SCIENCE AND TECHNOLOGY ORGANIZATION
CENTRE FOR MARITIME RESEARCH AND EXPERIMENTATION

Reprint Series CMRE-PR-2019-049

An autonomous underwater vehicle data-driven control strategy for target tracking

Gabriele Ferri, Andrea Munafò, Kevin D. LePage

May 2019

Originally published in:

IEEE Journal of Oceanic Engineering, vol. 43, issue no. 2, April 2018, pp. 323-343, doi: 10.1109/JOE.2018.2797558

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.

Copyright © IEEE, 2018. NATO member nations have unlimited rights to use, modify, reproduce, release, perform, display or disclose these materials, and to authorize others to do so for government purposes. Any reproductions marked with this legend must also reproduce these markings. All other rights and uses except those permitted by copyright law are reserved by the copyright owner.

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.

An Autonomous Underwater Vehicle Data-Driven Control Strategy for Target Tracking

Gabriele Ferri, *Member, IEEE*, Andrea Munafò, *Member, IEEE*, and Kevin D. LePage

***Abstract*—This paper presents a data-driven approach to control the movement of autonomous underwater vehicles (AUVs) operating as receivers of a multistatic sonar surveillance network. The algorithm adopts a nonmyopic receding horizon policy to control the AUV heading, with the objective of minimizing the expected estimate error in the target location produced by the onboard tracker. Minimizing this error is of the utmost interest in target state estimation since it is one way of maintaining tracks. The proposed policy builds on an automated perception module which produces a target track and on an acoustic model to estimate the target measurement signal-to-noise ratio. Driven by the mission objectives, it explores the best decisions for the AUV by evaluating the evolution of the tactical situation over a future time window. Nonmyopic approaches can offer increased performance if compared to myopic ones. The downside is an increase of the computational burden. Methods are described to simplify and solve effectively the resulting decision tree to allow the execution of these kinds of computational intensive algorithms onboard autonomous vehicles. At the same time, they aim at producing AUV paths robust to possible misleading target measurements, which might cause meandering trajectories. These are crucial factors to use data-driven strategies for AUVs which must work in real applications. Results from the COLLAB13 sea trial are reported and discussed. They show both the feasibility of executing the algorithm in realtime on an onboard computer and the benefits of using the proposed approach over conventional predefined paths. These results are, to the best of our knowledge, the first successful demonstration at sea of a complex nonmyopic algorithm running in realtime onboard AUVs in a realistic multistatic littoral surveillance scenario.**

***Index Terms*—Antisubmarine warfare (ASW), autonomous underwater vehicle (AUV), autonomy, data-driven strategies, nonmyopic control, optimal control, surveillance.**

I. Introduction

UNDERWATER surveillance consists in detecting, localizing, and classifying targets in an area of interest using sensors of different nature (fixed or mobile) [1].

Historically, one of the most studied underwater surveillance application is antisubmarine warfare (ASW) [2]. In a typical ASW scenario, a mobile target (a submarine), which manoeuvres to avoid the detection, must be localized. The ASW scenario can be considered as a representative case study for a wider range of underwater surveillance applications. ASW is a capital asset-intensive activity and has been carried out using fixed sensors such as sonobuoys [3], and/or using mobile assets such as marine patrol aircrafts and submarines or frigates with towed arrays [3], [4]. The ASW final objective is to infer from the large amount of collected data if a target is present in the area and to track it for its correct classification.

Recent advances in marine robotics have made small and low-cost AUVs a reality [5]–[8].

These results suggest that maritime unmanned systems can become part of underwater surveillance systems since they can guarantee the persistent monitoring of an area at lower costs than traditional assets and offer the possibility of complementing or even substituting current solutions [1], [9], [10]. Compared to traditional assets, these small, low power, and mobile units have usually limited processing and communication capabilities. However, when deployed in a spatially separated network, they can be interconnected to form an intelligent network characterized by scalability, robustness, reliability, and adaptability. Mobile units build upon collected data and use their mobility to adapt their mission to the changing environmental conditions and to the evolving tactical scene. This can improve the network performance in terms of target tracking/classification and of an increased area coverage (higher temporal and spatial resolution) in area search/patrol tasks.

A. NATO Science and Technology Organization CMRE Multistatic Antisubmarine Warfare Robotic Network

Pursuing this concept, the Centre for Maritime Research and Experimentation (CMRE) has been developing a hybrid multistatic robotic autonomous network for an ASW underwater surveillance application (see Fig. 1). Mobile gateways, such as the WaveGliders [11] (see Fig. 1), which exploit the wave energy to move, and fixed gateway buoys build the communications infrastructure of the network. AUVs towing a hydrophone array are mobile underwater platforms capable of detecting and tracking possible threats and communicating the acquired data to the command and control (C2) center. In this way, the support ship can stay far from the patrolled area increasing its safety, or doing different activities.

In CMRE's multistatic active sonar system (an example of bistatic geometry is shown in Fig. 2), one or more sonar sources

Manuscript received February 15, 2017; revised August 9, 2017 and November 17, 2017; accepted January 19, 2018. Date of publication February 14, 2018; date of current version April 12, 2018. This work was supported by the NATO Allied Command Transformation. *(Corresponding author: Gabriele Ferri.)*

Associate Editor: T. Maki.

G. Ferri and K. D. LePage are with the Research Department, NATO STO-Centre for Maritime Research and Experimentation, La Spezia 19267, Italy (e-mail: Gabriele.Ferri@cmre.nato.int; Kevin.LePage@cmre.nato.int).

A. Munafò was with the Centre for Maritime Research and Experimentation, La Spezia 19267, Italy. He is now with the Marine Autonomous and Robotic Systems, National Oceanography Centre, University of Southampton, Southampton SO14 3ZH, U.K. (e-mail: andmun@noc.ac.uk).

Digital Object Identifier 10.1109/JOE.2018.2797558

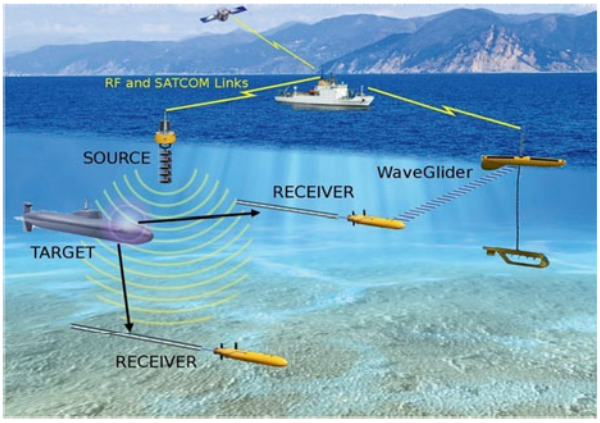

Fig. 1. CMRE cooperative ASW multistatic network concept. One or more sonar sources, which may be located on a stationary buoy or be ship deployable, transmit a sonar signal which reflects from objects and is collected by multistatic receivers, the Ocean Explorer (OEX) AUVs towing an hydrophone array. Mobile gateways, the WaveGlider autonomous surface vehicle (ASV), which exploit the wave energy for moving, and fixed gateway buoys build the communications infrastructure.

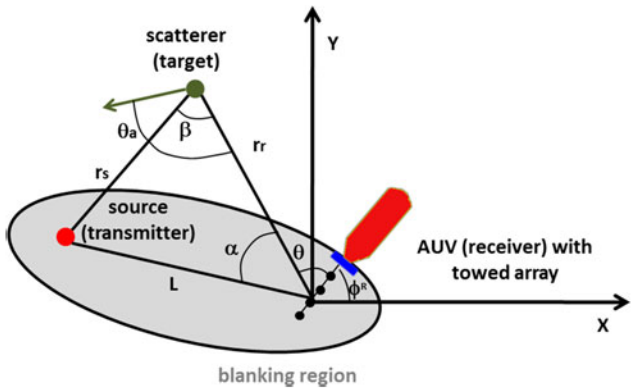

Fig. 2. Bistatic source–target–receiver geometry for a single source–receiver pair has shown the quantities of interest. The blanking region is also indicated. Inside this ellipse (with source and receiver as foci), targets are undetectable due to the direct blast of the transmitted waveform. We refer to *broadside* with a bearing angle θ close to 90°. We refer to *endfire* for θ angles close to 0° or 180°. The quality of contacts is highly related to the target bearing relative to the hydrophone array [16]. A target at broadside results in the highest bearing resolution, while a target at endfire will be poorly resolved.

(transmitters), either positioned on a stationary buoy or deployed from a ship, transmit a sonar signal (ping) which when reflected off objects can be collected by the multistatic receivers (the AUVs). The use of active sonar is dictated by today's "quiet" diesel–electric submarines [3], which make passive detection more difficult. Multistatic sonar systems have the potential to significantly increase the ASW coverage and performance [4]. The possibility of using several sources and receivers generates different geometric distributions of source–target–receiver. This increases the probability of detection (PD), tracking, and classification of a target [12]. Furthermore, with the proposed concept, the AUVs do not carry a source, which is beneficial for both increasing their endurance and their covertness (the AUV does not emit any sonar ping).

When the AUV receives an echo from the target, it builds a bearing-range contact that is fed into an onboard tracker [13]–[15]. The tracker combines spatially related contacts over time to produce tracks [14], which provide an estimate of the position of the target as well as of its velocity. In the addressed scenario, the quality of contacts is highly related to the target bearing relative to the hydrophone array (see Fig. 2).

B. Challenges

The underwater scenario, in particular, the littoral one, where the network is usually deployed, is challenging for several reasons. Sound propagation in shallow waters is generally characterized by significant multipath and poor propagation conditions. High level of clutter is typically present together with noise produced by passing ships and boats. The presence of "ghost" tracks due to the port-starboard ambiguity, typical of linear arrays,[1] exacerbates the problem. All this complicates target tracking [4].

Another challenge is represented by the limitations of underwater communications [1], [17], which do not allow a reliable exchange of information between the network nodes and between vehicles and the C2 center.

In these conditions, the autonomy of the vehicles becomes a key aspect. Decision making (e.g., if patrolling an area or prosecuting some tracks [1], [18], [19]) is essential to react to the evolving tactical situation and to increase the efficiency of the mission. Moreover, it is necessary when the vehicles cannot have a reliable communication with the C2 center.

To better meet mission objectives, the vehicles can make decisions to modify their trajectories based on the information collected in realtime, for instance, to increase the detection, classification, localization and tracking (DCLT) performance. These kinds of approaches, in which the behavior of the robot changes autonomously on the basis of the collected data, are called data driven [1].

In this paper, we focus on the latter aspect and we investigate how the AUVs can modify their paths to improve the tracking performance of a mobile target on the basis of the thus far collected sonar measurements.

C. Autonomy as a Key Driver: Prefixed Tracklines Versus Data-Driven Strategies

Traditionally, AUVs have used predesigned tracklines to conduct their work [20]. Predesigned tracklines offer an attractive solution for the deployment of robotics assets, since this guarantees the predictability of the AUV's operations. With this approach, the human operator still maintains most of the control since changes in the vehicle's behavior are only triggered by messages from the C2 center. A representative example can be found in [21], where the different phases of a mine countermeasures mission are performed sequentially with highly specialized robots covering pre-designed paths.

The increase in battery endurance associated with an ever-growing computational power has made possible for underwater robots to perform complex realtime signal processing and

[1]For linear array, once an echo is received, the beamforming [13] provides two possible directions of arrival since it is not possible to discriminate if the echo arrived from port or starbord. This leads to consider also false (ambiguous) measurements that can generate false (or "ghost") tracks.

decision making, with a resulting increase in system flexibility and adaptability, multiplying their data collection effectiveness. AUV-based data-driven approaches, as well as autonomous decision making [20], are becoming a reality [22], [23]. In data-driven approaches, the AUV modifies its path based on the data collected during the mission to meet some objectives. Data-driven policies are critical in communication limited environment when the robot need to react to events and a reliable link with C2 cannot be assured. However, adaptive approaches do not assure the uniform coverage of the area of interest guaranteed by prefixed tracklines. Mixed solutions trying to achieve a tradeoff between prefixed paths and data-driven trajectories can also be a viable solution as shown in [8].

Work carried out in different research communities such as optimal sampling [24], [25] in oceanographic applications or sensor management [26] has investigated how to deploy sensors and control nodes (robots) to optimize metrics of interest for the problem under investigation.

The problem, in general, can be formulated as a stochastic control problem in which the degrees of freedom of a sensor system (e.g., the path of an autonomous sensorized vehicle) are controlled to achieve some operational objectives quantified by a suitable cost function. An optimal or suboptimal policy is sought to achieve a desired sensor configuration on the basis of the available information from prior measurements and models of the environment, subject to the constraints of the problem under investigation.

Cost functions, in general, can be composed of a deterministic and a stochastic component. The deterministic part can include, for instance, the energy cost for a given movement and sensing action or bandwidth costs during the communication process. Stochastic costs can include the predicted tracking accuracy or the predicted entropy of a probability distribution of interest (e.g., target location) and/or the predicted gain in information [27].

Two approaches to sensor management can be defined: task- or mission-driven and information-driven sensor management. The former chooses sensor actions based on a given performance metric or error quantity directly related to the mission objectives. Mission-driven performance metrics for the problem of AUVs and multistatic active sonar include, for instance, maximizing the probability of detecting a target (if present) given some *a priori* target probability distribution or maximizing the PD or track formation over multiple platforms [1].

Alternative solutions for the stochastic component of the cost function are represented by the use of some norms of the predicted posterior covariance of a probability distribution of interest given the adopted sensor policy [24], [25], [28], [29].

Information theoretic costs have also been considered to design the objective functions. Information-driven sensor management chooses sensor placements and actions which maximize a measure of the information gain, or some function of the probability density function conditioned on the observations $\mathbf{z}$ and sensing actions $\mathbf{u}$, $p(\mathbf{x}|\mathbf{z}, \mathbf{u})$ [30], [31].

Several functions of the Fisher information matrix, such as its determinant and trace, have been adopted in different applications [26]. For instance, in [32], the posterior Cramér–Rao lower bound was used for multisensor scheduling, while in [33] the Fisher information gain is the basis for a sensor effectiveness metric driving the sensor assignment in a multisensor and multitarget tracking application.

Other information measures such as entropy and mutual information have become more common in the research community in the past years [26]. In [34] and [35], the objective of the optimization was the expected update in posterior entropy, called the information gain. In contrast to covariance-based objective functions, entropy quantifies areas of probabilities and not the average deviation to a single point. It is a measure of how much additional information is needed to infer the exact value from an estimate [36]. The information theoretic costs predict an average distance between the approximate predicted and filtered state densities for each sensing policy. The distance can be based on the Kullback–Leibler or the Rényi divergence between the prior and posterior distributions [26]. Another approach is to maximize mutual information between the state at the current time and the new observation [37]. Use of information theoretic costs to manage sensor devices (and/or autonomous vehicles) has been applied to different models and sensor management scenarios. Examples can be found in [38]–[41].

Independently of the type of the adopted cost function, one key issue is the planning horizon considered in the optimization [42], [43]. The sensor actions can in fact better achieve the overall objective by considering the evolution of the tactical scene in the future. Some actions that do not lead to a decrease in the cost function at the next step of prediction can be beneficial by considering a longer future horizon. Methods that take into consideration multiple steps into the future are called nonmyopic, in contrast with the approaches that try to optimize the cost function by looking only at one prediction step that are called myopic or greedy. Nonmyopic policies have been addressed with different theoretical frameworks, such as Markov decision process or partially observed MDP [44]. The resulting problem can be in theory solved optimally by using dynamic programming techniques [42]. Nevertheless, with the increase of the planning time horizon, the problem becomes prohibitive from the computation point of view, especially if we think to the limited resources generally available on a robot. Approximate solutions such as rollout or model predictive control [42], even if suboptimal, are used to find a tradeoff between results and required processing resources.

Myopic sensor management [26] policies have been investigated as low complexity alternatives to multistage policies. They only look ahead to the next stage to compute the expected utility in the immediate future to determine the best current action. Such greedy policies benefit from computational simplicity, at the expense of performance loss compared to multistage optimal policies. Often this loss is significant. However, there are cases where the myopic approach gives acceptable performance, and indeed is almost optimal in special cases.

One possible approach in the limited look-ahead control schemes is the so-called receding horizon (also called model predictive schemes) [42], [43] algorithms. Receding horizon approaches are cases of limited look-ahead control, in which at each decision step the optimal configuration sequence is

computed over a moving horizon. The crucial point in limited look-ahead approaches is the selection of the number of future time steps. Two competing mechanisms influence the choice of the time horizon: the exploration of the substructures of the decision spaces that may yield larger benefit (increasing the horizon), and the limitation of the effects of less accurate estimates in the control feedback loop (reducing the horizon). Therefore, increasing the size of the look-ahead does not always lead to superior performance [43] and may depend on the particular application under investigation.

As discussed, there exists quite a rich literature on sensor management. However, it must be underlined that not much work has addressed the use of AUVs for underwater surveillance and ASW, especially using active sonars. The interested reader is referred to [45]–[48] for recent work on passive sonars. Even fewer results are reported with experimental data from realistic ASW trials. This is due to the difficulties in operating AUVs in credible ASW operational scenarios, since this entails the robust integration of several enabling technologies into a system running flawlessly onboard the vehicles. Furthermore, data-driven strategies must be designed in a way that is robust to noisy measurements to be tested at sea. This is crucial to exploit the advantages that this paradigm can offer and makes porting algorithms from a simulation environment to a real application challenging and nontrivial.

In this respect, some early work using AUVs as active sonar receivers was done at CMRE, which demonstrated the feasibility of sonar realtime processing and decision making onboard AUVs [13]. Since then, more complex strategies have been developed to optimize different objective functions, and have been often tested at sea during ASW trials.

In [14], an AUV used a greedy approach to minimize the one-step predicted localization error of a target with the aim to improve the tracking performance of the onboard multiple hypothesis tracking tracker. Even if there are many practical cases in which myopic approaches give acceptable performance, nonmyopic policies have usually superior performance. One interesting example is the one reported in [49], in which simulation results demonstrate how to control the AUV heading over a future horizon to optimize the expected received signal-to-noise ratio (SNR).

A nonmyopic policy to optimize the AUV navigation based on a target track in a multistatic ASW scenario was presented in [19] and [50]. The solution minimizes the trace of the Cramér–Rao lower bound matrix computed at M prediction steps into the future (terminal scheduler cost). At each decision step, M heading decisions are computed and are executed before a new calculation is computed. This approach has the appeal of a low computational burden since the computation is carried out every M pings, but it may suffer from the occasional misleading measurements. In such a case, the AUV would navigate for M pings while optimizing an incorrect target trajectory, possibly compromising the mission since the track may break. Results are evaluated in simulation and no results of the algorithm running onboard the vehicles are reported.

A different approach, which operated directly at the contact level, is the one presented in [51]. In this case, the AUVs optimize the predicted PD of the target computed as a function of the hypothesized target position, receiver trajectory, and range-dependent environmental parameters. In [52], the approach was extended to control the AUV movement along the water column to reduce the uncertainty in the localization of the target depth.

D. *Contribution of this Paper*

In this paper, we describe a nonmyopic receding horizon strategy to control the heading of an AUV which tows a line array and that acts as a sonar receiver node of the CMRE's ASW network. The proposed algorithm, starting from a selected track, controls the AUV heading angle to minimize the expected target position estimation error of a tracking filter. To compute the expected error, an acoustic model [53] is used together with the computation of bistatic contact-localization statistics. Minimizing this error is typically of the utmost interest in target state estimation since it is one way of maintaining track.

The algorithm builds a decision tree to consider the possible future target–receiver–source configurations and the optimal sequence is sought over the tree. Since the resulting decision tree may present large dimensions (it is exponential in the number of heading decisions), we specifically aim at simplifying the tree and to solve it efficiently. At the same time, we introduce methods to avoid meandering AUV paths caused by possible misleading measurements (i.e., outliers) that would lead to a decrease in the acoustic signal processing performance. These are the critical features for data-driven policies that aim at being deployed in real-world applications.

We start describing the proposed algorithm in Section II. Then, we discuss the solutions used to solve the optimization problem in realtime onboard an AUV within the required time constraints. In Section III, after describing the CMRE multistatic network deployed in our trials, we present how we successfully integrated the algorithm in the OEX AUV control system based on MOOS-IvP middleware [54]. Then, we present and discuss results from tests at sea during COLLAB13 trial which demonstrate both the feasibility of running the algorithm in realtime on an onboard computer and the benefits of using the proposed algorithm over conventional predefined paths. These results represent, to the best of our knowledge, the first successful demonstration at sea of a complex nonmyopic algorithm running on an AUVs in a realistic multistatic littoral surveillance scenario.

II. Receding Horizon, Nonmyopic Optimization Algorithm

We propose a data-driven control policy (see Fig. 3) that uses the target state produced by the onboard automated perception module together with an environment model to search over a future horizon the sequence of heading decisions to meet the mission objectives. Specifically, the control algorithm uses a nonmyopic receding horizon policy to select the AUV heading to minimize, over the possible future AUV decisions, the expected posterior covariance matrices of the target location estimate errors of a Kalman tracking filter [44]. Minimizing the target position estimation error is typically of great interest in target state estimation since it can increase the track lifetime. To

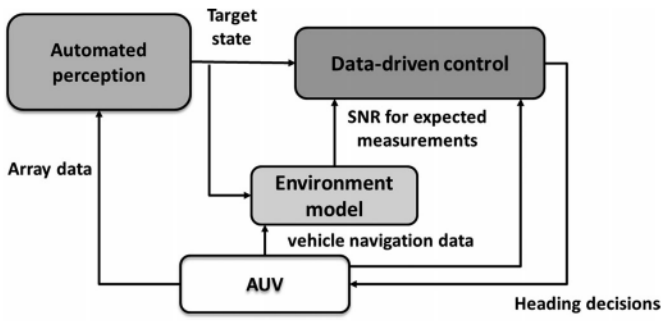

Fig. 3. Scheme of the data-driven algorithm. The data-driven control policy uses the target state produced by the automated perception modules, together with an environment model to search over a future horizon the sequence of heading decisions which meet the mission objectives, in this case reducing the target position estimate error. The selected heading decision is then passed to the vehicle for its actuation.

compute this quantity, an acoustic model [53] is first invoked to estimate the signal-to-noise ratio (SNR) for the expected measurements (see Section II-B1). SNR is a crucial quantity in our problem since it is directly proportional to the target probability of detection (PD) and inversely proportional to the measurement error.

The computed SNR is used together with bistatic contact-localization statistics [55] (see Section II-B) to compute the error covariance matrix $\mathbf{R}$ for the hypothetical measurement. $\mathbf{R}$ is needed for the correction step of the tracking filter and it is used to compute the posterior covariance matrix of the target location estimate error $\mathbf{P}$. We also assume that frequency modulated (FM) pulses are transmitted for simplicity. Other waveforms, such as continuous waveform, can be used as well by modifying the FM-related quantities accordingly.

The heading decision is discretized and the AUV can select a change $\Delta\psi_k$ from its current heading angle over a set of U possible choices. These possible motions define the set of possible source–target–receiver configurations given the prediction of the target and source position. The resulting optimization problem is solved by building a tree of possible configurations (see Section II-C). From each node, U descendants are produced, each one characterized by one of the possible $\Delta\psi_k$ decisions. The tree is built with a number of levels equal to the prediction steps M. The optimization seeks the sequence which minimizes the sum of a norm of $\mathbf{P}$ of the nodes belonging to the sequence. The first heading decision is then executed while the others are discarded according to the receding horizon policy. This approach can prove robust against possible misleading measurements since, at every computation step, it considers the information brought in by a new measurement. The algorithm is, therefore, run at each ping, when a new value for the considered track is available, and produces a new heading command. The downside of this approach lies in an increase of the computational burden since a new optimal sequence needs to be calculated at every ping.

More specifically, the optimization problem is exponential in the number of heading sequences. This complicates the solution of the problem in realtime on the vehicles and makes an exhaustive search not viable.

To solve the problem in realtime, we first simplify the tree (see Section II-C). Then, on the resulting tree, a branch-and-bound approach is used to find the optimal sequence. This way, the algorithm can be run at each ping time, sharp manoeuvres can be avoided, and the optimal sequence can be computed before the next pulse repetition interval (PRI) (i.e., a new sonar ping is emitted).

The algorithm achieves a tradeoff among different objectives considering the scene evolution over a future temporal window. In particular, it aims at:

1) moving the vehicle toward geometric configurations of low bistatic target measurement localization error which lead to good tracking performance (reduced error in the posterior target position estimate);
2) increasing the expected SNR which translates in an increase of the target PD;
3) avoiding configurations with target at array endfire or areas of high reverberation or the blanking region;
4) performing smooth manoeuvres to avoid large array bending to not deteriorate the beamforming results.

Achieving a tradeoff among these objectives is crucial to increase the AUV DCLT performance. In Sections II-A–II-C, the algorithm is described in detail.

A. Mathematical Formalism

We consider a target moving in two dimensions. A discrete-time model is adopted to describe its dynamics. The sampling time of the model is T, equal to the PRI of the acoustic source. The state vector at time k is $\mathbf{x}_k = [x_k, \dot{x}_k, y_k, \dot{y}_k]^T$, with Cartesian x–y position and relative speeds. The selected model assumes the target moving with a nearly constant velocity [56]. That is

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{w}_{k-1} \tag{1}$$

where the state transition matrix

$$\mathbf{F} = \begin{bmatrix} 1 & T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

and $\mathbf{w}_k$ is a zero-mean white Gaussian sequence with covariance matrix $\mathbf{Q}$.

At a certain time k, the information about the target position is retrieved by a track produced by the onboard tracker. Specifically, a track provides the estimate of the target state $\hat{\mathbf{x}}_k$ and the relative estimate error covariance matrix $\mathbf{P}_k$, being $\mathbf{P}_k = E[\mathbf{e}_k \mathbf{e}_k^T]$, with $\mathbf{e}_k = \mathbf{x}_k - \hat{\mathbf{x}}_k$ representing the estimate error, and E being the mathematical expectation operator.

To follow the target, the AUV can change its heading, that is the only control variable considered here. The decision is discretized and a change $\Delta\psi_k$ from the current heading angle can be selected over the set of U possible choices. These possible motions define the set of configurations given the prediction of the source and the target position. The process is then repeated for M steps (the planning horizon) into the future defining U^M possible sensor sequences. This creates a tree of possible future

TABLE I
COST COMPUTATION FOR THE CONFIGURATION $c_{i,j}$

1) For a certain configuration $c_{i,j}$, the configuration $c_{i-1,l}$ relative to the father node is retrieved. *// If $i = 1$, the information is retrieved from the root of the tree. The root is initialized with the information from a track produced by the tracker, the current estimated source position, and the measurements read from the vehicle's onboard navigation filter.*
2) Propagation step for the target:
   i) State of the target: $\hat{\mathbf{x}}^T_{i,j} = \mathbf{F}\hat{\mathbf{x}}^T_{i-1,l}$, with $\hat{\mathbf{x}}^T_{i-1,l}$ being the estimate of target state at the node $\{i-1, l\}$.
   ii) Target estimate error covariance: $\mathbf{P}^-_{i,j} = \mathbf{F}\mathbf{P}_{i-1,l}\mathbf{F}^T + \mathbf{Q}$, with $\mathbf{P}_{i-1,l}$ being the covariance error matrix relative to the node $\{i-1, l\}$.
3) Prediction of the position of the AUV/array according to the considered heading decision:
   i) Compute the new AUV heading $\hat{\psi}_{i,j} = \hat{\psi}_{i-1,l} + \Delta\psi_{i,j}$, with $\Delta\psi_{i,j}$ the heading change decision related to the node (i, j).
   ii) Predict the AUV position $\begin{bmatrix} \hat{x}^R_{i,j} \\ \hat{y}^R_{i,j} \end{bmatrix} = \begin{bmatrix} \hat{x}^R_{i-1,l} \\ \hat{y}^R_{i-1,l} \end{bmatrix} + \begin{bmatrix} \sin(\hat{\psi}_{i,j}) \\ \cos(\hat{\psi}_{i,j}) \end{bmatrix} u$, with u indicating the surge speed of the AUV.
   iii) Predict the array heading angle $\hat{\phi}^R_{i,j} = f(\hat{\phi}^R_{i-1,j}, \hat{\psi}_{i,j})$, according to the dynamic model of the towed array described in Section II-IV-B.
4) Estimate the source position, $\hat{x}^S_{i,j}, \hat{y}^S_{i,j}$. *// In our scenario we consider the source position fixed.*
5) Kalman Filter correction of the predicted covariance matrix of estimate error:
   i) Compute the measurement covariance matrix $\mathbf{R}_{i,j}$ given the computed configuration $c_{i,j}$ according to the procedure described in Section II-B.
   ii) Correct the prior covariance with the estimated information carried by the hypothetic FM measurement:
   $\mathbf{P}_{i,j} = [\mathbf{H}^T\mathbf{R}^{-1}_{i,j}\mathbf{H} + (\mathbf{P}^-_{i,j})^{-1}]^{-1}$, with $\mathbf{H}$ the measurement matrix for an FM contact. The measurement in this case provides the *x*–*y* coordinates of the target.
6) Cost computation for configuration $c_{i,j}$: $J(c_{i,j}) = \mathrm{tr}(\mathbf{P}_{\mathrm{i,j}})$

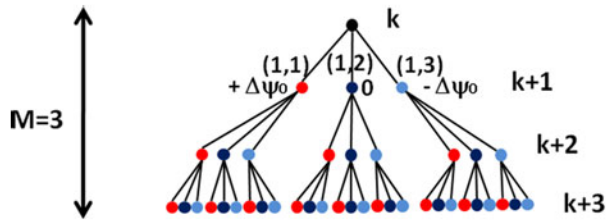

Fig. 4. Decision tree with lookahead prediction steps $M = 3$ and heading decisions $U = 3$ at time k. In the figure, the future times $(k+1, k+2, \ldots)$ of the predictions are reported. For the nodes, the relative indices and associated heading change decisions $(\Delta\psi)$ are shown. For the sake of visibility, only the indices of the first level of future predictions are reported.

source–target–receiver configurations as shown in Fig. 4. Each leaf of the tree represents one heading decision sequence.

Each node of the tree corresponds to a configuration $c_{i,j}$, where i stands for the step of future prediction and j runs on the number of nodes for a certain level of prediction i. The configuration $c_{i,j}$ is defined as $c_{i,j} = [\hat{x}^R_{i,j}, \hat{y}^R_{i,j}, \hat{\psi}^R_{i,j}, \hat{\phi}^R_{i,j}, \hat{x}^S_{i,j}, \hat{y}^S_{i,j}, \hat{x}^T_{i,j}, \hat{y}^T_{i,j}, \hat{\psi}^T_{i,j}]$, where $\hat{x}^R_{i,j}, \hat{y}^R_{i,j}, \hat{\psi}^R_{i,j}$ are the estimate/prediction of the position and heading of the receiver AUV, $\hat{\phi}^R_{i,j}$ is the prediction of the heading angle of the array (in general different from $\hat{\psi}^R_{i,j}$), $\hat{x}^S_{i,j}, \hat{y}^S_{i,j}$ is the estimate of the position of the source, and $\hat{x}^T_{i,j}, \hat{y}^T_{i,j}, \hat{\psi}^T_{i,j}$ is the estimated/predicted position and heading angle of the target.

Each configuration $c_{i,j}$ is characterized by a cost $J(c_{i,j})$. A sequence of decisions is, therefore, an M-tuple: $S_M = [c_{1,j_1}, c_{2,j_2}, \ldots, c_{M,j_M}]$. The total cost of a sequence of decisions is computed as the sum of the costs of the configurations that compose the sequence, that is

$$J(S_M) = \sum_{i=1}^{M} \alpha_i J(c_{i,j_i}) \tag{2}$$

with α being a vector of weights assumed equal to 1 in the following.

The optimization problem consists, therefore, in finding the sequence of heading decisions that minimizes $J(S_M)$

$$S^*_M = \arg\min_{S_M} \quad J(S_M). \tag{3}$$

Equation (3) is a discrete optimization problem, where the cost is minimized over the set of possible sequences of AUV heading decisions. The cost computation process for a configuration $c_{i,j}$ is detailed in Table I.

When the optimal sequence is computed, according to the adopted receding-horizon policy, the first heading of the sequence is used to control the vehicle. The process is repeated at the next ping when new values for the track are produced by the onboard tracker.

Used cost: As the cost of a configuration, we consider the trace of $\mathbf{P}$. The selection of the additive form (2) to compute the heading sequence cost instead of using a terminal scheduler cost (i.e., considering the cost of the leave configuration as that of the whole sequence [50]) is justified mainly by two reasons.

The first is that, as also underlined in [27], the considered cost is stochastic and predictive in its nature. As M increases, the accuracy with which the tracking performance is predicted is reduced. We consider the costs at the intermediate steps of the sequence instead of relying only on the terminal one that may be little informative due to the accumulated uncertainty in the prediction.

The second point is related to our specific scenario. Similar terminal costs may be produced by sequences that in the intermediate steps increase their cost drastically. For instance, this can be caused by manoeuvres that bring the target toward the array endfire. By considering the sum of the values, we want to avoid these decisions, since they can cause the track to break and hence to lose the target.

The cost for a certain configuration is also automatically increased in the following two cases.

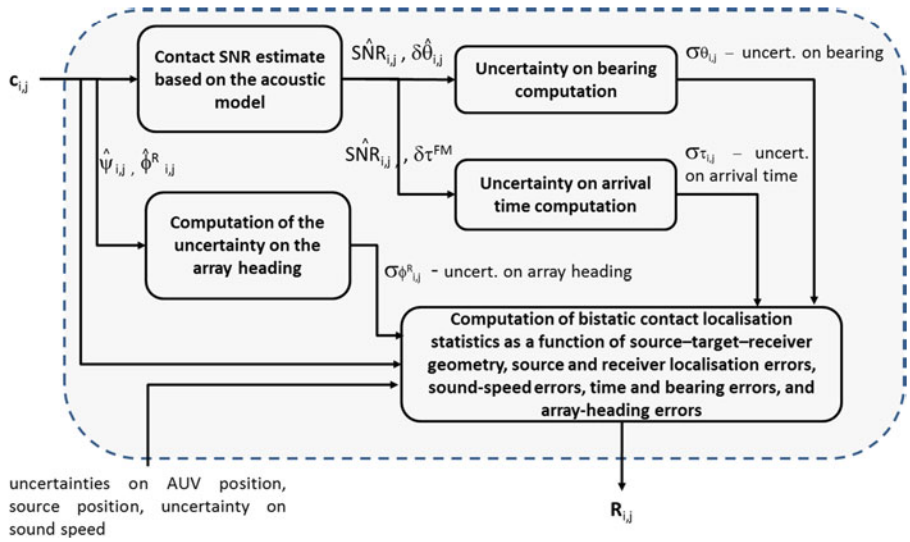

Fig. 5. Process of computation of the measurement error covariance matrix $\mathbf{R}$.

1) Blanking region: If $L + T_w c > r_r + r_s$, the cost of the node is increased. Where the different geometric quantities are described in Fig. 2, T_w being the time length of the transmitting pulse and c being the sound speed. This increment in the node cost is done to force the AUV to avoid manoeuvres causing the target to enter the blanking region.
2) Target at endfire: If one configuration is characterized by $|\theta_{i,j}| < \theta_0$ or $|\theta_{i,j}| > \theta_1$, we add a penalty to the relative cost. This mechanism is used to force the resulting behavior to not take into considerations movements which lead to the target being at array endfire. In general, this must be avoided, but there may be situations in which the AUV is forced to select this kind of manoeuvre given the current tactical scene. The proposed mechanism can, therefore, be used to tune the resulting behavior to make it more aggressive (increased probability of selecting manoeuvres leading the target being close at endfire) or more conservative (these kinds of manoeuvres are penalized due to the added cost penalty).

Finally, the vector α gives a way to discount the costs based on their future importance.

B. Computation of the Measurement Covariance Matrix R

In the procedure reported in Table I, it remains to define how to compute the estimate of the measurement error covariance matrix $\mathbf{R}_{i,j}$, which is the matrix $\mathbf{R}$ for a certain configuration $c_{i,j}$. The algorithmic steps are the following (see also Fig. 5).

1) A physics-based environmental acoustic model is invoked to compute the predicted $\widehat{\mathrm{SNR}}_{i,j}$ of the hypothetical contacts (see Section II-B1).
2) The acoustic measurement uncertainty on the bearing $\sigma_{\theta_{i,j}}$ and on the time of arrival (TOA) $\sigma_{\tau_{i,j}}$ is evaluated based on the predicted $\widehat{\mathrm{SNR}}_{i,j}$ (see Section II-B2).
3) The uncertainty on the knowledge of the array heading $\sigma_{\phi^R_{i,j}}$ is estimated.
4) $\mathbf{R}_{i,j}$ is finally calculated based on the bistatic contact localization statistics, geometry, and the evaluated uncertainties (see Section II-B3).

It is important to highlight that the $\mathbf{R}_{i,j}$ depends on the estimated SNR of the contacts since it depends on the uncertainties of the acoustic measurements. This is the link that connects the minimization of the trace of $\mathbf{P}$ to the maximization of the expected SNR.

1) Environmental Acoustic Model: A physics-based acoustic model is first invoked to compute the predicted $\widehat{\mathrm{SNR}}_{i,j}$ of the expected contact for a specific configuration. Developing an accurate model is challenging due to the complexity of the underwater acoustic environment and to the inaccurate knowledge of the input parameters. An increase in complexity of the model (e.g., using numerical models) requires an increased computational power, often not available with onboard computers. For these reasons, we adapt to our case the approach proposed in [53]. In [53], closed-form formulae are proposed for signal and reverberation in a wide variety of bistatic scenarios with variable bottom bathymetry. These formulae assume isospeed conditions and are based on an incoherent eigenray sum (expressed as the vertical grazing angle integral of the boundary reflection coefficient with spreading losses from 0° to the critical angle θ_c of the seabed) with (bistatically varying) Lambert's law scattering. The formulation was extended to handle general smoothly varying bathymetry being usable in range-dependent scenarios. The formulas are strictly closed form for limited geometries, which are of high interest in our case, specifically a flat seafloor, a tilted plane seabed, and a tilted plane with a shelf. The selected model has important features making it appealing for the use in our application. First, its computation onboard the vehicles is feasible in the required time given the low available computational power. Second, the used model provides us realistic estimates of the differences of SNRs between different configurations. More

than an exact estimate of SNR, this is the information needed by the optimization to compare different configurations. Further details on the used environmental acoustic model are reported in Appendix A.

2) Computation of Uncertainties in Acoustic Measurement and Array Heading Angle Knowledge: The SNR estimate produced by the acoustic model is used to compute the acoustic measurement uncertainties σ_τ (on TOA) and σ_θ [on direction of arrival (DOA)] (see Appendix B for details). We also quantify the uncertainty $\sigma_{\phi^R_{i,j}}$ in the knowledge of the array heading angle. We express σ_{ϕ^R} as dependent on the difference between the current AUV heading and the array heading angle, as predicted by an array dynamic model. The difference is higher when the AUV manoeuvres with larger heading changes, in that case the uncertainty on the knowledge of ϕ^R increases together with the deterioration of beamforming performance due to the array bending. Both phenomena are caught by the σ_{ϕ^R} value. It is an explicit way of weighting the importance of the turning decisions of the AUV. If σ_{ϕ^R} is increased, it will result in an increased norm of $\mathbf{R}$ incrementing the cost for that AUV manoeuvre. The AUV will, therefore, prefer smaller changes in heading avoiding large turns unless the manoeuvre is highly beneficial given the tactical scene. Details on how σ_{ϕ^R} is computed are reported in Appendix B.

3) Measurement Covariance Matrix Computation: It remains to compute the measurement covariance error matrix $\mathbf{R}$ defined as

$$\mathbf{R} = \begin{bmatrix} \sigma^2_{x^T} & \sigma_{x^T y^T} \\ \sigma_{x^T y^T} & \sigma^2_{y^T} \end{bmatrix} \tag{4}$$

with $\sigma^2_{x^T}$ and $\sigma^2_{y^T}$ being the variances on the x–y coordinate of the measured target position, respectively, and $\sigma_{x^T y^T}$ is the covariance value. This matrix represents the uncertainty in the localization of the target associated with the measurement τ, θ.

Given a certain configuration $c_{i,j}$ and the relative computed uncertainties $\sigma_{\tau_{i,j}}$, $\sigma_{\theta_{i,j}}$, and $\sigma_{\phi^R_{i,j}}$, the matrix $\mathbf{R}$ is calculated based on the approach reported in [57]. Further details on how to apply results from [57] to the specific scenario presented in this paper are reported in Appendix C.

C. Solving the Optimization in Realtime on an Embedded PC

Our aim is to provide an approach to solve the optimization problem onboard the AUV with the low available computational power. The adopted receding-horizon strategy forces us to solve an optimization problem within the PRI. Furthermore, the solution has to produce paths with limited zig-zag movements, which may be caused, at sea, by meandering tracks generated by the onboard tracker.

The optimization problem can be posed as a tree search problem, in particular as a single-source shortest path problem. We seek the sequence of configurations that minimizes the cumulative cost $J(S_M)$, as defined in (2). That sequence corresponds to the lowest cost branch of the tree.

In our case, the created decision tree can be characterized by a large number of possible sequences (U^M). For example, during COLLAB13, we used $U = 11$ and $M = 5$, which results in a tree with 177 155 nodes and with 161 051 possible heading sequences.

When we talk about the decision tree, we use the following terminology. One node identifies univocally a sequence of ancestor nodes. We define as the sequence corresponding to that node the sequence composed of all the node's ancestors. We define a node (i, j) open if its cost $J(c_{i,j})$ has been computed, expanded if all its children nodes have been opened. Finally, we define a node pruned if the node and its children have been removed from the tree. Pruning a node optimally means that the pruned node is guaranteed not to belong to the optimal decision sequence.

Different tree search techniques present in the literature can be used to exhaustively explore the tree so that once all the leaves have been opened, the optimal sequence can be found. Details about different techniques are reported, among others, in [27] and [58]. In breadth-first search (BFS), the nodes are expanded in depth order. Since each level of the tree must be stored to generate the next level, in general this approach requires a large quantity of memory. Depth-first search (DFS) remedies the memory limitation of BFS by always expanding the child of the deepest unexpanded node. Uniform-cost search (UCS) expands the lowest cost unexpanded node of a tree regardless of its depth. UCS is, in general, more time efficient than BFS to provide the optimum solution, but has a memory complexity similar to BFS [58].

Alternative approaches to the exhaustive search are based on greedy search or on heuristic approaches [58]. Greedy search does not search the tree exhaustively and does not guarantee the optimal solution. Heuristic algorithms try to exploit partial knowledge about the structure of the space to produce computationally efficient solutions. Their main disadvantage is that they are ad hoc solutions.

The approach followed in this paper is based on a heuristic rule to simplify the decision tree relying on the knowledge of the problem, then the additive nature of the cost is exploited to limit the memory complexity of UCS through a branch-and-bound algorithm.

1) Forbidden Decision Method: The forbidden decision method reduces the number of possible heading decisions based on the previous decision considered along the tree expansion. Specifically, selecting at a certain level j a heading change with an absolute value greater than $\Delta^{FD}\psi_{\max}$ implies to discard all the decisions at the next level $j + 1$, which lead to change the heading in the opposite direction.

The mechanism is shown in Fig. 6. In this case, $|\Delta^{FD}\psi_{\max}| = 2\psi_0$. This mechanism not only is used during the decision tree expansion, but also takes into account the executed AUV heading decisions at the previous decision time. For instance, if at a step k the AUV executes a manoeuvre which increases its heading more than $\Delta^{FD}\psi_{\max}$, at the next decision time all the decisions that decrease the current heading are discarded (all negative heading decisions) at the first decision level of the tree.

The forbidden decision method has the objective to reduce the number of nodes to open with a consequent decrease in the total computational time. At the same time, the mechanism reduces unwanted zig-zag trajectories and avoid large array bending and

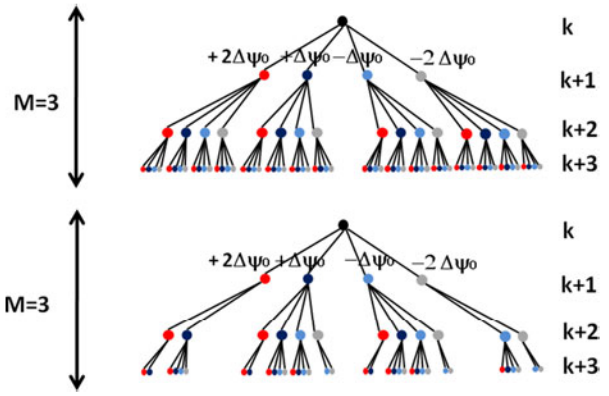

Fig. 6. (Top) Original decision tree with $U = 4$ and $M = 3$ built at time k. (Bottom) Simplified tree according to the forbidden decision criterion. If at a certain step j, the considered decision is $\pm 2\Delta\psi_0$, in the next step the heading changes decreasing/increasing the current heading are discarded.

possibly unsafe manoeuvres. This simplification is considered sufficient for our problem. In fact, it leaves a solution space large enough to find all the sequences of interest. As an example, the algorithm makes it possible to calculate heading sequences in which the AUV first turns clockwise and after some steps turns counterclockwise. This may happen when the AUV navigates parallel to the target and then decides to move closer to finally return to a parallel trajectory. The tree simplification must not remove these kinds of trajectories.

2) Uniform Cost Searching Branch-and-Bound Algorithm: A branch-and-bound-based method [27] is used to optimally prune the simplified decision tree. Reducing the search space is indeed mandatory to reduce the number of nodes to be opened thus decreasing the needed computational time. The key issue to apply a branch-and-bound method is the possibility of setting some lower bounds on the costs of all nodes easily. To do this, we exploit the tree structure and the additive nature of the sequence cost. This allows us to assign the lower bound on the cost of any unopened node as the cost of the sequence corresponding to its nearest open ancestor. Furthermore, the additive cost provides the guarantee that the optimal sequence cannot be generated from sons of any pruned node. A uniform cost search is used as the tree search technique in solving the tree. Details about the used branch-and-bound algorithm are reported in Appendix D.

This kind of pruning scheme, known in the literature as a best first branch-and-bound algorithm [59], was selected due to its relatively efficiency in processing time. Reducing the overall computational time is the main purpose. For instance, it opens fewer nodes (lower computational time) than a BFS-greedy search branch-and-bound scheme (see [29]) which, on the other hand, offers a more efficient memory usage. Branch-and-bound algorithms can reduce drastically the computational time but typically they do not eliminate exponential complexity [27]. As will be demonstrated in Section III, in our problem, the proposed approach finds the optimal sequence on the tree that is simplified by the forbidden decision mechanism and it does so respecting the AUV realtime constraints (execution ends before the next PRI).

One interesting feature of the approach is that the first candidate solution found is already a good compromise between finding a low-cost solution and computing it in a short time. In this way, a suboptimal solution for the control of the AUV is calculated in a short amount of time. If more time is available, the algorithm will explore further the tree to incrementally seek better solutions until the optimal one is found.

3) Adapting the Discretization of Heading Decisions: As already discussed, the possible AUV heading changes $\Delta\psi$ are discretized. The maximum range $[-\Delta\psi_{\max}, +\Delta\psi_{\max}]$ of heading changes is selected with the purpose of avoiding large turns at each decision step. The number of discretization steps U is then selected. Given U, the discrete heading changes $\Delta\psi_i$s are computed by equally dividing the range of available heading decisions

$$\left\{-\Delta\psi_{\max}, \ldots, -2\frac{\Delta\psi_{\max}}{\lfloor \frac{U}{2} \rfloor}, \frac{-\Delta\psi_{\max}}{\lfloor \frac{U}{2} \rfloor}, 0, \frac{\Delta\psi_{\max}}{\lfloor \frac{U}{2} \rfloor}, 2\frac{\Delta\psi_{\max}}{\lfloor \frac{U}{2} \rfloor}, \ldots, \Delta\psi_{\max}\right\}.$$

However, fixed $\Delta\psi_i$ values during the whole mission may not be satisfying given the different situations the AUV may face. There may be situations in which the AUV is navigating along a course close to the optimum solution and needs only small heading changes to improve further the tracking performance. In this case, a finer granularity is needed. In other cases, the AUV may need large turns to escape from a situation in which a high tracking error is expected potentially compromising the mission.

We address this issue by introducing an adaptive range for the possible heading change decisions. Leaving U fixed, we change the value of the adapted maximum range $\Delta\psi_{M_j}$ before each activation of the control algorithm on the basis of how the mission is progressing. In this way, the maximum range of possible heading decisions becomes $[-\Delta\psi_{M_j}, +\Delta\psi_{M_j}]$). The last H taken heading decisions are analyzed.

1) If low values of $\Delta\psi$ have been executed, the range of heading decision is decreased up to the $\Delta\psi_{\min}$ value.
2) Otherwise, the range is increased up to a maximum value of $\Delta\psi_{\max}$.
3) During the operations, the adapted range $\Delta\psi_{M_j}$ is utilized every two uses of $\Delta\psi_{\max}$.

This gives the AUV the opportunity to increase the granularity when it is more likely required (low heading changes in the recent past). On the other hand, large $\Delta\psi$s taken in the recent past increase the range since it is likely the AUV is going to need large turning also in the future. Finally, the adapted range selection policy (once every three optimization activations) leaves the AUV the possibility to continue exploring the solution space even when the adapted range shrinks.

III. Experimental Results

We validated the nonmyopic control algorithm during multiple trials at sea [60], [61], and we report here results from the COLLAB13 sea trial held in July 2013, where the algorithm

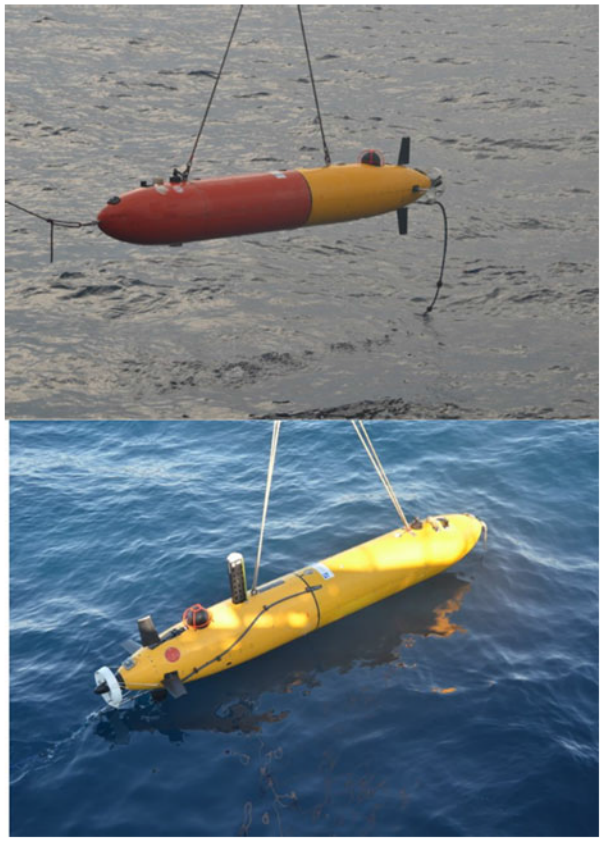

Fig. 7. Two CMRE OEX AUVs, (top) Groucho and (bottom) Harpo.

Fig. 8. (Left) DEMUS source during deployment. The eight free-flooded ring transducers are clearly visible. Below the rings are the electronics. (Right) The DEMUS radio buoy.

Fig. 9. CMRE echo repeater. It records the waveforms received from the DEMUS transmissions and then it retransmits with a user-specified amplitude gain after a user-specified delay.

was extensively tested in a realistic ASW scenario. This section describes the CMRE cooperative ASW system deployed during COLLAB13 and the implementation of the proposed algorithm in the CMRE AUV control system. Results during the at-sea tests are then reported and discussed.

A. Centre for Maritime Research and Experimentation Cooperative Antisubmarine Warfare Multistatic System

The vehicles used as receiving nodes in the CMRE network are two OEX AUVs (shown in Fig. 7). These vehicles are approximately 4.3 m long and 0.53 m wide. The endurance depends on the payload. They can reach 16 h of operations at a speed of 1 m/s. The maximum vehicle operating depth is 300 m. The onboard navigation filter fuses inertial navigation and a velocity estimate provided by an acoustic Doppler velocity log (DVL). The DVL used is capable of maintaining bottom lock up to a maximum altitude of 200 m, allowing operations in the shallow areas of interest with depths of 100–150 m. AUVs communicate with each other and with the C2 center via a 7/17-kHz Evologics low-frequency modem.

The OEX AUVs are both deployed with the BENS slim towed array (SLITA) [4]. The arrays are deployed approximately 3.5 m behind the vehicle. The BENS arrays have three nested sets of 32 hydrophones each. The hydrophone set used during the sea trials was optimized for frequencies up to 3.47 kHz (0.21 m spacing).

As the acoustic source, we used the deployable experimental multistatic undersea surveillance system (DEMUS) [4] (see Fig. 8).

The source is bottom-tethered and was not moved during the experiments. The DEMUS source was equipped with a modem to enable remote actuation via underwater acoustic communications, and a radio buoy to allow wireless actuation and monitoring of the source, as well as GPS synchronization for timing and accurate position estimates.

All the assets were deployed from the NATO Research Vessel (NRV) alliance. The vessel operated as the C2 center during the experiments. Furthermore, it towed the echo repeater (see Fig. 9) to act as the simulated target during the experiments.

The echo repeater recorded the waveforms received following the DEMUS transmissions and then re-transmitted the recorded signals with a user-specified amplitude gain after a user-specified delay. This gain serves as a substitute for the target sonar cross section or reflectivity. The depth of deployment of the echo repeater was 50 m.

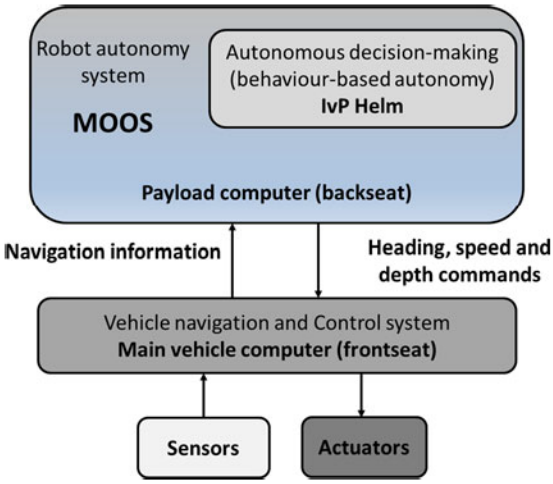

Fig. 10. Schematic representation of the backseat–frontseat autonomy paradigm as used in MOOS-IvP. The key idea is the separation of vehicle autonomy from vehicle control [54].

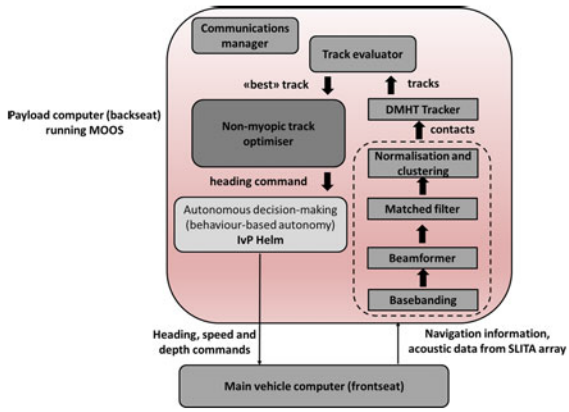

Fig. 11. OEX vehicles software architecture.

B. OEX Control System and Nonmyopic Control Algorithm Integration

The OEX AUV control system is based on the "backseat-driver" paradigm: a backseat computer executes the processes managing the mission and produces commands for a frontseat computer in charge of the vehicle low-level control [54]. In our AUVs, two CPU boards are present, one for the frontseat and one for the backseat computer. The nonmyopic control algorithm hence runs on the backseat computer. The backseat computer of OEX vehicles was an ETX-945-L2400 E, from Diamond Systems Corporation, Sunnyvale, CA, USA [62]. It is a 1.66-GHz, Intel Core 2 Duo board with 2-GB DDR2 SDRAM.

The software architecture running on the vehicles is based on MOOS-IvP [54] (see Fig. 10). MOOS-IvP is an open source C++ framework for providing autonomy to robotic platforms, in particular marine vehicles. MOOS-IvP is based on the publish/subscribe paradigm: a community of processes subscribe to receive and publish variables from/to a database. For the management and control of vehicles, the MOOS framework works according to the backseat driver and fits with the described hardware. MOOS processes which manage the robot mission run on the backseat and receive data/issue commands from/to the frontseat computer. The IvP Helm [54] is a MOOS application that enables behavior-based autonomy. Behaviors can run simultaneously and can be grouped into behavior sets, which are active based on certain conditions. Each behavior is given a weight which is defined at mission start but may also be adjusted dynamically during the mission by onboard processes. IvP, a mathematical interval programming technique, combines the objective functions produced by active behaviors to determine a globally optimal solution for each domain [54]. The IvP Helm, typically running four times per second, is able to reconcile the different active behaviors to produce the commands for the frontseat controller, specifically speed, heading, and depth commands.

The AUV software architecture is sketched in Fig. 11. Data from the array hydrophones are acquired by the frontseat computer and, at each ping time, are read by the backseat computer. Here, the time series are first basebanded. Then, the time series on each of the towed arrays hydrophones is mathematically transformed in such a way as to produce a time series per look direction (beam). For each beam, the time series is filtered using a replica of the sources outgoing pulse, resulting in a pulse-compressed time series. Finally, the matched filter output is normalized to be able to calculate the SNR per point and an SNR threshold is applied to the normalized data to form contacts (detections). Details about the CMRE signal processing chain can be found in [15]. Contacts (range and bearing data) are then processed by the distributed multiple hypothesis tracker [55]. The tracker processes the contacts, first geolocating them using the contact's range and bearing, the source location, and the array location and array heading. Then it fuses the contacts to generate tracks. The tracks are then evaluated and a score is given based on the track length and associated contact SNR [14]. A long track with high SNR associated contacts is considered as a track interesting to prosecute. High SNR contacts associated with the target are the typical case when an artificial target is used. The situation becomes indeed more challenging when a real vehicle is tracked due to the target strength dependence on the bistatic aspect angle. The "best" track is then passed to the nonmyopic algorithm which produces the heading command sent to the IvP Helm. The command is then reconciled with the other active behaviors (e.g., the safety behaviors) and sent to the frontseat computer autopilots.

The described algorithm was implemented in C++ and ran in realtime as a MOOS process. In this paper, with "realtime" we mean that the algorithm produces its output before new data are produced by the signal processing chain. To avoid possible excessive computation time during the optimization, a safe mechanism is introduced that aborts the optimization returning the best solution so far after 25 s of execution. In COLLAB13, PRI was 48 s, so 25 s is considered sufficient time to avoid overlapping the optimization at two consecutive pings considered all the other onboard running processes.

C. Experimental Results From the COLLAB13 Sea Trial

The COLLAB13 sea trial took place in the Ligurian Sea, off the coast of the city of La Spezia, Italy, in July 2013. One of the

TABLE II
PARAMETERS USED IN THE ACOUSTIC MODEL

Parameter	Value	Notes
Parameters related to the array		
N	32	number of hydrophones
d [m]	0.21	hydrophone spacing
Parameters related to the source and signal		
$SL^{(\mathrm{dB})}$ [dB]	217	source level
PRI [s]	48	Pulse repetition interval
Source position (North, East) [m]	[12172.427 11350.9551]	North-East position of the source
T_w [s]	1	pulse width
F_{bw}^{FM} [Hz]	200	bandwidth of the FM pulse
F_c^{FM} [Hz]	2700	carrier frequency of sonar FM pulse signal
Parameters related to the bathymetry and acoustic environment		
c [m/s]	1510	sound speed (isospeed assumption)
NL^{dB} [dB]	70	noise level
α [dB/rad]	2	reflection loss gradient
θ_c [rad]	0.488	critical angle (28°)
μ	$10^{(-27/10)}$	constant related to the bottom-surface albedo
Origin slope (North, East) [m]	[12172.427 11350.9551]	the tilted-plane starts at the source position in a North-East local coordinate reference system
ϵ_0 [m/m]	0.0306	bottom slope in the tilted plane direction
Maximum depth [m]	350	maximum depth considered in the operative area
ψ_s [rad]	3.9964	angle of the tilted-plane (from North) (227.26°)

TABLE III
MEASUREMENT UNCERTAINTIES

Parameter	Value	Notes
σ_{xR} , σ_{yR} [m]	15, 15	uncertainty on receiver position
σ_{xS} , σ_{yS} [m]	3, 3	uncertainty on source position
σ_c [m/s]	7.5	uncertainty on sound speed
σ_θ [deg]	variable [0.8, 3]	uncertainty on measured bearing. Computed as described in Section B
σ_τ [s]	variable [0.048, 0.1]	uncertainty on the TOA. Computed as described in Section B
$\sigma_{\phi R}$ [deg]	variable [2, 12]	uncertainty on the knowledge of array heading. Computed as described in Section B

main objectives was demonstrating the performance increase in AUV onboard tracking using autonomous decision making.

The bathymetry of the operational area was modeled as a tilted plane, with the slope starting at the DEMUS source location (90 m) and increasing moving toward south-west. The full set of parameters used in the environmental model is reported in Table II.

The parameters used for the computation of the measurement error covariance matrix (see Section II-B3) are reported in Table III. An example of the resulting tracker localization error

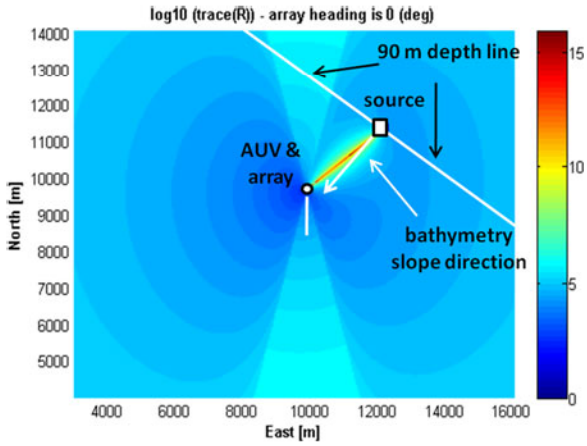

Fig. 12. Localization error shown as $\log 10(\mathrm{tr}(\mathbf{R}))$ at different target positions. The considered reflector is isotropic with a target strength of 20 dB. The SNR used to compute this figure is that shown in Fig. 20.

TABLE IV
OPTIMIZATION ALGORITHM PARAMETERS

Parameter	Value	Notes
M	5	number of prediction steps
T [s]	48	prediction step time length
U	11	number of heading possible choices
$\Delta\psi_M$ [deg]	[20, 60]	maximum heading change decisions (adapted as described in Section II-C3)

with varying target positions within the COLLAB13 scenario is reported in Fig. 12.

As already presented in [60], the effect of the various parameters on the algorithm's performance can be evaluated using simulations and historical at-sea data. Results from [60] allowed to select the following parameters to be deployed at sea. More specifically, $U = 11$ with $\Delta\psi_{\max} = 60°$, $\Delta\psi_{\min} = 20°$, with a step size in the adaptation of 10°. Furthermore, 20° was the value used for $\Delta^{FD}\psi_{\max}$ in the forbidden decision method. The prediction step is a time length equal to the PRI used in COLLAB13, 48 s. Concerning the other algorithm's parameters, we assume a Q matrix in the form reported in [63] with the process noise standard deviation $\sigma_v = 310^{-2}$ m/s^2 and T set equal to the PRI. The other key parameter to be selected is the number of prediction steps M in the optimization, since this deeply affects the behavior of the algorithm. Again, from simulation runs, the best results from the localization point of view were obtained with $M = 5$, while producing a decision tree of tractable dimensions. For this reason, the horizon length was set to $M = 5$ during all the experiments at sea. The algorithm's parameters are summarized in Table IV.

In the following, we report the results from two experiments in which the proposed algorithm controlled Groucho OEX and we discuss its performance. In both cases, the target was an echo repeater (with a gain of 20 dB) towed by the NRV Alliance, while the DEMUS acted as the acoustic source. For this reason, the target model for an isotropic reflector was used. In both the

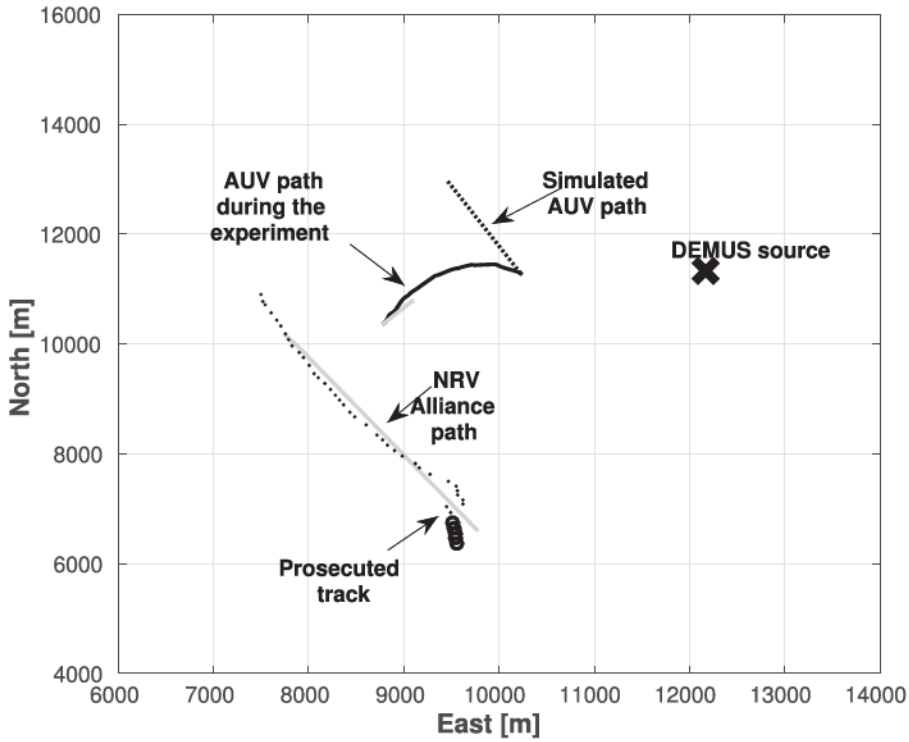

Fig. 13. Experiment 1 (May 7, 2013—Groucho) at ping 40 (pings are counted from the beginning of the behavior activation). The real path followed by the AUV controlled by the nonmyopic algorithm is visible in black. In gray the array heading is shown. The simulated trajectory of the AUV continuing along the trackline is shown in dashed black. In gray the NRV Alliance path is visible with the prosecuted track produced by the onboard tracker is drawn in dotted black. Black circles are the prediction of the algorithm for the target position over the planning horizon.

experiments, the source transmitted a 2.6–2.8 kHz linear FM sweep, with a duration of 1 s, and a PRI of 48 s.

1) Experiment 1 (May 7, 2013)—Groucho OEX: Experiment 1 is part of a trial held on July 5, 2013 in which the NRV Alliance navigated at a speed of 2 m/s toward south-east while towing the echo repeater. The Groucho OEX was driving in a (rectangular) racetrack heading to north-west at a speed of 1 m/s. The nonmyopic controller was then activated via an acoustic message from the NRV Alliance. The "best" track produced by the onboard tracker was selected and prosecuted by the algorithm. To evaluate the effectiveness of the nonmyopic controller in considering the effect of future actions, we compare the at-sea results with the equivalent path and measurements had AUV proceeded along the predefined trackline.

In Fig. 13, the paths of the AUV during the experiment and the simulated trackline are shown together. The target's real position and the track produced by the on-board tracker are also reported.

The adaptive algorithm controlled the AUV to turn toward south-west. The movement is beneficial for a number of reasons with the controller achieving a very good tradeoff between keeping the target at broadside (highest signal processing performance) and moving the AUV closer to the target. This is visible in Fig. 14, where we show the trace of the computed $\mathbf{R}$, the estimated covariance matrix of the measurement error given the true position of the target. We also show the trace of $\mathbf{P}$, the posterior covariance of the target location estimate error of the tracking filter. $\mathbf{R}$ and $\mathbf{P}$ are computed at every ping as described in Section II. These are important quantities to assess the performance of the AUV behavior. Low values of $\mathrm{tr}(\mathbf{R})$ mean accurate measurements. This brings both a lower error in the target position and velocity estimate (as visible by looking at the $\mathrm{tr}(\mathbf{P})$ in Fig. 14) and, more importantly, tracks with a longer life time, due to the increase of probability of measurement-to-track associations. After the beginning of the turning, a larger error in the nonmyopic case is present. This is caused by the estimated increase in the uncertainty of the array heading knowledge due to the array's bending. Later on (at ping 6), the error in the nonmyopic case becomes lower than trackline's. This can be achieved since the nonmyopic approach, differently from a myopic one, can make decisions that produce a decrease in the objective function at the immediate future steps, but that offer better reward later in the future. From that time on, the localization error in the racetrack pattern case increases more and results in a larger $\mathrm{tr}(\mathbf{P})$. At ping 40, when the track breaks, in the racetrack pattern case, the average measurement error in the x–y coordinates presents a standard deviation of 179.7 m, while in the nonmyopic case the standard deviation is 73.2 m. With the use of the nonmyopic algorithm, the error is 40.1% of that we would have had without adaptation of the AUV trajectory. The nonmyopic algorithm achieves a tradeoff between driving the AUV closer to the target and keeping it at a good DOA (close to broadside). On the other hand, when navigating along the trackline, the AUV moves farther from the target and, in addition, the target moves toward the array endfire. This can be seen in the left figure of Fig. 15, where the DOA is shown in both cases. While during the navigation along the trackline, the DOA increases since the target is moving toward endfire

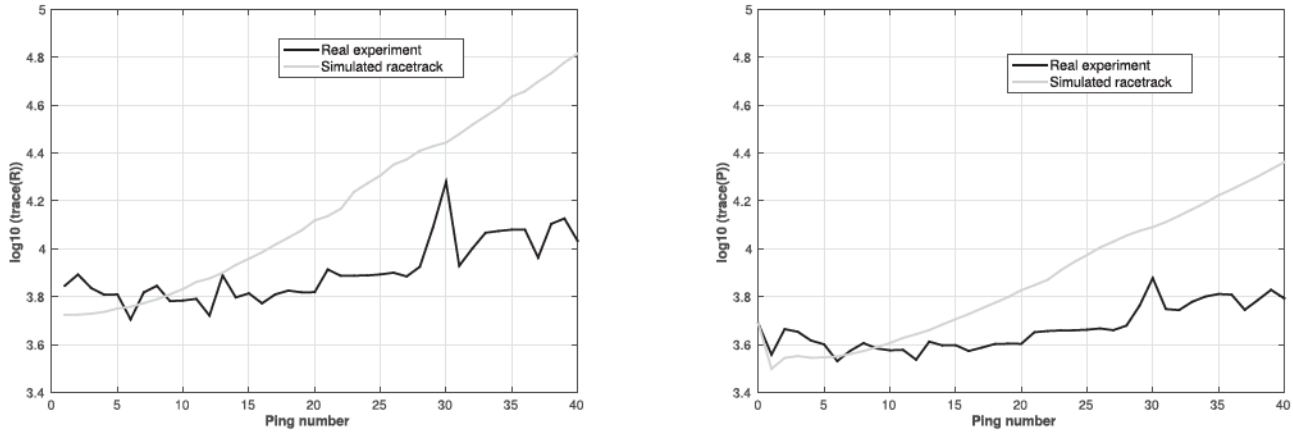

Fig. 14. At the beginning of the manoeuvre, the measurement localization error increases with respect to the trackline due to the estimated array bending. Then, the manoeuvre leads to a lower localization error if compared with the trackline, as visible in (a). This results in a reduced error in the estimate of the target position as can be seen in (b). (a) Experiment 1 (May 7, 2013—Groucho). $\log_{10}(\mathrm{tr}(\mathbf{R}))$ (computed using the real target position) in the case of the AUV controlled by the non-myopic algorithm and of the simulated trackline. (b) Experiment 1 (May 7, 2013—Groucho). $\log_{10}(\mathrm{tr}(\mathbf{P}))$ in the case of the AUV controlled by the non-myopic algorithm and of the simulated trackline.

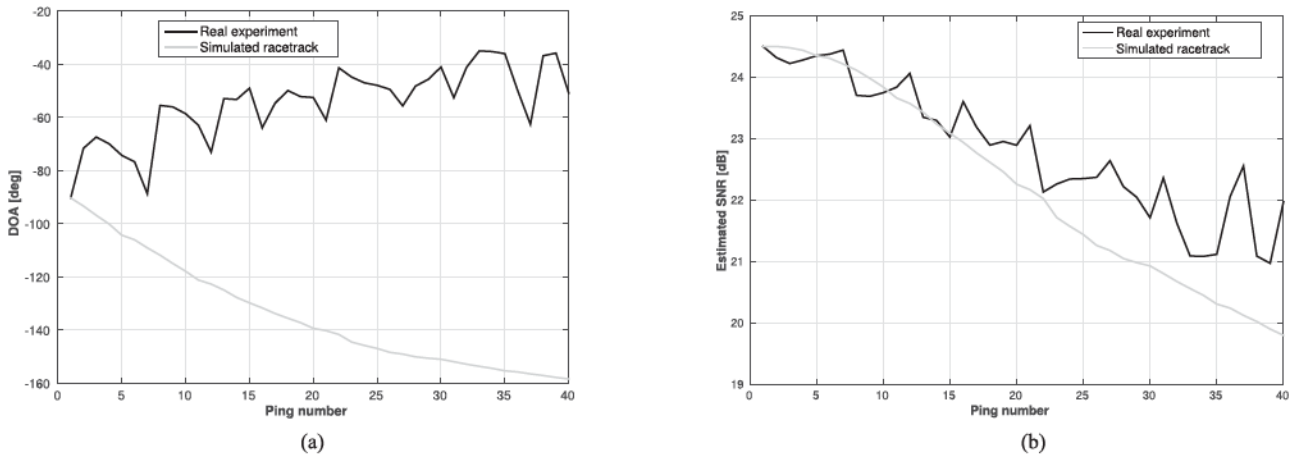

Fig. 15. (a) In the trackline case, the target moves towards endfire, while in the adaptive case the DOA is kept between $40°$ and $60°$. (b) SNR in the trackline case decreases since the AUV is moving farther from the target and the target is moving at endfire. With the data driven algorithm, the turning leads to higher contact SNR, which increases the target PD.

(reaching $\sim 20°$ of DOA), with the manoeuvring the DOA value remains between $40°$ and $60°$. The movement is not only beneficial from the localization error point of view, but also from the SNR perspective. This is visible in the right figure in Fig. 15. Here, the computed SNR is shown for the two cases. An increase of the SNR brought by the turning can be appreciated. A larger SNR is another important objective to achieve for target tracking, since it increases the target PD. These combined effects allowed the AUV to maintain the target track for long (40 pings). Increasing the life time of tracks is critical for a correct target classification.

2) Experiment 2 (February 7, 2013)–Groucho OEX: In the experiment 2, the target was moving at 2 m/s along a trackline heading to north-west (see Fig. 16). The AUV was moving at 1 m/s in a rectangular racetrack pattern heading south-west.

After its activation, the adaptive algorithm commands a turn toward west. As in the previous case, we also present a simulative analysis of what the vehicle would experience at-sea had it continued along the trackline it was following at the moment of the adaptive behavior activation.

The turning of the AUV avoids the target moving at endfire as it happens in the trackline case. The benefits of using the proposed algorithm in this experiment are even more evident than in the previous one. This is evident from Fig. 17. After ping 8, in the racetrack pattern, a critical increase in the measurement error occurs. This could result in breaking the track with consequent target loss. The manoeuvre, instead, allows us to reduce the error. The effects of the adaptive behavior are also appreciated by looking at Fig. 18. In the trackline case, the target moves at array endfire (DOA crossing the $0°$). With the adaptive behavior, the DOA increases reaching the broadside condition after an initial period in which remains close to $20°$ (also due to the error of the track in estimating the real target position visible in Fig. 16—see the situation at ping 5). This has important consequences on SNR (see Fig. 18). The manoeuvre not only avoids the drop of SNR experienced in the trackline

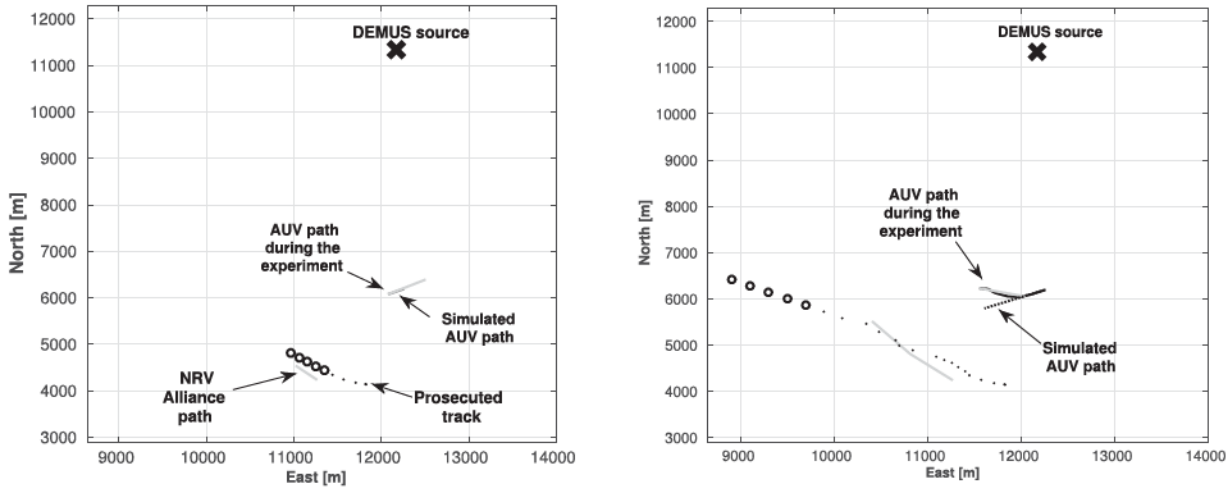

Fig. 16. Experiment 2 (February 7, 2013—Groucho) at (a) ping 5 and at (b) ping 17. The real path followed by the AUV controlled by the nonmyopic algorithm is visible in black. In gray the array heading is shown. The simulated trajectory of the AUV continuing the trackline is shown by dashed black. In gray the NRV Alliance path is visible with the optimized track produced by the onboard tracker is drawn in dotted black. Black circles are the prediction of the algorithm for the target position.

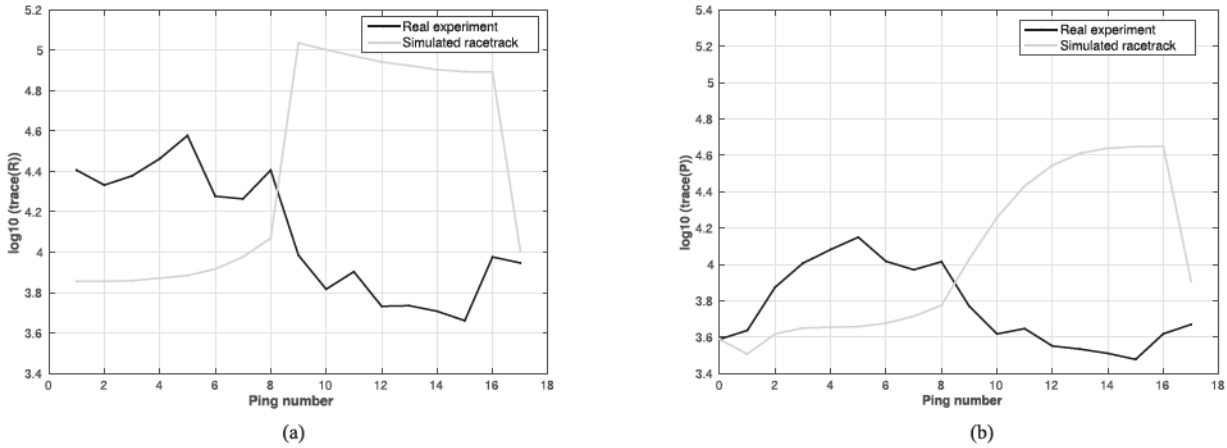

Fig. 17. (a) At the beginning of the manoeuvre, the measurement localization error increases with respect to the racetrack due to the estimated array bending . Then, $\log_{10}(\mathrm{tr}(\mathbf{R}))$ reduces thanks to the turn. The adaptive algorithm avoids the target moving at endfire as it happens in the racetrack pattern case. Target at endfire causes a significant increase of the localization error, which can break the track. (b) Effects of this increase in localization error are visible in the estimate of the target location.

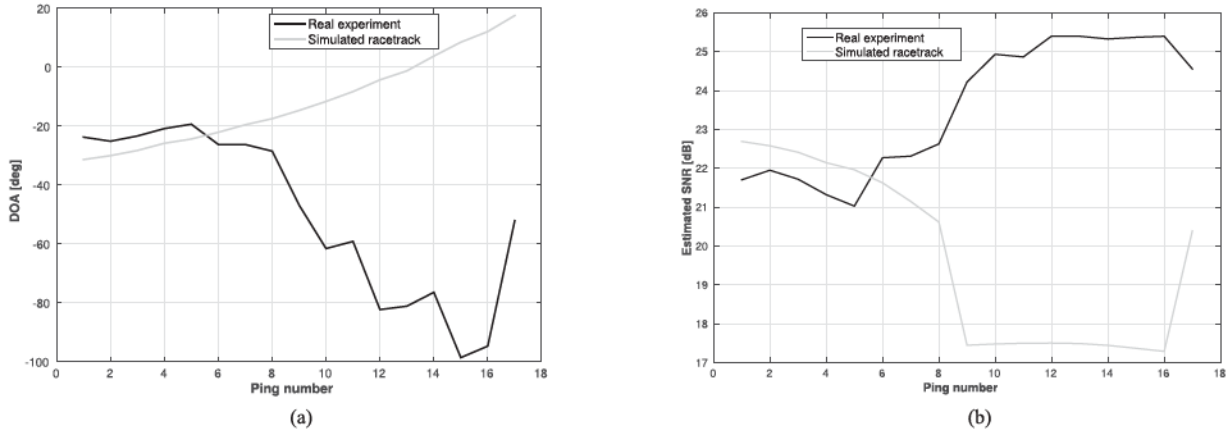

Fig. 18. (a) In the racetrack case, the target moves at endfire, while with the manoeuvre the target moves at broadside. This increases the SNR and reduces the localization error. (b) SNR in the racetrack decreases significantly after ping 8 since the DOA becomes lower than 20°. On the other hand, the SNR, in the case of the adaptive behavior, increases due to the effect of the manoeuvre.

case, but also increases it with respect to the beginning of the experiment.

3) Considerations About Computation Time: Execution time of the algorithm on the vehicles can be related to the number of the open nodes in the decision tree. The computation time on the backseat computer to open one node resulted on average 1.72 ms. During the trials, the number of opened nodes was on average 5027 with a standard deviation of 3225. The resulting execution time was $\mu = 8.49$ s with $\sigma = 5.46$ s. A timeout to allow a maximum computation time was set to 25 s to avoid problem during the cycle of operations. If the timeout expires, the best solution found up until that point is considered the final one. The timeout expired in less than 4% of the algorithm activations, showing that it is possible to execute on embedded computers.

These results demonstrate that the algorithm can be run on the embedded computers present onboard the vehicles.

IV. Conclusion

In this paper, we have tackled the problem of controlling an AUV in a multistatic surveillance scenario to increase its tracking performance. Results from sea trials demonstrate how the proposed data-driven algorithm is effective in controlling the AUV headings and robust in its ability to deal with outliers and computational constraints. The described approach couples an environment model with bistatic localization statistics and seeks optimal heading decisions over a future time horizon. The objective is to minimize the predicted estimate error of the target location. This increases the track survival time, which, in turn, can facilitate the target classification process, a critical aspect in surveillance applications. The deployment of the system done during the COLLAB13 experimental campaign, is, to the best of our knowledge, the first successful demonstration at sea of a complex nonmyopic algorithm running in realtime on AUVs in realistic multistatic ASW scenarios.

The resulting behavior optimally trades off among different competing objectives:

1) search for geometries of low bistatic target localization error;
2) maneouvre to increase the expected SNR which results in a higher target PD;
3) avoid configurations with target at array endfire or in areas of high reverberation or in the blanking region.

The usage of a nonmyopic solution improves the overall performance with respect to more traditional reactive and myopic algorithms [14]. With the proposed method, the AUV can make decisions which do not produce high short-term benefits, but that offer better rewards later in the future. The receding horizon approach showed its ability to improve the robustness of the algorithm in tackling with inaccurate measurements potentially producing nonsmooth target tracks. The downside of this approach lies in an increase in the computational overhead required since a new optimal sequence has to be calculated at every ping. For this reason, this paper proposed solutions to reduce the size of the search space of the algorithm and to decrease its computational requirements to run onboard AUVs, which are typically equipped with low computational power computers. Furthermore, by including explicitly an estimate of the array bending, turning decisions are only made if a large reward from a maneouvre is expected. This reduces large turning manoeuvres, which could bend the array with the consequent deterioration of the beamforming. These are essential features to port algorithm on underwater robots which must work robustly in a real world.

The algorithm can, therefore, be considered as a valid building block to create more complex and more general AUV tactical planners.

However, the tactical scene is often complex. The selection of the candidate track to be prosecuted may not be trivial. Several tracks may be present also after the execution of disambiguation manoeuvres to break the ghost track. It is also possible that the best candidate tracks break after only a few pings. The answer to provide to move toward a fully autonomous system is how to effectively manage the situation in these cases [18]. Future work will address this topic and develop a framework capable of dealing with multiple candidate tracks to suitably trigger and interrupt the nonmyopic optimization manoeuvres on the basis of the evolving tactical scene. Another direction of investigation is to use the information given by bistatic angle-dependent target strength models. These can be used when tracking a non-isotropic target to move the AUV toward glint configurations. In this case, the robot-target relative speed becomes critical, since only a fast vehicle (and this is not usually the case when AUVs are considered) can effectively move to effectively exploit glints. However, when multiple vehicles are deployed, this information could be passed on to other AUVs which are in better geometrical positions.

Appendix

A. Environmental Acoustic Model

The bistatic sonar equation is used to compute the SNR in decibels (dB). The equation for FM measurements and for a configuration $c_{i,j}$ can be written as [64]

$$\mathrm{SNR}_{i,j}^{(\mathrm{dB})} = EL_{i,j}^{(\mathrm{dB})} - \left[\begin{array}{l} \left(NL_{i,j}^{(\mathrm{dB})} - AG^{(\mathrm{dB})} - PG_{\mathrm{noise}}^{(\mathrm{dB})} \right) \oplus \\ \left(RL_{i,j}^{(\mathrm{dB})} - PG_{\mathrm{reverb}} \right) \end{array} \right] - \Delta^{(\mathrm{dB})} \quad (5)$$

with subscript i, j indicating quantities that depend on the considered $c_{i,j}$, $EL_{i,j}^{(\mathrm{dB})}$ denotes the echo-level, the first row of the matrix contains noise-related terms and the second one reverberation-related terms, $\oplus$ denotes the power summation, and $\Delta^{(\mathrm{dB})}$ is the additional sonar signal processing loss included in the FM processing.

The echo-level $EL_{i,j}^{(\mathrm{dB})}$ is computed as

$$EL_{i,j}^{(\mathrm{dB})} = SL^{(\mathrm{dB})} - TL_{s_{i,j}}^{(\mathrm{dB})} - TL_{r_{i,j}}^{(\mathrm{dB})} + TS_{i,j}^{(\mathrm{dB})} \quad (6)$$

where $SL^{(\mathrm{dB})}$ is the source level; $TL_{s_{i,j}}^{(\mathrm{dB})}$ is transmission loss from the source to the scatterer (target) (see Fig. 2); $TL_{r_{i,j}}^{(\mathrm{dB})}$ is transmission loss from the scatterer (target) to the receiver

(see Fig. 2); $TS_{i,j}^{(\mathrm{dB})}$ is bistatic target strength, in general function of the target aspect angle θ_a and of the bistatic angle β (see Fig. 2).

Several terms in (5) are related to the noise present in the environment and to the reduction of this noise through the signal processing executed onboard the vehicle once the sonar echo is received. Some quantities, therefore, depend on the particular beamforming/signal processing used.

Specifically, $NL_{i,j}^{(\mathrm{dB})}$ is the ambient noise level within the processing bandwidth at the baseband data. Typical noise value can be retrieved from the classic Knudsen's ambient noise measurements [64]. $AG^{(\mathrm{dB})}$ is the array gain against noise in the beamformer for the FM processing. The array gain becomes the directivity index under isotropic noise assumption [50], [64] and is defined as $10\log(N)$ with N as the number of the hydrophones composing the array.

$PG_{\mathrm{noise}}^{(\mathrm{dB})}$ represents the processing gain against noise in the matched filter, and is related to the time–bandwidth product of the matched filter [65].

Reverberation is the total acoustic intensity scattered from inhomogeneities and boundaries in the ocean measured at the receiver location. The problem of calculating the reverberation can be difficult due to the complexity in summing the contributions of a large number of scatterers insonified by many propagation paths due to the source and receiver beam patterns [50]. In [53], the principal contribution to the reverberation is assumed within the main lobe of the receiver, and scattered from the bottom-surface scattering patch at the contact position. In this way, the in-beam reverberation can be written as

$$RL_{i,j}^{(\mathrm{dB})} = SL^{(\mathrm{dB})} - TL_{sp_{i,j}}^{(\mathrm{dB})} - TL_{rp_{i,j}}^{(\mathrm{dB})} + BTS_{ij}^{(\mathrm{dB})} \qquad (7)$$

where $TL_{sp_{i,j}}^{(\mathrm{dB})}$ is the transmission loss from the source to the bottom-surface scattering patch; $TL_{rp_{i,j}}^{(\mathrm{dB})}$ is the transmission loss from the receiver to the bottom-surface scattering patch; $BTS_{ij}^{(\mathrm{dB})}$ is the bistatic bottom-surface strength for FM return signal.

The array beamwidth versus DOA is needed to compute the area of the bottom-surface scattering patch for the return signal [53]. A larger beamwidth causes a larger bottom-surface scattering patch that results in an increase in the reverberation level RL. An increase in the reverberation level reduces the overall SNR.

In Fig. 19, the computed half-power (3 dB) beamwidth (HPBW) as a function of the steering angle assuming rectangular array shading is shown for an array composed of $N = 32$ hydrophones spaced 0.21 m apart [16]. This is the typical configuration of the BENS SLITA, developed by CMRE [66], and towed by the CMRE AUVs during the described experiments at sea.

Example of SNR Computation: Details on the computation of the different terms of the previous equations can be found in [53].

An example of the computation of SNR at different target positions in the COLLAB13 scenario, given the fixed position of the source and the receiver is given in Fig. 20. An isotropic reflector is assumed with a target strength of 20 dB. This is the typical situation when we use an echo repeater as an artificial target (see Section III). The high SNR area at broadside of the array is visible, as the low SNR regions at endfire. The region of high reverberation between the source and receiver is evident with extremely low SNR.

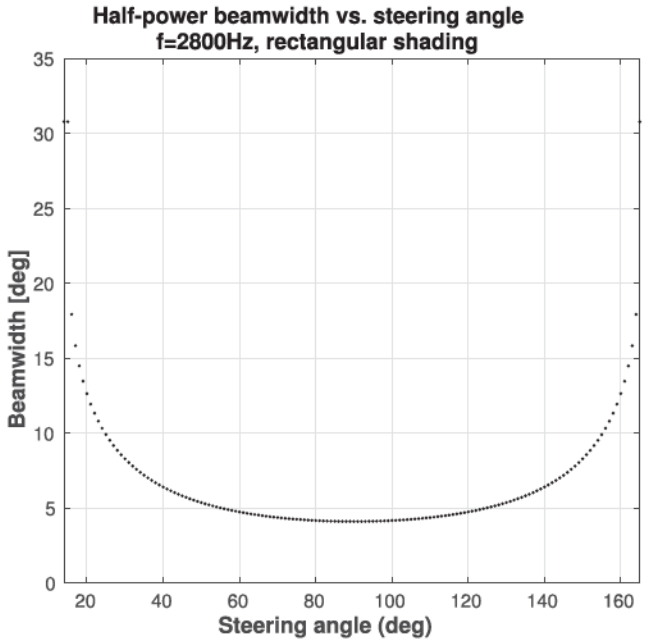

Fig. 19. HPBW as a function of array steering angle at 2800 Hz. The considered array is a linear array with $N = 32$ hydrophones positioned every 0.21 m. 90° of DOA corresponds to broadside, 0/180° correspond to endfire. The beamwidth varies from approximately 4.5° on broadside to over 30° at endfire.

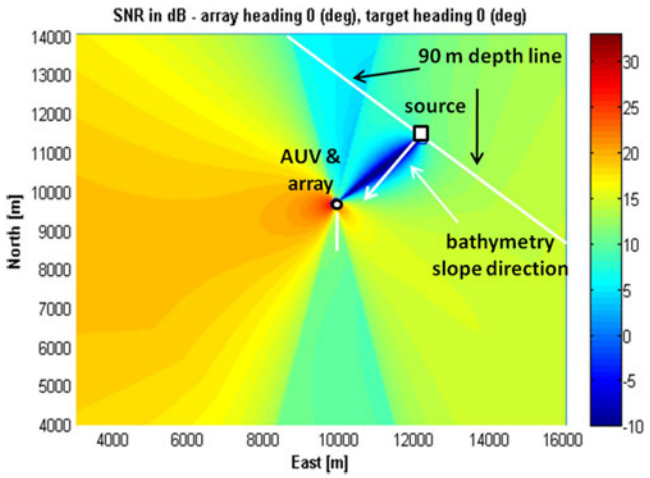

Fig. 20. SNR computed at different target positions by considering the COLLAB13 operation area. We assume an isotropic reflector with a target strength of 20 dB. This is a typical situation when we use an echo repeater as an artificial target. We notice the high SNR area at broadside of the array, as the low SNR regions at endfire. The region of high reverberation between the source and receiver is evident with extremely low SNR.

Target Strength: In a surveillance scenario, one important term is the target strength $TS^{(\mathrm{dB})}$. The target strength models the echo returned by the target [64], specifically the ratio between backradiated acoustic intensity from the object (at a certain reference distance from the acoustic center and in direction of the receiver) and the incident acoustic intensity.

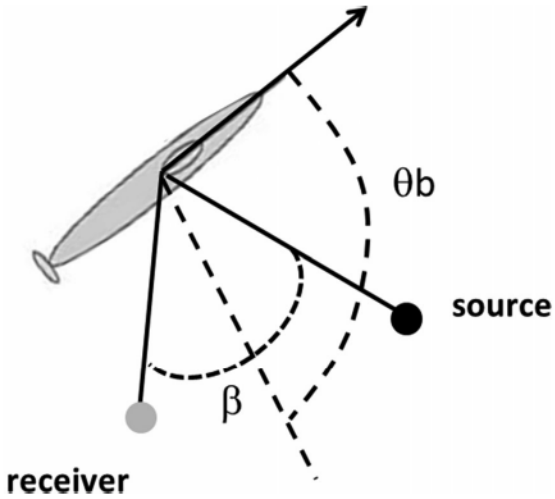

Fig. 21. Angles of importance for the target strength computation. The bistatic aspect angle θ_b is the angle from target's heading to the bisector of the bistatic opening angleβ. The "glint" condition occurs when θ_b is equal to 90° or 270°.

$TS^{(\mathrm{dB})}$ highly depends on the bistatic target aspect angle (see Fig. 21). The bistatic aspect angle is the angle from target's heading to the bisector of the bistatic opening angle β. When θ_b in the figure is $\pm 90°$, the specular condition (or beam aspect), named "glint," [12] occurs, largely increasing the probability of detecting the target.

Our experience is that the acoustic model detailed above tends to overestimate the contact SNR; however, it has a low computational cost and provides a consistent physics-based prediction of the difference between contact SNRs expected in different geometric configurations.

B. Computation of Uncertainties in Acoustic Measurement and Array Heading Angle Knowledge

For a high SNR target, these two quantities can be expressed according to the well-known Woodward relations [50] as

$$\sigma_{\tau_{i,j}} = \frac{k_1 \delta\tau^{FM}}{\sqrt{\mathrm{SNR}_{i,j}}} \tag{8}$$

$$\sigma_{\theta_{i,j}} = \frac{k_2 \delta\theta^{FM}_{i_{i,j}}}{\sqrt{\mathrm{SNR}_{i,j}}} \tag{9}$$

where $\delta\tau^{FM} \propto 1/F^{FM}_{bw}$ is related to the TOA estimation resolution of the FM matched filter in the signal processing chain, $\delta\theta^{FM}_{i_{i,j}}$ is the beamwidth of the FM beamforming, and k_1 and k_2 are proportionality constants.

These relations show that $\sigma_{\tau_{i,j}}$ is proportional to $\delta\tau^{FM}$, a quantity related to the TOA estimation resolution of an FM matched filter in the signal processing chain, and $\sigma_{\theta_{i,j}}$ is proportional to $\theta^{FM}_{i_{i,j}}$ (the beamwidth of the FM beamforming at a certain bearing [16]). Both quantities are inversely proportional to SNR.

Uncertainty in the Knowledge of Array Heading Angle and Array Dynamic Model: σ_{ϕ^R}, the uncertainty in the knowledge of array heading angle, is computed as proportional to the difference between the current AUV heading and the array heading angle as predicted by an array dynamic model.

In general, the description of the array dynamics is hampered by delays in the array angle response, nonlinear effects, and influence of environmental factors such as water currents. By fitting experimental results, we propose a model that is both simple and easy to compute onboard the vehicle and captures the main dynamic properties of the array. Our first-order continuous-time model computing the array heading based on the AUV heading angle is $\hat{\phi}^R = \lambda/s + \lambda\hat{\psi}^R$, with s being the Laplace variable, and λ equal to 0.02 rad/s. To tune the maximum array angular rate, the derivative of the output is saturated at 1.5°/s.

C. Bistatic Measurement Error Computation

We describe here how to compute the bistatic measurement [57]. In the following equations, we omit $\hat{}$ symbol to indicate estimated quantities and (i, j) indices indicating one particular configuration for simplicity

$$\begin{aligned}\sigma^2_{x^T} = {} & \sigma^2_{x^R} + \cos^2(\theta+\phi^R)\sigma^2_{r_r} + r_r^2 \sin^2(\theta+\phi^R)(\sigma^2_\theta + \sigma^2_{\phi^R}) \\ & + 2\cos(\theta+\phi^R)\sigma_{x^R r_r} - 2r_r \cos(\theta+\phi^R) \\ & \times \sin(\theta+\phi^R)(\sigma_{\theta r_r} + \sigma_{\phi^R r_r})\end{aligned} \tag{10}$$

$$\begin{aligned}\sigma^2_{y^T} = {} & \sigma^2_{y^R} + \sin^2(\theta+\phi^R)\sigma^2_{r_r} + r_r^2 \cos^2(\theta+\phi^R)(\sigma^2_\theta + \sigma^2_{\phi^R}) \\ & + 2\sin(\theta+\phi^R)\sigma_{y^R r_r} + 2r_r \cos(\theta+\phi^R) \\ & \times \cos(\theta+\phi^R)(\sigma_{\theta r_r} + \sigma_{\phi^R r_r})\end{aligned} \tag{11}$$

$$\begin{aligned}\sigma_{x^T y^T} = {} & \sigma_{x^R y^R} + \sin(\theta+\phi^R)\sigma_{x^R r_r} + \cos(\theta+\phi^R)\sigma_{y^R r_R} \\ & + \cos(\theta+\phi^R)\sin(\theta+\phi^R)\sigma^2_{r_r} + r_r \cos^2(\theta+\phi^R) \\ & \times (\sigma_{\theta r_r} + \sigma_{\phi^R r_r}) - r_r \sin^2(\theta+\phi^R)(\sigma_{\theta r_r} + \sigma_{\phi^R r_r})\end{aligned} \tag{12}$$

where $\sigma^2_{r_r}$ the variance in range between target and receiver; $\sigma_{x^R r_r}$, $\sigma_{y^R r_r}$ the covariances between, respectively, receiver x or receiver y position error and the target–receiver range error; $\sigma_{\theta r_r}$ the covariance between bearing error and target–receiver range error; $\sigma_{\phi^R r_r}$ the covariance between array heading error and target–receiver range error.

D. Uniform Cost Searching Branch-and-Bound Algorithm

The UCS branch-and-bound algorithm used to explore the decision tree is summarized in the following steps and is described in the Algorithm 1.

1) UCS expansion is executed until M depth level in the tree is reached. The lowest cost decision sequence is used as the first candidate solution, and $J_{\min}$ is initialized to its cost.
   Next actions are repeated until no open unexpanded nodes (of level $< M$) are present in the tree.
2) Perform a UCS on the tree beginning at the lowest cost open node. At each expansion, compare the cost of sequence relative to the children nodes with $J_{\min}$ and prune away the nodes whose relative sequence cost is larger or equal to $J_{\min}$. If the UCS gives a path with a terminal node whose corresponding sequence has a cost lower than $J_{\min}$, set $J_{\min}$ to this value and the best sequence to be this path.

Algorithm 1: UCS branch-and-bound algorithm.

Result: Optimal decision sequence;
Set $J_{\min} = \infty$;
Perform UCS until the terminating depth M is reached;
Set the candidate solution S_c equal to the lowest cost decision sequence of length M;
Set $J_{\min}$ equal to the cost of this sequence;
Store the open unexpanded nodes of depth$< M$ sorted in ascending order of cost in a list L;
while *There is one node in* L **do**
  Expand the first node and remove it from L;
  if *Depth of children nodes* $= M$ **then**
    if *The sequence corresponding to the lowest-cost child node has cost* $< J_{\min}$ **then**
      Set $J_{\min}$ equal to this cost;
      Set S_c equal to this sequence;
    end
  else
    Ordered insert of the children nodes in L to maintain the list
    in ascending order of costs of the relative decision sequences;
  end
  forall *Nodes in* L **do**
    if *The sequence corresponding to a node has a cost* $\geq J_{\min}$ **then**
      Remove the node from L;
    end
  end
end
Return S_c as the optimal decision sequence;

ACKNOWLEDGMENT

The authors would like to thank all the *NRV Alliance* crew members and the CMRE Engineering and Technological Department staff members for their help, willingness, and collaboration during the COLLAB13 cruise.

REFERENCES

[1] G. Ferri *et al.*, "Cooperative robotic networks for underwater surveillance: An overview," *IET Radar, Sonar Navig.*, vol. 11, no. 12, pp. 1740–1761, Dec. 2017.
[2] L. Gerken, *ASW Versus Submarine Technology Battle*. Valencia, CA, USA: Amer. Sci. Corp., 1987.
[3] N. Polmar and E. Whitman, *Hunters and Killers: Volume 2: Anti-Submarine Warfare From 1943*. Annapolis, MD, USA: Naval Inst. Press, Jun. 15, 2016.
[4] D. T. Hughes, F. Baralli, S. Kemna, K. LePage, M. Hamilton, and A. Vermeij, "Collaborative multistatic ASW using AUVs: demonstrating necessary technologies," NATO Undersea Res. Centre, La Spezia, Italy, Tech. Rep. NURC-PR-2009-007, 2009.
[5] T. Curtin, J. Bellingham, J. Catipovic, and D. Webb, "Autonomous oceanographic sampling networks," *Oceanography*, vol. 6, pp. 86–94, 1993. [Online]. Available: http://dx.doi.org/10.5670/oceanog.1993.03
[6] S. T. Tripp, "Autonomous underwater vehicles (AUVs) a look at coast guard needs to close performance gaps and enhance current mission performance," U.S.C.G. Res. Develop. Center, New London, CT, USA, Tech. Rep., 2006, http://www.dtic.mil/dtic/tr/fulltext/u2/a450814.pdf
[7] D. R. Yoerger, M. Jakuba, A. M. Bradley, and B. Bingham, "Techniques for deep sea near bottom survey using an autonomous underwater vehicle," *Int. J. Robot. Res.*, vol. 26, no. 1, pp. 41–54, 2007.
[8] G. Ferri, M. Jakuba, and D. Yoerger, "A novel trigger-based method for hydrothermal vents prospecting using an autonomous underwater robot," *Autonom Robots*, vol. 29, no. 1, pp. 67–83, 2010.
[9] A. Munafò, E. Simetti, A. Turetta, A. Caiti, and G. Casalino, "Autonomous underwater vehicle teams for adaptive ocean sampling: A data-driven approach," *Ocean Dyn.*, vol. 61, no. 11, pp. 1981–1994, 2011. [Online]. Available: http://dx.doi.org/10.1007/s10236-011-0464-x
[10] A. Caiti, V. Calabrò, A. Munafò, G. Dini, and A. Lo Duca, "Mobile underwater sensor networks for protection and security: Field experience at the UAN11 experiment," *J. Field Robot.*, vol. 30, no. 2, pp. 237–253, 2013. [Online]. Available: http://dx.doi.org/10.1002/rob.21447
[11] [Online]. Available: http://liquidr.com/
[12] D. Grimmett, C. Wakayama, and R. Plate, "Multistatic post-track classification using a target strength function," in *Proc. 15th Int. Conf. Inf. Fusion*, Jul. 2012, pp. 2339–2346.
[13] M. J. Hamilton, S. Kemna, and D. T. Hughes, "Antisubmarine warfare applications for autonomous underwater vehicles: The GLINT09 field trial results," *J. Field Robot.*, vol. 27, no. 6, pp. 890–902, 2010.
[14] S. Kemna, M. J. Hamilton, D. T. Hughes, and K. D. LePage, "Adaptive autonomous underwater vehicles for littoral surveillance the glint10 field trial results," *Intell. Service Robot.*, vol. 4, no. 4, pp. 245–258, 2011.
[15] G. Canepa, A. Munafò, M. Micheli, L. Morlando, and S. Murphy, "real-time continuous active sonar processing," in *Proc. OCEANS*, Genova, Italy, May 2015, pp. 1–6.
[16] H. L. V. Trees, *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*. Hoboken, NJ, USA: Wiley, 2002.
[17] A. Munafò and G. Ferri, "An acoustic network navigation system," *J. Field Robot.*, vol. 34, no. 7, pp. 1332–1351, Oct. 2017.
[18] G. Ferri, A. Munafò, R. Goldhahn, and K. LePage, "Towards fully autonomous underwater vehicles in ASW scenarios: An adaptive, data driven AUV mission manager layer," in *Proc. OCEANS*, Genova, Italy, May 18–21, 2015, pp. 1–9.
[19] H. Schmidt, M. Benjamin, S. Petillo, and R. Lum, "Nested autonomy for distributed ocean sensing," in *Handbook of Ocean Engineering*. New York, NY, USA: Springer-Verlag, 2016 pp. 459–480.
[20] K. Rajan, "The role of deliberation (and AI) in marine robots," in *Proc. Workshop Marine Robot., Looking Crystal Ball, Eurocast*, Gran Canaria, Spain, Feb. 11–13, 2013, pp. 1–7.
[21] D. Williams and M. Couillard, "Efficient dense sonar surveys with an autonomous underwater vehicle," in *Proc. 11th Eur. Conf. Underwater Acoust.*, 2012, pp. 676–683.
[22] E. Burian, D. Yoerger, A. Bradley, and H. Singh, "Gradient search with autonomous underwater vehicles using scalar measurements," in *Proc. Symp. Autonom. Underwater Vehicle Technol.*, Monterey, CA, USA, 1996, pp. 86–98.

[23] W. Li, J. Farrell, S. Pang, and R. Arrieta, "Moth-inspired chemical plume tracing on an autonomous underwater vehicle," *IEEE Trans. Robot.*, vol. 22, no. 2, pp. 292–307, Apr. 2006.
[24] A. Alvarez and B. Mourre, "Optimum sampling designs for a glider-mooring observing network," *J. Atmos. Ocean. Technol.*, vol. 29, pp. 601–612, 2012.
[25] G. Ferri, M. Cococcioni, and A. Alvarez, "Mission planning and decision support for underwater glider networks: A sampling on-demand approach," *Sensors*, vol. 16, no. 1, p. 28, 2016. [Online]. Available: http://www.mdpi.com/1424-8220/16/1/28
[26] A. O. Hero and D. Cochran, "Sensor management: Past, present, and future," *IEEE Sensors J.*, vol. 11, no. 12, pp. 3064–3075, Dec. 2011.
[27] A. S. Chhetri, "Sensor scheduling and efficient algorithm implementation for target tracking," Ph.D. dissertation, Dept. School Elect., Comp. Energy Eng., Arizona State Univ., Tempe, AZ, USA, 2006.
[28] M. Kalandros and L. Y. Pao, "Covariance control for multisensor system," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 38, no. 4, pp. 1138–1157, Oct. 2002.
[29] A. S. Chhetri, D. Morrell, and A. Papandreou-Suppappola, "Nonmyopic sensor scheduling and its efficient implementation for target tracking applications," *EURASIP J. Appl. Signal Process.*, vol. 2006, pp. 1–18, 2006.
[30] H. Wang, K. Yao, and D. Estrin, "Information-theoretic approaches for sensor selection and placement in sensor networks for target localization and tracking," *J. Commun. Netw.*, vol. 7, no. 4, pp. 438–449, Dec. 2005.
[31] A. Ryan, H. Durrant-Whyte, and J. Hendrick, "Information-theoretic sensor motion control for distributed estimation," in *Proc. ASME Int. Mech. Eng. Congr. Expo.*, 2007, pp. 725–734.
[32] M. L. Hernandez, T. Kirubarajan, and Y. Bar-Shalom, "Multisensor resource deployment using posterior Cramer-Rao bounds," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 40, no. 2, pp. 399–416, Apr. 2004.
[33] T. Kangsheng and G. Zhu, "Sensor management based on Fisher information gain," *J. Syst. Eng. Electron.*, vol. 17, no. 3, pp. 531–534, 2006.
[34] J. M. Manyika and H. F. Durrant-Whyte, "Information-theoretic approach to management in decentralized data fusion," in *Proc. Meeting Int. Soc. Opt. Eng.*, vol. 202, 1992, pp. 1828–1830.
[35] W. W. Schmaedeke, "Information-based sensor management," in *Proc. Meeting Int. Soc. Opt. Eng.*, vol. 156, 1993, pp. 1–9.
[36] M. Huber, "Probabilistic framework for sensor management," Ph.D. dissertation, Dept. Intelligent Sensor-Actuator-Syst. Lab., Univ. Karlsruhe, Karlsruhe, Germany, 2009.
[37] A. Logothetis and A. Isaksson, "On sensor scheduling via information theoretic criteria," in *Proc. Amer. Control Conf.*, San Diego, CA, USA, Jun. 1999, vol. 4, pp. 2402–2406.
[38] B. Grocholsky, "Information-theoretic control of multiple sensor platforms," Ph.D. dissertation, Dept. Aerospace, Mechatronic Mechanical Eng., The Univ. Sydney, Aust. Centre Field Robot., Darlington, NSW, Australia, 2002.
[39] C. Kreucher, K. Kastella, and A. O. Hero, "A Bayesian method for integrated multitarget tracking and sensor management," in *Proc. 6th Int. Conf. Inf. Fusion*, Cairns, QLD, Australia, 2003, pp. 132–139.
[40] G. Zhang, S. Ferrari, and M. Qian, "An information roadmap method for robotic sensor path planning," *J. Intell. Robot. Syst.*, vol. 56, no. 1, pp. 69–98, 2009.
[41] L. Paull, S. Saeedi, M. Seto, and H. Li, "Sensor-driven online coverage planning for autonomous underwater vehicles," *IEEE/ASME Trans. Mechatronics*, vol. 18, no. 6, pp. 1827–1838, Dec. 2013.
[42] D. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA, USA: Athena Scientific, 2007.
[43] S. Page, "Multiple-objective sensor management and optimisation," Ph.D. dissertation, Dept. Faculty Eng., Sci. Math. School Electron. Comput. Sci., Univ. Southampton, Southampton, U.K., 2009.
[44] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. Cambridge, MA, USA: MIT Press, 2005.
[45] A. Poulsen, D. Eickstedt, and J. Ianniello, "Bearing stabilization and tracking for an AUV with an acoustic line array," in *Proc. IEEE/MTS Oceans Conf.*, Boston, MA, USA, 2006, pp. 1–6.
[46] D. P. Eickstedt, M. R. Benjamin, H. Schmidt, and J. J. Leonard, "Adaptive control of heterogeneous marine sensor platforms in an autonomous sensor network," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2006, pp. 5514–5521.
[47] A. Gadre *et al.*, "Cooperative localization of an acoustic target using a team of AUVs," in *Proc. IEEE/OES Auton. Underwater Veh. Conf.*, Woods Hole, MA, USA, 2008, pp. 64–71.
[48] D. Maczka, D. Spinello, D. J. Stilwell, A. S. Gadre, and W. L. Neu, "Coordinated tracking of an acoustic signal by a team of autonomous underwater vehicles," in *Proc. AUVSI's Unmanned Syst. North America*, Washington, DC, USA, 2009, pp. 1–18.
[49] K. LePage, M. Hamilton, and S. Kemna, "Autonomous underwater vehicles for active multistatic undersea surveillance," in *Proc. 4th Int. Conf. Underwater Acoust. Meas., Technol. Results*, Kos, Greece, Jun. 2011, pp. 1471–1478.
[50] R. H. K. Lum, "Integrated perception, modeling, and control paradigm for bistatic sonar tracking by autonomous underwater vehicles," Ph.D. dissertation, Dept. Mech. Eng., Massachusetts Inst. Technol., Cambridge, MA, USA, 2012.
[51] R. Goldhahn, K. LePage, P. Braca, A. Munafò, and G. Ferri, "Environmentally sensitive AUV behaviors for collaborative multistatic surveillance networks," in *Proc. Underwater Acoust. Signal Process. Workshop*, West Greenwich, RI, USA, Oct. 16–18, 2013, pp. 1–6.
[52] A. Munafò, G. Ferri, and K. LePage, "AUV active perception: Exploiting the water column," in *Proc. OCEANS*, Aberdeen, U.K., 2017, pp. 1–8.
[53] C. H. Harrison, "Fast bistatic signal-to-reverberation-ratio calculation," *J. Comput. Acoust.*, vol. 13, no. 2, pp. 317–340, 2005.
[54] M. Benjamin, P. Newman, H. Schmidt, and J. Leonard, "An overview of MOOS-IvP and a users guide to the IVP helm autonomy software," Massachusetts Inst. Technol., Cambridge, MA, USA, Tech. Rep. MIT-CSAIL-TR-2010-041, 2010.
[55] S. Coraluppi and C. Carthel, "Distributed tracking in multistatic sonar," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 3, pp. 1138–1147, Jul. 2005.
[56] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation With Applications to Tracking and Navigation*. Hoboken, NJ, USA: Wiley, 2001.
[57] S. Coraluppi, "Multistatic sonar localization," *IEEE J. Ocean. Eng.*, vol. 31, no. 4, pp. 964–974, Oct. 2006.
[58] R. Korf, "Artificial intelligence search algorithms," in *CRC Handbook of Algorithms and Theory of Computation*. Boca Raton, FL, USA: CRC Press, 1998, pp. 36.1–36.20.
[59] T. M. Breuel, "A comparison of search strategies for geometric branch and bound algorithms," in *Proc. Eur. Conf. Comput. Vis.*, 2002, vol. 4, pp. 837–850.
[60] G. Ferri, A. Munafò, R. Goldhahn, and K. LePage, "A non-myopic, receding horizon control strategy for an AUV to track an underwater target in a bistatic sonar scenario," in *Proc. 53rd IEEE Conf. Decision Control*, Los Angeles, CA, USA, Dec. 15–17, 2014, pp. 5352–5358.
[61] G. Ferri, A. Munafò, R. Goldhahn, and K. LePage, "Results from COLLAB13 sea trial on tracking underwater targets with AUVs in bistatic sonar scenarios," in *Proc. Oceans*, Sep. 2014, pp. 1–9.
[62] 2015. [Online]. Available: http://www.diamondsystems.com/files/binaries/etx945-datasheet.pdf
[63] Y. Bar-Shalom and X. Li, *Multitarget-Multisensor Tracking: Principles and Techniques*. Storrs, CT, USA: Yaakov Bar-Shalom, 1995.
[64] R. P. Hodges, *Underwater Acoustics Analysis, Design and Performance of Sonar*. Hoboken, NJ, USA: Wiley, 2010.
[65] H. L. V. Trees, *Radar-Sonar Processing and Gaussian Signals in Noise: Part III of Detection, Estimation, and Modulation Theory*, I. Hoboken, NJ, USA: Wiley, 2001.
[66] A. Maguer, "Receiving and transmitting acoustic systems for AUV/gliders," NURC, La Spezia, Italy, Tech. Rep. NURC-PR-2009-004, 2009.

Gabriele Ferri (M'08) received the M.S. degree (*Laurea*) in computer engineering (with honors) from the University of Pisa, Pisa, Italy, in 2003 and the Ph.D. degree in robotics jointly from the Scuola Superiore S. Anna, Pisa, Italy, and IMT Advanced Studies, Lucca, Italy, in 2008.

From 2003 to 2005, he was a Software/System Consultant Engineer at the WASS company (a Finmeccanica group company), Livorno, Italy, developing a new system of control and guidance for a light-weighted torpedo. In 2007, he was a Visiting Researcher at the Woods Hole Oceanographic Institution, Woods Hole, MA, USA, working on deep water hydrothermal vents prospecting. After a period as a Postdoctoral Investigator with Scuola Superiore Sant Anna where he was the Project Leader of the HydroNet project, in 2012 he became a Research Scientist at CMRE, La Spezia, Italy, for investigating novel approaches for the autonomy of AUV networks in ASW applications. His research interests include robotic systems for surveillance and exploration, control theory, and robot navigation topics.

Dr. Ferri has been the Technical Director of the Student Autonomous Underwater Vehicle Challenge-Europe robotics competition since 2014, and was the Director of the European multidomain robotics competitions such as euRathlon 2014 and 2015 (www.eurathlon.eu) and the European Robotics League Emergency Robots 2017.

Andrea Munafò (M'09) received the B.Sc. degree in computer science engineering, the M.Sc. degree in automation engineering, and the Ph.D. degree in automation, robotics and bioengineering, all from the University of Pisa, Pisa, Italy, in 2002, 2005, and 2009, respectively.

From 2009 to 2013, he was a postdoctoral student with the interuniversity Centre of Integrated Systems for the Marine Environment, Italy. He was with the NATO STO Centre for Maritime Research and Experimentation, La Spezia, Italy, from 2013 to 2016 as a Research Scientist, working in the field of cooperative underwater robotics and underwater communications. Since 2017, he has been a Senior Scientist and an Engineer with the Marine Autonomous and Robotics System Department, National Oceanography Center, Southampton, U.K., where he is leading the development of novel onboard autonomy solutions for marine vehicles. His research interests include underwater robotics, underwater acoustics and sonar systems, adaptive planning, and oceanographic sampling.

Kevin D. LePage received the B.S. degree in naval architecture and marine engineering from Webb Istitute in 1983 and the M.S. degree and the Ph.D. degree in ocean engineering from Massachusetts Institute of Technology, Cambridge, MA, USA, in 1987 and 1992, respectively.

He is a Principal Scientist and the Program Manager for the Cooperative Anti-Submarine Warfare, NATO Science and Technology Organization Centre for Maritime Research and Experimentation, La Spezia, Italy, where his work is directed toward capturing requirements for future unmanned solutions for ASW capability under funding from NATO Allied Command Transformation. While with CMRE, he has been a Scientist in Charge of five scientific sea trails and the CMRE component of three NATO naval ASW exercises where unmanned systems were deployed and tested.

Dr. LePage was the recipient of the NATO STO Scientific Achievement Award in 2017 for the Development and Demonstration of Networked Autonomous ASW as a Team Leader. He is a Fellow and Past Chair of the Underwater Acoustics Technical Committee of the Acoustical Society of America.

Document Data Sheet

Security Classification		Project No.

Document Serial No.	Date of Issue	Total Pages
CMRE-PR-2019-049	May 2019	21 pp.

Author(s)

Gabriele Ferri, Andrea Munafò, Kevin D. LePage

Title

An autonomous underwater vehicle data-driven control strategy for target tracking

Abstract

This paper presents a data-driven approach to control the movement of autonomous underwater vehicles (AUVs) operating as receivers of a multistatic sonar surveillance network. The algorithm adopts a nonmyopic receding horizon policy to control the AUV heading, with the objective of minimizing the expected estimate error in the target location produced by the onboard tracker. Minimizing this error is of the utmost interest in target state estimation since it is one way of maintaining tracks. The proposed policy builds on an automated perception module which produces a target track and on an acoustic model to estimate the target measurement signal-to-noise ratio. Driven by the mission objectives, it explores the best decisions for the AUV by evaluating the evolution of the tactical situation over a future time window. Nonmyopic approaches can offer increased performance if compared to myopic ones. The downside is an increase of the computational burden. Methods are described to simplify and solve effectively the resulting decision tree to allow the execution of these kinds of computational intensive algorithms onboard autonomous vehicles. At the same time, they aim at producing AUV paths robust to possible misleading target measurements, which might cause meandering trajectories. These are crucial factors to use data-driven strategies for AUVs which must work in real applications. Results from the COLLAB13 sea trial are reported and discussed. They show both the feasibility of executing the algorithm in realtime on an onboard computer and the benefits of using the proposed approach over conventional predefined paths. These results are, to the best of our knowledge, the first successful demonstration at sea of a complex nonmyopic algorithm running in realtime onboard AUVs in a realistic multistatic littoral surveillance scenario.

Keywords

Antisubmarine warfare (ASW), autonomous underwater vehicle (AUV), autonomy, data-driven strategies, nonmyopic control, optimal control, surveillance

Issuing Organization

NATO 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