Human-Guided Management of Collaborating Unmanned Vehicles in Degraded Communication Environments

by

Daniel Noel Southern

S.B., Massachusetts Institute of Technology (2009)

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the Degree of

Master of Engineering in Electrical Engineering and Computer Science

at the Massachusetts Institute of Technology

May 2010

Author . . . .
Department of Electrical Engineering and Computer Science
May 21, 2010

Certified by . . . .
Mary L. Cummings
Associate Professor of Aeronautics and Astronautics
Thesis Supervisor

Accepted by . . . .
Arthur C. Smith
Professor of Electrical Engineering
Chairman, Department Committee on Graduate Theses

Report Documentation Page		Form Approved OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE **MAY 2010**	2. REPORT TYPE	3. DATES COVERED **00-00-2010 to 00-00-2010**
4. TITLE AND SUBTITLE **Human-Guided Management of Collaborating Unmanned Vehicles in Degraded Communication Environments**		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Massachusetts Institute of Technology,Department of Electrical Engineering and Computer Science,77 Massachusetts Avenue,Cambridge,MA,02139-4307**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)
12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**		
13. SUPPLEMENTARY NOTES		

14. ABSTRACT

Unmanned Aerial Systems (UASs) currently ful ll important roles in modern military operations. Present commitments to research and development e orts for future UASs indicate that their ubiquity and the scope of their applications will only continue to increase. For sophisticated UASs characterized by coordination of multiple vehicles, it is a formidable challenge to maintain an understanding of the complexities arising from the interaction of human supervisory control, automated planning, and network communication. This research investigates the robustness of UAS performance under degraded communication conditions through simulation with a particular futuristic UAS, the On-board Planning System for UxVs in Support of Expeditionary Reconnaissance and Surveillance (OPS-USERS) system. The availability of reliable communications is vital to the success of current UASs. This dependence is not likely to be diminished in future systems where increased inter-vehicle collaboration may actually increase reliance on communications. Characterizing the e ects of communications availability on the performance of a simulated UAS provides crucial insight into the response of UASs to communication failure modes which may be encountered in real-world implementations. Additionally, de ning a minimum tolerable level of communication availability which will allow a UAS to operate with acceptable performance represents the groundwork for designing engineering speci cations for communications systems as well as for de ning conditions under which such a system could be expected to operate e ectively. Experiments are designed and executed to investigate the impact of degraded communication conditions on the performance of UASs by sampling the performance of a simulated UAS under a variety of degraded communication conditions. These experimental conditions are based on a similar previous experiment, which utilized the same simulation testbed and investigated the impact of operator workload on system performance in experiments with human participants. However, this research seeks to collect data over a wider range of communication conditions than experimentation with human participants practically allows. Therefore, a human model is also developed to emulate the interaction of an average human operator with the system. After initial experiments validated that the human model produced results that were statistically indistinguishable from the results of the experimental data on which the model was based, it was employed in repeated simulations to collect data across a large number of experimental conditions. Communication availability was modulated by imposing various network connectivity topologies on the agents in the UAS, as well as by introducing arti cial delays into message transmissions between agents. Analysis of the simulation results suggests that the various functions of the system exhibit two main modes of

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**102**	

Human-Guided Management of Collaborating Unmanned Vehicles in Degraded Communication Environments

by

Daniel Noel Southern

Abstract

Unmanned Aerial Systems (UASs) currently fulfill important roles in modern military operations. Present commitments to research and development efforts for future UASs indicate that their ubiquity and the scope of their applications will only continue to increase. For sophisticated UASs characterized by coordination of multiple vehicles, it is a formidable challenge to maintain an understanding of the complexities arising from the interaction of human supervisory control, automated planning, and network communication. This research investigates the robustness of UAS performance under degraded communication conditions through simulation with a particular futuristic UAS, the On-board Planning System for UxVs in Support of Expeditionary Reconnaissance and Surveillance (OPS-USERS) system.

The availability of reliable communications is vital to the success of current UASs. This dependence is not likely to be diminished in future systems where increased inter-vehicle collaboration may actually increase reliance on communications. Characterizing the effects of communications availability on the performance of a simulated UAS provides crucial insight into the response of UASs to communication failure modes which may be encountered in real-world implementations. Additionally, defining a minimum tolerable level of communication availability which will allow a UAS to operate with acceptable performance represents the groundwork for designing engineering specifications for communications systems, as well as for defining conditions under which such a system could be expected to operate effectively.

Experiments are designed and executed to investigate the impact of degraded communication conditions on the performance of UASs by sampling the performance of a simulated UAS under a variety of degraded communication conditions. These experimental conditions are based on a similar previous experiment, which utilized the same simulation testbed and investigated the impact of operator workload on system performance in experiments with human participants. However, this research seeks to collect data over a wider range of communication conditions than experimentation with human participants practically allows. Therefore, a human model is also developed to emulate the interaction of an average human operator with the system.

After initial experiments validated that the human model produced results that were statistically indistinguishable from the results of the experimental data on which the model was based, it was employed in repeated simulations to collect data across a large number of experimental conditions. Communication availability was modulated by imposing various network connectivity topologies on the agents in the UAS, as well as by introducing artificial delays into message transmissions between agents. Analysis of the simulation results suggests that the various functions of the system exhibit two main modes of sensitivity to communication failures. In one mode, exhibited in searching the environment and discovering targets, performance gains associated with a high level of communication availability are relatively small. Performance did not continue to drop with the introduction of further communication failures,

indicating a robustness to communication failures. The other mode, observed in target tracking and hostile destruction performance, exhibits a negative correlation with increasing communication delays. The magnitude of the effect of communication delays is also significantly impacted by the connectivity of the network topology, with lower connectivity topologies amplifying the negative correlation.

Data collected through this experiment provided insight into the characteristics of an ideal minimum level of communication. In addition, the trade-offs between performance in different aspects of the system as well as the optimal allocation of communication resources were considered. This work also investigated the potential for the operator to mitigate performance losses incurred due to communication degradation through more frequent replanning. However, no evidence was found which supported this possibility. Although these results represent preliminary research into the effect of degraded communication on a complex autonomous system, they provide valuable principals to consider when designing future UASs.

Thesis Supervisor: Mary L. Cummings
Title: Associate Professor of Aeronautics and Astronautics

Acknowledgments

First I would like to thank Missy Cummings, my advisor, for guidance and support throughout the past year.

Thank you also to Luca Bertuccelli for providing invaluable feedback on this thesis.

Thank you to the Aerospace Controls Laboratory, Aurora Flight Sciences, and specifically Olivier Toupet for his work on the OPS-USERS software and support throughout all of our experiments.

To the Office of Naval Research for funding the work represented in this thesis.

To my fellow HALiens Andrew Clare and Christin Hart, I'm happy to have had your support throughout the experience of working with OPS-USERS and writing our theses. I would also like to thank Jackie Tappan for a thorough round of feedback in the final weeks of the semester.

To my Mother and Father for your tremendous support throughout my time at MIT.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

13

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

Acronyms

ANOVA Analysis of Variance

BAMS Broad Area Maritime Surveillance

CC Centralized Controller

CBBA Consensus Based Bundle Algorithm

CMM Centralized Mission Manager

GCS Ground Control Station

HITL Human-In-The-Loop

MANOVA Multivariate Analysis of Variance

OPS-USERS On-board Planning System for UxVs in Support of Expeditionary Reconnaissance and Surveillance

POMDP Partially Observable Markov Decision Process

RDTA Robust Decentralized Task Assignment

SCT Schedule Comparison Tool

UAS Unmanned Aerial System

UAV Unmanned Aerial Vehicle

UxV Unmanned Vehicle

WUAV Weaponized Unmanned Aerial Vehicle

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction

1.1 Research Motivation

Unmanned Aerial Vehicles (UAVs) continue to play an increasingly important role in modern military operations, with many countries either currently operating or developing Unmanned Aerial Systems (UASs), a term denoting an entire system, including the UAV(s) as well as required infrastructure. The well-known Predator system operated by the U.S. Air Force exemplifies the utility of unmanned systems. Predator drones have been employed extensively in intelligence gathering and combat operations in the past 16 years, logging over one million hours of flight time as of April 2010 since the first Predator was delivered in 1994 [1]. Other examples of unmanned vehicles currently in use include the Global Hawk, which is also operated by the U.S. Air Force, as well as the Shadow, operated by the U.S. Army. The U. S. Navy operates the vertical takeoff and landing Firescout, and has committed significant resources to the Broad Area Maritime Surveillance (BAMS) UAS which employs a maritime version of the Global Hawk surveillance UAV. With the success of unmanned systems in U.S. military applications, the Department of Defense has pledged an investment of $16 billion in UAS development funding from 2008 through 2013 [2].

Despite the limited autonomy of current UASs, their increasing use and the investment of resources towards the development of future unmanned systems demonstrates the efficacy of unmanned aircraft and emphasizes the importance of the role they will continue to play in future military and civilian applications. Current implementations follow an operational paradigm requiring multiple dedicated human controllers to operate a single unmanned vehicle; for example, the Predator system requires the full attention of a pilot and sensor operator for each aircraft in flight [3]. These unmanned systems represent significant improvements in operational effectiveness over traditional manned aircraft in terms of cost, pilot safety, and even mission capabilities. However, their limited autonomy continues to pre-

clude more advanced mission capabilities, especially missions which call for coordination between many vehicles operating simultaneously. To address these limitations and to develop the capability to apply these systems in a growing array of situations, all branches of the US military continue to invest heavily in new UASs, shifting the operational paradigm away from requiring multiple human operators for each aircraft. Improving flight autonomy represents an active area of research, and advances in autonomy are being applied to fly many UAVs simultaneously, allowing collaboration between the vehicles on a common mission. The applications of multiple UAV coordination are immediately apparent in new mission capabilities, such as cooperative mapping, coordinated search and track, and even coordinated combat operations. These abilities allow UASs to play an important role in network centric operations. A UAS can both provide intelligence to other networked agents, such as soldiers on the ground, and simultaneously empower them with the control of aerial resources. Making the benefits of UASs available to non-pilot operators through improved autonomy both augments the utility of the system and mitigates the necessity of expertly trained operators working in dedicated control stations.

The new mission scenarios in which UASs are being applied call for coordination of a heterogeneous fleet of vehicles with different capabilities, likely working on both disparate and closely related tasks to accomplish higher level goals. Just as current remotely-piloted vehicles are dependent on communication with a Ground Control Station (GCS), a distributed planning system for UAVs would require communications between vehicles to facilitate information sharing and message passing for collaboration, as well as a link to a GCS for additional guidance from human operators. In a practical implementation of a UAS, communication availability represents an integral factor in the system's success; therefore, understanding the impact of communication availability on system performance is a necessity. The performance requirements on a UAS will dictate the level of communication failure that the system can tolerate, and will have important implications in system design specification decisions, cost, and reliability. Communication availability has been shown to have serious theoretical performance implications in automated planning [4], and otherwise represents a practical concern in unmanned vehicle network applications. Improvements in the robustness of system performance to communication failures are valuable.

In real-world implementations, the challenge facing system architects will be to maximize the efficacy of a UAS and the human operator under constraints on personnel, equipment availability, and mission parameters. For example, design choices for forward deployed combat support UASs controlled by a soldier on the ground will differ dramatically from long-range surveillance systems managed through a sophisticated GCS. Understanding the impact of communication availability on system performance, as well as how communication availability interacts with the role of the human supervisor, will be critical to the design of a successful UAS.

1.2 Decentralized Unmanned Vehicle Control

Managing more complex unmanned vehicle missions requires a combination of unmanned aircraft containing significantly augmented autonomy and a human in a supervisory control role, known as Human-In-The-Loop (HITL) control [5]. In such systems, human operators monitor the state of the system and occasionally provide input critical to the completion of tasks within the system, although the autonomy may be capable of acting on a subset of tasks without human intervention. As increases in autonomy diminish the level of human intervention required to pilot each individual UAV, the paradigm for unmanned vehicle control shifts towards each operator managing a system of multiple vehicles rather than a single vehicle or a subset of the systems on a single vehicle. The supervisory role then evolves to encompass higher level functions of the UAS as a whole.

This research focuses on characterizing the effects of communications availability on the performance of UASs by working with one proposed decentralized control architecture for multiple unmanned vehicles, the On-board Planning System for UxVs in Support of Expeditionary Reconnaissance and Surveillance (OPS-USERS) system. OPS-USERS is a framework for coordinated search, track and destroy missions [6], and the term "UxV" (Unmanned Vehicle) refers to an arbitrary type of vehicle, so this system is not limited to aerial vehicles. However, the term UAV will be preferred over UxV in this work, with the understanding that systems may have the capability to support ground- or sea-based vehicles. A more detailed explanation of the OPS-USERS system implementation is provided in Section 3.1.

1.3 Research Questions

This research addresses how communication availability impacts the performance of a system of distributed, highly autonomous UAVs supervised by a human operator, and whether the parameters of HITL control can be adjusted to mitigate the effects of communication failures. A more complete overview of the system architecture under study is provided in Section 3.1, however for now it suffices to note that the system consists of a set of heterogeneous unmanned vehicles and a Ground Control Station (GCS) from which a human operator monitors the system and occasionally provides input. In such a system, the set of vehicles and the GCS will be referred to as the agents in the system, and communication in the system refers to communication between any pair of agents. To assess the performance of such a system, the following research questions were posed.

1.3.1 Performance Response to Communication Failures

Communication availability between the agents in a UAS can be expected to significantly impact performance. Specifically, the availability of communication between each pair of individual UAVs and also between each UAV and the GCS will impact the system's ability to generate efficient plans to complete

high-level goals in different ways, leading to the following question:

Research Question #1: *How does performance of the decentralized system with a human in a supervisory control role degrade with increasing communication failures?*

To investigate the relationship between communication availability and system performance, a baseline performance level is first established with perfect communication. Next various regimes of communication failure modes are simulated to estimate the impact of a set of specific communication failure modes. This information provides insight into the amount of communication loss which can be tolerated in a given mission before the performance of the system is significantly degraded.

Research Question #2: *Is the system more or less tolerant to specific types of communications failures?*

Several modes of communication degradation are possible. Communication links can be completely broken, requiring that information be passed through another route (if possible), or a partially degraded link may introduce latency as messages are passed between agents. Furthermore, the location of the link in the network (i.e. between two UAVs or between a UAV and the GCS) could affect the type of traffic the link carries. Real world conditions could produce communication networks that experience several failure modes simultaneously, and the various behaviors of the system will be impacted in different ways for each failure mode. Characterization of the performance degradation in such failure modes will serve as a basis for understanding the relative merits of various system architectures.

1.3.2 Impact of Human Supervisory Control on Performance

This research will also address the impact of human supervisory control on the tolerance of a UAS to communication failures.

The human operator interacts with the system to monitor the state of the mission and to make decisions to guide the automation. Specifically the operator may oversee prioritization of tasks assigned to the UAVs, as well as target identification, weapon launch approval, and waypoint task creation. Throughout the course of a mission, the operator may choose to or may be prompted by the automation to consider the most recent plan suggested by the automated planner; during this process, called "replanning", the operator evaluates the plans, may make modifications to the plans, and finally approves an updated plan for the vehicles' behavior in the near future.

Previous research with the OPS-USERS framework has shown that the frequency at which the operator engages in the replanning task significantly impacts system performance, with high replanning rates resulting in degraded system performance [6]. Choosing a replanning frequency therefore represents a trade-off between the benefits of re-optimizing the plan more frequently against the cost in operator situational awareness and cognitive capacity available to other tasks. Furthermore, it is not necessarily beneficial to replan frequently aside from the effects of increased operator workload. Replanning on short times scales with noisy information has been demonstrated to reduce system performance as the difference between several near-optimal plans may be small, causing consecutive replanning events to produce inconsistent results. This undesirable effect, referred to as "churning", reduces the efficiency of the system; it can be eliminated by actively filtering the data, but is also less prevalent when replanning occurs less frequently [7]. The "replan interval" also influences the trade-off between vehicle time spent on specific tasks and time spent searching over unmapped areas. This leads to the following question in the context of degraded communication environments:

Research Question #3: *To what extent does HITL control increase the robustness of autonomous planners to communication failure?*

This question deals primarily with whether the optimal replanning interval changes in degraded communication environments, or more broadly, whether the parameters of the HITL control can be adjusted in order to compensate for the negative effects of communication failures on performance. This research will focus on the effects of the HITL control parameters through simulation due to the infeasibility of performing a large enough set of live experiments.

1.3.3 Thesis Organization

This thesis is organized into the following chapters:

Chapter 1 Introduction: Describes the motivation and research questions for this thesis.

Chapter 2 Background: Reviews previous work in fields related to decentralized unmanned vehicle control, including automated planning systems, the effect of network topologies on distributed system performance, and human supervisory control.

Chapter 3 Experiment Design: Describes the methodology for gathering system performance data, including the implementation of a discrete event model simulating human interaction with the OPS-USERS system. Additional background on the OPS-USERS system implementation is also provided.

Chapter 4 Results: Demonstrates validity of the human model used in simulations, and details how performance was affected by communication degradation and the parameters of HITL control.

Chapter 5 Conclusions and Future Work: Revisits research questions in the context of the results, and suggests areas for further investigation.

Chapter 2

Background

Development of sophisticated decentralized UASs draws together a diverse set of fields including human supervisory control, automated planning for coordinating vehicle behavior, and network theory covering communication between agents in the system. This chapter discusses research conducted on automated planning algorithms for multi-agent systems, research in human supervisory control for operators in decision support roles, and research on networks and the effect of degraded communications on the performance of UASs or other distributed systems.

2.1 Automated Planning

The problem of coordinating multiple unmanned vehicles becomes more complex as the number of vehicles is increased or systems are tasked with more sophisticated missions. Because a human operator would be quickly overwhelmed in assigning tasks to individual vehicles throughout the course of a mission, an automated planner must be included to assist in coordinating and assigning tasks. Much work has been done to develop algorithms capable of efficient planning under sufficiently complex constraints to be useful in real-world implementations. This section discusses the automated planning problem formulation as well as a variety of planning algorithm implementations and their suitability for command and control networks.

2.1.1 Automated Planning Problem Statement

A mission is defined as a high level goal, such as searching a region, to be completed by a heterogeneous set of n UAVs as given by $\mathbf{V} = \{v_1, \ldots, v_n\}$, with guidance provided by a GCS, g. The set of agents in the system is defined to be the vehicles and the GCS: $\mathbf{A} = \mathbf{V} \cup g$. The high level goal is translated into a set of active individual tasks, $\mathbf{T} = \{t_1, t_2, \ldots\}$ dynamically maintained throughout the course of the mission. New tasks are added to $\mathbf{T}$ as they arise or are created through human intervention, and tasks

are removed as they are completed or expire. The goal of a task assignment algorithm is to distribute the tasks currently in $\mathbf{T}$ between the UAVs in $\mathbf{V}$, based on knowledge of the current state of the system including the location of each vehicle, location of tasks, and task priorities to maximize an objective function. This problem is complicated by the inherent uncertainty in any estimation of the system's state arising from sensor noise and imperfect communication. Each agent in the system may maintain its own local version of an estimate of the system state, which could be inconsistent across agents. In this thesis, an agent's local version of an estimate of the system state will be referred to as that agent's *beliefs*, similar to the Bayesian concept of beliefs although not necessarily considered in a Bayesian context.

The automated planning problem of assigning the active tasks in $\mathbf{T}$ to the agents in $\mathbf{V}$ can be decomposed into two distinct components:

(1) Data Fusion: Data fusion is a mechanism for agents to share information and update their own beliefs based on the observations of the other agents [8].

(2) Task Assignment: The agents compute assignments of tasks to specific vehicles by maximizing an objective function over possible assignments. During this process, a human supervisor may provide input to further prioritize the available tasks, coaching the automation and influencing the resulting task assignment. The task assignment step also involves communicating assignments to the specific vehicles [4].

2.1.2   Data Fusion

Information about the state of the system represents the basis for planning decisions. This information may include:

- Locations of the vehicles and known targets
- Known target priorities, estimated trajectories, and other information
- Environment map indicating the likelihood of undiscovered target locations

As the quality of the state estimate of the system directly impacts the quality of plans which can be generated based on this information, it is desirable to achieve the best system state estimate possible before computing task assignments [9]. Individual agents in the system make their own observations of the environment and maintain their own beliefs about the state of the environment. Due to possible heterogeneity of sensors per agent or dispersion of agents throughout the environment, one agent's observations may cover a different portion of the environment or otherwise contain information not present in any other agent's observations. Even if two agents make observations of the same area, it remains beneficial to share the information, as combining multiple noisy observations allows for a more accurate estimate of the system's state [10].

26

In an ideal system with a strongly connected communication network (i.e. two-way communication between all agents), unlimited bandwidth, and no latency, agents would not have trouble reaching aagreement over their beliefs by taking all available information into account. In real-world systems, the costs of sharing information must be considered, including time and bandwidth utilization. Constraints on the bandwidth of communication between vehicles, or on the amount of time afforded to data fusion may make it infeasible for all vehicles to converge to an identical set of beliefs. Alternatively it may only be possible for all agents to converge to the same set of beliefs if some vehicles' information is not taken into account; for example, if one vehicle can receive but cannot send information, then its beliefs cannot be taken into account by the remaining vehicles. Agreement on beliefs can only be reached by ignoring the observations of this vehicle. The performance of automated planning systems may depend closely on both achieving the best possible estimate of the system's state and converging to a single set of beliefs across all agents. Several data fusion approaches are discussed below.

Data fusion can be accomplished through a variety of information-sharing schemes, and can be performed either continuously or immediately preceding each task assignment computation. Furthermore, different implementations of the task assignment step may benefit from a higher degree of convergence such that all agents have similar beliefs about the system. Alternatively, an implementation might achieve the best performance when a high quality estimate of the system state is built in just one or some of the agents. For example, the Consensus Based Bundle Algorithm (CBBA) and Robust Decentralized Task Assignment (RDTA) algorithms employ an auction framework which allows vehicles to bid on the tasks that are easiest for them to complete. Any conflicts where two vehicles would prefer to perform the same task are eliminated in an iterative process, resulting in a conflict-free assignment. In this setting, the vehicles need not reach agreement on beliefs, but agents with higher-quality estimations of the system state can continue to improve the systems' performance [4, 11, 12]. Data fusion algorithms may be designed either to spend more time reaching an agreement on beliefs before replanning, or to terminate after a short time and more promptly initiate planning before agreement is reached.

Furthermore, the availability of communications directly impacts data fusion algorithm performance. Algorithms can be optimized for one static network topology, or can be designed to tolerate dynamically changing topologies. A consideration of the characteristics of the task assignment step along with the availability of communication between agents and time constraints drive data-fusion algorithm selection decisions.

Data Fusion Techniques

Although data fusion research is considered here only in the context of automated planning, data fusion algorithms are applicable in a variety of other multi-agent coordination problems [13]. Surveys of data fusion research provide an overview of simplified data fusion problems as the basis for more

sophisticated data fusion schemes. Simplified problems may assume static and reliable communication networks, infinite bandwidth for communication between agents, static final values over variables, or systems with state-spaces on the order of a few variables [9, 14]. These approaches can then be extended to operate under conditions with degraded communications and dynamically evolving variables. This is the case in the OPS-USERS system, where the data to be combined includes vehicle locations, target locations, and target trajectories.

Simpler problems for multi-agent coordination include vehicle formation control [15, 16], coordinated rendezvous of vehicles [17, 18], and coordinated task planning [19–21]. One specific application example includes coordinating actions within a network of decision makers in Ref. [22]. The decision makers represent buyers who attempt to set a threshold price based on the threshold prices of the other buyers so that all buyers act optimally. The agents effectively collaborate on decision making problems using a distributed data fusion protocol. This system requires only local communication, but is shown to achieve performance equal to a less sophisticated centralized approach where all buyers communicate their threshold strategy to a single agent who then performs a local optimization with complete information. A separate application focuses on distributed noisy sensors observing a single event and later collaborating to determine the maximum a posteriori (MAP) estimate about the event [10].

Another approach to data fusion building employs the ubiquitous Kalman filter [23]. Agents record both observations and a relative measure of confidence (covariance matrix) about the information. Agents communicate their best estimate of the system state and their confidence level to each other, and each agent updates its local estimate based on the other agents' estimates to improve agreement on beliefs. Under perfect communication (no latency, unlimited bandwidth, strongly connected communication network), this approach performs well; however, extensions to the Kalman filter are generally required to increase robustness to unreliable communication [24, 25].

Often, a simplification of the Kalman filter solution may be viable. One approach involves computing the average belief over the system, and has been shown to achieve complete agreement on beliefs over a variety of network topologies [26]. This particular procedure, referred to as *consensus propagation*, represents a protocol for distributed averaging. Confidence levels associated with beliefs are not taken into account explicitly, instead trading accuracy in state estimation for better agreement on beliefs. This trade-off may be desirable because the algorithm affords several advantages: the system can operate in certain degraded network topologies, the algorithm can be performed asynchronously, and each agent needs only to be aware of its neighbors. In situations where communication is expensive or otherwise limited, consensus propagation may be an appropriate choice. Another averaging technique is described in [15], where vehicles attempt to maintain a specific formation by communicating their measurements of relative vehicles positions to one another. This system has been shown to be robust to considerable communication loss. The vehicles were able to consistently converge into formation with as much as

two-thirds of all communication lost.

Building upon techniques discussed above, work has been done to develop extensions to handle the complexities inherent in UASs. This includes relaxing the assumption that perfect communication is available, and dealing with finite bandwidth and latency in communication. Theoretical research into several Kalman-like linear and non-linear data fusion protocols has demonstrated that agreement on beliefs is possible in the presence of communication latency or switching network topologies [27–29]. The analysis in this research provides conditions on the types of networks which allow these protocols to converge, for example, that a network be connected *on average*. In one example implementation of a data fusion algorithm in a UAS, the algorithm described in [9] was implemented in simulation in [4, 30], although the UAV mission considered represents a considerable simplification of the complex OPS-USERS mission considered in this work.

2.1.3   Task Assignment

The specification for any task assignment algorithm simply requires generation of a list of tasks assigned to each vehicle, and communication of assignments to corresponding vehicles. Implementations of task assignment algorithms vary in that the required computation may take place on one specific agent (for example, the GCS), distributed across several agents, or may occur simultaneously on all agents. There are two primary categories of task assignment algorithm implementations: centralized algorithms and decentralized algorithms. A third alternative, quasi-decentralized systems, is a variant which combines elements of both centralized and decentralized systems.

In centralized architectures, each vehicle communicates its beliefs to a single agent, which acts as a Centralized Controller (CC). This agent is presumably more capable of calculating optimized assignments, either through augmented computational power, or by representing a central node in terms of communication topology. In this case, the data fusion procedure is formulated such that the central agent receives the best possible estimation of the system state. The central agent then computes an optimized task assignment based on this estimate, and sends the resulting assignment to each vehicle. This architecture aggregates all available information together into the central agent, and considers the entire optimization problem at once. A human supervisor may be given oversight of the system by interacting with the central node, to provide external guidance to the centralized task assignment algorithm.

An alternative approach is a decentralized architecture where task assignment optimization takes place across several or all of the agents using on-board computers. Each agent may be responsible for computing their own task assignment, or a quasi-decentralized hierarchical approach may partition agents into groups (for example, based on location) with one agent from each group selected to plan for the entire group. Although decentralized systems do not rely on a CC for task assignment, some decentralized systems such as the OPS-USERS system include a centralized element which facilitates

human oversight of the decentralized planner. In this case the CC may never directly compute task assignments for the vehicles, but it rather facilitates external guidance when communication availability allows it. This will be discussed in greater detail in a subsequent section.

2.2 Network Topologies and Distributed System Performance

The communication networks present in current UASs contain just two or a few nodes, including a centralized controller plus one or more vehicles. Small networks such as these receive little benefit from the consideration of their network topologies, as routing of communications is of little practical value when the vehicles are required to be in contact with the CC and inter-vehicle communication is not emphasized. However, as UASs begin to incorporate more and more vehicles, and collaboration between vehicles becomes standard behavior, synergistic routing strategies may begin to impact system performance. For example, in networks with as few as 10 nodes, routing protocols between agents in a dynamic environment were demonstrated to significantly affect network throughput and connectivity [31]. The number of vehicles in future UASs will grow large enough to introduce challenges in managing the routing topology of the communication network in order to maintain acceptable performance. Current and future UASs also face challenges in connectivity and bandwidth concerns, time constraints, and dynamic link availabilities. Using techniques from network theory, the communication networks in UASs can be characterized numerically, facilitating investigation into the effects of changing network conditions on system performance.

Figure 2-1: Distinct subgroups of UAVs, decentralized architecture

Distributed planners respond to communication failures differently than centralized planners in ways which enhance robustness to certain types of communication failures. For example, when communication limitations partition UAVs into isolated groups with locally complete communications networks (Figure 2-1), vehicles continue to collaborate with their neighbors (i.e. vehicles 1, 2, and 3 will collaborate, and vehicles 4, 5, and 6 will collaborate). This illustrates one of the main advantages of distributed systems: if tasks are weakly coupled across long distances (as communications tend to be weaker across long distances), then solving several local task assignment problems will tend to produce results more closely matching results obtained by computing the entire task assignment in a central location [32]. In this sense, distributed systems scale more easily relative to their centralized counterparts, as no new hardware is

required beyond the hardware already present on each UAV, and the computational requirements on each agent in the system are not increased, as would be the case in a centralized system [33]. The computational requirements do not become burdensome as the size of the optimization performed in each small group can remain relatively constant as the system size is increased.

Where centralized planners attempt to take advantage of all available information about the state of the system in one large optimization problem, decentralized planners solve several smaller local optimization problems. In situations where communication between UAVs is perfect, and without computational limitations imposed by the computer platforms available on the vehicles themselves, we would expect the performance of the centralized and decentralized approaches to be the most similar. Indeed, much initial research on new task assignment algorithms is done under the assumption that communication is always available and instantaneous between all agents [34–36]. However, in real world applications the limitations of communication availability force the planners to operate under sub-optimal network topologies, introducing latency and inefficiency into the system.

For example, a system which understands its communication topology can dynamically adjust its behavior to maximize the usefulness of communication resources on hand. Given a specific network topology, we may wish to predict the runtime or communication requirements of a data fusion algorithm in order to manage the amount of time it will take to generate a task assignment [27]. With the ultimate goal of balancing the quality of the task assignments generated and the latency with which new task assignments can be computed, a runtime estimation could aid decisions for the type of data fusion algorithm to employ or the number of iterations to perform given the current network topology. Work has also been done to investigate convergence properties of data fusion algorithms; certain network topologies may guarantee convergence or make convergence impossible [29, 37]. Analysis of what level of agreement can be reached by data fusion algorithms and at what cost in time and communications informs the system of an expectation of performance and whether adjustments to data-fusion algorithms would be possible to increase performance (see Ref. [4]).

2.2.1   Measuring Communication Availability

In order to address the research questions outlined in Section 1.3.1, metrics must be calculated which numerically quantify the level of communications available. For example, in RDTA research, the amount of information communicated was limited during both the data fusion and task assignment phases separately, yielding performance plots showing how performance was degraded relative to the case of unlimited communication [30].

In another approach, artificial constraints on communication were defined in an otherwise realistic simulation [38]. This experiment modulated the distance over which communication was allowed while testing whether the robots were able to successfully traverse their environment. Results demonstrated that

successful navigation was possible under several regimes of limited communications, and that sacrificing optimality to simplify a planning problem leads to acceptable solutions in cases complex enough to be of practical interest. However, the work lacked a quantifiable metric for measuring the relative performance of the system at each level of communication, and therefore failed to characterize the relationship between performance and communication availability.

Another approach to characterizing the performance of information-sharing algorithms over a variety of network topologies applies probabilistic analysis of the possible network topologies which can arise in a network [39]. Using this method, the network characteristics represent the "performance" and the parameters of the underlying model for generating the networks (for example, the probability of a specific communication link being present in a given time step), represent the independent variables. As the networks generated by the model evolve through time, they may only periodically achieve a "useful" state where information can be exchanged. The fraction of time spent in useful states is investigated through simulation. The characteristics of each random network drawn are then evaluated and averaged to gain insight into the levels of performance that can be expected from stochastic networks given link failure probability characterizations.

2.2.2 Communication Topologies and Automated Planning

The network topology of a UAS will influence the preferred architecture for autonomous planning. Centralized and decentralized architectures dictate dissimilar communication requirements in terms of the type of information transmitted, which pairs of agents communicate, and the amount of information transmitted. The theoretical performance characteristics of the centralized and decentralized approaches will be considered in the following sections, accompanied by a consideration of several notable implementations from both algorithm classes. These characteristics have implications for each architecture's suitability in a command and control network setting, as they influence fragility to both intermittent communication availability and inconsistencies in beliefs across agents.

Centralized Task Assignment Architecture

Centralized planning architectures rely on one or a small number of agents in the system to act as a CC. Figure 2-2 represents a possible configuration of a centralized planning architecture. This example demonstrates the benefits of designating one agent in the system specifically suited for calculating task assignments. The central agent may trade mobility for augmented computational power, shifting computational burdens away from the remainder of the agents. This allows cheaper and lighter construction of the remaining agents, and further specialization for other purposes. Upon receiving the beliefs of each vehicle, the central agent generates task assignments and communicates them back to the remaining agents. An analysis of the types of network traffic generated by a centralized planning architecture, as well as how this traffic would be affected by a specific network topology follows in the next section.
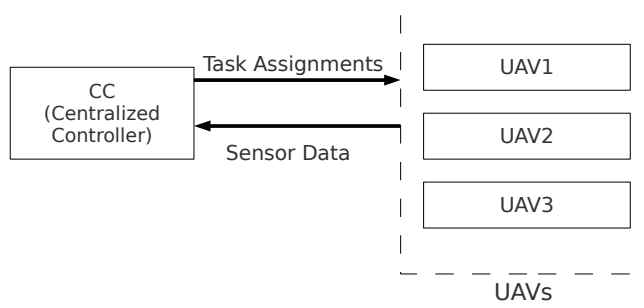
Figure 2-2: Centralized architecture

Centralized Communication Requirements In centralized planning, the central agent receives observation information from each agent and performs the information update step locally, reconciling the possibly conflicting beliefs from all other agents into a local best estimate of the system state. The content of the belief information received from agents determines the bandwidth and time requirements for the information update step; however, each vehicle's data needs to be sent across the network to the CC only once per planning interval. In the task assignment stage, the CC calculates a task assignment and transmits it to each vehicle; although this is a low-bandwidth communication, it must reach every agent.

One immediate advantage of a centralized architecture is that centralized data-fusion can be accomplished through a single transmission of belief information. This avoids latencies in the task assignment process associated with most decentralized algorithms, where multiple iterations of communicating and updating belief information across the network may be required [40]. However, the CC represents a single point of failure in the sense that each vehicle must be in communication with the CC when the planning algorithm executes so that its beliefs are taken into account and it can be assigned tasks. Although the risk of failure or loss of the central node may be mitigated by either adding redundancy, allowing another agent to step into the central agent role, or delegating some decision autonomy to the decentralized system in the event of a communication failure, this still implies sensitivity to temporally intermittent communication links. Asynchronous task assignment procedures that account for the characteristics of communication availability (such that out-of-contact agents may receive delayed task assignments rather than no task assignments at all) have also been proposed to mitigate this specific effect [26]. However, such an approach carries significant operational risks in command and control settings, especially when task assignments are time critical.

In a centralized setting, the characteristics of the communication channel between each vehicle and the CC is the dominant communication factor in determining system performance. This includes whether communication is available at all, and if so, the bandwidth, probability of data loss, and latency of the communication channel. In a command and control context, centralized systems represent a conventional hierarchy with decisions being made at a central location (in the CC) and handed down to the subordinate vehicles. In terms of network connectivity, the degree of the central node is high (many agents

communicating directly with central node), but each vehicle needs only to communicate with the CC and little or no inter-vehicle communication is necessary. This translates into a low communication density over the network of agents. However, the concentration of information at the centralized node imposes high computation requirements on the CC.

Centralized Network Topologies The actual planning algorithm implemented inside the CC may vary, but the characteristics of information flow in a centralized system will be similar across implementations. Such an architecture places constraints on the types of network topologies required. For example, vehicles experiencing communication failures may be unable to contribute their belief information to the CC or may be unable to receive task assignments. Also, noisy communication links will increase the time required to transmit belief data to the CC, especially for increasingly complex belief information including imagery or other information requiring high bandwidth communication links. For the task assignment algorithm to operate nominally, all agents must have access to a two-way communication link with the CC with sufficient bandwidth.

In a centralized architecture, communication with the CC need not be direct, or equivalently, the network of agents participating in the centralized task assignment must only be strongly connected; communications may be *routed* through one or more intermediate agents [9]. An important note about centralized architectures is that a subset of vehicles able to communicate locally but which are all unable to communicate with the CC cannot take advantage of their locally connected network to collaborate amongst themselves. For example, in Figure 2-3, vehicles 1, 2, and 3 will operate nominally with the CC, while vehicles 4, 5, and 6 will fail to collaborate with the CC or amongst themselves.

Figure 2-3: Distinct subgroups of UAVs, centralized architecture

There are three distinct types of directed communication links in a centralized system: UAV-to-UAV links, UAV-to-CC links, and CC-to-UAV links. Permanently breaking any link will affect the system differently depending on the type of link that is broken. Specification of a failure mode implies not only that the two agents cannot communicate directly, but also that no route exists through the remaining nodes. The implications of the failure of each type of communication link are outlined below.

- **CC-to-UAV link broken**: Vehicles that cannot receive messages from the CC cannot receive task assignments. They may either continue their mission under the assumption that contact with the CC will be restored eventually, or abort the mission and return to base. Assuming that the vehicles

continue the mission during periods of communication failure, they may continue to act on the last task assignment they received. Upon completion of assigned tasks, they may then carry out implicit tasks which do not require specific assignment from the CC. This may include searching for targets and attempting to send their own sensor data to the CC, as the directional link from the vehicle to the CC may not be broken. Vehicles in this state can be required to attempt to restore contact with the CC at regular intervals, for example by moving to a new location.

- **UAV-to-CC link broken**: Vehicles which cannot send messages to the CC will not be able to send updated belief information, and as a consequence sensor data from these vehicles cannot be taken into account in new task assignments. The CC has a choice in this case either to assume that the vehicle can presently receive messages from the CC, or to exclude the vehicle (and possibly vehicles with communications from the CC routed through it) from the set of available vehicles when optimizing the task assignment.

- **UAV-to-UAV link broken**: This link failure does not necessarily affect performance, but based on the assumption of network routing capability, this implies also that affected vehicles cannot have a two-way communication link to the CC; therefore this mode must be accompanied by one of the failure modes listed above.

Centralized Task Assignment Implementations Previous work includes several UAS simulation implementations with centralized task assignment architectures. Assuming that the data fusion step has already formed the best possible estimate of system state in the CC, the task assignment algorithm has only to do a local optimization to calculate the task assignment. The optimization can be posed as a network flow optimization [36], or a trajectory optimization using model predictive control [35]. Locally simulated auctions [4, 11, 34] have also been employed, centralizing the inherently distributed behavior of each vehicle bidding on tasks by fusing belief information at a single central location and then emulating the behavior of each agent in the system at the central location; this could be considered tantamount to assuming perfect information sharing.

A common simplification in theoretical work on task assignment algorithms is to assume perfect communication between vehicles [34–36]. This effectively abstracts away the algorithm execution and the model of the environment from the vehicles, and allows all information to be shared perfectly, minimizing the burden that a data fusion procedure places on task assignment. In the context of algorithm architecture, these approaches will be considered centralized, as the strong assumptions about communication network connectivity would allow a centralized approach, and the complexity of the planning algorithm often suggests the requirement of a powerful ground-based computer, as in [35]. However, in some cases such as [34, 36], the algorithm is simple enough to be run in on-board hardware; in this case the algorithm could be implemented in either a centralized or decentralized manner with identical results under perfect

communication.

One algorithm which applies the perfect information sharing simplification puts the Centralized Mission Manager (CMM) in the role of an auction mediator [34]. The vehicles first compute which task they prefer by minimizing the value of a simple cost function over each available task, and these costs are sent to the CC. The CC then processes these choices to resolve conflicts, preferring the vehicle which incurs the lowest cost if a task is selected by more than one vehicle. This procedure is iterated until all tasks have been assigned. Benefits of this approach include very light computational requirements, as the vehicles need only to compute their cost function for each task, which can be structured to account for UAVs with different capabilities and even varying levels of expertise in each capability. This algorithm is straightforward to decentralize if agreement on beliefs is attainable. However the simplicity of this approach renders the task assignments produced by this algorithm generally sub-optimal [34]. More sophisticated truly decentralized auction algorithms will be discussed below.

The previous approach represents a greedy algorithm to minimize the cost of completing as many tasks as possible; the vehicles consider only minimizing their own cost rather than the cost for the entire system. This approach trades speed and simplicity against searching only a very small portion of the possible task assignments. An improvement would be to allow the CC to consider the cost function for all vehicles and tasks simultaneously in the context of a linear programming problem. Then known linear programming techniques can be employed to optimize performance [36, 41].

Another approach considers the entire problem as a vehicle trajectory optimization where the dynamics of the vehicles are simplified in order to abstract them away from the high level planner [35]. In this case, trajectories are computed for each vehicle for some planning horizon into the future, long enough such that at least one vehicle will reach a target within the horizon. Rewards are given to trajectories bringing UAVs near targets, or for maximizing the likelihood of discovering new targets. In this representation, task assignments are implicit in the simplified trajectory planning rather than in an explicit assignment of tasks to vehicles. This approach allows application of existing trajectory optimization techniques, and another interesting feature of this method is direct consideration of the useful lifetime of plans generated. Planning happens as infrequently as possible; for example, the algorithm replans only if a new target is detected, or the system state deviates significantly from the expected state. If the current plan contains only long-term goals, then replanning may not be necessary for a long interval, minimizing communication requirements.

Decentralized Task Assignment Architecture

A decentralized planning solution relies on on-board computers to generate task assignments. Removing the constraint that all information must be transmitted to a single agent in the system (i.e. the role of the CC in centralized planning) requires a new strategy to account for each vehicle's beliefs

in computing task assignments. Figure 2-4 illustrates a simple decentralized architecture. Note that a centralized component akin to the CC in a centralized system may still be present to provide a facility for human oversight of the distributed system; this is the case in the OPS-USERS system. This centralized component will be referred to as the Centralized Mission Manager (CMM). A decentralized system may
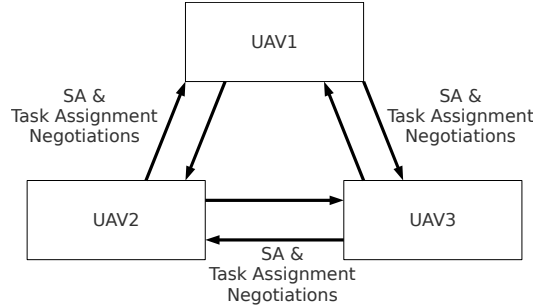
Figure 2-4: Decentralized architecture

either implement a data fusion procedure to reach agreement on beliefs to be used to calculate task assignments [23, 26, 42], or be allowed to tolerate some inconsistencies in beliefs in a scheme where the vehicles interactively distribute tasks using an auction framework [4, 11, 40]. Auction frameworks may require multiple iterations to allow each vehicle to bid on each task, but provide protection from inconsistencies in belief in some cases. Each implementation involves different communication requirements; however, a major contrast to centralized architectures is that an increase in total communication required can be traded off against network topology flexibility.

Decentralized Communication Requirements Decentralized systems offer more flexibility than centralized systems in their communication requirements, allowing trade-offs between toleration of limited communication availability and support for increasingly sophisticated optimization algorithms. Decentralized automated planning can still be considered as two components, data-fusion and task assignment. In contrast to centralized systems, decentralized implementations for the task assignment component exist which are optimized for a wide range of communication availabilities and levels of belief agreement. A data fusion algorithm can be chosen to best match the data fusion requirements of the task assignment algorithm, and both choices dictate communication requirements. Three possible approaches to choosing appropriate data fusion and task assignment implementations for a decentralized planning algorithm are outlined below.

One strategy is to require that the data fusion algorithm converges, such that all vehicles agree on the system state before performing task assignment. This approach guarantees conflict-free task assignments and allows implicit coordination, where each vehicle plans for itself; this is only possible when each vehicle runs the same planning algorithm over the same set of beliefs. Requiring convergence can incur significantly longer latencies, or may fail if an agreement on beliefs cannot be reached under current communication conditions. This strategy is most suitable when communication is reliable.

A second strategy could be developed that place a limit on the time or bandwidth committed to data fusion can control latency in the data fusion procedure. This can be implemented through the use of an *anytime* algorithm to perform data-fusion, where the algorithm can be terminated after a set time interval or a fixed number of iterations even if complete agreement on beliefs has not been achieved. This approach helps bound latency, while taking advantage of data-fusion convergence when the network topology is favorable for efficient convergence. However, the planning algorithm used in conjunction with this data fusion strategy must be capable of planning on inconsistent information across the vehicles.

Under constrained communications, it may be desirable to eliminate the data fusion component altogether in order to reduce communication requirements. Communication will only be required for the task assignment step, at the expense of the UAVs attempting to plan on conflicted beliefs. This strategy is suitable when communication is constrained. A careful consideration of the trade-off between agreeing on beliefs about the state of the system and the ability to plan rapidly dictates which approach is most desirable for a particular application. To this end, a variety of distinct implementations exist to assign tasks, outlined below.

Decentralized Network Topologies As for centralized planners, specifying a failure mode implies no direct communication nor indirect communication through other nodes. In the decentralized case, there is only a single type of link to consider, a UAV-to-UAV link. If a UAV-to-UAV link is broken, this failure condition implies a partitioning of the vehicles into two or more groups. Each group will continue to collaborate amongst themselves. Conflicting assignments may arise if two vehicles in different groups attempt to complete the same task.

Decentralized Task Assignment Implementations Decentralized task assignment implementations are diverse; implementations exist which are optimized to operate under a variety of communication network topologies, and with varying levels of agreement across the vehicles. Decentralized task assignment algorithms may be direct adaptations of centralized algorithms, where global belief agreement is reached and each vehicle takes the role of the CMM to calculate their task assignments independently; in this case, all vehicles arrive at the same task assignment since the algorithm is deterministic and in each case was run using the same beliefs [34–36]. Auction-style algorithms represent the other principal decentralized architecture, where plan optimality is sacrificed for reduced communication overhead and increased planning speed.

Many decentralized approaches can be interpreted as decentralizations of centralized approaches. A straightforward implementation utilizes a data fusion scheme to build global belief agreement such that each node has identical beliefs. Then each node performs the centralized planning algorithm locally to achieve the same result that a centralized planner would achieve. [34–36, 41].

Another decentralized approach considers the data fusion and task assignment stages of the algorithm

simultaneously using a probabilistic model for the system – a Partially Observable Markov Decision Process (POMDP). Similar to the Kalman filter, the POMDP framework handles uncertainty directly in the model of the environment. Since multi-UAV systems are highly dynamic and stochastic, it is natural to capture the uncertainty of the problem directly. Furthermore, rather than basing task assignment decisions on the most likely estimate of the system state calculated through a data fusion algorithm, the POMDP model captures each possible system state along with its likelihood in planning decisions. However, these characteristics come at a performance cost which causes POMDP algorithms to struggle to handle even small problem sizes due to enormous computational requirements. Recent work using decentralized POMDPs for task assignment includes the Multi-Agent A* algorithm (MAA*) [43]. As solving an optimization of the full POMDP process is intractable, the algorithm takes into account each agents' probabilistic model of the system to act as a heuristic for a dynamic programming search for solutions to the task assignment problem. Results have shown this approach to generate high-quality task assignments. However, this algorithm still suffers from robustness problems in that tolerance to any communication failures is not a feature of the algorithm, and perfect communication was assumed in Ref. [43].

Distributed auction-style algorithms represent another decentralized approach which trade-off optimality for speed. Auction-style algorithms provide quick solutions at the expense of the quality of task assignments generated. In certain implementations, no attempt to reach agreement on beliefs is made [44, 45]. Rather, the vehicles act greedily to take the tasks which they believe are easiest for them to complete. The vehicles iterate over the available tasks, with each vehicle bidding its estimated cost of completing the task. The task is given to the vehicle with the lowest cost. Despite their simplicity, auction algorithms offer competitive performance amongst available decentralized task assignment algorithms.

In one particular implementation, each UAV calculates the sets of tasks it could feasibly accomplish, and the cost of accomplishing that set of tasks (in both distance, time, or another cost metric). Then sets of tasks are chosen for each vehicle to maximize an objective function [46]. This is known as the petal algorithm, and forms the basis of two notable extensions: CBBA [11] and RDTA [4]. The CBBA algorithm is a light-weight version which minimizes computational and communication requirements, making few assumptions about network connectivity; vehicles must only be able to communicate with their neighbors, and information cannot be relayed between distant vehicles. The UAVs communicate their situational awareness to each other along with a suggested task assignment for themselves until an agreement on task assignment is reached. Agreement on beliefs is not necessary. CBBA provides excellent run-times for even large problems; however, the optimization performed is not substantial, and data fusion over vehicle beliefs is not a component of the CBBA algorithm. On the other hand, RDTA does fuse belief information and conflicting beliefs are resolved through further iterations of information sharing steps. After agreement is reached, the vehicles iterate on the auction of tasks, until a conflict-

free assignment is found. The emphasis on data fusion in RDTA comes with heavier requirements for communication between vehicles. This also results in longer times to convergence, but RDTA is guaranteed to converge to a conflict-free task assignment for all vehicles connected to the network – a guarantee CBBA doesn't make when discrepancies between each vehicles' beliefs are large. The task assignments of the RDTA algorithm have been shown to offer performance enhancements over CBBA's task assignments, representing the trade-off between planning speed and planner performance which is possible in auction algorithms.

Quasi-Decentralized Task Assignment Architecture

The quasi-decentralized task assignment architecture combines elements of both centralized and decentralized task assignment algorithms [47]. For example, a hierarchical approach instantiates the CMM on a subset of the vehicles, effectively splitting the vehicles into teams. Then leaders of each team collaborate among their peers and dictate to their subordinates, yielding a planning hierarchy [48].

Hierarchical methods reduce the dependency on data-fusion and consistent beliefs across vehicles [49–51]. By partitioning the vehicles into groups, collaboration occurs across the system in terms of higher-level goals between the groups. Within each group task-level optimization occurs across locally dense communication networks. This approach decouples portions of the task assignment optimization problem in order to reduce the communication requirements across the system and to reduce the complexity of each task assignment optimization performed. Furthermore, splitting the vehicles into teams with dense communication capabilities allows optimization across the subsets of vehicles best equipped to collaborate. Hierarchical approaches represent some trade-off between communication requirements and planner performance. In systems with larger numbers of vehicles, this trade-off is often desirable; many aspects of the system can be decoupled without affecting the results of the algorithm.

Decentralized automated planning with a HITL component represents another quasi-decentralized implementation. In such a system, task assignments are computed via any decentralized algorithm implemented across the vehicles, although the resulting assignment must be evaluated and approved by a human supervisor in a centralized location before it can be executed by the vehicles. From a command-and-control perspective, this architecture emulates a centralized point of command while simultaneously taking advantage of the robustness provided by distributed planning solutions in the case of communication failures.

2.3   Human Supervisory Control and Task Allocation

Human supervisory control of autonomous systems plays an important role in achieving high system performance. In previous research, a commercial airline flight planning system with a supervisor, which will also be referred to as an operator, interacting with an automated planner has demonstrated both

possible benefits and pitfalls of such interaction [52]. Benefits of the automation included increasing the probability that the best flight plan was selected in certain scenarios by assisting the user in searching the large solution space. However, the automation also made poor suggestions under circumstances where the automation's model was not sufficiently sophisticated or did not consider all relevant factors. These poor suggestions subsequently had a negative influence on the operator's ability to select the best flight plan as the operators relied too heavily on the automation. Furthermore, evidence suggested that presenting a larger amount of data to the user caused critical data to be overlooked.

Previous work has also demonstrated that modes of human-automation collaboration can significantly impact system performance [53]. The automation and the operator must communicate to achieve mutual understanding of intentions towards solving shared goals. By framing interaction between a human operator and the automation to minimize the disruption of the operator's normal cognitive processes, the requirements for this communication can be greatly reduced. Furthermore, in the context of supervisory control in a UAS, a system which allows a single operator to control multiple UAVs demonstrates a need both to develop systems with high levels of autonomy and simultaneously address the challenges multiple vehicle control presents in supervisory control [54]. These challenges include developing vehicles with a high level of autonomy, and addressing the cognitive abilities of a human operators interacting with this autonomous system. The automation must be designed to reduce the operator's workload, while simultaneously maintaining operator situational awareness to minimize the danger of over-reliance on the automation.

The majority of previous human-automation collaboration work addresses design considerations and models of human-automation collaboration. There has only recently been progress towards quantifying the extent to which humans can aid algorithms in a cooperative manner in the context of UASs [6].

Human Operators and Task Allocation When faced with a task assignment problem, human operators and automated algorithms apply significantly different methodologies towards generating a solution [55]. Acknowledging and understanding the strengths and deficiencies of automated algorithms provides guidance for the development of decision support utilities for humans solving task assignment problems. Specifically, controlling a system of multiple UAVs requires a method for assigning available tasks to individual vehicles to maximize collaboration. HITL control, also known as human supervisory control [5], represents one strategy for taking advantage of possible synergies between automation and a human operator.

The primary advantage of automated algorithms lies in their ability to efficiently consider a large set of possible solutions in a short period of time to provide task assignment solutions which are optimized over some quantifiable objective function. However, the limitation of any automated task assignment algorithm arises from the inability of the objective function to perfectly coincide with the operator's

goals in all cases. Furthermore, the automation can be *brittle*; it may make assumptions which do not hold in every situation, possibly leading to important factors being accounted for incorrectly. In these cases the automation's decisions may degrade system performance or may even be dangerous [52].

These two issues – objective function mismatch and brittleness – may result from a failure to anticipate specific situations and how the system would react to them, or even a deliberate oversimplification of automation design. If the automated planner makes poor decisions in situations outside of the range of possibilities accounted for in the automation's design, due to either brittleness or objective function mismatch, inclusion of a mechanism for oversight of the automation's decisions is required to ensure safe and reliable operation. HITL control is one method to introduce this oversight.

Humans approach task assignment problems from a more qualitative perspective, although they are also capable of quantitative optimization (to a significantly lesser degree than computers). However, the most valuable advantage human operators provide lies in the ability to adapt to any situation which may arise, even completely unanticipated situations where goals vary significantly from the original purpose of the system. In the case of multi-UAV systems, this may include adapting a system to unanticipated environments or to novel enemy behavior. Working without automation assistance, however, a human operator can quickly be overwhelmed by the requirement to manually specify the set of tasks for each vehicle, degrading overall system performance.

An automated system collaborating with a human operator in the decision-making loop can generate synergies such that the resulting task assignments benefit from the precise (yet possibly misguided) insight of the automation, along with the qualitative and adaptive power of the human. When the situation falls well within the domain of the automation's model, the system benefits from the optimized plans generated by the automation. Under extreme circumstances when the automation makes sub-optimal decisions, the human may be able to mitigate this effect by investigating alternative plans with additional input to the decision-making process by imposing additional constraints on the automation's optimization [6]. Of course, operator intervention does not necessarily yield an improvement in performance, and in computationally constrained systems, adding constraints to the optimization problem can exacerbate computational requirements to an extent that hinders planning performance.

Communication Between the Operator and Automation A human operator interacting with an automated system faces many challenges. Operators must choose how to allocate their attention, they must be able to communicate their intent to the automation in diverse situations, and they must be able to maintain global situational awareness while focusing on specific tasks. Previous studies with the OPS-USERS system has shown that the operator utilization rate significantly affects both situational awareness and system performance [6]. High operator utilization rates tended to degrade performance and reduce situational awareness as operators became overwhelmed with the level of interaction required

by the system. Operators performed better, and also reported feeling more confident, when working at a moderate pace.

Translation of operator intent into an automated system is a crucial component in the success of supervisory control systems. This can be accomplished by assigning priorities to objectives in the system or through direct specification of certain behaviors. However, in order to facilitate efficient supervisory control of complex systems, supervisors must be able to express intent at multiple levels of abstraction. One method to address this issue is the *playbook* approach to supervisory control [56]. In the *playbook* system, all entities understand a specific set of actions, or plays, which can be performed. The plays provide a language for communication of intent between the operator and the system, and they may be defined at various levels of abstraction. The operator is free to manage each aspect of the system at varying degrees of detail.

Attention allocation represents another important factor in the performance of decision support systems [57]. As a supervisor monitors the overall state of the system, he or she will choose to periodically address specific tasks. The tasks to which he or she chooses to allocate attention, as well as the amount of attention allocated to each task directly impacts the operator's effectiveness. Automated systems which provide suggestions to operator for where his or her attention is most needed may be able to improve system performance by applying the operator's attention in the most effective way possible.

2.4   Summary

There is a large body of research which addresses the many interesting aspects of UASs, including automated planning, network theory, and human supervisory control. Although each aspect has been considered individually in previous research, there is a lack of understanding of how to optimally assign communication resources in a complex UAS in order to maximize performance. In the next chapter, the work outlined above is used as a basis to design an experiment which will provide insight into the research questions posed in Section 1.3.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Experiment Design

This chapter outlines the development of a discrete event model simulating human interaction with the On-board Planning System for UxVs in Support of Expeditionary Reconnaissance and Surveillance (OPS-USERS) environment, with the goal of investigating the effect of communication degradation on system performance. This model approximates "average" human interaction with the system and facilitates the use of a Monte Carlo method by gathering data through repeated simulations on the OPS-USERS system. The approach would not be feasible in an experiment with human participants due to the number of repetitions required. Using this method, performance on a set of missions in the OPS-USERS environment is sampled under several pre-determined static communication networks and communication delays. A parameter which modulates operator workload and the frequency of replanning is also of interest, especially to the extent that it affects the impact of communication delays.

By subjecting the distributed system to a variety of communication conditions, ranging from perfect communication to communication degraded by both latency and network connectedness, this chapter addresses the research questions set forth in Section 1.3. By measuring performance across scenarios with communications degraded in multiple ways, the impact of overall degradation and the effects of each individual type of failure are investigated. The effect of the OPS-USERS replan-interval is chosen as another independent variable to investigate the interaction of HITL control with system performance across each level of communication degradation.

Network topology and communication delay independent variables were chosen to provide a multi-dimensional parameter space over communication parameters. This choice allows investigation of the relative importance of each variable in determining system performance, a topic which was posed in Research Question #2. Furthermore, both the communication delay and network topology independent variables represent compelling aspects of communications in real-world multi-agent systems. The replan

suggestion interval independent variable was included to maximize the extent to which results from this experiment could be compared to previous experiments with the OPS-USERS testbed. Verifying results between this experiment and the experiment performed in [6] serves as an important basis for validating the human model proposed in this research.

Implementation details of the OPS-USERS system are first provided, followed by a description of the specific experiments performed using the system.

3.1 OPS-USERS Implementation

The OPS-USERS system consists of a heterogeneous fleet of autonomous vehicles and a human operator located in a GCS who oversees the behavior of the system. In contrast to contemporary UASs, the system considered here requires no operator intervention for aircraft navigation, and furthermore, a single human operator controls the entire fleet. In each mission, the team of UAVs cooperates to search an environment for moving targets. When targets are found, they are classified by the human operator as either *friendly*, *hostile*, or *unknown*. Figure 3-1 illustrates the operators' map interface giving an overview of the state of the system. The locations of the UAVs, targets, tasks, etc. are represented with symbols corresponding to Military Standard 2525 [58]. Unknown and hostile targets are tracked by the UAVs, and hostile targets are then tasked for destruction by a Weaponized Unmanned Aerial Vehicle (WUAV). A WUAV is not allowed to engage a hostile target unless it is also being tracked by another UAV, and the human operator must approve the weapon launch.
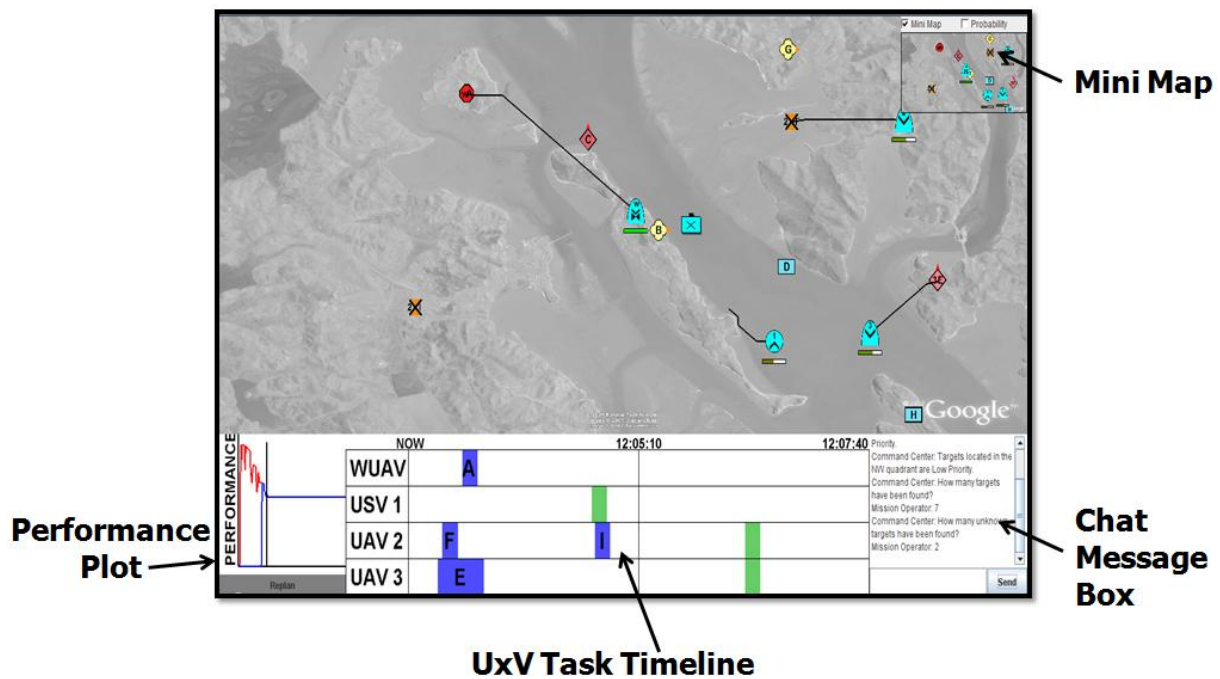
Figure 3-1: Map view in the OPS-USERS interface

Figure 3-2 illustrates the HITL role of the human operator in the OPS-USERS system. The system integrates a decentralized component, the distributed tactical planner, with a centralized planning system inside the GCS. The centralized planner consists of a planning loop involving the CMM and human input through the human interface. The CMM and the distributed tactical planner both maintain a list of tasks which are currently active in the system, including tasks to search specific regions of the environment, tasks to "track" known targets to monitor their location and trajectories, and tasks to destroy known hostile targets. Tasks may arise dynamically as the vehicles interact with the environment or they may be generated through human input.
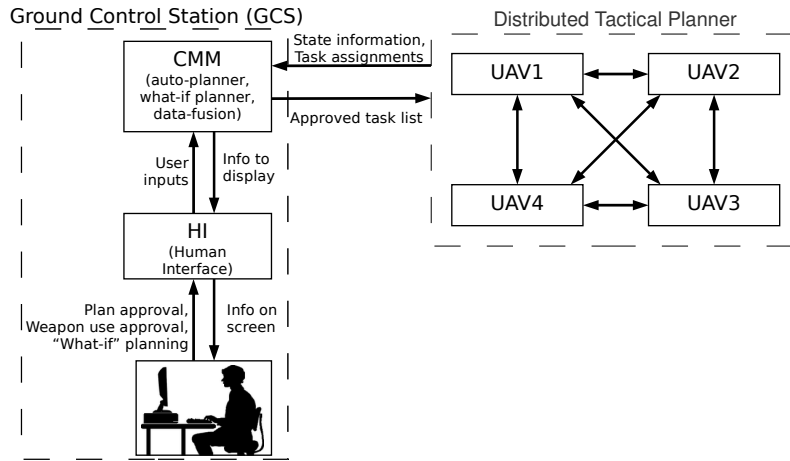
Figure 3-2: Human interaction in the OPS-USERS framework

3.1.1 Automated Planning

Based on observations of the environment made by the vehicles and transmitted to the GCS, the CMM performs a centralized planning algorithm to generate suggested task assignments. It may not be feasible for the vehicles to complete all of the current tasks, but as many tasks as possible will be assigned to vehicles in the suggested plan. Suggested plans are then presented to the human operator in the form of an *approved task list*, or schedule, which is comprised of all tasks which are assigned to a vehicle. The operator receives only the abstracted approved task list rather than the entire plan because the operator's role is to oversee system behavior at a high level and not to micromanage the individual vehicles' behaviors.

The automation prompts the operator to review new suggested schedules at specified time intervals or when an approved task list generated by the CMM differs dramatically from the current list. The operator may also independently choose to review the most recent proposed schedule at any time. The frequency of reviewing and approving proposed plans, or "re-planning", helps to determine the trade-off in the system of the amount of time the vehicles spend attending to individual tasks versus the amount of time the vehicles spend searching unexplored areas of the environment. A short replanning interval

will cause the vehicles to be assigned new tasks immediately following completion of old tasks, whereas a longer interval may provide more time to perform automated searching of unmapped areas in between task assignments. Searching can also be accomplished explicitly if the human operator provides waypoint tasks to the vehicles. However, this method is not preferable to the automated search algorithm unless the operator obtained exogenous intelligence indicating specific areas which should be given preference in searching the environment, or the operator believed that he or she could implement a more efficient search strategy by managing waypoint tasks.

Figure 3-3 shows the replanning mode of the OPS-USERS human interface, the Schedule Comparison Tool (SCT). The approved task list is presented to the user by displaying the set of tasks which the centralized planner could not feasibly assign to a vehicle. The SCT also graphically represents the proportion of high, medium, and low priority tasks which are assigned in the current schedule, the proposed schedule, and the working schedule, as well as the portion of the map which will be searched given the vehicles' trajectories in each schedule. The operator may attempt to coach the automation by dragging one of the unassigned tasks into the "assign" drop area. This causes the centralized planning algorithm to run again after introducing additional constraints to force the task selected by the operator to be assigned if at all possible. This may cause other tasks to become unassigned, and any resulting changes will be reflected in the working schedule. This process is called "what-if" planning and can be repeated until the operator is satisfied with the working schedule, or the operator can choose to return to either the current schedule or the schedule initially proposed by the automated planner. The operator accepts the approved task list in the working schedule when he or she is satisfied and the CMM then sends it to the vehicles, where the distributed tactical planner runs to optimize assignment of the approved tasks to the vehicles. Ultimately, the vehicles sort out the task assignments between themselves, with the centralized system providing guidance on the set of tasks to prioritize.
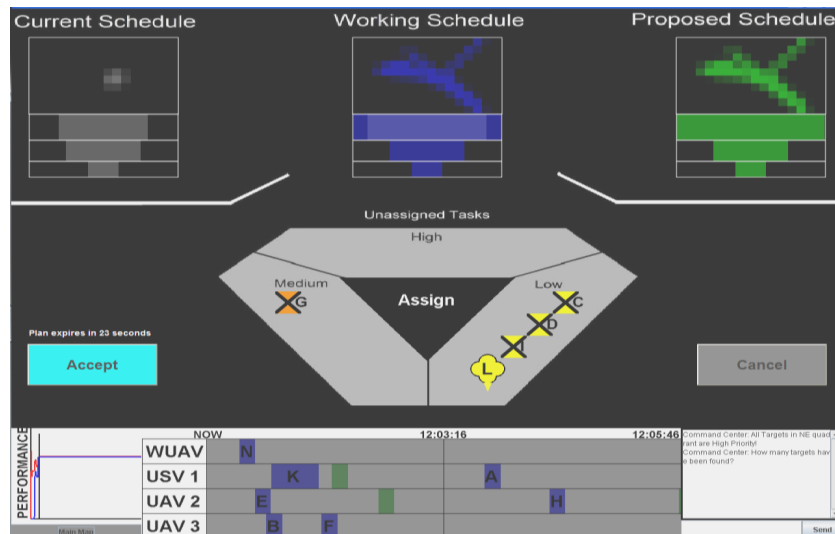
Figure 3-3: Replanning view in the OPS-USERS interface

Although the primary impetus to replan is through prompting by the automation, replanning is also somewhat event-driven in the OPS-USERS system. For example, when a search task is created by the human operator, the operator is aware that he or she needs to replan in order for the new search task to be assigned for a vehicle. The operator need not wait until the automation prompts them to replan. Other events which warrant an unprompted replan are discovery of a new target, or upon changing an unknown target identification to hostile.

3.1.2 Network Communication

In the OPS-USERS system, the individual communication links between two vehicles or between a vehicle and the GCS form a communication network, such that two vehicles may communicate via a third vehicle (or the GCS) if a direct communication link is unavailable. However, any indirect communication increases the latency of these time-sensitive communications and may also strain the bandwidth of a particular communication link if traffic from many vehicles is routed through a single link. To some extent, the increased latency introduced by communication failures may also impact the trade-off between the amount of time the vehicles spend attending to individual tasks, such as tracking a known target, against the amount of time spent performing the background automated searching algorithm. Modulating the interval at which the human operator is prompted to review approved task lists proposed by the distributed tactical planner may somewhat compensate for the effects of degraded communications on system performance.

3.1.3 Data Fusion

OPS-USERS employs Robust Decentralized Task Assignment (RDTA), which includes a linear update filter as its information update step [4]. Each agent updates their beliefs about each variable using the update rule given by:

$$\dot{x}_i = \sum_{j=1}^{N} \sigma_{ij} G_{ji} (x_j - x_i)$$

where agent A_i updates its local belief about variable x given information from each agent A_j (including it's own observations where $i = j$). The indicator variable G_{ji} takes the value 1 when a communication link exists from agent A_j to agent A_i and is 0 otherwise. The relative confidence of each agent about its information is accounted for in the update gain, σ_{ij} [9,27]. Although the OPS-USERS system is capable of performing data fusion, all experiments performed for the purposes of this research were simplified such that all sensors did not introduce noise into measurements. Therefore, the data-fusion problem is reduced to accepting the most recent observation of a variable in the system as the best estimate of that variable. Note that the vehicles can still maintain inconsistent beliefs about the state of the system if they are unable to communicate the most recent observations due to communication failures.

3.1.4 Mission Definition

The basic goals in an OPS-USERS mission are to search for targets and identify each as either friendly, unknown, or hostile. The unknown and hostile targets should be tracked, and the hostile targets should be destroyed. A more complete explanation of the tracking process and the hostile destruction process are provided below.

Target Tracking in the OPS-USERS System The purpose of the target tracking component of the OPS-USERS mission is to maintain accurate information about the position and trajectory of all targets with an unknown or hostile designation. The vehicles need not continuously observe a target in order for the system to consider the target to be in a known location. Based on the target's observed velocity and the sensing capabilities of the vehicles, a *revisit time*, which specifies how long the vehicle may wait before returning to the target's last known location with a high probability of observing the target, can be calculated for each vehicle. If the revisit times of all of the vehicles pass without any of the vehicles observing the target again, the target is considered "lost". The percentage of time targets are in the lost state forms the basis of an important performance metric, the ratio of time targets tracked metric. If a lost target is observed again, the human operator is prompted to re-confirm its designation as either unknown or hostile.

Targets are maintained in a known state by the CMM generating revisit tasks for all targets, which are scheduled to occur before the revisit time of a vehicle expires. Each time a vehicle observes a target, the revisit times are updated (i.e. pushed further into the future) to allow other tasks to be scheduled before the next revisit occurs. In order for the tracking task to be assigned again, it must also be re-approved by the operator in a new schedule after the target has been revisited.

Hostile Destruction in the OPS-USERS System The OPS-USERS system will generate a task for a WUAV to engage hostile targets when certain conditions are met. After a vehicle initially observes a hostile target, the human operator must confirm the hostile designation of the target. Once this confirmation is received, a destroy task is created if the target is not lost, otherwise a destroy task is created when the target is next relocated. In order for the target destroy task to be assigned to the WUAV, the human operator must approve a schedule which contains that task. Once assigned to a WUAV, if the target becomes lost before the destroy task can be completed, then the destroy task is canceled until the target can be located again. Furthermore, when the WUAV arrives at the hostile target, confirmation is requested from the human operator to approve weapons launch. During this time, there is no danger of the target becoming lost, as the WUAV is observing the target.

Target Identification in the OPS-USERS System As new targets are discovered by the vehicles, the human operator is prompted to complete a visual search task before identifying the target as either friendly, unknown or hostile based on an unambiguous image presented in the search task. The symbols

for each target are drawn from the Military Standard 2525 [58]. The human operator is not permitted to incorrectly identify a target upon completion of the visual search task. Therefore, it is primarily the amount of time taken for the operator to initially identify the target which may affect system performance.

Additionally, information is provided to the human operator throughout the mission which specifies the actual designation of unknown targets as either hostile or friendly in the form of messages which appear in the chat box in one corner of the interface. Arrival of new chat messages is accompanied by both a visual and an auditory cue. Experiments with the OPS-USERS system have shown that operators failed to consistently act on information presented through the chat box, preventing some unknown targets from being appropriately identified as either hostile or friendly [6]. These mistakes may drastically impact system performance, as failure to identify a target as hostile prevents the system from destroying it; failure to mark an unknown target as friendly causes the system to continue to expend resources to track the target, which is time that could be spent servicing other tasks or search the environment for new targets.

3.1.5 Previous Research

This work builds upon a previous experiment with the OPS-USERS framework [6]. In the experiment, 30 human participants were presented with three test sessions on the OPS-USERS system with different replan suggestion intervals. A replan suggestion included both an auditory and visual cue for the operator to enter the SCT and approve an updated approved task list. Each scenario had a duration of 10 minutes, and the operators were prompted to replan at intervals of 30-, 45-, and 120-seconds. The system performance was measured for each simulation in order to investigate the impact of the replan interval. The results of the experiment consistently showed that operators performed worse at a 30-second replan interval than at a 45- or 120-second replan interval. Further information about the OPS-USERS system can be found in Reference [59].

3.2 Human Model

Performing enough trials to gather statistically strong evidence is challenging for missions requiring HITL control. Thus, in order to explore the effects of communication degradation on system performance, the human in this research effort is simulated based on information gathered in previous experiments with actual human participants using the interface described previously in Section 3.1. Parameters describing ways in which the human operator interacts with the OPS-USERS system were gathered for the five modes of human interaction:

- **Target Identification:** When the vehicles first discover a new target, a visual search task is presented to the user. The user interacts by scrolling around a map until the target appears in the field of view. When the target is in the field of view, target identification buttons become active and the operator selects the button corresponding to the image shown on the map.

- **Weapons Launch Confirmation:** Targets identified as hostile generate a task for a WUAV to complete. When the WUAV approaches the target, the user is prompted to confirm that the target should be destroyed. This prompt consists of another visual search task as in the target identification task. Completion of this task activates a button to confirm weapons launch.

- **Replanning:** The automation constantly searches for optimal task assignments as the scenario progresses. Periodically, if an improved task assignment is found, the operator is prompted to review and possibly approve the task assignment.

- **Search Task Creation:** In the case where nothing in the simulation environment has changed since the last time the automation prompted the operator to replan, the automation instead prompts the user to create a search task in a specific quadrant of the map.

- **Incoming Chat Messages:** Periodically throughout each mission, operators receive chat messages with instructions to redesignate certain targets with unknown designations to either a friendly or hostile designation.

Each cognitive task can be sequential steps. The human operator first requires some time to process each cognitive task prior to providing an input to the system. Both the time delay and input of the operator have been analyzed in the context of the mission and in accordance with previous observations of human interaction in this system so that a model of human behavior can be built to interact with the system in the same way an operator would [6]. The human model serves as a replacement for the human operator in repeated simulations of mission scenarios. More detail about the specific methods for emulating each type of human interaction with the OPS-USERS system is provided below.

3.2.1 Target Identification

The target identification task is modeled by the time taken to complete the visual search task and to identify the target, drawing from data gathered in a previous experiment on operator utilization and replanning intervals indicating when the target ID window opened and closed. Each data point was defined as the duration that the target ID window remained open each time it appeared. A probability distribution was fit to the experimental data, resulting in a log-normal distribution with parameters given in Table 3.1. Figure 3-4 shows a histogram of the probability density function described by the data, along with the best fit log-normal distribution.

In order to simulate the human operator processing the visual search task to identify targets, a time delay is drawn from this best fit probability distribution, and the program simulates this input delay for that time before correctly identifying the target. This impacts the performance of the system as the delay before the target is identified can be significant. In some of the observed cases, this delay was greater than 20 seconds, which is significant compared to the total mission duration of 10 minutes. Further action

cannot be taken on the target until it is identified. For example, if the target is hostile, no destroy task could be generated or assigned until the target ID process is completed. This delay directly impacts the number of hostile targets destroyed during a mission. It is important to note that the emphasis during this task is this delay and not the correctness of the target identification. In the HITL experiment, there were no incorrect identifications as the task was presented with an unambiguous correct answer, and the user was prompted to select another identification until the correct identification was chosen.

Table 3.1: Log-normal distribution parameters for target identification times

Log-normal parameter μ	1.67
Log-normal parameter σ	0.71
Mean $= e^{\mu+(\sigma^2/2)}$ (seconds)	6.79
Std. Dev. $= \sqrt{(e^{\sigma^2}-1)e^{2\mu+\sigma^2}}$ (seconds)	5.47
K.S. Statistic	0.031
p-value	0.356
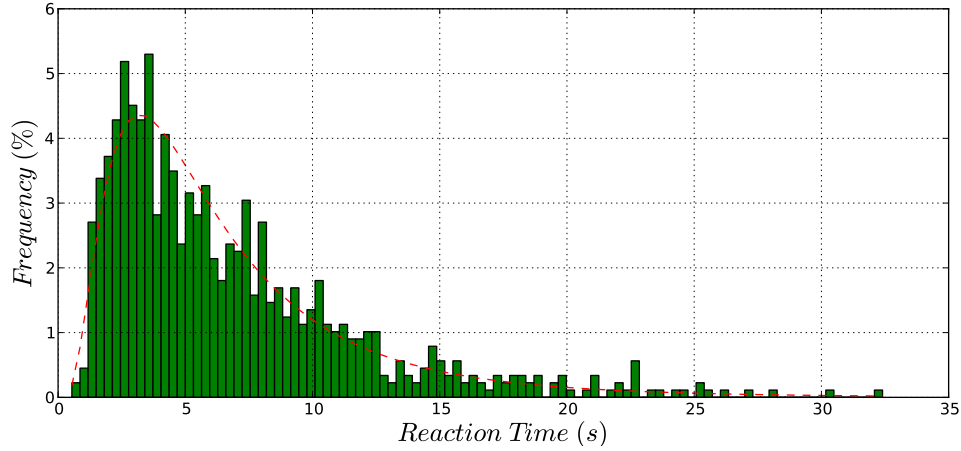
Figure 3-4: Target identification times fit to log-normal distribution

3.2.2 Weapon Launch Confirmation

The weapon launch confirmation dialog is similar in nature to the target identification dialog. The user is presented with a simpler version of the target identification visual search task which requires less time to process. When the target appears in the field of view, a button to confirm weapon launch is activated. The procedure for modeling this task is analogous to the model for the target identification task. A data point is defined as the length of time that the weapon launch confirmation dialog remained visible each time it appeared. Again a log-normal distribution was found to fit the data the best, with parameters given in Table 3.2. Figure 3-5 shows a histogram of the data and the best fit distribution. In this scenario there is never a situation which requires the operator to perform any action other than confirming the weapon launch, so the response to the confirmation question does not need to be modeled,

53

and will always be assumed to confirm the weapon launch.

Table 3.2: Log-normal distribution parameters for weapon use confirmation times

Log-normal parameter μ	1.15
Log-normal parameter σ	0.68
Mean $= e^{\mu + (\sigma^2/2)}$ (seconds)	3.97
Std. Dev. $= \sqrt{(e^{\sigma^2} - 1)e^{2\mu + \sigma^2}}$ (seconds)	3.03
K.S. Statistic	0.082
p-value	0.136
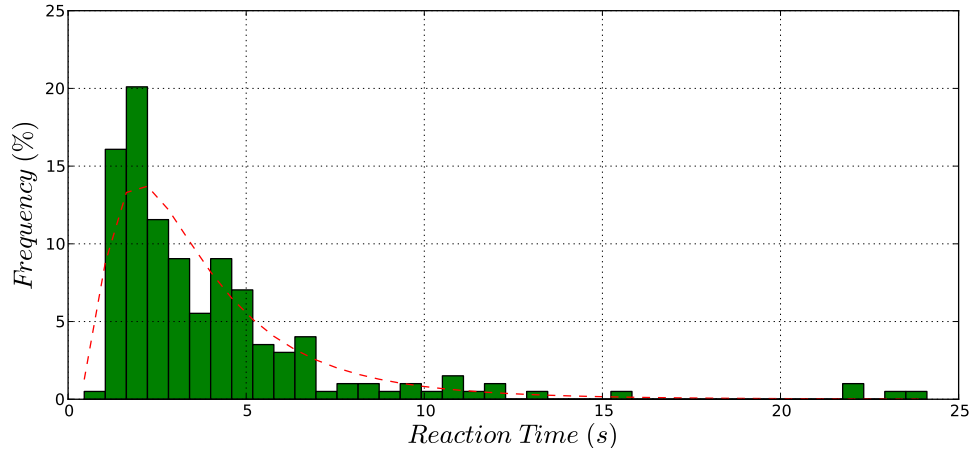
Figure 3-5: Target destroy confirmation times fit to log-normal distribution

3.2.3 Replanning

The replanning task presents the user with a view of the SCT shown in Figure 3-3. Here the user reviews the set of tasks which remain unassigned in the current schedule proposed by the automation. The operator can choose to explore other possible task assignments by adding constraints to the optimization requiring specific tasks to be assigned. Then the operator chooses to accept either the original task assignment suggested by the automation, or a new task assignment found after adding constraints to the optimization. This task is modeled using two independent components: the time taken to complete the replanning task, and the changes made to the plan.

Table 3.3: Log-normal distribution parameters for replanning times

Log-normal parameter μ	1.40
Log-normal parameter σ	0.72
Mean $= e^{\mu + (\sigma^2/2)}$ (seconds)	5.24
Std. Dev. $= \sqrt{(e^{\sigma^2} - 1)e^{2\mu + \sigma^2}}$ (seconds)	4.29
K.S. Statistic	0.037
p-value	0.044

Figure 3-6: Replan times fit to log-normal distribution

The model of the time taken to complete the replanning task was performed in much the same way as the previous two tasks. Data was drawn from the previous human experiment measuring the length of time that the replanning window was open, and a log-normal distribution was found to best fit the data. However, replan data was only drawn from participants who were classified as *conformist* operators [60], meaning that they collaborated effectively with the automation by consistently replanning when prompted. The other classifications were mixed-conformist and non-conformist. Mixed-conformist operators replanned when prompted to by the automation, but also engaged in significant unprompted interaction with the automation. Non-conformist operators replanned on their own schedule, largely ignoring prompts from the automation. This choice was made in order for the human model to best reflect the behavior of an expert conformist operator. The parameters of this distribution are shown in Table 3.3. A histogram of the probability density function for the data set along with the best fit model are shown in Figure 3-6.

Another component specific to the replanning task is the operator's modifications to the schedule proposed by the automation. Conformist human operators modified the automation's proposed schedule approximately 32% of the time, with a standard deviation of 4.5% [60]. Modeling whether or not the operator modifies the plan will be accomplished via a simple Bernoulli random variable with parameter $p = 0.32$. However, modeling how the plan is modified is not so straightforward.

When the realization of the Bernoulli random variable indicates that the plan should be modified, the model attempts to generate a new proposed schedule from the automation. This is accomplished by requesting that the planner assign one of the unassigned tasks, with a preference for the unassigned task closest to any of the vehicles. This process is repeated with tasks further from vehicles until a modified plan is found. If no attempted changes result in a feasible task assignment, or if there are no unassigned tasks, then the task assignment remains unchanged. This method is motivated by a heuristic

which the human operator typically employs if he or she perceives the proposed plan to be suboptimal. Assigning unassigned tasks to the closest vehicle has been exhibited in other single operator multiple-UV control simulations [57]. Unassigned tasks located near a vehicle capable of completing the task represent a reasonable candidate for modifying the plans proposed by the automation. It is important to note that the algorithm design does not place the operator in a situation where the suggested plans are intentionally suboptimal or that would otherwise necessitate human intervention in the replanning process. The operators intervene based only on their perception of the utility of the possible plans given the mission objectives [59].

3.2.4 Incoming Chat Messages

In the original human experiment, operators received chat messages via a chatbox in one corner of the interface (see the lower right corner of Figures 3-1 and 3-3). Approximately half of the messages prompted the operator to respond to a question regarding the state of the system (for example, to indicate the number of hostile targets which had been destroyed so far). The answers to these questions were recorded and the operators' response times and response accuracies were taken as a measure of their situational awareness. This aspect of chat message interaction was not modeled in the human model developed for this research, as it does not directly impact performance of the system, but rather measures the performance of the operator.

The rest of the chat messages prompted the operator to take actions to redesignate targets with unknown designations to either friendly or hostile status. These message represent intelligence gathered exogenously to the UAS, which is to be incorporated by the operator. Specifically, the message format would indicate that all unknown targets in a specific quadrant of the environment are to be designated as either friendly or hostile. The interesting aspect of this task is whether the operators took the appropriate action, as well any delay incurred before the action was taken.

Based on the results of the human experiment, it was discovered that performance in the redesignation task was significantly degraded at the 30-second replan suggestion interval, where operators instituted correct changes in target designation less accurately than at a 45- or 120-second replan suggestion interval [6]. To model the chatbox aspect of the OPS-USERS system in this model, the probability of correctly responding to chatbox messages at the 30-second interval was set reduced relative to greater replan suggestion intervals. This decision was justified as target redesignation performance at the 45-second replan interval was not significantly different from the 120-second interval, although both were significantly higher than the 30-second interval. In the model employed in this research, the simulated operator correctly responded to target redesignation requests 70% of the time when the replan suggestion interval was 45-seconds or more, and 40% of the time for the 30-second replan interval. These values were based on the accuracy with which participants responded to chat prompts in the previous experiment.

Figure 3-7: Network topologies for the fully connected network (left), groups network (center), and round-robin network (right)

3.3 Independent Variables

Mission parameters and limitations on communication between the agents in the system are specified for each experimental condition. Specifically, there are 3 independent variables: the set of communication links available between the vehicles and the GCS (also called the *network topology*), the level of communication latency, and the replan suggestion interval.

Communication Limitations The network topology dictates the presence or absence of each individual directed communication link in the system, and latency dictates a global time delay associated with communication over a link between any two agents. Three specifications of the network topology were tested: 1) a *fully connected* network where all possible communication links are present; 2) a *groups* network which splits the vehicles into two groups, and 3) a *round-robin* network which represents the minimal number of communication links possible while maintaining strong connectivity (i.e. two-way communication is still possible between all pairs of vehicles). In the groups network, communication links allow direct communication with the GCS for all vehicles, and direct communication between vehicles within each group; vehicles in different groups must communicate via the GCS. The three network topologies are shown in Figure 3-7. These networks were chosen to represent both ends of the spectrum of strongly connected network topologies – both the best case (fully connected) and the worst case (round-robin). The network characteristics are summarized in Table 3.4 to highlight the differences between the networks. The metrics used to characterize networks include density, clustering coefficient, and diameter, summarized below [61].

Table 3.4: Network characteristics

Feature	Fully Connected	Group	Round Robin
Density	1.0	0.6	0.25
Clustering Coefficient	1.0	0.2	0.0
Diameter	0.25	0.5	1.0

- **Density** refers to the portion of possible directed network links which are present in the network. Density is calculated as:

$$D = \frac{\text{Number of active connections}}{\text{Total possible connections}}$$

In the networks of 5 agents considered in this research, there are $2 \cdot \binom{5}{2} = 20$ possible directed connections. The round-robin network contains only 5 of these connections, giving it a density of one quarter of the possible connections, or $\frac{5}{20} = 0.25$.

- **Clustering Coefficient** Provides a measure of the tendency for the neighbors of a node to be neighbors themselves. Clustering coefficient is calculated as:

$$C = \frac{\text{Number of connected triangles in graph}}{\text{Total sets of three nodes in graph}}$$

- **Diameter** is length of the longest shortest path between two nodes. Two agents communicating indirectly always choose the shortest path through other vehicles to route their communications. The longest such path determines the diameter, which is presented as a ratio of the length of this path over the number of agents in the network.

Each of these topologies were further subjected to a set of induced latencies (where the latency is static and equivalent for each link the in the network). The latencies tested were chosen relative to the duration of the mission, as well as the time-frame of individual tasks the vehicles complete. The chosen latency values were (in seconds): $\{0, \quad 2, \quad 4, \quad 6, \quad 8, \quad 10\}$. Note that two vehicles which cannot communicate directly, but must communicate via other vehicles, will incur multiple latencies before messages are received. For example, in the round-robin network in Figure 3-7 with a latency of 2-seconds, a messages from UAV1 to UAV4 will incur a 6-second total latency.

The set of latencies was chosen to span the spectrum of communication availability with unlimited communication available at one-end of the spectrum to extremely limited communication at the other end of the spectrum. At the 0-second latency, which was also included to correspond with the conditions in the initial human experiment, the amount of communication is practically unlimited. A maximum latency of 10-seconds was chosen based on the total round-trip communication time in the round-robin network topology. All simulations in this research were performed with 5 agents, such that round trip communication between any two will take 50 seconds when the latency is 10 seconds. This represents a significant portion of the mission duration (just under 10%), and was considered extreme communication degradation within the context of the simulated mission.

Mission Parameters The mission definition was held constant throughout all simulations with the exception of initial target locations and the replan suggestion interval. The mission definition encompasses the mission environment, the number of vehicles, the capabilities of each vehicle, mission duration, and

the number and designation of targets. As with previous human experiments, the same three sets of initial target locations were used, and an equal number of simulations were performed with each set of initial target locations for each combination of independent variables. This strategy was deemed sufficient to randomize target locations since each target moves along a predefined track, but has a random motion as each mission progresses.

The replan suggestion interval controls the frequency with which the automation prompts the human operator to review and possibly approve the automation's current best proposed plan. This was the independent variable in previous work with the OPS-USERS system [6]. As noted in Section 3.1, this variable also modulates the trade-off between the two primary goals in the OPS-USERS system: 1) processing of tasks, and 2) performing the background search task for new targets. This variable is interesting in the context of its interaction with communication availability, specifically related to Research Question #3, stated in Section 1.3.2. The merits of choosing a particular trade-off between searching for new targets and completing tasks may change dramatically as communication in the system is degraded. The chosen replan intervals were (in seconds): {30, 45, 60, 75, 90, 105, 120}. This range of replan intervals spans the range of replan intervals tested in the original human experiments (which were 30-, 45-, and 120-seconds), but also includes extra resolution at intermediate replan intervals.

3.4   Dependent Variables

The dependent variables measure system performance as well as characteristics of the communication network. A consistent set of metrics relevant to system performance over a range of both missions and communication regimes is necessary to investigate the relationship between the communications available to the system and the system's performance. Data collected from the OPS-USERS system during each simulation consists of the following mission performance metrics:

- Percent targets found
- Percent environment searched
- Percent hostile targets destroyed
- Ratio of time targets tracked

Values of the metrics were captured at the end of each mission simulation, which indicates an aggregate measure of performance over the duration of the mission. However, the time-evolution of each variable throughout the duration of each variable is also relevant to the interaction of the independent variables and system performance. Therefore, the values of each variable were recorded at 10 second intervals through each simulation. Exact times were also recorded for target discovery events or target destruction events.

3.5   Gathering System Performance Data

The independent variable choices are summarized here:

- Communication delay values (seconds): {0.0, 2.0, 4.0, 6.0, 8.0, 10.0}.
- Replan suggestion interval values (seconds): {30, 45, 60, 75, 90, 105, 120}.
- Network topologies: {Fully Connected Network,   Groups Network,   Round-Robin Network}

As described above, 6 levels of communication latency were tested, 3 network topologies were tested, and 7 levels of the replan suggestion interval were tested. For each of the 126 combinations of the three variables, the simulation was run 30 times, resulting in 3,780 total simulations. Each simulation had a duration of 10 minutes, and simulations were run in real time. This represents over 26 days of processing time, performed simultaneously across 7 computers.

3.6   Summary

This chapter identified the challenge of collecting large sets of system performance data in human supervisory control settings. As collecting a data set of the size required for this research with human operators would be prohibitively complex from a logistical standpoint, a human model was proposed to simulate the behavior of the human operator during data collection. Using results from previous work with the same OPS-USERS simulation testbed as a basis, the human model emulates both the time an operator requires to process each task in the system as well as input from the operator. The next chapter presents results from repeatedly sampling the system performance using the human model described above.

Chapter 4

Results

This work both builds upon previous research in developing an understanding of the performance of UASs under human supervisory control and extends this understanding into the communication availability domain. Previous experimental results demonstrated that the replan suggestion interval significantly affected the performance of the OPS-USERS UAS [6]. The experimental results presented below provide further insight into the relationship between system performance, communication availability, and human supervisory control in the context of the research questions posed in Section 1.3.

In order to ensure that the substitution of the human model described in Section 3.2 for actual participants does not significantly affect system performance, the model is first validated by comparing experimental results from this research against the results from the research on which the model was based. This ensures that the human model affects the results in simulation in the same way the human operator affected results in experimental data. Specifically, the model determines how frequently to replan, when and how to alter plans suggested by the automation, and how promptly to address other tasks delegated to the human (identifying targets and approving weapons launch). Next, simulation results are analyzed to investigate the relationship between the network topology, communication delay, and replan suggestion interval on the four system performance dependent variables, percent area covered, percent targets found, ratio of time targets tracked, and percent hostiles destroyed.

4.1 Human Model Validation

An important aspect of this research was the development of a simulated operator model to facilitate generation of a large data set covering a variety of degraded communication conditions. The parameters of the model are based on the results of previous HITL experimentation measuring the impact of replan suggestion interval on human performance [6]. In order to ensure the validity of the data set generated

from the model, the results are compared against results from the original experiment. The human model is demonstrated to provide results consistent with the original experiment.

Communication Conditions Overview The original human-subject experiment did not introduce communication failures into the simulation. This condition directly corresponds to the fully-connected network with a 0-second communication delay (i.e. without lag). Furthermore, the original experiment tested only 3 replanning intervals, 30-, 45-, and 120-seconds, selected to measure the performance of the system when the operator experienced three different levels of utilization. Therefore, simulation data gathered from the fully-connected network topology without lag and the 30-, 45-, and 120-second replan interval will be compared with data from the original experiment. In both the original human experiments and this research, 30 simulations were conducted for each condition. For statistical analysis, the significance level, α, was set to 0.05. To test the null hypothesis that the data from human experiments and simulations come from the same distribution, independent samples t-tests were performed.

4.1.1 Percent Area Covered

Figure 4-1 shows the area covered metric from both the previous human experiments and the simulations. For each value of the replan interval, an independent samples t-test was performed to determine whether the mean of percent area covered metric values was significantly different when the human operator was replaced with the model developed for this research. Levene's test for equality of variances was also performed. The results of the tests are summarized in Table 4.1. At all three replan intervals, there was no significant difference in the mean or variance between the human experiment and the simulation data. Overall, percent area covered performance with a simulated human operator was consistent with the original OPS-USERS experimental results.

Figure 4-1: Comparison of percent area covered metric in original human experimental results with simulation results

Table 4.1: Independent samples t-test and Levene's test results for the percent targets found metric comparing performance with a real human operator vs. the simulated operator model

$\alpha = 0.05$ $*$ Significant result $\dagger$ Marginally significant result

Independent Samples t-Test					Levene's Test for Equality of Variance	
Replan Interval	t value	df_N	df_D	p	F	p
30s	0.359	1	58	0.721	0.025	0.876
45s	1.758	1	59	0.084	0.398	0.531
120s	0.116	1	58	0.908	0.272	0.604

4.1.2   Percent Targets Found

Figure 4-2 shows the percent targets found metric from both the previous human experiments and the simulations. For each value of the replan interval, an independent samples t-test was performed to determine whether the mean of percent targets found metric values was significantly different when the human operator was replaced with the model developed for this research. Levene's test for equality of variance was also performed. The results of the tests are summarized in Table 4.2. At all three replan intervals, there was no significant difference in the mean or variance between the human experiment and the simulation data. The percent targets found performance with a simulated human operator was consistent with the original OPS-USERS experimental results.

Figure 4-2: Comparison of percent targets found metric in original human experimental results with simulation results

4.1.3   Percent Hostiles Destroyed

Consistent with the original OPS-USERS human operator experiment, the percent hostiles destroyed metric values will be treated as count data. Figure 4-3 shows the area covered metric from both the previous human experiments and the simulations. As the hostiles destroyed metric was considered to represent count-data, non-parametric tests were preferred for this variable. For each value of the replan

Independent samples t-test and Levene's test results for the percent targets found metric comparing performance with a real human operator vs. the simulated operator model

$\alpha = 0.05$ $*$ Significant result $\dagger$ Marginally significant result

Independent Samples t-Test					Levene's Test for Equality of Variance	
Replan Interval	t value	df_N	df_D	p	F	p
30s	-0.875	1	58	0.385	0.786	0.379
45s	0.963	1	59	0.339	0.696	0.407
120s	1.293	1	59	0.201	1.332	0.253

interval, a non parametric Mann-Whitney U test was performed to determine whether the distribution of percent hostiles destroyed metric values was significantly different when the human operator was replaced with the model developed for this research. The results of the test are summarized in Table 4.3. There was not a significant difference between the human experiment and the simulation data at any of the three replan intervals. Overall, percent hostiles destroyed performance with a simulated human operator was consistent with the original OPS-USERS experimental results.


Figure 4-3: 95% confidence interval for mean comparison of the percent hostiles destroyed metric in original human experimental results and simulation results

Table 4.3: Mann-Whitney U test results for the percent hostiles destroyed metric comparing performance with a real human operator vs. the simulated operator model

$\alpha = 0.05$ $*$ Significant result $\dagger$ Marginally significant result

Replan Interval	Mann-Whitney U	Z	p
30s	416.0	-0.776	0.438
45s	377.0	-1.359	0.174
120s	427.0	-0.579	0.563

4.1.4 Summary

The human model produces results which corresponded relatively well to the original human experiments on which it was based. It has been demonstrated that the model produces values for the percent area covered, percent targets found, and percent hostiles destroyed metric values which do not significantly differ from the distributions recorded in the original experiment. Therefore, it is reasonable to assume that the data gathered from the simulations performed using this human model can provide insight for the performance of the OPS-USERS system with an actual human operator under degraded communication conditions. One important assumption made in this conclusion is that the model of human behavior would be invariant under degraded communication conditions. If the human operators adjusted their behavior in the face of large communication delays or limited network connectivity, the model developed in this research will likely fail to capture such effects on system performance. It is recommended that future work include experiments with participants interacting with the system under significantly degraded communication conditions in order to better understand the validity of this assumption.

4.2 Extension of Human Model to Degraded Communication Situations

Data gathered through simulation using the human model described in Section 3.2 is analyzed to determine whether degraded communication conditions produced statistically significant differences in system performance. For each performance metric dependent variable from Section 3.4, the final value was recorded at the end of every simulation. The network topologies tested in this experiment are reproduced below in Figure 4-4.

Figure 4-4: Network topologies for the fully connected network (left), groups network (center), and round-robin network (right)

This experiment was designed with 3 independent variables, outlined in Section 3.5 as network topology, communication delay, and replan suggestion interval, with 3, 6, and 7 factor levels respectively. The experiment measured 4 dependent variables, outlined in Section 3.4. Analysis of the data was therefore performed using a 3×6×7 Multivariate Analysis of Variance (MANOVA) test, detailed below. A MANOVA

test was chosen to first investigate whether the independent variables had an overall effect on the set of dependent variables, and whether any significant interactions occurred between independent variables. The significance level for each test was set to $\alpha = 0.01$ to account for family-wise error. Family-wise error was further controlled by selecting the two extreme levels and one intermediate level for pairwise tests of the communication delay and replan suggestion interval independent variables to reduce the number of tests performed on the dataset. This results in 3 factor levels for each of the 3 independent variables in pairwise testing.

4.2.1 MANOVA Results

A Multivariate Analysis of Variance (MANOVA) test was performed to determine whether any of the independent variables had a significant effect on any of the performance metrics in the system. The full table of results is provided in Appendix A. A summary of the main results are presented below in Table 4.4.

Table 4.4: Summary of MANOVA results

$\alpha = 0.01$ ∗ Significant result   † Marginally significant result

Source	Wilk's λ F Statistic	df	Sig.
Network Topology	65.984	8	0.000*
Communication Delay	54.836	20	0.000*
Replan Suggestion Interval	11.697	24	0.000*

Each independent variable is shown to have a significant effect on at least one of the dependent variables. In addition, the interaction term between the network topology and communication delay independent variables in the MANOVA results was also significant. A summary of the interaction results are presented in Table 4.5 below. The significant interaction will be investigated in the next Section.

Table 4.5: Summary of MANOVA interaction terms

$\alpha = 0.01$ ∗ Significant result   † Marginally significant result

Source	Wilk's λ F Statistic	df	Sig.
Network Topology × Communication Delay	7.456	40	0.000*
Network Topology × Replan Suggestion Interval	0.954	48	0.564
Communication Delay × Replan Suggestion Interval	1.026	120	0.404
Network Topology × Communication Delay × Replan Suggestion Interval	1.062	240	0.244

4.2.2 Follow-up Univariate ANOVA Tests

All three independent variables were shown to have a significant main effect in the MANOVA test, indicating they may have a significant impact on at least one dependent variable. This result justifies performing univariate Analysis of Variance (ANOVA) tests on each dependent variable over all three independent variables simultaneously to further investigate the effect. The results of the four univariate ANOVA tests are summarized in Table 4.6 below.

Table 4.6: Summary of univariate ANOVA results

$\alpha = 0.01$ * Significant result   † Marginally significant result

Source	Dependent Variable	df	F	Sig.
Network Topology	% Area Covered	2	0.715	0.489
	% Targets Found	2	1.736	0.176
	% Time Targets Tracked	2	107.686	0.000*
	% Hostiles Destroyed	2	190.782	0.000*
Communication Delay	% Area Covered	5	11.649	0.000*
	% Targets Found	5	7.617	0.000*
	% Time Targets Tracked	5	77.140	0.000*
	% Hostiles Destroyed	5	163.518	0.000*
Replan Suggestion Interval	% Area Covered	6	21.762	0.000*
	% Targets Found	6	9.004	0.000*
	% Time Targets Tracked	6	14.696	0.000*
	% Hostiles Destroyed	6	2.656	0.014†

Interaction Terms As mentioned above, the full MANOVA model also demonstrated a significant interaction between the network topology and communication delay independent variables. Follow-up univariate tests confirmed this interaction was significant in two of the four dependent variables outlined in Table 4.7 below.

Table 4.7: Significant interaction terms in univariate ANOVA results

$\alpha = 0.01$ * Significant result   † Marginally significant result

Source	Dependent Variable	df	F	Sig.
Network Topology × Communication Delay	% Area Covered	10	1.942	0.036
	% Targets Found	10	1.942	0.157
	% Time Targets Tracked	10	15.878	0.000*
	% Hostiles Destroyed	10	13.623	0.000*

Summary These results demonstrate that all three independent variables significantly impacted some or all measures of the UAS's performance. Individual analysis of each dependent variable follows in the subsequent subsections, where the results of univariate ANOVAs and post-hoc Tukey pairwise comparisons are presented for each significant effect. Investigation of the significant interactions between independent variables is also included for each dependent variable.

4.2.3   Percent Area Covered

Response plane visual representations of the results are shown in Figure 4-5, where the percent area covered within each network topology are presented across the communication delay and replan suggestion interval axes. The lowest area coverage performance appears to occur at small values of the replan suggestion interval. Consistent with these observations, the percent area covered dependent variable showed significance in the ANOVA over the replan suggestion interval independent variable, with $p < 0.001$. The communication delay was also demonstrated to significantly impact area coverage performance, with $p < 0.001$, although the precise relationship is not immediately apparent in the response plane visualization.

Examining the aggregate area coverage results across the three network topology plots show a consistent range of values between 55% and 65% in each network topology. Despite the presence of several of the lowest performance cases in the round-robin network for small replan intervals and moderate communication delays, the univariate ANOVA did not show the network topology variable to be significant and therefore pairwise comparisons were not performed for this independent variable.

(a) Fully Connected Network   (b) Groups Network   (c) Round-Robin Network

Figure 4-5: Percent area covered in fully connected betwork (left), groups network (center), and round-robin network (right)

Based on the significant results in the ANOVA test, post-hoc Tukey pairwise comparisons were performed for the communication delay and replan suggestion interval independent variables. Results are outlined in the subsections below, and full post-hoc Tukey pairwise comparison results tables are included in Appendix A.

Communication Delay Figure 4-6 shows a 95% confidence interval for the mean percent area covered across levels of communication delay. The percent area covered at a 0-second delay was significantly higher than at either a 4-second or a 10-second delay, with $p < 0.001$. There was not a significant difference between the 4- and 10-second delay, with $p = 0.601$.

Replan Suggestion Interval Figure 4-7 shows a 95% confidence interval for the mean percent area covered across levels of the replan suggestion interval. The percent area covered was shown to be sig-

Figure 4-6: 95% confidence interval for the mean of percent area covered over levels of communication delay

nificantly lower at the 30-second interval than at the 75-second or 120-second intervals with $p < 0.001$. The difference in percent area covered between the 75- and 120-second intervals was not significant with $p = 0.993$.

Figure 4-7: 95% confidence interval for the mean of percent area covered over levels of the replan suggestion interval

4.2.4 Percent Targets Found

Response plane visual representations of the percent targets found metric are shown in Figure 4-8, where the results within each network topology are presented across the communication delay and replan suggestion interval axes. The targets found results appear to have a higher variance relative to the other variables, which makes identification of trends in the data across any of the independent variables difficult to perceive visually.

Results are outlined in the subsections below, and full post-hoc Tukey pairwise comparison results tables are included in Appendix A. The results from the univariate ANOVA performed (Table 4.6), show that the network topology independent variable did not significantly affect the percent targets found

69

(a) Fully Connected Network   (b) Groups Network   (c) Round-Robin Network

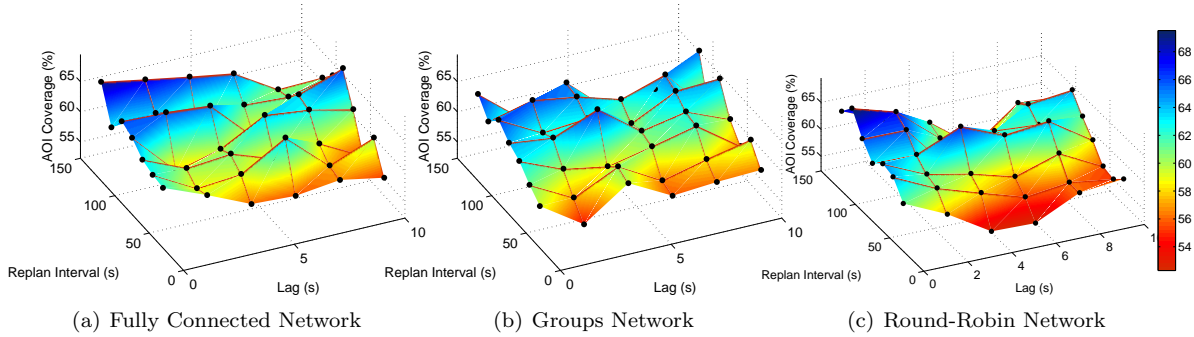Figure 4-8: Percent targets found in fully connected network (left), groups network (center), and round-robin network (right)

metric, with $p = 0.176$. The communication delay and replan suggestion interval were both shown to have a significant effect, with $p < 0.001$. Based on the significant results in the ANOVA test, post-hoc Tukey pairwise comparisons were performed for the communication delay and replan suggestion interval independent variables, which are outlined below.

Communication Delay Figure 4-9 shows a 95% confidence interval for the mean percent targets found across levels of communication delay. The percent targets found at a 0-second delay was significantly higher than at a 4- or 10-second delay, with $p < 0.001$. There was not a significant difference between the performance with a 4-second delay and with a 10-second delay, with $p = 0.997$.

Figure 4-9: 95% confidence interval for the mean of percent targets found over levels of communication delay

Replan Suggestion Interval Figure 4-10 shows a 95% confidence interval for the mean percent targets found across levels of the replan suggestion interval. The percent targets found was shown to be significantly lower at the 30-second interval than at a 75-second interval or 120-second interval, with $p = 0.005$ and $p < 0.001$ respectively. There was not a significant difference between the 75- and 120-second intervals, with $p = 0.088$.
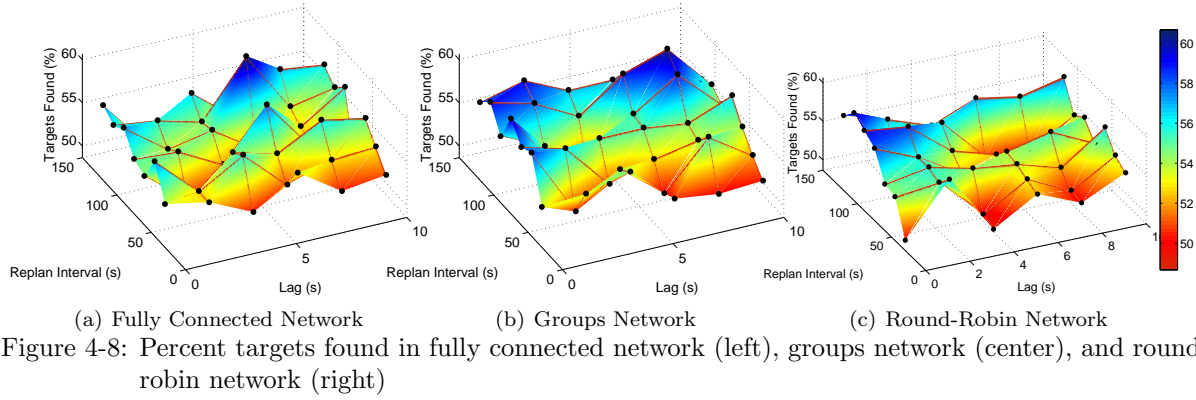
Figure 4-10: 95% confidence interval for the mean of percent targets found over levels of the replan suggestion interval

4.2.5 Ratio of Time Targets Tracked

Response plane visual representations of the results for ratio of time targets tracked are shown in Figure 4-5, where the ratio of time targets tracked results within each network topology are presented across the communication delay and replan suggestion interval axes. The lowest target tracking performance appears to occur in the round-robin network topology. There also appears to be a negative correlation between the ratio of time targets tracked and increases in either communication delay or the replan suggestion interval. The ratio of time targets tracked dependent variable showed significance in the univariate ANOVA over all three independent variables at $p < 0.001$ (see Table 4.6), which is consistent with these observations.

(a) Fully Connected Network (b) Groups Network (c) Round-Robin Network
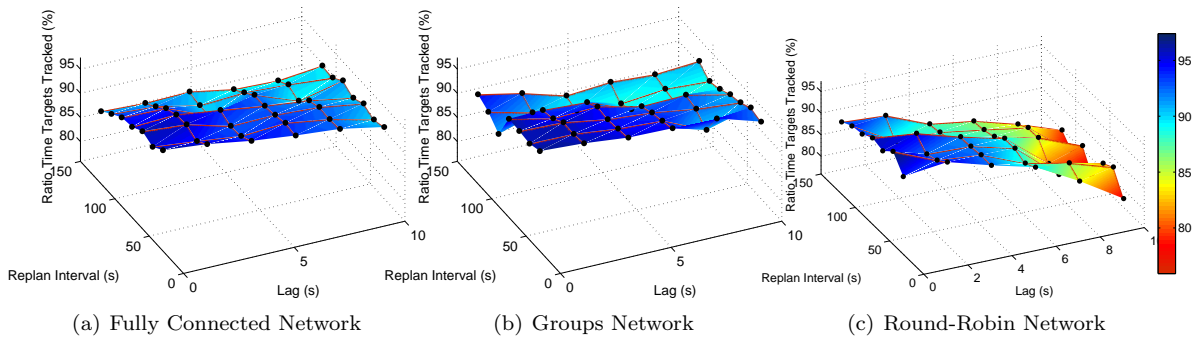
Figure 4-11: Ratio of time targets tracked in fully connected network (left), groups network (center), and round-robin network (right)

Based on the significant results in the ANOVA test, post-hoc Tukey pairwise comparisons were performed for the network topology, communication delay, and replan suggestion interval independent variables. Results are outlined in the subsections below, and full post-hoc Tukey pairwise comparison results tables are included in Appendix A.

71

Network Topology Figure 4-12 shows a 95% confidence interval for the mean ratio of time targets tracked across the three network topologies. The performance in the fully connected and groups networks were both significantly higher than in the round-robin network, $p < 0.001$. The fully connected network was not significantly different from the groups network, with $p = 0.869$.
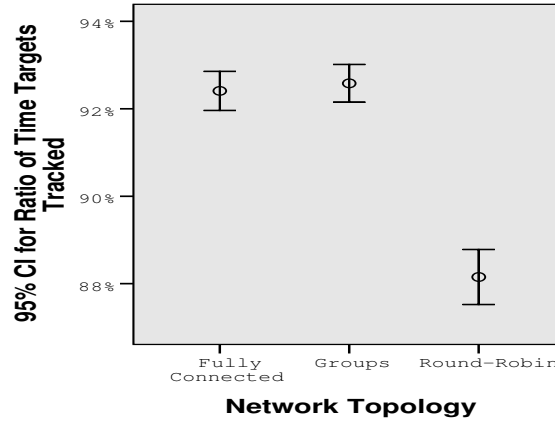
Figure 4-12: 95% confidence interval for the mean of ratio of time targets tracked over network topologies

Communication Delay Figure 4-13 shows a 95% confidence interval for the mean ratio of time targets tracked across levels of communication delay. The results show a statistically significant decrease in the ratio of time targets tracked metric at each increase in the communication delay. Performance with a 4-second or 10-second delay was significantly lower than with a 0-second delay at $p < 0.001$, and performance with at a 4-second delay was significantly lower than at a 10-second delay at $p < 0.001$ as well.
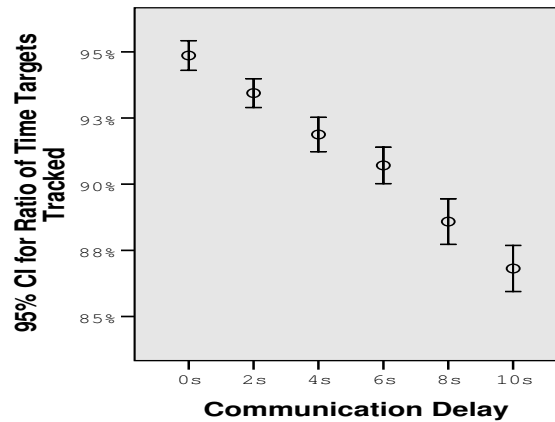

Figure 4-13: 95% confidence interval for the mean of ratio of time targets tracked over levels of communication delay

Replan Suggestion Interval Figure 4-14 shows a 95% confidence interval for the mean ratio of time targets tracked across levels of the replan suggestion interval. Mean performance was highest at a 30-second replan interval, but was not significantly higher than at 75-seconds with $p = 0.392$. However,

the 30-second interval did produce significantly higher performance than the 120-second interval, with $p < 0.001$. Performance at the 75-second interval was also significantly higher than at the 120-second interval with $p = 0.002$. These results indicate a clear trend in decreasing performance in the target tracking aspect of the system as the replan suggestion intervals is increased.
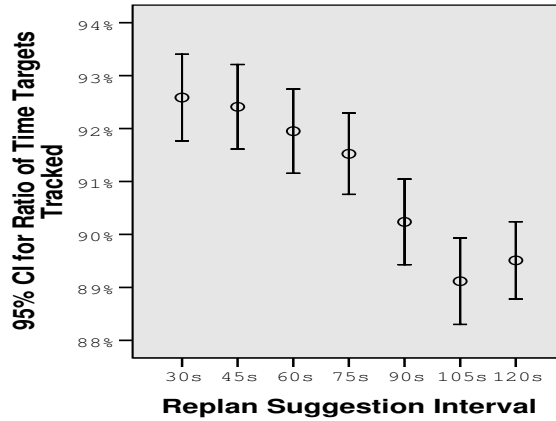
Figure 4-14: 95% confidence interval for the mean of ratio of time targets tracked over levels of the replan suggestion interval

Interaction Effects The MANOVA results in Table 4.7 demonstrated that a significant interaction occurred between the network topology and communication delay independent variables. To investigate this interaction, the ratio of time targest tracked was plotted across levels of the communication delay for each network topology. This plot is shown in Figure 4-15.

Figure 4-15: Estimated marginal means of ratio of time targets tracked vs. communication delay separated by network topology

From the means plot we can infer that effect of the communication delay is again amplified in the round-robin network, as the ratio of time targets tracked performance drops more rapidly in the round-robin network than in either the fully connected or groups networks.

(a) Fully Connected Network　　　　(b) Groups Network　　　　(c) Round-Robin Network

Figure 4-16: Percent hostiles destroyed in fully connected network (left), groups network (center), and round-robin network (right)

Figure 4-17: 95% confidence interval for the mean of percent hostiles destroyed over network topologies

4.2.6 Hostiles Destroyed

Response plane visual representations of the percent hostiles destroyed metric are shown in Figure 4-16, where the results within each network topology are presented across the communication delay and replan suggestion interval axes. Based on the response planes, the round-robin network appears to yield worse performance in hostile destruction. Increasing lag also decreases performance in all three network topologies, but especially in the round-robin network topology; this effect is consistent with the significant interaction found between the network topology and communication delay variables in the MANOVA analysis. Increasing the replan suggestion interval also appears to degrade performance, although to a lesser degree than the other independent variables. These observations are consistent with the results of the univariate ANOVA tests performed for the percent hostiles destroyed dependent variable. Both the network topology and communication delay were shown to have a significant effect with $p < 0.001$, while the replan suggestion interval was shown to have a marginally significant effect, with $p = 0.014$.

Based on the significant results in the ANOVA test, post-hoc Tukey pairwise comparisons were performed for all three independent variables. Results are outlined in the subsections below, and full post-hoc Tukey pairwise comparison results tables are included in Appendix A.

74

Figure 4-18: 95% confidence interval for the mean of percent hostiles destroyed over levels of communication delay

Network Topology Figure 4-17 shows a 95% confidence interval for the mean percent hostiles destroyed across the three network topologies. There was not a significant difference between the fully connected network and the groups network, with $p = 0.276$; however, both the fully connected and groups networks performed significantly higher than the round-robin network, with $p < 0.001$.

Communication Delay Figure 4-18 shows a 95% confidence interval for the mean percent hostiles destroyed across levels of communication delay. Percent hostiles destroyed performance was similar to the ratio of time targets tracked performance, with all factor levels producing performance in the metric at a level significantly higher values than all greater communication delays and significantly lower than all smaller communication delays, at $p < 0.001$.

Replan Suggestion Interval Figure 4-19 shows a 95% confidence interval for the mean percent hostiles destroyed across levels of the replan suggestion interval. Although this independent variable was demonstrated to have a marginally significant main effect in the univariate ANOVA performed (see Table 4.6), the pairwise Tukey comparisons did not show any significant or marginally significant differences between any of the factor levels. Although a significant effect was not observed, percent hostiles destroyed performance visually appears to trend downwards as the replan suggestion interval is increased from 45-seconds to 105-seconds. However, this trend is abruptly broken at a 120-second replan interval, where the percent hostiles destroyed performance achieves the highest value at any of the factor levels. This aspect of the results is discussed more closely in Section 5.2.2.

Interaction Effects

The MANOVA results in Table 4.7 demonstrated that a significant interaction occurred between the network topology and communication delay independent variables. A marginally significant interaction was also found between communication delay and the replan suggestion interval independent variables. Both interactions are explored below.

Figure 4-19: 95% confidence interval for the mean of percent hostiles destroyed over levels of the replan suggestion interval

Network Topology and Communication Delay Interaction To investigate this interaction, the percent hostiles destroyed metric was plotted across levels of the communication delay for each network topology. This plot is shown in Figure 4-20.
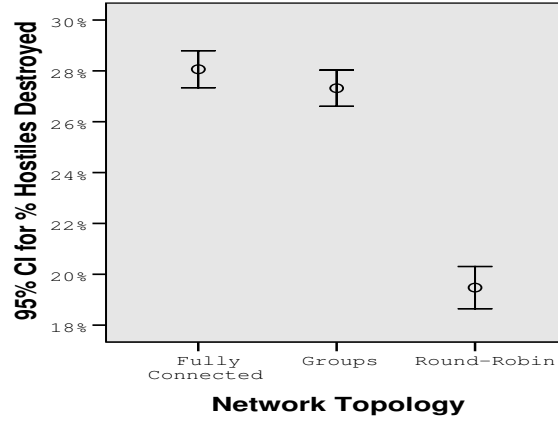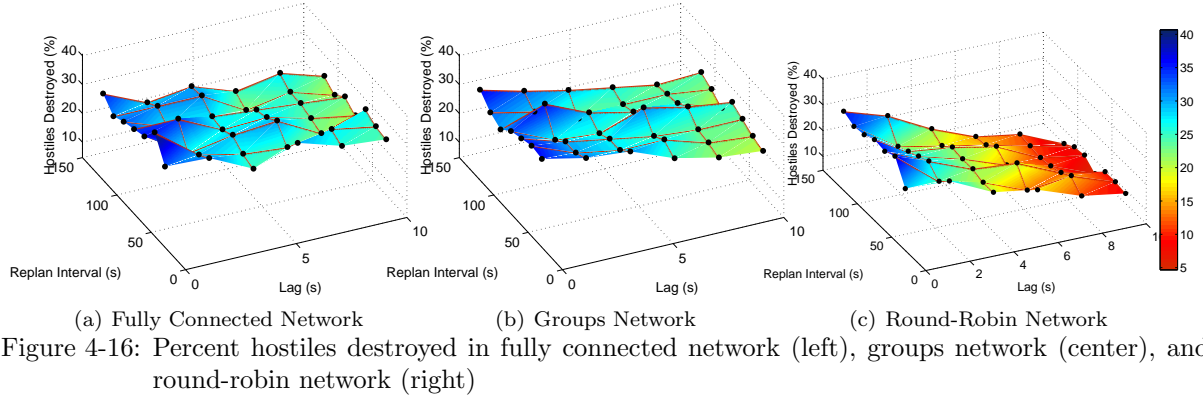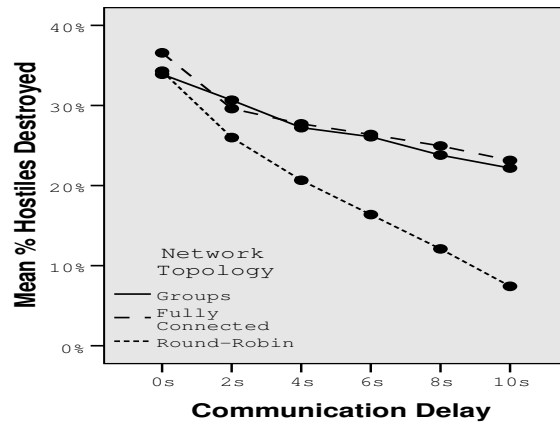

Figure 4-20: Estimated marginal means of percent hostiles destroyed vs. communication delay separated by network topology

For the percent hostiles destroyed dependent variable, the means plot in Figure 4-20 demonstrates that the effect of the communication delay is amplified in the round-robin network. Performance drops more rapidly in the round-robin network than in either the fully connected or groups networks.

4.3 Summary

Tables 4.8, 4.9, and 4.10 summarize the effects of the three independent variables on each performance metric. In Table 4.8, an increase in connectivity denotes moving from the fully connected network to the groups network and then to the round-robin network. These tables illustrate the broad trends observed in the data. Degrading the network topology resulted in the degradation of target tracking

and hostiles destroyed performance of the system, but did not affect the area coverage or targets found performance. The communication delay decreased performance in all functions of the system. However, performance leveled off in the area covered and targets found metrics while it continued to decrease in target tracking and hostile destruction. The replan suggestion interval increased performance in the area covered and targets found functions, but decreased performance in the target tracking and hostile destruction functions.

These results demonstrate that the communication availability in the system affected performance. Furthermore, the various system functions were affected in different ways by each independent variable. The next chapter will interpret these results in the context of the posed Research Questions to build an understanding of why performance in the various functions of the system responded to the changes in the independent variables in the manner observed in the data.

Table 4.8: Summary of significant results for network topology

∗ Significant result   † Marginally significant result

Source	Dependent Variable	Effect as Connectivity Decreases
Network Topology	% Area Covered	not significant
	% Targets Found	not significant
	% Time Targets Tracked∗	↓
	% Hostiles Destroyed∗	↓

Table 4.9: Summary of significant results for communication delay

∗ Significant result   † Marginally significant result

Source	Dependent Variable	Effect as Delay Increases
Communication Delay	% Area Covered∗	↓, plateaus
	% Targets Found∗	↓, plateaus
	% Time Targets Tracked∗	↓
	% Hostiles Destroyed∗	↓

Table 4.10: Summary of significant results for replan suggestion interval

∗ Significant result   † Marginally significant result

Source	Dependent Variable	Effect as Interval Increases
Replan Suggestion Interval	% Area Covered∗	↑, plateaus
	% Targets Found∗	↑
	% Time Targets Tracked∗	↓
	% Hostiles Destroyed†	↓

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Discussion and Future Work

The goal of this research was to investigate how degraded communications impact the performance of UASs. Specifically, the effects of two types of communication failures, network topology and communication delays, were investigated along with the effect of the replan suggestion interval. The research questions set forth in Section 1.3 were addressed through the following procedures:

- Research into the fields of data fusion algorithms, automated planning, human supervisory control, and network theory aided the design of an appropriate numerical experiment to answer the research questions (Chapter 2).

- A review of the results of a previous human experiment with the OPS-USERS system served as the basis for the design and implementation of a human model to emulate the types of interaction that a human operator would have with the interface (Chapter 3).

- Utilizing this model, data was gathered in simulation as a proxy for actual experiments with human operators to investigate the performance of the OPS-USERS system for a variety of degraded communication conditions and with a variety of replan suggestion intervals (Chapter 4).

The information reported in Chapter 4 addresses each research question, including the effect of communication failures on a decentralized system's performance, the tolerance of system performance to two types of communication failures, and finally the impact of HITL characteristics of the robustness of an autonomous planner to communication failures.

5.1 Conclusions for Research Questions #1 and #2

The research questions from Section 1.3 are restated and then discussed in the context of the experimental results.

Research Question #1:

How does performance of the decentralized system with a human in a supervisory control role degrade with increasing communication failures?

Section 4.2 demonstrated that both independent variables directly controlling the communication characteristics of the system, communication delay and network topology, significantly affected system performance. Furthermore, systematic trends appeared in the performance of the system as either the communication delay or the network connectivity incrementally worsened. In order to address Research Question #1, which asks how performance degrades with increasing communication failures, the trends in performance are analyzed in the context of the communication delay and network topology independent variables.

Research Question #2:

Is the system more or less tolerant to specific types of communications failures?

Research Question #2 addresses the relative effect of communication failures on the system's performance. The general trend in system performance across all metrics shows that the Fully-Connected network and the Groups network performed similarly in almost all cases. This result is intriguing as it implies that a significant fraction of the communication links in a system can be lost without a significant impact on aspects of system performance. However, this does not imply that it would be better in all cases to focus communication resources on minimizing delays.

5.1.1 Analysis of Area Coverage Performance

System performance in terms of percent area covered was affected by communication delays, with the condition of no delay producing the highest system performance in percent area. Since there was not a clear trend in the statistical significance of the results with higher levels of communication delay, it is difficult to characterize precisely what influence the delay has on the area covered performance. The best conclusion supported by the evidence provided in the experimental data suggests that there is a small increase in area coverage efficiency available to the system when the vehicles can communicate without delays. The introduction of any communication delay appears to quickly drop off area coverage performance by about 2 - 4%, with further delays (up to 10 seconds) not reducing area coverage performance any further.

This effect could be explained by the need for vehicles to collaborate on short time scales to optimize their trajectories to maximize search performance when operating near one another. This situation occurs, for example, when the weaponized UAV approaches a hostile target which another vehicle is simultaneously tracking. If the two vehicles are able to communicate their positions and intentions

promptly to one another, they may be able to more effectively minimize the overlap of the area they observe as they carry out their tasks. This effect may also be significant when two vehicles are performing the background search task in close proximity, where overlap of search effort can be reduced or eliminated by prompt communication.

Although the effect of the network topology independent variable was not found to be significant in the area covered context, it is interesting to note that this effect, where performance initially degrades slightly with communication failures, but levels off with further failures, would be expected to be applicable to the network topology as well. For example, vehicles in close proximity under a fully connected topology may be able to collaborate more promptly than under a round-robin topology. However, just as the maximum difference in area covered performance was less than 4% across levels of the communication delay, the difference across network topologies may be even smaller when results are aggregated across values of the replan suggestion interval and communication delay. The network topology could still affect the area covered results despite the univariate ANOVA failing to produce a significant result.

Summary

These performance characteristics suggest that area covered performance is relatively robust to communication degradation. When communication failure levels were pushed beyond a small threshold, the area covered performance dropped into a slightly degraded regime, performing at about 2 - 4% lower than with nominal communication levels, but then leveled off. Further degrading communication levels did not further degrade performance, suggesting that the overall environment search functionality of the system is relatively robust to communication failures. An analogous conclusion would be that the effect size of collaboration in terms of the area covered performance of the system is relatively small and sensitive to communication delays. The best performance was only achieved with very small communication delays, and even then was only slightly higher than when the vehicles' ability to collaborate over short time scales was significantly impaired.

In terms of the percent area covered metric, it then makes sense to minimize the communication delay between vehicles as much as possible, especially between vehicles which are operating in close proximity, where the need to collaborate quickly to optimize search path generation is greatest. This conclusion represents an argument for ensuring vehicles have short-range, low-latency connectivity to maximize their ability to collaborate effectively on area coverage.

5.1.2   Analysis of Target Search Performance

Overall, the system's performance in the target search task, as measured by the percent targets found dependent variable, corresponds closely with performance in the percent area covered dependent variable. As with area covered, the percent targets found metric did not show statistical significance across the

three network topologies, but did show significance across levels of the communication delay. Figure 4-9 shows that the mean target found performances all fell between 53% and 57%. Consistent with what was observed with the area covered metric, a 2 - 4% increase in the percent targets found was observed when the communication delay was 0 seconds, and there was not a significant difference among any of the factor levels with a delay greater than 0 seconds. This trend is identical to the trend observed in the area covered metric across values of the communication delay variable. Furthermore, the magnitude of the performance difference corresponds closely; in both variables, an approximate improvement of 2 - 4% was observed with the delay set to 0 seconds.

This correspondence can be explained by reasoning about the probability of finding a target during each mission scenario. Given that the vehicles assume that the targets will be distributed uniformly throughout the environment, it would be reasonable to expect an $x\%$ increase in the area covered metric to lead to the same $x\%$ increase in the targets found metric on average. In other words, the factors influencing the differences in the area covered metrics may translate directly into changes in the percent targets found metric as well; the best strategy to maximize the percent targets found metric is to search the largest percentage of the environment as possible, simultaneously maximizing the percent area covered metric.

Considered in this context, the similarity in the results for percent area covered and percent targets found is not surprising. Consistency across the effect of the communication delay and network topology independent variables reinforce this conclusion. One difference between the percent area covered performance and the communication delay performance is that the marginally significant interaction observed for the area covered metric was not present in the targets found performance. This difference is not too troubling, as the range of means for the percent targets found across values of the communication delay or network topology was less than 4%. This performance difference is much smaller than the performance differences observed for the ratio of time targets tracked and hostiles destroyed metrics.

Summary

The percent targets found metric provides another argument for minimization of the communication delay between agents in the system, as performance fell with the addition of any communication delay. The network topology variable ultimately was not demonstrated to affect performance. When the communication delay is near zero, it follows that the effects of the network topology would be minimized; even if communication had to be routed through several agents, each would incur only a small delay and the total delay would also ultimately be small. The results suggest that a delay of less than 2-seconds must be achieved to maximize performance, although it is unclear what the exact tolerable delay would be.

5.1.3 Robustness of Area Coverage and Targets Search to Communication Failures

Considering the large range of communication conditions tested throughout this experiment, it is relatively surprising to find that the range of the mean area covered and targets found performance averaged over multiple trials is relatively small. The range of values was less than 4% for the communication delay variable, and less than 1% for the network topology variable. Considering also that the range of values was observed to be almost twice as large across the replan suggestion interval independent variable, this experiment has provided evidence supporting the conclusion that the system's performance in percent area covered and percent targets found is robust to communication failures.

In order to understand why this would be the case, the dynamics of the simulation must be considered. For example, discovering a target in the OPS-USERS system represents an instantaneous increase in performance in terms of the percent targets found metric. However, the new target will increase the task load on the vehicles, making less time available for searching the environment as the target is tracked and, if it is hostile, engaged by the WUAV. With less time available to search the environment, the probability of finding more new targets is decreased.

In this sense, the system provides a negative feedback loop which may serve to regulate the number of targets that were found throughout each mission. Even if the vehicles happened to discover many targets near the beginning of a mission, they could be kept busy throughout the remainder of the simulation servicing these targets. The operator could intervene to further encourage search behavior, but in the mission specification used throughout these experiments, there was a relatively high number of targets compared to the number of vehicles – in each mission the four vehicles were searching for 10 targets. Therefore, if many targets were found, performing functions other than tracking the targets and engaging hostile targets with the WUAV represents a trade-off between a pair of metrics in the system. For example, attempting to increase the percent area searched would diminish the vehicles' ability to track the known targets, decreasing the ratio of time targets were tracked metric.

5.1.4 Analysis of Target Tracking Performance

The ratio of time targets tracked performance metric responded dramatically and consistently to both communication independent variables. The ability of the vehicles to track targets was clearly impaired by increasing communication delays, as every increase of 2 seconds in communication delay represented a significant average decrease in the ratio of time targets tracked metric of approximately 1.5%. Furthermore, the target tracking performance was diminished overall by approximately 4% in the round-robin network topology compared to the fully connected or groups topologies.

These results can be interpreted by considering the target tracking process in the OPS-USERS sys-

tem. As described in Section 3.1.4, target tracking tasks are generated by the Centralized Mission Manager (CMM) when the CMM's estimate of the uncertainty associated with a target's location reaches a threshold limit. In order for the tracking task to be assigned to a vehicle, the task must be placed on the approved task list by the human operator. This requires communication from the CMM to each vehicle for the vehicles to become aware that the tracking task is available. Furthermore, the vehicles must collaborate amongst themselves to decide which will perform the tracking task. Therefore, the delay between a tracking task becoming available and the task being carried out by one of the vehicles will include the amount of time taken for the CMM to communicate with each vehicle, and for the vehicles to collaboratively generate a conflict-free task assignment.

These delays in the assignment of tracking tasks are manifested in the vehicles' late arrival to target tracking assignments, causing the targets to become lost for short periods of time and degrading the ratio of time targets tracked performance. Inflicting further delays on the assignment of tracking tasks to vehicles would then be expected to further degrade the ratio of time targets tracked performance, and this is indeed the effect observed in the data across levels of the communication delay independent variable. Similarly, reducing the connectivity of the network topology will further delay communication between certain pairs of vehicles, which in turn lengthens the amount of time required for the vehicles to generate task assignments.

Interaction Effects

The effect of the network topology variable can also be explained by considering communication requirements of the target tracking process. The fully connected and groups topologies produced near-identical performance, suggesting that the amount of delay before tracking tasks were assigned to vehicles was similar in each network, or was not long enough to cause targets to become lost. This similarity in performance is reasonable because all vehicles have a direct communication link to the Centralized Mission Manager (CMM) in both the fully connected and groups networks. As these specific links have been identified as directly impacting the delay in assigning track tasks to vehicles, the fact that the groups network does not include failures in these links helps to explain the similarity in performance to the fully connected network. The inter-vehicle links which are missing in the groups network topology may have a smaller impact on the target tracking function of the system.

However, the degradation in performance is more prominent in the round-robin network. Given the links available in this topology, the total delay before a message originating at the CMM reaches every vehicle is 4 times longer than in either the fully connected or groups networks with the same level of communication delay. This is caused by the round-robin network needing to route messages through other agents, experiencing a communication delay at each agent and amplifying the total delay before a message reaches it's final destination agent. The total delay in communication between pairs of vehicles

is similarly amplified in the round-robin network.

This effect is demonstrated quite clearly by the interaction between the communication delay and network topology variables. Figure 4-15 clearly shows that the groups and fully connected networks respond to the communication delay variable identically, but that the round-robin network causes the ratio of time targets tracked performance to diminish more rapidly with increasing delays.

Summary

The experimental results suggest two important characteristics of the target tracking performance. First, target tracking performance is increased when communication delays are minimized. Second, the size of this effect is significantly affected by the network topology. The system appears capable of tolerating delays up to a point between the creation of a target tracking by the CMM and its assignment to a specific vehicle before performance is significantly impacted. For example, increasing the communication delay from 0 to 10 seconds in the fully connected or groups networks diminished the target tracking performance only slightly, by about 4%. However, when the network connectivity was more severely impaired in the round-robin network, the same increase in communication delay from 0 to 10 seconds caused a 15% drop in the ratio of time targets tracked performance.

In order to achieve the best possible target tracking performance, it seems clear that either a near-zero level of communication delay or a threshold level of connectivity in the network topology must be maintained. When the communication delay is low, less connectivity can be tolerated in the network topology without a negative impact on performance. When the connectivity of the topology is above some threshold, greater communication delays can be tolerated.

Unfortunately, in this research only three static communication network topologies were tested through simulation, and it is not clear at what threshold level of network connectivity the performance in target tracking would begin to drop significantly as it did in the round-robin network. We can only conclude that the threshold level falls somewhere in between the connectivity of the groups network and the connectivity of the round-robin network. Another factor mentioned above that would be of interest is the relative importance of each specific directed communication link in the system. For example, the number of communication links in a topology could be held constant, but the links could be shifted between various pairs of vehicles. This is tantamount to maintaining the structure of the networks pictured in Figure 4-4 but rearranging the labels on each node.

As previously mentioned, the communication link between each vehicle and the CMM likely impacts the performance of the target tracking significantly. If the rest of the specific communication links could be similarly analyzed to determine their relative importance, the information would be instructive as to where to focus communication resources in the system to achieve optimal performance.

5.1.5 Analysis of Hostile Destruction Performance

System performance in terms of the hostiles destroyed metric responded to communication failures in a manner similar to the ratio of time targets tracked metric. As both metrics represent attendance to task-driven events, modulating system parameters which affect changes in the system's ability to promptly service tasks simultaneously impacts both hostiles destroyed and ratio of time targets tracked performance. As there was not a significant difference in performance between the fully connected and groups network, but the round-robin network incurred a large 8-9% performance degradation, we can again conclude that system performance is robust to a threshold level of network topology connectivity. As noted before, this threshold level of connectivity lies in between the groups and round-robin networks, although the limitations of this experimental design preclude providing an estimate of the level.

Th effect of the communication delay proved consistent in the hostile destroyed metric, with increasing delays degrading performance. As the difference in performance was drastic across communication delays, with the system destroying 15% more of the available hostiles at a 0-second delay than at a 10-second delay, the hostiles destroyed metric provides one of the most compelling arguments to show that optimizing the communication system of a UAS must be a principal design consideration.

Analysis of the hostile destruction process provides insight into how a degraded communication environment could be expected to affect performance. As outlined in Section 3.1.4, the process for destroying a hostile target involves several steps, each of which may be affected differently by changes in the availability of communications. For example, a task to destroy a hostile target cannot be created unless the target is currently being tracked by the vehicles, and communication degradation has already been shown to affect the ability of the system to consistently track targets. The hostile destruction task cannot be assigned to a vehicle until the operator approves a schedule which contains it, a process involving approval of a new schedule which requires the task to be communicated from the CMM to all of the vehicles. In the final step the vehicles perform the distributed planning algorithm to assign the hostile destruction task to one of the WUAVs.

Given the reliance on communication among the vehicles or between the vehicles and the CMM, it is not surprising that the hostile destruction performance metric proved to be the most sensitive to the communication delay or network topology. Because of the increased complexity of the hostile destruction function of the system, there are more opportunities for degraded communications to affect some aspect of the process. For example, an increased communication delay may cause the messages indicating the initial discovery of a hostile target not to reach the CMM for some time. In the worst case under the round-robin network with a 10-second communication delay, the information could take up to 40 seconds to reach the CMM. This represents a significant portion of the mission duration, almost 7%. The time required for the vehicles to collaborate to generate a task assignment would also be increased under more

degraded communications.

In the context of the OPS-USERS system's ability to destroy hostile targets, the question of the relative importance of each communication link in the system is again raised. One difference between the hostiles destroyed and target tracking metrics is the explicit reliance on communication with a subset of the vehicles (the WUAVs) with the ability to engage hostile targets. Testing several network topologies which differed in the connectedness of the WUAV would provide more insight into this aspect of the system performance.

Interaction Effects

As with the target tracking performance of the system, the effect of the communication delay was significantly impacted by the network topology, with the round-robin topology amplifying the effects of communication delays. This interaction between the communication delay and network topology is demonstrated clearly in Figure 4-15. The groups and fully connected networks respond to the communication delay variable identically, but the round-robin network causes the ratio of time targets tracked performance to diminish more rapidly with increasing delays.

Performance in the hostiles destroyed metric was affected most dramatically of all the performance metrics under degraded communication conditions. The results further serve to reinforce the notion that, in order to achieve satisfactory system performance, either communication delays must absolutely be minimized or a threshold level of network connectivity must be maintained.

Summary

Two distinct modes of degraded performance were observed in the results. For the percent area covered and percent targets found metric, any small level of communication disruption caused performance to drop into a degraded regime. In this case, unless near-perfect communication with an emphasis on minimizing delays was achieved, there was not a strong dependence on the communication characteristics. Therefore, if communication degradation must be tolerated, it is reasonable to maximize performance in terms of the ratio of time targets tracked and percent hostiles destroyed metrics. The results clearly show that when network connectivity fell sufficiently low, as in the round-robin network, the marginal value of increasing the network connectivity is higher than the value of minimizing delays.

Because of the interaction observed between the communication delay and network topology independent variables, it is difficult to draw conclusions about the relative robustness of the system to either type of communication failure. In future experiments, research question #2 could be better investigated by choosing a different set of independent variables. For example, maintaining the same overall level of network connectivity but modulating which specific links are available. This may provide insight into which types of links (i.e. UAV-to-UAV or UAV-to-CMM) have a larger effect on system performance.

5.2 Conclusions for Research Question #3

The final research question is restated and then discussed in the context of the experimental results.

Research Question #3:

To what extent does HITL control increase the robustness of autonomous planners to communication failure?

Research Question #3 asks whether the characteristics of HITL control can be modulated to compensate for communication failures in the system. In the experiment carried out for this research, the replan suggestion interval independent variable represents the mechanism for influencing the human supervisory control aspect of the system. Performance variances associated with this independent variable are outlined below, along with an analysis of how this aspect of the supervisory control influenced the robustness of the system performance to communication failures.

Univariate ANOVA results showed that the replan suggestion interval significantly affected performance in terms of percent area covered, percent targets found, and the ratio of time targets tracked metrics with $p < 0.001$. The percent hostiles destroyed metric also showed a marginally significant effect with $p = 0.014$.

5.2.1 Area Covered and Targets Found Performance

The replan suggestion interval was found to significantly impact both percent area covered and percent targets found performance. In all three network topologies, a general trend of increasing performance was observed as the replan suggestion interval was increased. An explanation for this relationship is that decreasing the frequency of the replanning task can result in an increase in the amount of time that vehicles are not explicitly assigned a task. In this condition, the vehicles perform the "background" search task. Replanning less frequently is likely to increase the amount of time that vehicles spend performing the background search task, and can therefore be expected to increase performance in terms of the area covered metric or targets found metric.

Area Covered Performance Although increasing the replan suggestion interval from 30-seconds to 75-seconds results in an approximately 6% increase in performance, increasing the replan suggestion interval beyond 75-seconds did not continue to improve performance. One possible explanation for the apparent plateau in performance beyond a 75-second replan interval is that other events in the system may also trigger replanning events. When the replan interval is set sufficiently high (i.e. a long interval between replan suggestions), the frequency of event driven replanning may cause the effective replanning rate to be higher than the replan suggestion interval.

Targets Found Performance As postulated above in Section 5.1.2 when considering degraded communications, a close correspondence between the area cover and targets f18ound metrics would be expected as increasing the area covered increases the percent targets found in expectation. This was supported by decreases of the same magnitude across values of the communication delay independent variable.

This similarity in response applies to the replan suggestion interval as well, although the correspondence was not as close. One marginal inconsistency is the behavior at larger values of the replan suggestion interval. Where the performance in area coverage plateaus at the 75-second interval, with no significant difference between increase between the 75- and 120-second intervals, performance in the percent targets found metric may have continued to increase, and was marginally higher at the 120-second interval than at the 75-second interval. The range of values observed across replan suggestion intervals was consistent, however, with a 4-5% increase in performance in both metrics as the replan suggestion interval was increased from 30 seconds to 120 seconds.

5.2.2 Target Tracking and Hostile Destruction Performance

The effect of the replan suggestion interval on both target tracking and hostile destruction performance is discussed below. Although the response of these dependent variables to the overall communication degradation corresponded somewhat, there is not a clear congruency in response across the target tracking and hostile destruction metrics to the replan suggestion interval.

Target Tracking Performance Target tracking performance was clearly diminished by increasing the replan suggestion interval. As target tracking requires reassignment of the tracking task to a vehicle after each time a target is revisited, waiting longer before replanning would be expected to afford the vehicles less flexibility in accomplishing the tracking task as well as other tasks. In some cases where a sufficiently long interval has passed before replanning, it may not be feasible to complete a tracking task before a target becomes lost. In this way, a lower replan suggestion interval can be expected to increase performance in the ratio of time targets tracked metric.

Hostile Destruction Performance Although replanning remains an integral component of the hostile destruction process in the OPS-USERS system, replanning to complete hostile destruction tasks is not driven solely by the replan suggestion interval. As the addition of novel tasks into the system represent a situation where the operator is aware that replanning is necessary for the additional task to be taken into account, the replan suggestion interval will not affect the first attempt to assign the hostile destruction task to a vehicle. However, it's possible that the task will not be placed on the approved task list during replanning, or that the WUAV which is assigned the task may later decide that it cannot complete the task. Either situation will require further replanning before the hostile destruction task can be carried out. In these cases, the replan suggestion interval would significantly impact performance.

Although the replan suggestion interval was shown to have a marginally significant effect on the percent hostiles destroyed metric, there is no clear trend of either increasing or decreasing performance across the range of replan intervals tested. Across all values, the average performance varied by less than 3%, with no significant difference appearing between performance at the 30-second, 75-second, or 120-second intervals. One anomaly which appeared in the otherwise relatively continuous response is a sharp increase in hostile destruction performance at the 120-second interval after several small but consistent decreases in performance between the 45-second interval and 105-second interval. In fact, performance at the 105- and 120-second intervals represent that minimum and maximum mean performance values respectively observed across all values of the replan interval.

There is no reason to expect that increasing the replan suggestion interval from 105-seconds to 120-seconds would elicit such profoundly different performance results. This artifact in the data may be related to the specific configuration of missions utilized in this experiment. For example, it's possible that waiting slightly longer to replan near the beginning of a mission may significantly change the probability of finding a specific target, which happens to be hostile, early on in one of the three mission scenarios tested. There is no evidence to support a theory other than that this apparent pattern emerged by chance. Furthermore, the results observed at the 120-second replan interval are not inconsistent with observations made in previous experiments. For example, Section 4.1.3 demonstrated that performance at the 45- and 120-second intervals was similar.

5.2.3 Conclusion

Evidence to support any interaction between the replan suggestion interval and robustness to communication failures would have been expected to appear as an interaction effect between the replan suggestion interval and either the network topology or communication delay independent variables. As these interaction effects were not significant, there is no evidence to conclude that the parameters of the HITL control in the OPS-USERS system can be modulated to compensate for communication failures. This suggests that the OPS-USERS system provides a robust platform for human supervisory control of multiple UAVs under a variety of communication failures, and isolates the operators from the effects of communication failures.

5.3 Future Work

5.3.1 Further Human Model Validation Experiments

The relevance of the results generated from this work hinge on the validity of the human model utilized to generate the data set. Although the model was examined in some detail in Section 4.1 and found to match the original experimental results reasonably well, this does not speak to the ability of the model to

remain consistent with the performance of human operators under degraded communication conditions.

Therefore a natural follow-on experiment to this work would be to validate the results of the human model with similar human experiments. A small subset of the experimental conditions investigated in simulation could be chosen, for example at each extreme value of the lag and replan interval. Observing similar trends in performance in actual human subject data would serve to increase confidence in the human model and strengthen the evidence supporting conclusions made in this thesis.

5.3.2 Further Characterization of Area Coverage Performance

An important extension of this work would be to introduce incremental network topologies to characterize performance under communication failures with a greater resolution. This may be useful as the number of communication delays tested was twice that of the number of network topologies tested. Ultimately it would be ideal to prescribe a specific threshold of communication availability which, if maintained, would guarantee area coverage performance did not drop into the degraded regime. Another intriguing aspect of this extension would be to investigate the relative importance of each communication link the network topology to characterize where the important communication channels are in the context of each system function.

5.3.3 The Impact of the Replan Suggestion Interval on Performance

Experimental results demonstrated that increasing the replan suggesting interval improved performance in area coverage and targets found. However, target tracking and hostile destruction performance were adversely affected. These results suggest that the replan suggestion interval influences a trade-off between performance in these two sets of system functions. In the operation of the OPS-USERS system, replanning serves as a mechanism for shifting the vehicles' attention to either specific task-driven functions such as target tracking or hostile destruction or the background environment search task which influences area covered and targets found performance. More research is needed to characterize such trade-offs and to determine the optimal balance between the interrelated functions of a UAS.

5.4 Summary

This work has demonstrated that communication availability in UASs represents an important factor in determining system performance. Through the analysis of the effect of communication failures on several types of functions of a UAS including environment search, target search, target tracking, and hostile destruction, this thesis provides a basis for understanding the implications of communication availability on a diverse set of collaborative tasks. This research may serve as a guide to instruct future UAS system designers on the best way to allocate communication resources for UASs with a variety of mission objectives.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix A

Statistical results

Table A.1: MANOVA Results over all experimental data

α=0.01   * Significant result   † Marginally significant result

Source	Dependent Variable	Type III Sum of Squares	df	Mean Square	F	Sig.
Network	AC	.023	2	.011	.715	0.489
	TF	.045	2	.023	1.736	0.176
	RT	1.586	2	.793	107.686	0.000
	HD	5.703	2	2.851	190.782	0.000
Lag	AC	.930	5	.186	11.649	0.000*
	TF	.498	5	.100	7.617	0.000*
	RT	2.840	5	.568	77.140	0.000*
	HD	12.220	5	2.444	163.518	0.000*
Interval	AC	2.085	6	.347	21.762	0.000*
	TF	.706	6	.118	9.004	0.000*
	RT	.649	6	.108	14.696	0.000*
	HD	.238	6	.040	2.656	0.014†
Network * Lag	AC	.310	10	.031	1.942	0.036
	TF	.188	10	.019	1.439	0.157
	RT	1.169	10	.117	15.878	0.000*
	HD	2.036	10	.204	13.623	0.000*
Network * Interval	AC	.192	12	.016	1.000	0.446
	TF	.166	12	.014	1.058	0.392
	RT	.094	12	.008	1.062	0.388
	HD	.105	12	.009	.588	0.854
Lag * Interval	AC	.498	30	.017	1.041	0.406
	TF	.306	30	.010	.781	0.797
	RT	.183	30	.006	.830	0.730
	HD	.676	30	.023	1.509	0.037

Table A.2: Percent area covered over communication delay Tukey pairwise comparisons

Source	Mean Difference	Sig.
0-seconds vs. 4-seconds	0.0320	0.011†
0-seconds vs. 10-seconds	0.0433	0.000∗
4-seconds vs. 10-seconds	0.0114	0.601

Table A.3: Percent area covered over replan suggestion interval Tukey pairwise comparisons

Source	Mean Difference	Sig.
30-seconds vs. 75-seconds	-0.0590	0.000∗
30-seconds vs. 120-seconds	-0.0643	0.000∗
75-seconds vs. 120-seconds	-0.0053	0.993

Table A.4: Percent targets found over communication delay Tukey pairwise comparisons

Source	Mean Difference	Sig.
0-seconds vs. 4-seconds	0.033	0.004∗
0-seconds vs. 10-seconds	0.030	0.000∗
4-seconds vs. 10-seconds	-0.003	0.997

Table A.5: Percent targets found over replan suggestion interval Tukey pairwise comparisons

Source	Mean Difference	Sig.
30-seconds vs. 75-seconds	-0.025	0.005∗
30-seconds vs. 120-seconds	-0.044	0.000∗
75-seconds vs. 120-seconds	-0.019	0.088

Table A.6: Ratio of time targets tracked over network topology Tukey pairwise comparisons

Source	Mean Difference	Sig.
Fully Connected vs. Groups	-0.0017	0.869
Fully Connected vs. Round-Robin	0.0426	0.000∗
Groups vs. Round-Robin	0.0443	0.000∗

Table A.7: Ratio of time targets tracked over communication delay Tukey pairwise comparisons

Source	Mean Difference	Sig.
0-seconds vs. 4-seconds	0.0142	0.039†
0-seconds vs. 10-seconds	0.0298	0.000∗
4-seconds vs. 10-seconds	0.0506	0.000∗

Table A.8: Ratio of time targets tracked over replan suggestion interval Tukey pairwise comparisons

∗ Significant result   † Marginally significant result

Source	Mean Difference	Sig.
30-seconds vs. 75-seconds	0.0106	0.392
30-seconds vs. 120-seconds	0.0308	0.000*
75-seconds vs. 120-seconds	0.0201	0.002*

Table A.9: Percent hostiles destroyed over network topology Tukey pairwise comparisons

∗ Significant result   † Marginally significant result

Source	Mean Difference	Sig.
Fully Connected vs. Groups	-0.007	0.276
Fully Connected vs. Round-Robin	0.86	0.000*
Groups vs. Round-Robin	0.78	0.000*

Table A.10: Percent hostiles destroyed over communication delay Tukey pairwise comparisons

∗ Significant result   † Marginally significant result

Source	Mean Difference	Sig.
0-seconds vs. 4-seconds	0.097	0.000*
0-seconds vs. 10-seconds	0.173	0.000*
4-seconds vs. 10-seconds	0.076	0.000*

Table A.11: Percent hostiles destroyed over replan suggestion interval Tukey pairwise comparisons

∗ Significant result   † Marginally significant result

Source	Mean Difference	Sig.
30-seconds vs. 75-seconds	-0.001	1.000
30-seconds vs. 120-seconds	-0.011	0.749
75-seconds vs. 120-seconds	-0.010	0.995

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix B

Descriptive statistics

Table B.1: Descriptive statistics, all data

% Dependent Variable	N	Min.	Max.	Mean	Median	Std. Dev.
630 1890 3150 % Area Covered	3780	0.232	0.996	0.610	0.597	0.129
% Targets Found	3780	0.0	1.0	0.544	0.5	0.116
% Time Targets Tracked	3780	0.0	1.0	0.910	0.937	0.094
% Hostiles Destroyed	3780	0.0	0.8	0.250	0.2	0.143

Table B.2: Descriptive statistics, fully connected network

% Dependent Variable	N	Min.	Max.	Mean	Median	Std. Dev.
% Area Covered	1260	0.294	0.979	0.612	0.630	0.128
% Targets Found	1260	0.0	1.0	0.542	0.5	0.113
% Time Targets Tracked	1260	0.0	1.0	0.924	0.945	0.081
% Hostiles Destroyed	1260	0.0	0.8	0.281	0.2	0.132

Table B.3: Descriptive statistics, groups network

% Dependent Variable	N	Min.	Max.	Mean	Median	Std. Dev.
% Area Covered	1260	0.232	0.978	0.612	0.602	0.129
% Targets Found	1260	0.2	0.9	0.549	0.5	0.115
% Time Targets Tracked	1260	0.534	1.0	0.926	0.949	0.078
% Hostiles Destroyed	1260	0.0	0.8	0.273	0.2	0.129

Table B.4: Descriptive statistics, round-robin network

% Dependent Variable	N	Min.	Max.	Mean	Median	Std. Dev.
% Area Covered	1260	0.312	0.996	0.607	0.589	0.131
% Targets Found	1260	0.0	0.9	0.541	0.5	0.119
% Time Targets Tracked	1260	0.0	1.0	0.882	0.907	0.114
% Hostiles Destroyed	1260	0.0	0.8	0.195	0.2	0.150

References

[1] General Atomics Press Release, "Predator-Series UAS family achieves one million flight hours." http://www.ga.com/news.php?id=284&read=1. Accessed May 10th, 2010.

[2] U.S. Government Accountability Office, "Defense Acquisitions: Opportunities Exist to Achieve Greater Commonality and Efficiencies Among Unmanned Aircraft." http://www.gao.gov/new.items/d09520.pdf. Accessed May 10th, 2010.

[3] U.S. Air Force, "MQ-1 Predator Fact Sheet." http://www.af.mil/information/factsheets/factsheet.asp?fsID=122. Accessed May 10th, 2010.

[4] Alighanbari, M. and J. How, "Decentralized task assignment for unmanned aerial vehicles," in *Proceedings of the 44th IEEE Conference on Decision and Control*, (Seville, Spain), pp. 5668–5673, December 2005.

[5] Sheridan, T. B., *Telerobotics, Automation, and Human Supervisory Control*. Cambridge, Massachusetts, USA: The MIT Press, 1992.

[6] Clare, A., C. Hart, and M. L. Cummings, "Assessing operator workload and performance in expeditionary multiple unmanned vehicle control," in *48th AIAA Aerospace Sciences Meeting*, (Orlando, Florida, USA), January 2010.

[7] Alighanbari, M., L. Bertuccelli, and J. How, "Filter-embedded UAV task assignment algorithms for dynamic environments," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, (Providence, Rhode Island, USA), August 2004.

[8] Durrant-Whyte, H. and M. Stevens, "Data fusion in decentralised sensing networks," in *Proceedings of the International Conference on Information Fusion*, (Montreal, Quebec, Canada), July 2001.

[9] Beard, R. and V. Stepanyan, "Synchronization of information in distributed multiple vehicle coordinated control," in *Proceedings of 42nd IEEE Conference on Decision and Control*, (Maui, Hawaii, USA), pp. 2029–2034, December 2003.

[10] Alanyali, M., S. Venkatesh, O. Savas, and S. Aeron, "Distributed bayesian hypothesis testing in sensor networks," in *Proceedings of the American Control Conference*, vol. 6, (Boston, Massachusetts, USA), pp. 5369–5374, June 2004.

[11] Brunet, L., H. Choi, and J. How, "Consensus-based auction approaches for decentralized task assignment," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, (Honolulu, Hawaii, USA), August 2008.

[12] Choi, H., L. Brunet, and J. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 912–926, 2009.

[13] Cortés, J., "Analysis and design of distributed algorithms for chi-consensus," in *Proceedings of the 45th IEEE Conference on Decision and Control*, (San Diego, California, USA), pp. 3363–3368, 2006.

[14] Ren, W., R. Beard, and E. Atkins, "A survey of consensus problems in multi-agent coordination," in *Proceedings of the American Control Conference*, vol. 3, (Portland, Oregon, USA), pp. 1859–1864, June 2005.

[15] Glavaški, S., M. Chaves, R. Day, P. Nag, A. Williams, and W. Zhang, "Vehicle networks: Achieving regular formation," in *Proceedings of the American Control Conference*, (Denver, Colorador, USA), pp. 4095–4100, June 2003.

[16] Fowler, J. and R. D Andrea, "Distributed Control of Close Formation Flight (I)," in *Proceedings of the 41st IEEE Conference on Decision and Control*, vol. 3, (Las Vegas, Nevada, USA), pp. 2972–2977, 2002.

[17] Kingston, D., W. Ren, and R. Beard, "Consensus algorithms are input-to-state stable," in *Proceedings of the American Control Conference*, vol. 3, (Portland, Oregon, USA), pp. 1686–1690, June 2005.

[18] Beard, R., T. McLain, M. Goodrich, and E. Anderson, "Coordinated target assignment and intercept for unmanned air vehicles," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 6, pp. 911–922, 2002.

[19] Bellingham, J., M. Tillerson, M. Alighanbari, and J. How, "Cooperative path planning for multiple UAVs in dynamic and uncertain environments," in *Proceedings of the 41st IEEE Conference on Decision and Control*, vol. 3, (Las Vegas, Nevada, USA), 2002.

[20] Inalhan, G., D. Stipanovic, and C. Tomlin, "Decentralized optimization, with application to multiple aircraft coordination," in *Proceedings of the 41st IEEE Conference on Decision and Control*, vol. 1, (Las Vegas, Nevada, USA), 2002.

[21] Flint, M., M. Polycarpou, and E. Fernandez-Gaucherand, "Cooperative control for multiple autonomous UAV's searching for targets," in *Proceedings of the 41st IEEE Conference on Decision and Control*, vol. 3, (Las Vegas, Nevada, USA), pp. 2823–2828, 2002.

[22] Bauso, D., L. Giarre, and R. Pesenti, "Distributed consensus protocols for coordinating buyers," in *Proceedings of the 42nd IEEE Conference on Decision and Control*, vol. 1, (Maui, Hawaii, USA), December 2003.

[23] Ren, W., R. Beard, and D. Kingston, "Multi-agent kalman consensus with relative uncertainty," in *Proceedings of the American Control Conference*, (Portland, Oregon, USA), pp. 1865–1870, 2005.

[24] Alighanbari, M. and J. How, "An unbiased Kalman consensus algorithm," in *Proceedings of the American Control Conference*, (Minneapolis, Minnesota USA), pp. 3159–3524, June 2006.

[25] Ren, W., R. Beard, and E. Atkins, "Information consensus in multivehicle cooperative control," *IEEE Control Systems Magazine*, vol. 27, no. 2, pp. 71–82, 2007.

[26] Moallemi, C. C. and B. V. Roy, "Consensus propagation," *IEEE Transactions on Information Theory*, vol. 52, no. 11, pp. 4753–4766, 2006.

[27] Olfati-Saber, R. and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[28] Olfati-Saber, R. and R. M. Murray, "Consensus protocols for networks of dynamic agents," in *Proceedings of the American Control Conference*, (Denver, Colorado, USA), pp. 951–956, June 2003.

[29] Tahbaz-Salehi, A. and A. Jadbabaie, "On consensus over random networks," *IEEE Transactions on Automatic Control*, vol. 53, no. 3, pp. 791–795, 2008.

[30] Alighanbari, M. and J. How, "Robust decentralized task assignment for cooperative UAVs," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, (Keystone, Colorado, USA), August 2006.

[31] Broch, J., D. Maltz, D. Johnson, Y. Hu, and J. Jetcheva, "A performance comparison of multi-hop wireless ad hoc network routing protocols," in *Proceedings of the 4th annual ACM/IEEE International Conference on Mobile Computing and Networking*, (Dallas, Texas, USA), p. 97, ACM, 1998.

[32] Brunet, L., H. Choi, and J. How, "Consensus-based auctions for decentralized task assignment," Master's thesis, Massachusetts Insitute of Technology, Dept. Aeronautics and Astronautics, 2008.

[33] Meyer, K., M. Erlinger, J. Betser, C. Sunshine, G. Goldszmidt, and Y. Yemini, "Decentralizing control and intelligence in network management," in *Proceedings of the 4th International Symposium on Integrated Network Management*, (Santa Barbara, California, USA), pp. 4–16, 1995.

[34] Jin, Y., A. Minai, and M. Polycarpou, "Cooperative real-time search and task allocation in UAV teams," in *Proceedings of the 42nd IEEE Conference on Decision and Control*, (Maui, Hawaii, USA), December 2003.

[35] Cassandras, C. and W. Li, "A receding horizon approach for solving some cooperative control problems," in *Proceedings of the 41st IEEE Conference on Decision and Control*, vol. 4, (Las Vegas, Nevada, USA), 2002.

[36] Schumacher, C. and P. Chandler, "Task allocation for wide area search munitions," in *Proceedings of the American Control Conference*, vol. 3, (Anchorage, Alaska, USA), pp. 1917–1922, May 2002.

[37] Wu, C. W., "Synchronization and convergence of linear dynamics in random directed networks,," *IEEE Transactions on Automatic Control*, vol. 51, no. 7, pp. 1207–1210, 2006.

[38] Clark, C., S. Rock, and J. Latombe, "Motion planning for multiple mobile robots using dynamic networks," in *Proceedings of the IEEE International Conference on Robotics and Automation*, (Taipei, Taiwan), September 2003.

[39] Hatano, Y. and M. Mesbahi, "Agreement over random networks," *IEEE Transactions on Automatic Control*, vol. 50, no. 11, pp. 1867–1872, 2005.

[40] Alighanbari, M., "Task assignment algorithms for teams of UAVs in dynamic environments," Master's thesis, Massachusetts Insitute of Technology, Dept. Aeronautics and Astronautics, 2004.

[41] Schumacher, C., P. Chandler, M. Pachter, and L. Pachter, "UAV Task Assignment with Timing Constraints via Mixed-Integer Linear Programming," in *AIAA 3rd "Unmanned Unlimited" Technical Conference*, (Chicago, Illinois, USA), September 2004.

[42] Xiao, L. and S. Boyd, "Fast linear iterations forb distributed averaging," in *Proceedings of the 42nd IEEE Conference on Decision and Control*, vol. 5, (Maui, Hawaii, USA), December 2003.

[43] Szer, D., F. Charpillet, , and S. Zilberstein, "MAA*: A heuristic search algorithm for solving decentralized POMDPs," in *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence*, (Edinburgh, Scotland), 2005.

[44] Gerkey, B. and M. Mataric, "Sold!: Auction methods for multirobot coordination," *IEEE Transactions on Robotics and Automation*, vol. 18, no. 5, pp. 758–768, 2002.

[45] Bertsekas, D. P. and D. A. Castanon, "Parallel synchronous and asynchronous implementations of the auction algorithm," *Parallel Computing*, vol. 17, pp. 707–732, 1991.

[46] Bellingham, J., M. Tillerson, A. Richard, and J. How, "Multi-task assignment and path planning for cooperating uavs," in *Conference on Cooperative Control and Optimization*, (Gainesville, Florida, USA), November 2001.

[47] Ousingsawat, J., "Quasi-decentralized task assignment for multiple UAV coordination," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, (Keystone, Colorado, USA), 2006.

[48] Xu, L. and U. Özgüner, "Battle management for unmanned aerial vehicles," in *Proceedings of the 42nd IEEE Conference on Decision and Control*, (Maui, Hawaii, USA), pp. 3585–3590, December 2003.

[49] McLurkin, J. and D. Yamins, "Dynamic task assignment in robot swarms," in *Proceedings of Robotics: Science and Systems Conference*, (Cambridge, Massachusetts, USA), 2005.

[50] Flint, M., T. Khobanova, and M. Curry, "Decentralized control using global optimization," in *Proceedings of the AIAA Infotech Conference and Exhibit*, (Rohnert Park, California, USA), 2007.

[51] Chandler, P. and M. Pachter, "Hierarchial control for autonomous teams," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, (Toronto, Ontario, Canada), 2001.

[52] Smith, P., E. McCoy, and C. Layton, "Brittleness in the design of cooperative problem-solving systems: The effects on user performance," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 27, pp. 360–371, 1997.

[53] Lesh, N., C. Rich, and C. Sidner, "Using plan recognition in human-computer collaboration," in *Proceedings of the Seventh International Conference on User Modeling*, (Banff, Alberta, Canada), June 1999.

[54] Cummings, M. L., S. Bruni, S. Mercier, and P. Mitchell, "Automation architecture for single operator, multiple UAV command and control," *The International Command and Control Journal*, vol. 1, no. 2, 2007.

[55] Cummings, M. L. and S. Bruni, "Collaborative Human-Computer Decision Making in Network Centric Warfare," in *TTCP HUM TP7 Workshop on Aerospace Human Factors Issues in Network-Centric Warfare*, (Salisbury, UK), 2005.

[56] Miller, C., H. Funk, P. Wu, J. Meisner, and M. Chapman, "The playbook approach to adaptive automation," in *Human Factors and Ergonomics Society Annual Meeting Proceedings*, vol. 49, (Orlando, Florida, USA), pp. 15–19, Human Factors and Ergonomics Society, 2005.

[57] Crandall, J., M. L. Cummings, and C. Nehmez, "A predictive model for human-unmanned vehicle systems," in *MIT Humans and Automation Laboratory*, (Cambridge, Massachusetts, USA), (HAL2008-05). 2008.

[58] U.S. Department of Defense, "Interface standard, common warfighting symbology," *MIL-STD-2525B*, 1999.

[59] Fisher, S., "Replan Understanding for Heterogenous Unmanned Vehicle Teams," Master's thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 2008.

[60] Cummings, M. L., A. Clare, and C. Hart, "The role of human-automation consensus in multiple unmanned vehicle scheduling," *Human Factors and Ergonomics Society*, vol. 52, no. 2, 2010.

[61] Lewis, T. G., *Network Science: Theory and Practice*, Chapter 2, p. 24. Hoboken, New Jersey, USA: John Wiley and Sons, Illustrated ed., Mar. 2009.