Marine world representation and acoustic communication: challenges for multi-robot collaboration

Francesco Maurelli, Zeyn Saigol, Carlos C. Insaurralde, Yvan R. Petillot, David M. Lane
Ocean Systems Laboratory
School of Engineering & Physical Sciences
Heriot-Watt University
EH14 4AS Edinburgh
Email: f.maurelli@hw.ac.uk

***Abstract*—This paper presents some of the challenges related to multi-robot cooperation for the marine environment. Special attention is given to the world representation topic and to the communication challenges. Ontologies represent the tool to store and dynamically update world information. Due to the conditions of the underwater domain, communication among robots presents several issues. The exchange of information between the local world model of each robot, and those of the other robots needs to properly address specific points, such as limited bandwidth, reliability of the acoustic channel, selection of the information to be shared with other vehicles and information merging with previous knowledge of the world. Three scenarios will be then analysed: the Pandora project, with an emphasis on persistent autonomy, world modeling and failure management through appropriate ontologies; the Trident project, which deals with joint missions with an Autonomous Surface Vessel (ASV) cooperating with an Autonomous Underwater Vehicle (AUV), and the Arrows project, which envisages the use of a fleet of AUVs for underwater archaeology operations.**

I. Introduction

Day after day, autonomous robots are more and more used in maritime operations. However, their uses remain often constrained to very specific tasks. Autonomous systems routinely blindly execute a state machine. Adaptation to the environment or reaction to failures - being them on hardware, software or logical - is very limited. This is particularly problematic in the case of multiple vehicles where knowledge discovery is rarely shared, responses to sensor or platform failure rarely dealt with and architectures for collaborative planning in the presence of low communications bandwidth is a real challenge. The challenges related to this domain require the development and integration of more evolved embedded tools that can raise the platform's autonomy levels while maintaining the trust of the operator. Autonomous adaptation is crucial not just for the enhancement of the range of tasks and mission success, but also to release the operator from decision making tasks during the lifespan of the mission itself. As a consequence, less communication is required, which not only results in power saving, but it could be very much desirable in military scenarios, when communication is often best avoided.

Adaptation plays an important role in providing autonomy. Plans needs to be adapted when faced by unexpected or changing circumstances, whether environmental, new goals being required or sensor or components failure. Autonomous adaptation requires an autonomous understanding of the environment. Perception of the environment, comprehension of the situation and projection of the future status are key areas to deal to understand highly dynamic and complex environments. Increasing the levels of situation awareness can help the transfer from current full human control to fully autonomous unmanned capabilities.

However most current knowledge representation systems for marine robots are generally very simple and target mono-platform and mono-domain applications, therefore limiting the potential of multiple coordinated actions between multiple agents. Consequently, the main application for autonomous underwater vehicles is data gathering from sensors. The data are generally off-line manually processed to gather information at a later stage. However, in order to arrive to higher levels of autonomy and control, robots require access to higher levels of knowledge representation or abstraction. These higher levels will be required to provide knowledge representation for contextual awareness, temporal awareness and behavioural awareness. This paper focuses on distributed world modeling and presents some of the challenges of underwater acoustic communication, presenting then various application scenarios in which the Ocean System Lab (OSL) is involved.

II. Distributed Knowledge Representation

During the mission life, new information is normally gathered by the robots. Effectively sharing this new knowledge to other cooperating robots is very important. In terms of communication middleware, different message transfer protocols have been proposed to share information [17]. They enable embedded agents to read data from sensors, to write commands to actuators, and to configure devices on the fly. Blackboard approaches [12] are vulnerable to bottle-necking' at the server as the ASCII messages can generate considerable parsing overheads. On the other hand, distributed broadcasting

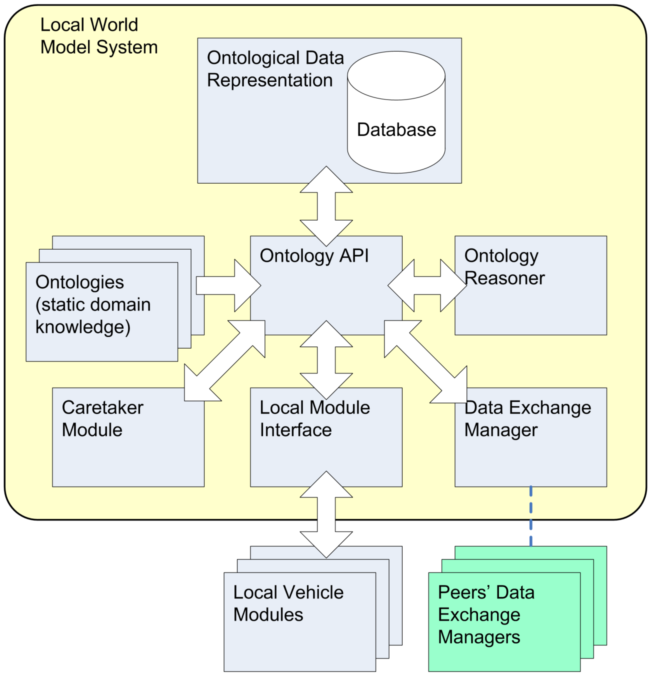

Fig. 1. Proposed World Model Architecture

approaches can flood the communication network. Neither of these approaches standardises the semantic content of the information transferred. The Joint Architecture for Unmanned Systems (JAUS) proposes a common set of concept and services, generic across the robotics domain (land, air, underwater). These concepts and services can make the link between the data observed; the information extracted from this data; and the domain knowledge (what to do with the data?, what does it mean for the mission?) . This is not currently done in JAUS. The OSL has successfully realised a semantic world model framework to combine sensor data with domain knowledge, based on JAUS [11]. Our representation of the knowledge is based on ontologies. They allow different embedded agents to communicate shared concepts whilst keeping sole responsibility and awareness of the services that they provide. The representation of the knowledge base is displayed in Fig. 1.

Each vehicle has its own world model storing the current state of the world and of the vehicle itself (including the vehicle's goals and intentions). The vehicles communicate using a subscription and priority-based information distribution protocol. This protocol takes into account the low-bandwidth of acoustic communications and information priority. Platforms only subscribe to information that they can use.

III. World Modeling and Ontologies

Several ways have been explored in order to represent world data in a connected system, which is able to express relations among different entities. Recently, there has been raised interest in using semantic frameworks, which can provide a hierarchical distributed representation of knowledge for multidisciplinary agent interaction [13]. They provide a common machine understanding representation of knowledge between embedded agents that is generic and extendible. They also include a reasoning interface for inferring new knowledge from the observed data and knowledge stability by checking for inconsistencies. These frameworks improve local (machine level) and global (system level) situation awareness and context for mission and trajectory behaviour. They can therefore act as enablers for autonomy and on-board decision-making. There are currently several institutions developing standards for knowledge representation under these frameworks. Particular attention is taking the effort describing the concepts and relationships for the domains of unmanned platforms [13] and the underwater environment [2].

Ontologies are a way of representing data in a structured manor, such that it is easy for both humans and machines to interpret the data contained in them. Conventionally the contents of an ontology are divided into the `TBox` (terminological box) and `ABox` (assertion box), where the `TBox` is equivalent to a class hierarchy in object-oriented programming, and the `ABox` would be a container for all instances of the classes.

The world modelling system proposed can be used with any `TBox` ontology, provided the same ontology is supplied to all platforms running the system. To enable efficient transmission of data, a unique but short UID will be created for every class in the `TBox`, and these UIDs used to instantiate or refer to instances of the classes.

While the `TBox` is fixed, the `ABox` can be supplemented at run-time by recording new observations or objects, and these will be automatically shared between all nodes in the system. Observations should be formatted according to a supplemental `TBox` which defines a compact way of storing data about objects that may change over time. This `TBox` is based on situation awareness ontology work by [7], and defines and `Attribute` class that can be associated with an object of any class. An `Attribute` can represent any property, for example “name” or “longitude”. Rather than storing a single value for an `Attribute`, we allow an `Attribute` to have many `PropertyValues`, where a `PropertyValue` stores both a value and a timestamp (which is taken to be the time the value was observed at). A `PropertyValue` can store either a data value (such as an integer or string) or a link to another object in the `ABox`.

There are widely available and proven tools for building ontologies in languages such as OWL [10]. Ontology reasoning engines are able to perform logical consistency checking of ontologies, and infer new information based on ontology axioms and rules. They can be key for multi-AUV communication, as the different vehicles exchange values with references to ontological concepts.

IV. Data Exchange Manager

A key part of any multi-robot operation is the communication among the different agents. In our case, the robots want to share world model information. In this section, different challenges are addressed focusing on the Data Exchange Manager, which orchestrates the exchange of information between the local world model for a vehicle, and those of the other cooperating vehicles. Tasks performed by the Data Exchange Manager include the selection of information to send to other vehicles based on utility and communications availability, and the merging of received data into the local world model whilst maintaining consistency. When accessing the local world model, the Data Exchange Manager uses the ontology API. Some of the issues surrounding the exchange of world model information between vehicles are discussed below.

- **Access control of broadcast medium**
  For communication between autonomous vehicles, wireless media are generally used, such as radio waves through the air, or acoustic signals underwater. These are broadcast media, so a mechanism for media access control (MAC) is required to reduce transmission collisions. For the purposes of this research it is assumed that a MAC facility is in place, such as in [14].
- **Bandwidth available per vehicle**
  In a simplistic broadcast scenario, if each vehicle is to be given an equal portion of bandwidth, on average the bandwidth available is inversely proportional to the number of operating vehicles. Communications bandwidth is especially limited in the underwater domain that this project focuses on; typical throughput is of the order of a few kilobytes per second. Thus it is essential to make the most effective use of this bandwidth.
- **Consistency/integrity of the knowledge base**
  When information is received from other vehicles, care must be taken to ensure that the vehicle's local world model remains consistent, without dangling links, etc. A simple example would be where a block of information is to be exchanged, but it will not fit in a single message, and when split, each part is incomplete. Upon receipt of part of this information block, the data exchange manager must cache it until a whole consistent block has been received.
  The question is, who determines that the block is complete? The transmitting data exchange manager, which would convey the boundaries of the block to the recipient; or the receiving data exchange manager. It is proposed that a combined approach is taken. Firstly, the sender must delimit the boundaries of a complete block, to be understood by the recipient. This handles the issue of consistency within the block, by requiring the whole block to be present before insertion. Secondly, the recipient is responsible for checking that any dependencies with other instances in the distributed world model are satisfied before the block is inserted into the local world model. The later check ensures that, for example, the sensor instance or even the vehicle that the information relates to is present in the local world model. Should this check fail, the local data exchange manager could either discard the block, or request the instance information it is lacking.
- **Selecting information to be shared**
  The type or level of information to be exchanged depends both on the needs of the recipients, and the capabilities of the communications channel. Given less communication bandwidth, either less frequent information or more concise, higher level information must be transmitted. If performing a covert mission, very little or no explicit communication will be possible. It is intended that information on the needs, capabilities and status of links to all the peers be maintained by the data exchange manager, using the world model for storage where appropriate. The data exchange manager utilises and updates this information during its communications.
  To perform selective push-based distribution, the sender must have some concept of the value of the information to its peers. To this end, each platform maintains its own set of information needs, where each need specifies a named attribute of either an ontological class or specific class instance, referenced by UID. For example, one information need could be "position of object of interest", and another "classification of object of interest". Each information need has an associated priority, which is used in the exchange selection algorithm. To bootstrap the exchange process, the system assigns an implicit need on all peers to receive the local node's information needs. This allows nodes to automatically discover the needs of their peers, while supporting online modification of information needs as the mission progresses.
  Observations that match an information need are referred to as "exchange candidates". A weighted queue is used to rank the exchange candidates. Each time the local node's slot transmission time is reached, the following actions are performed:
  1) Check for modified information needs, grouping the same needs from different peers.
  2) If a new peer is added to a need, mark previously completed candidates as potentially up for retransmission.
  3) Check for new exchange candidates matching needs, and ACK timeouts of previously transmitted candidates.
  4) Remove candidates that are now known by all peers that have the associated need.
  5) Increment the weight of all exchange candidates by need priority.
  6) Iterate through candidate list, sorted by decreasing weight, adding candidates to the exchange packet until no more will fit.
  7) Mark chosen candidates as in progress and start ACK timer.

  Whilst many works in the literature are concerned with reliable acoustic exchange mechanisms, they operate mainly at the packet level. By performing selection for (re)transmission on exchange candidates, our system only retransmits unacknowledged observations that have not become obsolete, or are known to have been successfully delivered by another peer. It also provides the option to seamlessly include repeated observations from another peer together with the node's own observations.

- **Reliability of data exchange**
  Wireless communication systems usually have some degree of loss; nowhere is this more true than of acoustic underwater communications. Data exchange between local world models must therefore be robust to this. Where information such as a regular status update is sent, it may be that data loss is merely tolerated. In the case of critical non-repeating transmissions, loss must be detected and remedied. The information exchange protocol includes the facility to request an acknowledgement of receipt, which will be used where necessary. When an ack is not requested, it will be the sender's responsibility to periodically repeat transmission of important information.
- **Merging of received information into local world model**
  It is not intended that information received from a peer vehicle will ever replace or overwrite information present in the local world model. Rather, each piece of information will be differentiated by its globally unique identifier (GUID), and merely add to the total knowledge store. For example, during an archaeological survey a possible wreck is detected by a transit AUV and its presence recorded with an "unknown classification"; the classification for this target could later be received from a hover-capable AUV, and added to the world model. This classification would be more recent, and thus take precedence. The origin of a piece of information may be tagged, to enable the flow of information to be traced. This would also prove useful in disregarding information if a particular vehicle is found to be malfunctioning.
- **Communications protocol and data encoding**
  In terms of communications protocol, the author considers agent communication languages such as KQML and FIPA-ACL to be unnecessarily complex for this application. In particular, it is not intended that communication be considered a plan-able act for the purposes of this system, as this would tie the distributed world model too much to the planning system of the vehicle. Thus a simple packet based protocol has been created, to allow ontological instance data to be transferred from one vehicle to another with the minimum of overhead, both of processing and bandwidth. TDMA-based exchange policies have been discussed in [9].

V. Application Domains

In this section, three different projects will be briefly described. The first one, Pandora, aims to vehicle persistently autonomous. In order to reach that, a key part is a semantic high-level representation of the reality which is then updated according to the new information gathered during the mission. The project Trident presents the cooperation among an Autonomous Underwater Vehicle (AUV) and an Autonomous Surface Vessel (ASV) and the associated challenges. Finally, the Arrows project envisages cooperation of multiple AUVs for archaeology operations.

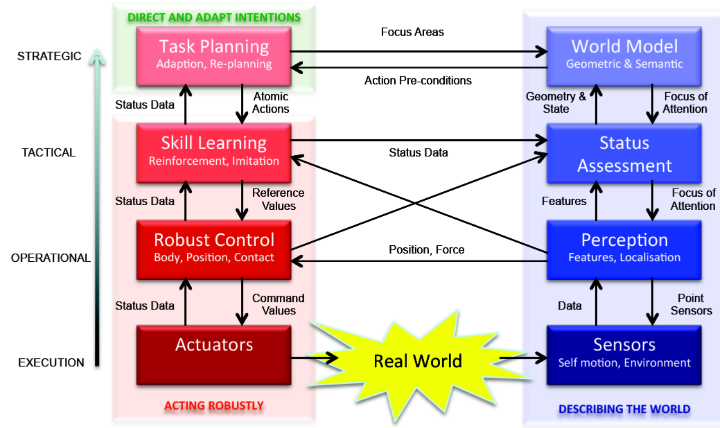

Fig. 2. PANDORA: Computational architecture to develop and study persistent autonomy

A. EU FP7 Pandora

The aim of the project is to extend the range of tasks that can be carried on autonomously and increase their complexity while reducing the need for operator assistances [4], [5]. Dynamic adaptation to the change of conditions is very important while addressing autonomy in the real world and not just in well-known situation. The key of Pandora is the ability to recognise failure and respond to it, at all levels of abstraction. In order to achieve these goals, ontologies are the primary instrument that is used in order to represent the reality.

Figure 2 outlines the computational architecture designed for development and study of persistent autonomy. Key is the notion that the robots response to change and the unexpected takes place at one or a number of hierarchical levels.

At an Operational level, sensor data are processed in Perception to remove noise, extract and track features, localise using SLAM, in turn providing measurement values for Robust Control of body axes, contact forces/torques and relative positions. One of the goals is to further explore some of the current approaches [1], [6] and integrate them on a real vehicle. In the cases where a map is given, localisation techniques will be used [15], with a specific attention to active localisation [8].

At a Tactical Level, Status Assessment uses status information from around the robot in combination with expectations of planned actions, world model and observed features to determine if actions are proceeding satisfactorily, or have failed. Alongside this, reinforcement and imitation learning techniques are used to train the robot to execute predetermined tasks, providing reference values to controllers. Fed by measurement values from Perception, they update controller reference values when disturbance or poor control causes action failure.

Finally at a Strategic level, sensor features and state information are matched with geometric data about the environment to update a geometric world model. These updates include making semantic assertions about the task, and the world geometry, and using reasoning to propagate the implications of these through the world description. Task Planning uses both semantic and geometric information as pre-conditions on possible actions or action sequences that can be executed. When Status Assessment detects failure of an action, Task Planning instigates a plan repair to asses best response, if any. Where there is insufficient data to repair, Task Planning specifies Focus Areas where it would like further sensor attention directed. These are recorded in the World Model and propagated through Status Assessment as Focus of Attention to direct the relevant sensors to make further measurements.A big emphasis of the Pandora project is therefore on the world representation area, on the possibility to update it dynamically and to reason on it, in order to respond to failures.

B. EU FP7 Trident

The project objective is the design and implementation of a methodology enabling multipurpose underwater intervention missions with a very high autonomy level [16]. The Trident approach is supported by a multi-vehicle marine platform. The OSL contribution to this project is the development of the intelligent control architecture for autonomous marine vehicles.

The main challenge when tackling a self-governing solution for a team of marine vehicles is to develop a system that can perform complex tasks reliably and with minimal operator intervention. A critical issue to achieve this is to design and build a system with the ability to deal with internal faults, changes in the environment as well as their impact on sensor outputs used for the planning phase. Therefore, this new generation of marine vehicle platforms requires a certain degree of autonomy (including coordination among vehicles and underwater manipulation facilities), and a collaborative operation mode in order to minimise the operator intervention. They also have the requirement of providing on-the-fly re-planning of activities when needed.

Trident proposes an Intelligent Control Architecture (ICA) to enable multiple marine vehicles to carry out autonomous multipurpose underwater intervention missions [3]. Thus, an Autonomous Surface Vehicle (ASV) and an Intervention Autonomous Underwater Vehicle (I-AUV) are required to work cooperatively. They are capable of cooperating towards the execution of complex activities since they have different but complementary capabilities. The proposed ICA moves away from fixed mission plans and very elementary diagnostics scheme currently used in the community to a more robust architecture to deal with the above missions. It is able to handle unexpected faults within vehicles as well, e.g. at the sensor and sensor processing levels based on either hardware failure or environmental changes. The architectural foundation to achieve the ICA lays on the flexibility of Service-Oriented Computing (SOC). Each vehicle module provides basic services which advertise their capabilities to the system. The service also publishes regular updates on its current status. In addition,

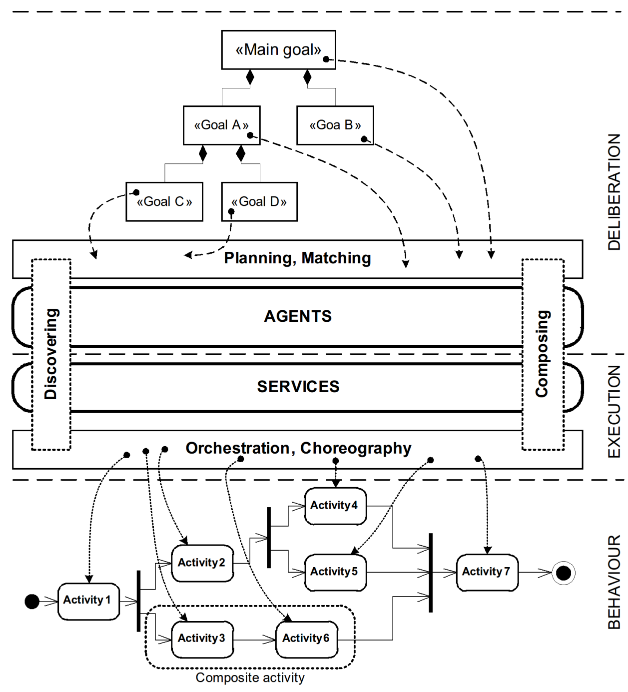

Fig. 3. Service-oriented agent-based architecture

a knowledge-based database (ontology) captures the domain specific skills of the human expert (how to perform a specific task) as well as the dynamic information concerning the environment and platform capabilities. This makes possible to include small atomic plans to test and validate the performance of specific services. The knowledge captured enables high-level reasoning agents to monitor, refine or adapt mission plans based on the current situation.

The resulting architectural solution is a service-oriented agent-based approach which is suitable for integrating the vehicle modules (from project partners) as well as the capabilities of each marine vehicle in a collaborative manner, as outlined in Fig. 3. In the top of the figure, the system deals with hierarchical mission goals that are achieved by the execution of agent plans (sequence of command messages that invoke agents' activities). The planning of activities is carried out by matching agents' capabilities with required activities. The agents can discover each other's and monitor them in near real-time. In the bottom of Fig. 3, the activities can be seen as service processes (execution of services). They can have a basic or composite structure. The basic ones are indivisible and the composite ones can be decomposed in other activities. This composition of activities or service processes is called orchestration of services. It plays an important role in the system architecture since it can define different encapsulation levels to execute services. On the other hand, choreography of services deals with the message exchanges among services that are executed in parallel (collaborative nature). Orchestration and choreography are terms from SOC. The ICA implementation is based on the Robot Operating System (ROS) middleware. Services or capabilities from different project partners (navigation, mapping, vision, communication, manipulation, guidance, etc.) are managed as messages, services, and actions in ROS. The approach has been tested in computer simulations, and the first tests of partial functionalities in real marine vehicles have already begun.

C. EU FP7 Arrows

Arrows proposes to adapt and develop low cost autonomous underwater vehicle technologies to significantly reduce the cost of archaeological operations, covering the full extent of archaeological campaign. One of the most innovative areas of development of the Arrows project concerns the investigation of algorithms for planning and adaptation of missions for the optimisation of undersea surveys by two or more autonomous vehicles. This approach allows to optimize the coverage of large areas of seabed, reducing the time and cost of research missions. These algorithms will be tailored to the needs of the archaeological missions. As an example, the existing database of historical data already developed by DBCIS for the Seas around Sicily, allows to associate the investigation area with one or more probability maps, related to the possible presence and to the typology of wrecks in that area, thus allowing to initialize the mission. During the mission, anomalous events, autonomously identified by the real-time analysis software installed on vehicles, allow to modify these probability maps, allowing the reconfiguration and online replanning of the mission by each vehicle. The identification of an anomalous event (and of its type: optical, acoustic, magnetic) is communicated from the vehicle that carried it out the other vehicles through acoustic communication. As described in the previous sections, the underwater acoustic communications are limited both in distance and in bandwidth. The cooperation algorithm envisaged for Trident must therefore manage the relative positioning of vehicles, so that the connectivity of communication is maintained, and the management of this loss of connectivity. Moreover, the approach to the re-planning is fully distributed: each vehicle operates through elementary behaviour algorithms such to ensure the collective satisfaction of the mission algorithms without the need of a centralized decider. In this way the robustness of the mission will be guaranteed even in case of operating anomalies.

Field trials are scheduled in the Mediterranean Sea and in the Baltic Sea.

VI. Conclusions

This paper has presented two key challenges related to underwater cooperation of a team of robots: semantic world representation, using ontologies, and communication issues in the underwater domain.

Three scenarios have been presented, coming from three different current projects OSL is currently involved in. The vision for the success of those projects is the increase of the range of operations autonomous robots can perform underwater, without a direct, continuous, expensive link with the human operator. The Trident and the Pandora projects have already shown promising results, whilst the Arrows project has just started and has an exiting research time in front of it, with a well committed team.

Acknowledgment

The authors would like to thank all the members of the Ocean Systems Laboratory at Heriot-Watt University. The research leading to these results has received funding by the EU Seventh Framework Programme FP7/20072013 Challenge 2 Cognitive Systems, Interaction, Robotics under grant agreement No 288273 PANDORA; FP7-TRIDENT Project SCP8-GA-2010-248497; FP7 ARROWS - under grant agreement No 308724.

References

[1] Josep Aulinas, Marc Carreras, Yvan R. Petillot, Joaquim Salvi, Xavier Llado, Rafael Garcia, and Ricard Prados. Feature extraction for underwater visual slam. In *IEEE/OES OCEANS 2011*, 2011.

[2] National Science Foundation. Marine metadata interoperability. In *www.marinemetadata.org*, 2004-2012.

[3] C. C. Insaurralde, J. J. Cartwright, and Y. R. Petillot. Cognitive control architecture for autonomous marine vehicles. In *Proceedings of 6th IEEE Int. Systems Conf., Vancouver, Canada*, 2012.

[4] David M. Lane, Francesco Maurelli, Petar Kormushev, Marc Carreras, Maria Fox, and Konstantinos Kyriakopoulos. Persistent autonomy: the challenges of the pandora project. In *Proceedings of IFAC MCMC 2012 - Manoeuvring and Control of Marine Craft*, 2012.

[5] David M. Lane, Francesco Maurelli, Tom Larkworthy, Darwin Caldwell, Joaquim Salvi, Maria Fox, and Konstantinos Kyriakopoulos. Pandora: Persistent autonomy through learning, adaptation, observation and replanning. In *Proceedings of IFAC NGCUV 2012 - Navigation, Guidance and Control of Underwater Vehicles*, 2012.

[6] Chee Sing Lee, Daniel E. Clark, and Joaquim Salvi. Slam with single cluster phd filters. In *Proceedings of the 21st International Conference on Robotics and Automation, ICRA 2012*, 2012.

[7] Christopher J. Matheus, Mieczyslaw M. Kokar, and Kenneth Baclawski. A Core Ontology for Situation Awareness. In *Proceedings of the Sixth International Conference on Information Fusion*, pages 545–552, 2003.

[8] F. Maurelli, A. Mallios, S. Krupinski, Y. Petillot, and P. Ridao. Speeding-up particle convergence with probabilistic active localisation for auv. In *IFAC IAV*, 2010.

[9] F. Maurelli, Z. Saigol, J. Cartwright, D.M. Lane, A. Bourque, and B. Nguyen. Tdma-based exchange policies for multi-robot communication of world information. In *Proceedings of IFAC MCMC 2012 - Manoeuvring and Control of Marine Craft*, 2012.

[10] DL McGuinness and F Van Harmelen. Owl web ontology language overview. *W3C recommendation*, 2004.

[11] Emilio Miguelanez, Pedro Patron, Keith E. Brown, Yvan R. Petillot, and David M. Lane. Semantic knowledge-based framework to improve the situation awareness of autonomous underwater vehicles. *IEEE Transactions on Knowledge and Data Engineering*, 23:759–773, 2011.

[12] P. Newman. Under the hood of the moos communications api. Technical report, Oxford Robotics Research Group, 2009.

[13] P. Patrón, E. Miguelanez, J. Cartwright, and Y. Petillot. Semantic knowledge-based representation for improving situation awareness in service oriented agents of autonomous underwater vehicles. *IEEE Oceans Quebec*, 2008.

[14] Borja Peleato and Milica Stojanovic. A MAC Protocol for Ad-Hoc Underwater Acoustic Sensor Networks. In *WUWNet '06: Proceedings of the 1st ACM international workshop on Underwater networks*, pages 113–115, 2006.

[15] Y. Petillot, F. Maurelli, N. Valeyrie, A. Mallios, P. Ridao, J. Aulinas, and J. Salvi. Acoustic-based techniques for auv localisation. *Journal of Engineering for Maritime Environment*, 224(4):293–307, 2010.

[16] Pedro J. Sanz, Pere Ridao, Gabriel Oliver, Claudio Melchiorri, Giuseppe Casalino, Carlos Silvestre, Yvan Petillot, and Alessio Turetta. Trident: A framework for autonomous underwater intervention missions with dexterous manipulation capabilities. In *in proceedings of the 7th Symposium on Intelligent Autonomous Vehicles IAV-2010*. IFAC, 2010.

[17] M. Somby. A review of robotics software platforms. online, August 2007.