power microcontroller operates at full capacity and enables the power rails supplying the various sensor subsystems. In between sampling periods, the low power microcontroller goes into a sleep mode, only maintaining power to the onboard RAM and internal real-time clock with its alarm circuitry that reawakes the microcontroller at a preset time. If wave spectra measurements are enabled, the microcontroller will periodically wake for short periods of time between the main sampling cycles to record data from the onboard IMU. Every few sampling cycles, the microcontroller will power on the single-board computer to run the data processing algorithms and transmit data over Iridium. As the single-board computer is the largest power draw on the system, it is powered down outside of these data processing cycles.

E. Sensors

Each drifter comes equipped with a standard set of sensors that are primarily focused on environmental sensing (see Table 1) and a differentiated mission sensor. The mission sensor typically has a much higher cost, higher energy consumption, and/or higher signal processing requirements. The options for the mission sensors include: a pair of horizon facing cameras with a combined near-360 deg field of view, a software defined radio (SDR), a microphone, a hydrophone, a high-sensitivity magnetometer, a volatile organic compound sensor (VOC), or a bio-optical sensor designed to detect Chlorophyll-a fluorescence and bioluminescence.

TABLE I. STANDARD SENSORS AVAILABLE ON ALL DRIFTER CONFIGURATIONS AND THE CORRESPONDING DATA OUTPUT

Standard Sensor	***Data Product***	***Type***
Environmental	Air temperature	Time Series
Environmental	Relative Humidity	Time Series
Environmental	Air Pressure	Time Series
IMU	Wave Spectra	Time Series
GPS	Drifter Location	Time Series
GPS	Drifter Velocity	Time Series
IMU	Drifter Orientation	Time Series
IMU	Acceleration	Anomaly
IMU	Magnetic	Anomaly
IMU	Gyroscope	Anomaly
AIS Receiver	AIS contact	Contact Report
Thermistor	Water Temperature	Time Series
Water Ingress	Water Intrusion	Anomaly
Solar Panel	Solar Irradiance	Time Series

To evaluate sensor performance during deployment, we compared data generated from drifters while in close proximity, within 50 km, to a National Data Buoy Center operated buoy with a SCOOP payload located in the western Gulf of Mexico[11]. Fig. 5 and Fig. 6 show the results of this comparison over a two-week period in later half of February 2022. Generally, the data align well with a notable excursion towards the end of that period where the nearby drifters were in a small eddy. Differences in water temperature may be due to the rather shallow sampling depth (base of drifter chassis, ~0.1 m) compared to that of the reference buoy (3 m). The air temperature data is highly influenced by solar irradiation during the daytime and by the water temperature during the nighttime, mainly due to the location of the environmental sensor, which also measures relative humidity and air pressure, under the solar panel and the overall compact design of the drifter. The environmental sensor uses a set of factory-calibrated compensation algorithms to correct the relative humidity and air pressure outputs based on the air temperature measurement. These compensation algorithms, when coupled with the solar heating seen by the air temperature measurement, may be responsible for some of the small spikes in air pressure that do

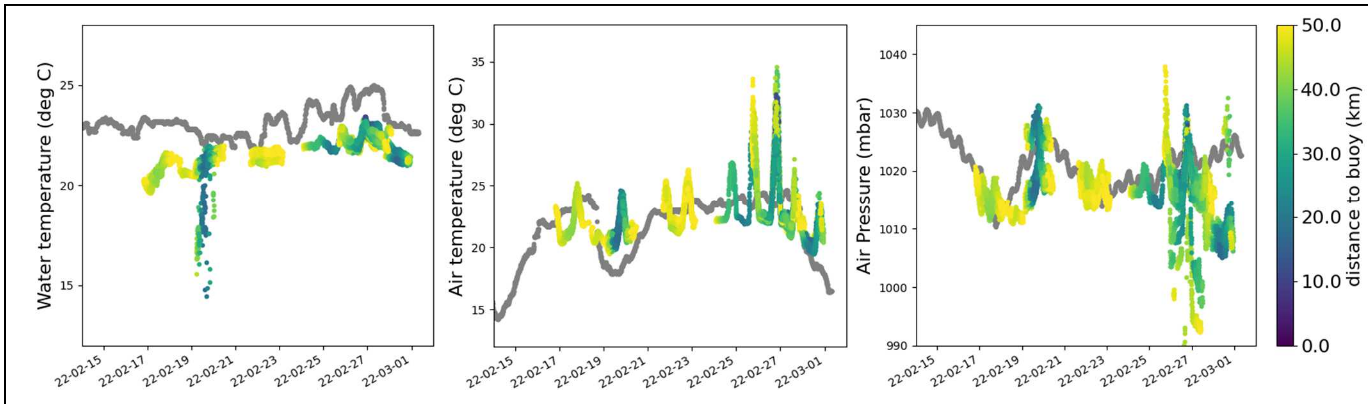

Fig. 5: Comparison of environmental data measured by NDBC data buoy (gray dots) versus data from nearby drifters over a two-week period in February 2022 in the Gulf of Mexico. Drifter measurements are colored based on proximity to buoy.

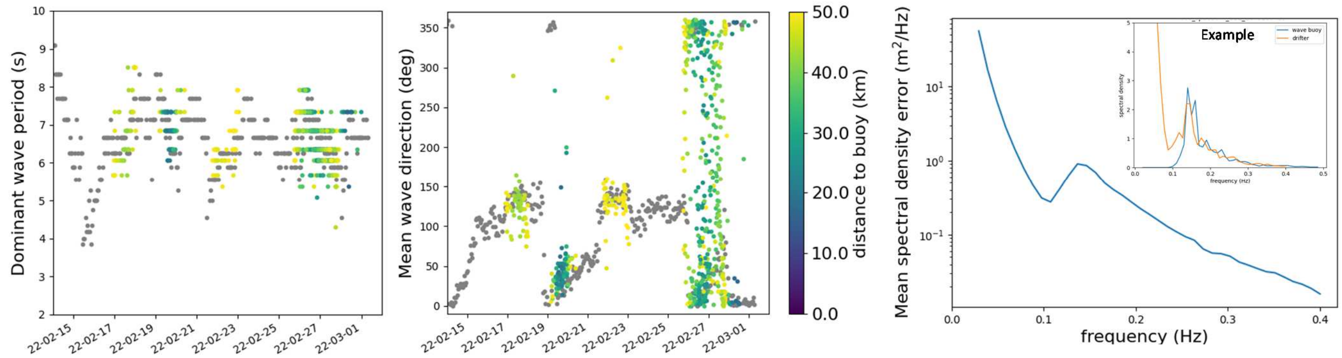

Fig. 6: [Left, Center] Comparison of wave data measured by NDBC data buoy (gray dots) versus data from nearby drifters over a two-week period in February 2022 in the Gulf of Mexico. Drifter measurements are colored based on proximity to buoy. Right, Average measurement error of wave spectral density reported by drifters compared with a nearby (within 50 km) NDBC data buoy. Inset shows comparison of single drifter report (orange) versus that from the reference data buoy (blue).

not correlate well with those observed by the reference buoy. With some additional processing however, this sensor could likely provide an accurate readout of solar irradiation. Also, the wave spectral density data, derived from the onboard IMU using the methods described in [12][13], show a large error in lower frequency range particularly below 0.1 Hz (Fig. 6). While further processing is required to improve the overall accuracy of the environmental data, as is the data already shows interesting events that would otherwise go unobserved, like the sharp dip in water temperature shown in Fig. 5.

IV. Information Management

During operation, a low-power ARM-based microcontroller manages data collection from the onboard sensors, which is then stored in a local data store (SD card). When sufficient data is collected, a single board computer is powered on to process the stored data through a set of anomaly detection algorithms. Identified anomaly events, associated meta data, and background summaries are packed into Iridium short-burst data messages and sent via satellite to a cloud server for further analysis. Messages from the drifters are processed at the fleet level to track signals of interest across vast areas of the ocean (see Fig. 1, top right for example).

A. Compression

Data from each sensor is processed by an onboard algorithm that compresses the information content sufficiently to transmit the results over the limited bandwidth available via the Iridium satellite connection of 340 bytes per message. While the anomaly detection and compression algorithms used for higher data rate sensors are outside the scope of this paper, compression ratios of greater than 10^4:1 are typical and Fig. 7 and Fig. 8 give some idea of what is possible. The methods used for low data rate sensors, however, are straightforward.

For low data output rate sensors, such as the standard environmental sensors, the data are sent up as short duration time series using a simple differential encoding scheme. To implement this scheme, the initial value is encoded with a high precision bit-depth tuned to cover the expected observational dynamic range of the sensor at the desired resolution. Subsequent values are encoded as the difference from the last value with the bit-depth tuned to match the expected rate of change for the particular data type over the time between samples. During encoding, in the rare case that a differential value exceeds the value range covered by the assigned bit-depth, the maximum allowable difference is encoded and the remainder is carried over to the next value. This carry-over technique preserves the encoding accuracy of values later in the time series if a larger than expected change in value is recorded, but acts as a smoothing filter on the data when those events occur. In applications where this smoothing is undesirable, the encoding bit-depth for the differential values can be increased using the parameter updating system (see Section V) at the cost of more bandwidth devoted to that particular data type.

B. Data Priority Queue and Modular Encoding

To ensure the most valuable information gets exfiltrated during each transmission session, the drifters use a modular encoding scheme for outgoing messages that, besides a short message header, allows for any combination of data packets to be included in a given message, instead of predefining sections of the message to specific data outputs. To manage which data packets to include a message, a priority value and, optionally, a minimum reporting frequency value are assigned to each data

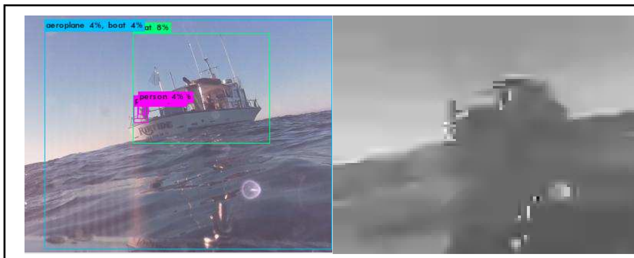

Fig. 7: [Left] Image object detection using a lightweight neural network, condensed to run on drifter hardware. [Right] Ship image compressed to within a single Iridium SBD message (<300 bytes).

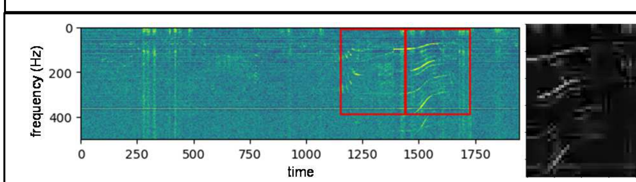

Fig. 8: [Left] Acoustic spectrogram overlaid with example anomaly detection of passing aircraft (red boxes) from a microphone equipped drifter. [Right] Transmitted spectrogram chip compressed to fit within a single Iridium SBD message (<300 bytes).

type. During processing, each new data packet generated from recent observations is placed onto a data queue. After processing is complete, the message packing algorithm follows a multi-step method to fill the available message space. First, the algorithm adds any data packet available on the queue whose type has a minimum reporting frequency and has a time since last data packet transmission that exceeds that minimum reporting frequency. After that step, any space remaining is filled using an optimization algorithm that chooses which data packets available on the queue maximally use the space based on the priority value assigned to their data types. If message transmission fails, any data packets included in the failed message are returned to the data queue to be made available again during the next transmission cycle. Both the priority and minimum reporting frequency values for each data type can be adjusted by a remote operator after deployment if a different set of data priorities is desired.

C. *Message Decoding*

Encoded Iridium messages received from the fleet are decoded using a restful API running on a cloud server that returns the message in JavaScript Object Notation (JSON) format, broken down into individual data packets for storage in a cloud database. This decoding API is integrated with the Parameter Updating API used for Command and Control (C2) that enables the decoding API to adjust to variations in encoding from one drifter to the next (e.g., due to variations in software version or compression bit-depth settings). The JSON output can be easily reformatted for ingestion into other data repositories. For example, as part of the OoT program, data generated by our drifters has already been integrated with the U.S. Integrated Ocean Observing System (IOOS) and made publicly available online[14][15]. See [16] for more information on the available data.

V. Configurability

After deployment, drifter operation can be modified using an extensive set of adjustable parameters (>2600) and commands deliverable via Iridium. These parameters enable control of sensor sampling, onboard algorithm tuning, and data exfiltration prioritization. This functionality enables both configurable operation for a particular mission or real-time adjustment of sampling uptime in response to anticipated event or object of interest. A cloud-based API provides access to this command set along with a queryable database that keeps track of the current configuration for each drifter in a deployed fleet. Sets of parameter modifications that correspond to a particular mission objective can be bundled together and applied in bulk using built-in capabilities of the Parameter Updating API.

VI. Deployment and At Sea Performance

A. *Packing*

Our drifters are designed for simple, large volume deployment from a ship in transit. The drifters are delivered in palletized packaging with 64 drifters per 42”x42” pallet yielding 1280 drifters per 20 ft shipping container (Fig. 9). Pallet packaging is UN38.3 certified for land and ocean shipment. The packaging design uses only cardboard to minimize environmental impact if any packaging material falls overboard during deployment. To simplify deployment, the top panel of each layer of the packed drifters can be exposed at once for activation while the drifters remain held in place.

Fig. 9: [Left] Arrangement of drifters in pallet-level packaging showing ease of access for device startup onboard deployment vessel. [Right] High density arrangement of palletized drifters (1280 units) in 20-foot shipping container for large scale deployment.

B. *Deployment Method and Drift Patterns*

During deployment, a single user can power-on each system by applying a magnet to a specific location on the chassis and observe the LED startup sequence through a service window to indicate successful initialization. After powering on, drifters can be tossed directly into the water from a height of up to ten feet. Drifters deployed in a small area, typically covered by a ship transit over a single day, quickly expand to cover large areas. Fig. 10 shows how our deployment in the northern Atlantic Ocean scaled to >10^6 km^2 after several weeks with spacing between drifter steadily increasing in a similar manner. Fig. 11 shows how the drifters tend to follow mesoscale ocean features as they evolve over time, providing a high-resolution view of ocean surface currents.

VII. Applications

Designed to be a low-cost, highly adaptable sensor platform, our drifters have the potential to support multiple applications. The ability of ocean models to provide accurate forecasts is directly related to the spatio-temporal resolution of the observations used to drive them[17] and our drifters are

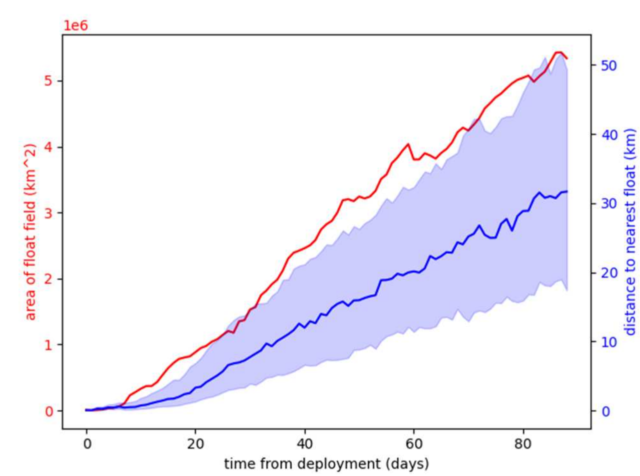

Fig. 10: Evolution of drifter field over time from a deployment in the Gulf Stream in the North Atlantic Ocean. Area of the field (red line) determined by the convex hull encapsulating all deployed drifters. The distribution of drifters within the field is depicted as the distances to the nearest neighbor for each float with the blue line showing the median and the filled blue region spanning from the 25th to the 75th percentile.

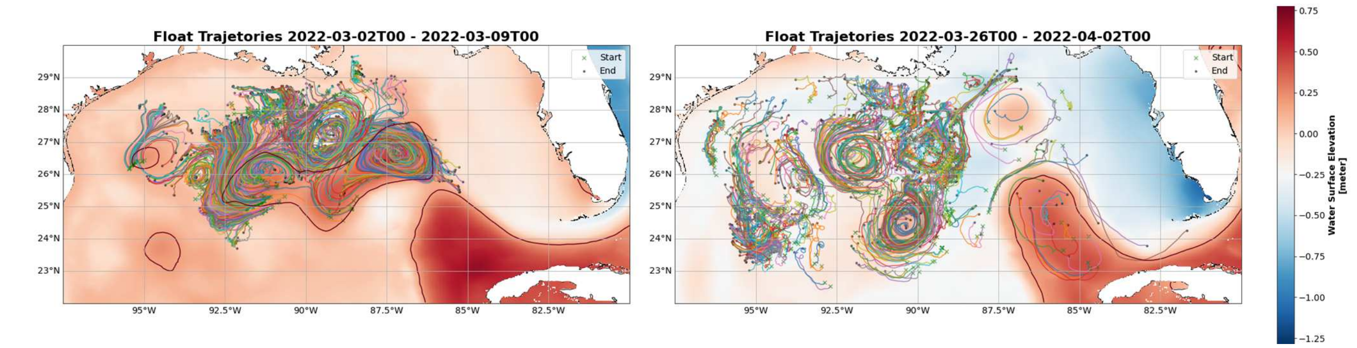

Fig. 11: 7-day drifter trajectories (colored lines) in the Gulf of Mexico over two time periods in March 2022 overlaying sea surface height output from the Global Navy Coastal Ocean Model (NCOM). The drifter trajectories clearly capture the evolution of eddies generated by the loop current.

unparalleled in both the spatio-temporal resolution and the variety of observations that they can provide. Improvements in ocean forecasting provides benefits to a variety of stakeholders including coastal communities as well as the aquaculture, oil and gas, and shipping industries. For defense and security applications, data from the onboard AIS receiver and the variety of optional mission sensors provide multi-modal maritime domain awareness that can be used to track activities of interest like illegal fishing and smuggling.

VIII. CONCLUSION

This paper presents an overview of our drifter sensor platform along with a brief discussion of the onboard environmental sensors. Currently, two large-scale deployments, funded by the DARPA Ocean of Things program, are underway and we are looking forward to seeing what can be learned with the unprecedented data set being collected. Looking beyond these deployments, we are exploring the integration of new sensors, particularly in the biogeochemical category, to support a broader set of applications. We are also interested in exploring the C2 capabilities to open new methods of operation such as turning the drifter field into a single large-scale sensor through coordinated observations. We are just starting to explore the capabilities our drifter platform has to offer, and we are excited to see what insight might be gained in the future.

ACKNOWLEDGMENT

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions, and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. Distribution Statement "A" (Approved for Public Release, Distribution Unlimited).

The authors acknowledge Keith McLaughlin, John Wolf, Andrew Wilson, Alexander Sell, and Louise Xu from Leidos for supporting magnetometer and acoustic sensor development, and Dimitri Deheyn, Noah Martin, and Evelien De Meulenaere from Scripps Institute of Oceanography for supporting at sea testing of early prototypes and qualification of the bio-optical sensor.

REFERENCES

[1] D. Roemmich *et al.*, "The Argo Program: Observing the global ocean with profiling floats," *Oceanography*, vol. 22, no. SPL.ISS. 2, pp. 34–43, 2009.

[2] "Global Drifter Program." [Online]. Available: https://www.aoml.noaa.gov/phod/gdp/.

[3] J. Waterston, J. Rhea, S. Peterson, L. Bolick, J. Ayers, and J. Ellen, "Ocean of things: Affordable maritime sensors with scalable analysis," *Ocean. 2019 - Marseille, Ocean. Marseille 2019*, vol. 2019-June, 2019.

[4] "Argo." [Online]. Available: http://www.argo.ucsd.edu/index.html.

[5] "Sofar Spotter." [Online]. Available: https://www.sofarocean.com/products/spotter.

[6] A. C. Poje, T. M. Özgökmen, B. L. Lipphardt, B. K. Haus, E. H. Ryan, and A. C. Haza, "Submesoscale dispersion in the vicinity of the Deepwater Horizon spill," vol. 111, no. 35, pp. 12693–12698, 2014.

[7] G. Novelli *et al.*, "A biodegradable surface drifter for ocean sampling on a massive scale," *J. Atmos. Ocean. Technol.*, vol. 34, no. 11, pp. 2509–2532, 2017.

[8] C. Eswaraiah, T. Kavitha, S. Vidyasagar, and S. S. Narayanan, "Classification of metals and plastics from printed circuit boards (PCB) using air classifier," *Chem. Eng. Process. Process Intensif.*, vol. 47, no. 4, pp. 565–576, 2008.

[9] W. F. Holmgren, C. W. Hansen, and M. A. Mikofski, "pvlib python : a python package for modeling solar energy systems," vol. 2014, no. 2014, pp. 2015–2017, 2018.

[10] T. Huld and R. Mu, "A new solar radiation database for estimating PV performance in Europe and Africa," vol. 86, pp. 1803–1815, 2012.

[11] "NDBC Station 42002." [Online]. Available: https://www.ndbc.noaa.gov/station_page.php?station=42002.

[12] M. D. Earle, "Nondirectional and Directional Wave Data Analysis Procedures. NDBC Technical Document 96-01," vol. 002, no. January, p. 37, 1996.

[13] D. E. Barrick, B. J. Lipa, and K. E. Steele, "Comments on 'Theory and Application of Calibration Techniques for an NDBC Directional Wave Measurements Buoy': Nonlinear Effects," *IEEE J. Ocean. Eng.*, vol. 14, no. 3, pp. 268–272, 1989.

[14] "DARPA Ocean of Things (OoT) - Gulf of Mexico - February 2022 - Environmental Data." [Online]. Available: https://data.ioos.us/dataset/darpa-ocean-of-things-oot-gulf-of-mexico-february-2022-environmental-data.

[15] "DARPA Ocean of Things (OoT) - Atlantic / Gulf Stream - March 2022 - Environmental Data." [Online]. Available: https://erddap.secoora.org/erddap/tabledap/oot-lant-deployed-environmental.html.

[16] J. Ellen, "In-situ data from two 1500+ sub-mesoscale near-Lagrangian float deployments across 14 different sensor types," in *OCEANS 2022 - Hampton Roads*, 2022.

[17] G. A. Jacobs, J. M. D'Addezio, B. Bartels, and P. L. Spence, "Constrained scales in ocean forecasting," *Adv. Sp. Res.*, vol. 68, no. 2, pp. 746–761, 2021.