Naval Mine Countermeasure Missions

A Distributed, Incremental Multirobot Task Selection Scheme

BY SANEM SARIEL, TUCKER BALCH, AND NADIA ERDOGAN

© DREAMSTIME

© DREAMSTIME

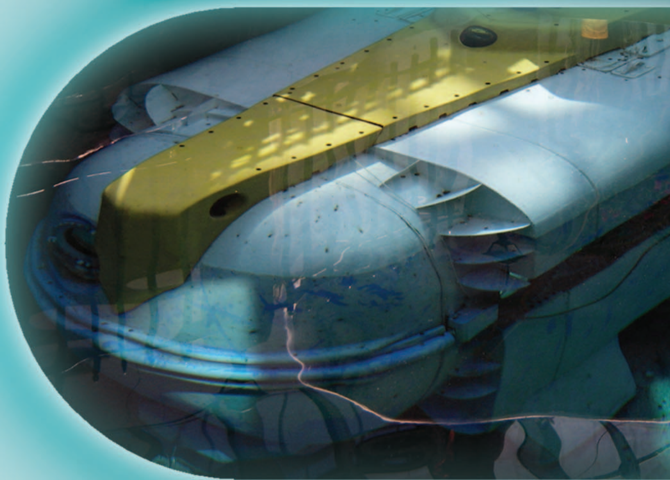

© ISTOCKPHOTO

Mobile Multirobot Systems

Undersea operations using autonomous underwater vehicles (AUVs) provide a different and in some ways a more challenging problem than tasks for unmanned aerial vehicles and unmanned ground vehicles. In particular, in undersea operations, communication windows are restricted, and bandwidth is limited. Consequently, coordination among agents is correspondingly more difficult. In traditional approaches, a central planner initially assigns subtasks to a set of AUVs to achieve the team goal. However, those initial task assignments may become inefficient during real-time execution because of the real-world issues such as failures. Therefore, initial task allocations are usually subject to change if efficiency is a high concern. Reallocations are needed and should be performed in a distributed manner. To provide such flexibility, we propose a distributed auction-based cooperation framework, distributed and efficient multirobot-cooperation framework (DEMiR-CF) [1], which is an online dynamic task allocation (reallocation) system that aims to achieve a team goal while using resources effectively. DEMiR-CF, with integrated task scheduling and execution capabilities, can also respond to and recover from real-time contingencies such as communication failures, delays, range limitations, and robot failures. It has been implemented and tested extensively in the multirobot multitarget exploration domain [2] and in complex missions of interrelated and resource constrained tasks [3]. In this article, we report the performance of the framework against real-world difficulties encountered in multi-AUV coordination for the naval mine countermeasure (MCM) mission obtained through several experiments on the U.S. Navy's Autonomous Littoral Warfare Systems Evaluator-Monte Carlo (ALWSE-MC) simulator [4]. DEMiR-CF supports a distributed strategy for real-time task execution and is designed to use the advantages of auction-based approaches. Additional precaution routines are integrated into the framework to enhance solution quality. Other works in auction-based coordination research include M+ [5], MURDOCH [6], TraderBots [7], and the allocation scheme by Lemaire [8]. According to the review given in [9], existing auction-based systems are not fully capable of replanning task distributions, redecomposing tasks, rescheduling commitments, and replanning coordination during execution. Our approach aims at filling these gaps. We propose an integrated cooperation framework for multirobot task execution and analyze the performance of the precaution routines and solution quality maintenance schemes for single-item auctions in a multi-AUV coordination context [10]. Experiments are performed in a realistic simulation environment with real-time constraints and events such as AUV failures and limitations, and delays in communication range. Precaution routines embedded into the framework not only recover from failures but also serve to

Digital Object Identifier 10.1109/M-RA.2007.914920

Additional precaution routines are integrated into the framework to enhance solution quality.

maintain a high solution quality. Our experiments show that communication delays significantly influence the solution quality and should be analyzed in multirobot systems, especially working in harsh environments. As the experiments and scenarios demonstrate, online task handling performance of DEMiR-CF is considerably promising.

Naval MCM Missions

Naval MCMs are actions taken to counter the effectiveness of underwater mines. MCM operations include finding and seizing mine stockpiles before they are deployed, sweeping desired operational areas, identifying mined areas to be avoided, and locating and neutralizing individual mines [11]. Our research is focused on the subset of MCM operations that involve locating and mapping all individual mines in an operational area. In general, recognizing proud mines on the seafloor is not overly difficult; the difficulty arises with the abundance of nonmine objects on the seafloor that possess mine-like characteristics (e.g., geologic outcroppings, coral, manmade debris) [12]. This ample supply of false alarms has necessitated the following strategy typically employed by the navy: detect and classify mine-like objects (MLOs) with high-coverage rate sensors (e.g., sidelooking sonar), employ advanced signal processing techniques for maximal false alarm reduction, and then revisit the remaining MLOs with identification-quality assets (e.g., electrooptic sensors) to confirm them as mines or dismiss them as false alarms. A sample image in which an MLO remains is illustrated in Figure 1.

The reference mission in this research is to detect, classify, and identify underwater mines in a given operational area simulated in ALWSE-MC [4], an analysis package designed to simulate multiple autonomous vehicles performing missions in the littoral regions, including mine reconnaissance, mapping, surveillance, and clearance. This mission employs two types of vehicles: unmanned underwater vehicles (UUV), which are free-swimming AUVs and possess large-footprint sensors (e.g., side-scan sonar) for detection and classification (D/C) of mines, and seafloor crawlers equipped with short-range, identification-quality sensors (e.g., camera). The crawlers have the ability to stop at an object and take a picture with a camera.

Figure 1. Sidelooking sonar sensors may fail in correctly classifying mines because of their similarities to some nonmine objects in undersea habitat [12].

The MCM domain has important similarities to some of the well-known domains where the use of a multirobot team is usually beneficial. The search and rescue domain where different types of robots are required is one example. Searching for victims in the disaster area is similar in nature to the detection of mines. Rescue operations in which first aid is provided to victims are also similar to the classification tasks. Another interesting domain, the space exploration mission, has a high resemblance in form also. The mission can be divided into two submissions: searching for important points to reconsider and revisiting the sample points determined in the first phase to further investigate specific locations and collect scientific data with more specialized robots. Therefore, we believe that the solutions offered to carry out the MCM mission can be successfully applied to these domains also.

DEMiR-CF

The MCM mission is performed undersea and in real time. Managing the overall robot team by a central authority is not usually possible because of the limitations of the real-world environment. Therefore, each individual robot should find a way to solve the global problem from its local perspective while assuming a global approach is possible in a distributed setting.

To meet the real-world limitations, we propose a dynamic and distributed task allocation scheme, DEMiR-CF, to coordinate robots that cooperate to fulfill different parts of a mission. DEMiR-CF is designed for complex missions including interrelated tasks that require diverse (heterogeneous) capabilities and simultaneous execution [1], [13]. Dynamism is achieved through incremental selection and allocation of the targets. By means of the distributed characteristic of the proposed allocation scheme, each robot is allowed to select a candidate task for itself and, next, the robots proceed to cooperate in the process of selecting the most suitable robots for the tasks. A time-extended view is considered while selecting tasks after forming rough schedules. The framework combines the dynamic priority-based task selection scheme, distributed task allocation procedures and coalition formation schemes as cooperation components, and Plan B precaution routines, some of which are implemented by the coalition maintenance or dynamic task selection scheme. These components are integrated into a single framework to provide an overall system that finds near-optimal solutions for real-time task execution. The modules that embody the framework and information flow among them are given in Figure 2. Each robot keeps a model, which includes information on current status, of the other robots and the mission tasks. The model update module, the (system) consistency checking module, and the dynamic task selector module perform Plan B precaution routines by either updating the model maintained by the robot or activating the warning mechanisms. Model updates are initiated by either incoming information from the other robots or information perceived by the robot itself. If a system inconsistency arises, the consistency checking module is responsible for initiating warning mechanisms and informing the corresponding robots. The dynamic task selector module selects the most suitable task by considering the model of the robot. The distributed allocation scheme ensures the distributed task allocation by executing the required negotiation procedures for the selected task. The execution or coalition scheme implements synchronized task execution and coalition maintenance procedures. Task models are updated according to the selected task and the task currently in execution. A sample flow of the operations in the framework (as depicted in Figure 2) is summarized as follows:

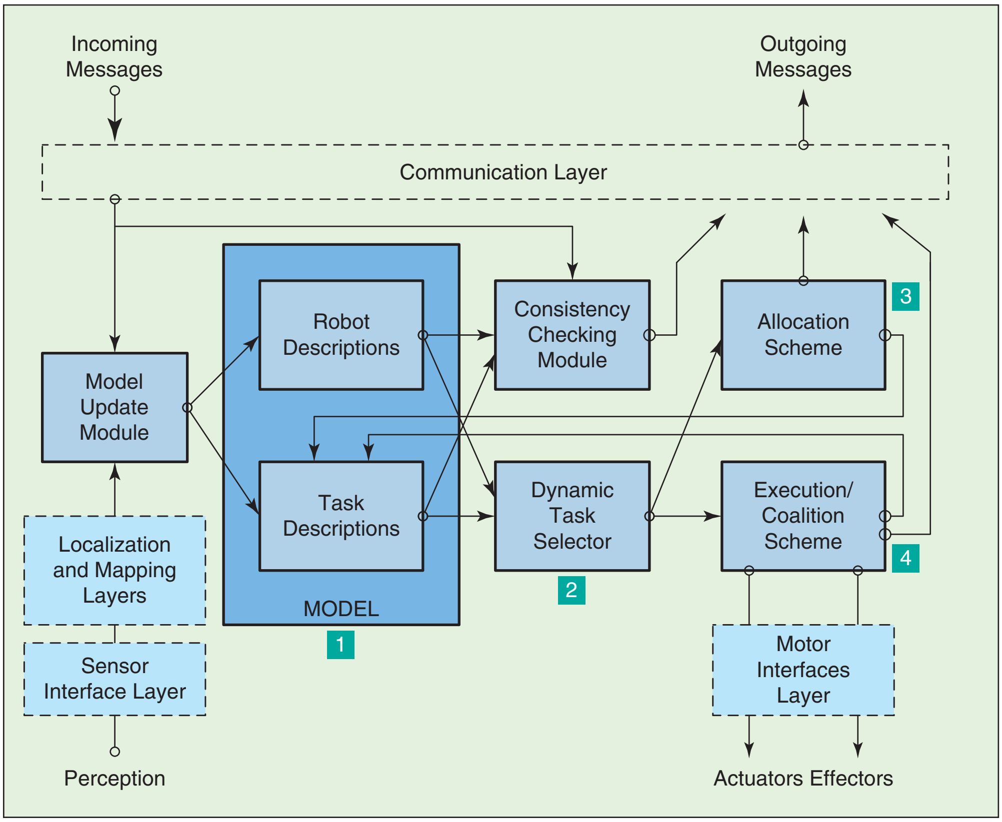

Figure 2. DEMiR-CF modules.

- Initially, robots are delivered the mission task definitions.
- Each robot selects the most suitable candidate task to execute through global cost consideration (dynamic task selection or switching).
- Robots offer auctions for the tasks they have selected. During auction steps, inconsistencies are cleared and conflicts are resolved.
- Task assignments are made for the announced tasks, making sure that each robot takes part in the most suitable execution when the global solution quality is considered.
- Dynamic task selection or switching proceeds simultaneously with task execution. This allows the robot to switch between tasks when executing the candidate task becomes more profitable than continuing with the current task and handling real-time contingencies at the same time. Hence, corresponding auction and selection procedures (second through fourth items) are applied continually.

DEMiR-CF is designed with the capability to deal with real-time situations. The framework can efficiently respond to these events and maintain the solution quality simultaneously with real-time task execution.

Plan B Precautions

Plan B precautions are taken in DEMiR-CF by the model update module, which updates the system model of the robot, and the system consistency checking module. The model update module uses incoming information from the other robots and its own perception data to update the world model. The system consistency checking module provides warning that initiate actions to keep the system consistent.

Recovery operations may include warning other robots about the problem or changing the model accordingly. Inconsistencies usually arise when robots are not informed about tasks that are achieved, under execution, or under auction in real-world operations. To keep system consistency, robots use explicit communication and broadcast the information as follows:

- Tasks known to be achieved in predefined time periods to prevent redundant executions. (This feature provides a bucket-brigade type of information sharing that enables information transition from one robot to another where point-to-point access is not possible, and consequently communication range limitations are resolved.)
- Newly discovered online tasks that are not yet achieved.
- Task execution messages in predefined time periods. (These messages contain the updated cost value and the estimated task achievement deadline information. Therefore, they serve as clues, meaning that the executer robot is still alive and the task is under execution.)
- Task achievement message when the task is achieved.
- Cancellation message if the task execution is canceled.
- Task invalidation message when an invalidity is detected.

Incoming messages from other robots are taken as clues for being marked as running properly. Some misleading beliefs such as setting the state of a robot as failed although it is running

In undersea operations, communication windows are restricted and bandwidth is limited.

properly may cause parallel executions. This is a desired feature from the point of view of the completion of mission. Designed precautions resolve these kinds of inconsistencies if communication resources permit in later steps. In designing the precautions, it is assumed that robots are trusted and benevolent.

Task Representation for the MCM Mission

Our general task representation is capable of describing complex tasks with interdependencies [1]. However, in this particular case study, tasks do not have interdependencies. Two types of tasks are defined for vehicles: visit waypoint (w) and identify MLO (t). The task representation includes the capabilities required for each type of task: reqcap_w contains side-scan sonar and reqcap_t contains cameras besides the standard capabilities of AUVs common in both types of vehicles. The coverage mission (M_C) contains predefined number of waypoints ($w_i \in M_C$, $0 < i \leq ||M_C||$) to be visited by all UUVs ($R_{\text{UUV}} \subset R$). One way to represent a task is to directly assign it for each waypoint. However, this representation has a drawback of high communication requirements for efficient completion of the mission. Instead, we represent tasks as interest points of regions or search areas ($W_k = \cup w_i$, $\forall w_i$ is unvisited, and $W_k \subseteq M_C$). These regions (and the corresponding centers) are determined by the robots during runtime dynamically although the waypoint locations are fixed at known coordinates. Therefore, both the allocation of the waypoints to the robots and the paths constructed to traverse these waypoints are determined online by negotiations. Negotiating the interest points (regions) instead of the individual waypoints reduces the communication overhead. Regions determined by different UUVs may vary during runtime and may sometimes overlap. However, the uncertainty related to the region determination is within an acceptable range, especially when the cost is compared with the requirements of complete knowledge sharing by representing each waypoint as a task. Before defining the regions, the relative distance values, $\text{reldist}(r_j, w_i)$, are determined for each unvisited waypoint w_i using (1), where function dist returns the Euclidean distance between points. r_k locations are the latest updated locations of the robots. If there is no known active robot assumed to be running properly, $\text{reldist}(r_j, w_i)$ is the value of the distance between the robot and the waypoint

$$\text{reldist}(r_j, w_i) = \text{dist}(r_j, w_i) - \min_{\forall k \neq j}(\text{dist}(r_k, w_i)). \tag{1}$$

Each robot defines its regions (W_{jk}, $1 \leq k \leq ||R_{\text{UUV}}||$). The number of regions equals the number of UUVs believed to be running properly. After sorting the $\text{reldist}(r_j, w_i)$ values of the unvisited waypoints in descending order as an array, the array is cut into subarrays that represent the regions. Each region contains approximately an equal number of waypoints. Each robot specifies the region of highest interest as its first region. If the robots are closely located, the regions of highest interest may overlap. In this case, negotiations are needed to resolve conflicts and to assign only one robot for each region.

The identification mission (M_I) contains an unknown number of tasks for the MLO locations ($t_i \in M_I, 0 < i \leq ||M_I||$) to be visited by the crawlers. Therefore, the tasks in M_I are generated online during runtime.

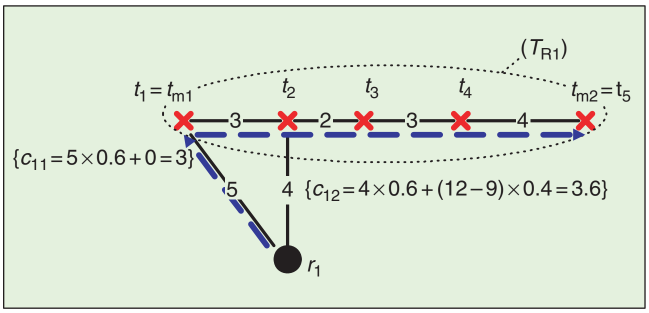

Figure 3. *Target selection strategy by the FAC heuristic function.*

Exploration for Detection of MLO Locations

To begin the mission, the UUVs survey the operational area following waypoints determined a priori; however, corresponding regions containing waypoints may be reassigned by the negotiations among UUVs autonomously. After determining regions, each UUV proposes an auction for the region of highest interest (interest point). After negotiations on several auctions, each UUV is assigned to the closest region (interest point). If more than one robot is almost at the same distance from the interest point, the one with the smaller id number is assigned to the region. The other UUVs continue to offer auctions for the remaining regions. Allocations of the regions may also change during run time to maintain higher solution quality. Whenever UUVs detect failures or recoveries from failures, they change their region definitions accordingly and offer new auctions. After the region assignments are completed, each robot visits waypoints in its region (W_j) in a sequence identified by an ordering of their cost values from the smallest to the largest:

$$\begin{aligned} c(r_j, w_i) = {} & \alpha \cdot \text{dist}(r_j, w_i) \\ & + (1 - \alpha) \cdot [\text{dist}(w_{f1}, w_{f2}) \\ & - \max(\text{dist}(w_i, w_{f1}), \text{dist}(w_i, w_{f2}))] \\ \{\text{dist}(w_{f1}, w_{f2}) = {} & \max(\text{dist}(w_k, w_l)), w_{l,k,j;f_1,f_2} \in W_j\}. \end{aligned} \tag{2}$$

This heuristic function considers boundary targets, w_{f1} and w_{f2} in W_j, which are the targets with the maximum distance value. The basic idea of this function is to forward the robot to one of these boundary targets since these targets determine the diameter of the region (W_j) and both of them should be visited. If the robot initially heads toward one of the boundary targets, the diameter (the longest path) can be traveled by visiting other targets along the path. A sample illustration of this cost function is given in Figure 3. In this figure, although t_2 is closer to r_1 than t_1, with the farthest addition cost (FAC) heuristic applied, t_1's cost value is smaller than that of t_2 ($3 < 3.6$), which results in a better route shown by the dashed arrows. The cost penalty applied to forward the robot to the boundary targets is limited to a small degree. By introducing a constant (α), this degree of direction can be adjusted. When α is assigned a value of 2/3, this heuristic function produces close to optimal results for the multirobot multitarget allocation domain [2]. If more than one pair of boundary targets exist, the pair that has a member at the smallest distance from the UUV is selected.

An illustrative example of the generation of the search regions (areas) and the traversed path patterns by the robots are depicted in Figure 4. Since there are three robots in this figure, three search regions are determined and covered by the robots.

As UUVs detect the MLOs on their way, they broadcast these estimated target positions to all AUVs (hence, tasks for crawlers are generated online during execution). Then, MLO information can propagate to all other AUVs in the group that can possibly be reached. Periodic broadcasting of important information (coming from either owned sensors or external agents) is a way to handle communication range limitations.

Identification of MLOs

When the crawlers are informed about the MLO locations, they update their world knowledge and dynamically select the best MLO targets to visit and propose auctions. Therefore, they can switch among tasks when new tasks appear if it is more profitable. It is also possible that a crawler may inadvertently discover a mine without being informed of its position by a UUV. In this case, the crawler identifies the target, adds it to its task list as an achieved task, and broadcasts achievement information to maintain the system consistency. Crawlers determine their bid values by using the cost functions proposed for the multirobot multitarget exploration domain [2].

In the identification task, when crawlers are within an area close to an MLO location, they begin keeping time while surveying the MLO location. Whenever the time limit is reached, they set the task status as achieved and broadcast this information. If a detection event occurs during this time period, the MLO location is considered to be an actual mine; otherwise, it is determined as a false alarm after deadline. In either case, the task is marked as achieved.

Experimental Results on the MCM Mission

The performance of our framework and the precaution routines is evaluated in ALWSE-MC. Three sample scenarios in the simulator are given to illustrate the performance of our framework for the naval MCM mission. The MCM mission movies are available online at [14]. UUVs are equipped with sensors capable of detecting mines within 30 ft from the skin of a target. However, they are not able to correctly identify them. The crawlers are equipped with cameras that can both detect and identify mines within 20 ft. None of the AUVs have predefined search patterns. UUVs have internal navigation errors; therefore, their estimated location values are

Naval MCMs are actions taken to counter the effectiveness of underwater mines.

different from actual locations in most cases. Two AUVs can communicate each other whenever the receiver AUV is in the sender AUV's transmitter range, within its transmitter beam width, and the sender AUV is within the transmitter AUV's receiver beam width.

All UUVs and crawlers begin execution from a deployment area. There is no a priori information about mine locations. Around 121 waypoint locations (environment size: 200 × 200) are known but are not assigned initially. UUVs begin negotiations and divide the overall mission area into three (known number of UUVs) regions. Since they are within the line of sight, they can communicate their location information. Therefore, initially defined regions are nearly the same for all UUVs. Figure 5 illustrates a successful mission scenario

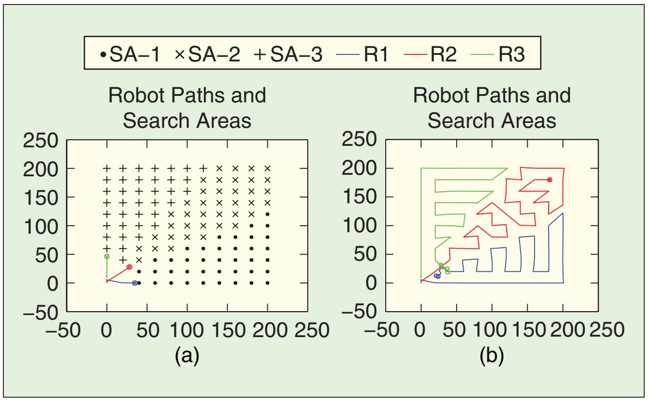

Figure 4. *(a) Mission execution begins. The overall area is divided into regions. (b) Robots patrol the area in the corresponding regions.*

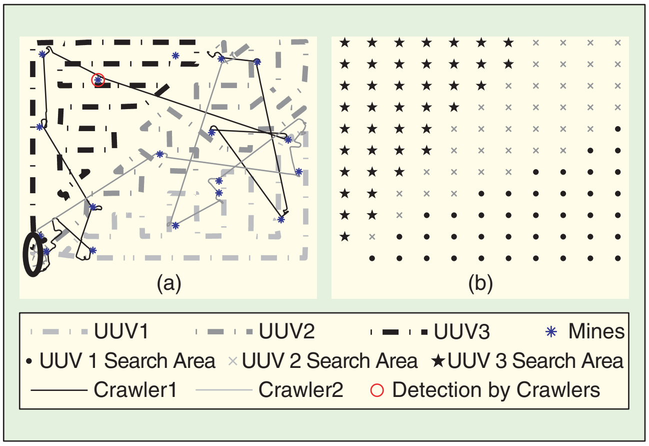

Figure 5. *Scenario 1: (a) The UUVs cover the area, and the crawlers visit the MLO locations. (b) The UUV regions are illustrated.*

DEMiR-CF is designed with the capability to deal with real-time situations.

with three UUVs and two crawlers. Allocations of waypoints after negotiations can be seen in Figure 5(b). Since there are no failures, waypoint assignments do not change during run time. However, the crawlers sometimes switch among tasks if they are not informed about tasks that are being executed, and sometimes parallel executions occur. Whenever they are in communication range, they can resolve the conflicts efficiently by means of the precaution routines. As shown in Figure 5(a), the crawlers can also detect mines without being informed (red circled in the figure). The routes of the crawlers may seem somewhat random. However, it should be noted that the tasks related to the MLO locations appear online during run time when they are discovered, and the communication range is limited.

In Scenario 2, UUV3 fails in the same setting of Scenario 1 (Figure 6, the location of the failure is indicated with a red arrow in the figure). Initial regions for all UUVs change after UUV3 fails [Figure 6(b)]. The other UUVs revise their region definitions and, after negotiations, they share the full area as indicated in the figure. The visited waypoints are not in their region coverage. Because of the uncertainties, some waypoints may remain uncovered in the schedules (indicated with the red diamond in the figure). However, this uncertainty-related problem is resolved by UUV2, and the mission is completed.

In the Scenario 3 (Figure 7), UUV3 fails and the other UUVs detect the failure and they negotiate the remaining unvisited waypoints and new schedules are determined as in Figure 7(b). While these UUVs execute their tasks, UUV4 is released from the deployment area. Detecting the arrival of a new UUV, the other UUVs change their region definitions accordingly [Figure 7(d)] and offer auctions for these areas. Initially UUV4 is not informed about the visited waypoints

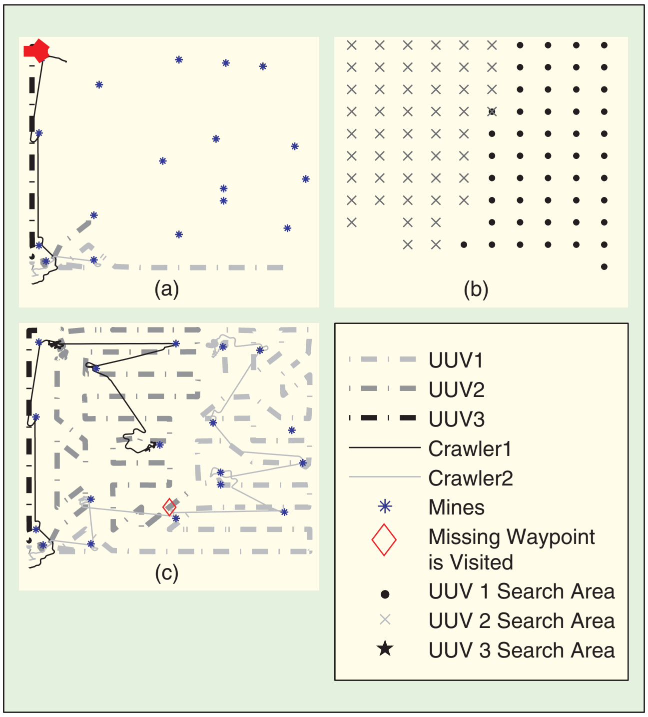

Figure 6. *Scenario 2: (a) Initially, all UUVs begin execution, UUV3 fails, and other UUVs take responsibility of all unvisited waypoints. (b) Region assignments are changed for UUV1–2 after detecting the failure. Because of an uncertainty, one waypoint is left uncovered. (c) UUV2 completes its region coverage task and adds the waypoint missing in (b) to its schedule after detecting that it is not visited.*

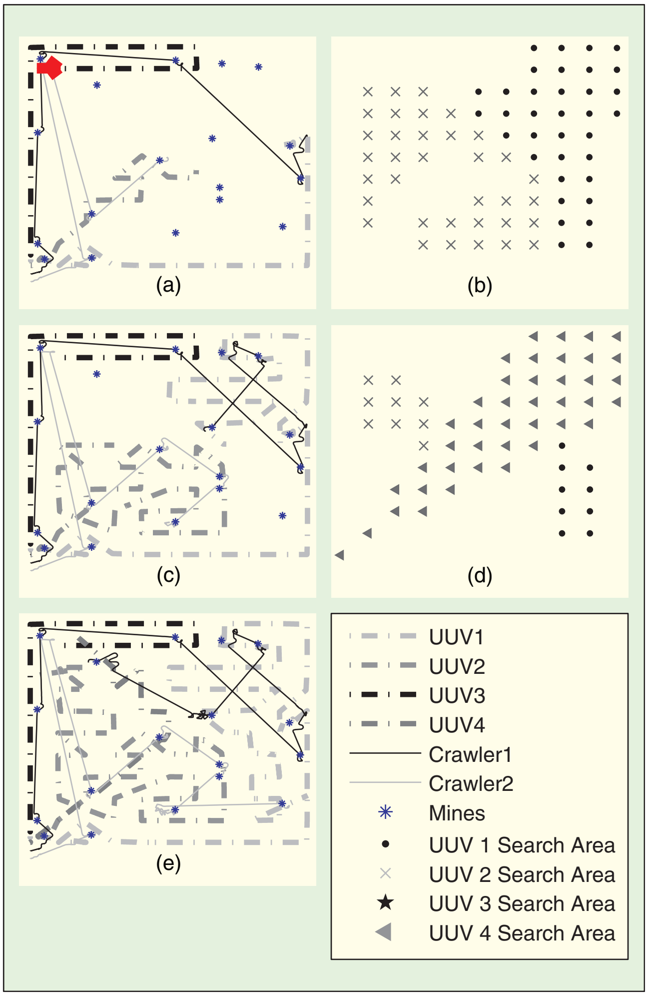

Figure 7. *Scenario 3: (a) UUV3 fails and other UUVs take responsibility of the waypoints initially assigned to UUV3. (b) Region assignments are changed for UUV1–2 after detecting the failure. (c) Another UUV(4) is released from the deployment area. (d) Schedules are changed accordingly after negotiations. However, UUV4 is not informed about visited waypoints and form regions by considering all waypoints. (e) After being informed about visited waypoints, UUV4 only visits unvisited waypoints in its schedule.*

and defines its regions with this incomplete knowledge. After negotiations, the regions are assigned and the schedules are formed. Entering into the communication range, UUV4 redefines its regions by considering incoming information for the visited waypoints.

In the same settings, another experiment is conducted to evaluate the message loss rate effects on the success of the completion of mission. Table 1 illustrates the results ($\mu|\sigma$) averaged over ten runs. When the message loss rate is different from 0, as expected, the mission completion time performance of the system degrades but linearly. It should be noted that, even for a rate of 0.75, the overall mission (M_C and M_I) by the final identification of the mines is completed. The average of the first visit times of the waypoints increases linearly because of the delays occurring by redundant visits of the targets. The number of waypoint (w) visits increases with high message loss rates. When the message loss rate is one, there is no communication among AUVs, and they cannot correctly reason about the region portions. Therefore, each UUV searches the full area completely. The crawlers detect and identify 12.8% of mines by their local detection in a small area (MLO target information cannot be communicated in this case). Since the identification mission is not complete, the overall mission is not completed. This table illustrates the performance of our framework against message losses. As a final remark, auction generation and clearing in an environment with communication delays desires special attention. Especially, auction deadlines should be determined by considering communication delays that may vary during the run. Plan B precautions can resolve these kinds of problems. Precautions for delayed messages on out-of-date situations prevent the system from getting stuck into further inconsistencies and deadlocks.

Further Extending MCM Mission to Prevent Hostile Attacks

The MCM mission can be further extended with the presence of possible threats from hostile vehicles. We analyze this situation in a dynamic simulation environment where the mission consists of the online tasks, whose generation times are not known in advance by the robots (AUVs). The overall mission is to search a predefined area as a part of the MCM mission and additionally protecting the deployment ship from any hostile intent [1].

The MCM mission is performed undersea and in real time.

Discussion and Conclusions

In this article, we presented the performance of a new framework, DEMiR-CF, in the context of a naval MCM mission in the realistic NAVY simulator ALWSE-MC. DEMiR-CF is a distributed framework for multirobot teams that integrates incremental task selection schemes, distributed allocation methods, and several precaution routines to handle failures and limitations of the real-world task execution. It maintains high solution quality with available resources. Precaution routines can respond to several failures as illustrated in the scenarios presented in this article. Evaluations reveal the high performance of DEMiR-CF on online task and situation handling. Since the framework is a single-item auction method, it can be used for environments with limited, delayed, or unreliable communication. In general, the framework is designed for more complex missions of interrelated tasks. We have implemented the DEMiR-CF framework on Khepera II real robots for the allocation of tasks of the multirobot multitarget exploration mission that can be treated as the classification tasks. Since the proposed approach is computationally cheap, its implementation on even very small robots has been possible, which makes the approach broadly applicable for different robot platforms. Accordingly, as the realistic simulation results reveal, limiting the assumptions in the design of the approach facilitates its porting to the real underwater vehicles. The naval MCM domain has appropriate characteristics to deploy teams of robots and let them cooperate to achieve the overall mission. It should be noted that the objectives and the limitations of this domain are similar to those of both search and rescue and space exploration domains. Therefore, we believe that research in this work can be useful for these domains as well.

Future work on the presented research includes considering the coverage and the detection strategies of the MCM mission together to improve the performance of the system. Especially, if the communication range is known a priori, this information can also be used in region determination and in constructing the paths of the robots to improve the responses of the system to robot failures.

Table 1. Performance results ($\mu|\sigma$) for different message loss rates.

	0		0.25		0.5		0.75		1	
Message Loss Rate	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
M_C completion (%)	100.0	0.0	100.0	0.0	100.0	0.0	100.0	0.0	100.0	0.0
M_I completion (%)	100.0	0.0	100.0	0.0	100.0	0.0	100.0	0.0	12.8	4.1
M_C completion (t)	3,349.4	60.5	3,683.2	167.1	4,909.0	430.1	5,141.2	938.1	6,304.2	139.0
M_I completion (t)	2,852.8	35.3	3,227.6	205.3	4,205.0	836.9	5,021.2	692.7	N/A	N/A
w first visit	1,380.1	6.1	1,390.0	16.3	1,922.0	92.8	2,256.6	334.5	2,936.0	104.5
w number of visits	1.0	0.0	1.0	0.0	1.01	0.01	1.09	0.04	3.0	0.0

Acknowledgment
We thank Jason R. Stack from Naval Surface Warfare Center, USA, for his constructive comments on this research.

Keywords

Multirobot cooperation, naval mine countermeasures, incremental task selection, robustness.

References

[1] S. Sariel, "An integrated planning, scheduling and execution framework for multi-robot cooperation and coordination," Ph.D. dissertation, Istanbul Technical Univ., Turkey, 2007.
[2] S. Sariel and T. Balch, "Efficient bids on task allocation for multi-robot exploration," in *Proc. 19th Int. FLAIRS Conf.*, 2006, pp. 116–121.
[3] S. Sariel, T. Balch, and N. Erdogan, "Incremental multi-robot task selection for resource constrained and interrelated tasks," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2007, pp. 2314–2319.
[4] U.S. Navy. (2008, Jan.). ALWSE. [Online]. Available: http://nswcpc.navsea.navy.mil/analysis/capabilities.asp
[5] S. Botelho and R. Alami, "M+: A scheme for multi-robot cooperation through negotiated task allocation and achievement," in *Proc. IEEE Intl. Conf. Robotics and Automation (ICRA)*, 1999, pp. 1234–1239.
[6] B. Gerkey and M. J. Mataric, "Sold!: Auction methods for multirobot coordination," *IEEE Trans. Robot. Automat.*, vol. 18, no. 5, pp. 758–768, 2002.
[7] M. B. Dias and A. Stentz, "Opportunistic optimization for market-based multirobot control," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems (IROS)*, 2002, pp. 2714–2720.
[8] T. Lemaire, R. Alami, and S. Lacroix, "A distributed task allocation scheme in multi-UAV context," in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2004, pp. 3622–3627.
[9] M. B. Dias, R. M. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," Robotics Inst., Carnegie Mellon Univ., Pittsburgh, PA, Tech. Rep. CMU-RI-TR-05-13, 2005.
[10] S. Sariel, T. Balch, and J. R. Stack, "Empirical evaluation of auction-based coordination of AUVs in a realistic simulated mine countermeasure task," in *Distributed Autonomous Robotic Systems (DARS)*. vol. 7, M. Gini and R. Voyles, Eds. New York: Springer-Verlag, 2006, pp. 197–206.
[11] J. R. Stack and R. C. Manning, "Increased autonomy and cooperation in multi-AUV naval mine countermeasures," in *Proc. Undersea Defence Technology Conf.*, 2004.
[12] G. J. Dobeck, "Image normalization using the serpentine forward-backward filter: Application to high-resolution sonar imagery and its impact on mine detection and classification," in *Proc. SPIE (Special Issue on Detection and Remediation Technologies for Mines and Minelike Targets X)*, 2005, vol. 5794, pp. 381–391.
[13] S. Sariel and T. Balch, "A distributed multi-robot cooperation framework for real time task achievement," in *Distributed Autonomous Robotic Systems (DARS)*. vol. 7, M. Gini and R. Voyles, Eds. New York: Springer-Verlag, 2006, pp. 187–196.
[14] S. Sariel. (2008, Jan.). MCMMovies. [Online]. Available: http://www2.itu.edu.tr/~sariel/videos/MCM-Movies.html

Sanem Sariel received the B.S. degree in computer and control engineering and M.Sc. degree in computer engineering from Istanbul Technical University (ITU) in 1999 and 2002, respectively. She was awarded a research scholarship for her Ph.D. studies in the United States in 2004, where she worked with Prof. Tucker Balch in the BORG Laboratory at Georgia Institute of Technology. She was jointly advised by Prof. Tucker Balch and Prof. Nadia Erdogan and received the Ph.D. degree in computer engineering from ITU in 2007. She is an instructor in the Computer Engineering Department of ITU, where she served as a research and teaching assistant from 1999 to 2007. Her research interests include distributed problem solving, multirobot coordination or cooperation, and intelligent agents.

Tucker Balch is director of the Institute for Personal Robots in Education (IPRE) and associate professor in the School of Interactive Computing at Georgia Tech. Tucker's robotics research interests include machine learning for robot navigation and large-scale multirobot systems. His recent work with IPRE is focused on understanding how robots can be used to help students learn more effectively. Balch received the B.S. and Ph.D. degrees in computer science from Georgia Tech in 1984 and 1998, respectively. He flew F-15s in the U.S. Air Force from 1988 to 1995. He has published more than 80 technical articles and two books.

Nadia Erdogan received her B.S. degree in electrical engineering and M.Sc. degree in computer science from Bosphorus University in Istanbul, Turkey, in 1978 and 1980, respectively. She received her Ph.D. degree in computer engineering from Istanbul Technical University, Turkey, in 1987, where she is now a professor in the Computer Engineering Department. Her current research areas include distributed computing and execution environments, mobile agent systems, and parallel programming.

Address for Correspondence: Sanem Sariel, Department of Computer Engineering, Istanbul Technical University, Turkey. Phone: +90 212 285 38 52. Fax: +90 212 285 36 72. E-mail: sariel@itu.edu.tr.