



SCIENCE AND TECHNOLOGY ORGANIZATION
CENTRE FOR MARITIME RESEARCH AND EXPERIMENTATION



Reprint Series

CMRE-PR-2017-006

A data-driven control strategy in synergy with continuous active sonar for littoral underwater surveillance

Gabriele Ferri, Andrea Munafò, João Alves, Kevin LePage

November 2017

Originally presented at:

OCEANS'16 MTS/IEEE Monterey

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, 2016. 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.

A Data-Driven Control Strategy in Synergy With Continuous Active Sonar for Littoral Underwater Surveillance

Gabriele Ferri, Andrea Munafò, João Alves, Kevin LePage

Abstract—In this work, we describe a data-driven Mission Management Layer (MML) running on-board AUVs which manages the phases of a littoral surveillance mission and exploits the characteristics of Continuous Active Sonar (CAS) signal processing. The MML selects for further investigation the tracks which are likely originated by a target. In this case, the MML launches a receding horizon, non-myopic control algorithm which controls the AUV’s heading to improve the tracking performance to ease the target classification. The algorithm minimises the expected target position estimation error over a prediction time window by achieving a trade-off amongst different objectives: keeping the target at broadside, reducing the distance to the target, avoiding areas of high reverberation and searching for geometric configurations with low bistatic target localisation error.

We report at-sea experiments obtained during the LCAS15 sea trial, which demonstrated, for the first time, that the proposed autonomy architecture can be executed together with real-time Continuous Active Sonar (CAS) processing on-board the AUVs. CAS has recently gained interest for littoral Anti-Submarine Warfare, since it offers the promise of multiple detections per waveform cycle. This can potentially improve the quality/length of tracks, thus increasing the adaptive behaviour’s performance, which, in turn, can increase the detection and tracking capabilities of the processing chain.

I. INTRODUCTION

The capabilities of today’s AUVs in terms of precise navigation, autonomy and endurance make them appealing assets for littoral surveillance and Anti-Submarine Warfare (ASW) scenarios. Traditionally, the task of ASW has been carried out by means of sensors such as sonobuoys and submarines or frigates with towed arrays [1]. 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. Existing traditional approaches are expensive and manpower intensive.

The envisioned solution we have been pursuing at the NATO-STO Centre for Maritime Research and Experimentation (CMRE) is the use of sensorised AUVs acting as autonomous mobile nodes in a multistatic network [2], [3]. AUVs can provide effective ASW capabilities at a fraction of the cost of traditional assets. In the CMRE multistatic sonar system, a sonar source (transmitter), which may be located on a stationary buoy or ship deployable, transmits a sonar signal (ping) which reflects from objects and is collected by multistatic receivers, the Autonomous Underwater Vehicles (AUVs) towing an array in this case. Multi-

static sonar systems have the potential to greatly increase ASW coverage and performance [2], [3]. The possibility to use multiple receivers and sources generates different geometric distributions of source-target-receiver increasing the probability of detection and classification for a target. The acoustic data received by the array are processed [4] to create bearing/range contacts which are fed into an on-board tracker [3], [5] based on a kinematic model of the target of interest. The tracker combines (spatially) related contacts over time to produce tracks.

Multiple nodes provide the system *redundancy* increasing the *robustness* of the network to failures and can *share information* between each other. Sharing information is vital to create a *common tactical picture* and to fuse the collected data. Contacts [6] or tracks [7] can be fused to improve the tracking/classification performance and to identify clutter-generated or ambiguous tracks. These are key points in real operations due to the difficulties of the littoral scenarios from the point of view of track generation and classification.

High clutter is in fact present along with noise from the surface produced by passing ships and boats. Littorals are characterised by poor sound propagation conditions and the sound speed profile usually changes during the day modifying the acoustic channel features. In a typical littoral scenario, several clutter-generated tracks may be present simultaneously. Some of these tracks may also be persistent and last for several pings. In addition, the target may not be observable for some time due to particular sound speed profile conditions or low probability of detection. In addition, the presence of “ghost” tracks due to the port-starboard ambiguity, typical of linear arrays [3], exacerbates the problem.

In the addressed scenario, tracking a possible target for long time periods is of the utmost importance for its correct classification. The high number of present tracks raises the question of how selecting one or few tracks for further investigation. Even if (real-time) classification [8], [9] can provide some results, the issue of correctly selecting which tracks are likely to be related to a target still remains open.

To address these issues, we envision two approaches: increasing the autonomy of the nodes and adopting novel waveform processing which offers the promise of improved detection and tracking.

The degree of autonomy of the vehicles is crucial, above all considered the limited communications bandwidth and range of the underwater sound channel. These factors make communications with the vehicles sparse and sometimes impossible. To be really effective, AUVs need to make decisions autonomously on the basis of the acquired data

Gabriele Ferri, Andrea Munafò, João Alves and Kevin LePage are with the Research Department, NATO Centre for Maritime Research and Experimentation, Viale San Bartolomeo, 400, 19126 La Spezia (SP), Italy Gabriele.Ferri@cmre.nato.int

and of the evolving tactical scene.

Recently, CMRE has been testing Continuous Active Sonar (CAS) in littoral scenarios as an alternative to the traditionally used Pulsed Active Sonar (PAS). Unlike PAS, which listens for echoes in between short-burst transmissions, CAS detects echoes amidst the continual interference from source(s) transmitting with nearly 100% duty cycle. The potential advantage of CAS is an increased number of continuous detection opportunities, leading to improved target detection, localisation, tracking, and classification [10]–[12].

In this work, after describing the control system which manages the autonomy of the CMREs Ocean Explorer (OEX) AUVs, we report the results at sea of a non-myopic, data-driven control algorithm using as input the data produced by the CAS processing running on-board the vehicles. This algorithm, proposed in [3], [13] and previously tested at sea with PAS processing, receives a track as input and controls the heading of the vehicle to improve the target tracking performance. Results from LCAS15 sea trial show how using autonomy and a data-driven behaviour in synergy with an advanced signal processing technique such as CAS can be beneficial for target tracking.

II. AUTONOMY AS THE MAIN DRIVING FACTOR

In littoral surveillance missions two different levels of autonomy can be assumed. The high level autonomy encompasses the AUV decision making process based on the tactical scene (e.g. which tracks to select for further inspection, etc.). Once the high level decisions are made, a lower level of autonomy can be identified. This level consists in executing the actions of the vehicle. At this level, for instance, the vehicle adapts its path to optimise some objective functions of interest.

A. The Mission Manager Layer

The CMRE control scheme shown in Fig. 1 follows this concept. An adaptive, data-driven Mission Management Layer (MML) [14] is proposed. MML runs on-board the vehicles and manages the high level autonomy. The MML receives the tracks and contacts produced by the signal processing chain, selects in real-time which tracks are interesting to be investigated and commands the vehicle Control Layer operations.

To make effective decisions on the tracks, a metric is needed to quantify the track quality, defined as the probability of existence of the target corresponding to the track. In [14] a *track scoring* is proposed to quantify the track quality. The scoring is based on the quality of the measurement-to-track associations. The method uses an acoustic model and the kinematic features of the tracks and does not rely on the knowledge of often difficult to estimate parameters such as the probability of detection. The track score is computed in real-time and is used to classify the tracks as confirmed (worthy of being further investigated by the AUV). The threshold of the classifier is selected by analysing the historical data collected at sea. Experimental results show that tracks related to a target are more likely to produce

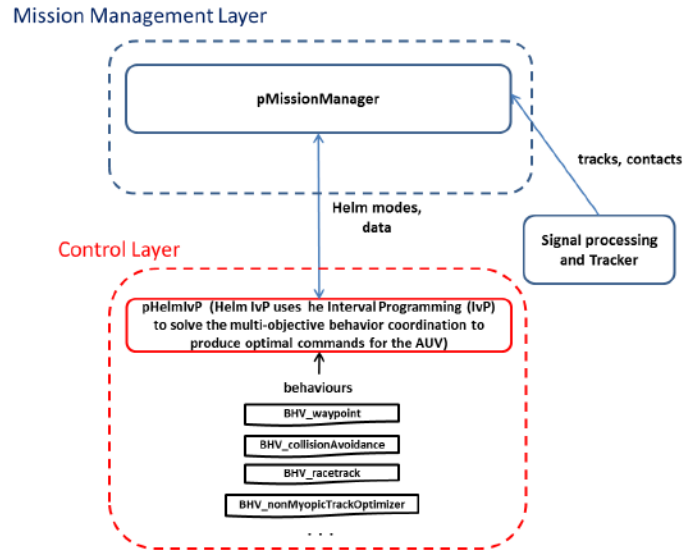


Fig. 1. Diagram of the control architecture for the OEX AUV’s autonomy.

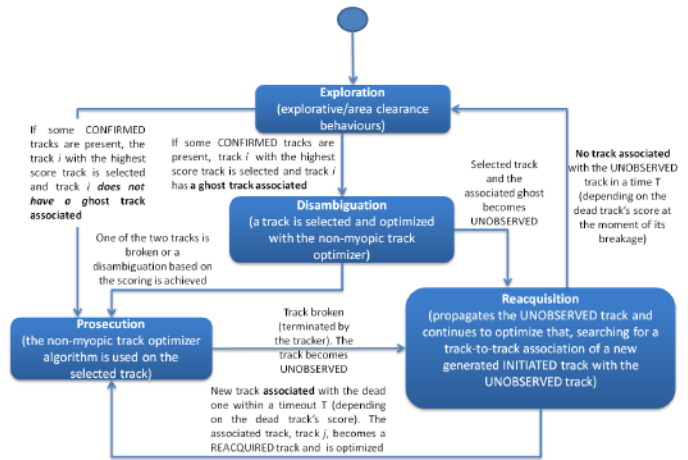


Fig. 2. MML state diagram.

a larger increase of scoring in time with respect to tracks produced by clutter.

The MML is characterized by several logical states covering the various phases of an ASW mission (see Fig. 2): exploration of an area of interest to search for cues of the target (candidate tracks); disambiguation between a track and its “ghost” [14] once a confirmed track is present; optimisation of a selected confirmed track by using *ad hoc* control algorithms (we call this task track prosecution) and target reacquisition if a confirmed track breaks. The MML controls the switching between one phase to another (with a change of the active set of behaviours) in function of the occurring events. Specifically, the MML analyses the history of the detections and tracks produced by the signal processing chain to select amongst them the candidates to be prosecuted.

The MML finds a balance between the exploration objective and the exploitation of cues about possible targets. It drives the mission by prosecuting only the tracks likely being target-originated increasing the benefits of data-driven

approaches.

The architecture shown in Fig. 1 is implemented on vehicles in a MOOS-IvP [15] framework. 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 subscribes to receive and publish variables from/to a database (MOOSDB). The MML controls the vehicle Control Layer which is in charge of managing the different behaviours which control the vehicle operations. The Control Layer is managed by the pHelmIvP [15], a MOOS application which enables behaviour-based autonomy. Behaviours can run simultaneously and can be grouped into behaviour sets, which are active based on certain conditions. IvP, a mathematical interval programming technique, combines the objective functions produced by active behaviours to determine a optimal solution for each domain [15]. The IvP Helm, typically running four times per second, is able to reconcile the different active behaviours to produce the commands (speed, heading and depth commands) for the frontseat controller which controls the AUV actuators.

The proposed layered architecture decouples the planning/deliberation activities managed by the MML from the executive actions conducted by the Control Layer.

B. A non-myopic, receding horizon track prosecution control algorithm

A non-myopic control algorithm [3], [13] is launched to prosecute the track selected by the MML. The candidate track is used to control the AUV navigation to achieve favourable target-source-receiver geometries for target tracking. Specifically, a receding horizon policy is adopted to control the AUV heading angle to minimise the expected target position estimation error of a tracking filter by considering the future positions of the source, receiver and target. To compute the expected error, we use bistatic contact-localisation statistics and we add environmental information by utilising an acoustic model. Minimising this error is typically of the utmost interest in target state estimation since it assists maintaining tracks, which in turn increases the probability of correct classification. At each step, the optimal sequence is computed for the planning horizon by solving the resulting decision tree. According to the receding horizon paradigm, the first heading decision is executed while the others are discarded. This approach proves robust against possible misleading measurements since it considers at every computation step the information brought by a new measurement. A branch and bound approach is adopted to solve the optimisation, making its execution feasible on low computational power on-board computers. The tree is also simplified to avoid AUV's sharp maneuvers that, by causing the bending of the array, can deteriorate the array processing performance. Results at sea with PAS processing [3] demonstrate that the algorithm can be executed in real-time on CMRE vehicles outperforming a myopic approach. The generated AUV's paths are also smooth enough to avoid the deterioration of array processing.

The effectiveness of the proposed approach shows how data-driven schemes, in which the AUVs modify autonomously their path based on the collected data, offer some operational benefits [5], [6], [16] not easily achievable by non-autonomous vehicles. It also demonstrates how planning over a future horizon can be of fundamental importance.

Data-driven approaches, however, do not guarantee the uniform coverage provided by pre-designed tracklines, traditionally used to plan the vehicle missions. Pre-designed tracklines, usually used to conduct oceanographic/military [17] surveys, provide uniform coverage but at a fixed resolution. In our system, the MML finds a trade-off between the required needs of area exploration and the adaptation of the AUV trajectory on the output of the signal processing chain.

III. IMPLEMENTING CAS PROCESSING CHAIN ON AUVS

The recent advances in transducer and computing technology have been made the use of CAS an appealing approach for ASW [11] in shallow waters.

The features of CAS in shallow waters are under investigation since it can potentially improve ASW sonar performance in several ways [11], [12]. The total transmitted energy can be increased by extending pulse duration with a constant source level, thus increasing the target detection range. Furthermore, a certain target energy at sea can be achieved by extending the pulse duration with a lower source level [10]. This has several advantages such as reduced power and size requirements for the amplifiers and transducers; reduced risk of transducer cavitation at shallow depth; a reduction in transmission non-linearities and increased transmission bandwidth available for most transducers [10].

However, transducers may not be able to achieve the performance required for a High Duty Cycle (HDC) waveform in practice. Additionally, the ideal processing gain may not be achieved for waveforms characterised by high time-bandwidth products, especially in shallow waters where sound propagation is complex and time coherence often short. Rather than processing the entire pulse coherently, suitable waveforms such as linear FM (LFM) sweeps can be segmented and treated as a series of short, non-interfering pulses, which are processed individually. This type of sub-band processing is the basis of another performance improvement that CAS offers: increasing the update rate of sonar contacts while maintaining the same Pulse Repetition Interval (PRI) and corresponding search radius [12]. Though the probability of detection and ranging accuracy may be lower for HDC sonar than for PAS (due to reductions in source level and processed bandwidth), the great advantage of sub-band processing is an important increase in the measurement update. This can contain the growth of the target's area of uncertainty [11] within a kinematic tracker, opening new horizons for tracking algorithms and robot autonomous decision making.

The major difficulty in executing CAS sub-band processing on-board an AUV is the increase of computational power required to process the several sub-bands at the same PRI of PAS.

Work has been done at CMRE [18] to extend the CMRE signal processing chain for PAS, *CAS-CAINPro*, to CAS, and to increase its execution speed. By reorganising the code and using vectorization and multi-threading to fully exploit the computer on-board the CMRE AUV (a Seco S9920-5000-1110-C0 pico-ITX SBC - AMD eKabini, from Seco, with a quad core processor @2.0 Ghz with 8GB RAM), the new signal processing chain, named *CAS-CAINPro* is capable of processing the several sub-bands in a time compatible with on-board processing (e.g. 9 sub-bands processed in ~ 8 s [18]). The produced clusters in each sub-band are passed to the Distributed Multi-Hypothesis Tracker (DMHT) running on-board the vehicles [3], [5]. The produced tracks are then passed to the Mission Manager Layer which scores and classifies the tracks. When a track is considered confirmed, the non myopic track optimiser is triggered and starts to control the vehicle navigation.

IV. THE LCAS15 SEA TRIAL

The LCAS15 experiment was carried out in October 2015, in the Ligurian sea, in front of Palmaria Island, Italy. Part of the experiment was devoted to test the AUV autonomous decisions making and the performance of the non-myopic track optimiser with on-board CAS waveforms processing.

The deployed CMRE ASW network is visible in Fig. 3. The vehicles used as receiving nodes of the network are two Ocean Explorer (OEX) AUVs. OEXs are approximately 4.3 m long and 0.53 m wide. The endurance depends on the payload. In usual operative conditions they can reach 16 hrs of operations at a speed of 1 m/s. AUVs communicate between each other and with the Command and Control (C2) centre via a 7/17 kHz Evologics low-frequency modem. The OEX AUVs are both deployed with the BENS slim towed array (SLITA) [19]. 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). The network infrastructure is composed of gateway buoys and Wave Gliders [20] surface vehicles. These nodes act as communication relays managing the communication under and above the water by using their acoustic modems and radio modules. They receive messages via radio from the C2 centre broadcasting them underwater to the vehicles. Vice-versa, they receive acoustic messages from the vehicles and transmit them to the C2 centre. As the acoustic source, we used the mid-frequency ATLAS source towed by the NRV Alliance. The target was represented by the CRV Leonardo which towed the CMRE Echo Repeater (E/R) [5]. 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. All the assets were deployed from the NATO Research Vessel (NRV) Alliance. The vessel operated as the C2 centre during the experiments.

We report here the result of Groucho OEX's mission of October 3, when the AUV was controlled by the MML.

During the sea trial, a 20 s, 1800-2600 Hz LFM, was transmitted using the mid-frequency ATLAS source towed by NRV Alliance. The PRI was 20 s, and the *CAS-CAINPro* processed the received signal in 9 sub-bands. The MML started to control Groucho's operations, upon activation via an acoustic message, at ping 508, while Groucho was following a pre-defined trackline heading to north-east. The vehicle's operations were controlled by MML for a total of 4 hours until the end of the experiment.

A. Results

In Fig. 4, the situation at ping 530 is shown. CRV Leonardo (its path is indicated in cyan and the label *TGT* shows its current position) is heading to south-east at a speed of 3.5 m/s. The ATLAS source (grey label *TX*) is towed by the NRV Alliance (black line) at a speed of 2 m/s towards south-east along a path parallel to Leonardo's. The AUV path is indicated in light green (the label *RX* shows the position of the AUV at the current ping). Groucho was moving at 1 m/s towards north-east along a fixed racetrack when the MML was activated at ping 508. The exploration phase consisted in following predefined tracklines.

In order to evaluate the effectiveness of the non-myopic controller, we also compute the path of the AUV had it proceeded along the racetrack path. The racetrack case is analysed to compare the tracking performance of the non-myopic controller against a non-adaptive trajectory. This allows to better evaluate the advantages and the improved performance achieved by using the non-myopic algorithm. In the figures, we report with the label *RX_SIM* in dark green the position of the simulated path of the AUV. This path continues the fixed racetrack Groucho was following in the exploration phase.

After its activation at ping 508, the MML evaluates the scoring of the tracks and confirms two of them, namely those with id 107066 and 107094, one of them being the "ghost" of the other. This means that one of them is created by the left-right ambiguity of the linear arrays [14]. The disambiguation state is therefore entered and one of the two tracks (in this case the first one) is selected to be prosecuted by the non-myopic track optimiser. In the experiment, the non-myopic optimiser was activated every two PRIs (every 40 s). The AUV starts maneuvering by turning towards north-east. The maneuver allows the track scoring to discriminate between the real and the ghost track. The AUV's maneuver makes the quality of the measurement-to-track associations for the real track better than the ghost track's, as explained in [14]. This increases the score for the real track and allows the MML to correctly select the track 107066 as confirmed entering the prosecution state. Groucho continues to prosecute track 107066 (see Fig. 4) by heading towards north-east until the track breaks at ping 548. MML enters the reacquisition state and associates a newborn track, the track 117424, to the propagation in time of the 107066 reacquiring the target at ping 549 [14]. Then, it continues to prosecute the track 117424 until ping 575 when the track breaks. After an

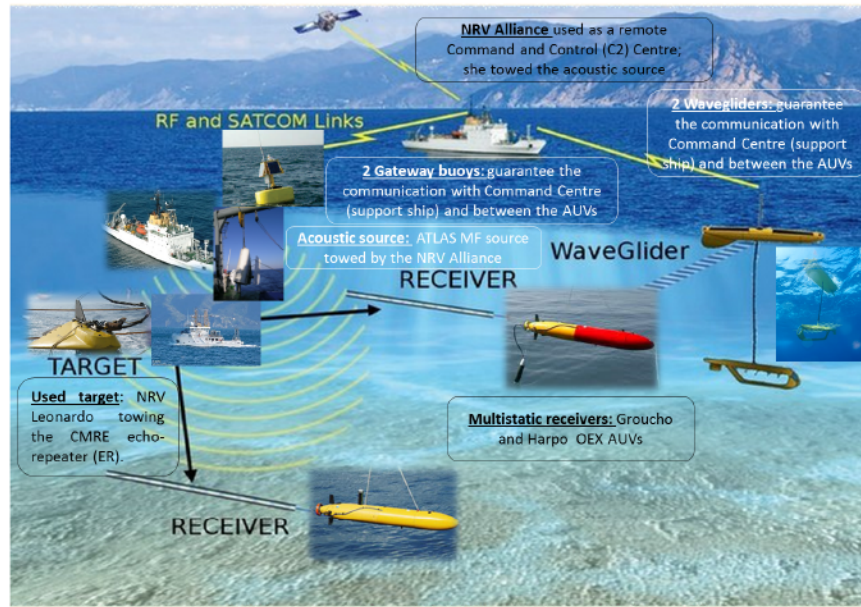


Fig. 3. CMRE cooperative ASW multi-static network, as deployed during LCAS15 trial.

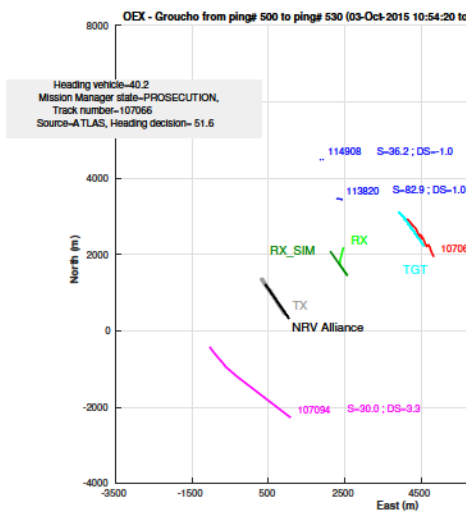


Fig. 4. Ping 530 - Positions of Groucho OEX in the experiment (light green), simulated racetrack (dark green), target (cyan), NRV Alliance (black) and tracks with scoring. Blue tracks are the initiated ones, while the red one is the confirmed one prosecuted by the AUV.

attempt of target reacquisition, the MML comes back to the exploration state coming back on the trackline (see Fig. 5).

The described results show how the MML can manage the autonomy of the AUV in an ASW mission. Starting from the exploration phase, it starts investigating a confirmed track. In case of presence of a “ghost” track, first it accomplishes the disambiguation without the need of harsh maneuvers and selects the track originated by the target. Then, the MML prosecutes it reacquiring the target if the track breaks. Groucho was able to prosecute a track originated by the target for a long time (~ 70 pings) which is fundamental for the target classification. The MML is essential to select the most interesting tracks amongst the many created by the

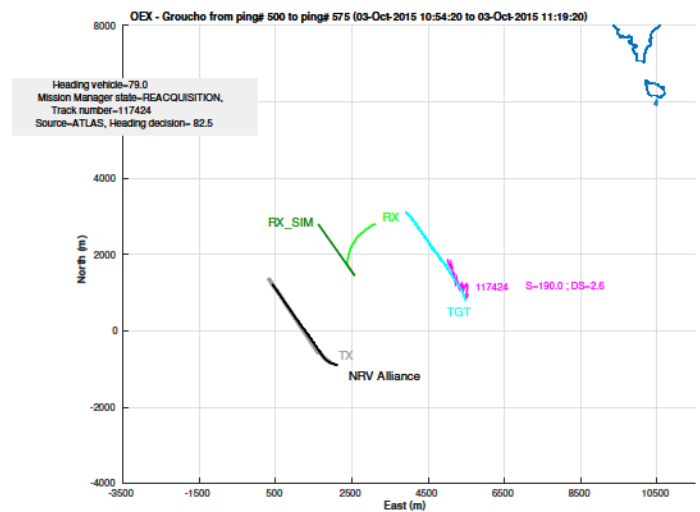


Fig. 5. Ping 575 - Positions of Groucho OEX in the experiment (light green), simulated racetrack (dark green), target (cyan), NRV Alliance (black) and tracks with scoring. In magenta, the track broken and propagated in the future for a possible target reacquisition.

tracker (in 4 hours, it selected 11 confirmed tracks, 8 of them being related to the target). This allows to really exploit the data-driven, non-myopic control policy, avoiding the AUV triggering a data-driven approach on a large set of tracks which may lead to inconsistent and unsatisfactory situations.

The non-myopic behaviour shows several benefits which are crucial for target tracking. The AUV maneuvers to optimise the expected tracking error by finding a trade-off amongst keeping the target at broadside¹, getting closer to it and searching for source-receiver-target geometries

¹ for “broadside” we mean a direction of arrival of 90 degrees relative to the array; a target at broadside improves the quality and SNR of measurements

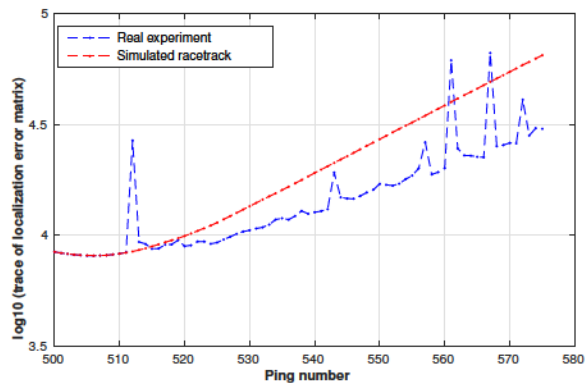


Fig. 6. $\log_{10}(\text{trace}(\mathbf{R}_{true}))$ for the experimental data and for the simulated racetrack (no adaptation of the path). \mathbf{R}_{true} is the localisation error matrix computed by considering the real target position.

favourable from an SNR perspective [3]. This translates in a reduction of the localisation error as shown in Fig. 6. Here we compute the \mathbf{R}_{true} matrix for the data from the experiment and the simulated racetrack. \mathbf{R}_{true} is the covariance matrix of the localisation error given the true position of the target [3]. We report the $\log_{10}(\text{trace}(\mathbf{R}_{true}))$ quantity. This is a measure of the estimated localisation error. The localisation error is lower in the adaptive case and increases at a lower rate than racetrack's. The peaks in error visible in the figure are due to the increase in the uncertainty of the array heading knowledge due to its bending. At the end of the prosecution, in the racetrack case the average localization error in the x-y coordinates has a standard deviation of 180.1 m while in the non-myopic case the standard deviation is 122.75 m. With the use of the non-myopic algorithm, we achieve a reduction of 32% of the localisation error. Accurate measurements imply both a lower error in the target position estimate and, more importantly, tracks likely with a longer life time. The adaptive behaviour achieves a trade-off between getting closer to the target and keeping it at broadside. While in the racetrack case the Direction Of Arrival (DOA) continues to increase since the AUV's movement causes the target moving towards the array endfire, in the adaptive case the DOA remains limited, at around 60 degrees. Groucho, by staying closer to the target, position herself in a better position by considering possible future maneuvers of the target. Finally, the movement is also effective from an SNR perspective (see Fig. 7). The SNR for the racetrack case decreases due to the combined effect of the increasing distance and movement towards the array endfire. In the adaptive case, the SNR remains higher. This is highly beneficial since this not only reduces the localisation error, but also increases the probability of detection.

V. CONCLUSION

In this paper, we described the use of a data-driven Mission Management Layer (MML) running on-board the vehicles managing all the phases of a littoral surveillance mission in the presence of CAS processing. In particular, the MML finds a trade-off between exploration of the area and exploitation of the cues (tracks) by selecting the most

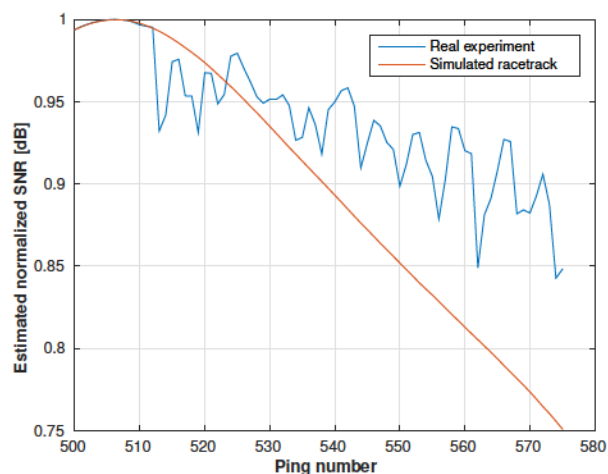


Fig. 7. Estimated normalized SNR for the real experiment and for the simulated racetrack.

likely tracks originated by the target for further prosecution. In this case, MML launches a receding horizon, non-myopic algorithm which controls the AUV's heading to improve the tracking performance.

The experiments at sea during LCAS15 demonstrated that the proposed autonomy architecture can be executed on-board the AUVs with real-time CAS processing. CAS, which has gained recently attention for ASW in littoral scenarios, offers the promise of multiple detections per waveform cycle. This can potentially improve the quality/length of tracks, thus increasing the adaptive behaviour's performance in terms of achievable detection and tracking performance. The non-myopic, adaptive behaviour, in turn, can provide clear benefits from the detection and tracking point of view, such as increasing the SNR, reducing the localisation error and maintaining the tracks for longer. Further investigation is required to evaluate the performance of CAS in littoral scenarios. Our future work will focus on extending the MML concepts to a multi-agent system.

ACKNOWLEDGEMENTS

This work has been supported by the NATO Allied Command Transformation under the Collaborative Anti-Submarine Warfare research programme. The authors thank Vittorio Grandi (Engineering Coordinator) and all the CMRE Engineering and Technological Department staff for their help, willingness and collaboration during the cruise. The authors gratefully acknowledge and thank also the NRV Alliance and CRV Leonardo crews and everyone involved in the LCAS15 trial.

REFERENCES

- [1] R. Been, D. T. Hughes, and A. Vermeij, "Heterogeneous underwater networks for ASW: technology and techniques," Nato Undersea Research Centre, Tech. Rep., 2008.
- [2] D. T. Hughes, F. Baralli, S. Kemna, K. LePage, M. Hamilton, and A. Vermeij, "Collaborative multistatic ASW using AUVs: demonstrating necessary technologies," Nato Undersea Research Centre, Tech. Rep., 2009.

- [3] 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 *Proceedings of CDC2014, Los Angeles, December 15-17, 2014*, 2014.
- [4] M. J. Hamilton, S. Kemna, and D. Hughes, "Antisubmarine warfare applications for autonomous underwater vehicles: the GLINT09 sea trial results," *J Field Robot (Special Issue: State of the Art in Maritime Autonomous Surface and Underwater Vehicles, Part 1)*, vol. 27(6), pp. 890–902, 2010.
- [5] S. Kemna, M. J. Hamilton, D. T. Hughes, and K. D. LePage, "Adaptive autonomous underwater vehicles for littoral surveillance the glint10 field trial results," *Intelligent service robotics*, vol. vol. 4, No. 4, pp. 245–258, 2011.
- [6] R. Goldhahn, K. LePage, P. Braca, A. Munafò, and G. Ferri, "Environmentally sensitive AUV behaviors for collaborative multistatic surveillance networks," in *2013 Underwater Acoustic Signal Processing Workshop October 16-18, 2013 West Greenwich, RI, USA*, 2013.
- [7] P. Braca, R. Goldhahn, G. Ferri, and K. LePage, "Distributed information fusion in multistatic sensor networks for underwater surveillance," *IEEE Sensors*, 2015.
- [8] J. Sildam, K. LePage, P. Braca, and M. Micheli, "On unsupervised track classification based on entropy distribution estimated along track related detections," in *1st International Conference and Exhibition on Underwater Acoustics*, 2013.
- [9] G. Ferri, A. Munafò, R. Goldhahn, and K. LePage, "Towards fully autonomous underwater vehicles in ASW scenarios: an adaptive, data driven AUV Mission Management Layer," in *OCEANS*, 2015.
- [10] P. Hines, "Experimental comparison of continuous active and pulsed active sonars in littoral waters," in *1st International Conference and Exhibition on Underwater Acoustics, Corfu, Greece*, 2013.
- [11] D. Grimmett, "Target AOU growth containment using LFM high duty cycle sonar," in *Underwater Acoustics 2014, Rhodes, Greece, 22-27 June*, 2014.
- [12] A. Munafò, G. Canepa, and K. LePage, "Continuous active sonars for littoral undersea surveillance," *Submitted to IEEE Journal of Oceanic Engineering*, 2016.
- [13] G. Ferri, A. Munafò, R. Goldhahn, and K. LePage, "Results from COLLAB13 sea trial on tracking underwater targets with AUVs in bistatic sonar scenario," in *Oceans - St. John's, 2014*, Sept 2014, pp. 1–9.
- [14] G. Ferri, "Towards fully autonomous underwater vehicles in ASW scenarios: an adaptive AUV mission management layer," CMRE-FR-2015-005, Tech. Rep., 2015.
- [15] <http://oceanai.mit.edu/moos-ivp/pmwiki/pmwiki.php>.
- [16] K. LePage, "An SNR maximization behavior for autonomous AUV control," in *Proc. 10th European Conf. Underwater Acoustics, Istanbul, Turkey*, 2010.
- [17] K. Rajan, "The role of deliberation (and AI) in marine robots," in *Workshop on Marine Robotics: Looking into the Crystal Ball, Eurocast 2013, Gran Canaria, Spain*, February 11-13, 2013.
- [18] G. Canepa, A. Munafò, M. Micheli, L. Morlando, and S. Murphy, "Real-time continuous active sonar processing," in *OCEANS 2015*, 2015.
- [19] A. Maguer, R. Dymond, M. Mazzi, S. Biagini, and S. Fioravanti, "SLITA: a new slim towed array for AUV applications," in *Acoustics08, Paris, France, pp 141–146*, 2008.
- [20] <http://liquidr.com/>, Std.

Document Data Sheet

<i>Security Classification</i>		<i>Project No.</i>
<i>Document Serial No.</i> CMRE-PR-2017-006	<i>Date of Issue</i> November 2017	<i>Total Pages</i> 7 pp.
<i>Author(s)</i> Ferri, G., Munafò, A., Alves, J., LePage, K.		
<i>Title</i> A data-driven control strategy in synergy with continuous active sonar for littoral underwater surveillance.		
<i>Abstract</i> <p>In this work, we describe a data-driven Mission Management Layer (MML) running on-board AUVs which manages the phases of a littoral surveillance mission and exploits the characteristics of Continuous Active Sonar (CAS) signal processing. The MML selects for further investigation the tracks which are likely originated by a target. In this case, the MML launches a receding horizon, non-myopic control algorithm which controls the AUV's heading to improve the tracking performance to ease the target classification. The algorithm minimises the expected target position estimation error over a prediction time window by achieving a trade-off amongst different objectives: keeping the target at broadside, reducing the distance to the target, avoiding areas of high reverberation and searching for geometric configurations with low bistatic target localisation error. We report at-sea experiments obtained during the LCAS15 sea trial, which demonstrated, for the first time, that the proposed autonomy architecture can be executed together with real-time Continuous Active Sonar (CAS) processing on-board the AUVs. CAS has recently gained interest for littoral Anti- Submarine Warfare, since it offers the promise of multiple detections per waveform cycle. This can potentially improve the quality/length of tracks, thus increasing the adaptive behaviour's performance, which, in turn, can increase the detection and tracking capabilities of the processing chain.</p>		
<i>Keywords</i>		
<i>Issuing Organization</i> 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