

Underwater Intervention Robotics: An Outline of the Italian National Project MARIS

By the MARIS Team

G. Casalino¹, M. Caccia², S. Caselli³, C. Melchiorri⁴, G. Antonelli⁵, A. Caiti⁶, G. Indiveri⁷
G. Cannata¹, E. Simetti¹, S. Torelli¹, A. Sperindè¹, F. Wanderlingh¹, G. Muscolo¹
M. Bibuli², G. Bruzzone², E. Zereik², A. Odetti², E. Spirandelli², A. Ranieri²
J. Aleotti³, D. Lodi Rizzini³, F. Oleari³, F. Kallasi³
G. Palli⁴, U. Scarcia⁴, L. Moriello⁴,
E. Cataldi⁵

Corresponding author: G. Casalino (giuseppe.casalino@unige.it)

Abstract - **The Italian national project MARIS (Marine Robotics for InterventionS) pursues the strategic objective of studying, developing and integrating technologies and methodologies enabling the development of autonomous underwater robotic systems employable for intervention activities, which are becoming progressively more typical for the underwater offshore industry, for search-and-rescue operations, and for underwater scientific missions. Within such an ambitious objective, the project consortium also intends to demonstrate the achievable operational capabilities at a proof-of-concept level, by integrating the results with prototype experimental systems.**

Keywords - **underwater manipulation; cooperative underwater manipulation systems; autonomous underwater vehicles; task priority based control**

I INTRODUCTION

Research activities in autonomous underwater robotics have been so far mainly focused on autonomous underwater vehicles (AUVs) performing exploration and observation missions,

1 - ISME, Genova node at DIBRIS, University of Genova, Via Opera Pia 13, 16145 Italy

2 - CNR-ISSIA, Genova node, Bari, Italy

3 - DII, University of Parma, Italy

4 - ISME, Bologna node at DIIIEI, University of Bologna, Italy

5 - ISME, Cassino node at DIE, University of Cassino and Southern Lazio, Cassino, Italy

6 - ISME, Pisa node at DII, University of Pisa, Italy

7 - ISME, Lecce node at DIE, University of Salento, Lecce, Italy

with important applications in the fields of oceanography, marine geology, archaeology, environmental sciences, as well as security applications.

Within these fields, the need for autonomy has naturally emerged for several different reasons:

- need for avoiding the use of umbilical cables because of the maneuvering problems they generally pose;
- need for achieving operational efficiency at depth, without resorting to highly risky and costly manned vehicles;
- need for faster and adaptable means for environmental monitoring and security applications, including the possibility of employing multiple cooperating AUVs.

In order to employ autonomous vehicles in such fields, research has been geared towards some fundamental aspects, for instance localization, communication and the use of vision technologies, which are much more difficult to tackle underwater than within ground or space applications. Nowadays, different successful demonstrations within these fields have been carried out, and some operational systems are already employed.

Another important field in autonomous underwater robotics is the so-called Autonomous Underwater Intervention field, where operations like grasping, manipulation and transportation activities, as well as assembly and disassembly, are the main issues to be faced, on top of the aforementioned basic issues.

However, research on Autonomous Underwater Intervention has not been as intense as that of the vehicular field, remaining for a number of years confined within only a few, if significant, research projects launched in the 1990s. Certainly, this occurred because

executing, even semi-autonomous, underwater interventions are not easy, but also because the needs for autonomous interventions were previously not as evident. As a matter of fact, the dramatic accident that recently occurred in the Gulf of Mexico and the previous airline crash in the Atlantic Ocean have demonstrated the importance of the eventual availability of smart underwater intervention robots.

Other fields where underwater intervention tasks carried out by autonomous robots are envisaged include:

- the high-depth offshore industry, in particular plant/infrastructure maintenance and inspection, and possibly fabrication;
- underwater rescue missions, which can be expressed by both civil and military organizations;
- the marine-science community (viz. oceanography, marine biology, geology, archeology, etc.), which might sensibly benefit from the availability of robotized intervention means in their activities, e.g. excavation, coring, handling scientific instruments, sample collection;
- the professional diving community, for similar reasons to those currently proposed for future grounded factories. Symbiotic undersea human-robot cooperation might even be envisaged.

Accordingly, research in autonomous underwater intervention robotics is now receiving renewed interest. This is also favored by the availability of several results and technologies that have been developed for underwater monitoring-only missions, in the fundamental topics mentioned earlier.

Within the context outlined above, the Italian National project MARIS originated as the joint initiative of five Departments of the University members of the National Inter-University Research Centre ISME ("Integrated Systems for the Marine Environment") and the unit of Genoa of ISSIA (Institute for Studies on Intelligent Systems for Automation) of CNR (Italian National Council of Research).

The two institutions qualify for their long-term research experience in the application of ICT to the marine environment, with particular interests in underwater robotics. The project consortium also includes an additional University Department (Parma), which has been selected for the highly qualified contributions it can provide to the MARIS objectives.

In this framework, the MARIS consortium has established the strategic objective of further developing and integrating the main technologies and methodologies enabling the realization of autonomous systems for underwater manipulation and transportation activities. Furthermore, the consortium also intends to provide proof-of-concept demonstrations of the achievable abilities, by integrating the results within prototypes of autonomous floating vehicle-manipulator systems realized within the consortium itself.

II PROJECT GENERAL OBJECTIVES

In regard to the aforementioned strategic goal, some advances regarding the main technologies and methodologies to be integrated must be achieved; in particular, the following topics:

- Reliable guidance and control of the floating basis (during long-range motions) on a multi-sensory basis, via the integration of inertial sensors, Doppler velocity loggers (DVL), external acoustic supports to localization (USBL or LBL), and

possibly via real-time SLAM techniques based on sea-floor observations.

- Stereo-vision techniques and systems, devoted to object recognition, also including their position and pose estimation.
- Reliable grasp, manipulation and transportation of objects by part of the manipulators operating on a floating basis; automatic reasoning for gripper pre-shaping, integrated with appropriate reactive control techniques (based on visual, force/torque and possibly tactile sensing) coordinating the whole system motions (i.e. vehicle, arm, gripper).
- Coordination and control methods for large object grasp and transportation by part of the cooperative floating manipulator systems, based on information exchange among the agents that must be compatible with the low bandwidth of the acoustic communication channels.
- High-level mission planning techniques, including the automatic decomposition and distribution of cooperative tasks among the agents.
- Underwater communication techniques among the agents (acoustic or optical) to the aims expressed earlier.

One of the main goals of the development of such technologies and methodologies is to have each individual floating manipulator exhibit dexterity, agility, and internal coordination capabilities. In addition, such activities have the additional ambitious goal of enabling adequate cooperative capabilities between the individual systems.

Finally, a further objective of the MARIS project is the design and realization of prototype systems, allowing experimental demonstrations of the integration of the results from the

previous objectives.

III PROJECT OBJECTIVES WITHIN THE STATE-OF-THE-ART

The place occupied by the MARIS project within the current international state-of-the-art is now described, with reference to the various topics composing the project itself.

1. *Control of individual floating manipulators*

The first experiences in equipping AUVs with robotic arms date back to the 1990s, when the AUV's ODIN (at University of Hawaii) and OTTER (at MBARI) were equipped with arms with very few degrees of freedom (d.o.f.); later, the AUV VORTEX (at IFREMER) was endowed with a seven d.o.f. arm.

Then, at the end of the 1990s, the EU-funded AMADEUS project (Lane et al., 1997) represented a step forward in the coordination of robotic structures, culminating in the realization of an underwater dual-arm work-cell exhibiting cooperative capabilities in manipulating objects grasped by both arms (Casalino et al., 2001). Despite operating from a fixed base, the system de-facto allowed the development of one of the first coordination techniques, whose effectiveness was at that time demonstrated by different underwater experiments.

Slightly later, within the SAUVIM project at University of Hawaii (Yuh et al., 1998; Yuh & Choi, 1999; Marani & Yuh, 2014) a large AUV was endowed with one of the same types of arms of the AMADEUS work-cell, which employed vehicle-arm coordinated control techniques partially derived from those of the AMADEUS system.

In both AMADEUS and SAUVIM, the introduction of a clear separation between the roles of the higher levels of kinematic control was successful: it was devoted to motion

coordination, and the dynamic control lower ones, simply used for tracking the real-time velocity references provided by the higher kinematic levels.

Such a hierarchical decomposition later greatly facilitated the development of more advanced coordination techniques for floating manipulators. Most subsequent research has in fact adopted this decomposition. Consequently, research efforts have been divided into two parts. On the one hand, they have been directed toward the development of different specific dynamic control techniques characterized by various levels of sophistication (see for instance Antonelli, 2006). On the other hand, they have addressed further developments of kinematic-based control techniques, devoted to vehicle-arm motion coordination problems.

With reference to the last point, overcoming some a-posteriori noted limitations exhibited by the techniques derived from those of the AMADEUS project (i.e. the fact that, since operating by switching among different coordination schemes, they actually tended to proliferate) was aided by the introduction of the so-called task-priority based kinematic control techniques (Antonelli & Chiaverini, 1998a,b). Such task-based techniques were later organized into a modular and computationally distributable form (Casalino & Turetta, 2003; Casalino et al., 2005; Casalino et al., 2009). Finally, they also now include the management of tasks of inequality type (typically imposed by system safety requirements, operability ranges, obstacle avoidance needs, etc.) to be generally achieved with higher priorities, even with respect to the ultimate mission purposes (Casalino et al., 2010; Casalino et al., 2012a, Casalino et al. 2012b, Simetti & Casalino, 2015).

The resulting algorithmic framework for such a kinematic layer is quite simple, computationally efficient, and capable of automatically adapting to changes in the task

priorities, possibly dictated by changing control objective sets and/or operational contexts, without requiring any structural change (Simetti & Casalino, 2016). Moreover, the possibility of assigning different priorities to the mobility of the internally composing mechanical parts (i.e. the vehicle, the arm and the hand) is also allowed.

As a matter of fact, the earlier version of such an algorithmic framework constituted the supporting kinematic control layer that allowed the agile control and coordination of the floating manipulator TRIDENT (Simetti et al. 2013, Simetti et al. 2014) developed within the recently closed homonymous EU-funded project (Sanz et al. 2012).

Currently, activities devoted to the extension of such a framework, also toward the ease management of interaction control aspects at gripper level, are currently within the research scopes of MARIS. Such activities are also within the scope of other two EU-funded projects, namely the DEXROV project (Gancet et al. 2015) and the very recently launched ROBUST project. These projects are focused on the employment of autonomous intervention robots within underwater oil and gas and sea-floor mining applications, respectively.

The employment of the aforementioned task-priority-based kinematic control techniques requires integration with everything necessary for the extraction, from the sensory apparatuses endowing each system, of the information needed for their closed-loop functioning. In particular, we recall the following fundamental blocks: object pose estimation upon recognition via stereo vision; real-time grasp planning; integrated force-torque-tactile (and also vision) sensing; guidance and control for long-range motions; localization via integrated inertial and acoustic sensing; ambient reconstruction via SLAM-based techniques; acoustic communication for long distances, and possibly optical ones for

short-range ones.

The state-of-the-art of such aspects, once clustered into three substantially homogeneous topics in relation to the MARIS scopes, can be summarized as follows.

2. Object recognition, pose estimation, and grasp planning

Growing interest in the use of artificial vision, both stereo and mono (Nicosevici et al., 2009; Campos et al., 2011) has been recently seen in the underwater field.

For the recognition of underwater objects, monocular vision was formerly suggested, based on the detection of invariant features with respect to the observation point (Isaac & Srivastava, 2010). However, stereo vision technology, since enabling the generation of three-dimensional object models from disparity images (Brandou et al., 2007; Campos et al., 2011) has been considered with greater interest. In this case, perceptual data are interpolated, filtered, and then split into suitable topological representations, facilitating object recognition and its pose estimation, even under conditions of partial visibility, facilitating the planning of the grasp of objects of interest.

For grasp planning, the classical metric of robustness (Ferrari & Canny, 1992; Borst et al. 2003) can be efficiently integrated within a stereo vision system, with criteria capable of also indicating, whenever possible, the preferential grasp directions and hand pre-shaping (Aleotti & Caselli, 2010).

Moreover, the availability of a 3D model of an object can be used to take into account, when planning its grasp, possible constraints directly related to the grasping task itself. These constraints include, for instance, the presence of surface portions to be kept free, other than the preferential directions of approach (Aleotti et al. 2012), especially when conditioned by the possible presence of obstacles in the scene (Berenson et al., 2007). The

development of the hardware architecture and real-time software for underwater stereo vision systems allowing 3D model reconstructions, pose estimations and grasp planning, for the most typical types (templates) of objects encountered within underwater environments and/or plant infrastructure, represents a fundamental part of the activities within MARIS (Oleari et al., 2015; Kallasi et al., 2015; Lodi-Rizzini et al., 2015).

3. Integrated force-torque and tactile sensing

The problem of providing robots with the sense of “touch“ is fundamental in order to have the possibility of implementing fine and dexterous manipulation tasks, since feedbacks from tactile sensing enable the detection and the safe control of the interaction of the robot with objects. Tactile sensing technology has, until now, received significant attention only for terrestrial applications; studies about tactile sensors for underwater applications have been quite limited (Cannata & Bruno, 1999; Tan et al., 2008; Kampman & Kirchner, 2012). One of the main reasons for this is that, even at shallow depths, water pressure can significantly affect the performance of systems based on the transduction devices commonly used at sea level, where the effects of atmospheric pressure can generally be neglected.

Part of the consortium has been involved in the recently closed EU-funded project ROBOSKIN, devoted to the study and development of distributed tactile sensors for ground and/or space applications. Within MARIS, investigations regarding the possible extension to the underwater field of the technologies developed within ROBOSKIN (based on capacitive transduction) are currently very active (Muscolo & Cannata, 2015). The goal is to develop distributed tactile sensors to be integrated within the underwater grippers, and

possibly even distributable on the arm-link surfaces, as a whole-arm artificial underwater skin.

4. *Guidance, navigation, control, ambient reconstruction, localization and communication*

GNC architectures for vehicles' long-range motions allow the functional distinction among the subsystems relevant to dynamics (Control), its kinematics (Guidance) and its motion estimation (Navigation). Although the GNC literature for individual non-manipulative vehicles is actually extensive (e.g. Fossen, 1994; Roberts & Sutton, 2006), this is not the case for teams of cooperating floating manipulators (especially when strictly cooperating, as briefly outlined in the next point 5).

For this reason, relevant activities within MARIS currently take place along four main directions, each starting from the current state-of-the-art, as hereafter indicated.

Modeling and identification techniques originally developed for single vehicles (Caccia et al., 2000) have to be extended to the case of the floating manipulators, for low-level dynamic control purposes.

The problem of environmental reconstruction is instead tackled via SLAM techniques exploiting computer vision and sensor fusion (Ferreira et al. 2012).

As for localization, traditional techniques are to be used, based on distance and bearing measurements, in the case of costly employment of USBL systems, or based on acoustically obtained distance information from several points (Caiti et al. 2005), or even from a single point (e.g. Arrichiello et al. 2011; Parlangeli et al., 2015; Indiveri et al., 2016). Since these techniques, especially within underwater cooperative contexts, are strongly dependent on the adopted acoustic communication system and related architecture and protocols (Caiti et al. 2012), their performance optimization represents an additional

important intensive activity within the MARIS scope.

5. *Coordination of cooperating floating manipulators.*

As already mentioned in the introduction, the MARIS scope includes the investigation of enabling a couple of underwater floating manipulators to cooperate in manipulating and/or transporting a common object (Fig. 1).

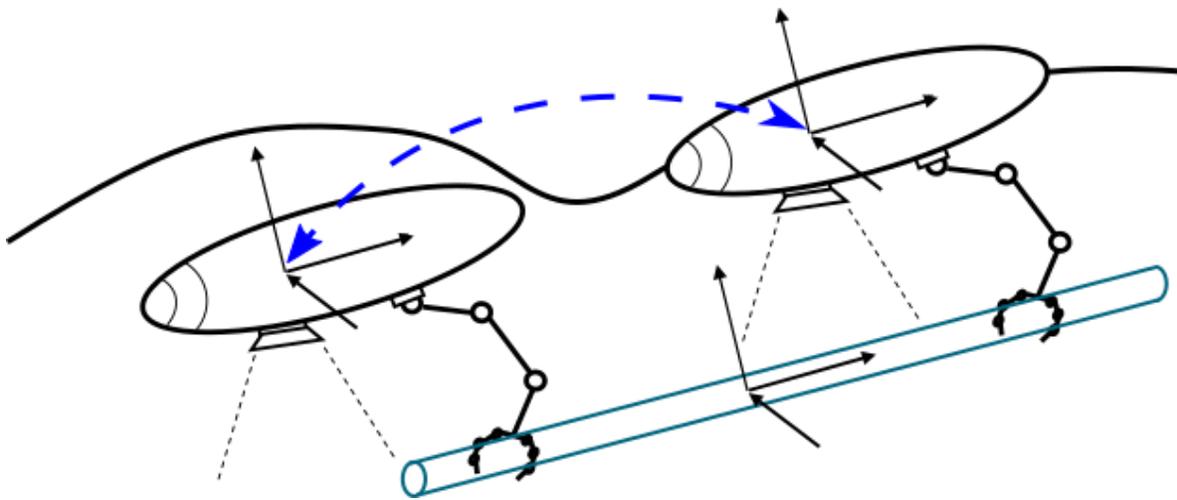


Fig. 1: Sketch of two strictly cooperating underwater floating manipulators

In this context, cooperation also exists during the approach phase to the object, when, still in the absence of manipulative operations, the agents must coordinate themselves for grasping the common object. Of course, the required level of cooperation becomes much more demanding when the agents must work close to each other for transporting the object. We can therefore speak about weak coordination at the beginning of the approach phase, and of strict coordination at its end. Weak coordination is also necessary when an agent

must transfer a grasped object to the other, or when they must coordinately navigate to reach a common operational zone.

As can be easily argued, both cases require each agent to exhibit suitable agile and dexterous operative capabilities, further motivating the emphasis that was assigned to all these aspects for individual agents.

As regards weak coordination, we can note how it is closely related to what is required of AUV teams used for non-manipulative tasks (i.e. distributed patrolling, sampling, exploration, etc.), for which many studies now exist (see for instance Chapter 9 of Antonelli, 2006). In this area, the most recent and effective proposed techniques (Antonelli & Chiaverini, 2003a,b; Antonelli 2006) actually gain their efficiency from being based on priority task concepts and relevant algorithmic structures, structurally similar to those employed for the redundancy resolution within generic kinematic chains. This fact demonstrates how internal coordination within single floating manipulators, as well as external coordination, whenever accomplishing weak cooperative tasks, can be included with a common conceptual and algorithmic framework.

It is also remarkable that strict cooperative tasks can be traced back to the same conceptual and algorithmic task priority framework developed for individual floating manipulators (this was formerly noted within Simetti et al., 2009, for strict coordination between grounded mobile manipulators).

In both weak and strict cooperative tasks, a necessary requisite is that the agents can real-time share at least the minimal amount of data that is deemed strictly sufficient for guaranteeing a conflict-less execution of cooperative tasks.

The extension to the decentralized cooperative case of the overall task priority-based

algorithmic framework is currently at an encouraging stage of development, where the main research efforts are oriented to reducing the required amount of information exchange for conflict-less cooperation.

IV PROJECT CURRENT ACHIEVEMENTS

Within the project, methodological research activities are running in parallel with the integration ones. Fig. 2 shows a CAD representation of one of the AUVs, named R2, which has been prepared and made available by the partner CNR-ISSIA-Genova node, once integrated with one of the seven-d.o.f. arms instead realized within the ISME-Genova node partnership.

This represents the first of two systems that were planned, while the second is currently under finalization. The R2 vehicle weighs about 300 kg in air and is neutral in water. The arm is instead 30 kg in air and about 10 kg in water, thus it requires the addition of a small buoyancy force to the whole system for achieving its global neutrality.

Fig. 3 shows the physical assembly of the first prototype, endowed with the stereo-vision sub-system designed and realized by the partnership of University of Parma (Kallasi et al, 2015), as well as the dexterous gripper designed and realized by the partnership of the University of Bologna (Palli et al., 2014, 2015).

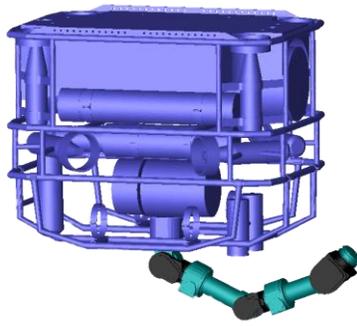


Fig. 2: CAD representation of the first prototype of underwater floating manipulator



Fig. 3: Assembly of the first prototype of underwater floating manipulator

The integration activities that led to the first system prototype required much effort in terms of time and manpower, with a considerable part of effort devoted to the software architecture design and implementation. The interfacing of the various resulting software processes was performed mainly via ROS, apart from the hard-real time parts internal to the arm and vehicle control levels. Further effort was devoted to the different separate

functionality and validation tests, which were performed on each composing subsystem. This work was carried out before proceeding toward their integration and before performing the validation trials in a pool environment, where good performance during individual grasping tasks was exhibited by the prototype.

In particular, Figures 4 and 5 show the prototype system while autonomously grasping a pipe during the mentioned pool trials; a video clip can be seen at: <https://youtu.be/b3jZUoeFTo> .

At the present time, two activities are proceeding in parallel: the ongoing realization of the second prototype system, and the extension to the decentralized cooperative case of the same task: priority-based coordination and control techniques internally employed by individual agents, which were validated by different simulation experiments.



Fig. 4: Underwater view of single-agent experimental grasp operation



Fig. 5: View from surface of single-agent grasp operations

In particular, in this regard, attention has been mainly focused on the case of strict cooperation, to be exhibited by two agents when manipulating and/or transporting a shared grasped object.

In this case, in order to limit the possible induction of uncontrolled object stresses, the vehicle and joint velocity *commands* (i.e. reference signals) provided to each agent, computed as a result of their individual task-priority based kinematic control layers, must at each time instant translate themselves into the *same* object Cartesian velocity. Such a velocity reference must also converge toward the one that, at each time instant, is required for having the shared object asymptotically reach its final position, as ultimately specified by the mission purposes.

At least in principle, satisfying the above requirements would certainly be possible if the agents were allowed to exchange rapidly all the information needed for de-facto transforming the decentralized problem into a centralized one, to be at each time instant

solved by each agent based on such a complete and shared information set.

Since, in practice, such an ideal situation cannot be considered due to communication bandwidth and latency limitations, for the time being, the best-effort decentralized coordination policy (Simetti et al., 2015; Casalino et al., 2015) described below was proposed, investigated, and tested on a simulative basis.

a) At each time instant, the task-priority based control layer of each agent produces its own arm joint and vehicle velocity references, as if each were actually the sole transporting agent and, consequently, by having each solely take into account its own internal list of prioritized tasks. Following this procedure, each agent can compute the associated resulting Cartesian velocity for the object frame. However, these resulting object velocities may differ from each other; for example, when one of the agents is engaged in achieving its own higher priority objectives (i.e. safety conditions, operability ranges, obstacle avoidance, etc.). If the corresponding joint velocity commands were applied, unwanted stresses on the object would certainly be created. Thus, with the aim of avoiding such possible occurrences, the following additional step is executed *before* applying any reference command by part of each agent.

b) The produced individual object-frame reference velocities are exchanged at each sampling period, and their common weighted mean separately evaluated by each agent.

c) Such weighted means are used as a new common reference velocity for the object frame, which is now the *highest priority task* of each agent. This new task hierarchy is now solved by both agents, and the resulting vehicle and arm-joint velocity references are now actually applied.

This best-effort coordination policy de-facto arises as a reasonable way for allowing both

systems to carry out their own safety and operational-enabling tasks separately, while also avoiding uncontrolled object stresses. Furthermore, it must be noted how the coordination scheme requires a very small amount of information exchange at each sampling period, since only six real numbers actually have to be exchanged, even if in a full duplex way.

For the time being, extensive simulation campaigns have provided quite encouraging results, also taking into account latencies in the information exchanges, to be hopefully later confirmed via experimental trials, once the second system prototype is available.

However, there is still the risk that currently available state-of-the-art acoustic modem technology might be not able to support the required full-duplex communication with reasonable (for control purposes) rates and latencies. For this reason, investigation activities regarding the alternative use of high bandwidth optical modems (Cossu et. al, 2013) are currently running in parallel with the others. Note how, in this situation, the physical distance between the cooperating agents, which for the most part of the foreseeable cooperative manipulation and transportation tasks should not exceed a few meters, is very reasonable for the adoption of optical communications. Finally, note how, in this case, an acoustic link could still be maintained, for the possible exchange of “tactical information” of useful semantic significance (i.e. mainly of symbolic type) for the overall coordinated mission, to be possibly exchanged when the agents are still far from each other at the mission start. This represents an additional aspect that will be addressed by the project in the near future.

For the time being, Fig. 6 shows the different stages of a simulated transportation mission, and final precise positioning, of a pipe carried out by two cooperating systems (the pipe icon appearing in the bottom-right represents the final goal position to be reached; while

rendering was obtained by employing the open-source package UWSim: Prats et al., 2012). Further details and discussion of the simulations can be found elsewhere (Manerikar et al., 2015a,b; Simetti et al., 2015).

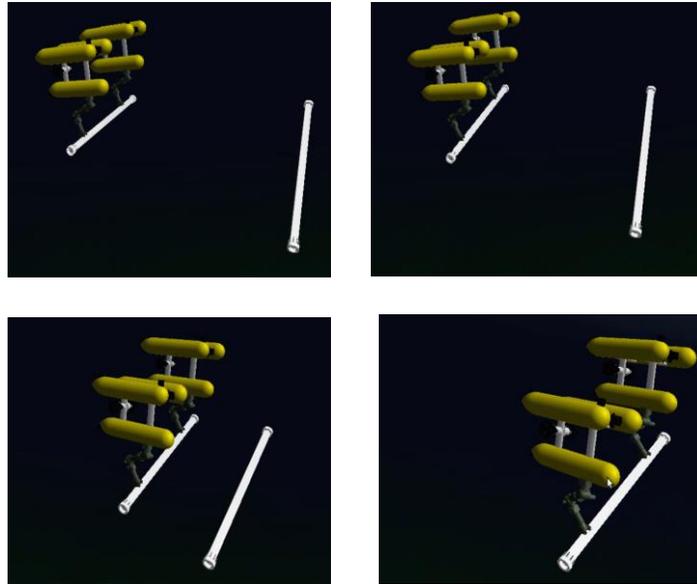


Fig. 6: Initial, intermediate, and final phases of a simulated cooperative transportation task

Fig. 7 shows the achievable dexterity in executing tasks of a more preponderant manipulative nature than transportation.

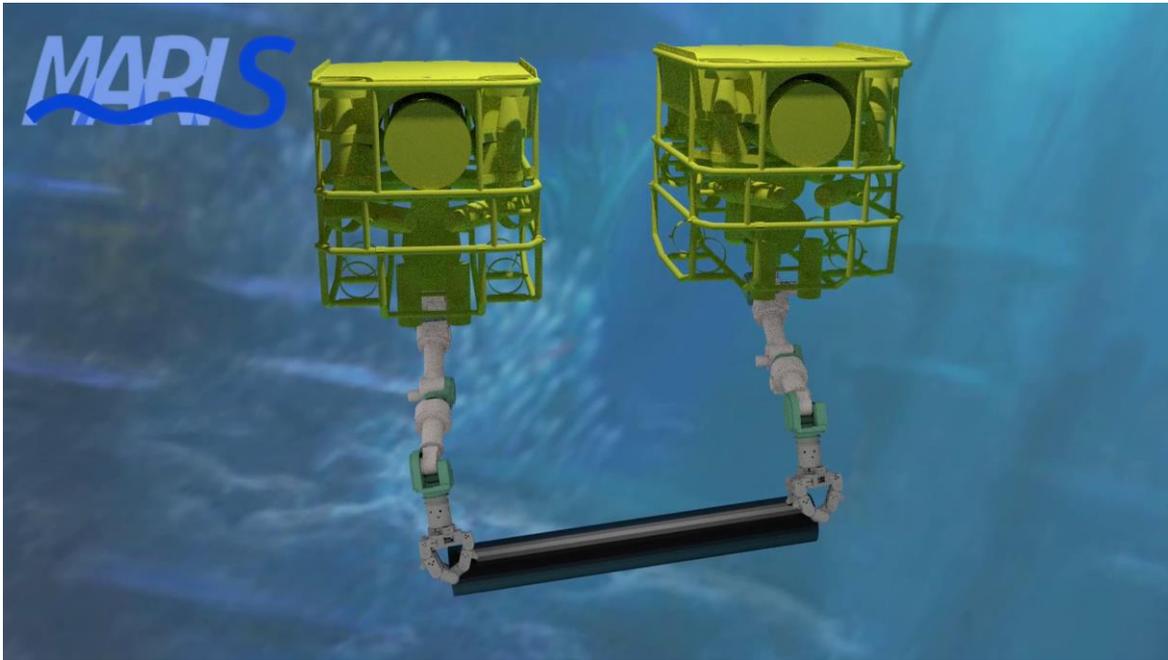


Fig. 7: Simulated cooperative object-manipulation task

Some video clips showing such simulated cooperative manipulation and transportation missions, also including the associated approach and final release phases, are available at <https://youtu.be/9WRRUotcjmM>.

V CONCLUSIONS

The Italian national project MARIS is devoted to the integrated study and development of all methodologies and technologies that enable the realization of agile autonomous robotized systems for (cooperative) intervention missions within underwater operative scenarios. Such systems have the opportunity to become, in the future, progressively typical for many underwater, scarcely structured, applications of different types (e.g. the underwater off-shore industry, sea floor mining, search-and-rescue operations, scientific

missions, etc.). This possibility is underlined by the existence of different EU-funded projects on the same subject, even if they are more heavily focused on specific applications, in which many partners of the MARIS consortium are involved.

This study has presented the general objectives of the MARIS project, each detailed within its relevant state-of-the-art. Furthermore, our current achievements have been presented. Experimental trials on single agent free-floating manipulation have been shown, while the cooperative manipulation and transportation tasks were validated, at this point, through extensive simulation campaigns. The consortium is now focusing its activities on the finalization of the second prototype and the subsequent execution of cooperative transportation trials.

ACKNOWLEDGMENTS

This work was supported by the MIUR (Italian Ministry of Education, University and Research) within the MARIS project, protocol number 2010FBLHRJ.

LIST OF FIGURE CAPTIONS

Fig. 1: Sketch of two strictly cooperating underwater floating manipulators

Fig. 2: CAD representation of the first prototype of underwater floating manipulator

Fig. 3: Assembly of the first prototype of underwater floating manipulator

Fig. 4: Underwater view of single-agent experimental grasp operation

Fig. 5: View from surface of single-agent grasp operations

Fig. 6: Initial, intermediate, and final phases of a simulated cooperative transportation task

Fig. 7: Simulated cooperative object-manipulation task

