THREATDETECT: AN AUTONOMOUS PLATFORM FOR SECURING MARINE INFRASTRUCTURES *

Paolo Casari (paolo.casari@imdea.org)
IMDEA Networks Institute, Madrid, Spain

Lutz Lampe (lampe@ece.ubc.ca)
University of British Columbia, Vancouver, Canada

Stephane Martinez (smartinez@univ.haifa.ac.il)
Morris Kahn Marine Research Station, Department of Marine Biology, Leon H. Charney School of Marine Sciences, University of Haifa, Israel

Roee Diamant (roee.d@univ.haifa.ac.il)
University of Haifa, Israel

Abstract. The NATO SPS multi-year project ThreatDetect investigates an autonomous platform for securing marine infrastructures by reliably detecting divers and mines in real time. Our system combines acoustic remote detection with verification using pattern recognition on underwater imagery. For diver detection, we rely on active acoustics from a single transceiver, and analyze the acoustic reflections to detect and localize a target that fits the pattern of a diver. For mine detection, we segment sonar images from an autonomous underwater vehicle (AUV) to differentiate between background, highlight, and shadow. In case of detection, we steer the AUV's trajectory so as to closely observe the target, and transmit segmented sonar images to a surface station via underwater acoustic communications. At the time of writing, the project is performing final technology tuning and integrated sea experiments.

Key words: Underwater threat detection, infrastructure security, diver detection and tracking, AUV, submerged mine detection, underwater acoustic communications, sea experiments

1. Outline of the project

The global marine industry represents a multi-billion-dollar per year business, whose sensitive infrastructure may include oil and gas facilities, offshore rigs, submerged pipes and cables, as well as harbor installations. Such strategic infrastructure must be protected against hostile intruders: once an unwanted access event is detected, surface control units must be promptly and securely

* This research has been sponsored in part by the NATO Science for Peace and Security Programme under grant G5293.

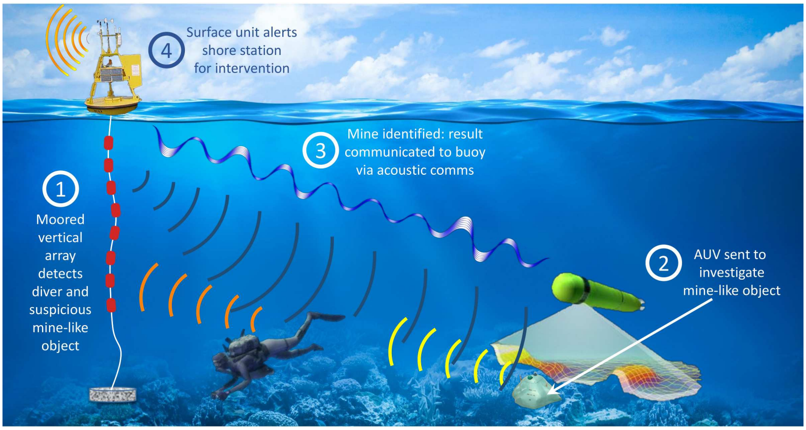

Figure 1. Concept of our solution for the identification of intruding divers and mines.

alerted. Currently, these tasks are performed manually, requiring considerable manpower investment (for example, about 100 people for a single gas rig in Israel), restricting real-time response capabilities. In particular, no holistic solution is currently available to secure facilities against direct diver attacks or against the deployment of submerged mines. While there exist active acoustic diver detection and autonomous underwater vehicle (AUV)-based mine detection systems, they often suffer from significant false alarm rates and require human interaction. The main reason is that most algorithms to identify acoustic reflections from a target remain suboptimal under high reverberation, without a full characterization of the target's reflection pattern.

In this project, we design a novel autonomous early detection and identification system for divers and mines that combines active acoustic remote detection with verification through an approaching AUV (see the system's structure in Fig. 1). Verified detections are communicated acoustically to a control station. We rely on a single acoustic transceiver, which greatly improves the applicability of our solution to small platforms, boats and buoys.[1]

The acoustic detection component processes reflected acoustic signals over time via a track-before-detect approach, and extract patterns through dynamic programming. This method distinguishes a moving target (e.g., a scuba diver) even when its acoustic footprint is deeply buried in noise or outweighed by strong static reflections. To localize the target using a single transceiver, we match the range and bearing of the target with a model-based reproduction of acoustic propagation under water.

The detection of submerged mines from sonar images is based on a combination of detection and segmentation algorithms that offer a good tradeoff

[1] Visual project overview: `https://tinyurl.com/NATO-SPS-ThreatDetect-video`.

between true positive and false negative detection rates. All the above systems are integrated into an automatic solution, where an AUV scans the area around some marine infrastructure through a sonar, and automatically adapts its course to approach detected targets and scan them more accurately. Relevant sonar segments including detections are compressed and transmitted acoustically from the AUV to a surface station. Such communications are robustified against anthropogenic and natural interference via interference suppression algorithms that leverage the characterization of shipping noise and other simultaneous acoustic transmissions.

At the current stage of the project, we have developed all required subsystems and tested them in multiple sea experiments, both in the Mediterranean sea and in the Red sea. These experiments include detections of real scuba divers from boats and buoys, and detection of submerged mines from our AUV and other surface vehicles. The AUV detection, course adaptation and image communication chain has also been fully integrated. In the following, we provide a high level description of the system's components.

2. Detection and localization

2.1. ACOUSTIC DETECTION, CLASSIFICATION AND TRACKING

Our method to detect scuba divers addresses the limitations of currently existing solutions. These limitations include the need to setup a rigid array of multiple receiving elements, the false alarms caused by mistaking marine fauna as scuba divers, and the use of high-power acoustic transmissions that may harm marine animals. Instead, we rely on a single transceiver, small enough to be deployed from any platform, and transmit at lower source levels by leveraging wideband acoustic signals. A full description of the system is available in (Diamant et al., 2019).

Since our transmission level is low, we accumulate reflections from multiple acoustic transmissions and arrange them into a time-distance matrix, where each row contains the reverberation pattern corresponding to one transmission (see, e.g., the data in Fig. 2a, collected near the Haifa shore, Israel). Detections are then achieved based on pattern recognition from all reflections (track-before-detect approach). Specifically, we assume that clutter is random, whereas reflections from a real target are stationary. We single out the latter from the former through probabilistic analysis, based on a modified Viterbi algorithm informed with the maximum speed and orientation change rate that are compatible with diver motion.

To reject static reflections from such objects as rocks, weights, and chains, we embed a clustering step based on expectation-maximization (EM). This

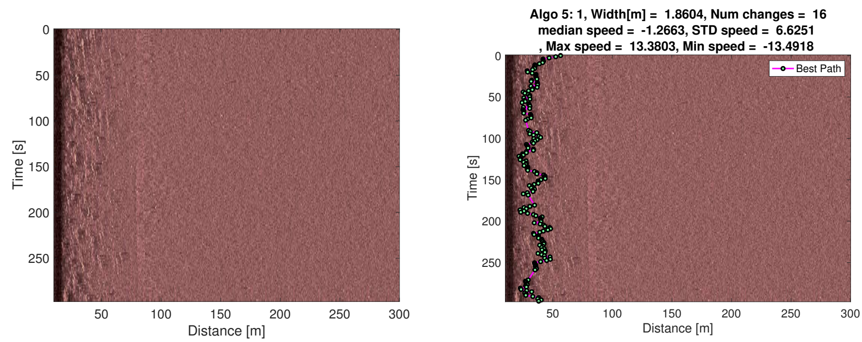

a. Time-distance matrix b. Detected track

Figure 2. A detected track (b) over a time-distance matrix with reflections from a diver (a).

scheme separates static reflections (whose energy arrives from the same location over time) from non-static ones, by operating column-wise on the time-distance matrix. The application of this method to Fig. 2a results in detecting the diver trajectory highlighted in Fig. 2b.

2.2. LOCALIZATION METHODS

We developed two different localization methods: one aimed at low-complexity systems with a single receiver, and a second one targeting systems that embed a vertical array. In the former case, we rely on the diversity of the bathymetry around a fixed receiver. This induces a different multipath propagation profile depending on where the source of an acoustic signal: therefore, knowing the bathymetry around the receiver and the sound speed profile allows us to infer the location of the source from the multipath distortion.

We proceed as follows. For an arbitrarily dense grid of possible source locations, we predict the channel impulse response (CIR) that acoustic signals would be subject to using a numerical acoustic propagation model. Every time we receive an acoustic signal we estimate the CIR, and correlate it with our CIR database to extract a set of candidate source locations. We then collect multiple transmissions and fuse the corresponding source location estimates using a trellis search algorithm. This yields the most likely source path with a lower complexity than the Viterbi algorithm. Fig. 3a shows an accurate trajectory estimate for a moving source using environmental data from the San Diego bay area, USA. The algorithm correctly selected a sequence of locations from a grid composed of about 4 million points.

For localization via a vertical array, we employ wideband acoustic beamforming helped by side information. We observe that underwater equipment often imposes mounting constraints that may impede to fully control the

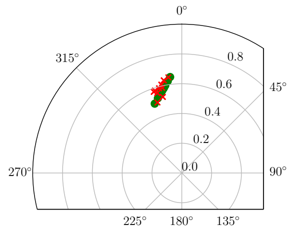

a. Path detection

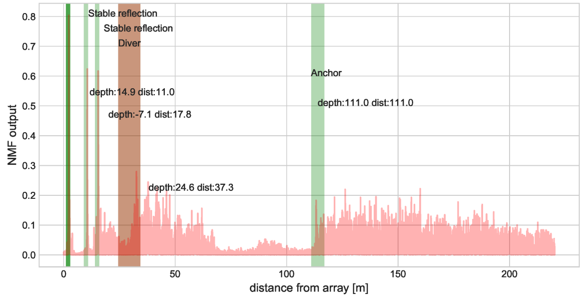

b. Angle-of-arrival estimation

Figure 3. Path detection using a single receiver (a); angle of arrival estimation for different targets and a diver using a vertical array (b).

shape of an acoustic array. For systems working at centimeter wavelengths λ, it may be even impossible to guarantee that the array preserve the $\lambda/2$ spacing. We solve these challenges through a processing chain that involves matched filtering, clustering of match filtering peaks, elimination of stationary arrivals, and wideband direction of arrival estimation. The latter is helped by time-difference-of-arrival information in order to bound the angle of arrival search and remove spatial ambiguity. This system was tested in a sea experiment involving two divers. The results in Fig. 3b show the output signal recorded by one of the array's hydrophones, where relevant peaks are tagged with range and depth estimates. We observe that our scheme correctly identifies the diver (brown) at about 25 m of depth and 37 m from the array.

2.3. SEGMENTATION OF SONAR IMAGES FROM AUVS

We employ a multi-stage chain for the detection of submerged mines from sonar images. We start with a rough detection of regions of interest (ROIs) within the sonar image, and then segment the ROI to separate background, highlight, and shadow regions. We report a detection based on prior knowledge about the target (e.g., minimum number of pixels, height above the bottom, and so forth). Specifically, we use likelihood ratio combining of highlight and shadow identifications. The former is based on segmenting a higher moment of the sonar image, whereas the latter is based on a blind classification using a support vector machine (SVM).

The segmentation of the highlights relies on the assumption that acoustic reflections from the target have homogeneous characteristics, whereas for shadows we assume that the distribution of the shadow pixels is the same throughout the image, hence the shadow's intensity level in the ROI can be learned from other sonar image sections. Fig. 4a shows an example of ROI. A

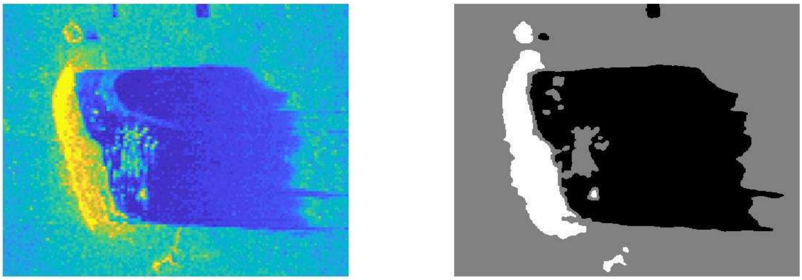

a. Detected ROI b. Segmented image

Figure 4. A region of interest including a target (a), and segmentation results (b).

full description of the segmentation process is available in (Abu and Diamant, 2019b).

We perform fine-grained target detection by clustering the ROI via two strategies. The first is based on clustering the ROI using an EM algorithm for a mixture of alpha distributions, which offers sufficient flexibility. For better performance, we allowed hard decisions to be made during EM iterations: this way, we can exploit the expected dependencies among the image pixels related to the target. For full details, we refer the reader to (Abu and Diamant, 2018). The above solution is accurate but computationally demanding. Therefore, we designed a second segmentation strategy based on fuzzy logic. Here, we manage the inhomogeneity within the sonar image through two fuzzy terms that reflect the location of the segmented pixels within the image (Abu and Diamant, 2019a). For both solutions, we designed a de-noising filter that smooths the image before segmentation. An example for a segmented image using our approach is shown in Fig. 4b.

We tested the above solutions via a specifically designed simulator that generates sonar images with targets over different types of seabed (e.g., sand ripples, grass, or rocks). Furthermore, using our own AUV, we collected a database of more than 1000 sonar images, and hand-labeled them for targets. Since these images were taken from different sea environments (in France, as well as in southern, central, and northern Israel), the database enables a sufficiently robust statistical exploration of ROI detection performance.

2.4. AUTOMATIC SCAN BY AN AUV

A distinctive feature of our system is a fully automatic behavior without humans in the loop. In this vein, we have designed a protocol that allows the AUV to change its predefined course in real time upon detection events. Such capability is essential to detect submerged mines, which often can be

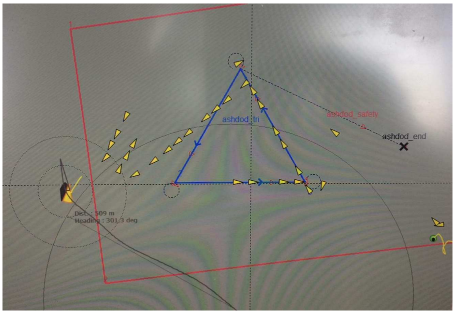

a. Image of the track record

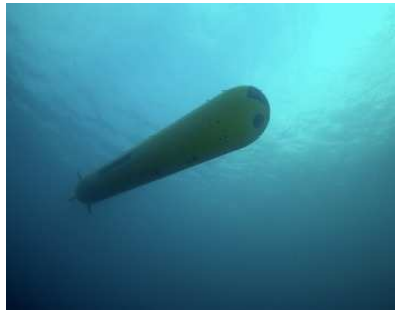

b. The AUV performing the mission

Figure 5. AUV (a) track record from a sea experiment (b). Yellow arrows show the actual track of the vehicle. Blue lines show the pre-determined mission track.

identified as such only after viewing them from different points of view. An automatic procedure is paramount in this case, as navigation errors may prevent an operator to successfully steer the AUV towards the suspected target, and would imply significant delays for surfacing and re-diving the AUV.

We implement the automatic adaptation of the AUV mission by acting on its backseat driver. For each sonar image collected, the AUV detects and segments ROIs in real time. It then determines the geographic location of the ROI, and correspondingly sets a new mission, which includes back-tracking to observe the object in the ROI from the opposite direction, and returning to the previous location to resume the original mission. During this last step, the AUV also hovers above the ROI, so as to profile the target via optical or sub-bottom scans. Finally, the AUV fuses all collected information, compresses the best ROI image, and sends it to an operator.

Fig. 5a shows an example of the above procedure operated by the AUV in Fig. 5b, in southern Israel in June 2019. Blue lines represent the AUV's pre-determined mission, whereas the yellow arrows shows the actual position of the AUV over time. Once it reaches the bottom-right corner of the trajectory, the AUV drops its pre-defined mission and loops around a suspect location, in real time without any intervention by an operator.

3. Noise characterization and interference cancellation

As we foresee our system to operate in the proximity of underwater installations or harbors, we investigate methods to make communications robust against noise generated by shipping activities as well as co-existing signals, which we expect to constitute major impediments in our scenarios of interest. To model shipping noise, we evaluated acoustic recordings from the Ocean Networks Canada (ONC) database available at (ONC, 2019), which includes

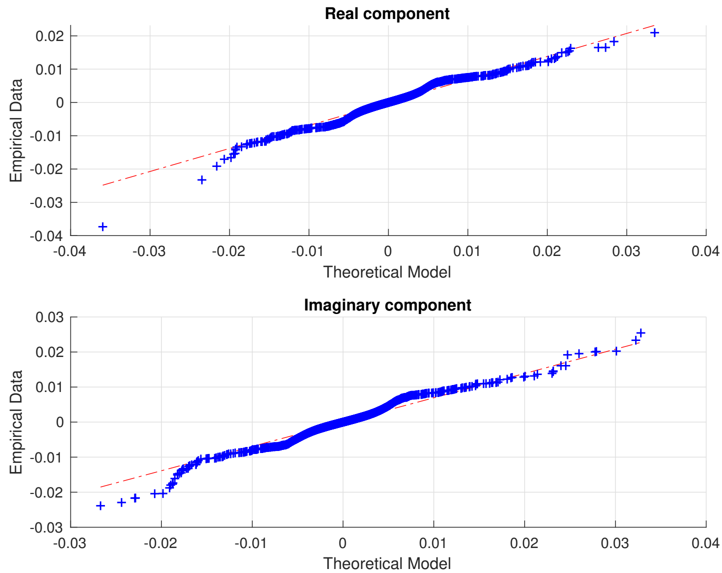

Figure 6. QQ-plot of empirical quantiles from a complex baseband noise signal in the 8–16 kHz band versus theoretical quantiles from a Gaussian mixture random variable.

a large number of measurements from hydrophone arrays, some of which are located close to busy shipping routes. As a starting point for our modeling attempt, we considered the impulsiveness of shipping noise as also reported in previous work such as (Chen et al., 2017; Mahmood and Chitre, 2016), and processed data to fit a Gaussian mixture model via the EM algorithm. Fig. 6 compares measured noise and noise according to a two-term Gaussian mixture model trained with the EM algorithm. We observe that shipping noise could be modelled with a two-term mixture having a 15 dB higher variance for the impulse noise component.

3.1. INTERFERENCE CANCELLATION

3.1.1. *Shipping Noise*

The structure inherent to shipping noise can be exploited for noise mitigation or cancellation. We consider the latter, making use of null-subcarriers in a multicarrier underwater acoustic communication signal, which enables noise (or interference) estimation. For this purpose, the fact that the impulse component of the noise is sparse can be exploited via compressed-sensing based signal recovery (Caire et al., 2008). This can be extended by also considering the noise distribution or by learning a sparse representation (Aharon et al., 2006). Using 192 ship noise recordings obtained from the ONC database and considering 25% of null-subcarriers, we conclude that communication

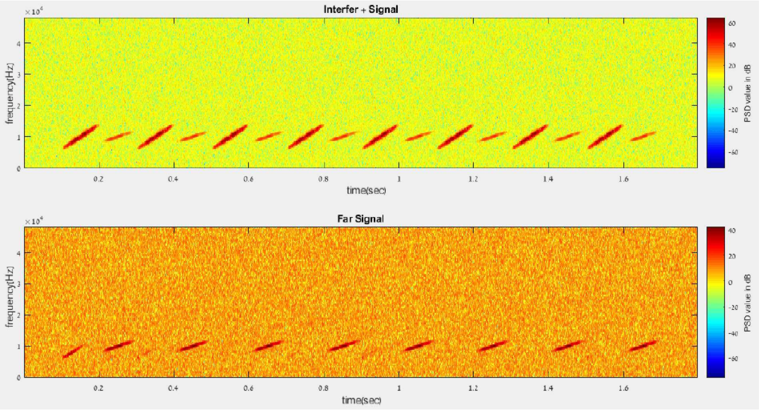

Figure 7. A spectrogram showing two signals before (upper panel) and after (lower panel) interference cancellation.

systems operating from 0.1 to 16 kHz can achieve relative cancellation gains of about 1 to 3 dB (Atanackovic et al., 2019).

3.1.2. *Co-signal Interference*

Co-existing acoustic transmissions constitute one of the strongest interference sources for underwater acoustic communications. Such signals can originate from acoustic systems (e.g., echo sounders for depth measurement, ADCP water current meters) or from transmissions by nearby modems. When these transmissions intersect in the same bandwidth of the desired signals, the signal-to-interference ratio (SIR) may be too low to correctly receive the message. Considering this challenge, we have designed an interfering cancellation filter that specifically handles the case of low SIR.

Our algorithm modifies the traditional noise cancellation filter in order to remove a reference signal from the signal to be cleaned. The solution first detects the strong interference and finds its analytic form (i.e., its time-frequency pattern). Then, we employ an adaptive filter to equalize the channel from the interference source to the receiver. The result is subtracted from the received signal, and is fed back into the the adaptive filter. To discriminate between the two channels (i.e., from the interference and from the designed source), we lock onto the channel taps of the interferer and zero-force all other taps. This way, the energy of the desired signal is not effected. The full details of this approach are available in (Diamant, 2018).

We have tested our system in several sea experiments including two interferers and a common receiver. An example of the performance of our algorithm is provided in Fig. 7, where we show the spectrogram of the received signal before and after interference cancellation for a sea experiment conducted in Ashdod, Israel. We observe that the strong interference is almost completely removed from the signal.

4. Conclusions

The ThreatDetect project targets the protection of sensitive and strategic marine infrastructures from intruding divers and submerged mines. In this paper, we reported on the progress of our research activities in the context of diver detection, threat localization, fully automatic detection of mines by an AUV, as well as noise and interference cancellation for robust communications. A number of integrated sea trials demonstrate the feasibility of our solution.

References

Abu, A. and Diamant, R. (2018) Unsupervised Local Spatial Mixture Segmentation of Underwater Objects in Sonar Images, *IEEE Journal of Oceanic Engineering*.

Abu, A. and Diamant, R. (2019)a Enhanced Fuzzy-based Local Information Algorithm for Sonar Image Segmentation, *IEEE Transactions on Image Processing*.

Abu, A. and Diamant, R. (2019)b A Statistically-based Method for the Detection of Underwater Objects in Sonar Imagery, *IEEE Sensors Journal*.

Aharon, M., Elad, M., Bruckstein, A., et al. (2006) K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation, *IEEE Transactions on signal processing* **54**, 4311.

Atanackovic, L., Zhang, R., Lampe, L., and Diamant, R. (2019) Statistical Shipping Noise Characterization and Mitigation for Underwater Acoustic Communications, In *Proc. MTS/IEEE OCEANS*.

Caire, G., Al-Naffouri, T. Y., and Narayanan, A. K. (2008) Impulse noise cancellation in OFDM: An application of compressed sensing, In *IEEE International Symposium on Information Theory*, pp. 1293–1297.

Chen, P., Rong, Y., Nordholm, S., He, Z., and Duncan, A. J. (2017) Joint channel estimation and impulsive noise mitigation in underwater acoustic OFDM communication systems, *IEEE Transactions on Wireless Communications* **16**, 6165–6178.

Diamant, R. (2018) Robust Interference Cancellation of Chirp and CW Signals for Underwater Acoustics Applications, *IEEE Access* **6**, 4405–4415.

Diamant, R., Kipnis, D., Bigal, E., Scheinin, A., Tchernov, D., and Pinchasi, A. (2019) An Active Acoustic Track-Before-Detect Approach for Finding Underwater Mobile Targets, *IEEE Journal of Selected Topics in Signal Processing* **13**, 104–119.

Mahmood, A. and Chitre, M. (2016) Robust communication in bursty impulsive noise and Rayleigh block fading, In *11th ACM International Conference on Underwater Networks & Systems*, p. 13.

ONC (2019) Search Hydrophone Data, https://data.oceannetworks.ca/SearchHydrophoneData.