Autonomous, Localization-Free Underwater Data Muling using Acoustic and Optical Communication

Marek Doniec, Iulian Topor, Mandar Chitre, Daniela Rus

Abstract We present a fully autonomous data muling system consisting of hardware and algorithms. The system allows a robot to autonomously find a sensor node and use high bandwidth, short range optical communication to download 1.2 MB of data from the sensor node and then transport the data back to a base station. The hardware of the system consists of an autonomous underwater vehicle (AUV) paired with an underwater sensor node. The robot and the sensor node use two modes of communication - acousic for long range communication and and optical for high bandwidth communication. No positioning system is required. Acoustic ranging is used between the sensor node and the AUV. The AUV uses the ranging information to find the sensor node by means of either stochastic gradient descent, or a particle filter. Once it comes close enough to the sensor node where it can use the optical channel it switches to position keeping by means of stochastic gradient descent on the signal quality of the optical link. During this time the optical link is used to download data. Fountain codes are used for data transfer to maximize throughput while minimizing protocol requirements. The system is evaluated in three separate experiments using our Autonomous Modular Optical Underwater Robot (AMOUR), a PANDA sensor node, the UNET acoustic modem, and the AquaOptical modem. In the first experiment AMOUR uses acoustic gradient descent to find the PANDA node starting from a distance of at least 25m and then switches to optical position keeping during which it downloads a 1.2 MB large file. This experiment is completed 10 times successfully. In the second experiment AMOUR is manually steered above the PANDA node and then autonomously maintains position using the quality of the optical link as a measurement. This experiment is performed two times for 10 minutes. The final experiment does not make use of the optical modems and evaluate the performance of the particle filter in finding the PANDA node. This experiment is performed 5 times successfully.

Marek Doniec, Daniela Rus
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology - Cambridge, Massachusetts
Iulian Topor, Mandar Chitre
ARL, Tropical Marine Science Institute
National University of Singapore - Singapore
e-mail: doniec@mit.edu, rus@csail.mit.edu, iulian@arl.nus.edu.sg, mandar@nus.edu.sg

1 Introduction

Our goal is to develop technologies that enable users to interact with ocean observatories. In an ocean observatory robots and and in-situ sensors collect information about the underwater environment and deliver this information to remote users, much like a web-cam delivers remote data to users on the ground. In this paper we focus on developing effective technologies for wireless data transmission underwater. When the amount of data from an ocean observatory is large (for example in the case of image feeds), low-bandwidth acoustic communication is not adequate. We instead propose using optical data muling with a robot equipped with an optical modem that can retrieve data fast from underwater nodes with line-of-site connection to the robot. An important problem is locating the underwater sensor node. When distances between the robot and the nodes are large, and their locations are unknown, positioning the data muling robot within optical communication range is challenging. In this paper we present a solution to autonomous data muling underwater, where the node's location is unknown. The algorithm has three phases. In the first phase, acoustic communication is used to bring the data muling robot within some close range of the desired sensor where it can detect the optical signal. In the second phase, the robot does a local search using the optical signal strength to precisely locate the sensor and position itself within communication range. In the third phase the robot uses optical communication to collect the data from the sensor. In practice, phase two and three overlap once the signal strength becomes strong enough to transmit data.

Previous work has looked at the theoretical performance of data muling [10] and the optimization of the path taken between nodes [7]. In both cases the locations of the nodes are assumed known. Data muling with an underwater robot has been previously shown in [5]. The nodes were found using a spiral search that looked for a valid optical signal. A method for homing to a single beacon using acoustic range measurements based on an Extended Kalman Filter with a fixed robot maneouver for initialization is presented and evaluated in simulation in [13]. In [9] the authors present an Extended Kalman Filter approach to localizing a moving vehicle using range-only measurements to a group of beacons. They use particle filters to initialize the beacons location. In [1] a high-frequency acoustic network is suggested, that offers range and bandwidth performance between conventional acoustic and optical rates.

We implemented and experimentally evaluated the data muling system described in this paper. This work uses a new version of the Pop-up Ambient Noise Data Acquisition sensor node called UNET-PANDA, which is presented, along with the acoustic modem used, in [2]. For simplicity UNET-PANDA is referred to as PANDA in the remainder of the paper. The optical modem has been described in [3]. The Autonomous Modular Optical Underwater Robot (AMOUR) was presented in [4]. Our implementation of the data muling system was repeatably able to acoustically locate the sensor node from distances of 25m and 100m and to download a 1.2 MB data file optically once the node was found.

2 Problem Statement

We consider a sensor node that is deployed at a fixed location on the seafloor. We assume that the sensor node is equipped with an acoustic modem and an optical modem. We use the acoustic modem for low data rate ($\leq$ 1 Kbps) and long range ($\geq$ 100 m) communications. We use the optical modem for high data rate ($\geq$ 1 Mbps) and short range ($\leq$ 100 m) communications. We do not require a precise external positioning system but we assume that a coarse location estimate of the node exists. By coarse we mean that the margin of error for this position estimate is within the acoustic communication range. This is usually on the order of hundreds of meters to a few kilometers, though acoustic communication ranges of over 100 km are possible [11]. Examples in which such a situation can arise are (1) when a node is deployed in deep waters from a boat and drifts before it finally reaches the ocean floor; (2) when a node is deployed by an autonomous underwater vehicle using dead reckoning the placement can have a large error; (3) when a node is not rigidly moored and its position changes with time because of water currents.

Further, we assume that an autonomous hovering underwater vehicle (AUV) is equipped with identical acoustic and optical modems and capable of communicating through these with the sensor node when in range. Hovering enables the vehicle to hold its attitude and depth statically and to execute surge (forward / backward) velocity commands.

Our problem statement concerns the case in which the sensor node is collecting data at rates in excess of what can be transmitted using the long distance acoustic channel. Kilfoyle *et al.* show empirically that the product of acoustic communication rate and bandwidth rarely exceeds 40 km-Kbps for state of the art acoustic modems [6]. For a single sensor separated by 5 km from the user this would result in a communication rate of 8 Kbps. If we consider an application that collects ambient acoustic signals or even video our data stream will far exceed the available acoustic channel capacity.

3 Technical Approach

Algorithm 1 Acoustic Stochastic Gradient Descent

$YAW_{robot} \Leftarrow \pi * Random(-1.0...1.0)$
$SPEED_{robot} \Leftarrow 0.25$ Knots forward
$RANGE_{th} \Leftarrow \text{inf}$
while No optical link available **do**
  Receive $RANGE_m$
  if $RANGE_m \geqslant RANGE_{th} + 1m$ **then**
    $YAW_{robot} \Leftarrow YAW_{robot} + \pi + \pi * Random(-0.5...0.5)$
    $RANGE_{th} \Leftarrow RANGE_m$
  end if
  if $RANGE_m < RANGE_{th}$ **then**
    $RANGE_{th} \Leftarrow RANGE_m$
  end if
end while
Begin Optical Gradient Descent

Algorithm 2 Optical Stochastic Gradient Descent

Retain YAW_{robot} from Algorithm I.
Retain $SPEED_{robot}$ from Algorithm I.
$SSI_{th} \Leftarrow 0$
while Optical Link Established **do**
  Wait 0.25 Seconds. Measure SSI_m
  if $SSI_m \leqslant 0.9 * SSI_{th}$ **then**
    $YAW_{robot} \Leftarrow YAW_{robot} + \pi + \pi * Random(-0.5...0.5)$
    $SSI_{th} \Leftarrow SSI_m$
  end if
  if $SSI_m > SSI_{th}$ **then**
    $SSI_{th} \Leftarrow SSI_m$
  end if
end while
Switch back to Acoustic Gradient Descent

We developed a combined acousto-optical communication network capable of large scale data recovery that does not require precise localization of the robot nor the sensor node. The robot uses acoustic communication and ranging to come close to the sensor node. High bandwidth optical channel allows the robot to download the payload data. More specifically, our approach to data muling is as follows:

1. We use acoustic ranging between the robot and the PANDA sensor node. The acoustic modems use a carrier frequency of 27 kHz, transmitting at a bandwidth of 18 kHz with a power level of 180 dB measured at 1m. We use Orthogonal Frequency Division Multiplexing together with a 12/23 Golay code and 1/3 rate convolution code. The acoustic modem on the PANDA transmitts a 18 byte long ranging beacon every 6 seconds that is received by the acoustic modem on the robot and provides it with a range measurement. The robot uses the stochastic gradient descent algorithm shown in Algorithm 1 to travel close to the PANDA.
2. At all times the PANDA is streaming the payload data using the optical modem and random linear rateless erasure codes known as Luby transform (LT) codes [8]. Each optical packet transmitted contained 576 bytes payload data plus 32 bytes of configuration data (i.e. source and destination address, packet size, 32 bit CRC checksum, degree and seed used for the LT-Codes). The test file is a random data file consisting of 2048 blocks of 576 bytes. The LT-Codes require on average an overhead of 3%, so about 2109 packets have to be received by the robot to decode the entire file. Because of the nature of the LT-Codes it does not matter which packets are received.
3. Once the robot is close enough to the sensor node to receive an error-free packet it switches from acoustic gradient descent to maintaining position using the optical gradient descent algorithm described in Algorithm 2. If the optical connection breaks at any point in time we return to step 1.
4. While AMOUR is in optical communication range every packet is used to (a) measure the signal strength and (b) decode the payload data if the CRC matches.
5. The experiment is considered to have completed successfully once the entire 1.2 MB file has been received and decoded by the robot. In a real world scenario the robot would now continue on the approximate location of the next sensor node to begin acoustic gradient descent there.

4 Performance Improvement with a Particle Filter

The stochastic gradient descent approach described in Section 3 has no memory of previous decisions. The only state variables are the current heading and a range threshold used to make the decision whether to keep going straight or to turn. When the algorithm encounters an increasing range it changes the direction of the robot in a random direction at least 90 degrees different from the current direction of travel. This new choice of direction takes into account only the most recent few measurements as reflected in the threshold stored. Because so little information is taken into account, bad choices are made frequently. Further, even when the robot is moving in the right direction a single spurious measurement caused by noise can make it veer of the correct course. The algorithm will recover from this mistake with high probability as the ranges will keep increasing from here on, but this comes at the cost of time and energy. It also causes a large variance in the time that it takes to find the target sensor node.

A more effective algorithm should keep a belief of where the robot is relative to the sensor node and update this belief with every measurement. An Extended Kalman Filter (EKF) delivers such a behavior. It represents the current belief of the robot's location as a mean and covariance. Because of this it needs to be initialized, for example by performing a circular maneuver such as in [13]. Further, because we are representing the robot's state with a multidimensional Gaussian, we cannot represent multimodal distributions, for example when we have a baseline ambiguity because our vehicle has been traveling straight.

In order to represent multimodal distributions we implement a particle filter algorithm, as shown in Algorithm 3. The filter represents the current belief using N particles. The sensor node position is assumed to be fixed at the origin. The filter localizes the robot relative to the sensor node. Each particle stores one guess $[X_k\ Y_k]$ of the robot's location. The algorithm initializes the particles when the first range measurement R_m is received by randomly placing all of the particles on a circle of radius R_m around the origin. This is done in lines 4 to 6. Once the particles have been initialized, we perform a prediction step every $100ms$, taking into account robot movement noise with a standard deviation σ_{robot} (lines 8 to 11). When a new range measurement is received we compute the probability of observing such a range for every particle taking into account the measurement noise σ_{range}. The particles are then re-weighted according to the algorithm presented in Thrun *et al.* [12]. Finally, at the end of every iteration we compute a new heading for the robot. This is done

Algorithm 3 Acoustic Particle Filter

1: $YAW_{robot} \Leftarrow \pi * Random(-1.0...1.0)$
2: $SPEED_{robot} \Leftarrow 0.25$ Knots forward
3: Receive first measurement $RANGE_m$
4: **for** $k = 1 \dots N$ **do**
5: $\quad \begin{bmatrix} X_k \\ Y_k \end{bmatrix} = RANGE_m \cdot \begin{bmatrix} \cos(\alpha_k) \\ sin(\alpha_k) \end{bmatrix}$, where $\alpha_k = Random(-\pi \dots \pi)$
6: **end for**
7: **while** No optical link available **do**
8: $\quad$ **for** $k = 1 \dots N$ **do**
9: $\quad\quad$ Independently draw e_x and e_y from $\mathcal{N}(0.0, \sigma_{robot})$
10: $\quad\quad \begin{bmatrix} X_k \\ Y_k \end{bmatrix} = \begin{bmatrix} X_k \\ Y_k \end{bmatrix} + SPEED_{robot} \cdot \begin{bmatrix} \cos(YAW_{robot}) \\ sin(YAW_{robot}) \end{bmatrix} + \begin{bmatrix} e_x \\ e_y \end{bmatrix}$
11: $\quad$ **end for**
12: $\quad$ **if** Received new measurement $RANGE_m$ **then**
13: $\quad\quad$ **for** $k = 1 \dots N$ **do**
14: $\quad\quad\quad W_k = \frac{1}{\sqrt{2 \cdot \pi \cdot \sigma_{range}}} \cdot \exp(\frac{-D_k^2}{2 \cdot \sigma_{range}})$, where $D_k = RANGE_m - \sqrt{X_k^2 + Y_k^2}$
15: $\quad\quad$ **end for**
16: $\quad\quad$ RESAMPLE particles using weights $W_k, k \in \{1 \dots N\}$.
17: $\quad$ **end if**
18: $\quad \begin{bmatrix} \bar{Z}_X \\ \bar{Z}_Y \end{bmatrix} = \frac{1}{N} \sum_{n=1}^{N} \begin{bmatrix} -X_k \\ -Y_k \end{bmatrix}$, $\mu_\theta = \text{atan}(\bar{Z}_Y, \bar{Z}_X)$, $\sigma_\theta = \sqrt{\ln(1/\|\bar{Z}\|)}$
19: $\quad YAW_{robot} = \mu_\theta + \sigma_\theta / 4$
20: **end while**
21: Begin Optical Gradient Descent

by computing the heading required for every particle to travel towards the node. We assume these headings form a wrapped normal distribution and we compute the mean and standard deviation in line 18. Setting the robot heading to the mean would result in the particles traveling directly towards the node and can create a baseline ambiguity. Thus, we chose to set the new headings as the mean plus the standard deviation divided by a factor of 4. The more uncertain the particle filter is, the more the robot will deviate from the straight path, and this in turn helps resolve the baseline ambiguity.

5 Theoretical Performance

Figure 1(a) shows the upper bounds of achievable data rates for our approach plotted against the travel distance for the AUV (one way). The black line shows the achievable data rates using the acoustic channel computed as

$$R_{acoustic}(d) = \frac{40,000\text{m-Kbps}}{d}.$$

according to Kilfoyle *et al.* [6]. In color are shown the upper bounds for different time intervals that the AUV spends above the sensor nodes. We compute the upper bounds of the effective data rates for data muling as follows. We assume the AUV travels on a direct path (best possible case) to the sensor node at a speed of $v_{AUV} =$ 1 m/s. This is the speed at which AMOUR can travel. Further, we assume the optical data rate to be $r_{OPT} = 4$ Mbps. This is the data rate of AquaOptical as used during the experiments. Finally, we call d the distance of the sensor from the user (one way trip in meters) and t_{hover} the time that the robot hovers above the sensor node to download the data. Under these assumptions the total travel time of the robot from the user to the node and back, including the time spent hovering above the sensor node, is

$$t_{travel}(d) = t_{hover} + \frac{2 \cdot d}{v_{AUV}}$$

This results in an effective data rate of

$$r_{OPT,t_{hover}}(d) = \frac{t_{hover} \cdot r_{OPT}}{t_{travel}(d)}.$$

Figure 1(b) shows the resulting latencies under the same assumptions. For the acoustic communication we compute the latency as distance over speed of sound in water, i.e.

$$L_{acoustic}(d) = \frac{d}{1,500\text{m/s}}$$

The optical latency is equivalent to $t_{travel}(d)$. Figure 1(a) shows that the achievable data rate even using data muling far exceeds the currently available acoustic data

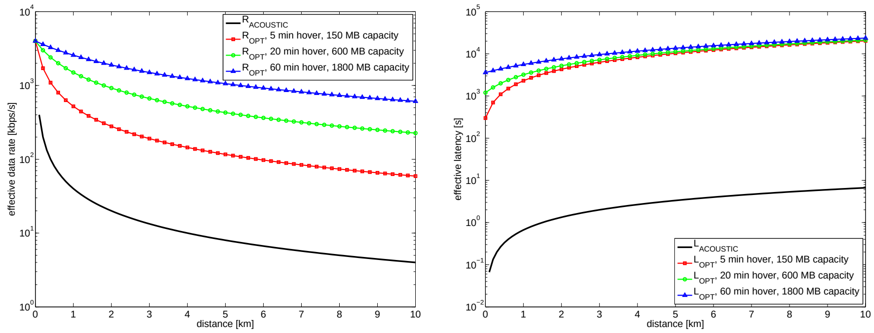

(a) Effective data rates over distance.

(b) Effective latencies over distance.

Fig. 1 The left graph shows effective data rates for given distances between the sensor node and the user. The x axis shows the distance in km and the y axis the data rate. Black shows data rates for acoustic communication. The colored lines show data rates for data muling with different times spent hovering above the sensor to download the data. The right graph shows effective latencies for the same theoretical cases as in the left figure. The x axis shows the distance in km and the y axis the resulting latency in seconds. The latencies are reported as worst-case latencies for data mulling (i.e. the entire trip time).

rates. This effect can even be amplified by using multiple AUVs which can travel in parallel to either a single or multiple sensor nodes. Using acoustic communication neighboring nodes often have to share the medium further reducing the effective data rate per node. The disadvantage of data mulling is its higher latency as seen in Figure 1(b).

6 Simulations

We evaluated both the acoustic stochastic gradient descent algorithm (Algorithm 1), and the acoustic particle filter algorithm (Algorithm 3) in simulation. In each simulation the robot state was represented as $[X_{robot}\ Y_{robot}\ YAW_{robot}]$. The robot was simulated with a constant speed of $SPEED_{robot} = 1m/s$. Independent white Gaussian noise with a standard deviation σ_{robot} in meters was added to the robot's position every second to simulate movement errors. Thus, every second the new robot position was computed as

$$\begin{bmatrix} X_{robot}(t+1) \\ Y_{robot}(t+1) \end{bmatrix} = \begin{bmatrix} X_{robot}(t) \\ Y_{robot}(t) \end{bmatrix} + SPEED_{robot} \cdot \begin{bmatrix} \cos(YAW_{robot}) \\ sin(YAW_{robot}) \end{bmatrix} + \begin{bmatrix} e_x \\ e_y \end{bmatrix}$$

where e_x and e_y are independently drawn from $\mathcal{N}(0.0, \sigma_{robot})$. Measurements were simulated every second with added Gaussian noise with a standard deviation σ_{range}. Each new measurement $RANGE_m$ is computed as

$$RANGE_m = \sqrt{X_{robot}^2 + Y_{robot}^2} + e_r$$

where e_r is drawn from $\mathcal{N}(0.0, \sigma_{range})$.

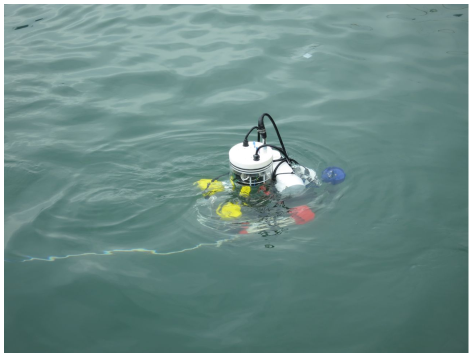

(a) Picture of AMOUR 6 [4].

(b) Picture of experimental site.

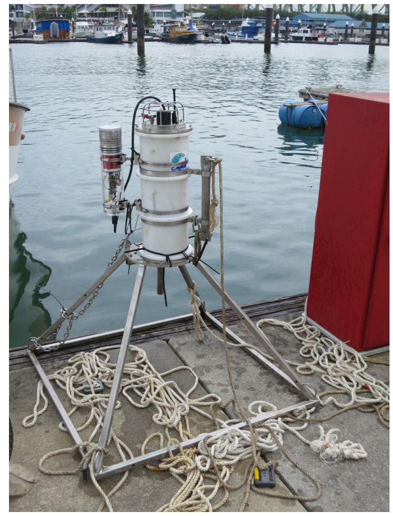

(c) Picture of PANDA with Optical Modem.

Fig. 2 (a) AMOUR 6 in the water with acoustic and optical modems attached. (b) Experimental site. The PANDA was deployed in the middle of the basin enclosed the dock. (c) PANDA node (white cylinder on tripod) with Optical Modem attached on the left.

Figure 4 shows two sets of 10 simulated paths taken by the robot using stochastic gradient descent (Fig. 4(a)) and a particle filter (Fig. 4(b)). The simulations were performed using Algorithm 1 and Algorithm 3. The parameters for these simulations were $\sigma_{range} = 1m$ and $\sigma_{robot} = 0.1m$. These plots visualize the characteristic difference in paths generated by the stochastic gradient descent and the particle filter. When the stochastic gradient descent encounters an increasing range, it picks a new direction almost entirely at random. The particle filter, on the other hand, continuously merges all gathered information about the sensor node location and continuously updates the robot's heading resulting in a more direct path. Plotted in Figures 4(c) and 4(d) are the corresponding distances of the robot to the sensor node over time.

Finally we conducted six sets of simulation runs with parameters chosen as $(\sigma_{range}, \sigma_{robot}) \in \{0.1m, 1.0m\} \times \{0.01m, 0.1m, 1.0m\}$. For each set we simulated 1000 runs using the stochastic gradient descent algorithm and 1000 runs using the particle filter. Figure 5 shows the results for all these runs grouped in six plots according to parameter choices. Each plot shows the mean distance over time (solid line) with 1σ boundaries (dashed lines). The stochastic gradient descent results are plotted in red, the particle filter results are plotted in blue. In all six cases the particle filter outperforms the stochastic gradient descent. When the noise is low (Fig-

(a) Panoramic picture of experimental site at Pandan Reservoir.

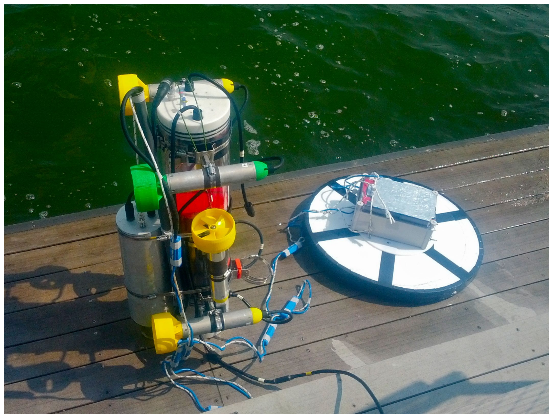

(b) Picture of AMOUR 6 and floating Wifi.

(c) Picture of PANDA.

Fig. 3 (a) Experimental site at Pandan Reservoir in Singapore. (b) AMOUR 6 with acoustic modem attached with the WiFi bouy next to the robot. (c) PANDA node with floats and weight.

ures 5(a) and 5(c)), the particle filter takes on average 104 seconds for the robot to come to within $5m$ of the sensor node. This is only 10% more then the theoretical minimum, which is 95 seconds since the robot starts $95m$ away and travels at $1m/s$.

7 Experiments

We conducted two sets of experiments to demonstrate the system's ability to localize a sensor node using a robot and recover data from it. A third set of experiments was conducted to evaluate the performance of the particle filter. In this work we did not focus on the return of the robot to the base station.

The first two sets of experiments were conducted from a dock at the Republic of Singapore Yacht Club (Figure 2(b)). The water depth was about 7m and we estimated visibility at about 2m. The PANDA with the optical modem was mounted on a tripod to guarantee that they should be pointing upright after being lowered to the ground. This setup can be seen in Figure 2(c). Our vehicle AMOUR carrying the acoustic and optical modems can be seen in Figure 2(a). It was tethered for data collection and security, but operated autonomously during the experiments. The robot speed was set at about $0.25m/s$ to ensure safe operation and to keep distance changes at a reasonable rate between updates. Generating the LT-Codes as described in Section 3 requires substantial computation. Because of this we needed to reduce the number of packets transmitted from the optical modem on the PANDA to 392 packets a second. Including overhead this corresponds to a bit rate of 1.24 MBit per second. The remainder of the optical channel (2.76 MBit per second) was not utilized.

The first set of experiments consisted of manually placing the robot close to the PANDA node and using optical gradient descent to maintain a position close to the PANDA node. This experiment was conducted two times, one time with the robot at the water surface and the second time with the robot keeping a depth of 1.5m under the water surface. Given a water depth of about 7m, a height of the PANDA of about 1m, and a robot height of slightly below 1m, the first and second experiment had a minimum distance of 5m between the optical modems and 3.5m, respectively.

In the second set of experiments we manually positioned the robot at a distance of about 25m away from the PANDA, dove it to 2m depth where we started the acoustic gradient descent algorithm. This experiment was conducted 12 times, of which two were aborted because the robot's tether got entangled with obstacles in the har-

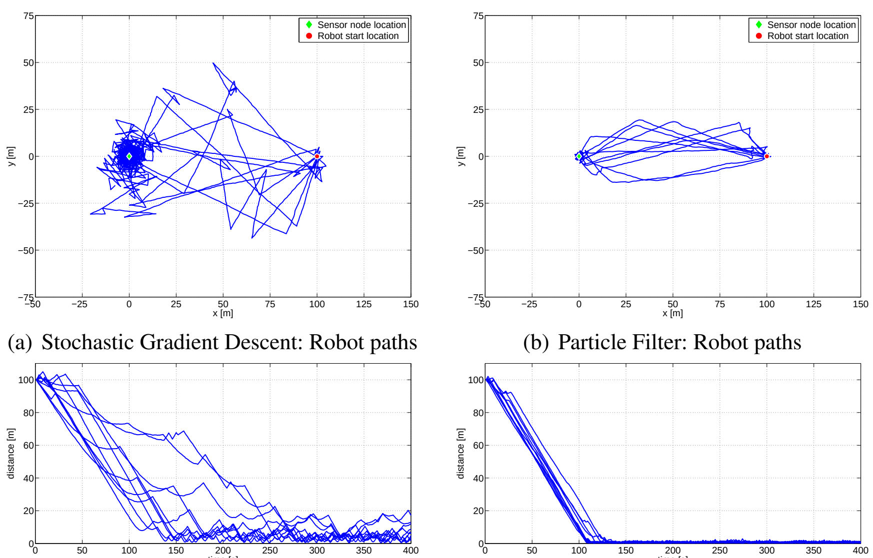

(a) Stochastic Gradient Descent: Robot paths (b) Particle Filter: Robot paths

(c) Stochastic Gradient Descent: Distances (d) Particle Filter: Distances

Fig. 4 Sample simulation results of 10 stochastic gradient descents and 10 particle filter node localizations using acoustic ranging in both cases. All simulations were performed with a robot speed of $1m/s$, measurement noise $\sigma_{range} = 1m$ and robot motion noise $\sigma_{robot} = 0.1m$. Range measurements occurred every $1s$. Plots (a) and (b) show the resulting robot paths. The x and y-axes show displacement in meters. The sensor node is located at the origin (green circle) and the robot starts at location (0m, 50m) (red diamond). Each continuous blue line denotes one simulation run. Plots (c) and (d) show the corresponding distances of the robot from the sensor node in m on the y axis over time in seconds on the x axis.

bor (the robot would dive under the docks sometimes due to the random nature of stochastic gradient descent). We used the other 10 experiments for evaluation.

The third set of experiments to measure the performance of the particle filter was conducted at the Pandan Reservoir in Singapore (Figure 3(a)). The reservoir covers an area of over $1km^2$ and has a depth of about $4m$ close to the shore. During this set of experiments we used a floating bouy that carried a long range WiFi. The buoy was tethered to the robot with a $5m$ long ethernet cable and additionally secured with a rope. This allowed for remote operation of the robot while it was able to move freely through the water without having a long tether that could get entangled in the many buoys that are present at the reservoir. To prepare the experiments the robot was used to transport the PANDA about $40m$ off shore and drop it there. A weight attached to the bottom of the PANDA together with floats attached at the top ensured that it would sink to the bottom of the reservoir but remain upright (Figure 3(c)). A rope

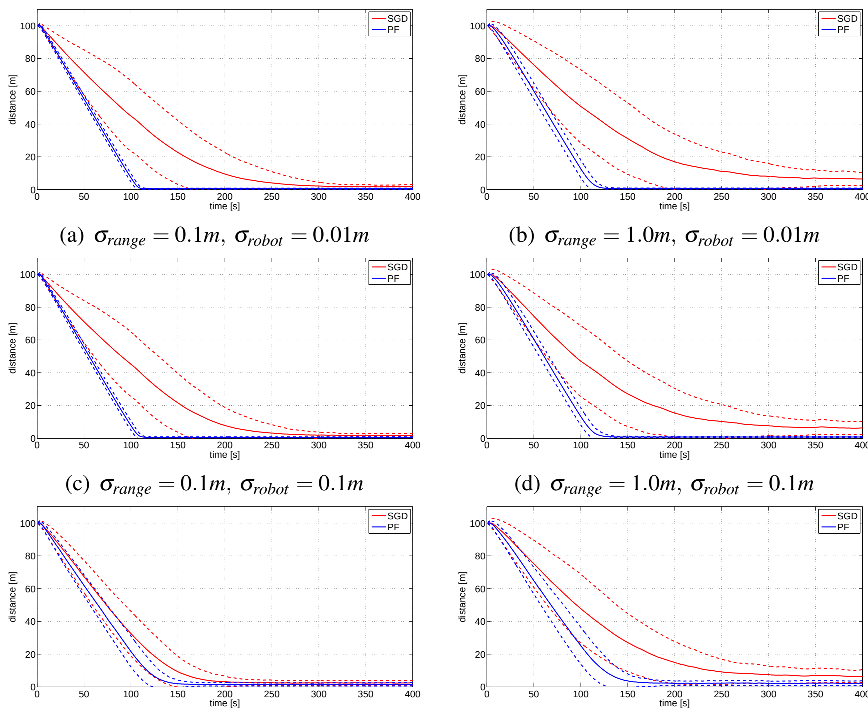

(a) $\sigma_{range} = 0.1m$, $\sigma_{robot} = 0.01m$

(b) $\sigma_{range} = 1.0m$, $\sigma_{robot} = 0.01m$

(c) $\sigma_{range} = 0.1m$, $\sigma_{robot} = 0.1m$

(d) $\sigma_{range} = 1.0m$, $\sigma_{robot} = 0.1m$

(e) $\sigma_{range} = 0.1m$, $\sigma_{robot} = 0.5m$

(f) $\sigma_{range} = 1.0m$, $\sigma_{robot} = 0.5m$

Fig. 5 Simulation results. Each plot corresponds to one of six different choices for $(\sigma_{range}, \sigma_{robot})$, the measurement noise and robot motion noise used in the simulation. All simulations were performed with a robot speed of $1m/s$ and range measurements occurred every $1s$. The x-axis corresponds to time in seconds and the y axis corresponds to the distance in meters of the simulated robot to the sensor node. Plotted in red are the results of stochastic gradient descent and in blue the results of the particle filter. For each algorithm and choice of $(\sigma_{range}, \sigma_{robot})$ 1000 simulations were performed for a total of 12000 simulations. The solid lines correspond to the mean, the dotted lines represent the 1σ boundaries.

was permanently attached to the PANDA that allowed us to recover it manually after the experiments. During each experiment the robot was manually positioned at a distance of at least $100m$ away from the PANDA and the node localization algorithm based on a particle filter was started. At all times the robot traveled with a speed of $0.5m/s$ at a depth of $1m$. The experiment was conducted 5 times, all of which were used for evaluation.

8 Results

The results of the optical gradient descent experiments can be seen in Figure 6. In the first run the robot maintained position for over 6 minutes before it lost track of the optical signal. During this time the optical modem on the PANDA transmitted 52.6 MB of payload data and on AMOUR received 37.5 MB of which 23.87 MB were error-free packets (one packet was 576 bytes large). In the second run it maintained position successfully for 11 minutes after which we stopped the experiment. During this time the PANDA transmitted 93.4 MB of payload data and AMOUR received

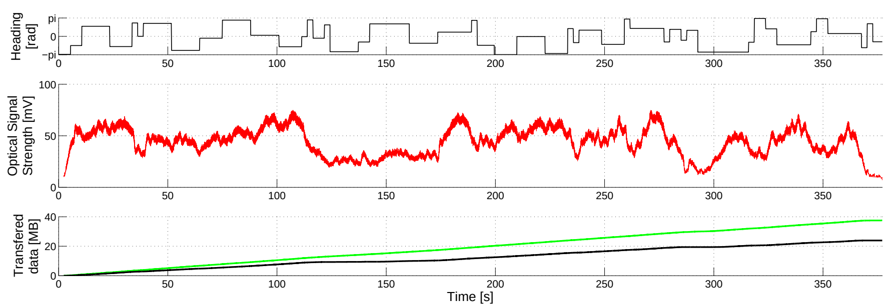

(a)

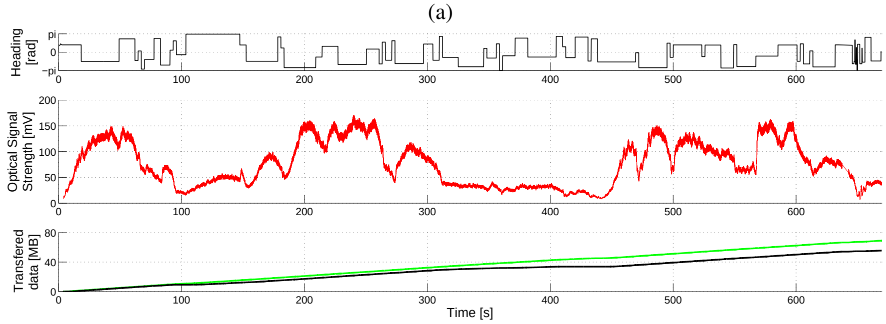

(b)

Fig. 6 Results of two optical gradient descent experiment runs. The first run is plotted in (a), the robot operated at the water surface. The second run is plotted in (b), the robot operated at a depth of 1.5 m. At the beginning of each experiment the robot was manually steered close to the PANDA to establish an optical link. The optical gradient descent algorithm then controlled the robot to stay close to the PANDA. In all plots the x-axis indicates the time in seconds since the beginning of the experiment. The top graph for each experiment (black) shows the heading of the robot as computed by the optical gradient descent algorithm. The middle graph (red) shows the measured signal strength. The bottom graph shows the amount of data received in MB (green) and the amount of data received error-free in MB (black). Packet size was 576 bytes with a 4 byte CRC.

69.1 MB of which 55.6 MB were error-free packets. The rate of error-free packets was higher in the second run because the robot was operating at a lower depth closer to the transmitter, which resulted in a higher signal strength at the receiver.

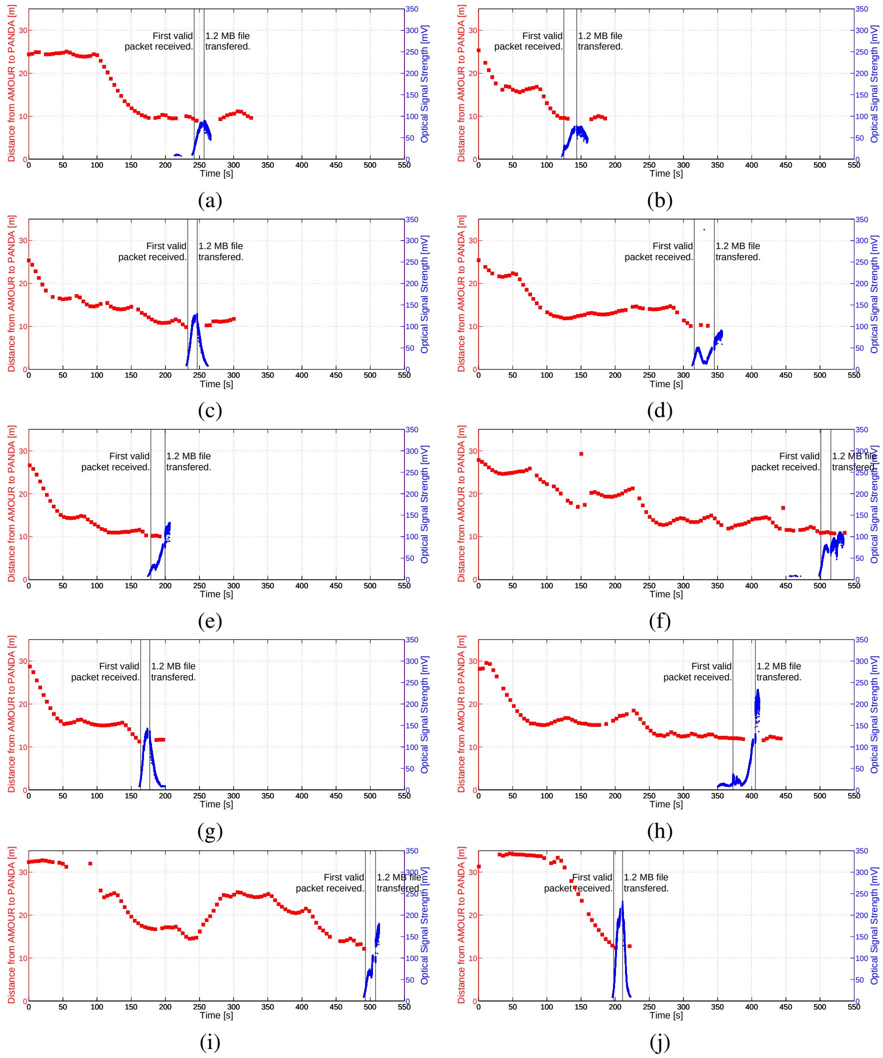

Fig. 7 Results of data muling experiments. Each graph (a)-(j) shows one experiment. The x-axis shows time in seconds since beginning of the experiment. The red curve (left y-axis) shows the distance between AMOUR and the PANDA node. Each red square corresponds to one range measurement received by AMOUR. The blue curve (right y-axis) shows the optical signal strength between PANDA and AMOUR. A link was established whenever there is a non-zero signal strength. Two horizontal black lines mark the receipt of the first error free packet (left line) and the receipt of the final error free packet needed to decode the 1.2 MB test file (right line).

The results of the second set of experiments can be seen in Figure 7. In all 10 experiments the robot successfully found the PANDA within 2.5 to 8 minutes and proceeded to download the 1.2 MB file within an additional 10 to 35 seconds.

Figure 8 shows the results of the third set of experiments in which a particle filter was used. In all 5 experiments the robot successfully found the PANDA within 4.2 to 9 minutes. This is a significant improvement over stochastic gradient descent when considering that the robot was coming from a distance 4 times larger than in the second set of experiments.

9 Main Experimental Insights

The proposed data muling system using bi-modal acousto-optical communication allows for large scale data recovery and eliminates the need for precise localization of the node and robot. It allows quick in-situ deployment of nodes and successive autonomous data recovery. In all gradient descent experiments the robot successfully found the underwater sensor node within a few minutes using acoustic gradient descent and proceeded to download a 1.2 MB file within 10 to 35 seconds using the

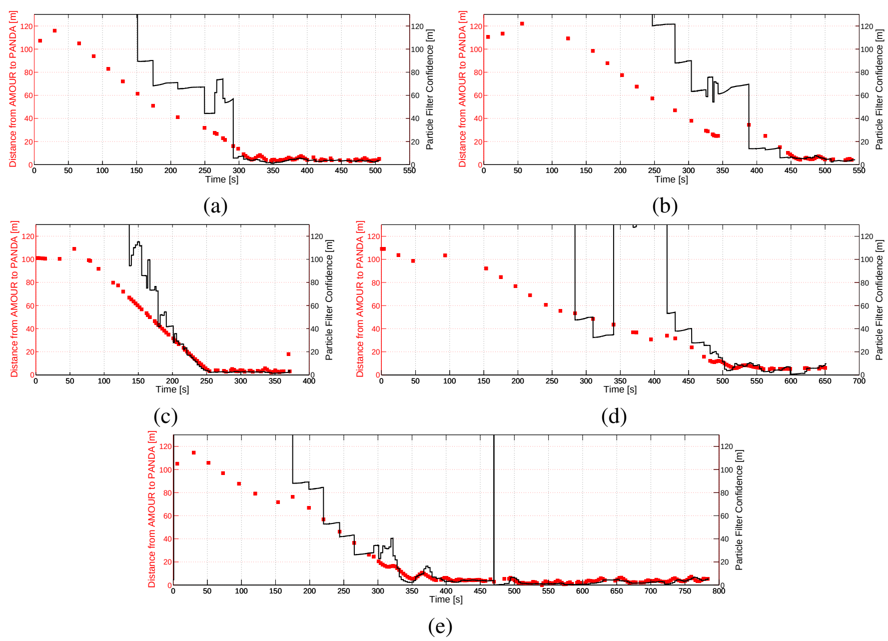

Fig. 8 Results of acoustic particle filter experiments. Each graph (a)-(e) shows one experiment. The x-axis shows time in seconds since beginning of the experiment. The red curve (left y-axis) shows the distance between AMOUR and the PANDA node. Each red square corresponds to one range measurement received by AMOUR. The black curve (right y-axis) shows the confidence of the particle filter (square root of the determinant of the covariance of all particle positions, lower is better). The spike visible at second 420 of plot (e) was caused by a spurious range measurement.

optical link. Further we demonstrated that we can use the optical signal strength to maintain the robot's position close to the position of the sensor node. We also experimentally evaluated the use of a particle filter to locate the node using only acoustic ranging. Both in simulation and experimentally the particle filter performed better than stochatic gradient descent.

If the PANDA had been able to generate LT-Codes at the full rate of 4 MBit/sec then our throughput would have been 3.2 times higher than measured. It should be noted that this was purely a limitation on the computational side and not a limitation of the optical or acoustic modem itself. Also, since we do not use error coding and correction, all packets with a single bit error were discarded. This amounted to 13.6 MB of 37.5 MB and 13.5 MB of 69.1 MB in payload data lost. With the expense of more computational resources this bandwidth can be almost entirely utilized.

Future improvements of the system include the usage of the acoustic link to turn on and off both the optical modem and the acoustic beacons (or at least reducing their frequency) to save battery life while the robot is not in range. We also plan to extend the presented data muling system to three dimensions which will allow for the nodes to be deployed at greater depths. The gradient descent algorithm can be extended to include state in order to speed up the robot's successful approach towards the node. Further, we plan to extend experiments into the open ocean where the algorithm can be tested at distances of multiple kilometers.

Acknowledgments

The authors would like to acknowledge the help of Mohan Panayamadam, Andy Marchese, TeongBeng Koay, and Puthenpurayil Unnikrishnan Saneesh.

References

1. Benson, C., Dunbar, R., Ryan, M., Huntington, E., Frater, M.: Towards a dense high-capacity underwater acoustic network. In: Communication Systems (ICCS), 2010 IEEE International Conference on, pp. 386 –389 (2010). DOI 10.1109/ICCS.2010.5686512
2. Chitre, M., Topor, I.: The unet-2 modem - an extensible tool for underwater networking research. In: OCEANS 2012 - Yeosu, pp. 1 –9 (2012)
3. Doniec, M., Rus, D.: Bidirectional optical communication with aquaoptical ii. In: Communication Systems (ICCS), 2010 IEEE International Conference on, pp. 390 –394 (2010). DOI 10.1109/ICCS.2010.5686513
4. Doniec, M., Vasilescu, I., Detweiler, C., Rus, D.: Complete se(3) underwater robot control with arbitrary thruster configurations. In: In Proc. of the International Conference on Robotics and Automation. Anchorage, Alaska (2010)
5. Dunbabin, M., Corke, P., Vasilescu, I., Rus, D.: Data muling over underwater wireless sensor networks using an autonomous underwater vehicle. In: Proc. IEEE ICRA 2006, pp. 2091–2098. Orlando, Florida (2006)
6. Kilfoyle, D., Baggeroer, A.: The state of the art in underwater acoustic telemetry. Oceanic Engineering, IEEE Journal of **25**(1), 4 –27 (2000). DOI 10.1109/48.820733

7. Li, K., Shen, C.C., Chen, G.: Energy-constrained bi-objective data muling in underwater wireless sensor networks. In: Mobile Adhoc and Sensor Systems (MASS), 2010 IEEE 7th International Conference on, pp. 332 –341 (2010). DOI 10.1109/MASS.2010.5664026
8. Luby, M.: Lt codes. In: Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on, pp. 271 – 280 (2002). DOI 10.1109/SFCS.2002.1181950
9. Olson, E., Leonard, J., Teller, S.: Robust range-only beacon localization. In: Proceedings of Autonomous Underwater Vehicles, 2004 (2004)
10. Shah, R., Roy, S., Jain, S., Brunette, W.: Data mules: modeling a three-tier architecture for sparse sensor networks. In: Sensor Network Protocols and Applications, 2003. Proceedings of the First IEEE. 2003 IEEE International Workshop on, pp. 30 – 41 (2003). DOI 10.1109/SNPA.2003.1203354
11. Stojanovic, M.: Recent advances in high-speed underwater acoustic communications. Oceanic Engineering, IEEE Journal of **21**(2), 125 –136 (1996). DOI 10.1109/48.486787
12. Thrun, S.: Probabilistic robotics. Commun. ACM **45**(3), 52–57 (2002). DOI 10.1145/504729.504754. URL http://doi.acm.org/10.1145/504729.504754
13. Vaganay, J., Baccou, P., Jouvencel, B.: Homing by acoustic ranging to a single beacon. In: OCEANS 2000, vol. 2, pp. 1457 – 1462 (2000)