The PANDORA project: a success story in AUV autonomy

Francesco Maurelli
Chair of Robotics and Embedded Systems (I6)
Technische Universität München
Garching bei München, DE
Email: f.maurelli@ieee.org

David Lane
Ocean Systems Laboratory
Heriot-Watt University
Edinburgh, Scotland, UK
Email: d.m.lane@hw.ac.uk

Petar Kormushev, Darwin Caldwell
Advanced Robotics Department
Italian Institute of Technology
Genoa, Italy

Marc Carreras, Joaquim Salvi
Depart. dArquitectura i Tecnologia de Computadors
Universitat de Girona
Girona, Spain

Maria Fox, Derek Long
Department of Informatics
Kings College Londons
London, UK

Kostas Kyriakopoulos, George Karras
Control Systems Laboratory
National Technical University of Athens
Athens, Greece

***Abstract*—This paper presents some of the results of the EU-funded project PANDORA - Persistent Autonomy Through Learning Adaptation Observation and Re-planning. The project was three and a half years long and involved several organisations across Europe. The application domain is underwater inspection and intervention, a topic particularly interesting for the oil and gas sector, whose representatives constituted the Industrial Advisory Board. Field trials were performed at The Underwater Centre, in Loch Linnhe, Scotland, and in harbour conditions close to Girona, Spain.**

I. Introduction

Whilst humans and animals perform effortlessly doing complicated tasks in unknown environments, our human-built robots are not very good at being similarly independent. Operating in real environments, they easily get stuck, often ask for help, and generally succeed only when attempting simple tasks in well-known situations. We want autonomous robots to be much better at being autonomous for a long time (persistent autonomy), and to be able to carry out more complicated tasks without getting stuck, lost or confused. Following the Deep Water Horizon disaster in the BP Macondo oilfield in the Gulf of Mexico in 2010, Oil Companies are developing improved ways to cost effectively and safely carry out more frequent inspection, repair and maintenance tasks on their subsea infrastructure. This is particularly challenging in deep water. To date, Autonomous Underwater Vehicles (AUVs) have been deployed very successfully for various forms of seabed and water column transit survey. First commercial units will soon be applied to simple hovering inspection tasks, with future units expected to address much harder intervention where contact is made to turn a valve or replace a component. Because these vehicles reduce or remove the need for expensive ships, their adoption is expected to grow over the next 5 to 10 years. To be successful commercially, these hovering AUVs must operate for extended periods (12-48 hours +) without the continual presence of a surface vessel. They must therefore demonstrate persistent autonomy in a challenging environment. We therefore choose this application focus to evaluate the projects research, with guidance from BP, Subsea7 and SeeByte Ltd. on the projects Industrial Advisory Group. We have identified four essential research areas where we believe core research is required in order to provide the essential foundations for Persistent Autonomy:

- Describing the World
- Directing and Adapting Intentions
- Skill Learning
- Acting Robustly

We have therefore applied the results of the research in these core areas into three tasks of industrial interest: structure inspection, valve turning and chain cleaning. Overall, the PANDORA project (persistentautonomy.com - [1], [2]) was very successful, with the technology developed demonstrated in the field, with great potentials for industrial applications. The following sections will summarise R&T contributions and will give an overview of the experimental validation in the field.

II. Describing the World

A. The role of the knowledge base

A significant obstacle to intelligent robots being used in real-world scenarios is their limited ability to cope with unexpected events and environments, to deal with faults, and to make smart decisions in response to changes in the world. This section presents a semantic knowledge-representation system that will make it easier for robots to handle such situations, and therefore achieve more persistent and long-term autonomy. The domain we focus on is the inspection of sub-sea structures such as oil platforms using autonomous underwater vehicles (AUVs). In this domain, the original plans for structures can generally be made available to the AUV in advance of a mission. However, frequently these structures have changed or moved when the AUV actually reaches them. If the robot has a pre-programmed survey path, unexpected obstacles can cause it to abort the mission. Reactive obstacle avoidance algorithms can improve the robustness but a proper knowledge representation framework, able to incorporate newly discovered information, is desirable when the aim is for the robot

978-1-4673-9724-7/16/$31.00 ©2016 IEEE

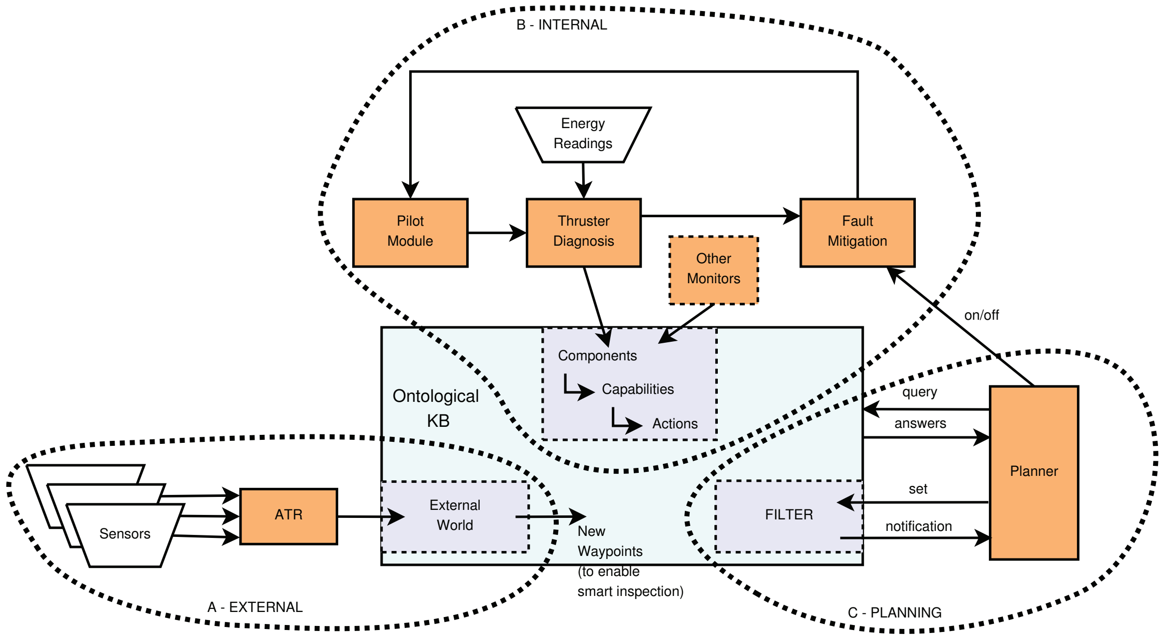

Fig. 1. The role of the knowledge base in the vehicles architecture, and its relations with other components of the vehicle.

to be intelligently autonomous, situation-aware and able to use dynamic world knowledge. Moreover, fault awareness - let alone fault management strategies - has been explored in research but very rarely integrated in a more complex system. In order to address those areas, a semantic knowledge base is introduced, to enable the robot to model and reason with the uncertainty in the world and to overcome possible faults. An ontological representation for semantic data is used for several reasons:
- it is easy to create and examine the concepts used by the system, and the attributes available for each concept, being a standard language with clear rules and widely adopted;
- many tools exist to perform logical reasoning within the knowledge base;
- an ontology represents a well-specified central data store that all the software components comprising the agent can make use of;
- common ontologies (such as the IEEE robotics ontology currently under development) are very useful for exchanging data between robotic systems created by different organisations.

A proper knowledge representation system is pivotal for long term autonomy as it allows to store information, to reason on the acquired information and to augment the knowledge through reasoning. This further allows to have smarter planning and control systems which can query the knowledge base and be notified about important information relevant to the current plan. Figure 1 shows the connections among the knowledge base and the various subsystems. Three subsystems are highlighted for clarity of presentation:
- subsystem A: external knowledge
- subsystem B: internal knowledge
- subsystem C: planner interaction

The subsystem A deals with representing the external world, and handles the full flow from signal to symbol [3]. The subsystem consists of sensors, which output data into an ATR System (Autonomous Target Recognition), which interact with the knowledge base to add new instances and update the uncertainty. However, the knowledge base is not only receiving data, but can be used to actively insert new waypoints, linked to the inserted instances. The semantic knowledge of a possible object can help to trigger a new smart inspection in order to reduce the ambiguity and correctly classify the object.

The subsystem B deals with the internal knowledge of the vehicle. For the knowledge base, this is represented mainly by the relation components $\rightarrow$ capabilities $\rightarrow$ actions. Each possible action of the vehicle can be performed only in presence of predefined capabilities. Those capabilities are dependent on the availability of specific components. The fault management subsystem, implemented for the case of a thruster failure, interacts with the knowledge base to signal specific components which are out of order. In the analysed case, only thrusters are considered as components, but the system is general and can be easily extended.

The subsystem C encompasses the planning system. The knowledge base is an essential tool for the planning system as a way to represent symbolic information that can be used by a planner. Under this subsystem there are all the usual communications with the planning system querying the ontology to get information about the world, in order to be able to plan, and the ontology replying to the queries of the planner. However, the system proposed also deals with the necessity to replan when the world changes and invalidates the current plan. This is done through a structure in the knowledge base, named Filter, which filters the operations happening in the knowledge base and sends notification to the planner system only if any operation is related to any instance of class of interest for the planning system. The planning system therefore interacts with the knowledge base to set / clear the filter, and receives notifications from the knowledge base about its topics of interest.

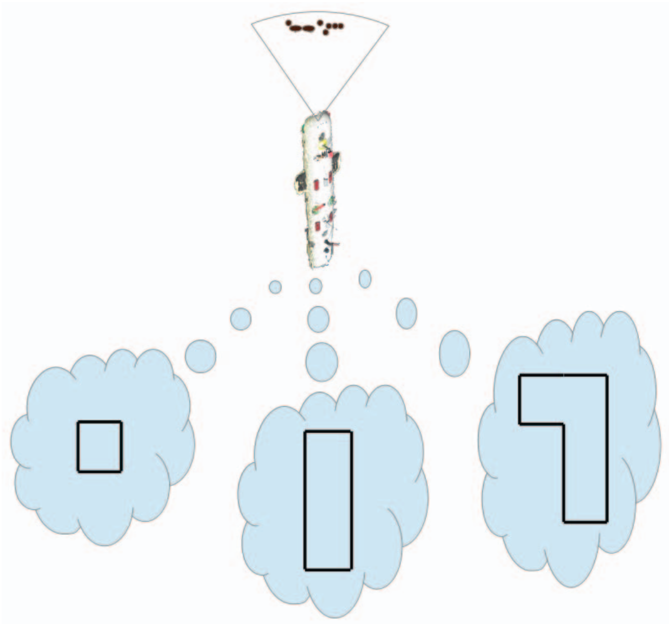

Fig. 2. A single measurement from a sensor is not enough to correctly label the world. The same sensor reading in the Figure can belong to any of the three objects below the robot. The actions undertaken by the robot can significantly vary according to the type of object and therefore a specific active navigation strategy needs to be employed.

B. Smart Inspection

A key function of subsystem A is to process a first signal to symbol analysis. We have developed probabilistic sonar processing techniques to insert information in the knowledge base, implementing Bayesian mechanisms to represent and propagate probabilities that quantify the extent to which semantic statements in the world model are correct. However, the module processes sensor data as they arrive, whereas it is also interesting to analyse possibilities for the vehicle to take decisions based on partial information. For example, if the task to be accomplished is to inspect all pillars, as soon as the sensor processing module shows something which could be a pillar, the planning system might want to generate a new plan that uses this information to better explore the unclassified object. In this section we show how the initial ontological knowledge with uncertainty can help the planning system to define an inspection path, to correctly label the world. This is still part of the Subsystem A, but it is a module which does not deal with external sensor data any more, but already with parsed symbolic data in the knowledge base, linked to the generation of new waypoints. The planning system is then notified through the use of the Filter structure, as previously described. As outlined in Figure 2, after one measurement (or several measurements from the same position) the robot does not know what the structure in front of it is, as the same sensor measurement can represent any of the three objects in the figure. In order to correctly label the object as efficiently as possible, the robot should employ an active path planning module. In this case the vehicles state is known, but the world is represented by a finite set of possibilities. The active path planning module is triggered when the AUV needs to understand which object or structure is in the world and when the initial sensor data are not enough to discriminate among several options. The module outputs a specific path to follow in order to discriminate between them. The general principle of the module is to find a set of actions that can maximise the information gain [4].

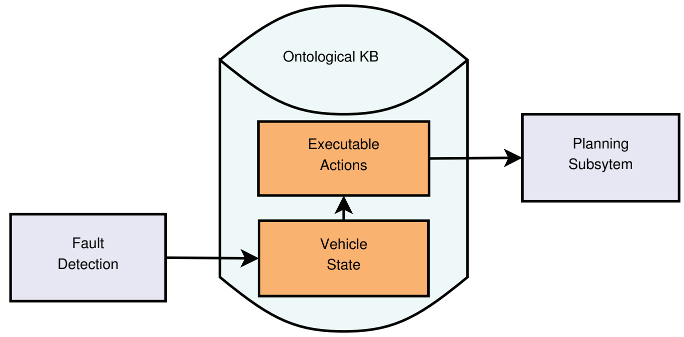

Fig. 3. The role of the Knowledge Base among the planning system and the faults that can happen.

C. Coping with faults

Another application of the knowledge framework is to track information related to the state of the vehicle and be the point of connection with the planning system, based on symbolic representation, and specific capabilities linked to specific actions. Keeping the architecture diagram previously described, this section now describes the subsystem B, and the role of the knowledge base. The Figure also shows a block named Other Monitors. This is to clarify that the thruster failure management has been actually developed and tested, but the system and the interaction with the knowledge base are generally enough that other modules can be plugged in, if the effect on actions available and capabilities is known. Figure 3 shows the role of the knowledge framework. Specific actions can be executed only if the vehicle has a specific configuration. In case of faults some actions might not be available any more.

III. Action planning

A ROS-based architecture was used to integrate all of the components of PANDORA. At the top level, an automated planning system (POPF) builds a mission plan to achieve a set of inspection and maintenance goals, represented symbolically. The structure of the mission plan is used to coordinate the behaviour of the AUV over a long period, in a persistently autonomous way. We have created a Plan-based ROS architecture, ROSPlan [5]. The mission plan contains a number of inter-dependent tactical plans, which are planned in detail, by POPF, when the mission plan execution reaches them. All planned activities are executed through the ROSPlan architecture. When a planned activity fails, because of unexpected changes in the world leading to world model updates, replanning is initiated. We use a background temporal-metric planning system, POPF, to construct mission plans and tactical plans for execution in the foreground ROSPlan plan execution system. POPF starts with a symbolic representation of the world, and the position of the AUV within the world, and a collection of goals to be achieved. Based on this information, it finds a time-efficient collection of actions that will achieve the goals, and an ordering on the starts and ends of these actions. A mission plan is a proof that a given set of goals can be achieved within a given timeframe. To construct it, POPF uses conservative estimates of the costs of each of the component tactical plans. These estimates are then interpreted as resource constraints at the tactical level, and the failure to execute a tactical plan within these constraints leads to plan failure and replanning. As a plan is executed, unspent resources accumulate because of the conservative estimates used in the mission plan. We make use of these accumulated resources by exploiting opportunities. Whenever the AUV sees something of potential value to explore during plan execution, it uses unspent resource to investigate it.

During the execution of a plan, failures can occur because of errors or missing structures in the world model. The ontology is responsible for detecting changes, interpreting them, and informing the planner of changes to its world model [6]. For example, the AUV might approach a structure that was not modelled, and that blocks its path to its destination waypoint. When this occurs, the plan under execution is terminated and the world model is updated by the ontology. Then a new plan to achieve the mission goals is constructed by POPF, using the new model. It is the responsibility of the ontology to determine whether the new structure is permanent, and should be added to the model, or whether it is transitory (a passing sea animal) and can be ignored.

IV. Skill learning

A. Learning by Demonstration (LbD) for valve turning

In order to program the intervention tasks in a flexible and intuitive way, we propose to use a machine learning algorithm known as learning by demonstration (LbD). Rather than analytically decomposing the problem and manually programming a desired behaviour, the LbD infers the knowledge from a set of user demonstrations. This kind of algorithms follow three sequential phases: a Demonstration, where a batch of task examples performed by a human pilot are recorded; a Learning, where a model is created by generalizing all the demonstrations; and a Reproduction, where the model is used to accomplish new instantiations of the learned task. Using this technique, the framework becomes easily extensible and new intervention tasks (involving both the AUV and the manipulator motions) can be added effortlessly just from operator demonstrations. The intervention in free-floating mode, without being attached to the panel, is a challenging problem in which the control of the AUV and the control of the manipulator are tightly coupled. This technique has been extensively applied into this task and has showed its feasibility even in the presence of water current perturbations. Figure 4 shows some images about the different phases of the learning process, and Figure 8 shows how the reproduced trajectories of the AUV and manipulator correctly followed the demonstrated ones [7].

B. Hierarchical and Reactive Skill Learning

We have proposed a learning method for reactive robot behaviour to deal with the challenging task of autonomous valve turning. The autonomous valve turning consists of two main phases: reaching and turning. Imitation learning is used to learn and reproduce the reaching phase. A hybrid force/motion controller is devised to accomplish the turning phase. In order to increase the autonomy of the system a reactive fuzzy decision maker is developed. This module evaluates the dynamic behaviour of the system and modulates the robot's movements reactively. The validity and performance of our approach is demonstrated in a real-world valve-turning experiment both in the lab and in underwater environment. Figure 5 shows the three layers of the developed hierarchical reactive skill learning approach.

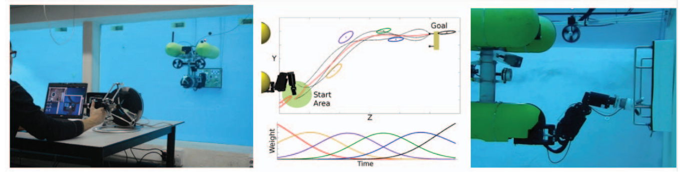

Fig. 4. Learning by Demonstration for valve turning with a free floating AUV. A human operator (left) demonstrates the task. A model of the task is learned (middle) and the AUV is able to reproduce the same trajectory, even in the presence of perturbations, to fulfill the same task (right).

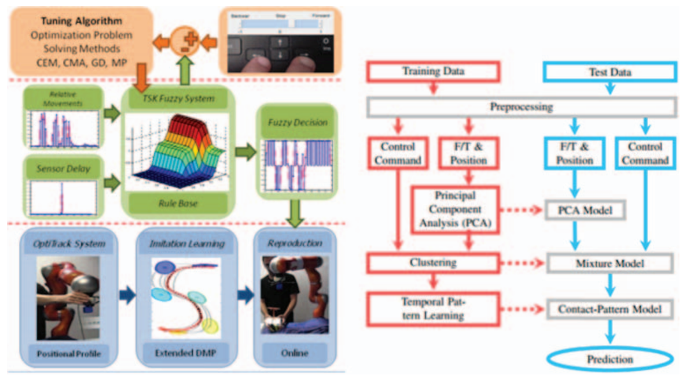

Fig. 5. Left: The three layers of the developed hierarchical reactive skill learning approach. Right: Flow diagram of the developed machine learning method for contact state estimation.

V. Vehicle control

During the last few decades underwater vehicles are being extensively used for complex operations and dexterous tasks such as surveillance and mapping of underwater structures, ship hull inspection, handling of underwater equipment (e.g control panels, valves), search and rescue missions, etc. In most of these cases, the vehicles operate under the influence of strong external disturbances caused from ocean currents and waves. The motion control problem for autonomous underwater vehicles has been an active research field for the past two decades and continues to pose considerable challenges to control designers especially when the vehicles exhibit model uncertainty and are affected by environmental disturbances [8]. A typical motion control problem is way-point tracking which is concerned with the design of control laws that force a vehicle to reach and follow a reference point. In this work, we address two distinct motion control strategies for underwater robotic vehicles: a) A model - based control approach where the dynamics of the vehicle are being exploited and b) A model - free motion control approach. The proposed strategies have been applied to the underwater vehicles involved in the PANDORA Project, namely Nessie VI and Girona 500.

Experimental results prove the efficacy of both motion control schemes .

VI. Perception

In addition to the previous scientific topics, in order to demonstrate them in the envisaged scenarios, other perception modules needed to be developed. Here we will show the algorithms behind the perception of a chain - one of the scenarios in Pandora.

A. *Horizontal chain detection and following*

In order to detect and inspect the chain in horizontal position, lying on the seabed, we use an AUV with a high resolution imaging sonar (Soundmetrics ARIS 3000) which delivers acoustic images at near-video frame rate. This technology allows the detection of the chain links regardless of the visibility conditions and the suspended marine fouling. The use of such a system arises several challenges that must be addressed. First, the automatic detection of the chain links in Forward-Looking Sonar (FLS) images becomes a complex problem due to the inherent characteristics of the sonar modality. Besides, the control of the AUV must be adapted to take into account the imaging geometry of the sonar. The vehicle location at a given instant differs from the point that is being inspected, which is located few meters ahead depending on the sonar's range configuration. Thus, to successfully follow the chain, the detection and control schemes must be tightly coupled and be able to react in real-time. Otherwise, chain links can easily drop of the sonar's narrow FOV resulting in the vehicle losing track of the chain. Hence, the system presented here deals with the problem of autonomously detecting and also following the underwater chain. Figure 6 shows the implemented framework with some results on the FLS image processing.

1) Acoustic mapping: Besides using the sonar for chain detection and following we can take advantage of the FLS images gathered along the chain trajectory to generate an acoustic mosaic that can be useful to perform an overall inspection of the chain state at a general level. To achieve this acoustic mosaicing capability, a complete mapping pipeline has been implemented in order to register consistently the different sonar images. It must be noted that the particularities of FLS imagery, such as low resolution, low Signal-to-Noise Ratio and intensity alterations due to viewpoint changes, pose a significant challenge to the techniques typically used in photomosaicing and therefore new techniques have had to be developed for the main stages of this pipeline. We have proposed a 2D registration methodology for the robust alignment of FLS frames coping with the challenging characteristics of this type of imagery, dealing with translational and rotational motions and applicable to a wide variety of environments. The pairwise registration of sonar images is a key step in the mosaicing pipeline which is typically addressed using feature-based techniques when working with video images. However, on sonar images, registration through feature matching becomes unstable, and therefore we have proposed to use a Fourier-based methodology that, by involving all image content into the registration, offers robustness to noise and the different artifacts associated with the acoustic image formation [9]. Figure 7 shows an acoustic mosaic of the chain which is used for inspection.

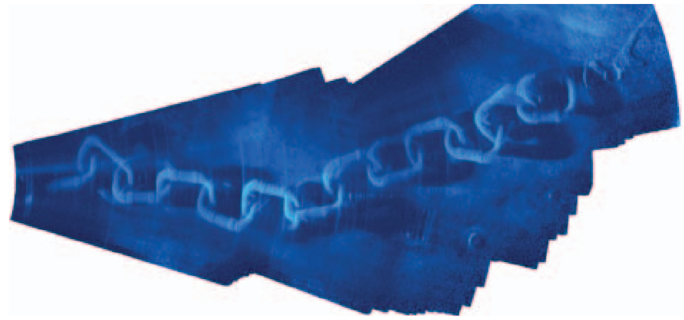

Fig. 7. Acoustic mosaic of the mock-up chain.

VII. Integrated Experiments

The experimental part of the project consisted in demonstrating the research activities of the project in 3 distinctive scenarios:

- Structure inspection: leaded by HWU with Nessie AUV.
- Chain inspection and cleaning: leaded by UdG with Girona 500 AUV.
- Valve turning: leaded by UdG with Girona 500 AUV.

A. *Structure inspection*

The combination of geometric and semantic representation allows the PANDORA vehicles to adapt on-the-fly their mission and to carry out their task in presence of uncertainty about the environment. The goal of these experiments is to evaluate the behaviour of a PANDORA vehicle in the context of the Autonomous Structure Inspection task. Different software modules are integrated and tested in real sea condition. The inspection task is executed with the support of the Autonomous Target Recognition (ATR), Knowledge Representation, Planning, and Robust Control software modules. All partners were involved and actively contributed bringing their core experience in the various areas addressed. Three main modules have been used in these experiments. The first is the Planning architecture, used to drive the behaviour of the AUV during the inspection task. The second is the knowledge representation module (Ontology) used to represent the environment in which the vehicle is operating. Third is a target recognition module used to transform sonar images into a symbolical representation to be stored into the knowledge base and allow the higher level modules to adapt the vehicles mission. Supporting these is the Robust Control architecture to allow the vehicle to cope with unknown disturbances of a real sea environment. An observe action is made available in the planning domain allowing the vehicle to use its sensors (in this case the on-board sonars of Nessie AUV) to inspect a portion of the environment in the search for underwater structures in a partially known environment. The vehicle is provided with some prior knowledge about the presence of objects in the mission area and it is tasked to conduct autonomously a close inspection of the underwater structure. At the begin of each mission the planning architecture calculates the optimal plan based on the semantic information provided. The vehicle starts its inspection by reacquiring the known objects while building its geometrical representation of the environment.

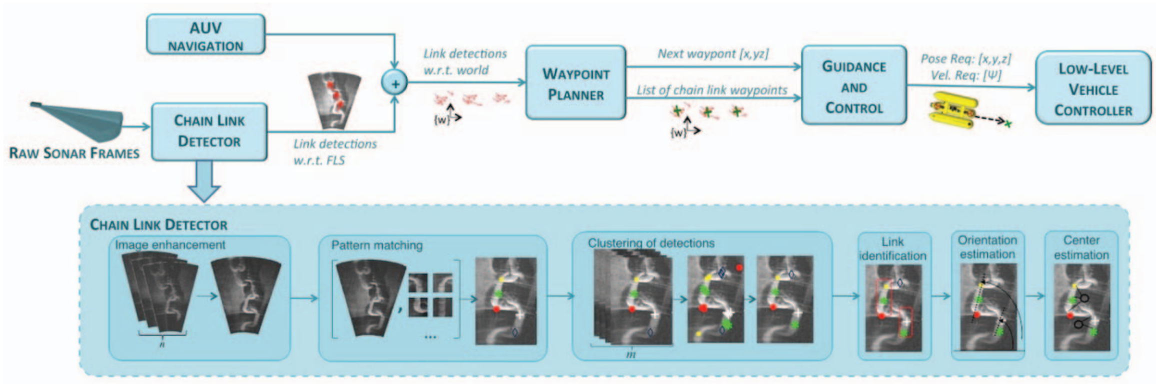

Fig. 6. Diagram illustrating the different steps of the chain detection and following framework.

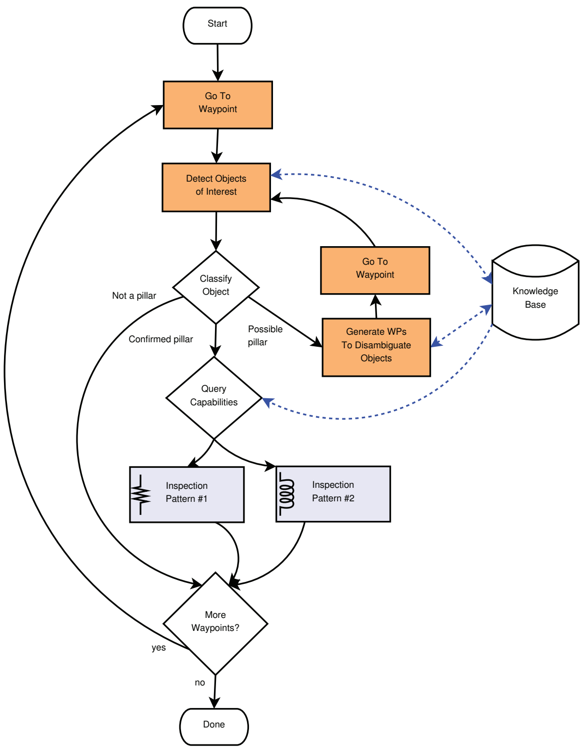

Fig. 8. Acoustic mosaic of the mock-up chain.

This allows the vehicle to increase its awareness of the environment, updating the prior information if needed, and identifying new objects encountered along its route. In order to support the autonomous inspection task the vehicle is provided with a geometric mapping framework, based on the Octomap software. This allows the vehicle to geometrically represent the external environment as the inspection task is conducted. Figure 8 shows the model of system performing an inspection task, in relation to new items discovered and faults, which trigger the execution of alternative actions. Field trials were performed at The Underwater Centre, in Scotland, in a mission to inspect the pillars of the pontoon.

B. Chain inspection and cleaning

Second scenario covers the inspection and cleaning of a mooring chain with an autonomous underwater vehicle. For that, we have used a high resolution imaging sonar which delivers acoustic images at near-video frame rate, in order to detect each of the links and follow the chain. In this way, the system can operate regardless of the visibility conditions and the suspended marine fouling that may arise during cleaning. However working with sonar data introduces several challenges (noisy data, insonification artifacts, narrow field of view, etc.) that had to be addressed. We have tackled the problem in two different configurations: a chain lying on the seafloor and a chain suspended vertically in the water column. For each of these configurations we have provided solutions for chain detection and for chain following using forward-looking sonar and also multibeam data in the vertical case. Moreover, for the vertical case we attempted also the case of a chain in motion, however the detection and following could only be performed reliably in the case of a static chain. Besides, we have evaluated the feasibility of performing autonomous chain cleaning using electrical pressure washer as a water jet system. Using a robust controller that is able to compensate for the forces originated by the waterjet we achieved also to simulate vertical chain cleaning with the AUV.

After successful performance of the chain detection and following algorithms in the water tank, we attempted the same procedures at sea, thus performing a final demonstration one step closer to a real operational environment, and exposing the system to more challenging conditions (larger environment, worse visibility, water currents, etc). Finally, we have also developed a system for forward-looking sonar mapping to perform a first evaluation of the chain state at a high level. This allows to see an overall view of the spatial layout of the links in the environment as well as provides a map of increased signal-to-noise ratio with respect to the individual frames in which features on the range of few centimeters can be appreciated. That might be enough to distinguish large issues in the chain links or the presence of marine growth. Although the developed work has focused on underwater chain following, we believe that by further developing on the involved disciplines (sonar image processing, acoustic mapping, planning and control) the system can become more widely applicable. For instance it could contribute to the assessment of other structures such as pipelines, harbor or dam infrastructures where visibility conditions typically restrict the use of optical inspection methods. Figure 9 shows the Girona 500 AUV with the payload required for chain detection and the diagram of

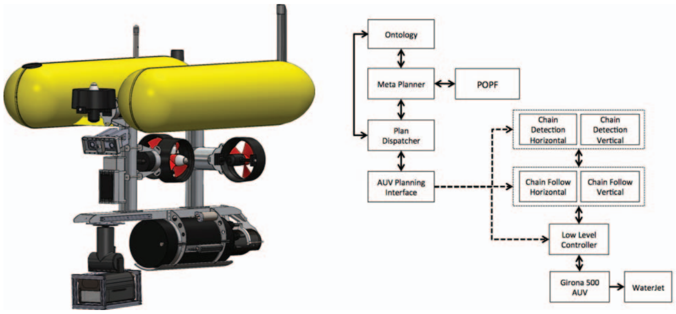

Fig. 9. Girona 500 AUV with the chain detection payload: Forward Looking Sonar, Pan & Tilt unit, Multibeam Sonar, Stereo Vision Camera. On the right, the diagram of the control architecture for the chain inspection task.

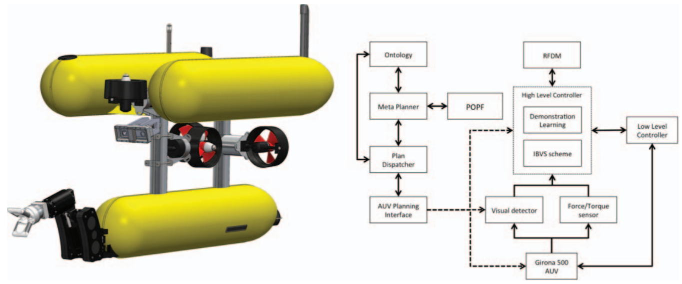

Fig. 10. Girona 500 AUV with the valve turning payload: 4 DOF manipulator, End-Effector with camera and F/T sensor and Stereo Vision Camera. On the right, the diagram of the control architecture for the valve turning task.

all systems that were in charge of fulfilling the task.

C. Valve turning

In the valve turning scenario, we have implemented and tested a framework in which persistent autonomy goes beyond survey and inspection capabilities to autonomous intervention capabilities. We have combined different disciplines involving: visual-based detection and localization systems, robust control schemes, learning-by-demonstration techniques for intervention and a temporal planning architecture, to operate many times a valve panel mock-up in front of several disturbances: water currents, blocked valves, panel unknown position and panel occlusion. All these systems and methodologies have been validated individually and in an integrated long term experiment, in which the AUV locates an intervention panel and performs multiple valve turning operations while handling several failures (either spontaneous or induced) that take place along the mission time. Experiments have only been performed in water tank, in which exhaustive experiments were possible testing differents conditions and configurations. The positive results encourage the use of, in the short-term, autonomous robots operating in subsea facilities performing interventions with a cost benefit, in comparison with teleoperated vehicles. However there are still fundamental requirements to be addressed such as improving underwater electrical manipulator's capabilities with those of their above-water counterparts, and improving or replacing vision-based techniques by others techniques which can act more robustly.

The long-term mission was carried out in a water tank of 16x8x5 meters using the Girona 500 AUV equipped with a 4 Degrees of Freedom electrical manipulator, a Stereo Camera and a specifically designed end-effector, which had a camera in-hand and force and torque sensor. A mock-up panel of 0.8x0.5 meters with 4 valve handles was used to emulate a subsea panel from the offshore industry. To make the environment more realistic, two propellers, able to generate up to 14Kg of thrust each, were placed close to the panel in order to generate lateral water currents. Experiments were performed in a completely autonomous mode, the vehicle ran on its own batteries and all required processing was performed with the on-board computers.

In the experiment, the vehicle had to locate the intervention panel among different locations and modify the valve handles to achieve different panel configurations. The planning algorithm generated the inspection points to locate the panel, which was located by a vision-based detection system. The vision system was also determining the state of the valve handles, which was used by the planning system to generate the actions to turn the valves and to achieve the desired panel configuration. In the process of turning a valve, several systems were working: vision system for panel and valve detection; robust controller for vehicle and manipulator control; Learning By Demonstration for moving the AUV and the manipulator; reactive fuzzy decision maker for deciding if the task could be completed; and Force and Torque sensor processing for determining the contact and turning of the valves. After attempting a valve turning, the planning checked again the state of the valves, and decided new actions in case the action failed due to perturbations [10]. During more than 3 hours, the AUV changed the panel to 9 different configurations, which required the turning of 29 valves. In order to accomplish these valve turnings, the planner generated 37 valve turning actions: 23 were successful; 10 failed because the valve was blocked; and 4 failed because the platform could not execute the action due to the high water current perturbations. Figure 10 shows the Girona 500 AUV with the payload required for valve turning and the diagram of all systems that were in charge of fulfilling the task.

VIII. CONCLUSIONS

The EU-funded PANDORA project represented a significant advancement in AUV autonomy, addressing the topic from various points: knowledge representation and cognitive aspects, autonomous planning and dynamic replanning, skill learning, control in presence of strong disturbances. The scientific results of the project were field validated in the context of three scenarios, using vehicles available at Heriot-Watt University and Univeristat de Girona.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Union Seventh Framework Programme FP7/2007-2013 - Challenge 2 - Cognitive Systems, Interaction, Robotics - under grant agreement No 288273 - PANDORA

REFERENCES

[1] 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.

[2] David M Lane, Francesco Maurelli, Petar Kormushev, Marc Carreras, Maria Fox, and Konstantinos Kyriakopoulos. Persistent autonomy: the challenges of the pandora project. *Proceedings of IFAC MCMC*, 2012.

[3] Francesco Maurelli, Zeyn Saigol, and David Lane. Cognitive knowledge representation under uncertainty for autonomous underwater vehicles. In *Proceedings of IEEE ICRA'14 IEEE Hong Kong, Workshop on Persistent Autonomy for Underwater Robotics*, 2014.

[4] F. Maurelli, Z. Saigol, D. Lane, M. Cashmore, B. Ridder, and booktitle = Proceedings of IEEE-MTS Oceans'14, St. John's, Canada year = 2014 D. Magazzeni, title = On AUV actions to correctly label world information.

[5] Michael Cashmore, Maria Fox, Derek Long, Daniele Magazzeni, Bram Ridder, Arnau Carrera, Narcís Palomeras, Natàlia Hurtós, and Marc Carreras. Rosplan: Planning in the robot operating system. In *Proceedings of the 23rd International Conference on Automated Planning and Scheduling (ICAPS'13)*, 2015.

[6] Michael Cashmore, Maria Fox, Derek Long, Daniele Magazzeni, Bram Ridder, and Francesco Maurelli. Dynamically Extending Planning Models using an Ontology. In *Proc. Workshop on Planning and Robotics (PLANROB) at ICAPS*, 2015.

[7] A. Carrera, SR Ahmadzadeh, A. Ajoudani, P. Kormushev, M. Carreras, and DG Caldwell. Towards autonomous robotic valve turning. *Cybernetics and Information Technologies*, 12(3), 2012.

[8] C. P. Bechlioulis, G. C. Karras, S. Heshmati-alamdari, and K. J. Kyriakopoulos. Trajectory tracking with prescribed performance for underactuated underwater vehicles with model uncertainties. *IEEE Transactions on Control Systems Technology*, in press.

[9] Natàlia Hurtós, Xavier Cufí, and Joaquim Salvi. A novel blending technique for two-dimensional forward-looking sonar mosaicing. In *IEEE/MTS OCEANS 2013 San Diego*, 2013.

[10] N. Palomeras, A. Carrera, N. Hurtós, G. C. Karras, C. P. Bechlioulis, M. Cashmore, D. Magazzeni, D. Long, M. Fox, K. J. Kyriakopoulos, P. Kormushev, J. Salvi, and M. Carreras. Toward persistent autonomous intervention in a subsea panel. *Autonomous Robots*, in press.