SCIENCE AND TECHNOLOGY ORGANIZATION

CENTRE FOR MARITIME RESEARCH AND EXPERIMENTATION

Reprint Series **CMRE-PR-2019-012**

Autonomous underwater surveillance networks: a task allocation framework to manage cooperation

Gabriele Ferri, Jeffrey Bates, Pietro Stinco, Alessandra Tesei, Kevin LePage

May 2019

Originally presented at:

2018 OCEANS - MTS/IEEE Kobe Techno-Oceans (OTO),
doi: 10.1109/OCEANSKOBE.2018.8558813

About CMRE

The Centre for Maritime Research and Experimentation (CMRE) is a world-class NATO scientific research and experimentation facility located in La Spezia, Italy.

The CMRE was established by the North Atlantic Council on 1 July 2012 as part of the NATO Science & Technology Organization. The CMRE and its predecessors have served NATO for over 50 years as the SACLANT Anti-Submarine Warfare Centre, SACLANT Undersea Research Centre, NATO Undersea Research Centre (NURC) and now as part of the Science & Technology Organization.

CMRE conducts state-of-the-art scientific research and experimentation ranging from concept development to prototype demonstration in an operational environment and has produced leaders in ocean science, modelling and simulation, acoustics and other disciplines, as well as producing critical results and understanding that have been built into the operational concepts of NATO and the nations.

CMRE conducts hands-on scientific and engineering research for the direct benefit of its NATO Customers. It operates two research vessels that enable science and technology solutions to be explored and exploited at sea. The largest of these vessels, the NRV Alliance, is a global class vessel that is acoustically extremely quiet.

CMRE is a leading example of enabling nations to work more effectively and efficiently together by prioritizing national needs, focusing on research and technology challenges, both in and out of the maritime environment, through the collective Power of its world-class scientists, engineers, and specialized laboratories in collaboration with the many partners in and out of the scientific domain.

Autonomous underwater surveillance networks: a task allocation framework to manage cooperation

Gabriele Ferri, Jeffrey Bates, Pietro Stinco, Alessandra Tesei, Kevin LePage

Abstract— **The design of efficient task allocation schemes is essential to manage autonomous underwater robotic surveillance networks. The network has to assign the most suited robots to the tasks which compose the mission, in spite of the unreliable and intermittent underwater communications.**

In this paper, we describe a market-based policy for this. It works in a completely distributed way and, through periodic auctions, sequentially allocates the robots to the tasks. There is no central auctioneer and each robot can resolve the current auction. These features increase the robustness to poor communications.

The proposed scheme can manage two kinds of tasks, continuous tasks which never terminate and are always available in the task pool and sporadic ones, created upon the occurrence of some events and requiring a rapid response of the network. An Anti-Submarine Warfare scenario is considered to validate the task allocation scheme. Surveillance of areas of interest are the addressed continuous tasks, while sporadic tasks are created whenever a track produced by the tracker on-board a robot is confirmed as likely related to a target. In this case the team has to rapidly react to the event and to increase the tracking performance. Results from non-trivial Matlab simulations are reported and demonstrate the effectiveness of the allocation scheme in degraded communications conditions.

Index Terms— **Task Allocation, Autonomous Robotic Networks, Underwater Surveillance, Autonomous Underwater Vehicle, Anti-Submarine Warfare.**

I. INTRODUCTION

Underwater surveillance is the requirement to detect, localise and classify targets (e.g. divers, manned or unmanned vehicles) using sensors of different nature [1].

Recent advances in marine robotics have made small and low-cost AUVs a reality [1]. These novel assets can guarantee the desired monitoring at a fraction of the cost of traditional approaches, which consist in statically deployed sensors, or in expensive and time-consuming ship (or submarine) based operations. Today's AUVs are characterised by lower communication, sensing and computational capabilities with respect to traditional assets, but they can be deployed to form robotics underwater surveillance networks [1], [2]. These kinds of networks can exploit the synergies that mobile robots can offer to achieve high mission performance.

The NATO STO-Centre for Maritime Research and Experimentation (CMRE) has been pursuing this approach [3]–[5] developing and testing at sea a hybrid Anti-Submarine Warfare (ASW) network composed of static nodes and mobile nodes. Static nodes constitute the backbone of ad hoc communications infrastructure. Mobile nodes are surface vehicles such as WaveGlider [6] Autonomous Surface Vehicles (ASVs) which operate as mobile communications gateways and the Ocean Explorer Class Version C (OEX-C) Autonomous Underwater Vehicles (AUVs). OEX-C AUVs tow a linear hydrophone array [7] and act as the receivers in a multi-static active sonar system. One or more active sources (transmitters) transmit a sonar signal (ping), which once reflected off some objects can be recorded by the receivers (see Fig. 1). The robot processes these data to produce contacts (range and bearing) and finally tracks [1]. Mobile robots can use their mobility to extend the operational area and to adapt mission objectives to the changing environmental and tactical conditions. This can increase network adaptability, mission performance and robustness.

Communications become crucial to manage the network operations and to improve their effectiveness. Data fusion [8] between contacts and tracks produced by different robots is a way to improve target detection and tracking. The robots use also communication to find an agreement on which tasks of the ASW mission to execute.

This problem, known as Multi Robot Task Allocation (MRTA), consists in finding an agreement in the robotic team on how to assign a certain task to one or to some subset of the robots to achieve the overall missions goals in an efficient manner [9], [10]. In its simplest formulation (each task can be assigned to exactly one robot and one robot cannot be assigned to more than one task), the MRTA becomes an instance of the Optimal Assignment Problem (OAP), which can be solved in a centralised way by using the Hungarian method [9] or in a distributed fashion by using the auction algorithm [11]. In general, however, MRTA is not a one-time assignment and becomes a dynamic decision problem, since utilities may vary or tasks may be terminated or created. The static assignment can no longer be considered applicable and iterative procedures must instead be sought [9], [10], [12]. Changing the relations between utilities, the type of the group utility function and relaxing the OAP constraints, make the MRTA more complex leading also to NP-hard problems [10]. Specific reviews on the topic can be found in [9], [10], [13], [14].

In this paper, we investigate the MRTA problem for a robotic network similar to the described CMRE system. We consider the network deployed in a littoral environment for an ASW surveillance scenario. The scenario is characterised by the limitations of the underwater acoustic communication channel. Communications are unreliable, with low range and low bandwidth [15]. This makes MRTA underwater

Gabriele Ferri, Jeffrey Bates, Pietro Stinco, Alessandra Tesei and Kevin LePage are with the Research Department, NATO Centre for Maritime Research and Experimentation, Viale San Bartolomeo, 400, 19126 La Spezia (SP), Italy Gabriele.Ferri@cmre.nato.int

challenging and prevents to adopt centralised policies [13] (which can in theory compute the optimal allocation).

Several approaches which decentralise the existing classical approaches have been proposed in the robotics community [9], [10], [13]. Usually, they produce sub-optimal solutions, but they can handle poor communications and require reduced computational power. Among the solutions available in the literature and of interest to our scenario, behavioural approaches can provide ad hoc solutions for MRTA, such as the ALLIANCE architecture solution proposed in [16]. It represents one of the earliest demonstrations of iterated assignment architectures to MRTA [9]. These approaches are easy to implement in a team of robots, but generally do not have a detailed evaluation of the optimality of the solutions found. Probabilistic frameworks, such as Markov Decision Process (MDP) or Partially Observable MDP as those presented in [10], [13] can be used to provide an optimal control strategy in tightly coupled domains. However, the main limitation to their use is that they quickly become intractable even for small problems. Approximation techniques need to be further investigated for an effective use of these frameworks in real scenarios [10], [13].

One other family of algorithms is represented by auction/market-based approaches. These methods are distributed in nature and are flexible. They can be adapted to changing conditions also by incorporating some features of centralised approaches [17]. The underlying philosophy of *market-based* methods is that of distributing common resources among the team members taking inspiration from human market economies where individual pursuit of profit leads to the redistribution of resources and to an efficient production of output [18]. In this virtual economy, tasks are traded as commodities and virtual money acts as currency, while robots compete to be assigned to a task by participating in auctions. When the system is correctly designed (i.e. costs, revenues and auctions mechanisms), each robot acts to maximize its own profit while, at the same time, moving towards an increase in the group efficiency [10], [12], [17], [18].

Most of the research on multi-robot task allocation has been confined so far to terrestrial or aerial robotics [9]. The engineering difficulties in deploying teams of vehicles at sea, make examples of experiments at sea in which multi-robot coordination is demonstrated rare [19], [20]. Due to these difficulties, task allocation schemes in underwater multi-robot systems are usually tested in simulation environments [2], [21]–[23], in general without the required attention to the communications perspective.

In this paper, we extend the PADA (Periodic Auctions Distributed Algorithm) task allocation framework [24] for an underwater robotic network. PADA works by using periodic negotiations among neighbouring robots and requires only local communication. Particular attention has been paid to make the auction scheme robust to intermittent communications. This work adds the possibility to manage two kinds of tasks, the continuous ones which never terminate and are always available in the task pool (e.g. area surveillance) and

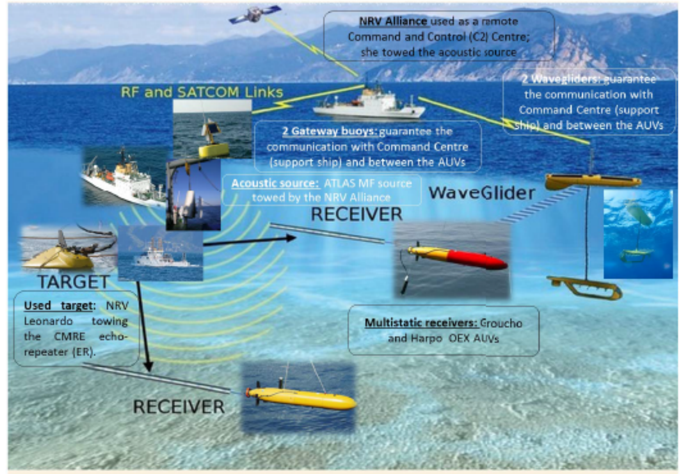

Fig. 1. The CMRE cooperative ASW multi-static network.

sporadic tasks, created upon the occurrence of some events and requiring a rapid response of the network (e.g. response to some confirmed cue of a target). A task creation chain mechanism is also introduced. This allows the creation of ancillary tasks to better exploit the synergies the robots can offer to increase the flexibility of the proposed framework. The allocation framework is now capable to manage the different tasks which compose a complex ASW surveillance mission.

II. Mission tasks

An ASW mission [1] can be decomposed in two kinds of tasks, continuous and sporadic. Continuous tasks never terminate while sporadic ones are created by some of the robots when some events occur and terminate once the triggering conditions are not true any more.

A. Continuous tasks: area surveillance tasks

An area surveillance task involves to guarantee an adequate surveillance of some areas to detect possible cues of targets. The AUVs are requested to survey some areas of interest periodically due to the dynamic nature of the problem (the possible intruders move). Each task t_i is associated with a geographical area A_i to survey and a fixed rectangular path which the robots cover when allocated to t_i. We allow that more than one robot can be assigned to the same task. This violates the assumption of the classical OAP and makes our problem more difficult. The task utility u_{ki}, for a robot k, r_k, to perform an area surveillance task i and used in the auction-based MRTA algorithm is:

$$u_k^i = w_k^i - c_k^i \tag{1}$$

w_k^i being the reward to accomplish the task as computed by r_k and c_k^i the cost associated to r_k to reach the path associated to t_i. w_k^i is the reward the robot would have to be assigned to the task. This quantity drives the robots to decide on which task to bid. The cost c_k^i is a scaled distance from the current position of the r_k to the closest point of the path associated to A_i and gives a measure of how difficult it is for the robot to reach the area.

w_k^i is defined as:

$$w_k^i = \alpha_k^i \frac{N_i}{N_{max}} I_k^i(t) \quad (2)$$

with α_k^i the current coverage index of A_i (a value of 0 means the area is fully surveyed) computed by r_k with its knowledge of the other robot actions, N_i is a measure of the dimensions of the area A_i and N_{max} a measure of the largest area in the tasks to allocate. $I_k^i(t)$ is the urgency index for t_i as computed by r_k. The coverage index for A_i is built based on the cumulative probability of detection on the patrolled area, computed over a moving, relatively narrow, temporal window. Large values of this quantity means that a target, if present, will be likely detected, leading to successful tracking and final classification. The urgency factor, $I_k^i(t)$, is introduced to quantify how urgently an area needs to be surveyed. $I_k^i(t)$ increases over time if α_k^i is larger than a certain pre-fixed threshold, meaning the area is not adequately covered. It decreases, instead, if the area is adequately surveyed by the robots. As said, these tasks never terminate. Details on how the coverage index and the urgency factor are computed are reported in [24].

B. Sporadic tasks: track prosecution at broadside and glint seek

A sporadic task is created upon the occurrence of some events. The task is then communicated to the network and the allocation process starts. Once the condition which triggered the task creation ceases to be valid, the task terminates and is removed from the task pool.

In our scenario, two types of sporadic tasks are considered, both related to the team reaction to a track that becomes confirmed. A track is confirmed whenever an onboard classifier considers it as likely related to target activity [25]. In this case, the robots modify their behaviours [1], [7] to improve the tracking/classification performance (i.e. they prosecute the track). Generally, no robot other than the task creator has information about the track. Thus the task creator sends continuously information about the track to the other robots as it evolves over time. Each robot propagate the received track into the future in case of missing updates from the task creator, or updates it with the received information. Even if complex behaviours have been used for track prosecution [7] in CMRE trials, in this paper we consider simpler behaviours for simplicity reasons. A *broadside* behaviour, in which the AUV controls its heading to keep the estimated position of the target (as retrieved from the track) at broadside (with the angle $\theta = 90°$ - see Fig. 2) (to increase the detection performance) of its sensing array. The related task is named *track prosecution at broadside*. A second task is produced as ancillary to the previous one. In this second task, named *glint seek*, the AUV moves at a speed greater than the cruise velocity to reach the glint condition (see Fig. 2). If the glint condition is reached, the AUV switches to navigate to keep the target at broadside. A *glint seek* task is produced once a robot loses an auction for the track prosecution at broadside and no more robots can be assigned to that task.

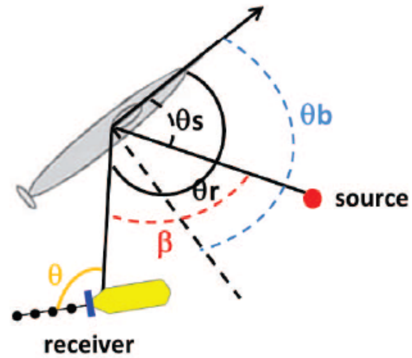

Fig. 2. Important angles for bistatic target strength computation.The glint conditions which increases considerably the target strength is the specular reflection of the acoustic wave. This happens when the bistatic aspect angle, θ_b, is equal to 90° or 270°. θ_b is the angle from targets heading to the bisector of the bistatic opening angle β. In the figure also the bearing θ is shown. Broadside condition occurs if $\theta = 90°$.

The utility of the two sporadic tasks for a robot r_k on the task t_i are defined as $u_{ki} = -c_{ki}$, with c_{ki} being the cost for a robot k to be allocated to task i. The cost is created as $\psi_{ki} + a\ d_{ki}$, where ψ_{ki} is a scaled heading angle error (from 0 to 1) between the current robot's heading and the objective one, d_{ki} is the scaled distance error (from 0 to 1) from the current robot position to the track (in the case of the broadside behaviour) or to the closest position on the glint line (in the glint seek behaviour) and a a weighting factor.

III. Problem Definition

We assume a set of n robots (denoted $r_1, \dots, r_n$) and a set T of m tasks, denoted $t_1, \dots, t_m$. Our aim is to determine distributed control laws to partition the robots into m groups of size l_k associated to each task, t_k, with $k = 1, \dots, m$. The types of tasks considered are the ones previously described. For area surveillance tasks, we make the assumption that each robot can only sense the area related to the surveillance task to which is currently assigned.

The proposed problem can be formulated as to find mn non-negative integers γ_{ij} (which must be either 0 or 1, 1 meaning that r_i is assigned to t_j), that minimise the following group objective function J:

$$J = \sum_{i=1}^{m} w_i + \beta \sum_{j=1}^{n} c_j \quad (3)$$

where w_i is the reward related to t_i computed considering the measurements of all the team robots, β is a weight factor and c_j is the cost of r_j relative to the task which is currently assigned to. A task allocation which minimises J lowers the rewards for the tasks w_i (task are adequately accomplished and there is less need to execute them), using the robots with the lowest costs c_j (the most suited to execute the tasks). The minimisation being subject to:

$$\begin{aligned} &\textstyle\sum_j^n \gamma_{ij} = 1, \quad 1 \le i \le m \\ &\textstyle\sum_i^m \gamma_{ij} \le l_j, \quad 1 \le j \le n \end{aligned} \quad (4)$$

Our problem is more complex than the OAP previously described for several reasons. The second constraint allows that up to l_j robots can be assigned to t_j. This defines the problem as a Multi-Robot Task (MR) problem [9], [10]. This is known in the multi-agent community as a coalition formation problem [26].

In the case of area surveillance tasks, the group members must survey the areas A_is (reducing w_is) to minimise the group utility in (3). The areas must be surveyed according to the required time constraints (dictated by the urgency factor $I(t)$ - see (2)). At the same time, the robots must reduce c_js, their cost on the currently assigned task. This means that they must reduce the switching from one area to another which is the major cause of the increase of c_js in continuous tasks, or negotiating the best allocation to the track to be prosecuted in the case of execution of a sporadic task. One robot decides to change the area to survey computing its reward and the cost to reach that area: a trade-off has to be reached between task switching and the coverage of the areas with highest rewards.

The management of sporadic tasks creates additional constraints, since these tasks are assumed to have the priority of execution over the continuous ones. A quick allocation is required for a prompt response of the network to the occurring events.

We underline that the robots during the MRTA use approximate information due to the limitations in communications.

IV. The Periodic Auctions Distributed Algorithm (PADA)

We propose a distributed market-based MRTA algorithm, the Periodic Auctions Distributed Algorithm (PADA), in which robots bid periodically for task assignment. Two slightly different auctioning systems are run in parallel, one for continuous and one for sporadic tasks. Each robot has a *primary* and a *secondary* task for which it bids and executes the actions related to the primary task. We make the assumption that each robot has knowledge of the surveillance tasks present in the system and of the maximum number of robots which can be allocated to each task. The sporadic tasks, instead, are created upon the occurrence of some events. A robot commences to execute a sporadic task and to bid for it as soon as receives information about it from a collaborator. The aim is to guarantee a quick response of the network.

PADA offers a distributed approximation to (3) and (4) to be implemented on underwater vehicles. There is no presence of a central auctioneer, typical in market-based policies. Having a central, even if local, decision maker may be problematic in limited communications environments. The robot selects the highest utility task to bid for and then commences bidding for that task. Sporadic tasks, created by itself or received from collaborator, have the priority of execution over the continuous ones. Periodically, each robot evaluates the bids received and a winner of the auction is awarded (the agent with the highest bid). The robots which are not the winners of the auction select another task to bid for and initiate other negotiations. The assignment process consists in a sequence of such auctions taking place locally among the robots which can communicate with each other. The auctions are periodic to cope with unreliable communications. When new information from some bidders is available, it is taken into account in the ongoing negotiation.

The task assignment is performed considering the status and knowledge of the world attainable at the assignment time (Instantaneous Assignment (IA) according to the taxonomy in [9], [10]). Any situation unknown at the negotiation time is not taken into account. However, PADA introduces means to allow renegotiations (task reallocations) for continuous tasks to reallocate dynamically to better achieve the group's objective. This is necessary to handle events which may occur, such as changes in the number of team members, creation of new tasks or with an evolving tactical scene (some tasks become more important than others). PADA also has some precautions to reduce the number of task switches (reallocations). A trade-off between surveying the most urgent areas and reducing the number of task switches is sought to improve the group performance.

The same algorithm runs on each robot. At the beginning of the mission, each robot selects the highest utility surveillance task. The selected task becomes both the primary and secondary task. The AUV then heads towards the rectangular path relative to the selected area.

We report in Algorithm 1 the algorithm running on r_i. Messages are exchanged between the robots via the acoustic channel. We assume that the messages contain the status of the AUVs (e.g. the position, heading and speed) and information about the task allocation. This includes the primary and secondary task for which a robot r_k is bidding for, the bid values and, for the area surveillance tasks, the urgency factor at the time of message creation t_c of the task t_j as computed by a robot r_k.The information is kept minimal to make feasible the porting of the algorithm to the underwater scenario.

If messages from other robots are present, the Allocation Table structure A_T is updated by the function `UpdateAllocationTable`. A_T contains all the information regarding the vehicles and the knowledge about their allocation to the tasks. The function stores the allocation of a certain robot to the task in execution (both the primary and secondary).
`UpdateUrgencyFactors` updates the urgency factors $I(t)^i_k$ for each surveillance task t_i present in the system with the minimum of the received values from other assets bidding for task t_i, only if the minimum value is lower than the stored one. A received $I(t)$ lower than the current one means that the remote asset has likely more information to compute a more accurate $I(t)$. Finally, `StorePositionAssets` writes in A_T the information about the positions of the assets.

At each ping time, as a new measurement is made available, the robot estimates the positions of the other vehicles using the received information (`ComputePositionAssets`) and assuming a constant speed model.

The utilities of all tasks are then computed

Algorithm 1: PADA algorithm for r_i.

```
if IsMsgPresent()=true then
    UpdateAllocationTable() ;
    UpdateUrgencyFactors() ;
    StorePositionAssets() ;
end
if IsPingTime()=true then
    ComputePositionAssets();
    ComputeTaskUtilities();
end
if IsAuctionResolutionTime()=true and SecondaryTask()=continuous task then
    [resultAuction,isToCheckTaskReallocation]=ResolveAuctionContinuousTasks();
end
if IsAuctionResolutionTime()=true and PrimaryTask()=sporadic task then
    [resultAuction]=ResolveAuctionSporadicTasks() ;
end
[t_j, t_k]=SelectPrimarySecondaryTask(resultAuction,isToCheckTaskReallocation);
ProduceMsg(t_j, t_k);
UpdateWinnerTable();
UpdateAllocationTable();
with IsMsgPresent returns true if any message has been received from other assets;
IsPingTime returns true if it is time at which a ping has been performed;
IsAuctionResolutionTime returns true if it is time to resolve an auction;
SecondaryTask returns the secondary task for robot k;
PrimaryTask returns the primary task for robot k;
resultAuction being a string stating the results of an auction;
isToCheckTaskReallocation being a flag commanding to check a possible task reallocation;
ProduceMsg(t_j, t_k) function producing an acoustic message to be sent to collaborators;
```

(`ComputeTaskUtilities`) using the estimated positions/heading of the robots (and the related measurement locations).

A. *Continuous task auctions*

Periodically, every robot resolves the auctions. If the secondary task is a continuous one, the negotiation is resolved by `ResolveAuctionContinuousTasks` function. The period with which the auctions are resolved is set large enough to increase the probability for each robot to receive at least one bid from neighbours for the current task in execution. The auction resolution time is different for the two categories of tasks. For sporadic tasks, the auction time window is smaller than that continuous task auctioning's to allow a quicker resolution of negotiations. The bid of r_i for t_j, b_{ij} is computed as the cost with a negative sign, that is $-c_{ij}$, with c_{ij} being the scaled distance from the position of r_i to the closest point of the racetrack associated to t_j [24]. If no bids are present (only r_i is bidding on the task), r_i assumes to be the winner of the auction, and returns *isToCheckTaskReallocation*=1, to command a possible task reallocation. The robot evaluates all the bids received during the current auction time window for the task t_j in execution on r_i. If there are some bids received from other nodes, we remove from this set all the bids placed by the past winners on auctions on t_j which are stored in the structure W_i^j. In this case, r_i then assumes to be the winner of the auction if:

- No other bids are present (after the removal of previous auction winners).
- r_i has bid equal to the highest received bid(s) and the robot i is less suited than the other collaborators to execute other tasks in the pool. (higher costs to execute the surveillance tasks not in execution). With this rule, in case of parity during, the auction winner is the robot which would have the highest cost to execute the other task(s).
- r_i has bid higher than the maximum received bid.

Otherwise it assumes to have lost the negotiation and stores the id of the winner of the auction in W_i^j. The winners will not be considered as participants in future auctions by r_i for a certain temporal window. This avoids contrasting results in the auction due to imperfect (intermittent) communications. Further details of the `ResolveAuctionContinuousTasks` algorithm can be found in [24].

B. *Sporadic task auctions*

Sporadic tasks, once created and received, becomes primary tasks (see Sec. IV-C). This triggers the auction resolution system for sporadic tasks `ResolveAuctionSporadicTasks`. The bid of r_i for the task t_j, b_{ij}, is computed as the cost with a negative sign, that is $-c_{ij}$. The robot evaluates all the bids received during the current auction time window for the task t_j in execution on r_i. The robot creator of the task is always a winner of the auction. The task creator is indeed the only robot which has the knowledge about the task (i.e. the track). If r_i realises that the task creator is one of the bidders, the task creator is added to the set of winners on the task. In the case there are some bids received from other nodes, we remove from this set all the bids placed by the past winners on auctions on t_j as in the auction system for continuous tasks. In case of victory in the auction, if the robot is the task creator or if the robot has been bidding for the same task for a certain number of auctions, a particular condition is set, setting the secondary equal to the primary task. This means that the task has been "confirmed". Robots in this conditions are winning in auctions over others which have their primary task not equal to the secondary. If no bids are present, the robot assumes to be the winner of the auction. Otherwise it evaluates the received bids. If primary task is equal to secondary, r_i is the winner of the auction if:

- No bids are present produced by robots with the primary and the secondary task equal.
- Its bid is equal to some of those transmitted by robots with the primary and the secondary task equal, and its id is the highest one.
- Its bid is higher than the maximum of the bids received from robots with the primary and the secondary task equal.

Otherwise, in the case the primary task is different from the secondary one, the robot is the winner of the auction if there are no bids from robots with primary equal to secondary task and:

- Its bid is equal to the highest received bid(s) (received from robots with the primary and the secondary task different).
- Its bid is higher than the maximimum of the bids eceived from robots with the primary and the secondary task different.

Otherwise it assumes to have lost the negotiation and stores the id of the winner of the auction in W_i^j. As for the auctions with the continuous tasks, the winners will not be considered as participants in future auctions by r_i for a certain temporal window.

C. Task management

The `SelectPrimarySecondaryTask` function (in Algorithm 1 at line 16) selects the primary and secondary task in execution for r_i, using the results of the auctions. Let us define the set of available tasks as computed by robot i at time t as $T_i^a(t) = \{t_k \in T | \omega_k(t) > 0\}$ where $\omega_k(t)$ is the number of robots which can still be assigned to t_k at time t, given the maximum number of allocable robots l_k and the current known assigned robots. If no task is currently under execution on the robot, the highest utility task $t_j \in T_i^a(t)$ is selected. Let us assume that at the mission start no sporadic task is present in the task pool. The highest utility surveillance t_j task is therefore selected *both as primary and as secondary task.*

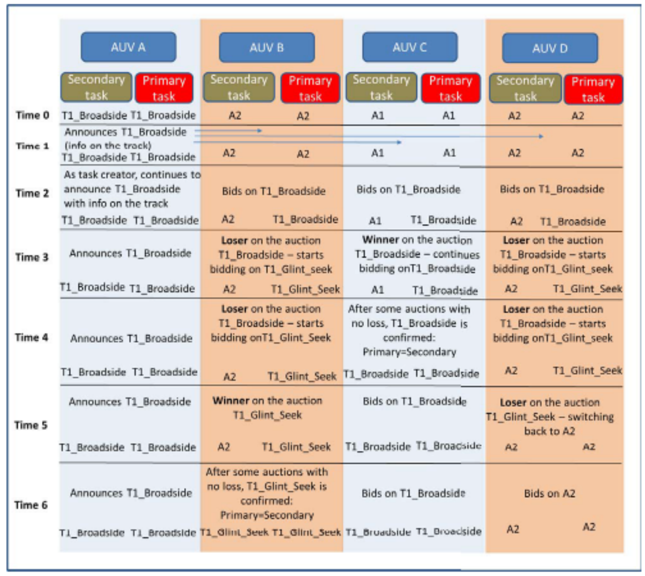

Fig. 3. Example of one auction including one sporadic task. The events of the auction are reported together with the primary and secondary task for each of the four robots involved. The continuous area surveillance tasks, A_1 and A_2, and the track prosecution in the example can be allocated to a maximum of two robots. The track prosecution has associated an ancillary glint seek task, which can be allocated to only one robot.

Then, the auctioning system influences the task in execution on the robot. If an auction is won on t_j, t_j continues to be the task in execution, both primary and secondary. If r_i loses the auction or if the bid of r_i is the only one for t_j (*isToCheckTaskReallocation*=1) (no other bids are received from other agents) a possible new task is sought. The use of the flag *isToCheckTaskReallocation* provides the opportunity to a robot which is the only one assigned to a certain t_j to evaluate a possible reallocation strategy. The robot excludes from the selection of a new task to bid for the tasks t_k for which it knows that they are assigned to a number of robots greater than l_k. The task $t_k \in T_i^a(t)$ with the highest utility u_{ik} is selected. If t_k is different from the task currently in execution t_j, it becomes the new primary and secondary task only if its utility $u_{ik} > u_{ij} + \eta_r$, with η_r being a reallocation threshold.

Sporadic tasks

As visible in Algorithm 1 where the variable *resultAuction* is overwritten, the results of auctions on sporadic tasks have priority of execution over those of negotiations on continuous tasks. If one robot creates a sporadic task t_s, it sets it *both as primary and secondary task*. The robot then announces it to the other robots of the network and executes it. The robots which receive the task (within a certain triggering distance from the track) change their *primary task* to t_s. From that moment on, auctions will be executed also on the sporadic task as previously described. One robot having t_s as its primary task and an area surveillance as secondary one is in a transition state until the sporadic task is definitely "confirmed". The sporadic task becomes confirmed if a certain number of auctions have passed with the robot continuing bidding for it.

In the case r_i loses an auction for t_s and t_s can produce an ancillary task t_s^a, the robot creates t_s^a and sets it as primary task and commences to bid for it as a new sporadic task. Otherwise, if no ancillary task are produced by t_s or if the sporadic task ends, the robot selects the next task from the pool. If sporadic tasks are present, it selects the highest utility one as its primary task while selecting the highest utility continuous task as secondary. In case no sporadic tasks are presents, the highest utility continuous task is set both as primary and secondary.

After the management of the task in execution, the acoustic message with the new bid(s) is created (`ProduceMsg`). The structure containing the winners for all the tasks, W_i, is then updated removing, for each task, the robots which have not resulted as winners in auctions for that task for a certain amount of time (`UpdateWinnerTable`). A robot is also removed as assigned to a task from the allocation table A_T if no information has been received about that assignment for some time (`UpdateAllocationTable`). This allows to negotiate again a task after some time to better adapt the general team allocation as the tactical scenario evolves.

Given n robots, each implementing the PADA algorithm, the team is always guaranteed to reach a feasible assignment to the tasks. By assuming perfect communications, the algorithm reaches a feasible assignment after n auctions, in the worst case scenario. Assume there are n tasks, with each one having $l_i = 1$. Without loss of generality, we consider that the n robots start bidding for the first task. The PADA algorithm assumes that one of the robot is recognised by the team as the winner at each auction. The remaining $n - 1$ robots bid for the second task and a new winner is declared. Proceeding in this way, after n auctions each robot is assigned to one task.

Example of auction

In Fig. 3 we report a diagram of an auction involving a sporadic task and four robots. The continuous tasks and the track prosecution task in the example can be allocated to a maximum of two robots. In the diagram, A_1 and A_2 are two generic area surveillance tasks. The track prosecution has associated an ancillary glint seek task which can be allocated to a single robot. AUV A confirms a track and switches both its primary and secondary task at $T1\_Broadside$ and manoeuvres to keep the target at broadside. Then, it announces the task to the team. The other robots receive the announcement, and switch their primary task to $T1\_Broadside$. All of them change their behaviour to move the target at broadside and bid for the task. AUV C is the winner of the auction since it has the lowest cost for the task. Since $T1\_Broadside$ can be allocated to a maximum number of two robots, AUV B and AUV D needs to change the task to execute and they create $T1\_Glint\_seek$, the ancillary task of $T1\_Broadside$. They bid for that task and navigate to reach a position favourable to detect the target glint. AUV C after some time confirms $T1\_Broadside$ and its secondary task is set to $T1\_Broadside$. At time 5, finally, AUV B wins the auction and AUV D changes task in execution setting back both its primary and secondary task to A2.

V. Results

Matlab (©MathWorks) simulations were carried out to evaluate the algorithm's performance and to study how changes in its parameters can influence the overall team's behaviour. In particular, we were interested to study the effect of the communications quality on the allocation performance. An ASW network with characteristics similar to CMRE's (see Fig. 1) was simulated in different scenarios. Communications are simulated in a realistic way taking into account packet loss and a Time-Division Multiple Access (TDMA) channel access method [15] is used. A detailed qualitative and quantitative analysis of task allocation with area surveillance tasks is reported in [24]. We focus in this work on simulations which investigate the performance of PADA when sporadic tasks are injected in the task pool. Four mobile robots, acting as receivers in a multi-static ASW network, were considered in the reported simulations.

Two area surveillance tasks are present in the task pool: t_1 related to the more southern area, A_1, while t_2 related to an area A_2 located in the north-west part of the operational area. Both tasks have the same parameters in terms of coverage index and urgency factor, but A_1 is larger than A_2. This implies that A_1 requires a survey more accurate than A_2 to reach the mission objectives. Both tasks can be simultaneously executed by a maximum of two vehicles. The related areas are associated with two rectangular path patterns. A robot, after selecting one surveillance task to execute, moves at its cruise speed to the associated path and then covers it until the task in execution is changed. At a certain time, one of the robots confirms a track as likely related to the target. A ***track prosecution at broadside*** task is then created and announced to the team together with information related to the confirmed track. The network reacts to it and allocates the best suited robots to the track prosecution. Also the ancillary ***glint seek*** task is created. Two different acoustic communications conditions were investigated:

- Scenario 1 - perfect communications.
- Scenario 2 - 40% of message transmission success rate up to 3 km of range between the communicating assets. Then the success rate decreases linearly up to 10% at a 6 km range.

The parameters used for this simulation session are reported in Tab. I. One of the key point in handling the sporadic tasks is a quick response of the network. For this reason the auction period is reduced for the auctions regarding sporadic tasks, to allow a faster allocation.

A typical simulation with perfect communications is described in the following. The task allocation of the team robots during this simulation is shown in Fig. 6, both for the primary and for the secondary tasks. The AUVs commence their mission at the positions visibile in Fig. 4 (left). They

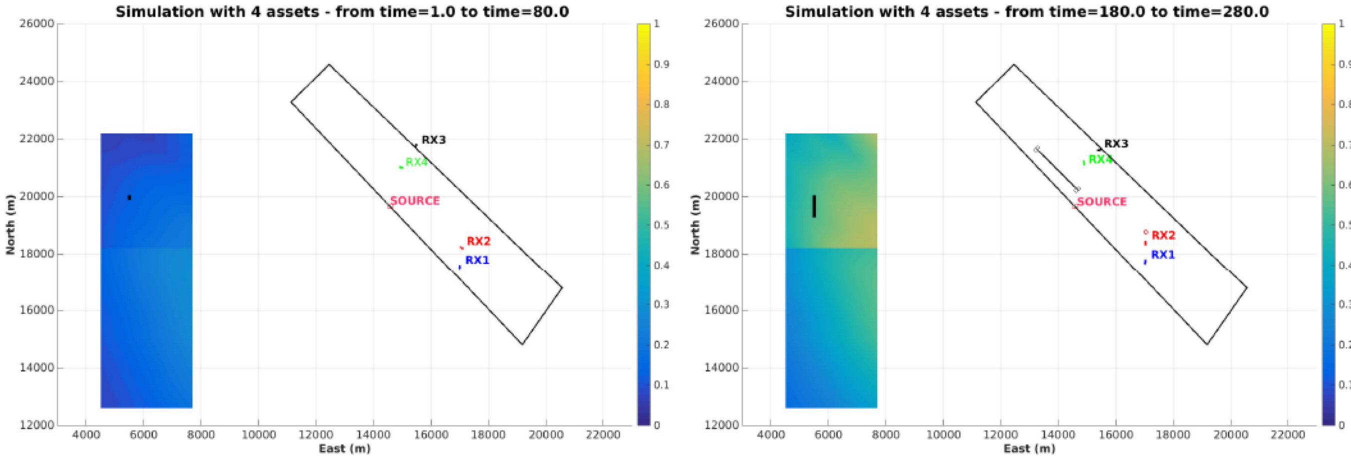

Fig. 4. (Left) Scenario 1 - situation of the assets at mission time 80 s. The acoustic source considered at a fixed position is indicated by a red square. RX_2 is announcing the *track prosecution at broadside* task to the network. The two areas to survey, A_1 to the south and A_2 to the north, are visibile. The colour is related to the cumulative probability of detection computed by the robots assigned to the survey. The black line in the more northern area is the track confirmed by the on board tracker of RX_1. (Right) Scenario 1 - situation of the assets at mission time 280 s. RX_2 is reaching the glint location.

TABLE I
PARAMETERS USED FOR THE SIMULATIONS.

Parameter	Value	Notes
Parameters related to the assets		
Used robots	4	
Robot cruise speed [m/s]	1	
Robot glint seek speed [m/s]	3	
Target speed [m/s]	3	
PRI [s]	20	Pulse Repetition Interval
Parameters related to the acoustic communications		
Time-slot for each asset [s]	6.5	5 s to send messages
TDMA frame length [s]	46	
Parameters related to the auction		
Auction resolution period periodic [s]	230	auction resolution period- 5 TDMA frames
Auction resolution period sporadic [s]	92	auction resolution period - 2 TDMA frames
Parameters related to the tasks		
Number of tasks	4	2 area surveillance tasks, one track prosecution broadside and one ancillary glint seek
l_k for periodic tasks	2	Maximum number of vehicles which can be assigned to t_k
l_k for track prosecution broadside	2	
l_k for glint seek	1	

select the surveillance tasks: RX_1, RX_2 and RX_3 select t_1, and RX_4 chooses t_2. A target moving (the black thick line in Fig. 4) to the south is detected in A_2 area and a track is created by the tracker on board RX_1. The track, named track 57 by the system, is confirmed as likely related to the target. RX_1 creates a *track prosecution at broadside* task for track 57 (t_3), and announces it to the network together with the related information about the track. The other robots receive the task and select it as their primary task (see Fig. 6). They begin to bid for the task and manoeuvre to keep the target at broadside of their array. At the first auction resolution, RX_4 results the winner (RX_1 does not participate being the task creator). RX_3 and RX_2, since there are two allocated robots to the task (RX_1 and RX_4), create a ***glint seek*** task (t_4) for track 57 and bid for that (see Fig. 6 at 92 s of mission time). At the same time, they move towards the closest point on the line which offers the possibility to detect the glint, that is the specular reflection from the target.

In the following auction, RX_3 loses an auction for the ***glint seek*** task and comes back to survey area A_2 (see Fig. 6 at 184 s of simulation time). RX_2 continues to move towards the glint line. The glint line moves rapidly towards RX_2 due to movement of the target towards south. This can be seen in Fig. 4 (right) where the situation at simulation time 280 s is described.

Then, the target turns towards the south-east. This changes immediately the area of where the glint can be detected (see the red diamond in Fig. 5). This shows how difficult is for the ASW network to reach a position to detect the glint. The large variation in space of the glint locations suggest, for the surveillance network, to increase the number of its sensing nodes to create more opportunities of glint detections.

Finally, when the track breaks, each team robot selects an area to survey coming back to auctions on continuous tasks.

The task allocation of the robots during a simulation in the Scenario 2 is shown in Fig. 7, both for their primary and secondary task. In this case, the degraded communications performance create a delay in the response of the network to the track announcement. However, the network is capable of responding to the the situation. As can be seen in Fig. 7, RX_1 is the first to change its primary task since it creates the track prosecution task related to track 57. Differently from the case with perfect communications, in which all the other robots react to the announce (see the switch of the three robots to t_3 at $\circ$ 40 s in Fig. 6), in this case only RX_4 receives the task announcement of t_3. RX_2 receives the announcement of the task later and finally RX_4 receives

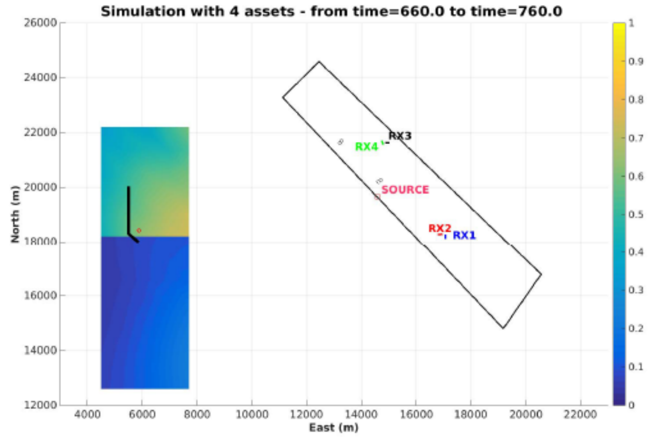

Fig. 5. Scenario 1 - situation of the assets at mission time 760 s. The turning of the track to south-east moves immediately the glint to a different area of the region (see the red diamond).

it. Despite these delays, the allocation is correctly performed by the network, with RX_1 and RX_4 assigned to the *track prosecution at broadside* and RX_2 allocated to the *glint seek task*.

VI. Conclusion

We described PADA, a market-based policy for task allocation specifically designed for environments characterised by unreliable, intermittent communications, such as the underwater one. The algorithm works in a completely distributed way and, through periodic auctions, sequentially allocates the robots to the tasks. It is based on negotiations among neighbouring robots and requires only local communication, with no need for a central auctioneer. The algorithm waits for messages for several TDMA frames before making a decision. These features increase the robustness to poor communications. The task allocation policy takes into account task reallocations to better adapt to the evolving conditions.

We extended our previous work [24] introducing the possibility to handle two kinds of tasks, the so-called continuous tasks which never terminate and remain always available in the task pool and sporadic ones, created upon the occurrence of an event, and that requires a rapid network response. We considered as a case study an ASW robotic surveillance network in which AUVs operate as autonomous, mobile receivers in a multi-static sonar system. The ASW mission is decomposed in continuous tasks, that is surveillance of areas of interest, and in sporadic tasks created whenever a track produced by an on board tracker algorithm is considered as likely related to a target. The areas of interest have to be surveyed periodically by the robots due to the dynamic nature of the problem (targets move) and a specific utility was designed [24]. When a cue of the target is detected, the AUVs have to react rapidly adapting their navigation to increase the tracking performance for a final target classification. To enable cooperative strategies more than one robot can be allocated to each task. This possibility increases the mathematical difficulty of the assignment problem making it a coalition formation problem [9].

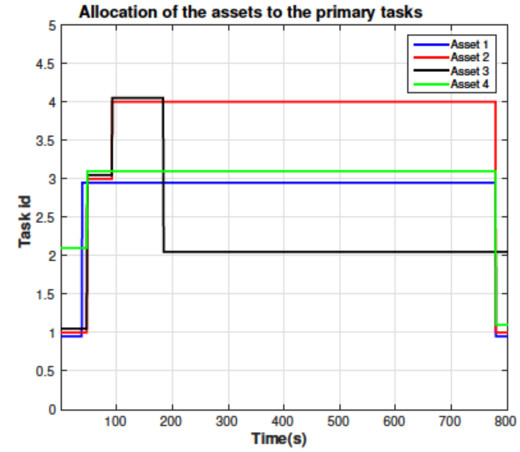

(a) Task allocation (primary tasks)-Scenario 1 with sporadic tasks.

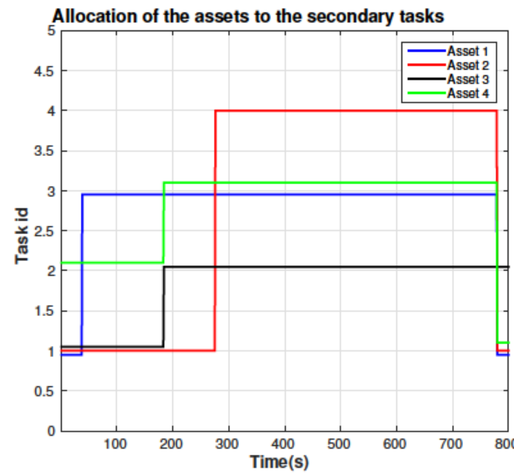

(b) Task allocation (secondary tasks)-Scenario 1 with sporadic tasks.

Fig. 6. Task allocation (primary (a) and secondary (b) task) in the Scenario 1 (perfect communications) with sporadic tasks. Task with id 1, t_1, is the area surveillance for the most southern area, t_2 is the area surveillance task for the other region of interest. t_3 is the track prosecution at broadside task for track 57 and t_4 is the glint seek task for the same track. t_4 is an ancillary task of t_3.

The proposed algorithm was tested in non-trivial Matlab simulations. In this paper we reported results from allocations with sporadic tasks produced by some robots. Results demonstrate that the proposed framework can manage the task allocation, providing a quick response to the occurred events even when communications are degraded.

The capability to handle both continuous and sporadic tasks in the allocation offers the possibility to manage the tasks of different nature that compose a complex ASW mission. Future work will expand the framework introducing communications and localisation optimisation tasks. These kinds of tasks can either be managed as continuous or as sporadic tasks (e.g. ancillary of a task prosecution task) considered the current tactical scene.

We will also further investigate situations in which one robot cannot communicate with other nodes at all for long time. Increasing the scheme robustness in these extreme

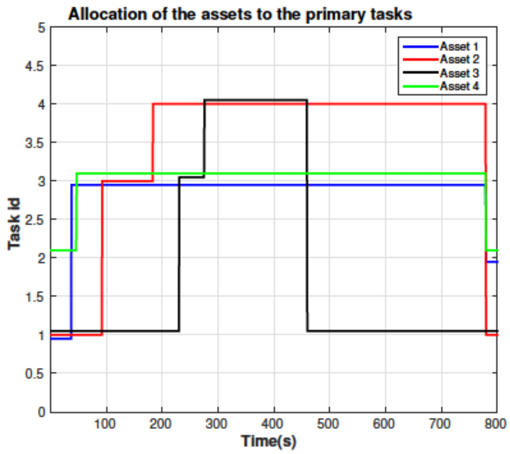

(a) Task allocation (primary tasks)-Scenario 2 with sporadic tasks.

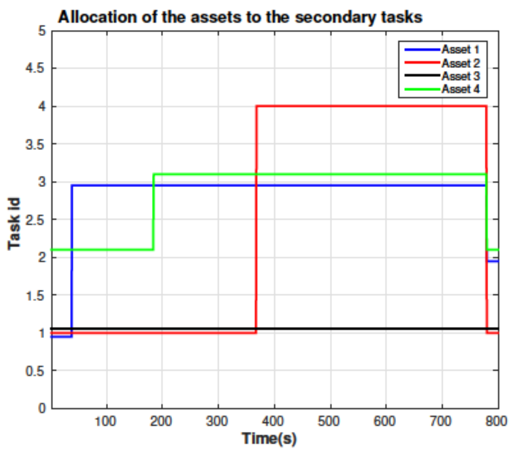

(b) Task allocation (secondary tasks)-Scenario 2 with sporadic tasks.

Fig. 7. Task allocation (primary (a) and secondary (b) task) in the Scenario 2 (degraded communications) with sporadic tasks. Task with id 1, t_1, is the area surveillance for the most southern area, t_2 is the area surveillance task for the other region of interest. t_3 is the track prosecution at broadside task for track 57 and t_4 is the glint seek task for the same track. t_4 is an ancillary task of t_3.

conditions is indeed a strict requirement before its fully use at sea in the CMRE network.

ACKNOWLEDGEMENTS

This work has been supported by the NATO Allied Command Transformation under the Collaborative Anti-Submarine Warfare research programme.

REFERENCES

[1] G. Ferri, A. Munafò, A. Tesei, P. Braca, F. Meyer, K. Pelekanakis, R. Petroccia, J. Alves, C. Strode, and K. LePage, "Cooperative robotic networks for underwater surveillance: an overview," *IET Radar, Sonar &. Navigation. Special Issue: Sonar multi-sensor applications and techniques*, vol. 11, no. 12, pp. 1740 – 1761, 2017.

[2] S. Sariel, T. Balch, and N. Erdogan, "Naval mine countermeasure missions," *IEEE Robotics Automation Magazine*, vol. 15, no. 1, pp. 45–52, March 2008.

[3] D. T. Hughes, F. Baralli, S. Kemna, K. LePage, M. Hamilton, and A. Vermeij, "Collaborative multistatic ASW using AUVs: demonstrating necessary technologies," Nato Undersea Research Centre, Tech. Rep., 2009.

[4] G. Ferri, A. Munafò, R. Goldhahn, and K. LePage, "A non-myopic, receding horizon control strategy for an auv to track an underwater target in a bistatic sonar scenario," in *in Proceedings of CDC2014, Los Angeles, December 15-17, 2014*, 2014.

[5] K. LePage, R. Goldhahn, J. Alves, C. Strode, P. Braca, G. Ferri, A. Munafò, M. Oddone, J. Sildam, F. Baralli, S. Biagini, G. Canepa, M. Colombo, V. Grandi, G. Grenon, M. Mazzi, M. Micheli, G. Parisi, P. Saia, A. Vermeij, and G. Zappa, "Autonomous networked anti-submarine warfare research and development at CMRE," in *OCEANS 2015 - Genova*, May 2015, pp. 1–6.

[6] http://liquidr.com/.

[7] G. Ferri, A. Munafò, and K. LePage, "An autonomous underwater vehicle data-driven control strategy for target tracking," *IEEE Journal of Oceanic Engineering*, online, 2018.

[8] P. Braca, R. Goldhahn, G. Ferri, and K. LePage, "Distributed information fusion in multistatic sensor networks for underwater surveillance," *Sensors Journal, IEEE*, 2015.

[9] B. Gerkey and M. Mataric, "A formal analysis and taxonomy of task allocation in multi-robot systems," *International Journal of Robotics Research*, vol. 23, no. 9, pp. 939–954, 2004.

[10] G. Korsah, A. Stentz, and M. Dias, "A comprehensive taxonomy for multi-robot task allocation," *The International Journal of Robotics Research*, vol. 35, pp. 514–527, 2013.

[11] D. Bertsekas, "Dynamic programming and suboptimal control: a survey from ADP to MPC," in *CDC Proceedings*, 2005.

[12] N. Michael, M. Zavlanos, V. Kumar, and G. Pappas, "Distributed multi-robot task assignment and formation control," in *ICRA 2008*, 2008.

[13] A. R. Mosteo and L. Montano, "A survey of multi-robot task allocation," University of Zaragoza, AMI-009-10-TEC, Tech. Rep., 2010.

[14] Z. Yan, N. Jouandeau, and A. Cherif, "A survey and analysis of multi-robot coordination," *International Journal of Advanced Robotic Systems, 10*, 2013.

[15] A. Munafò and G. Ferri, "An acoustic network navigation system," *Journal of Field Robotics*, 2017.

[16] L. E. Parker, "ALLIANCE: An architecture for fault-tolerant multi-robot cooperation," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 2, pp. 220–240, 1998.

[17] M. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," *Proceedings of the IEEE*, 2006.

[18] A. Viguria, "Market-based distributed task allocation methodologies applied to multi-robot exploration," Ph.D. dissertation, University of Seville, 2009.

[19] N. Leonard, D. Paley, R. Davis, D. Fratantoni, F. Lekien, and F. Zha, "Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in monterey bay," *Journal of Field Robotics 27(6), 718740 (2010)*, 2010.

[20] J. Kalwa, M. Carreiro-Silva, F. Tempera, J. Fontes, R. S. Santos, M.-C. Fabri, L. Brignone, P. Ridao, A. Birk, T. Glotzbach, M. Caccia, J. Alves, and A. Pascoal, "The MORPH concept and its application in marine research," in *OCEANS 2013 MTS/IEEE Bergen, Norway*, 2013.

[21] Y. Deng, P. Beaujean, E. An, and E. Carlson, "Task allocation and path planning for collaborative AUVs operating through an underwater acoustic network," in *OCEANS 2010*, 2010.

[22] I. Kulkarni and D. Pompili, "Task allocation for networked autonomous underwater vehicles in critical missions," *IEEE Journal on Selected Areas in Communications, 28, 5*, 2010.

[23] N. Tsiogkas, G. Papadimitriou, Z. Saigol, and D. Lane, "Efficient multi-AUV cooperation using semantic knowledge representation for underwater archaeology missions," in *Oceans - St. John's, 2014*, 2014.

[24] G. Ferri, A. Munafò, A. Tesei, and K. LePage, "A market-based task allocation framework for autonomous underwater surveillance networks," in *OCEANS 2017, Aberdeen, UK*, 2017.

[25] G. Ferri, A. Munafò, R. Goldhahn, and K. LePage, "Towards fully autonomous underwater vehicles in ASW scenarios: an adaptive, data driven AUV mission manager layer," in *OCEANS 2015, Genova, Italy, 18-21 May, 2015*, 2015.

[26] B. Gerkey and M. Mataric, "A framework for studying multi-robot task allocation," in *In Multi-Robot Systems: From Swarms to Intelligent Automata, Volume II*, 2003.

Document Data Sheet

Security Classification		Project No.
Document Serial No. CMRE-PR-2019-012	**Date of Issue** May 2019	**Total Pages** 10 pp.

Author(s)

Gabriele Ferri, Jeffrey Bates, Pietro Stinco, Alessandra Tesei, Kevin LePage

Title

Autonomous underwater surveillance networks: a task allocation framework to manage cooperation

Abstract

The design of efficient task allocation schemes is essential to manage autonomous underwater robotic surveillance networks. The network has to assign the most suited robots to the tasks which compose the mission, in spite of the unreliable and intermittent underwater communications. In this paper, we describe a market-based policy for this. It works in a completely distributed way and, through periodic auctions, sequentially allocates the robots to the tasks. There is no central auctioneer and each robot can resolve the current auction. These features increase the robustness to poor communications. The proposed scheme can manage two kinds of tasks, continuous tasks which never terminate and are always available in the task pool and sporadic ones, created upon the occurrence of some events and requiring a rapid response of the network. An Anti-Submarine Warfare scenario is considered to validate the task allocation scheme. Surveillance of areas of interest are the addressed continuous tasks, while sporadic tasks are created whenever a track produced by the tracker on-board a robot is confirmed as likely related to a target. In this case the team has to rapidly react to the event and to increase the tracking performance. Results from non-trivial Matlab simulations are reported and demonstrate the effectiveness of the allocation scheme in degraded communications conditions.

Keywords

Task Allocation, Autonomous Robotic Networks, Underwater Surveillance, Autonomous Underwater Vehicle, Anti-Submarine Warfare

Issuing Organization

NATO Science and Technology Organization
Centre for Maritime Research and Experimentation

Viale San Bartolomeo 400, 19126 La Spezia, Italy

[From N. America:
STO CMRE
Unit 31318, Box 19, APO AE 09613-1318]

Tel: +39 0187 527 361
Fax:+39 0187 527 700

E-mail: library@cmre.nato.int