MANgO: federated world Model using an underwater Acoustic NetwOrk

Gordon Frost*, David M. Lane, Nikolaos Tsiogkas*
Daniele Spaccini†‡, Chiara Petrioli†‡
Maarja Kruusmaa§, Victoria Preston§, Taavi Salumäe§
* Ocean Systems Laboratory, School of Engineering & Physical Sciences, Heriot Watt University
{gwf2,D.M.Lane,nt95}@hw.ac.uk
† Computer Science Department, University of Rome "La Sapienza", Rome, Italy
{petrioli, spaccini}@di.uniroma1.it
‡ WSENSE s.r.l., Rome, Italy
§ Centre for Biorobotics, Tallinn University of Technology
{taavi.salumae, maarja.kruusmaa}@ttu.ee

***Abstract*—The usage of Autonomous Underwater Vehicles (AUVs) in scientific and industrial applications has increased in the last decade. In many applications the usage of fleet of heterogeneous assets is a key requirement. Their coordination and cooperation require high performance in terms of communication and information sharing. *MANgO* presents a federated world model operating over an underwater acoustic network through the networking capabilities provided by the SUNSET Software Defined Communication Stack (SDCS). It provides a method to transparently synchronise the local *World Model* (WM) of a vehicle with the rest of the heterogeneous team. The performance of the proposed architecture is demonstrated and compared with previous approaches using simulation and hardware-in-the-loop tests. Results show an improvement in the performance of at least 30% as well as the ease of integration with a variety of heterogeneous platforms.**

Keywords*—*MANgO, Distributed World Model Service, SUNSET, Acoustic communications.

I. Introduction

The use of robots to perform complex tasks underwater is of paramount importance. This is specially due to the harshness of the underwater domain that makes difficult or even unfeasible the execution of several tasks by humans, such as pipeline monitoring, discovery and protection of marine archaeology, port security, etc. Usually underwater robots are tethered to a base station, e.g., a ship, allowing the remote operators to have full control but at same time, limiting their movements. Using autonomous assets allows to overcome the limits imposed by the use of cables and therefore to enable novel application scenarios. Multiple Autonomous Underwater Vehicles (AUVs) can be used in several cooperative mission scenarios to perform tasks in parallel and decrease the overall mission time. Collaborative robots can be used in several application scenarios, such as mine counter-measures operations, surveying, deep-water oilfield inspection-and-repair and search-and-inspect tasks [1] [2] [3].

In order to coordinate efficiently, vehicles usually have to communicate and exchange important information regarding the execution of the given mission. Communication is even harder in the underwater environment since it is usually performed via acoustic waves that suffer from high latencies, varying channel quality and low bandwidth [4], [5], [6]. The authors in [7] face such problem by proposing a novel method to reliably exchange information among the members of an AUV team.

The *MANgO* project, presented in this paper, has been part of the FP7 SUNRISE project [8], [9]. It proposes a refined application of an open device-independent Distributed World Modelling Service (DWMS) [7] for multiple static and mobile autonomous underwater assets. In the scope of the *MANgO* project, the acoustic communication, among the members of the team, is performed using the SUNSET Software Defined Communication Stack (S-SDCS), which is an enhanced version of SUNSET developed in the scope of the SUNRISE project. By combining the DWMS and the S-SDCS an asynchronous device-agnostic method for synchronising a mission's relevant information among the team's assets, is provided. A local World Model (WM) in the form of a database is maintained with up-to-date information of it's environment, task, and other asset's intentions. Each vehicle's local WM is kept up-to-date globally by a Data Exchange Manager (DEM) which selects key information to transmit acoustically, based on their priority. The synchronised WMs allow each vehicle's planner to optimize selected actions and to minimise mission execution time, improving the efficiency. Additionally, the operator situational awareness is increased as he can obtain information regarding the mission execution status. Moreover, the operator can intervene and provide actions that the collective should execute using the same mechanisms for synchronising his knowledge and intentions.

The rest of the paper is organised as follows. In section II the DWMS architecture is presented. Section III presents the SDCS that the DWMS relies upon. In section IV the experimental setup that validates the architecture developed in the scope of the *MANgO* project is shown. Section V presents the results and discusses on them. Finally, in section VI useful conclusions are drawn and discussion on future work is performed.

II. *MANgO* DWMS Architecture

As mentioned in Section I, attempts for an underwater network that synchronises the knowledge among a team of vehicles were made and presented in [7]. Results from in-field experiments of such architecture are presented in [10]. In the proposed approach a custom communication scheme was implemented based on Time division Multiple Access (TDMA) medium access control. Based on this work in [11] were presented methods for efficient collaboration among a team of AUVs. Such implementation made easy the integration of various task scheduling schemes.

In the scope of the *MANgO* project, a new architecture is instead proposed. Experience with the previous architecture suggested that TDMA increased the mission execution time. This was even more evident in cases when errors in communication were high and large amounts of messages had to be retransmitted. Additionally, it was observed that many of the timeslots were empty as there was no new information to be transmitted. This fact made the TDMA inapplicable in larger teams of robots, as most of the time the communication channel was not utilised reducing the effective bandwidth. Instead, the MANgO DMWS was re-designed to utilise the S-SDCS asynchronous communication functionality. The DEM module of the DWMS, shown in Figure 1, has hence transitioned from a TDMA to an asynchronous push-based method.

The push-based method is implemented by using a priority queue to store data to be transmitted in the DEM. This priority queue is filled by the ontology database as data are inserted or updated from the other software modules of the vehicle. The DEM periodically checks if any data has to be transmitted. In such a case, it selects and transmits the one with the highest priority. Similarly, whenever the DEM receives data from another vehicle it inserts or updates such information in the ontology database, which is in turn pushing the information to all the interested software modules of the vehicle. After the successful reception of data, the DEM is responsible for generating appropriate acknowledgement messages.

An important component of the *MANgO* architecture is the ontology that is used to store the information of the environment. In the architecture described in [7], the ontology was implemented using Jena [12]. In the scope of *MANgO* it was decided to replace it with a simple database as it would allow for easier and faster integration with the rest of the system. The database of choice was *RethinkDB* [13], a *NoSQL* database that is widely used in production environments. One of the most appealing features of RethinkDB, is the *changefeed*. This feature allows software modules to subscribe to specific pieces of information that they are interested in. The database then takes care of pushing any changes to the subscribed software modules. The DEM uses this feature to learn about the changes in the world that are observed by the vehicle and are inserted in its knowledge base. In the same manner, the database automatically pushes any changes received by the DEM to the mission executor. This feature allowed for easy integration and development of the whole architecture saving the time needed to develop a polling mechanism with the related complexity to handle synchronisation.

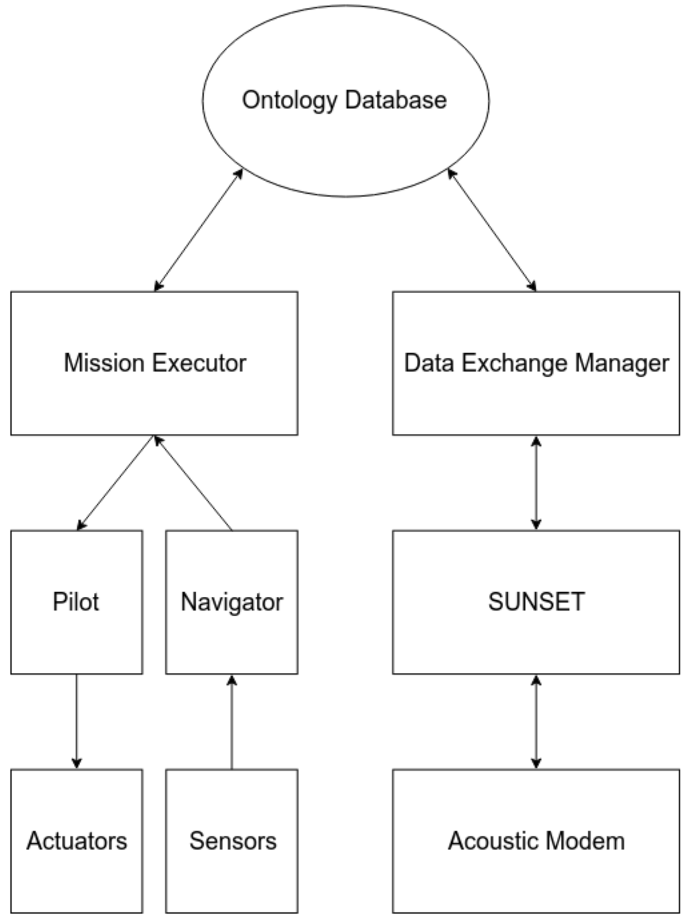

Figure 1: Architecture Diagram of an AUV modules including the *MANgO* Distributed World Model modules.

Another useful feature of RethinkDB is the integrated web interface it comes with. It gives the user the ability to view and query the database in real-time. This can be used to increase the user awareness regarding the mission execution. For instance, as the DEM synchronises all vehicles' DBs throughout the mission, a user on the shore can view the execution progress, such as detected targets, by querying the local DB through the web interface. Similarly, the user can input information into the DB using this web interface which will then be transparently shared between all vehicles by the DWMS using the S-SDCS. Additionally, using the bindings provided a user can create easily an interface that would allow for better status visualisation and control of the robotic team.

III. SUNSET SDCS

The MANgO DWMS relies on the networking capabilities provided by the S-SDCS [14]. The S-SDCS is an improved version of the SUNSET framework developed by the University of Rome "La Sapienza". The Sapienza University Networking framework for underwater Simulation Emulation and real-life Testing framework [14] (SUNSET) is a framework that provides networking and communication capabilities to underwater nodes. Its design allows to easily implement novel protocols and algorithms and integrate external hardware, such as sensors, modems and mobile platforms. One of the key feature of SUNSET is that the same code can be used

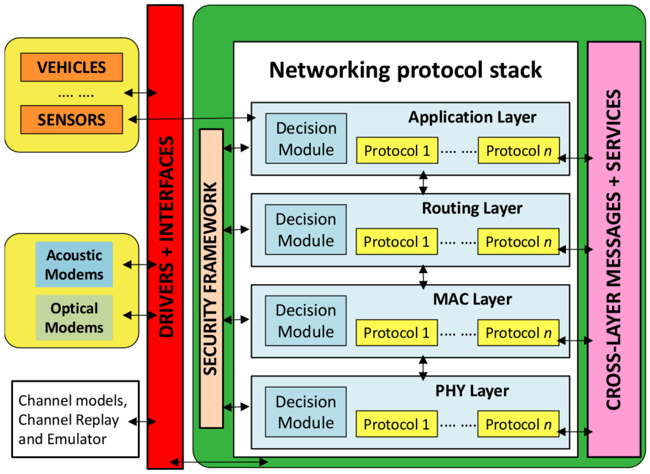

Figure 2: SUNSET SDCS architecture.

in simulation, in lab emulation using real hardware and in field without any code rewriting. SUNSET has recently been extended to support the innovative concept of Software Defined Communication Stack (SDCS) [15] allowing the use of different network protocols and modems on demand, according to the application scenario or the environmental conditions thus improving the overall performance of the system. The selection of the best protocol or modem to be used is dynamically and adaptively selected according to the network conditions, and applications requirements. A simplified architecture of the S-SDCS is shown in Figure 2.

The MANgO DWMS relies on the networking capabilities provided by the SUNSET SDCS [14]. A new module has been developed for the S-SDCS to provide the use of all it's capabilities to external software or third parties through simple APIs, named S-SDCS APIs. For instance, it can be selected which protocol and modem among those available in the stack will be used to transmit and receive data and tune their parameters on demand in real-time and remotely, if needed. The transmission/reception of data can be performed directly using the acoustic modem or through the protocols available in the network stack according to the user's needs. This module therefore allows to third parties that want to use the SUNSET SDCS capabilities without the need to develop their solutions inside the framework itself. In this way, novel networking solutions can be designed and implemented outside the SUNSET architecture but at the same time they can be easily tested using all the features of such framework. New modules for the ROS ecosystem have also been developed to allow the interaction between the DWMS and S-SDCS. Specifically, the S-SDCS ROS module acts as an interpreter between the S-SDCS APIs and the DWMS. In this way, all the information needed by the DWMS can be exchanged underwater through the S-SDCS protocol stack.

All these new modules have been designed and developed to support also multiple and different modems at the same time in the protocol stack thus allowing users to design and test in field multi-modal protocols.

IV. Experimental Setup

To evaluate the performance of the novel DWMS architecture proposed in the scope of the *MANgO* project, a set of simulated and hardware in-the-loop experiments were performed. The following sections describe the setup used to perform the tests.

A. Simulation setup

The simulated experiment consisted of a multi-vehicle mine-counter-measure mission, where the DWMS globally synchronises each vehicle's local world model. In this type of missions the robots discover targets in an area and then they classify them as being a mine or not. In the current work the robotic team is assembled with two types of robots: (a) robots capable of searching for targets (SAUVs) and (b) robots capable of inspecting and classifying the detected targets (IAUVs). The mission goal is to search a specific area for targets to be inspected and then inspect and classify those targets. The mission is completed when all of the discovered targets are classified and the classification information has been propagated to the whole team.

A decentralised task allocation scheme is implemented as described in [11] to complete the mission. In the decentralised task allocation there is one SAUV that provides potential mine information to the other AUVs of the robotic team. As the mission progresses and new targets are discovered, the SAUV communicates the targets to the other vehicles. Whenever a new target appears, each IAUV decides autonomously if it should inspect the target or not. This is done in a greedy manner where each vehicle calculates the distance of the target from all the inspection vehicles in the fleet and leaves the target for the closest one.

The reuse of the scenario and allocation scheme allows the direct comparison between the newly developed asynchronous DEM and the previous TDMA-based DEM. DEMs performance have been assessed through the investigation of the total mission time, defined as the time between the first target inserted into the SAUV WM to the final target being acknowledged to be inspected by an IAUV.

For the simulation part, the S-SDCS channel emulator was used to emulate the underwater channel behaviour by introducing propagation delays and per-link Bit Error Rates (BERs) in order to mimic the real underwater communication channel. In addition, the emulator takes in account also acoustic modem features, such as bitrate and additional delays. The vehicles taking part in the mission were simulated using the dynamics simulator designed for the NESSIE VII AUV [16], which was developed by HWU. To acquire statistically important results, both DWMS architectures were tested using the same sets of randomly generated targets. Specifically, ten sets of ten randomly generated targets were used, allowing to run ten different missions. In each mission the SAUV was discovering a target every 15 seconds. Additionally, in each mission we considered decreasing link qualities to affect nodes communication and therefore the system performance. Packet Error Rates (PER) started from 0% packet loss up to 60% packet loss. A visualisation of the simulation can be seen in Figure 3.

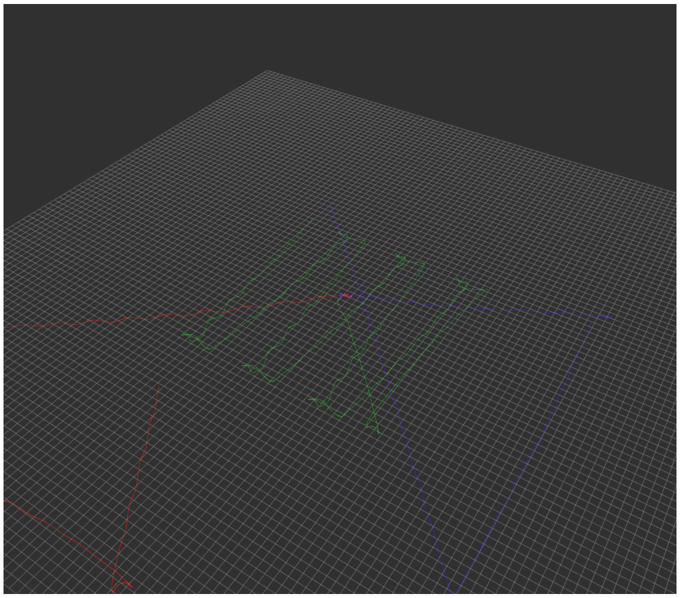

Figure 3: Simulated inspection mission consisting of one SAUV and two IAUVs. The SAUV detects synthetic inspection targets for the IAUVs. The IAUVs in turn minimise the time to visit all the targets using information from their local view of their world. This view is updated and reaches a global view of the world using the MANgO provided DEM.

B. Hardware-in-the-loop setup

In addition to the simulated tests, hardware-in-the-loop experiments have been performed to show the architecture's independence on the asset it is used on, the ease of integration on different platforms, and ultimately to validate the performance on realistic communication conditions.

In these tests, we used stand-alone communication nodes using low cost ARM-based embedded computers, such as the RaspberryPi and the Beaglebone Black, and underwater acoustic modems, developed by Evologics GMBH [17]. These nodes were used to simulate SAUVs and IAUVs performing the aforementioned simulated mission, while using the actual acoustic channel for the communication. One of the stand-alone nodes is shown in Figure 4.

To even further demonstrate the flexibility of the MANgO architecture, we also have integrated the DWMS and S-SDCS on U-CAT, a biomimetic vehicle. U-CAT is a low cost platform developed for autonomous inspection of confined areas [18], [19]. It is one of the outcomes of the FP7 Arrows project [1], where the goal was to develop robotic tools to reduce the cost of underwater archaeological campaigns. We demonstrate that the DWMS and S-SDCS can be quickly integrated onto a vehicle having relatively low computational resources. The U-CAT is used in simulated experiments being an inspection vehicle in the previous mine-counter-measure scenario, as well as a real vehicle in the hardware-in-the-loop tests. For these tests the vehicle was tethered to an Evologics modem. The U-CAT vehicle is shown in Figure 5.

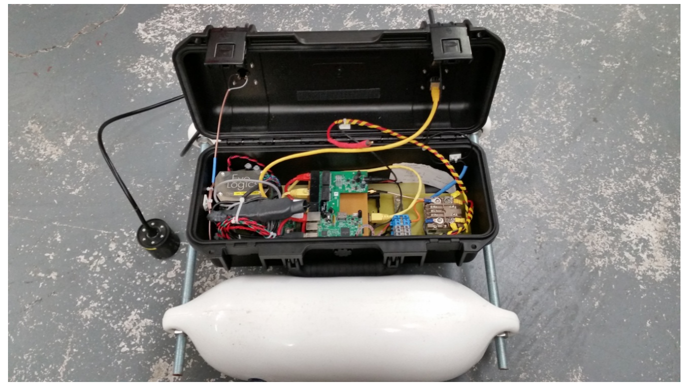

Figure 4: Acoustic Modem Buoy showing Evologics OEM S2CR 18/34 modem, computing and networking inside. Such buoys were used to test the architecture in real life communication conditions.

Figure 5: The U-CAT Autonomous Underwater Vehicle. Designed as a low cost platform to operate in dangerous conditions where vehicle loss is a possibility.

V. Results

This section presents the results obtained during simulations and hardware-in-the-loop tests. Simulations were performed using software provided by the SUNSET SDCS, while the hardware in the loop tests were performed in the testing facilities of the Ocean Systems laboratory and in real world conditions in Loch Earn, Scotland.

A. Simulation results

Figure 6 shows the simulation results for the comparison of the Ontology-based DWMS and the new *MANgO* DWMS. It is clear that the new DWMS outperforms the Ontology-based DWMS by a large margin. The increased efficiency of the *MANgO* DWMS is more evident when the packet error rate increases thanks to the lower execution time. In fact, the asynchronous communication of S-SDCS provides a huge benefit in mission execution allowing the underlying networking infrastructure to arrange the transmission scheduling.

When the error rates is higher, the new DWMS could have benefited of features firstly introduced in the old DWMS architecture. As an example, the use of cumulative acknowledgement messages would have allowed to reduce the number of message retransmissions and therefore the waste of bandwidth when the channel quality is poor. This would improve even more the already better performance of the new DWMS.

B. Harware-in-the-loop results

Multiple trials have been performed at the Loch Earn to evaluate the performance of the proposed approach. The achieved results are completely in line with the simulated results presented in the previous section as the packet loss experienced in Loch Earn was below 10%. In particular, we estimated the MANgO DWMS average mission execution time of about 17.25 minutes. In addition, the estimated average time between a target discovery and a target transmission was of about 8 seconds. Figure 7 shows a timeline of when targets are published by the search vehicle, received by an inspection vehicle and then subsequently visited.

The U-CAT performed the same mission that was simulated and described in Section IV. Modems were deployed at a distance of about 20 meters from each other. One modem represented an SAUV detecting targets, while the U-CAT was inspecting them. In this mission, the topology used was slightly different with respect to the simulated scenarios. In particular, we used only two nodes and the simulated inspection targets were in closer distances since the U-CAT was operating in tethered mode. In the simulated and hardware-in-the-loop cases the targets were in 100 by 100 meters box, while in the U-CAT case the targets were in a 20 by 20 meters box.

In Figure 8 a mission executed by the U-CAT is shown. The SAUV detected targets are depicted in red, while the U-CAT's odometry when navigating to these targets is represented by the blue line. It can be seen that the U-CAT does not navigate precisely to the waypoints and considers them reached from relatively large distance (approximately 2 meters). This happens due to localisation errors and to low navigation precision.

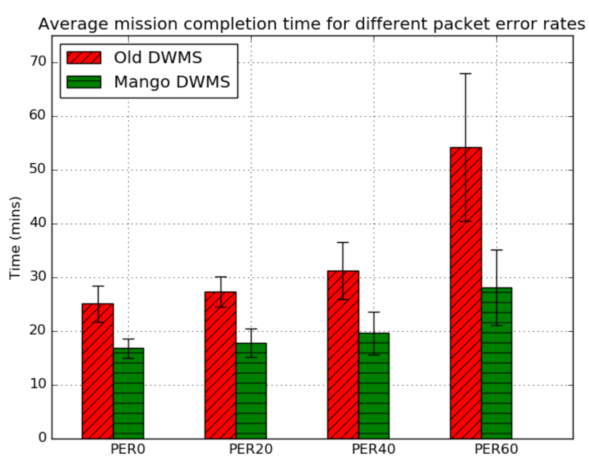

Figure 6: Analysis of the Ontology-based DWMS and the new MANgO DWMS with respect to Packet Error Rates (PER). The new DWMS performs consistently better. Specially in high error rates the performance is increased as less bandwidth is wasted due to the asynchronous way of communication.

Nevertheless, the U-CAT is able to receive all the targets from the SAUV and navigate to all of them.

The real execution can also be compared with the execution of the same mission by a simulated U-CAT shown in Figure 9. It can be seen that the paths followed by the vehicle are similar. In the real execution the path is different due to localisation errors, but the path is still really close to the optimal.

VI. Conclusion

The *MANgO* project has shown the initial version of a platform-agnostic, distributed world model service that successfully shares information between three AUVs in a simulated mine-counter-measures mission. We benchmark the *MANgO* DWMS, which uses the S-SDCS for communication, against previous TDMA-based strategy [11] demonstrating more efficient exchange of information in both simulation and hardware-in-the-loop experiments.

As future steps it is planned to integrate with a module that would assign priorities to data dynamically based on the value of information. The current implementation only supports static priorities assigned at the start of the mission. Additionally, evolving and integrating some of the previous work on acknowledgement policies, first presented in [7], would improve the performance in high error conditions. Finally, the integration with a state of the art ontology [20] and its planning and reasoning capabilities can improve mission performance and allow the utilisation of the underwater assets in more complex scenarios.

References

[1] B. Allotta, R. Costanzi, A. Ridolfi, C. Colombo, F. Bellavia, M. Fanfani, F. Pazzaglia, O. Salvetti, D. Moroni, M. A. Pascali *et al.*, "The arrows project: adapting and developing robotics technologies for underwater archaeology," *IFAC-PapersOnLine*, vol. 48, no. 2, pp. 194–199, 2015.

[2] C. C. Sotzing, J. Evans, and D. M. Lane, "A multi-agent architecture to increase coordination efficiency in multi-auv operations," in *OCEANS 2007 - Europe*, June 2007, pp. 1–6.

[3] V. Yordanova and H. Griffiths, "Synchronous rendezvous technique for multi-vehicle mine countermeasure operations," in *OCEANS 2015 - MTS/IEEE Washington*, Oct 2015, pp. 1–6.

[4] M. Roth, D. Vail, and M. Veloso, "A real-time world model for multi-robot teams with high-latency communication," in *Intelligent Robots and Systems, 2003. (IROS 2003). Proceedings. 2003 IEEE/RSJ International Conference on*, vol. 3, Oct 2003, pp. 2494–2499 vol.3.

[5] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke, "Data collection, storage, and retrieval with an underwater sensor network," in *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems*, ser. SenSys '05. New York, NY, USA: ACM, 2005, pp. 154–165. [Online]. Available: http://doi.acm.org/10.1145/1098918.1098936

[6] Z. Saigol, G. Frost, N. Tsiogkas, F. Maurelli, D. Lane, A. Bourque, and B. Nguyen, "Facilitating cooperative AUV missions: experimental results with an acoustic knowledge-sharing framework," in *Proceedings of IEEE-MTS Oceans'13, San Diego, USA*, 2013.

[7] F. Maurelli, Z. Saigol, J. Cartwright, D. Lane, A. Bourque, and B. Nguyen, "Tdma-based exchange policies for multi-robot communication of world information," *IFAC Proceedings Volumes*, vol. 45, no. 27, pp. 357 – 361, 2012. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1474667016312551

[8] C. Petrioli, R. Petroccia, D. Spaccini, A. Vitaletti, T. Arzilli, D. Lamanna, A. Galizia, and E. Renzi, "The SUNRISE GATE: Accessing the sunrise federation of facilities to test solutions for the internet of underwater things," in *Proceedings of UCOMMS 2014*, Sestri Levante, Italy, September, 3–5 2014.

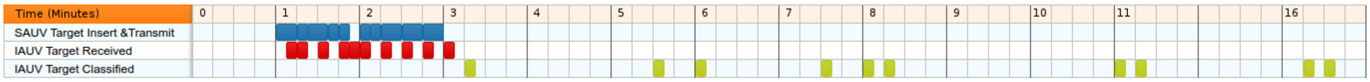

Figure 7: Timeline of when targets are received and inspected during a mission. Results were obtained using a harware-in-the-loop setup in real world communication conditions. The mission execution time is dominated by the inspection process.

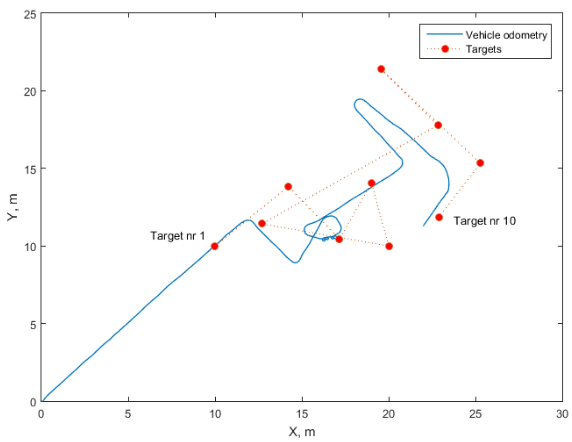

Figure 8: Odometry of the real U-CAT vehicle whilst it visited targets obtained through the DWMS in Loch Earn.

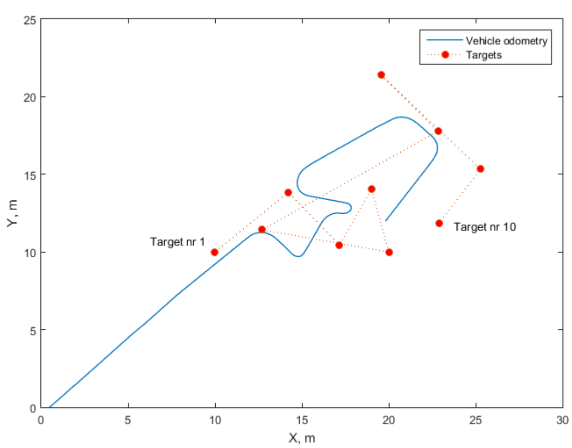

Figure 9: Odometry of a simulated U-CAT vehicle visiting targets received through the DWMS.

[9] "SUNRISE: Sensing, monitoring and actuating on the UNderwater world through a federated Research InfraStructure Extending the Future Internet," http://fp7-sunrise.eu/.

[10] F. Maurelli, Z. Saigol, G. Frost, N. Tsiogkas, and D. Lane, "Control architecture for multiple autonomous unmanned vehicle operations–trial results," 2013.

[11] N. Tsiogkas, Z. Saigol, and D. Lane, "Distributed multi-auv cooperation methods for underwater archaeology," in *OCEANS 2015 - Genova*, May 2015, pp. 1–5.

[12] A. Jena, "A free and open source java framework for building semantic web and linked data applications." 2007. [Online]. Available: https://jena.apache.org/

[13] L. Walsh, V. Akhmechet, and M. Glukhovsky, "Rethinkdb-rethinking database storage," 2009.

[14] C. Petrioli, R. Petroccia, J. R. Potter, and D. Spaccini, "The SUNSET framework for simulation, emulation and at-sea testing of underwater wireless sensor networks," *Ad Hoc Networks*, vol. 34, pp. 224–238, 2015.

[15] V. Di Valerio, F. Lo Presti, C. Petrioli, L. Picari, and D. Spaccini, "A Self-Adaptive Protocol Stack for Underwater Wireless Sensor Networks," in *Proceedings of MTS/IEEE OCEANS 2016*, Shanghai, China, April, 10–13 2016, pp. 1–8.

[16] M. Radojevic, M. M. Nawaf, F. Maurelli, T. Tomassson, C. Wang, J. Clay, and Y. Petillot, "Nessie vi autonomous underwater vehicle," 2011.

[17] E. GMBH, "Evologics s2cd 18/34," 2017. [Online]. Available: https://www.evologics.de/en/products/acoustics/s2cr_18_34.html

[18] T. Salumäe, R. Raag, J. Rebane, A. Ernits, G. Toming, M. Ratas, and M. Kruusmaa, "Design principle of a biomimetic underwater robot u-cat," in *Oceans-St. John's, 2014*. IEEE, 2014, pp. 1–5.

[19] T. Salumäe, A. Chemori, and M. Kruusmaa, "Motion control architecture of a 4-fin u-cat auv using dof prioritization," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2016.

[20] M. Tenorth and M. Beetz, "Knowrob: A knowledge processing infrastructure for cognition-enabled robots," *The International Journal of Robotics Research*, vol. 32, no. 5, pp. 566–590, 2013.