SCIENCE AND TECHNOLOGY ORGANIZATION
CENTRE FOR MARITIME RESEARCH AND EXPERIMENTATION

Reprint Series **CMRE-PR-2014-008**

Efficient dense sonar surveys with an autonomous underwater vehicle

David P. Williams, Michel Couillard

January 2014

Originally published in:

Proceedings of the 11th European Conference on Underwater Acoustics, Edinburgh, July 2012, pp. 676-683.

Proceedings of Meetings on Acoustics, Vol. 17, 070083 (2013).

About CMRE

The Centre for Maritime Research and Experimentation (CMRE) is a world-class NATO scientific research and experimentation facility located in La Spezia, Italy.

The CMRE was established by the North Atlantic Council on 1 July 2012 as part of the NATO Science & Technology Organization. The CMRE and its predecessors have served NATO for over 50 years as the SACLANT Anti-Submarine Warfare Centre, SACLANT Undersea Research Centre, NATO Undersea Research Centre (NURC) and now as part of the Science & Technology Organization.

CMRE conducts state-of-the-art scientific research and experimentation ranging from concept development to prototype demonstration in an operational environment and has produced leaders in ocean science, modelling and simulation, acoustics and other disciplines, as well as producing critical results and understanding that have been built into the operational concepts of NATO and the nations.

CMRE conducts hands-on scientific and engineering research for the direct benefit of its NATO Customers. It operates two research vessels that enable science and technology solutions to be explored and exploited at sea. The largest of these vessels, the NRV Alliance, is a global class vessel that is acoustically extremely quiet.

CMRE is a leading example of enabling nations to work more effectively and efficiently together by prioritizing national needs, focusing on research and technology challenges, both in and out of the maritime environment, through the collective Power of its world-class scientists, engineers, and specialized laboratories in collaboration with the many partners in and out of the scientific domain.

NOTE: The CMRE Reprint series reprints papers and articles published by CMRE authors in the open literature as an effort to widely disseminate CMRE products. Users are encouraged to cite the original article where possible.

Proceedings of Meetings on Acoustics

Volume 17, 2012 http://acousticalsociety.org/

ECUA 2012 11th European Conference on Underwater Acoustics
Edinburgh, Scotland
2 - 6 July 2012
Session UW: Underwater Acoustics

UW228. Efficient dense sonar surveys with an autonomous underwater vehicle

David Williams* and Michel Couillard

***Corresponding author's address: Science & Technology Organization -- Centre for Maritime Research and Experimentation, La Spezia, 19126, SP, Italy, williams@nurc.nato.int**

An algorithm for the in situ adaptation of the survey route of an autonomous underwater vehicle (AUV) equipped with side-looking sonars was recently proposed. This algorithm immediately exploits the through-the-sensor data that is collected during a mission in order to ensure that quality data is collected everywhere in the area of interest. By introducing flexibility into the survey of the AUV, various limitations of pre-planned surveys are overcome. In particular, the need to re-deploy the AUV (to fill gaps in the data coverage) is obviated. In turn, the time and costs of the data-collection mission are significantly reduced. In this work, we improve the aforementioned algorithm by introducing an additional constraint to the survey track-selection process. This modification significantly increases the efficiency of a survey by further reducing both transit time and the overall number of tracks executed. In particular, the revised algorithm more closely approximates the optimal survey route that would be executed if perfect knowledge of the future sonar performance during the mission were known a priori.

Published by the Acoustical Society of America through the American Institute of Physics

Received 26 Sep 2012; published 6 Feb 2013

1 INTRODUCTION

An algorithm for the *in situ* adaptation of the survey route of an autonomous underwater vehicle (AUV) equipped with side-looking sonars was recently proposed in [1]. This algorithm immediately exploits the through-the-sensor data that is collected during a mission in order to ensure that quality data is collected everywhere in the area of interest. By introducing flexibility into the survey of the AUV, various limitations of pre-planned surveys are overcome. In particular, the need to re-deploy the AUV (to fill gaps in the data coverage) is obviated. In turn, the time and costs of the data-collection mission are significantly reduced.

In this work, we improve the algorithm in [1] by introducing an additional constraint on the objective function that determines the selection of survey tracks. This modification significantly increases the efficiency of a survey by further reducing both transit time and the overall number of tracks executed. In particular, the revised algorithm more closely approximates the optimal survey route that would be executed if perfect knowledge of the *future* sonar performance during the mission were known *a priori*.

The high-resolution imaging of underwater environments afforded by synthetic aperture sonar (SAS) [2] has proven useful in a wide range of applications. Thanks to breakthroughs in marine robot technology, the sonar data is now often collected by an AUV equipped with two side-looking sonars, one on the port side and one on the starboard side.

These sensors image in directions orthogonal to the direction in which the AUV travels. Because of the geometry of the problem, a “dead zone” — from directly below the AUV out to a certain range on either side — between the two sonar swaths will lack sonar coverage.

The standard AUV survey plan used in practice [3] consists of a series of equidistant parallel tracks — often referred to as a “lawnmower pattern” — such that the sonar swaths of consecutive tracks interleave, resulting in sonar coverage for the entire area of interest.

The implicit assumption made in this purely pre-planned approach is that quality sonar coverage will be achieved consistently to a certain range on either side of the AUV’s track-centerline. However, if this coverage is underestimated (so the area of seabed interrogated on consecutive tracks overlaps), the coverage rate decreases, causing more time and resources to be expended to survey the mission area. If this coverage is overestimated, no usable sonar data will be collected for certain portions of the mission area, resulting in gaps in coverage.

Designing a survey based on assumed coverage is a real problem because data collected at sea with an AUV is rarely flawless [4]. For example, sonar image quality at long range is often poor because of signal attenuation or excessive undesired vehicle motion (caused by currents) that can result in defocused imaging.

A rigid pre-planned survey route is suboptimal because it has no flexibility to adapt to the environmental variability encountered during the data collection process at sea if the *a priori* coverage assumptions that determine the set of tracks are not satisfied.

Previous work related to adaptive underwater path planning has not focused on ensuring the *quality* of the data being collected, as we do here. Instead, the main body of work has been interested in finding selected regions of particular interest to survey or routes that satisfy certain objectives, rather than performing an exhaustive search of an area. The work involving dense, exhaustive surveys of an area with an AUV has tended to focus on navigation correction and localization [5] or the fusion of multiple images into a mosaic of the seabed [6] for subsequent object detection tasks.

The remainder of this paper is organized as follows. The proposed strategy for performing adaptive data collection based on the quality of the through-the-sensor sonar data is outlined in Section 2. Experimental results of the algorithm are shown in Section 3. Concluding remarks and directions for future work are provided in Section 4.

2 ADAPTIVE SURVEY ALGORITHM

2.1 OVERVIEW

In [1], an algorithm was proposed to adaptively determine, *in situ*, the optimal placement of AUV survey tracks that would ensure quality sonar data was collected over the entire area of interest. The peak correlation of consecutive ping returns, as a function of range, was used to quantify the sonar image quality. To predict the area that would be covered with quality data by a hypothetical survey track, a full Bayesian model of this correlation quantity was continuously updated during the mission based on the ongoing performance of the sonars. The objective function that then drove the track selection was defined in terms of the expected gain in coverage and the transit cost of each potential track.

This work presents an improvement to the adaptive algorithm presented in [1]. Specifically, an additional constraint that improves survey efficiency is introduced in the objective function.

2.2 IMAGE QUALITY

A sonar image with good image quality is one in which objects and seabed features appear more or less as they would when observed under the ideal operating conditions of the sonar system.

Synthetic aperture sonar (SAS), which provides an order-of-magnitude improvement in resolution over simple (real aperture) side-scan sonar data, works by coherently summing received sonar signals, or ping-returns, of overlapping elements in the array. The peak correlation between successive pings at a given range (*i.e.,* distance from the sonar) provides a measure of the success of the SAS processing — and in turn, the image quality — as this quantity is directly related to the signal-to-noise ratio (SNR) [7].

Although full SAS processing is a computationally intensive procedure, the ping-to-ping correlation can be easily calculated onboard an AUV in real-time — with minimal computational power — as data is collected.

Analysis [8] of a detection algorithm on a large data set of sonar images containing 480 man-made targets indicated that the performance was noticeably affected by the image quality when the peak correlation was less than $\tau_\rho = 2/3$. Therefore, in this work, we deem data to be of insufficient quality below this threshold.

2.3 QUALITY COVERAGE

For a given track, let $\rho(r,u)$ be the peak correlation value between consecutive pings at range r on sonar side $u \in \{port, starboard\}$, averaged over all pings associated with the track.

Let $\nu(r,u)$ be the historical success rate of collecting quality sonar data during the mission. This success rate will essentially be a running tabulation of the proportion of tracks at each range that resulted in quality data (*i.e.,* $\rho(r,u) \geq \tau_\rho$) during the mission, and will be updated via a full Bayesian model after each track is executed.

The likelihood of coverage success (a binary variable) achieved during a track for a particular (r,u)-pair can be considered to follow a Bernoulli distribution with probability ν of success. The prior expressing our *a priori* beliefs about the sonar's performance is assumed to follow a Beta distribution, $\mathcal{B}(\alpha,\beta)$,

$$p(\nu) = \nu^{\alpha-1}(1-\nu)^{\beta-1}C(\alpha,\beta), \tag{1}$$

where $\alpha > 0$ and $\beta > 0$ are shape parameters and C is a normalization constant to ensure the total probability integrates to unity. Specifically, the prior we employ is $\nu(u,r) \sim \mathcal{B}(\alpha = 1, \beta = 1)$ for $r \in [r_{\min}, r_{\max}]$ (and otherwise, outside these valid ranges, 0), where $r_{\min}$ and $r_{\max}$ are the nominal operation limits of the sonar; this Beta prior is simply a uniform distribution over $(0,1)$.

Because the Beta prior is a *conjugate prior* for a Bernoulli likelihood, the posterior distribution over ν is another Beta distribution, with parameters $\alpha = Q+1$ and $\beta = T-Q+1$, where Q is the number of tracks that successfully achieved quality coverage at the given (r,u), out of T total tracks (*i.e.,* attempts).

That is, $Q(r,u) = \sum_{i=1}^{T} Q_i(r,u)$, where $Q_i(r,u)$ is a binary variable indicating whether the i-th of T tracks in the present mission resulted in quality coverage at range r from sonar side u (*i.e.,* whether $\rho(r,u) \geq \tau_\rho$ during the i-th track).

The expected value of a random variable ν following a Beta distribution with parameters α and β is

$$E[\nu] = \frac{\alpha}{\alpha+\beta}. \tag{2}$$

Therefore, the expected value of the posterior for successful quality coverage is

$$E[\nu(r,u)] = \frac{Q(r,u)+1}{T+2}. \tag{3}$$

This posterior expectation is used in the objective function calculation for determining the next track to execute.

By tying the expected future performance in the remainder of the mission to the historical *in situ* coverage performance of the sonars, the anticipated coverage by a hypothetical survey track can be more accurately predicted.

2.4 NEW TRACK SELECTION

Given an area of seabed to be surveyed, a series of parallel tracks will be executed over that area. The proposed algorithm adaptively determines, *in situ*, the specific locations of those tracks based on the through-the-sensor data collected.

The mission area is discretized into a set of non-overlapping swaths. Associated with each swath, s_i, is a binary coverage variable, $\gamma(s_i)$. (For a given track, there exists a unique mapping, based on simple geometry, from each (r,u)-pair to some particular s_i.) If $\rho(r,u) \geq \tau_\rho$ after executing a track, the corresponding swath coverage variable, $\gamma(s_i)$, is set to unity. Then the utility of executing each potential new track is calculated.

The utility of running a new track, t_i, having just executed track t_j, is simply the difference between the expected benefit and the cost,

$$U(t_i|t_j) = [B(t_i) - C(t_i|t_j)]\,\mathbf{1}_{B(t_i)>0}, \tag{4}$$

where the indicator function $\mathbf{1}_{B(t_i)>0}$ is included to ensure only tracks with positive benefits are considered.

The benefit of running a particular track, t_i, as the next track is defined to be the weighted sum of the expected coverage improvement,

$$B(t_i) = \sum_r \sum_u E[\nu(r,u)](1 - \gamma(r,u)), \tag{5}$$

where $\gamma(r,u)$ is the binary coverage quantity for the swath of seabed at range r on sonar side u, computed with respect to the centerline of track t_i; $E[\nu(r,u)]$ is obtained from (3).

The cost of running a particular track, t_i, is a function of the transit distance required to reach the new track from the vehicle's present location. Let $d(t_j, t_i)$ be the distance between the most-recently executed track, t_j, and a potential track t_i. Let ω denote twice the turning radius of the vehicle. The cost of executing track t_i, having just executed track t_j, is defined to be

$$C(t_i|t_j) = (d(t_i,t_j)/\omega)^2 - (2d(t_i,t_j)/\omega) + 1, \tag{6}$$

which is the equation of a parabola with a vertex at $(d, C) = (\omega, 0)$, indicating that the transit cost is a minimum (namely, zero) when two tracks are separated by twice the turning radius of the vehicle.

In [1], the track t_i for which the utility in (4) is maximized was then selected as the next track to be executed. In this work, we introduce a new additional constraint that can significantly alter the track selection process.

2.5 PROPOSED NEW CONSTRAINT

In [1], one constraint on the utility function was included for determining the next track to select. Only those potential tracks for which the benefit, B, was positive (non-zero) were considered. This constraint ensured that every track selected contributes toward covering the mission area. (Otherwise, a case could arise in which a track that made no progress in covering new swaths of seabed but possessed a low transit cost had the highest overall utility of any track.)

In this work, we introduce a new stronger constraint that renders the previous constraint superfluous. The new constraint is that only those tracks that would have a non-zero probability above some threshold, τ_c, of covering the "first" uncovered swath (*i.e.,* assuming the AUV begins the mission on the left side of the survey area, the left-most uncovered swath) are considered.

With this constraint, the mission progresses in such a way that the subsequent track is selected so that a gap in coverage would be immediately filled. Therefore, it effectively eliminates the need to travel back across the survey area to fill gaps after an initial pass. The tradeoff for this behavior, which reduces transit time, is that the coverage rate does not increase as quickly. However, because we are interested in *dense* surveys in which complete coverage must be achieved – rather than maximizing the coverage rate in an exploratory mission – this constraint is justified and worthwhile.

The overall utility of running a new track, t_i, having just executed track t_j, is then simply the difference between the expected benefit and the weighted cost,

$$U(t_i|t_j) = [B(t_i) - C(t_i|t_j)]\,\mathbf{1}_{p(\alpha(s')=1|t_i)>\tau_c}, \tag{7}$$

where the indicator function $\mathbf{1}_{p(\alpha(s')=1|t_i)>\tau_c}$ is included to ensure only tracks that have non-zero probability of covering the "first" uncovered swath, denoted s', above some threshold, τ_c, are considered.

The track t_i for which the utility is maximized is then selected as the next track to be executed. This entire process is repeated until $\alpha(s_i) = 1\ \forall s_i$ (*i.e.,* all coverage variables are unity), indicating that quality sonar data has been collected for the entire mission area.

3 EXPERIMENTAL RESULTS

3.1 EXPERIMENTAL SET-UP

The NATO Undersea Research Centre (NURC) possesses an AUV called MUSCLE that is equipped with a 300 kHz sonar with a 60 kHz bandwidth. The proposed adaptive track-spacing algorithm is intended for use on the MUSCLE AUV at sea. In standard mode operation, the AUV collects sonar data from $r_{\min} = 40m$ to $r_{\max} = 150m$ away from the AUV in both the port and starboard directions.

Experiments were conducted for a simulated, rectangular mission area that was $2km \times 5km$. The maximum range at which quality sonar data was assumed to be achieved, r_{max}, for each track was treated as fixed but unknown *a priori*.

We present experimental results for three cases. In the first case, it is assumed that quality sonar coverage is always achieved from $r_{\min} = 40m$ to $r_{\max} = 70m$ away from the AUV in both the port and starboard directions. In the second case, $r_{\max} = 110m$ instead for both the port and starboard sonars. The third case simulates the scenario in which the starboard sonar has malfunctioned (so no useable sonar data is collected by it), but the port sonar is functioning perfectly with $r^p_{\max} = 150m$. For each of these three cases, the survey performance of several strategies were compared.

It has been shown in [9] that the optimal placement of survey tracks can be derived analytically given knowledge of the sonar swath limits, $r_{\min}$ and $r_{\max}$. Therefore, one method tested follows a pre-planned survey in which the track locations are given by this optimal placement, assuming sonar swath coverage is always achieved from $r_{\min} = 40m$ to $r_{\max} = 150m$ away from the AUV in both the port and starboard directions. That is, the pre-planned tracks are designed with minimal overlap, fully exploiting the interleaving nature of the swath coverage on consecutive tracks. A second method follows a pre-planned survey in which the track locations are selected when assuming perfect knowledge of the sonar performance is known *a priori*. Therefore, this optimal method is by construction the most efficient, and it provides a bound on achievable performance.

Additionally, three adaptive track-selection methods based on the algorithm in Section 2 are considered. Each differs in the constraint that is used with the objective function. The first of these adaptive methods represents the algorithm presented in [1], which does not employ an additional constraint (beyond that of a non-negative benefit); this approach is denoted “no constraint.” The second method requires the probability of covering the most extreme uncovered swath to be non-zero, so $\tau_c = 0$; this approach is denoted “mild constraint.” The third method requires the probability of covering the most extreme uncovered swath to be greater than $\tau_c = 0.5$; this approach is denoted “strong constraint” and represents the modified algorithm proposed in this work.

3.2 ANALYSIS

In the first set of experiments, $r_{\max} = 70m$, meaning the realized sonar coverage is significantly smaller than anticipated. (Such a scenario can be caused, for example, by currents that introduce undesired vehicle instability and motion.)

The locations of the tracks selected by each of the aforementioned methods are shown in Figure 1. The progression of the fraction of the mission area for which quality data was collected, as a function of the number of tracks executed by each method, is shown in Figure 2.

Because of poor quality data over a significant portion of the sonar’s operating range, the pre-planned survey achieved quality coverage over only 22% of the mission area. In contrast, the adaptive survey approaches adjust their tracks based on the data collected *in situ* to ensure quality coverage is obtained over 100% of the mission area.

Of particular interest in Figure 1 is the effect of the constraint on the adaptive methods. With no constraint, the adaptive method attempts to maximize the area coverage with each track via a purely myopic strategy. As a result, gaps in coverage (due to poor data quality) are created during the first

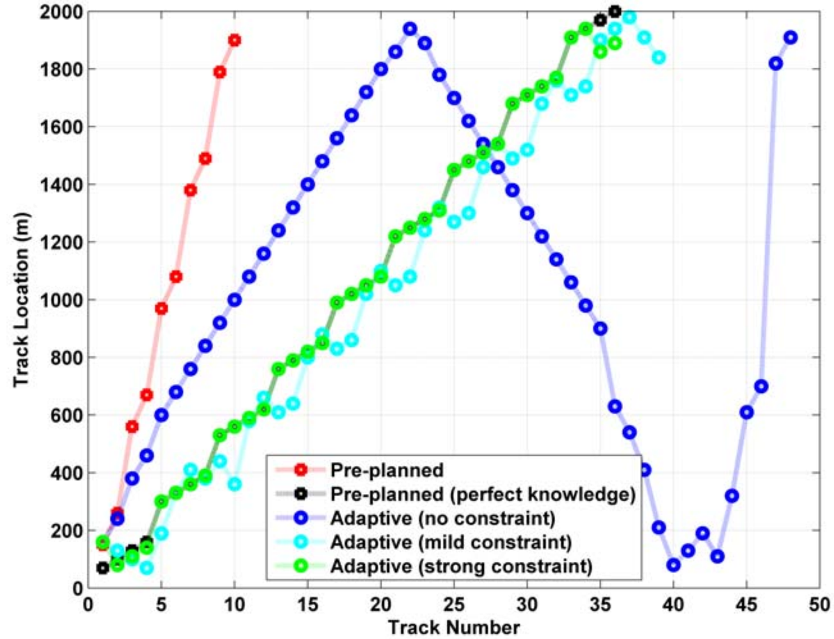

Figure 1: Locations of the tracks selected by each approach, when the true maximum range for quality data was $r_{\max} = 70m$.

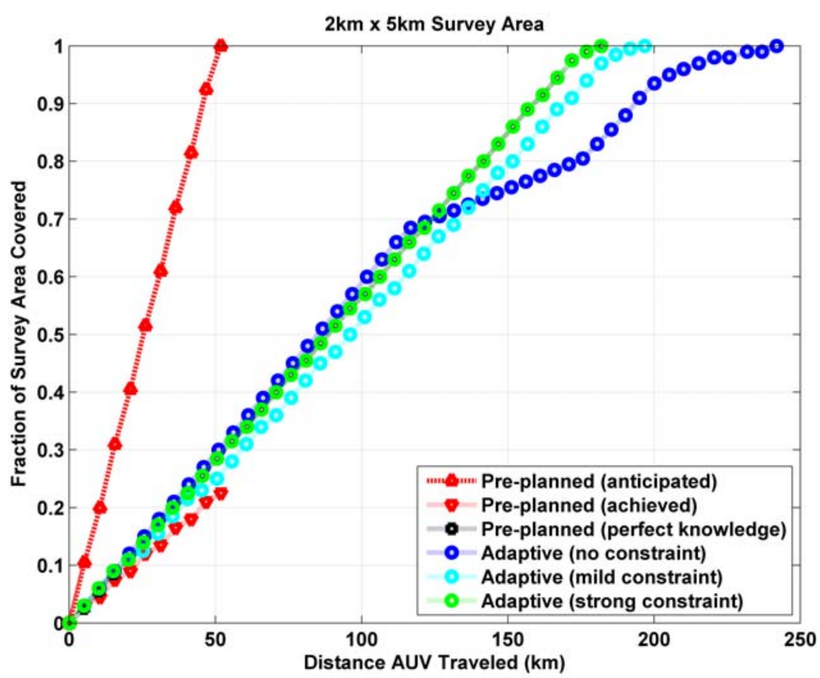

Figure 2: Progression of the fraction of the mission area for which quality data was collected by each approach, when the true maximum range for quality data was $r_{\max} = 70m$. Each marker on the plot indicates the coverage accumulated at the end of a track.

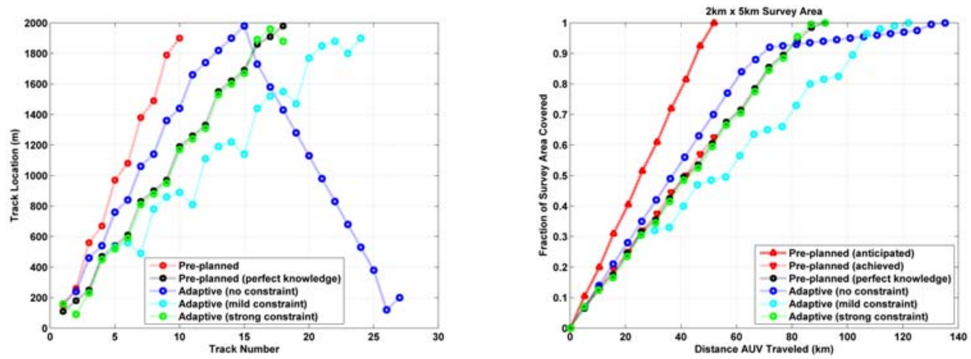

(a) Track locations (b) Area coverage

Figure 3: When the true maximum range for quality data is $r_{\max} = 110m$, (a) the locations of the tracks selected, and (b) the progression of the fraction of the mission area for which quality data was collected.

pass across the mission area. A second pass back across the mission area is required to fill in these gaps (and in fact, in this particular scenario, a third pass was needed to eliminate some remaining gaps in coverage). Even though this approach achieves high coverage rapidly in the early portion of the survey (*cf.* Figure 2), the additional transit time and track execution ultimately result in an inefficient survey.

In contrast, it can be observed that the inclusion of the proposed constraint related to the most extreme uncovered swath reduces the overall transit time and distance. In the case of the mild constraint, a weak attempt is made to ensure that the next track would fill coverage gaps as the mission progresses. However, because the uncovered swath must simply have a non-zero probability – no matter how small – of being covered by a track, there is little guarantee that a track will actually successfully fill the coverage gap. As a result, multiple tracks are sometimes required to fill a gap that in theory could have been covered by a single, more cautiously placed track (*i.e.,* one that does not ambitiously assume quality data will be collected over the full nominal operating range of the sonar).

The strong constraint, which requires a track to have a minimum probability of covering the most extreme uncovered swath, is more efficient at filling coverage gaps. As can be seen in Figure 2, this method achieves complete coverage using virtually the same travel distance that would be required if perfect *a priori* knowledge was possessed. In fact, the proposed approach is nearly optimal, as it requires the same minimum number of tracks that the optimal case would employ.

In the second set of experiments, $r_{\max} = 110m$, meaning the realized sonar coverage is significantly larger than in the first scenario, though still not perfect. The locations of the tracks selected by each method and the progression of the fraction of the mission area for which quality data was collected, as a function of the number of tracks executed by each method, are shown in Figure 3. Similar trends to those observed in the first case are also observed here.

Finally, the third set of experiments provides insight into the coverage performance when one sonar is malfunctioning. The locations of the tracks selected by each method and the progression of the fraction of the mission area for which quality data was collected, as a function of the number of tracks executed by each method, are shown in Figure 4 for this scenario. With a purely pre-planned approach, roughly half of the mission area would be left uncovered, as the AUV would be oblivious to the malfunctioning sonar. But with an adaptive approach, the AUV can immediately adjust its survey route to compensate for the mechanical failure and still achieve complete, dense coverage. In this scenario, it can again be seen how the proposed constraint reduces the number of tracks required in the survey.

Document Data Sheet

Security Classification		Project No.
Document Serial No. CMRE-PR-2014-008	**Date of Issue** January 2014	**Total Pages** 8 pp.

Author(s)

Williams, D.P., Couillard, M.

Title

Efficient dense sonar surveys with an autonomous underwater vehicle.

Abstract

An algorithm for the in situ adaptation of the survey route of an autonomous underwater vehicle (AUV) equipped with side-looking sonars was recently proposed. This algorithm immediately exploits the through-the-sensor data that is collected during a mission in order to ensure that quality data is collected everywhere in the area of interest. By introducing flexibility into the survey of the AUV, various limitations of pre-planned surveys are overcome. In particular, the need to re-deploy the AUV (to fill gaps in the data coverage) is obviated. In turn, the time and costs of the data-collection mission are significantly reduced. In this work, we improve the aforementioned algorithm by introducing an additional constraint to the survey track-selection process. This modification significantly increases the efficiency of a survey by further reducing both transit time and the overall number of tracks executed. In particular, the revised algorithm more closely approximates the optimal survey route that would be executed if perfect knowledge of the future sonar performance during the mission were known a priori.

Keywords

Issuing Organization

Science and Technology Organization
Centre for Maritime Research and Experimentation

Viale San Bartolomeo 400, 19126 La Spezia, Italy

[From N. America:
STO CMRE
Unit 31318, Box 19, APO AE 09613-1318]

Tel: +39 0187 527 361
Fax:+39 0187 527 700

E-mail: library@cmre.nato.int