

Flood Disaster Mitigation: A Real-world Challenge Problem for Multi-Agent Unmanned Surface Vehicles

Paul Scerri, Balajee Kannan, Pras Velagapudi, Kate Macarthur, Peter Stone, Matthew E. Taylor, John Dolan, Alessandro Farinelli, Archie Chapman, Bernadine Dias, and George Kantor

Abstract. As we advance the state of technology for robotic systems, there is a need for defining complex real-world challenge problems for the multi-agent/robot community to address. A well-defined challenge problem can motivate researchers to aggressively address and overcome core domain challenges that might otherwise take years to solve. As the focus of multi-agent research shifts from the mature domains of UGV and UAVs to USVs, there is a need for outlining well-defined and realistic challenge problems. In this position paper, we define one such problem, *flood disaster mitigation*. The ability to respond quickly and effectively to disasters is essential to saving lives and limiting the scope of damage. The nature of floods dictates the need for a fleet of low-cost and small autonomous boats that can provide situational awareness (SA), damage assessment and deliver supplies before more traditional emergency response assets can access an affected area. In addition to addressing an essential need, the outlined application provides an interesting challenge problem for advancing fundamental research in multi-agent systems (MAS) specific to the USV domain. In this paper, we define a technical statement of this MAS challenge problem based and outline MAS specific technical constraints based on the associated real-world constraints. Core MAS sub-problems that must be solved for this application include coordination, control, human interaction, autonomy, task allocation, and communication. This problem provides a concrete and real-world MAS application that will bring together researchers with a diverse range of expertise to develop and implement the necessary algorithms and mechanisms.

Keywords: multi-agent systems, challenge, communication, autonomy, path-planning, coordination, task-allocation

1 Introduction

Robotics challenge problems like FIRST, DARPA Challenges, BotBall, MAGIC, etc [17, 24, 23, 7] have shown to be an effective motivational tool for invigorating robotics researchers at all levels, from high-school to experienced professionals, while solving real-world problems. Such challenges offer an incredible opportunity to shorten the time-cycle required to advance the state of the art in autonomous vehicle technology. The success of the DARPA challenges, Grand and

Urban, are a testament to this. Robots such as Stanley and Boss have become part of the robotics lore, while the developed technological solutions have become the backbone for translating commercially developed autonomous vehicles on our roads from a dream to reality. Furthermore, as we look to translate the developed technology and success over different applications, there is a need to define real-world challenges in alternate domains like USVs.

According to the United Nations, annual flooding currently impacts in excess of 500 million people, costs the world up to \$60 billion USD and the number of casualties exceed 20,000 in Asia alone [22]. Changes to the environment, such as mining, deforestation, and general industrialization are likely to worsen the problem worse over time. Unfortunately, flooding disasters disproportionately effect people in under-developed countries due to lack of early warning systems, flood control and emergency response infrastructure. The scope and application of the problem have far-ranging implications. Currently, despite large scale flooding disasters world over, in the immediate aftermath, victims are largely left to fend for themselves. The lack of relief aid is in part due to a limited knowledge of the affected areas and specific needs of the victims. Thus, any reliable solution, regardless of efficiency, will have immediate real-world benefits, while further research and development can increase the value of the system over time. In most cases, floods occur over large areas and over relatively long time-scales. Often their occurrence can be cyclical in nature and can be predicted well in advance, e.g., monsoonal or hurricane flooding can be expected annually during a well-defined season. We believe that the inherent properties and the scale of impact of a flooding disaster make it an ideal problem to be addressed with robot teams with multi-agent technology playing a central role. Finally, the cyclical nature of flooding means that it is likely that solutions can be iteratively evaluated and improved in real disaster environments over time. We believe that flood mitigation might be the *seminal* challenge for MAS because it is an important real-world problem for which MAS appears to be an ideal and essential technology.

In order to effectively address the problem, small, autonomous watercrafts are ideal for flood mitigation and response. Relative to other types of vehicles, watercraft are simple, robust and reliable. By keeping the vehicles small, most safety issues can be avoided simply by ensuring that if there is a collision, it can cause at most very minor damage. Influenced in part by field experiences of Murphy et al. [18], we believe unmanned surface vehicles (USV's), such as airboats, rather than unmanned underwater vehicles (UUVs), are better suited to this operational domain. Airboats are flat-bottomed boats that use an above-water fan to propel themselves forward safely and effectively through shallow or debris-filled water.

The challenge, then is to be able to construct and deploy small and capable airboats at a low cost. The low cost is particularly important for feasibly deploying sizable teams capable of covering large areas. In the immediate future, three or four cooperative boats might be deployed to provide some situational awareness over a small area, but as algorithms become more scalable, hundreds or thousands of vehicles could be used to provide detailed situational awareness

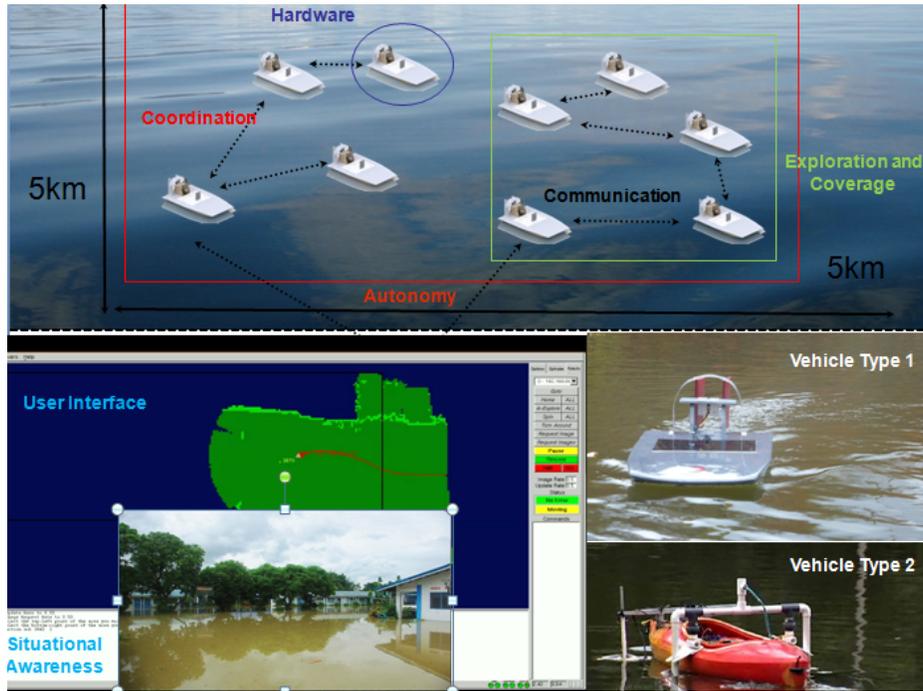


Fig. 1. Flood Disaster Mitigation Challenge

over dramatically larger areas. Longer term, a combination of autonomous aircraft and surface boats might be able to quickly cover a large area, being robust to obstacles and debris and able to safely navigate in places where large numbers of civilians are moving around.

2 Problem Definition

The contribution of this paper is not a description of algorithms for controlling boats for flood mitigation, but rather a technical description of the problem from the perspective of multi-agent systems. We present constraints descriptions for sub-areas of MAS so as to provide target problem descriptions for algorithm developers to overcome. Specifically, we pose the challenge as a sustained distributed situational monitoring problem in flood affected areas over a relative large area (25 sq km) using a large team (25-50) of autonomous watercrafts with minimal human oversight. This primary system objective leads naturally to a number of specific technical challenges that must be overcome to successfully complete the challenge (see Figure 1). Our challenges to the research community are:

Vehicle Design: *How to develop a robotic watercraft platform specific to the demands and requirements of flood disaster mitigation activities?*

Vehicle Intelligence: *How to robustly control a single boat and perform fundamental tasks of way-point following, collision avoidance, information gathering, energy management, payload management, etc?*

Team Intelligence: *How to develop efficient techniques for coordinating a team of airboats, identifying the scale of autonomy, situational awareness, establishing and maintaining communication network, path planning and task allocation?*

We need to bring together mature technologies from different areas including *autonomy, robotic watercraft, coordination, networking, fault-tolerance, coverage, exploration and human-robot interaction* towards building an integrated, large scale, autonomous system capable of monitoring and payload delivery in dynamic environments over an extended period of time.

3 Vehicle Design

For the principal task of providing situational awareness, identifying potential victims, and augmenting current first responder capabilities, any prototype design should address essential components pertaining to *cost, sensing, payload, power*. Relative to other types of vehicles, surface boats or *airboats* are simple, robust and reliable. Cognizant to the operating domain, we argue that the overall size and weight of the boats should be relatively small. By keeping the vehicles small, most safety issues can be avoided simply by ensuring that if there is a collision, then the relatively low operational speeds cause only minor damage. Specifically, the following issues must be considered:

- Sensors are a major component of the cost of the boat and are most likely to fail. Hence, there is significant advantage to be able to use simple and cheap sensors. However, this typically makes control more difficult. For effective operation, the vehicles should have four essential capabilities: *adequate range, communications, navigation, and environmental sensing*.
- The surface boats must be inexpensive and reliable so that we can build, maintain, and deploy a large numbers of them at a fraction of the cost of commercial alternates. The challenge is to find the right balance between using low-cost simple sensors and making the boat useful.
- Using large numbers of boats over extended periods will inevitably result in individual failures. The large number of hours that boats will be in the water in this project will provide data that quantifies the type, nature and rate of failures, providing key input to future development and for evaluating the effectiveness of the overall solution. Hence, the boats must be easy to construct, maintain and to repair.
- Power will be a limiting factor in boat performance, having a direct impact on attributes such as range and speed. The boats must be capable of traversing a set distance within a given amount of time in order to preserve the quality of the acquired information. Based on initial analysis of the problem domain, we estimate the operational range for exploration to be in the range of 5 km

over the course of several hours. This translates to average operating speeds of 2 – 3 km/hr, which confirms to operational safety standards.

- The boats must have some payload capability for dropping off essential supplies as well as picking up water samples for contamination analysis.

Vehicle navigation in the intended environment involves two basic challenges pertaining to obstacle avoidance, negotiating surface debris, and adapting to “waterscape” changes as flood waters rise or recede. Development of novel individual and joint control schemes to handle these requirements is therefore a key area of research. Control in an aquatic environment has been studied widely in larger surface vehicles, primarily at oceanic scales, and complex active control problems in semi-autonomous vehicles. In the former, a “follow-the-carrot” style approach is often sufficient, as the larger vehicle size coupled with high confidence measures in vehicle and obstacle position makes trajectory errors small. However, there are several aspects of the flood mitigation domain that warrant further research. The small size of the boat and unstructured nature of the environment mean that a-priori planning is limited and necessarily uncertain. It is impractical to precisely map and localize the boat with sufficient precision to avoid all obstacles, making reactive control a necessity for short-range obstacle avoidance. Second, while complete avoidance is necessary for safety at higher speeds, it is primarily a matter of efficiency at lower speeds. Low-speed collisions with obstacles often do not lead to damage or entanglement, and it is possible for boats to bump objects occasionally at low speeds. Finally, the high speed of the boats and water currents and eddies, relative to boat size, make vehicle dynamics significantly less predictable. Large discrepancies in trajectory may be effected by small changes in ambient current and vehicle hydrodynamics. Thus, building accurate motion models through physics alone may not be practical or possible.

In developing control strategies, it is important to note that given the nature of these vehicles, the objectives of the strategy are to minimize power consumption and travel time under the constraint of safe traversal of the environment. The uncertainty and complexity of realistic aquatic motion models make learning-based approaches particularly attractive. Methods such as reinforcement learning offer a number of advantages, including the ability to adapt to changing environmental conditions such as water currents, changes in payload, and the possibility of transferring learning between vehicles. While the environment is unstructured at a fine scale, it is evident that particular classes exist by virtue of human development. Areas such as towns provided a semi-structured but dense lattice where potential for human interaction is high, while rural areas or existing waterways will be less dense in terms of static obstacles, but potentially rich in manned vehicle traffic. In order to move efficiently in these and other classes of environment, it may be necessary to implement hybrid or layered control approaches that explicitly model the different areas.

4 Vehicle–Level Intelligence

Flood environments are uniquely cluttered, and for practical purposes, unknown a-priori. Situation awareness constitutes the collection of data that enable operators to better characterize the state of flooded areas for the prioritization of emergencies, the allocation of resources, and the establishment of further relief infrastructure. Information such as imagery of affected areas, traversability roadmaps or obstacle maps, and sampling of environmental factors such as water quality and temperature can play a role in creating this representation. The on-board sensor package must ensure vehicle safety, and, moreover, provide useful situation awareness to human operators to facilitate rescue and response activities. Safe traversal of the environment requires ego-motion estimation and obstacle detection. The former problem is made challenging by the small scale of the vehicles and obstacles, and the chaotic nature of water currents. Unlike ground robots, aquatic surface vehicles drift with prevailing currents that are hard to measure and model, resulting in motion that cannot easily be predicted. Absolute positioning using GPS can significantly assist in correcting for this, but commercial receivers at the price point of the vehicle cannot resolve to the require accuracy. In addition, bearing cannot be directly estimated from these measurements, and is also subject to drift. Obstacle detection requires sensors to look out over the surface of the water for potential hazards, both above and at the surface. The distance at which this can be done safely is a critical factor in determining the maximum rate of travel of the vehicle.

The sensory constraints can be overcome by novel estimation and filtering methods that will enable useful streams of information to address these three tasks while incorporating only data provided from relatively low-cost sensors. Such sensors include many of the staples of ground robotics, including MEMS accelerometers, gyros, and magnetometers, physical contact switches, local optic flow sensors, in-air IR and sonar range finders, and stereo or monocular vision. Interestingly, many of these sensors are already present in modern mobile devices such as smart phones, making these embedded platforms even more attractive as a low-cost, integrated solution to consider.

A key area of study for the MAS community lies in using vehicles jointly to improve perception of local and environmental features. While boats cannot control many aspects of the environment, they can (a) exchange information that allows other boats to reduce uncertainty or correlate features between vehicles and (b) use other boats for relative localization through direct or indirect relative tracking (i.e. boat-mounted fiducials). Since large teams of these vehicles are expected, this may be a very powerful alternative to more expensive local sensing strategies. Successful joint sensing strategies can bring down the cost of the vehicles by requiring less of individual sensors.

5 Team–Level Intelligence

This section outlines the key sub-problems that arise at the level of the team – the boats and the human controllers. We break the challenge problem into

six key areas, each of which is an important problem in its own right and has already received significant attention from the MAS research community. These are:

1. *Autonomy and human interface,*
2. *Situational awareness,*
3. *Communications & networking,*
4. *Path planning,*
5. *Team planning,*
6. *Task allocation.*

The Airboats Challenge brings with it the domain specific constraints of operating in the midst of a flood disaster and the added complexity of integrating the sub-problems to produce a coherent system. This is a key reason why Airboats Challenge is so appealing as a challenge problem: although big, the global MAS problem is relatively modular, and does not need to be treated as a whole. Instead, each of the sub-problems we have identified can be tackled by separately, and this paper is written in the hope that many separate research groups with different interest and expertise will all be able to contribute.

5.1 Autonomy and Human Interface

One of the most interesting research challenges for the Airboats Challenge is to design a system that will be able to provide the appropriate *level* of autonomy for the agents. A set of exemplar tasks follow to demonstrate the necessary types of autonomy, as well as identifying the underlying research challenges imposed by these behaviors.

Task 1: Autonomous exploration. Prior knowledge about the terrain might be useful for identifying potentially traversable areas. However, realistically in such a fluid environment, many areas may become (un)traversable by the boats, necessitating alternate agent behavior. Exploration will be essential towards updating the local and, subject to communication constraints, the global maps. As a consequence the team, on an individual and sub-team levels, must be capable of performing autonomous exploration of its surrounding environments. Interestingly the autonomy level for exploration could be varied from fully-autonomous to tele-operated on a case-to-case basis.

Task 2: Human interaction. The agents must be capable of detecting and interacting with civilians towards providing accurate situational awareness to the responders. To this end, agents should be able to integrate information coming from the on-board sensors (e.g. cameras and microphones) for detecting survivors. Civilians should then be approached to provide information to the base station about their position, photographs of the area, voice transmissions, etc. This throws up very interesting challenges in identifying the most effective manner of interaction with humans. Moreover, the airboats must be able to autonomously identify adversarial behaviors. Potential countermeasures include sending a “SOS” signal, identifying potentially hostile people by taking pictures for later identification, etc.

Task 3: Self awareness. The unpredictability of communication range dictates that the agents be capable of autonomously returning to the base station or to dynamically determine an alternate rendezvous point. The ability to return to the base station would be necessary for enabling agents to return with collected information in the absence of communication connectivity. Direct communication to base station may be infeasible due to limited network structure and the agents will likely need to coordinate to construct a network infrastructure. This might require autonomous task-switching on a sub-team level. Furthermore, in Airboats Challenge, operations will be carried out over a long time scale; as a consequence the agents should be self aware and capable of switching states to operate for days, up to approximately two weeks. Agents will need to return to the base station for recharging, resupply, repair, etc.

Task 4: Situational awareness. The agents should be able to recognize dangerous situations and activate specific behavior. For example, the agents should be able to recognize when the agent is going to hit an obstacle, or when the agent's localization has failed. Actions to address such situation may include sending broadcasting an alarm signal, quickly changing direction, or activating a search routine. It is important to note that "recognizing a situation" here refers to the agents ability to reason about abstract concepts such as "the agent is in trouble" and "the agent is having difficulty navigating to point X". Moreover the agent must be able to identify these crucial situations quickly, reliably and using inexpensive sensors.

While the above autonomous behaviors have been broadly addressed in the autonomous agents and robotics community, we believe that deploying a systems which is able to perform this kind of autonomous tasks in Airboats Challenge poses many interesting research challenges. Among the many issues which prevents direct application of off-the-shelf solutions in this domain, the most important address are: (a) algorithms must work in real time with low-power devices and will be unlikely to find guaranteed optimal solutions; (b) humans operators not experts in controlling nor repairing the agents – any proposed solution must be very simple and reliable; (c) agents should use adjustable autonomy: when human operators can provide help, the agents should try to take advantage of their expertise; (d) agents must work in a broad range of non-optimal conditions: for instance, if the weather or lighting conditions change, the agents should continue to work, even if it reduces their efficacy.

5.2 Situational Awareness

Providing situational awareness (SA) for human operators is a primary goal of the system. The task involves collecting information about the environment and getting it to the operators to allow them to understand the disaster that they are dealing with. SA also has a role to play in informing the networking, path planning and task allocation problems of the Airboats Challenge. Specifically, SA is used to put constraints on the set of feasible paths, and consequently network configurations, that the agents can take, and also to assign levels of importance to different tasks. From a MAS perspective, we are only interested

in the task of collecting information and communicating it back to the human operators and among the agents members; We are not interested in the important human factors issues related to its presentation.

The rate of change of different parts of the environment will be very different, with some requiring new information be collected and transmitted regularly and others only requiring an occasional visit for new information. Primarily, we anticipate that cameras will provide the majority of the data for getting situational awareness. Often still imagery will be sufficient, however under certain situations video data might be essential to do situational awareness, for e.g. the rate of water movement. However, other sensors such as microphones or wind-gauges might provide useful information. In flood disasters where water sits for long periods of time, it may be necessary to collect water samples to allow for checking for diseases. This would necessitate boats bringing samples all the way back to operators for analysis.

Two additional factors make the SA problem more complex from the MAS perspective. First, not all areas are as important as others. For example, areas that are likely to have high population density or will be critical for moving humans around the environment are more important than open areas where humans are not expected. The relative importance of different areas maps to preferences on locations to visit. Second, incoming data might be ambiguous or unclear and humans might request clarification in the form of additional information about an area. Both of these factors link SA to task allocation, in we expect that SA information will be used to identify tasks of high importance or value.

The problem can be formalized as follows. Consider the world to be made up of a set of locations, $L = \{l_1, \dots, l_n\}$. For each $l_i \in L$, a cost function $C_{l_i}(t) \rightarrow \mathcal{R}$ defines the value of not getting information on that location for a length of time t . Each time the location is visited, the function resets and might change. For example, areas found to have nothing of interest, will reset to a function that increases very slowly over time, while areas with a lot of interest will reset to functions that increase very rapidly over time. The system will not know in advance how the function will reset after it is visited, but we assume it will know as soon as it is visited. Another function, $V_{l_i}(t) \rightarrow \mathcal{R}$ gives the relative value of that location over time. The overall optimization is to minimize the cost of not seeing locations multiplied by the value of the location over time. That is,

$$\min \sum_{t=t_s}^{t=t_e} \sum_{l_i \in L} C_{l_i}(t) V_{l_i}(t)$$

where we assume time is discretized and t_s and t_e represent the start and end of the mission respectively.

5.3 Communications & Networking

The networking aspect of the flood mitigation problem is working out how to configure the boats to form an ad hoc network, in addition to making use of

any available infrastructure such as cell phone networks, to allow communication among the team and human operators. This component of the Airboats Challenge is of fundamental importance, since without a functioning communication network, the other team-level sub-problems – the situational awareness, path planning and task allocation capacities – of the system will be severely curtailed. Furthermore, in addition to running the algorithms that address these problems, maintaining a communications network itself places hard constraints on the solutions to the path planning and task allocation problems.

Regarding the SA goals of the system, we anticipate that the boats will be collecting a lot of potentially useful information, and will benefit from tight coordination with other boats when possible. However, this is likely to lead to there being far more data than communication bandwidth. The physical locations of the boats will create the physical network, hence the networking challenge is fundamentally to work out how the boats should move get the “best” network structure. Clearly, the positions of the boats cannot be dictated solely by the requirements of the networking, since this will impede their ability to do their primary task. However, it may be possible or necessary to dedicate some boats solely to the task of being network routers. Low-level issues of how to efficiently communicate data or to create more powerful transmitters are considered beyond the scope of the multi-agent problem. Similarly, it is anticipated to be the case that energy use for receiving and sending data is negligible versus energy costs of moving the boats around.

One network concept sometimes used in environments without wireless infrastructure is the idea of *delayed communication*, where robots will hold onto information and actively plan to get back to a location to transmit that information at some later time. Delayed communication is likely to be a useful mechanism in the flood mitigation problem, especially since delays on the order of minutes are unlikely to be important.

The wireless network connecting the boats and ground stations is required for sending three types of messages. First, messages are required to get information from the robot sensing the information to the boat or human who can utilize that information. Second, messages are required to facilitate coordination between the boats. Third, messages are required for human override of autonomous action, e.g. tele-operation. Appropriately designed coordination algorithms should mean that no particular message is absolutely critical to overall operation, instead each message will have some value to the team. Messages should only be delivered once and may pass through intermediate nodes to get to their destination. Naturally, there will be some time before which a message has no value and often a time after which a message has no value, e.g. information has become stale or opportunity for coordination has passed.

The movement of the boats around the environment and the availability of infrastructure, e.g. mobile phone towers, induces a network that changes over time. Because small vehicles moving in a complex environment, carefully placed mobile phone towers and human operators will have dramatically different communications equipment, it is not reasonable to assume that links are symmetric.

There will be constraints on edges in the network which restrict traffic on that edge. For example, a boat may have links to four other boats, but it cannot communicate with them at the same time, since the same wireless medium is being used for each link. More complex models might include constraints that capture interference between links degrading capacity, e.g. two different boats cannot broadcast on the same channel at the same time, but we believe these details are practically unimportant for this domain.

The network aspect of this problem is focused on providing the infrastructure to allow message delivery, other parts of an overall system will actually determine which messages are delivered. Thus, we have to think about the problem of optimizing the network structure as one of optimizing the potential for message delivery. This optimization must include the possibility that messages fail to be delivered and that the coordination is inefficient.

5.4 Path Planning

Path planning sits at the interface of vehicle- and team-level intelligence. For example, some path plans can be generated independent of other agents, such as return routes to a base station, while others require tight coordination of the actions of several agents, as when network connectivity requirements are paramount. Furthermore, solutions to a path planning problem may be constrained by environmental conditions (garnered from SA), network considerations, task requirements and vehicle power constraints.

The path planning component for a boat will be impacted by all other parts of the system, e.g. the networking component will tell it how it must move to maintain an appropriate network and the task allocation component will tell it what it must achieve in the environment. In an ideal solution, feedback from a path planner would impact other parts of the system, e.g. by indicating that it is expected to take the boat a long time to perform a particular task, hence it is better allocated elsewhere.

To generate even an independent path, the boat must deal with partial observability, because the environment is not perfectly known and action uncertainty since movement through the environment is inherently uncertain. Given multiple tasks, e.g. places to take observations or deliver supplies, the robot must appropriately order its tasks for best overall performance. It must also carefully balance risks, e.g. taking unknown but potentially more direct routes or moving at higher speed, time to complete time-sensitive tasks and the need to keep the boat intact for future efforts. The environment will not be completely static, making it necessary for the path planner to reason intelligently about the impact of any possible obstacles in advance and planning around them when they occur. Planning will need to occur over significant amounts of time, since boats may travel to tasks that take on the order of hours to reach. Therefore, in its most general formulation the single vehicle path planning problems can be considered as a Partially Observable Markov Decision Process (POMDP).

The path planning will be mostly individual but cooperation could dramatically improve overall performance. For example, if it is not known whether a

particular route is traversable, it may be optimal for one boat to first go down that street while others wait or take longer, safer routes. This type of exploit versus explore tradeoff is often studied in the literature, but not in the context of such complex individual planning. Cooperation will also be required to avoid hindering progress of other boats, e.g. impeding progress down a narrow alley. A natural and general framework for the multi-agent path planning problem is that of Decentralized POMDP. However, Dec-POMDPs are known to be intractable in general settings [1]. Therefore, a main research issue here is to find alternative formalizations or approximate techniques that can provide good solutions while meeting the real time constraints of the application.

Finally, while we anticipate that the primary focus of the path planner will be coming up with a path that achieves all the objectives of the boat at a minimal cost, some attention will need to be paid to actually being able to move the boat around the environment. While we consider issues of control outside of the scope of the MAS problem, environment features such as currents in the water, winds and narrow passageways will significantly effect what the boat can achieve and should be considered as a part of the path planning process. For example, in an area expected to have significant currents, it is not reasonable to plan or expect a fast path directly across or against the current.

5.5 Team Planning

Disaster response domains, like the one discussed in this paper, typically involve multiple sub-teams of agents working together towards achieving a common goal, saving lives and disaster mitigation. Each team-member has specific capabilities particularly suited to certain tasks. While some tasks are independent of each other, other tasks may be related by different constraints. As agents move about the environment, they have a direct influence on other team members from tightly-coupled scenarios [2, 8, 19] where multiple agents are required to complete a task, to loosely-coupled ones where the action of one agent might block the movement of others [6]. Team planning addresses the problem of decomposing a high-level set of goals into smaller independent, primitive tasks.

5.6 Task Allocation

Task allocation impacts the performance efficiency of teams in significant ways. Allocating vehicles to different tasks in an efficient and effective way is a crucial issue for the Airboats Challenge . More than any other sub-problem, task allocation connects together the components of the team-level intelligence of the system: The set of tasks may represent both SA and networking goals; The cost and benefits of completing tasks are computed using outputs from path planning and SA problems, and may be constrained by network considerations; and, we expect there to be human oversight of the weights attributed to tasks.

Task allocation is a very well known and widely studied problem in MAS, and many solutions have now been proposed, however, in the Airboats Challenge scenario, the task allocation problem is particularly challenging as the system

is composed of a large number of vehicles that will be equipped with cheap and low power devices and will have to coordinate in a highly dynamic and partially unknown environment.

Task allocation is usually formalized considering a set of tasks $\mathcal{T} = \{T_1, \dots, T_m\}$, a set of agents $\mathcal{A} = \{A_1, \dots, A_n\}$ and a reward matrix $R = \{r_{ij}\}$ where r_{ij} indicates the reward achieved by the system when agent A_i execute task T_j . An allocation matrix $A = \{a_{ij}\}$ defines the allocation of agents to task with $a_{ij} \in \{0, 1\}$ and $a_{ij} = 1$ if agent A_i is allocated to task T_j . The goal of the system is then to find

$$\arg \max_A \sum_{i=1}^{|\mathcal{A}|} \sum_{j=1}^{|\mathcal{T}|} r_{ij} a_{ij}$$

Moreover, a set of constraints \mathcal{C} usually describes valid allocations of agents to tasks, for example, one task could be executed at most by one agent or exactly k agents, or completing a task could be outright infeasible because of constraints on the actions of an agent. Therefore the above optimization must be performed subject to \mathcal{C} .

A first important challenge for the task allocation approach is to deal with a dynamic environment, where tasks appear, disappear and the reward to execute them may change during the mission execution: in the Airboats Challenge domain, vehicles will deal with tasks such as searching for civilian in a predefined area, approaching a group of detected civilians, collaborating with a set of other vehicles to relay information to the base station and so forth. These tasks are not known before hand and will be discovered during the mission; in addition, failures of vehicles should be taken into account: vehicles could be potentially stolen or the communication infrastructure could experience temporary break down. Hence, the above problem formulation must take time into account and one way to express this is to have that agents, tasks, reward matrix and consequently allocation matrix dependent on time and then find a series of allocation, one for each time step, such that the sum of reward over time is maximized:

$$\arg \max_{\{A^{t_s}, \dots, A^{t_e}\}} \sum_{t=t_s}^{t_e} \sum_{i=1}^{|\mathcal{A}^t|} \sum_{j=1}^{|\mathcal{T}^t|} r_{ij}^t a_{ij}^t$$

Therefore, the solution algorithm should be capable of continuously monitoring the environment and adapt the task allocation solutions to unexpected changes.

Second, in the Airboats Challenge vehicles should be able to take decision on their own, without necessarily relaying on information, or directives, from the base station; moreover, such decentralized task allocation approach must be designed to run on low power, cheap devices (such as smart phones). The low cost devices combined with the large scale operational domain, eliminates the use of intense computation and communication resources, typical of complete

algorithms¹, as their coordination overhead (computation time, message number and size), would be simply unacceptable in this scenario.

Third, since vehicles act in the real world without a complete knowledge of the environment, the benefit that the whole system would acquire for a given allocation of tasks is very hard to predict: vehicles are uncertain of their action outcome (e.g. a boat might be stalled while traveling towards an interesting area) and, more important, even if a task is completed successfully the reward for the team might be different than what is expected (e.g. it could be very hard to decide which group of civilians is more in need of help without having accurate information about their situation).

Finally, vehicles might need to form coalitions to execute tasks. Consider the example where agents might need to form sub-teams to approach a group of civilians while maintaining connectivity with the base station, or to search a given area of the environment where there is a high chance of discovering civilians. In our formulation coalition effects can be expressed by representing rewards as a set of functions instead of as a matrix: $R^t = \{r_j^t(a_{1j}^t, \dots, a_{|\mathcal{A}|^t j}^t)\}$ and considering the following objective function²:

$$\arg \max_{\{A^{t_s}, \dots, A^{t_e}\}} \sum_{t=t_s}^{t_e} \sum_{j=1}^{|\mathcal{T}^t|} r_j^t(a_{1j}^t, \dots, a_{|\mathcal{A}|^t j}^t)$$

Coalition formation is known to be a very hard problem to solve and current solutions can find optimal coalitions only for relatively small number of agents (in the order of 30) [20], so there is a clear need for approximate solutions in this context.

As mentioned above, there exists many potential approaches to address our task allocation problem, that range from approximate DCOP solution techniques [21, 5, 11, 4], to decomposing the problem as mixed integer linear programming problems [12, 15], market based approaches [9, 16], hybrid approaches [14, 13], etc and that have been used in similar application domains. Despite the fairly rich suite of algorithms for addressing team planning, the dynamic and complex environments, continuous configuration and observation spaces, and relative large team sizes coupled with limited computing and sensing far exceed the complexity handled by many existing approaches. Deciding how to represent the problem and determining which classes of algorithms are effective remains an open area of research.

6 Discussion

In order for the problem to be accepted in the MAS community as an open challenge problem, a case-study and subsequent feasibility analysis of the various de-

¹ With complete algorithms we are guaranteed to find the optimal solution

² If we aim to solve this problem using linear programming techniques we need to represent the reward for each possible coalition, this results in a combinatorial element in the complexity of the problem.

scribed components including vehicle design, intelligence and team-intelligence for the outlined problem must be performed. The feasibility analysis would allow us to identify system bias and weight associated with individual components as it affects overall system performance. This subsequently would allow us to formulate the challenge as a mathematical problem that can then be modeled for a simulator or real-world system. Genuine practical success may require that a modular open source architecture is developed into which various algorithms can be inserted. The development of such an architecture would also separate the hardware development from the software development and allow for researchers to collaborate and focus on specific domain. As part of the development framework, we are working on building a realistic simulator for the project as well as developing a prototype vehicle model. The simulator is intended as an open source resource that will allow the community to test and evaluate individual component algorithms as well as a full-system model on a common platform.

Furthermore, the feasibility study will also address an important component of any multi-robot system, *evaluation metrics*. The mission critical nature of the operating domain dictates the need for a high operating efficiency for the Airboats Challenge. In order to objectively evaluate operational performance, there is a need to have a well-defined and detailed set of metrics. Based on observations from earlier work in developing metrics for multi-robot teams [10], we believe that for the challenge problem the success metric should be a combination of qualitative and quantitative measures that can be used to analyze, evaluate, and subsequently improve performance of a team of airboats towards the overall goal of mitigating disasters during flooding. The goal therefore is to identify a set of flexible tools for researchers to use for in-depth system analysis. In addition, it is important to identify evaluation criteria that can help determine the quality of a metric in terms of the domain specific constraints, comprehensive understanding, construct validity, statistical efficiency, and measurement technique efficiency [3]. The idea of identifying generalizable classes allows researchers to independently evaluate specific sub-problems that constitute the challenge.

Finally, the unfortunate prevalence of floods will give many opportunities for solutions to be field tested, requirements to be updated and new designs to be explored. Beyond constrained environment testing, real world evaluation in places like the Philippines are essential for extended evaluation.

7 Conclusions

In this position paper, we present a challenge problem of using cooperative airboats to perform flood disaster mitigation. Floods are the natural disaster with the biggest annual impact and disproportionately affect the economically backward. We have outlined the key technical challenges and argued that the research from the MAS community is well suited to tackle many of the technologies that are necessary to develop a low-cost, high-impact solution. We are currently developing prototype simulators and robots to work on this problem and anticipate initial testing to occur in the near future. We plan to make the

simulation environment open for anyone in the community to test and contribute algorithms. It is also planned to make it possible for anyone in the community to provide code for key MAS functions on the robots themselves. This will provide both a realistic and important test for the algorithms and allow the MAS community to make a genuine contribution to the world.

References

1. Daniel S. Bernstein, Shlomo Zilberstein, and Neil Immerman. The complexity of decentralized control of markov decision processes. In *Proc. of UAI-2000*, pages 32–37, 2000.
2. M Bernardine Dias, R. Zlot, N. Kalra, and A. Stentz. Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE*, 94(7):1257–1270, July. 2006.
3. B. Donmez, P. Pina, and M. L. Cummings. Evaluation criteria for human-automation performance metrics. In *Proceedings of Performance Metrics for Intelligent Systems Workshop*, 2008.
4. A. Farinelli, A. Rogers, A. Petcu, and N. R. Jennings. Decentralised coordination of low-power embedded devices using the max-sum algorithm. In *Seventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-08)*, pages 639–646, May 2008.
5. S. Fitzpatrick and L. Meetrens. *Distributed Sensor Networks A multiagent perspective*, chapter Distributed Coordination through Anarchic Optimization, pages 257–293. Kluwer Academic, 2003.
6. Owen Holland and Chris Melhuish. Stigmergy, self-organization, and sorting in collective robotics. *Artif. Life*, 5(2):173–202, 1999.
7. Geoffrey Hollinger and Sanjiv Singh. Towards experimental analysis of challenge scenarios in robotics. In *12th International Symposium on Experimental Robotics*, December 2010.
8. Dapeng Jiang, Yongjie Pang, and Zaibai Qin. Coordinated control of multiple autonomous underwater vehicle system. *Intelligent Control and Automation (WCICA), 2010 8th World Congress on*, pages 4901–4906, July. 2010.
9. E. Jones, M Bernardine Dias, and A. Stentz. Learning-enhanced market-based task allocation for oversubscribed domains. In *International Conference on Intelligent Robots and Systems, 2007. IROS 2007.*, November 2007.
10. B. Kannan and L. E. Parker. Metrics for quantifying system performance in intelligent, fault-tolerant multi-robot teams. *International Conference on Intelligent Robotics and Systems*, November 2007.
11. C. Kiekintveld, Z. Yin, A. Kumar, and M. Tambe. Asynchronous algorithms for approximate distributed constraint optimization with quality bounds. In *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-10)*, pages 133–140, Toronto, ON, Canada, 2010.
12. Mary Koes, Illah Nourbakhsh, and Katia Sycara. Heterogeneous multirobot coordination with spatial and temporal constraints. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI)*, pages 1292–1297. AAAI Press, June 2005.
13. Mary Koes, Katia Sycara, and Illah Nourbakhsh. A constraint optimization framework for fractured robot teams. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 491–493, New York, NY, USA, 2006. ACM.

14. G. Ayorkor Korsah, Balajee Kannan, Imran Fanaswala, and M Bernardine Dias. Improving market-based task allocation with optimal seed scheduling. *Intelligent Autonomous Systems 11 (IAS-11)*, 0:249 – 259, August 2010.
15. G. Ayorkor Korsah, A. Stentz, M Bernardine Dias, and Imran Fanaswala. Optimal vehicle routing and scheduling with precedence constraints and location choice. In *ICRA 2010 Workshop on Intelligent Transportation Systems*, May 2010.
16. M. Lagoudakis, E. Markakis, D. Kempe, P. Keskinocak, A. Kleywegt, S. Koenig, C. Tovey, A. Meyerson, and S. Jain. Auction-based multi-robot routing. In *Proceedings of the International Conference on Robotics: Science and Systems*, pages 343–350, 2005.
17. David P. Miller, Charles Winton, and Jerry Weinberg. Beyond botball. In *AAAI Spring Symposium, Robots and Robot Venues: Resources for AI Education*, 2007.
18. Robin R. Murphy, Eric Steimle, Chandler Griffin, Charlie Cullins, Mike Hall, and Kevin Pratt. Cooperative use of unmanned sea surface and micro aerial vehicles at hurricane wilma. *Journal of Field Robotics*, 25(3):164–180, 2008.
19. Lynne E. Parker and Andrew Howard. Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection. *International Journal of Robotics Research*, 25:431–447, 2006.
20. T. Rahwan, S. Ramchurn, N. Jennings, and A. Giovannucci. An anytime algorithm for optimal coalition structure generation. *Journal of Artificial Intelligence Research (JAIR)*, 34:521–567, April 2009.
21. Paul Scerri, Alessandro Farinelli, Steven Okamoto, and Milind Tambe. Allocating tasks in extreme teams. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 727–734, 2005.
22. United Nations University. *Two Billion People Vulnerable to Floods by 2050; Number Expected to Double or More in Two Generations Due to Climate Change, Deforestation, Rising Seas, Population Growth*. United Nations University, 2004.
23. Christopher Urmson, Joshua Anhalt, Hong Bae, J. Andrew (Drew) Bagnell, Christopher Baker, Robert E. Bittner, Thomas Brown, M. N. Clark, Michael Darms, Daniel Demitrish, John M. Dolan, David Duggins, David Ferguson, Tugrul Galatali, Christopher M. Geyer, Michele Gittleman, Sam Harbaugh, Martial Hebert, Thomas Howard, Sascha Kolski, Maxim Likhachev, Bakhtiar Litkouhi, Alonzo Kelly, Matthew McNaughton, Nick Miller, Jim Nickolaou, Kevin Peterson, Brian Pilnick, Raj Rajkumar, Paul Rybski, Varsha Sadekar, Bryan Salesky, Young-Woo Seo, Sanjiv Singh, Jarrod M. Snider, Joshua C. Struble, Anthony (Tony) Stentz, Michael Taylor, William (Red) L. Whittaker, Ziv Wolkowicki, Wende Zhang, and Jason Ziglar. Autonomous driving in urban environments: Boss and the urban challenge. *Journal of Field Robotics Special Issue on the 2007 DARPA Urban Challenge, Part I*, 25(1):425–466, June 2008.
24. Christopher Urmson, Joshua Anhalt, Daniel Bartz, Michael Clark, Tugrul Galatali, Alexander Gutierrez, Sam Harbaugh, Joshua Johnston, Hiroki Kato, Phillip L. Koon, William Messner, Nick Miller, Aaron Mosher, Kevin Peterson, Charlie Ragnusa, David Ray, Bryon K. Smith, Jarrod M. Snider, Spencer Spiker, Joshua C. Struble, Jason Ziglar, and William (Red) L. Whittaker. A robust approach to high-speed navigation for unrehearsed desert terrain. *Journal of Field Robotics*, 23(1):467–508, August 2006.