Unmanned Vehicle Collaboration Research Environment for Maritime Search and Rescue

William Roberts*, Kelly Griendling*, Anthony Gray*, Dimitri N. Mavris*
***Aerospace Systems Design Laboratory (ASDL)**
School of Aerospace Engineering
Georgia Institute of Technology, Atlanta, GA, 30332-0150, U.S.A.

Keywords: *Unmanned Vehicles, System of Systems, Modeling and Simulation, Search and Rescue*

Abstract

With the proliferation of unmanned air vehicles in the last few decades, there has been a series of applications proposed around the concept of collaborative unmanned systems. However, it is difficult to assess at what point collaboration actually results in true benefits to mission performance and also difficult to determine the degree of benefit possible to achieve. This paper describes a framework which allows decision makers to analyze the effectiveness of a coordinated group of UAVs as a function of the physical mission parameters, showing benefit over a single UAV or uncoordinated multi-aircraft approach. To demonstrate the use of the framework, a case study is presented for a maritime search and rescue mission on a parameterized mission space. From the analysis, critical points in the mission space were found where the coordinated group of UAVs brought no additional benefits to the mission. However, at other points they were able to cut down the mission time by more than one hour. This framework can be used to guide decision makers and help estimate the effectiveness of coordinated UAVs in a parameterized mission space.

1 Introduction

Analysis of aircraft operations has become an increasingly difficult task as technology has developed in recent decades. Single aircraft are now often viewed as a single component of a complex system of systems (SoS) which cannot be well understood from simple analysis. A primary challenge facing decision-makers is the ability to predict the behavior and effectiveness of a SoS in the complex mission space facing operators on a daily basis. This is especially true in the case of unmanned systems, where collaborative tactics have been proposed for a number of missions. However, it is not necessarily clear that collaboration is necessarily better in all cases. While there are many promises of the performance benefits of collaborative behavior, these potential benefits come at the cost of increased overall complexity and risk. Increased reliance on the ability of vehicles to share information can lead to significant lapses in performance in the event of communications degradation. Software complexity for systems that need to create a shared information picture and make decisions based on this shared picture is generally higher than for systems operating in isolation. Furthermore, there is an increased probability of emergent behavior from these systems, making them more difficult to test and certify.

Therefore, a framework is needed which would allow decision-makers to understand the impact of various collaboration strategies on mission performance. Furthermore, the decision makers must be able to assess the performance of the mission in the event of degradation that inhibits collaboration. This paper hypothesizes that not every mission will have a measurable benefit in performance from collaboration, and that mission characteristics can be identified which describe the “tipping points” at which collaboration truly gives a significant bonus in overall mission performance.

To test this hypothesis, it is necessary to create a framework to model operations, quantify and evaluate the metrics of effectiveness, and determine which operational paradigm is most effective for a given mission and scenario. With advances in computing power and development in data analysis techniques, such an analysis is now possible. This paper will showcase such a framework used to evaluate the effectiveness of a group of coordinated unmanned aerial vehicles assisting in maritime search and rescue operations under varying mission parameters. The goal of the framework will be to allow decision makers to determine what combination of mission characteristics and collaboration tactics will actually lead to overall mission improvement, as well as identify mission scenarios in which mission performance is insensitive to the addition of collaborative tactics.

1.1 Autonomous Coordination

An autonomous system of UAVs is said to be coordinated if it consists of a group which communicates and shares information which is then used to make decisions leading toward a pre-determined goal. The type and frequency of information shared is dependent on the design of the system and would vary based on the application.

Currently, UAVs have largely been used individually and in an uncoordinated manner. However, recent technological advances are enabling the maturation of coordinated groups of UAVs. The main drivers have been advances in computer science which have improved data storage and processing capabilities. At the same time, many leaders in industry, government, and academia have included UAV coordination in their scope of future research and development. In their 2011 to 2036 Unmanned System Integrated Roadmap, the U.S. Department of Defense has stated that, "… autonomous systems need the capability to interact and work together with other autonomous systems … and do so safely and reliably."[2] Additionally, the U.S. Navy has stated that, "The area of autonomy and control is a major research area of all UVs … aspect of autonomy is the cooperative or collaborative coordination among multiple vehicles."[3] There is therefore a clear desire to implement coordinated multi-UAV systems. Furthermore, the technology to enable these advances is being matured, even though the advances themselves have yet to be realized.

However, there has not been a significant amount of research on defining the characteristics of missions and scenarios in which coordination would truly provide a measureable benefit in mission performance over simply using an uncoordinated group of multiple UAVs. Much of the literature and vision documents assert that there is much benefit to be had from collaboration, but often gloss over the cost of achieving that collaboration. Given that the development and implementation of such a technology is neither trivial nor inexpensive, it is necessary to understand the cases in which the achieved benefits from collaboration are worth the cost of implementing the technology. This paper presents a methodology and framework understanding the sensitivity of the collaborative benefits to changes in the scenario under which a mission is being executed. This approach provides a way to assess the value of unmanned vehicle collaboration for missions and scenarios of interest by comparing the coordinated mission performance to a baseline, uncoordinated mission performance. Using a sensitivity analysis, the characteristics of the scenario in which collaboration is necessary to achieve significant mission performance improvements are identified for the case study presented. This framework is presented through the application to a case study, described in the following section.

1.2 Case Study Introduction: Maritime Search and Rescue

Maritime search and rescue (SAR) has a very simple problem statement: a missing person is lost at sea and must be recovered as soon as possible. SAR is a relevant engineering problem for a number of reasons. First, the search area must be explored in minimum time. This can be framed as an optimization problem, and many possible methods may be devised to accomplish this goal. Second, uncertainty about the missing person's location and the impacts of the weather environment may make the search area very large. Finally, a person or life raft in the water is a challenging object for most sensors to detect. The combination of these challenges makes this a particularly difficult mission, and one in which any operational edge is critical.

Further challenges are introduced by assuming the search area is large. A larger area requires more time to explore when only a few vehicles are available for the search effort. Once the time required becomes large, persistence becomes an issue and planners must begin to consider refueling needs. If the number of search vehicles is increased to mitigate these challenges, more personnel will also be required. These challenges imply a need to search the area more quickly, with a larger number of vehicles having higher persistence, and with fewer personnel required for the search effort. This implies a possible solution to the problem may be a coordinated group of autonomous UAVs. The goal of the case study presented in this paper is to determine under what characteristics of the search area will a coordinated group of UAVs provide a significant benefit to mission performance (in terms of mission completion time) as compared to one or more uncoordinated UAVs,

2 Case Study Objectives

Several questions immediately arise when considering whether to employ collaborative UAV groups in maritime SAR. The first question asks whether the group is truly needed, or under what circumstances such a group may be useful. For instance, many SAR missions do not involve a large search area. If the search area is small, it is not clear that an autonomous group would bring added value. However for large search, there is little debate that, given the necessary technological advancements, such a group would be of immense use. The question still remains as to how much benefit can be gained from coordinated operations over simply using a large number of vehicles in an uncoordinated manner. Answering these questions is impossible without the means of such a framework as we propose. The framework must allow the physical characteristics of the search problem to be varied parametrically.

Therefore, the framework proposed here is designed to aid in answering the following general questions:

- Under what scenarios does a group of UAVs provide measurable benefit over a single UAV for a given mission?
- Under what scenarios does a coordinated group of UAVs provide measurable benefit over an uncoordinated group of UAVs for a given mission?

Thus, the goal of the work presented in this paper is to create a framework that can be used to evaluate the effectiveness of conventional maritime SAR operations and a hypothetical group of autonomous UAVs in a physically parameterized mission space. Here, the problem of optimizing the search methods is not undertaken. Instead, we create a framework to determine under what circumstances a coordinated group is useful.

Some of the mission specific questions answered in the case study include:

- How large must the search area be before a group brings added value?
- What are the impacts of sensor technology and weather conditions on the metrics of effectiveness?
- What are the cost impacts of utilizing a coordinated group?

With this framework, a decision maker can identify critical points (if any) inside the mission space where collaborative search becomes more effective than the current operational paradigm. The framework also allows the decision-maker to identify the more effective operational plan for a particular set of mission space input parameters.

This maritime search and rescue case study is not the end-goal of this framework and related research. This case study was merely a starting point in the research of exploring the implication of large mission spaces to the operations and technologies used within them. The same approach may be applied to any similarly parameterized mission, or extended to evaluate alternative collaboration methods to the single method considered in this study.

3 Methodology

3.1 General Approach

The proposed framework is based around the application of a decision-making process to answering the questions posed in the previous section. Thus, the general steps taken include are: (1) Clearly define the parametric mission space of interest, (2) Clearly define the baseline against which collaborative operations will be evaluated, (3) Determine coordination methods of interest and group configurations to be evaluated, (4) Develop and implement an appropriate modeling and simulation testbed, (5) Evaluate the baseline and proposed alternatives using the testbed, and (6) synthesize and visualize results to gain insights and support decision making. In order to describe the application of these steps to a problem, the case study is used as a demonstration. Note that this same overall framework can be applied for any other mission in which similar questions arise.

3.2 Application of Methodology to Case Study

The first step in the development of the framework was to describe the problem of interest. In this case, a person issues a distress call from a known point in the ocean. The coast guard immediately deploys one or more vehicles to search for and rescue the person in distress. The primary measure of success will be the time to locate and recover the distressed person. Other metrics, such as the cost to complete the mission, will also be tracked. The dimensions of the search area are linked to the estimated uncertainty of the missing person's initial location and the expected drift. Once the search area is defined, both the conventional and collaborative search operations must be identified. With these methods in place, evaluation of the operational alternatives requires an agent based modeling and simulation environment. In our environment, each search vehicle is modeled as an agent having a set of attributes and methods which guide its motion, behavior, and interactions with other agents. The environment enables input of mission parameters, dynamic visualization of a mission, measurement of metrics of effectiveness, and comparison between operational methods. The two operational methods considered in this study are the conventional SAR operations as described by the US Coast Guard in [4], and an algorithm for autonomous coordinated search described in [6] and [7].

3.2.1 Mission Definition

SAROPS [9] is a major operational planning and decision support tool used by the US Coast Guard. The tool automatically generates a search area based on a probability of mission success (POS), probability of containment (POC), and probability of detection (POD). This search area is comprised of an outer perimeter (an irregular polygon) and a probability density function overlaid on the interior space. Within SAROPS, the outer perimeter is varied until the POC is under a certain threshold. Thus, the definition of the search area is itself an optimization problem which depends on the rescuers' best intelligence about the behavior of the missing person prior to distress and upon weather conditions affecting drift.

For our study, we chose to adopt a similar probabilistic approach to defining the search area, but simplified its shape. We assume a datum point has been reported in a distress call, and a circular area whose radius is equal to the uncertainty attached to the reported position. The search area is then lengthened and widened in the direction of the expected drift. The widening angle is related to the uncertainty of the drift estimate, and the length is equal to the expected drift velocity multiplied by the total time allotted for the search effort. A sketch of a representative search area used in our simulation model is shown in Figure 1.

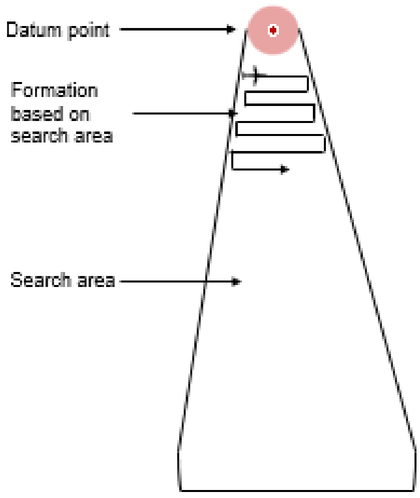

Figure 1: Search area and standard ladder pattern from the U.S. Coast Guard

3.2.2 Baseline Development

Although the Coast Guard uses a probability density map to define a search area, the rescue vehicle makes no use of the distribution after the mission has begun; it is used only as a planning tool to define the search area prior to mission start. Once a datum point is defined, a deterministic search pattern is flown over the area until the missing person is located and recovered. Two typical search patterns are the ladder (Figure 1) and the expanding square. In our implementation, we choose to model the expanding square pattern as the baseline.

3.2.3 Definition of Alternatives

The next step in the process was to identify coordinated operations to model and compare to the current operations of maritime search and rescue. As discussed previously, the current SAR-Ops generates a probability distribution based on inputs such as distress location certainty, and wind and drift velocity. The operators then use this data to determine a static search pattern. This method, while simple, abandons a lot of information in the distribution that could be further used to conduct a search using a more complex, but robust method. It was therefore sought to identify a search method which utilized such information and could be used by a group of UAVs to, in a coordinated fashion, intelligently search the area for the lost person.

After a literature search on coordination methods for finding a lost object in a dynamic environment, a method called “an optimal search in a Bayesian world” was identified as a potential approach for collaboration [5][6]. In such a scheme, a probability distribution function is overlaid on a physical environment based on known information. The probability over each cell is the probability or “belief” that the target is in that location. A group of UAVs are then released to explore the environment each with their own belief map. At each time step, the UAVs check their sensors and update their local belief map with the new information along with a time stamp. Then, they share their new belief maps with other agents within communication matrix and update their maps based on information received. Additionally, due to uncertainty, wind, and drift, the probability in the cells of each agents map begin to “grow back” if the cell has not been visited or updated for an extended amount of time, and the values in cells shift in space at the same rate at the drift; this is similar to the “spreading fire” problem. An example simulation from a group that helped developed this method is shown in Figure 2. Using this coordination method, a group of UAVs can search a complex space in a non-redundant manner while ensuring the areas of highest probability are searched first.

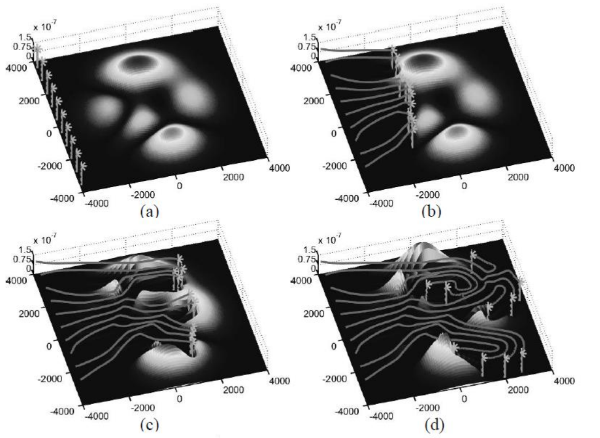

Figure 2: Example of multiple agents searching for a lost target using an optimal search in a Bayesian world [5] [6]

It was therefore required to generate a probability distribution for the agents in the model to use to search the space. The U.S. Coast Guard handbook was again consulted to attempt to reproduce a rough approximation of what their statistical model may output. A 2-dimensional Gaussian distribution was created based on statistical information provided by the U.S. Coast Guard data. The positional uncertainty of the location reported in the distress call was used as a guideline for creating the standard deviation of the distribution. The standard deviation stretches in the direction of the drift velocity. This stretching is time-dependent to reflect the real-world insight that as time passes, the true location of the missing person becomes more uncertain. The equations for a standard Gaussian distribution, the mean, and the standard deviation are shown below and a graphical representation of the distribution changing in time is shown in Figure 3. In this example, the bearing error is 4 degrees (as defined in the U.S. Coast Guard S&R manual) and the radius is of the distribution at mission start is 15 nautical miles.

Equations 1-5 describe this implementation.

$$f(x|\mu,\sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \tag{1}$$

$$\sigma_x = \frac{34.1}{100} * (DriftVel * time + radius) \tag{2}$$

$$\sigma_x = \frac{34.1}{100} * (DriftVel * time * \sin(bearingError) + radius) \tag{3}$$

$$\mu_x = DriftVel * time + radius \tag{4}$$

$$\mu_y = radius \tag{5}$$

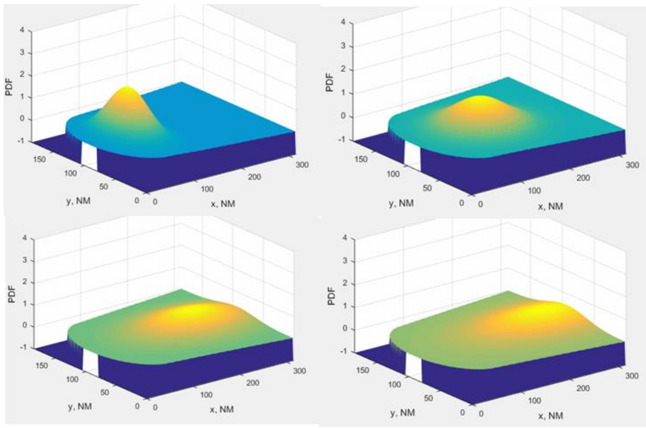

Figure 3: Gaussian distribution evolution at 0, 64, 150, and 240 min into after the distress call, read left to right, top to bottom

3.2.4 Modeling and Simulation

In order to penetrate the complexities of a SoS, a modeling and simulation environment is often required. While visualization of the system is important, the real benefit of such an environment is the ability to explore and identify behaviors in the SoS which were unforeseen until revealed by exploitation of the model. The environment required for this problem is an agent based simulation capable of modeling multiple, coordinated UAVs. In addition, the environment must be able to handle other aspects of the problem such as the missing person(s), weather condition, and the Bayesian model of the search area.

Before constructing a new environment from scratch, existing solutions were explored, and one was found that fit most of the requirements from the previous paragraph. The Java-based environment is called the Unmanned Vehicle Collaboration Research Environment (UV-CoRE) [10], shown in Figure 4. Although UV-CoRE was already an agent-based simulation, it lacked inter-agent coordination and the Bayesian search method. Therefore, these missing pieces were added to the environment.

Figure 4: Unmanned Vehicle Collaboration Research Environment (UV-CoRE)

3.2.5 Evaluation of Mission Effectiveness

In order to be able to evaluate the coordinated operations, the metrics of success for this mission had to be defined. When searching for a person lost at sea, especially in cool water, the most important thing is to find them as soon as possible. In waters as warm as 50 degrees Fahrenheit, hypothermia can become a serious risk in as little as 30 to 60 minutes, with the expected survival time being as low as 1 to 3 hours [8]. Another metric of success in the cost of the mission. Currently, the United States Coast Guard spends approximately 50 million dollars annually on SAR operations [4]. Therefore, the two main metrics of success identified for a maritime search and rescue mission were the total time to find the target, and the total cost of running the mission. The time to find the target is the sum of the response time (the time it takes the rescue team to take-off after receiving a distress call) and the amount of time to find the target. While the time metric is straight forward to calculate, the cost metric is much more difficult to estimate accurately. The mission cost is broken down in Table 1 and discussed in the following paragraph.

Although we acknowledge the acquisition cost of the rescue platform, we do not model this cost in the simulation environment. . As the acquisition cost of vehicles are amortized over its entire lifespan, a proper estimate for the cost of ownership during a single mission was not able to be found. Therefore, to avoid making poor assumptions, it was not included in the model (however, a user with such knowledge of their system could easily modify the environment to include any acquisition cost they chose). The direct operating cost consists of fuel cost and crew support. At each time step, the thrust required was calculated from a simple flat plate drag model. The fuel burn was determined from the thrust required and an estimated thrust-specific fuel consumption. The cost of fuel was estimated by the current average price of Jet-A. The crew support is an input of the cost per hour for the helicopter crew to monitor the group of UAVs and participate in the search. Lastly, much like the acquisition cost, there once again was difficulty properly estimating maintenance costs per vehicle per mission. Therefore, maintenance costs are acknowledged, but not included in the model.

Table 1: Cost Modeling Breakdown

Total Cost		
Acquisition Cost	Cost of vehicles and sensors	Not included
Direct Operating Cost	Fuel and crew support	Fuel: $2.00 per gallon
		Crew: $2800 per hour
Indirect Operating Cost	Maintenance of vehicles	Not included

4 Experimental Setup

With all of the pieces of the framework identified, the maritime search and rescue case study can be implemented. A design of experiments (DoE) was performed on the parameterized input space and a surrogate model of the metrics of success was generated. With the surrogate mode, an analysis of the data could be performed with substantially fewer simulation runs.

4.1 Mission Variables

To set up the parameterized mission space, a list of the mission parameters to be varied was defined. Some of the parameters are best explained with a visualization and are shown below in Figure 5. All input parameters are discussed below:

1. Response time
The response time is the time between receipt of a distress call and the dispatch of a rescue team to recover the missing person.

2. Radius of uncertainty
This parameter quantifies the uncertainty of the initial location reported in the distress call. The magnitude depends on the technology used by the distress vessel to generate the latitude and longitude coordinates.

3. Wind velocity / Ocean current velocity
The wind velocity and ocean currents affect the drift of the target throughout the mission. This parameter allows the user to perform sensitivity analysis based on atmospheric and oceanic conditions.

4. Estimated drift velocity
The drift velocity is estimated by a superposition of the wind velocity and ocean currents.

5. Offset factor/angle
The offset factor is the distance from the missing person's actual location to the location reported in the distress call. Small vessels with low-quality equipment may report a location which is inaccurate by tens of nautical miles. The offset angle defines a direction to the missing person relative to the drift velocity. These parameters are known to the simulation object, but are neither known nor accessed by the search and rescue team during a simulated mission.

6. Drift angle to base
The drift angle to base is the relative angle between the estimated drift velocity of the target and a vector to the take-of location of the rescue team. This parameter enables sensitivity analysis based on the effects of approaching the search area from a downwind, upwind, or crosswind direction.

7. Type of UAV
The type of UAV selected for this project was a MQ-1 Predator. The UV-CoRE environment is modular, so any existing or hypothetical platform may be modeled by the user.

8. Number of UAVs
The number of UAVs in the coordinated group.

9. Sensor effective sweep radius
This parameter defines the maximum radius at which the UAV may be reasonably expected to detect the missing person. The simulation does not model signal degradation. Instead, a probabilistic model is used to handle the possibility of the search vehicle not detecting the target and reporting a false negative.

10. Probability of detection (UAV)
Even if the target is within the sweep radius of the sensor, it is difficult for recognition algorithms to detect the person, especially in adverse conditions and high sea states. This parameter introduces stochasticity into the simulation.

11. Probability of detection (manned vehicle)
This is the probability that a trained search and rescue pilot will not identify a person even if the person is within the pilot's field of vision or sensor radius.

In addition to the mission parameters, a model of the current operational paradigm had to be defined. For this purpose, a search helicopter was simulated to follow an approximation of present-day U.S. Coast Guard operational procedures. The helicopter flies directly to the peak of the PDF. After reaching the peak, it discards the PDF and begins an expanding square pattern.

For a very practical reason, a helicopter was also included in the simulation of the coordinated group of UAVs. In a real mission, it is extremely unlikely that a group of fixed-wing UAVs will be deployed alone since they have no way to recover the person once located. So to reflect the fact that a person needs not only to be found but rescued, a conventional rescue helicopter is dispatched by the simulation along with the group of UAVs. Upon arrival, the UAVs continue to use the PDF and perform the steps shown in Figure 6. After arriving at the peak of the distribution, they check their sensors at each time step and use the information to update their belief map. If the missing person is not found at that location, the probability is lowered at that location. Then, each UAV shares its map with all other UAVs within communication range. After sharing data, each UAV moves in a steepest ascent direction to the next-highest probability. The step where they share information is what makes the operations coordinated and not just autonomous.

A sample visualization from an actual simulation run is shown in Figure 7. In this run, seven UAVs and a search helicopter are dispatched to find a missing person (identified by the circular icon). In the figure, the search helicopter and the UAVs are indistinguishable from one another.

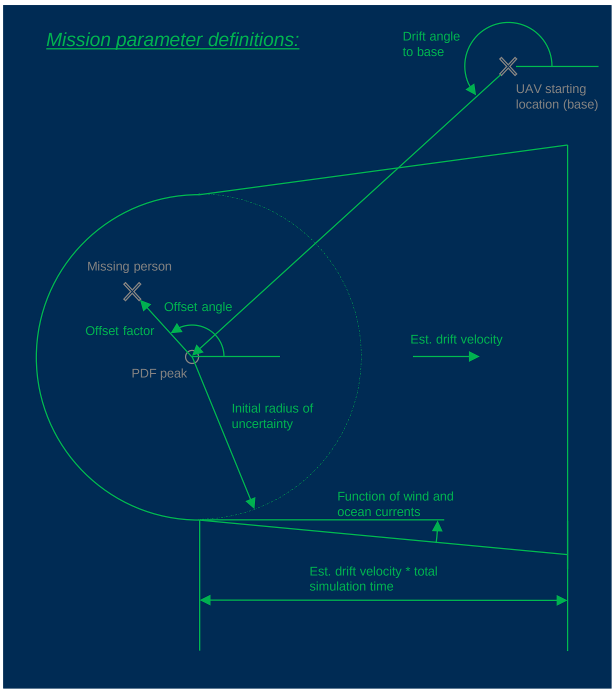

Figure 5: Search and rescue model mission parameter definitions

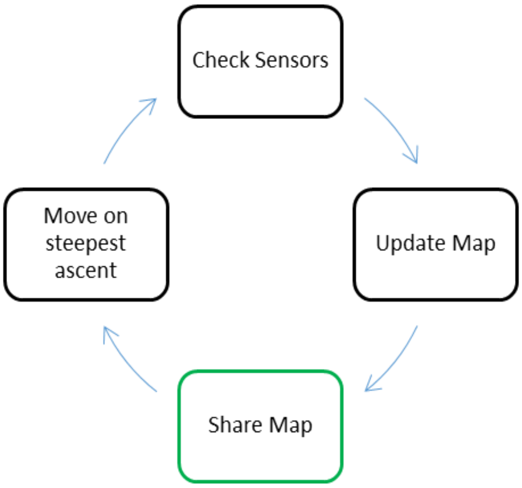

Figure 6: Overview of how the UAVs were modeled. Note the step outlined in green is where the coordination came into play

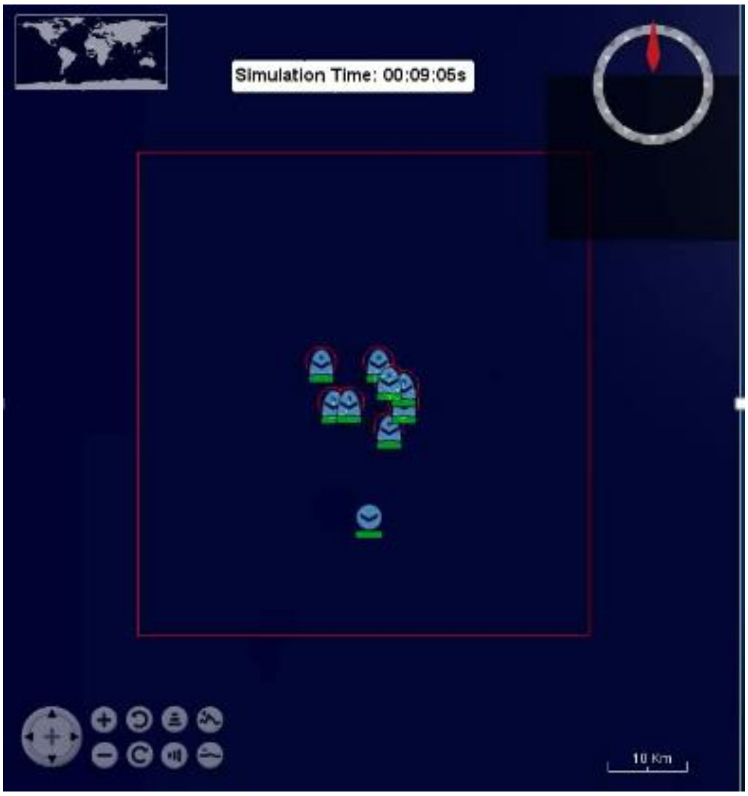

Figure 7: Screen shot of the finished maritime search and rescue model.

4.2 Design of Experiments

The modeling and simulation environment can easily be exploited for analysis of the mission space in its entirety and not just for analysis of a single simulation run. To perform this kind of analysis, many possible combinations of the mission parameters must be queried. The upper and lower bounds considered in this study are shown below in Table 2. Even with a resolution of one unit on every parameter (0.1 units in the case of the sweep radius), the required number of runs would be 6.9984e11 cases. At one minute of computation time per case, this would take about 1.3 million years to run! An appropriate design of experiments (DoE) was therefore needed to extract the most information from the least possible computation time. Even with a design of experiments to reduce the required number of runs, the probabilities of false negatives introduces stochasticity into the simulation and necessitates repetition of individual runs until the moving average remains within some tolerance. For the sake of computational time, this was estimated to be ten repetitions per case.

The DoE selected was a Latin hypercube with additional corner points to reduce error at the extremes of the mission space. The resulting hybrid DoE consisted of 880 cases plus 70 corner point cases Each case was repeated ten times.

Table 2: DoE Ranges

Mission Inputs	Parameter	Lower Bound	Upper Bound
Search Space	Response time (min)	0	30
	Initial radius of uncertainty (nmi)	5	30
	Wind velocity (kts)	0	20
	Ocean current velocity (kts)	0	3
	Offset factor (0-50% of uncertainty rad.)	0	1
	Offset angle (dew)	0	360
	Drift angle relative to base (deg)	0	360
Vehicle	Type of UAVs	MQ-1 Predator (test case)	
	Number of UAVs	0	10
	Sensor effective sweep radius (nmi)	0.3	1.5

	Probability of missed detection (UAV)	30%
	Probability of missed detection (manned)	20%

4.3 Creation of Surrogate Model

The results from the cases run in the DoE were used to create a surrogate model. This was done in order to enable a sensitivity analysis of the mission parameters without excessively large computational requirements. Secondly, the surrogate models are able to provide a continuous representation of the design space, which can then be used to visualize the actual shape of the space and identify key critical points where a particular architecture (i.e. collaborative or non-collaborative) becomes the preferable concept of operations.

The surrogate model selected for this problem was a neural network. The neural network was selected because the problem is highly non-linear, has continuous and discrete inputs, and has a large number of input variables with large ranges. A software tool developed the Georgia Tech Aerospace System Design Lab (ASDL) called BRAINN 2.4 was used to generate the model. The neural network architecture had a single layer with eight hidden nodes with a logistic sigmoid activation function.

Initial attempts to create the neural network resulted in overfitting, so subsequent runs employed early stopping during the training process. 220 random points were then used to validate the model, and the resulting validation report is shown in Figure 8.

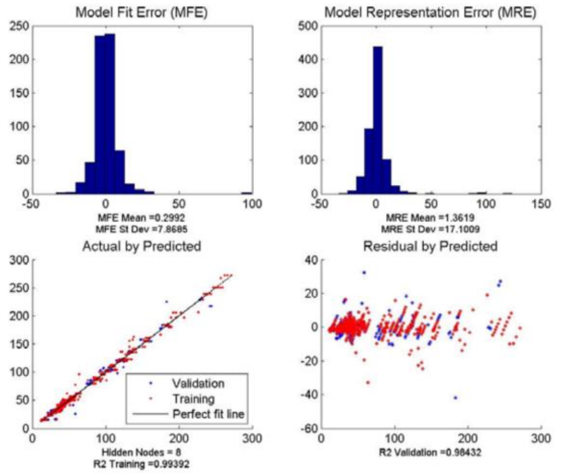

Figure 8: Goodness of Fit Measures for Surrogate Model

5 Case Study Results

From the surrogate model, a deterministic equation was obtained from which an eleven dimensional response surface could be generated. This equation was exported to the statistical analysis software JMP to be analyzed. The goal was to identify critical points in the large mission space where coordinated operations added no benefit to the success of the mission compared to the conventional operations. Additionally, it was sought to find areas of the mission space where the coordinated group of UAVs brought substantial benefit.

This discussion will center on the primary metric of success for this mission: time to locate the target. JMP was used to create a set of prediction profilers to visually represent the sensitivity of the response to each of the inputs. A screenshot of the JMP profiler is shown below in Figure 9, and can be used to assist in sensitivity analysis. The inputs of the model are shown on the x-axis with the outputs on the y-axis. The lines in each box represent the trace of the function (i.e. the trend line of that output against the given parameter when all other parameters are held constant at their displayed value). The values for the inputs shown in red are a critical point identified in the mission space where the metrics of success are insensitive to the addition of coordinated UAVs.

A search and rescue dispatcher would have control over only three inputs: the sensor radius, the response time, and the number of agents sent out. All others are either unknown to the dispatcher or uncontrollable. From the profiler below, it can be seen that for this level of uncertainty about the initial position and weather conditions, a single search helicopter could perform equally as well as a coordinated group of UAVs. However, the time to recover the missing person could be substantially reduced by investing in better sensor. An interesting tradeoff can be seen here between the number of agents and the strength of the sensors being used. There is a tension between having one vehicle with very advanced sensors and having many vehicles with cheaper sensors. Selecting the right choice not only could save more lives, but could save money. This framework can allow for such a tradeoff analysis to be performed before investing valuable resources into what may be a less effective platform.

The framework was also used to discover regions of the mission space where it is beneficial to dispatch a group of coordinated UAVs. Figure 10 shows the prediction profiler at one of these regions. To clearly show the impact on the time metric, the profiler is shown twice: once with one agent (the search helicopter) and again with eleven agents (the search helicopter plus ten coordinated UAVs.) The profilers indicate the target was found more than one hour sooner with a coordinated group. This is not an insignificant amount of time and could be the difference between life and death.

Even in this region of the mission space, it is again noted that the largest gains still come from investing in sensor technology. Also, an unintuitive result is visible at this point. Notice that increasing the response time actually decreases the time required to find the target. This is due to the drift angle relative to base (which is 175 degrees for this case). Referring back to Figure 5, one can see that such an angle implies that the person is drifting towards the base. It would be nonsense to advocate waiting longer to respond when the target is drifting towards the base, but this result does give insight on how the scenario itself can have a significant impact on the overall mission outcomes, as well as an example of the care which must be taken when interpreting the results of a framework such as the one presented here.

From these results, it has been demonstrated that the framework is capable of answering the questions which motivated the development of the problem. Based on the available technology and characteristics of the mission space, a user of this framework can clearly identify under what conditions the application of a coordinated group of UAVs can provide a quantifiable benefit in mission success over a conventional concept of operations.

A final demonstration of the utility of our modeling and simulation environment is its ability to support decision making. The decisions discussed here are representative of what may be encountered in the requirements definition phase of a new design, in selection of sensor packages for an existing vehicle, and in assembly of a rescue team during mission planning.

Constraint analysis may be performed during requirements definition by plotting contours of mission input parameters. For example, plotting the radius of uncertainty versus the sensor radius allows designers to set requirements for sensor packages if a fixed cost or time is allotted for the design mission. Figure 11 shows such a contour plot for a coordinated group of six UAVs and a maximum of 60 minutes allowed to locate the target. Varying the other mission parameters constitutes a different design point and allows rapid tradeoff studies to be performed.

Two decisions, selection of sensor packages and mission planning, can be performed with a contour plot similar to Figure 12. Here, the plot can be used in one of at least two ways. Given a maximum mission time and a number of UAVs, requirements for sensor packages can be set during the design phase of a new vehicle or system of vehicles. Conversely, the required number of UAVs can be determined in mission planning if the sensor radius is a given. Similar constraint analyses may be performed on the other mission parameters.

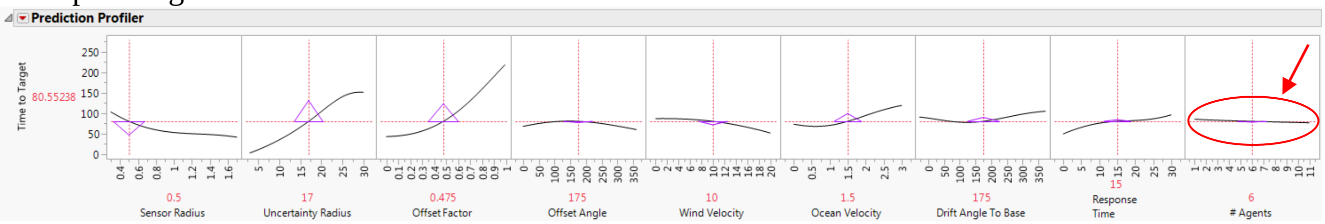

Figure 9: Prediction profiler at point which is insensitive to changes in the number of UAVs. The arrows indicate the magnitude and direction of the partial derivative at this point in the mission space.

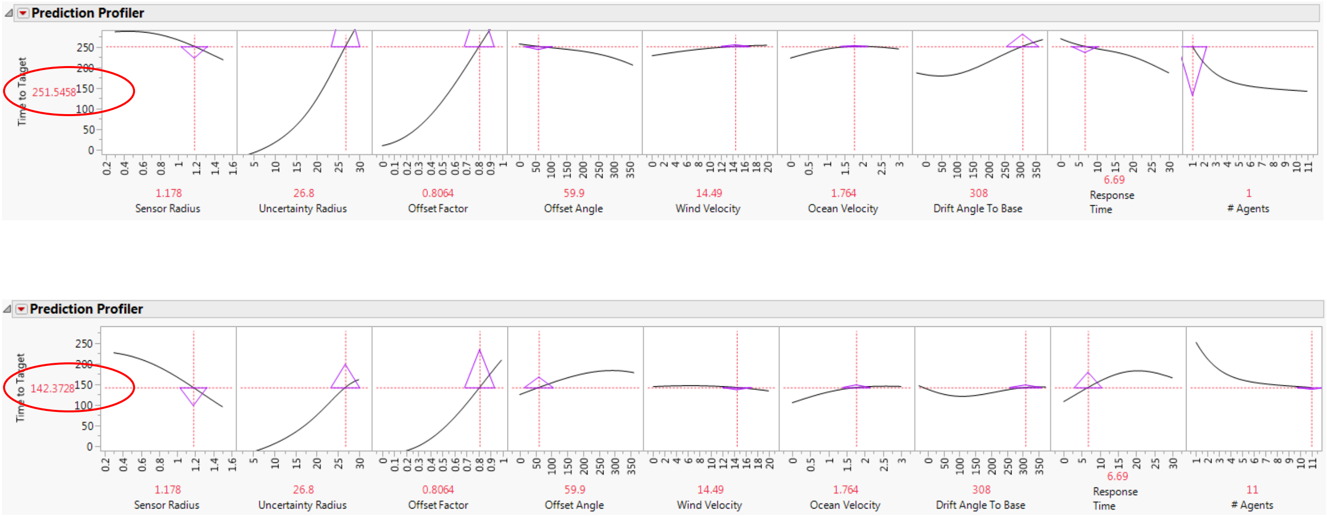

Figure 10: Point in mission space where the number of agents significantly impact the time to find the target (note the times in red to the right of the vertical axis).

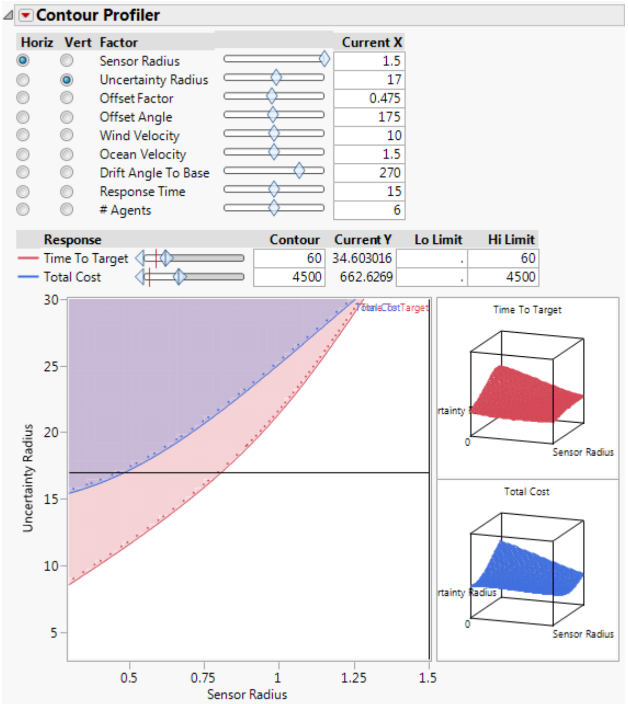

Figure 11: Contour plot of uncertainty radius versus sensor radius with cost and time constraints

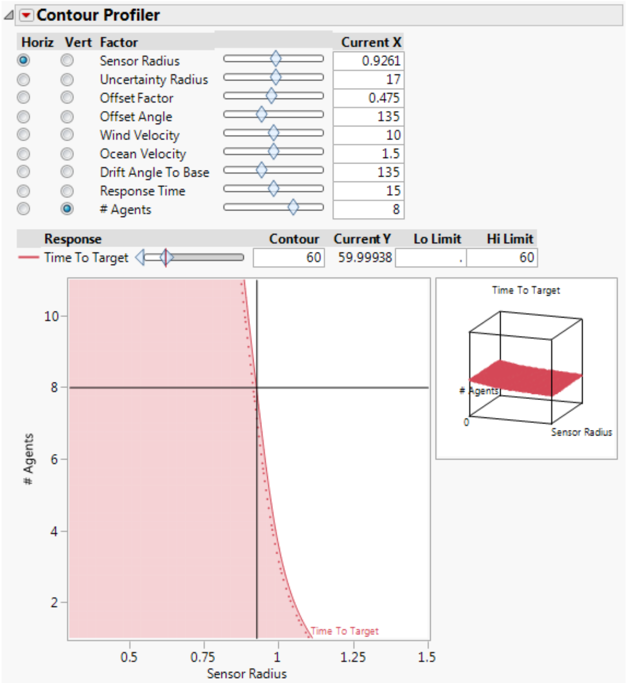

Figure 12: Contour plot of number of agents versus sensor radius with a requirement of 60 minutes allowable to complete the mission

6 Concluding Remarks

This paper has described and demonstrated an approach for understanding the types of mission scenarios for which coordinated UAV operations can be beneficial. The framework created for evaluating the effectiveness of coordinated operations on a large mission space has potential for assisting decision makers with understanding when and where their SoS will be effective. For the first implementation of the framework, a maritime search and rescue mission was modeled, and the case study has been presented here. From the results, critical points within the mission space where coordinated operations brought no additional benefit to the mission were able to be found, as well as critical points where coordination provided clear benefit to the mission. Additionally, the framework can allow for other tradeoff analysis such as examining trades between sensor performance and the benefits of coordinated operations. Future work include the inclusion of additional missions and collaboration schemes into the framework, and will consider approaches for improved cost modeling. The eventual goal will be to have a robust framework to perform virtual experiments on the operations of coordinated UAVs across a broad range of missions.

References

1 J. Hoog, *Role-Based Multi-Robot Exploration*, Doctor of Philosophy Thesis, University of Oxford, 2011.

2 Department of Defense, *Unmanned Systems Integrated Roadmap*, 2013

3 United States Navy, *The Navy Unmanned Surface Vehicle (USV) Master Plan*, 2007

4 U.S. Coast Guard Addendum to the US National Search and Rescue Supplement (NSS) To the International Aeronautical and Maritime Search and Rescue Manual (IAMSAR) (COMDTINST M161). January 7, 2013.

5 F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte, *Coordinated Decentralized Search for a Lost Target in a Bayesian World*, Intl. Conference on Intelligent Robots and Systems, October 2003

6 F. Bourgault, T. Furukawa, and H.F. Durrant-Whyte, *Optimal search for a lost target in a Bayesian world.* In *Int. Conf. on Field and Service Robotics (FSR'03), 2003*

7 S. Waharte, N. Trigoni, and S.J. Julier, *Coordinated Search with a Swarm of UAVs*

8 U.S. Coast Guard, *Cold Water Survival* , http://www.ussartf.org/cold_water_survival.htm [retrieved April 29, 2016]

9 T. Kratzke, L. Stone, J. Frost, "Search and Rescue Optimal Planning System", 13th Conference on Information Fusion, IEEE, July, 2010, Edinburgh.

10 Bays, M.P., Jim; Coker, Ayodeji; Cruzen, Christian; Braden, Julianne; Fouts, Cody; Griendling, Kelly; Mavris, Dimitri, *Modeling and Simulation of Unmanned Vehicle Sentry Missions to Assess Communications in a Maritime Environment.* Naval Engineers Journal, 2015. **127**(2): p. 105-110.

Copyright Statement