

Autonomous Data Collection from Underwater Sensor Networks using Acoustic Communication

Geoffrey A. Hollinger, Urbashi Mitra, and Gaurav S. Sukhatme

Abstract—We examine the problem of planning paths for an autonomous underwater vehicle (AUV) to collect data from an underwater sensor network. The sensors in the network are equipped with acoustic modems that provide noisy, range-limited communication. The AUV must plan a path that maximizes the information collected while minimizing travel time or fuel expenditure. This problem is closely related to the classical Traveling Salesperson Problem (TSP), but differs in that data from a particular sensor has a probability of being collected depending on the quality of communication. We propose methods for solving this problem by extending approximation algorithms for variants of TSP, and we compare our proposed algorithms to baseline strategies through simulated experiments with varying levels of communication quality. Our simulations utilize a realistic model of acoustic communication to determine the probability of acquiring data from each sensor. The results demonstrate that planning the tour for the entire network while exploiting the communication model during planning improves performance versus myopic methods.

I. INTRODUCTION

The use of sensor fields to monitor phenomena in underwater environments is of growing interest. Examples include algal blooms [1], seismic activity, and intrusion of enemy submarines [2]. In underwater monitoring scenarios, many standard methods of communication are no longer feasible (e.g., WiFi, cellular, satellite). Acoustic modems can provide communication underwater, but they suffer from severe range limitations and channel variations [3].

Without reliable communication, collecting data from underwater sensor networks becomes a challenging problem. A potential solution is the use of a mobile autonomous underwater vehicle (AUV) equipped with an acoustic modem to gather data from the sensors [4]. The problem now becomes one of planning the AUV’s path to minimize its travel time and maximize information gathered. We will refer to this as the Communication-Constrained Data Collection Problem (CC-DCP). The CC-DCP is closely related to the classical Traveling Salesperson Problem (TSP) [5]. The key difference is that information is gathered from sensors through a noisy channel, whose reliability decreases with distance and can be modeled probabilistically. Thus, we are dealing with a new problem of TSP with probabilistic neighborhoods.

G. Hollinger and G. Sukhatme are with the Dept. of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA, {gahollin,gaurav}@usc.edu

U. Mitra is with the Dept. of Electrical Engineering, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA, ubli@usc.edu

This research has been funded in part by the following grants: ONR N00014-09-1-0700, ONR N00014-07-1-00738, NSF 0831728, NSF CCR-0120778 and NSF CNS-1035866.



Fig. 1. Example sensor deployment on the ocean floor to monitor environmental conditions. Such sensors remain in place for many months. Retrieving data from the sensors during deployment is challenging due to communication limitations underwater. A mobile AUV equipped with an acoustic modem can act as a data collection device in this scenario by traversing a path that minimizes travel time and maximizes information gathered.

In this paper, we design algorithms for an AUV gathering data from an underwater sensor network. The novelties of this paper include: (1) the formulation of the Communication-Constrained Data Collection Problem (CC-DCP) as a Traveling Salesperson Problem (TSP) with probabilistic neighborhoods, (2) the extension of approximation algorithms for the deterministic Prize-Collecting TSP and TSP with Neighborhoods for use in the CC-DCP, and (3) the validation of proposed solutions to the CC-DCP through simulated experiments at varying levels of communication.

The remainder of this paper is organized as follows. We first examine related work (Section II) and then formulate the CC-DCP (Section III). We next propose a number of algorithms for solving the CC-DCP approximately (Section IV), and we provide theoretical analysis demonstrating performance guarantees in certain instances (Section V). We then validate our approach through simulated deployments (Section VI) before concluding and discussing avenues for future work (Section VII).

II. RELATED WORK

The underwater data collection problem is closely related to the Traveling Salesperson Problem (TSP). In TSP, a mobile agent must visit a number of locations in the minimal amount of time. This problem is known to be NP-complete and has a long history of both approximate and optimal solutions. With current methods, it is possible to solve regular TSP instances optimally with more than 10,000 locations [5]. The two key differences between our work and the classical TSP are: (1) the locations are associated with sensor measurements that may provide different amounts of

information regarding the phenomena of interest, and (2) information is gathered from locations based on a probabilistic communication model.

The case where locations are associated with different “prizes” has been studied by a number of researchers. The Prize-Collecting TSP was originally introduced by Balas [6] and has been extended to a number of related variants [7], depending on the type of path required, restrictions on the prizes, and whether new locations appear during the tour [8]. Goemans and Williamson proposed an approximation algorithm for Prize-Collecting TSP using a LP primal/dual schema that achieves a factor of two guarantee [9]. Slightly better approximation guarantees are also possible at the cost of computation and implementation complexity [10].

Researchers have also studied the TSP with deterministic neighborhoods. In this formulation, the locations are visited when the agent moves within a fixed radius around the location. Constant factor approximation algorithms are available for TSP with neighborhoods for both disjoint and overlapping regions of the same size [11] as well as regions of varying sizes [12]. To our knowledge, the combination of TSP with Neighborhoods and Prize-Collecting TSP has not been studied, and the use of probabilistic communication models to describe the neighborhoods has also not previously been examined.

Related problems have been studied in the context of robotic data mules. Bhaduria and Isler derived approximation algorithms for multiple data mules that must traverse a sensor field and download data [13]. In their work, downloading time is considered as part of the tour, and the communication radii are assumed to be uniform, fixed, and deterministic (i.e., data from a sensor is known to be accessible at a given location). In the present paper, we utilize a probabilistic acoustic communication model that degrades with distance [14]. Such models, developed by researchers in acoustic communication, have been shown to provide accurate predictions for experimental data [3].

Vasilescu et al. developed a system of mobile and stationary nodes for underwater data collection based on the use of both optical and acoustic communication [4]. They described the networking architecture and sensor specifications necessary for underwater data collection, and they presented experiments in the field on a mobile network. These experiments showed the feasibility of utilizing AUVs for underwater data collection, but the authors left open the problem of path planning for the mobile nodes in large networks.

III. PROBLEM SETUP

We are given a pre-deployed network of N sensors located in \mathbb{R}^d . For this paper, we limit analysis to $d \in \{2, 3\}$, which yields the 2D and 3D problems respectively. We assume that the location $x_n \in \mathbb{R}^d$ is given for each sensor $n \in [N]$, where $[N] := \{1, \dots, N\}$. Each sensor n contains data for retrieval, which we denote as Y_n . We define the *information quality* of the data as $I(Y_n)$. $I(x)$ corresponds to the expected value of information received at location x (e.g., information

gain in an inference problem [15], or variance reduction in a regression problem [16]). Generally, $I(x)$ incorporates both information received from the stationary sensors and also information sensed by a vehicle at location x . For this paper, we will not consider sensing actions taken by the vehicle, although these can be incorporated.

In the general case, coupling between the sensor measurements can lead to information being subadditive or super-additive. In the context of data collection, we will assume that information is additive (i.e., $I(Y_i, Y_j) = I(Y_i) + I(Y_j)$ for all $i \neq j$) and that multiple observations of any Y_n do not provide additional information. Relaxing either of these assumptions leads to interesting extensions (see Section VII).

The sensors are assumed to be low-cost with limited capabilities. Each sensor is capable of transmitting packets of data over a noisy channel. A single mobile vehicle has the capability to communicate with the sensors. The location $x_v \in \mathbb{R}^d$ of the vehicle is controlled and may be subject to constraints, such as obstacles or vehicle kinematics. Based on these constraints, a *traversal cost* $c(x_1, x_2)$ is defined for all pairs of points $x_1, x_2 \in \mathbb{R}^d$. We assume the traversal cost is metric and that the location of the AUV is known. The *communication quality* of the vehicle degrades with distance, $C(x_v, x_n) = f(D(x_v, x_n))$, where $D(x_v, x_n) = |x_n - x_v|$ and f is some relationship function. We will assume that f decreases monotonically with distance.

The optimization problem is to plan a path for the vehicle that retrieves data from the sensors and minimizes the traversal cost of the path. Let $P(x_v, x_n) \sim C(x_v, x_n)$ represent the probability that the observation made by the sensor at x_n is received by a vehicle at x_v . This represents the probability that the AUV has successfully pinged the sensor and received the entire data packet as a reply. We will assume that the AUV is always sending ping messages to each sensor, and that the sensors reply with a data packet whenever a ping is received. More general communication architectures are also possible, but they are not examined in this paper. Given this model, the expected received information quality at point x_v can be written as

$$R(x_v) = \sum_{\mathcal{N} \in 2^K} P(x_v, \mathcal{N}) I(Y_{\mathcal{N}}), \quad (1)$$

where 2^K is the powerset of K sensors with uncollected data, $I(Y_{\mathcal{N}})$ is the joint information from a set of $\mathcal{N} \in 2^K$ sensors, and $P(x_v, \mathcal{N})$ is the probability of observing \mathcal{N} . Using the independence assumption on sensors, this reduces to

$$R(x_v) = \sum_{n \in [N]} P(x_v, x_n) I(Y_n). \quad (2)$$

With the independence assumption, we can also calculate the expected received information along a path $\mathcal{P} = [x_v(1), x_v(2), \dots, x_v(T)]$ as

$$R(\mathcal{P}) = \sum_{n \in N} \left[1 - \prod_{t \in [T]} [1 - P(x_v(t), x_n)] \right] I(Y_n). \quad (3)$$

The Communication-Constrained Data Collection Problem (CC-DCP) can now be written formally.

Problem 1: Given path costs c , information quality I , communication quality C , a set of possible AUV paths Ψ , and a starting AUV location x_s , find

$$\mathcal{P}^* = \operatorname{argmin}_{\mathcal{P} \in \Psi} \sum_{t \in [T]} c(\mathcal{P}(t-1), \mathcal{P}(t)) \text{ s.t. } R(\mathcal{P}) \geq \tau, \quad (4)$$

where T is the index of the last location on the path, and τ is a threshold on information quality.

In some cases, we may want to set the hard constraint to gather information from all sensors, which makes the traversal distances stochastic. For instance, it may take a short amount of time to gather data if the transmissions are received early or a longer time if they are received later. For this case, the optimal solution is a mapping $2^N \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ from the 2^N possible *received* and *not received* states of the N sensors and the current location of the AUV to the next best location of the AUV. This solution can be encoded as a policy for an underlying Markov Decision Process (see Section IV).

IV. ALGORITHMS

We now present algorithms for solving the CC-DCP both optimally, at the cost of high computation, and heuristically, based on existing approximation algorithms for TSP variants. In the following section, we will demonstrate that our algorithms exhibit performance guarantees for some instances of the CC-DCP.

A. Optimal MDP Solution

The optimal solution to the CC-DCP can be encoded as a policy mapping from states to actions. To see this, note that the problem can be formulated as a Markov Decision Process (MDP) [17]. In the 3D problem, the states in the MDP are defined as $\mathcal{S} = \mathcal{X} \times \mathcal{Y} \times \mathcal{Z} \times 2^N$, where \mathcal{X} , \mathcal{Y} , and \mathcal{Z} are the coordinate spaces for the location of the AUV. The 2^N states represent whether or not data has been collected from each of the N sensors. The reward function $\mathcal{R}(s, s') = \sum_{n \in \mathcal{F}} I(n) - c(s, s')$, where \mathcal{F} is the set of sensors that go from uncollected to collected between s and s' , and $c(s, s')$ is the movement cost. The state transitions $\mathcal{T}(s, s')$ are defined by the communication model $C(s, s')$ and the motion model of the AUV.

If the environment is discretized, the MDP above can be solved using any standard method (e.g., value iteration or policy iteration) [17]. However, it is important to note that the number of states is exponential in the number of sensors and polynomial in the size of the environment (to the power of 2 or 3 depending on the dimension). Thus, we can expect optimal solutions to become infeasible for any instance of more than a few sensors. In addition, there may be additional

error that arises from discretizing the possible locations of the AUV. This formulation requires the location of the AUV to be known. Some position uncertainty can be incorporated in a conservative manner by employing the communication probability associated with the worst-case location of the AUV. Otherwise, a Partially Observable Markov Decision Process (POMDP) is required [18], which leads to a further increase in computational complexity.

B. Approximate Solutions

The CC-DCP is a generalization of the TSP, which makes it NP-hard. Due to the computational intractability of the optimal solution on large instances, we develop heuristics for solving the CC-DCP approximately. The key difference between the CC-DCP and the TSP with neighborhoods [11] is that communication is modeled probabilistically. Our approach is to generate contours of equal probability around the sensors and utilize these as if they were deterministic neighborhoods.

We define a *probabilistic neighborhood* \mathcal{M}_n for sensor n determined by $P(x_m, x_n) > \rho$ for all $x_m \in \mathcal{M}_n$. The value of $\rho \in (0, 1)$ determines how conservative the probabilistic neighborhood is. As $\rho \rightarrow 1$, it will be near certain that information will be received from sensor n if the AUV is within the neighborhood. As $\rho \rightarrow 0$, the AUV may need to query a sensor n many times before receiving data from it.

Once the probabilistic neighborhoods are defined, we can generate a maximal independent set of neighborhoods by greedily choosing sensors and removing adjacent sensors with overlapping neighborhoods. A valid tour can then be found by circling the neighborhoods in the maximal independent set (see Figure 3) [11]. In the case of probabilistic neighborhoods, it may be necessary to wait for information to be received from the entire set of sensors before moving to the next set.

For the case where all nodes must be visited and have equal informativeness (i.e., $I(Y_i) = I(Y_j)$ for all i, j), we design the following heuristics for planning the path of the AUV:

- **Myopic:** Move towards closest sensor. Once data is received, move to next sensor.
- **TSP Solution:** Find an optimal TSP ordering of the sensors using the Concorde solver [5]. Visit sensors in that order. Shortcut sensor once data is received.
- **TSP Solution with Neighborhoods:** Find a maximal independent set (MIS) of probabilistic neighborhoods. Find the optimal TSP ordering of the MIS using Concorde. Visit the sets in that order. Shortcut once data is received.

The myopic strategy is simply to move towards the closest sensor. This is a reactive strategy and will perform well when communication quality is high. In such cases, the benefit of long-term planning is negated by the homogeneity of the expected received information across the environment. The second solution is to solve the TSP of the sensors and ignore the communication model. This technique will perform well when communication is poor, since this situation requires the

AUV to move near all sensors to gain information from them. Finally, the TSP Solution with Neighborhoods incorporates the communication model as a probabilistic neighborhood.

For the case where the sensors provide different quantities of information, $I(Y_i) \neq I(Y_j)$, we propose the use of a Prize-collecting TSP algorithm (PC-TSP) to improve performance. The PC-TSP assigns a penalty $\zeta(i)$ to each location based on its informativeness. The tour now has the option of neglecting some locations and paying the required penalty. The total cost of the tour is then the movement cost plus $\zeta(i)$ for all i not visited. In our case, $\zeta(i) = \lambda I(Y_i)$, where λ is a scale factor. A binary search on λ can be performed if a threshold on information quality is desired, as in Equation (4). We employ the following strategies for the prize-collecting case:

- Myopic: Ignore the penalties and act as above. Terminate if the sum of remaining benefit is less than minimum distance to a sensor's probabilistic neighborhood.¹
- PC-TSP Solution: Use the primal/dual algorithm from [9] to determine sensors to visit. Find optimal ordering of this subset using the Concorde solver [5]. Visit sensors in that order. Shortcut once data received.
- PC-TSP Solution with Neighborhoods: Find a maximal independent set (MIS) of the sensors using probabilistic neighborhoods. Use the primal/dual algorithm on the MIS to determine a subset of the MIS to visit. Find the optimal ordering of the subset of the MIS using Concorde [5]. Visit in that order. Shortcut once data is received.

The non-myopic algorithms for the case of differing sensor information utilize an existing PC-TSP approximation algorithm to determine which sensors (or neighborhoods) to visit in the tour. The selected locations are then treated in a similar fashion to the problem with equal sensor information. We will now show that this approach provides performance guarantees relative to optimal in some cases.

V. THEORETICAL ANALYSIS

Difficult problems in NP, such as TSP variants, cannot be solved in time polynomial in the size of the instance, under the assumption that $P \neq NP$. A common approach is to develop efficient algorithms with bounded solution quality relative to optimal [19]. Such algorithms provide an *approximation guarantee* on solution quality. A constant factor approximation guarantee states that for any problem instance, the quality of the approximate solution $|\mathcal{P}^A|$ is bounded as $|\mathcal{P}^A| \leq \eta |\mathcal{P}^*|$, where η is a constant value (known as the approximation factor), and \mathcal{P}^* is the optimal solution. Clearly a lower η is better, and η will never be less than one (the approximate solution cannot be better than optimal).

In the case of the TSP with Neighborhoods (TSPN), the algorithm presented above of determining a maximal independent set of neighborhoods and then calculating a

¹An alternative myopic strategy is to move to the node with highest benefit/cost ratio. In practice, this did not perform as well as the simple nearest neighbor.

near-optimal TSP tour of this set is known to provide constant factor approximation guarantee. This guarantee is $\eta = (\pi + 8)(1 + \epsilon)$, where $1 + \epsilon$ is the approximation guarantee for the algorithm used to solve the TSP on the maximal independent set [11].

We now show that this approximation algorithm for the TSPN extends to the Prize-Collecting TSP with Neighborhoods (PC-TSPN) without a loss in approximation quality given a constraint on the prizes. We will first examine the case where the neighborhoods are deterministic (i.e., communication quality is measured as a step function). In addition, we limit the theoretical analysis to the problem in \mathbb{R}^2 and the case where the neighborhoods all have diameter d (see reference [11]).

Theorem 1: Let \mathcal{P}^A be the approximate solution to a PC-TSPN instance found by generating a maximal independent set I of N neighborhoods and then finding a PC-TSP tour that visits a subset of I . If $\zeta(i) \geq 2d$ for all $i \in N$, the approximate solution is bounded by $|\mathcal{P}^A| \leq (\pi + 8)(1 + \epsilon)|\mathcal{P}^*|$, where $1 + \epsilon$ is the approximation guarantee for the PC-TSP tour.

Proof: The proof follows that of standard TSPN closely [11]. The key difference is ensuring that choosing not to visit neighborhoods, by paying the required penalty, does not lead to loss in approximation quality. For the case of disjoint disk neighborhoods (i.e., no neighborhoods overlap), let \mathcal{P}^C be the path generated by running a PC-TSP tour with a $(1 + \epsilon)$ approximation guarantee on the center points of the disks. Note that this tour does not necessarily visit all neighborhoods. Let \mathcal{P}^R be the optimal PC-TSPN tour. For any path, let $|\mathcal{P}|$ be the cost of that path (i.e., traversal costs plus penalties for neighborhoods not visited). We claim that

$$|\mathcal{P}^C| \leq \left(\left(1 + \frac{8}{\pi}\right) |\mathcal{P}^R| + 8 \right) (1 + \epsilon). \quad (5)$$

Let $l = |\mathcal{P}^R|$, let $M \subset N$ be the neighborhoods visited by \mathcal{P}^R , and let $m = |M|$. W.l.o.g. scale the distances so that the size of all disk neighborhoods are radius 1. Since \mathcal{P}^R visits a subset of the disks, the area A_l swept by a disc of radius 2, whose center moves along \mathcal{P}^R is bounded as

$$\pi m \leq A_l \leq 4l + 4\pi, \quad (6)$$

and thus $m \leq 4 + 4l/\pi$. A tour \mathcal{P}^M that visits the center of the M disks can be found by going along \mathcal{P}^R and making a detour of length at most 2 to visit the center of each disk. This yields

$$|\mathcal{P}^M| \leq l + 2m \leq l + 2 \left(4 + \frac{4l}{\pi}\right) = \left(1 + \frac{8}{\pi}\right) l + 8. \quad (7)$$

Finally, we note that $|\mathcal{P}^C| \leq (1 + \epsilon) |\mathcal{P}^M|$, since $|\mathcal{P}^M|$ is a valid center tour. Note that \mathcal{P}^C and \mathcal{P}^M may visit different subsets of the neighborhoods. Thus, \mathcal{P}^C is bounded as in Equation (5). As $|\mathcal{P}^R|$ grows large, the approximation factor becomes $\left(1 + \frac{8}{\pi}\right) (1 + \epsilon)$.

Now, for the case of overlapping discs, we find a maximal independent set $I \subset N$ of the disks. Let \mathcal{P}^J be the PC-TSP tour that visits $J \subseteq I$ and is computed by running a $(1 + \epsilon)$

approximation algorithm on I , such that the prize of any neighborhood $i \in I$ is the sum of prizes of all neighborhoods covered by i . Note that \mathcal{P}^J may not be a valid disc tour of J . Next we compute a valid disc tour \mathcal{P}^A that circles the entire neighborhoods visited by \mathcal{P}^J . Let \mathcal{P}^* be the optimal PC-TSPN tour, and let \mathcal{P}^{I*} be the optimal PC-TSPN tour of I . We can derive an approximation guarantee from the following inequalities:

$$|\mathcal{P}^A| \leq \pi |\mathcal{P}^J| + 2\pi, \quad (8)$$

$$|\mathcal{P}^J| \leq \left(\left(1 + \frac{8}{\pi} \right) |\mathcal{P}^{I*}| + 8 \right) (1 + \epsilon), \quad (9)$$

$$|\mathcal{P}^{I*}| \leq |\mathcal{P}^*|. \quad (10)$$

Equation (10) follows from the fact that $I \subset N$ and that $\zeta(j) \geq 2d$ for all $j \in N$. The requirement on $\zeta(j)$ ensures that if one neighborhood covered by $i \in I$ is visited in the optimal tour, all neighborhoods covered by i are visited in the optimal tour. Equation (9) follows from the case of disjoint disks described above. Equation (8) is found through the same geometric analysis as in the regular TSPN [11], which is omitted here. From these inequalities, we have that

$$|\mathcal{P}^A| \leq ((\pi + 8) |\mathcal{P}^*| + 8\pi) (1 + \epsilon) + 2\pi. \quad (11)$$

As $|\mathcal{P}^*|$ becomes large, we have

$$|\mathcal{P}^A| \leq (\pi + 8) (1 + \epsilon) |\mathcal{P}^*|. \quad (12)$$

For the case of probabilistic neighborhoods, the guarantee above provides a bound on the solution quality of the algorithm when run on p contours. For instance, the solution returned for 80% contours is guaranteed to be less than a factor of $(8 + \pi)(1 + \epsilon)$ worse than the optimal solution that visits the 80% contours. This guarantee holds for both TSPN with probabilistic neighborhoods (from reference [11]) and PC-TSPN with probabilistic neighborhoods (from Theorem 1). Note that this guarantee does not account for solutions that may not visit all p contours (e.g., by sending many transmissions in a low communication quality area to achieve a high probability of receiving data).

VI. SIMULATIONS

A simulation environment was implemented in C++ running on Ubuntu Linux to test our proposed CC-DCP algorithms. The experiments were run on a 3.2 GHz Intel i7 processor with 9 GB of RAM. We now describe the acoustic communication modeling used in the simulation

A. Acoustic Communication Modeling

We utilize a probabilistic underwater acoustic communication model developed in prior work [2], [14]. This model accounts for noise factors in the environment, such as wind and shipping activity, as well as the distance and frequency dependent attenuation of the signal in the water medium. The model also accounts for both thermal noise and turbulence.

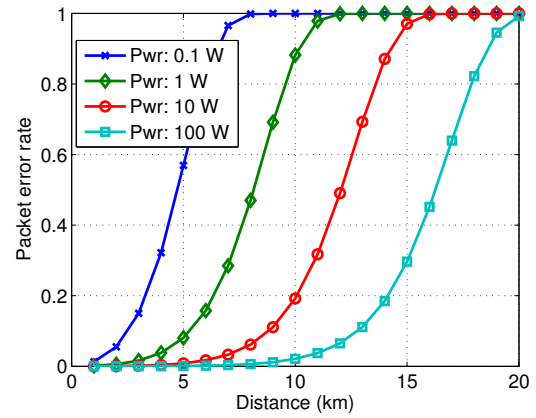


Fig. 2. Modeled data error rates for varying transmission power. The curves give the probability that a packet will not complete a successful round-trip between the AUV and a sensor node.

TABLE I
PARAMETERS FOR UNDERWATER ACOUSTIC COMMUNICATION

Parameter	Value
Transmission Frequency	13 kHz
Transmission Power	0.1 W - 10 kW
Bandwidth	1 kHz
Packet Size	256 symbols
Wind speed	5 m/s
Shipping activity	0.5 (moderate)

The following approximation is used to determine a symbol error rate based on a signal-to-noise-ratio (SNR):

$$P(e) \approx \frac{1}{4SNR}. \quad (13)$$

The above equation gives the probability of symbol error, $P(e)$, as a function of SNR for the case of Rayleigh fading with known channel states at the receiver. Ultimately, the SNR is a function of the transmitted power, frequency of transmission, wind speed, shipping factor, the distance and the bandwidth used (see reference [2] for more detail). For a packet with Q symbols, the probability of transmission error can then be computed as

$$P_{packet} = 1 - (1 - P(e))^Q. \quad (14)$$

Table I gives the parameters used for the communication model in the simulated experiments. Figure 2 shows the packet error rates at different levels of transmitted power. Note that the communication quality cannot be accurately modeled by a simple discrete neighborhood due to the intermediate communication levels.

B. Simulations with Random Deployments

We first implemented the optimal MDP solution using value iteration. This method was able to solve problems with up to 3 nodes on a 15×15 environment with a 1 km grid resolution. For these small problem sizes, the myopic nearest neighbor heuristic performed competitively with the optimal solution. Based on this finding, we scale up the size

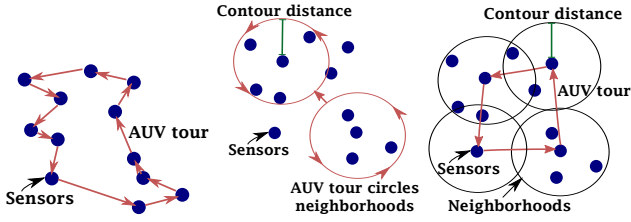


Fig. 3. Example tours using different neighborhood types. **Left:** Standard traveling salesperson tour [5]. **Center:** Tour circling a maximal independent set of neighborhoods [11]. **Right:** Tour visiting the center of a covering set of neighborhoods. All tours travel within the contour distance of all nodes, but the covering set tour is shortest.

of the instances to compare the approximate methods, and we remove the infeasible optimal method from consideration.

We now compare the myopic nearest neighbor strategy to the TSP strategy and the TSP with probabilistic neighborhoods (using 80% contours). We make one heuristic modification to the algorithm with probabilistic neighborhoods. Instead of generating an MIS and circling the entire neighborhoods, we generate a covering set at half the contour distance. This allows us to visit all sensors by simply planning a tour of the locations in the covering set. This modification improves performance in practice and also allows straightforward extension to 3D environments. Figure 3 gives a visualization of planning with varying neighborhood types, and the video attachment provides an animation of the simulated AUV executing data gathering tours.

The next experiments utilize random deployments of 100 sensors in $100 \text{ km} \times 100 \text{ km}$ 2D environments. The simulated underwater vehicle moves at 5 km/hr. The node utilities are set to be uniform, and the requirement is to visit all nodes. Figure 4 shows results from these trials. At low transmission power (poor communication), the value of utilizing the neighborhoods is minimal, and the problem reduces to the classical TSP. At high transmission power (good communication), the value of planning is reduced, and the simple myopic strategy moves closer to the quality of the non-myopic strategies. At moderate communication levels, there is some improvement from both solving the TSP and utilizing neighborhoods.

The same experiments were run with random information values from 0 to 25 were added to the 100 sensors. The total cost is calculated by summing traveled distance plus the scaled information values of sensors not collected. The scale factor was set as $\lambda = 1$ for these experiments. Modifying λ would allow for tuning the number of sensors visited. The PC-TSP approximation algorithm was compared with and without neighborhoods to the nearest neighbor strategy.

Figure 5 shows the numerical results for the prize-collecting case. Solving the underlying PC-TSP without neighborhoods does not perform well, even when compared to the myopic strategy. Since the PC-TSP algorithm cannot account for the cost of the neighborhoods when determining which sensors to visit, it chooses to ignore a number of sensors that would actually improve the final cost. In contrast, the combination of the neighborhoods and the

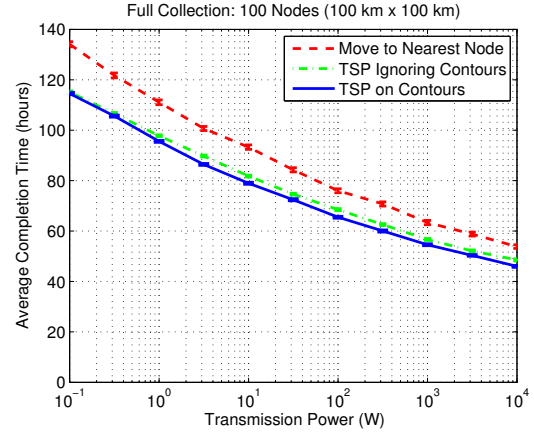


Fig. 4. Communication-constrained data collection experiments. Error bars are one SEM, and averages are over 100 random deployments. Solving the TSP provides improvement over myopic techniques, and utilizing neighborhoods provides some additional improvement.

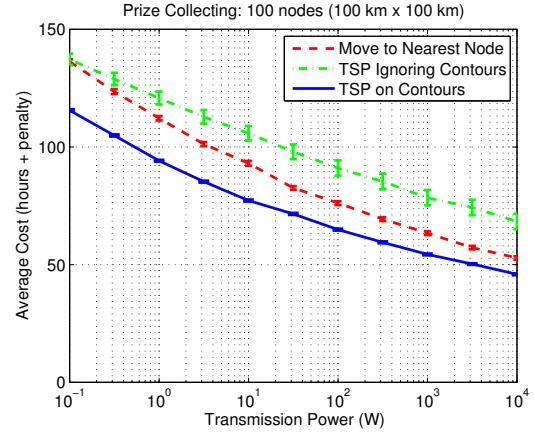


Fig. 5. Prize-collecting communication-constrained data collection experiments. Data is not collected from some sensors, and the required penalty is paid. Error bars are one SEM, and averages are over 100 random deployments. Considering neighborhoods significantly improves performance.

underlying PC-TSP approximation algorithms performs well. Putting Figures 4 and 5 together, we see that considering neighborhoods helps marginally in the case without prizes and more significantly when prizes are considered.

The running time of the algorithms is dominated by the cost of calculating the TSP tour with the Concorde solver. In the worst-case, this computation time can grow exponentially in the number of nodes. In practice, typical instances of 100 nodes were solved in 10-100 ms. For a more extensive discussion of the running time of Concorde, see reference [5]. When probabilistic neighborhoods are taken into account, the TSP solver uses a reduced set of nodes (those in the covering set). Thus, utilizing neighborhoods actually reduces the running time of the algorithm.

Additional simulations were run to determine the effect of changing the size of the probabilistic neighborhoods, determined by the parameter p . Figure 6 shows completion cost for prize-collecting problems using varying p values.

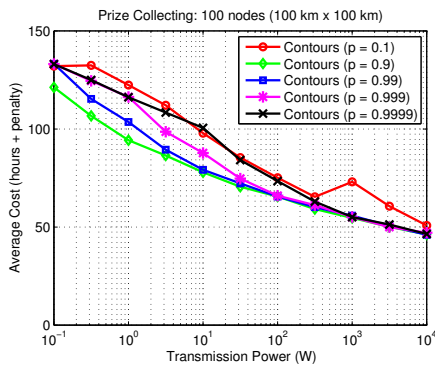


Fig. 6. Simulations with varying sizes of probabilistic neighborhoods. Each data point is averaged over 100 random deployments. At low transmission power (poor communication), the choice of neighborhood size has a larger effect on performance.

The results show that orders of magnitude changes in p are required to see significant differences in completion cost. This is an intuitive result, since the communication model predicts small changes in neighborhood size for large changes in transmission probability (see Figure 4). At low transmission power (i.e., when communication quality is poor), the value of p must be chosen more carefully. These results again demonstrate the importance of network scale and communication quality on the consideration of neighborhoods in planning.

VII. CONCLUSIONS AND FUTURE WORK

This paper has shown that communication-constrained data collection is feasible with an autonomous underwater vehicle. When communication quality is high (relative to the sparsity of the network), planning is less essential. When communication quality is low, it is important to solve the underlying TSP, rather than using a fully reactive approach. Finally, at moderate levels of communication, it is beneficial to consider the communication neighborhoods during planning. In addition, when information quality of the sensors is considered and all sensors do not need to be visited, considering neighborhoods improves performance. This analysis provides insight into the level of planning required to optimize information gathering at different levels of network sparsity, communication, and information quality. Such insight motivates the use of realistic communication models in the development and analysis of planning algorithms.

A number of interesting extensions provide avenues for future research. The case where sensor communication quality varies between sensors results in probabilistic neighborhoods of different sizes. While it is possible to apply our techniques directly to such cases, it is not clear if additional methods are necessary to provide good performance. Another interesting extension is the use of non-metric and asymmetric distances between sensors. For instance, in the case of ocean currents, it may be easier to travel in one direction than another. Such cases have been examined for the classical TSP [20], but not in the case of neighborhoods. Finally, dependencies between

sensors can result from correlations between gathered data. Exploiting the locality and submodularity of the information functions [16] may allow us to extend our algorithms to such cases as well.

VIII. ACKNOWLEDGMENTS

The authors gratefully acknowledge Jonathan Binney, Arvind Pereira, Hordur Heidarsson, and Srinivas Yerramalli at USC for their insightful comments.

REFERENCES

- [1] R. N. Smith, Y. Chao, P. P. Li, D. A. Caron, B. H. Jones, and G. S. Sukhatme, "Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model," *Int. J. Robotics Research*, vol. 26, no. 12, 2010.
- [2] G. Hollinger, S. Yerramalli, S. Singh, U. Mitra, and G. S. Sukhatme, "Distributed coordination and data fusion for underwater search," in *Proc. IEEE Conf. Robotics and Automation*, 2011.
- [3] M. Stojanovic, "On the relationship between capacity and distance in an underwater acoustic communication channel," in *Proc. ACM Int. Workshop on Underwater Networks*, 2006.
- [4] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke, "Data collection, storage, and retrieval with an underwater sensor network," in *Proc. Int. Conf. Embedded Networked Sensor Systems*, 2005.
- [5] D. L. Applegate, R. E. Bixby, V. Chvátal, and W. J. Cook, *The Traveling Salesman Problem: A Computational Study*. Princeton Univ. Press, 2006.
- [6] E. Balas, "The Prize Collecting Traveling Salesman Problem," *Networks*, vol. 19, pp. 621–636, 1989.
- [7] G. Ausiello, V. Bonifaci, S. Leonardi, and A. Marchetti-Spaccamela, "Prize-Collecting Traveling Salesman and related problems," in *Handbook of Approximation Algorithms and Metaheuristics*, T. F. Gonzalez, Ed. CRC Press, 2007.
- [8] G. Ausiello, V. Bonifaci, and L. Laura, "The Online Prize-Collecting Traveling Salesman Problem," *Information Processing Letters*, vol. 107, no. 6, pp. 199–204, 2008.
- [9] M. Goemans and D. P. Williamson, "A general approximation technique for constrained forest problems," *SIAM J. Computing*, vol. 24, no. 2, pp. 296–317, 1995.
- [10] A. Archer, M. H. Bateni, M. T. Hajiaghayi, and H. Karloff, "Improved approximation algorithms for Prize-Collecting Steiner Tree and TSP," in *Proc. IEEE Symp. Foundations of Computer Science*, 2009.
- [11] A. Dumitrescu and J. Mitchell, "Approximation algorithms for TSP with neighborhoods in the plane," *J. Algorithms*, vol. 48, no. 1, pp. 135–159, 2003.
- [12] M. de Berga, J. Gudmundsson, M. J. Katz, C. Levcopoulos, M. H. Overmars, and A. F. van der Stappen, "TSP with neighborhoods of varying size," *J. Algorithms*, vol. 57, no. 1, pp. 22–36, 2005.
- [13] D. Bhaduria and V. Isler, "Data gathering tours for mobile robots," in *IEEE Int. Conf. Intelligent Robots and Systems*, 2009.
- [14] F. Arrichiello, D. N. Liu, S. Yerramalli, A. Pereira, J. Das, U. Mitra, and G. S. Sukhatme, "Effects of underwater communication constraints on the control of marine robot teams," in *Proc. Int. Conf. Robot Communication and Coordination*, 2009.
- [15] A. Krause and C. Guestrin, "Near-optimal nonmyopic value of information in graphical models," in *Proc. Uncertainty in Artificial Intelligence*, 2005.
- [16] A. Krause, C. Guestrin, A. Gupta, and J. Kleinberg, "Near-optimal sensor placements: Maximizing information while minimizing communication cost," in *Proc. Information Processing in Sensor Networks*, 2006.
- [17] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Optimization*. Wiley, 1994.
- [18] L. Kaelbling, M. Littman, and A. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, vol. 101, pp. 99–134, 1998.
- [19] V. V. Vazirani, *Approximation Algorithms*. Berlin, Germany: Springer-Verlag, 2001.
- [20] H. Kaplan, M. Lewenstein, N. Shafir, and M. Sviridenko, "Approximation algorithms for asymmetric TSP by decomposing directed regular multidigraphs," in *Proc. IEEE Symp. Foundations of Computer Science*, 2003.