



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

THESIS

**A DECISION PROCESS FOR SURFACE MEDICAL
EVACUATION ROUTING UNDER ADVERSARY THREAT
AND UNCERTAIN DEMAND USING ONLINE
OPTIMIZATION**

by

Kenneth M. Marler

September 2022

Thesis Advisor:

Second Reader:

Ruriko Yoshida

Chrysafis Vogiatzis,

University of Illinois at Urbana-Champaign

Approved for public release. Distribution is unlimited.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, 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, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC, 20503.			
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE September 2022	3. REPORT TYPE AND DATES COVERED Master's thesis	
4. TITLE AND SUBTITLE A DECISION PROCESS FOR SURFACE MEDICAL EVACUATION ROUTING UNDER ADVERSARY THREAT AND UNCERTAIN DEMAND USING ONLINE OPTIMIZATION		5. FUNDING NUMBERS	
6. AUTHOR(S) Kenneth M. Marler			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A		10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.		12b. DISTRIBUTION CODE A	
13. ABSTRACT (maximum 200 words) <p>Emerging threats have focused U.S. Navy operating concepts on agile, distributed tactical forces in the littoral and maritime zones. Given the nature of the threat and its location, medical evacuation via air may be infeasible due to hostile conditions or distance, requiring a shift to a surface or subsurface strategy. Medical demand and adversary actions are unpredictable in warfare, therefore a decision process for routing that accounts for uncertainty is required. Using the principles of the U.S. Marine Corps Rapid Response Planning Process and online optimization, we propose a decision process for surface medical evacuation routing against an adversary given uncertain demand that can be applied to manned and autonomous transport operations. We showcase the computational tractability of the decision process by developing an algorithm to route a medical transport through a network then implement the algorithm as a simulation model in Python. The base case of the model is compared to two modified cases under perfect information to discuss the risks of modeling with inappropriate assumptions. Multiple runs of the simulation model are then used to propose a process to develop a distance multiplier to estimate the impact of adversary presence in existing simulation models without a complete re-design.</p>			
14. SUBJECT TERMS distributed maritime operations, medical evacuation, unmanned		15. NUMBER OF PAGES 71	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

**A DECISION PROCESS FOR SURFACE MEDICAL EVACUATION ROUTING
UNDER ADVERSARY THREAT AND UNCERTAIN DEMAND USING ONLINE
OPTIMIZATION**

Kenneth M. Marler
Lieutenant, United States Navy
BSBA, Hawaii Pacific University, 2012
MBA, University of Scranton, 2014
DHA, Central Michigan University, 2022

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
September 2022**

Approved by: Ruriko Yoshida
Advisor

Chrysafis Vogiatzis
Second Reader

W. Matthew Carlyle
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Emerging threats have focused U.S. Navy operating concepts on agile, distributed tactical forces in the littoral and maritime zones. Given the nature of the threat and its location, medical evacuation via air may be infeasible due to hostile conditions or distance, requiring a shift to a surface or subsurface strategy. Medical demand and adversary actions are unpredictable in warfare, therefore a decision process for routing that accounts for uncertainty is required. Using the principles of the U.S. Marine Corps Rapid Response Planning Process and online optimization, we propose a decision process for surface medical evacuation routing against an adversary given uncertain demand that can be applied to manned and autonomous transport operations. We showcase the computational tractability of the decision process by developing an algorithm to route a medical transport through a network then implement the algorithm as a simulation model in Python. The base case of the model is compared to two modified cases under perfect information to discuss the risks of modeling with inappropriate assumptions. Multiple runs of the simulation model are then used to propose a process to develop a distance multiplier to estimate the impact of adversary presence in existing simulation models without a complete re-design.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

1	Introduction	1
2	Background	5
2.1	Challenges in Wartime Medical Routing	5
2.2	Operational Vignette	7
2.3	U.S. Marine Corps Rapid Response Planning Process	8
2.4	Network Optimization	10
3	Proposed Decision Process	13
3.1	Process Overview	13
3.2	Proof of Concept	18
4	Simulation Results and Discussion	33
4.1	Proof of Concept Simulation Results	33
4.2	Distance Multiplier	37
4.3	Discussion	42
4.4	Medical Considerations for Autonomous Systems	43
5	Conclusion and Future Research	45
5.1	Conclusion.	45
5.2	Future Research.	45
	List of References	49
	Initial Distribution List	51

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

Figure 1.1	Expeditionary Fast Transport.	2
Figure 2.1	U.S. Marine Corps (USMC) deliberate and rapid response planning process comparison.	9
Figure 3.1	Visual depiction of node transitions in the simulation model. . . .	20
Figure 3.2	Process flow for proof of concept simulation execution.	27
Figure 4.1	Visualization of baseline proof of concept simulation results. . . .	34
Figure 4.2	Visualization of proof of concept simulation modification (1) results.	35
Figure 4.3	Visualization of proof of concept simulation modification (2) results.	36

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

CAT	Crisis Action Team
COA	Course of Action
DOD	Department of Defense
NM	nautical mile
R2P2	Rapid Response Planning Process
SAG	Surface Action Group
T-EPF	Expeditionary Fast Transport
USMC	U.S. Marine Corps
USN	U.S. Navy

THIS PAGE INTENTIONALLY LEFT BLANK

Executive Summary

New Naval operating concepts are focused on distributed, smaller, lighter, and more agile tactical forces operating in the littorals and the maritime. This change necessitates a shift in medical evacuation strategy at-sea. In future operating environments, evacuation by air may be infeasible due to hostile conditions or proximity to land. Technological changes, including unmanned and autonomous supply capabilities, change the dynamic for patient transport within wartime logistics and medical networks. Collectively, these challenges limit the response options for the commander in scenarios where time is critical. As surface medical evacuation capabilities emerge to meet the challenges of new operating concepts, routing strategies tailored to the new environment will be more important than ever to meet mission requirements.

Emergency medical demand is, by nature, unpredictable. The dynamic wartime environment and uncertainty of adversary actions add complexity. In this thesis, we explore the challenges of routing against uncertainty and propose a decision process, based on online optimization, to develop an algorithm to route maritime medical evacuation capabilities. This research is a continuation of previous research, which identified online optimization as a better strategy than Deep Q -Learning when prize collecting against a random adversary. The improved performance of online optimization from this study serves as a foundation for the strategic approach in this thesis. The direct application of the methodology in the previous study, which uses a specific application of online learning, does not appropriately handle the evolving conditions found in the problem context we address, hence the proposed decision process. The overarching goal is to further inform the development of Naval strategies for logistics routing, under which medical transportation resides. To provide context and a point of common reference, we use a notional operational vignette throughout this thesis, summarized below:

The transport, under the theater logistics task force, is assigned a 350×350 nautical mile operating area with a waypoint every 10 nautical miles. The operating area contains friendly and adversary units, including expeditionary advanced bases. The mission is to route through the operating area to retrieve casualties and transport them to the hospital. The fuel capacity red line is 1200 nautical miles and the adversary will not actively pursue the transport unless

it comes within 50 nautical miles.

We begin the study by exploring the challenges of routing against uncertainty, resulting in the proposed decision process. Broadly, the proposed process for conceptualizing logistics routing under uncertain demand and adversary movement consists of the following five steps:

- Step 1. Obtain the tactical picture and mission data.
- Step 2. Develop a routing strategy.
- Step 3. Evaluate the routing strategy against decision criteria.
- Step 4. Route to the next waypoint.
- Step 5. Repeat until success or failure criteria met.

We discuss the proposed process using the U.S. Marine Corps Rapid Response Planning Process to promote understanding and confidence in human-machine teaming. To ground the research in existing theory, we also compare the proposed process to the online optimization process used in the previous research. To show the proposed process is computationally tractable, we develop an algorithm as a proof of concept. We implement the algorithm as a simulation model using Python and route a single transport against unknown demand and a single adversary, for simplicity.

The environment we develop is a 35×35 grid to simulate the operating area assigned to the medical transport. The process in the algorithm uses Boolean logic to control the decision flow. We use a single objective routing strategy, greedy on distance to the closest request, to route the transport through the network. We simulate the adversary using a random path to ensure coherent movement within the model. As the adversary traverses the model, a penalty is assigned to five levels of surrounding edges to represent a danger area, which affects the computation of the transport's shortest path. Three different variations of the environment are used to compare the adversary base model results with two modifications where the adversary is not present and information is perfect. The first modified environment routes the transport on the shortest path against the requests that are fulfilled in the base model. The results show the difference in distance traveled by the transport, given the quality of information. The second modified environment routes the transport on the shortest path against the list of all the requests generated in the simulation run. The results show the difference in the number of requests fulfilled given the quality of information. We present

the results of a single simulation run in the context of the operational problem to show an example of the interpretation and the importance of adversary modeling in warfare simulations.

Current medical modeling strategies, such as the Joint Medical Planning Tool, use either straight line Euclidean or Great-Circle distance, which does not appropriately model the increased time required or distance traveled by transport in a contested environment due to maneuver. This presents a risk when planning against the results of the simulation, either in the underestimation of distance traveled or the overestimation of requests completed. To offset risk, we propose a data farming approach using the proof of concept algorithm to develop a distance multiplier. This multiplier could be used in simulation modeling to estimate the presence of an adversary without a complete redesign.

While the simulation model results in this thesis are not usable for military planning, they showcase the computational tractability of the proposed decision process. We see our process and algorithm as a starting point to develop implementable autonomous system algorithms for use in wartime logistics routing that are understandable to the commander. Further, the structure of the grid environment offers a base to explore communication strategies for these systems in denied environments. Finally, the distance multiplier proposed requires refinement but represents a cost-effective interim solution to provide a better estimate of adversary impact until improvements can be made in existing simulation models.

THIS PAGE INTENTIONALLY LEFT BLANK

Acknowledgments

First, a heartfelt thank you to my wife DiAnna and my daughter Kara for their unyielding support. I could not have accomplished what I have thus far without you.

To Dr. Ruriko Yoshida, Dr. Chrysafis Vogiatzis, and CAPT (Ret) Jeff Kline, thank you for the opportunity to pursue the research that I have been passionate about. Your time, the opportunities you provided, and the mentorship over the last two years has been invaluable. It is my hope that the results we have achieved in this academic endeavor and the organizations we have informed will contribute to casualty survival at-sea.

Finally, I would like to thank Dr. Bonnie Johnson and the staff at the Naval Engineers Journal for the opportunity to contribute to the conversation on the Future Hybrid Fleet.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:

Introduction

The economic rise of China and a more aggressive Russia have prompted a shift in U.S. Government focus from the Small Wars of the 2000s to Great Power Competition (White House 2017, 2021). As a result, new Naval operating concepts are focused on distributed, smaller, lighter, and more agile tactical forces operating in the littorals and the maritime. This change necessitates a shift in medical evacuation strategy at-sea. In future operating environments, evacuation by air may be infeasible due to hostile conditions or proximity to land. In maritime mass casualty scenarios, aviation may not have the capacity to assist due to the volume of casualties, fouled decks, or number of Sailors in the water. Technological changes, such as air, surface, and subsurface unmanned and autonomous supply capabilities, change the dynamic for both urgent and routine patient transport within wartime logistics and medical networks due to the design of these systems. Collectively, these challenges limit the response options for the commander in scenarios where time and capacity are critical. The advent of the Expeditionary Fast Transport (T-EPF) Flight II with medical mission capability (Figure 1.1) is an example of the change in approach needed to provide operational commanders with options to mitigate medical risk, both in transportation and proximity to surgical intervention, however, employment strategies for these capabilities are still in their infancy (Gillingham et al. 2022).

As surface medical evacuation capabilities emerge to meet the challenges of new operating concepts, routing strategies tailored to the new environment will be more important than ever to meet mission requirements. Emergency medical demand is, by nature, unpredictable and made more complex by the dynamic wartime environment. Adversary actions, while anticipated to a certain extent, remain largely uncertain. In this thesis, we explore the challenges of routing against uncertainty and propose a decision process, based on online optimization, to develop an algorithm to route maritime medical evacuation capabilities given unknown demand, perishable cargo, and the presence of an adversary (Belmega et al. 2018). We conceptualize the proposed decision process using the U.S. Marine Corps (USMC) Rapid Response Planning Process (R2P2) (U.S. Marine Corps 2022, App H). The goal is to further inform the development of Naval strategies for logistics routing, under

which medical transportation resides. This research is a continuation of Liu et al. (2022), which identified online optimization as a better strategy than Deep Q -Learning when prize collecting against a random adversary. The improved performance of online optimization given a random adversary serves as a foundation of the strategic approach in this thesis. The direct application of the methodology proposed in Liu et al. (2022), which uses a more specific application of online learning, does not appropriately handle the evolving conditions found in the problem context we address.



Figure 1.1. Expeditionary Fast Transport (T-EPF). Source: Aistrup (2021); The appearance of Department of Defense (DOD) visual information does not imply or constitute DOD endorsement.

We begin with a discussion of the challenges associated with medical routing in the wartime environment. We use an operational vignette as a point of common reference throughout the thesis and to provide context for the implementation of the proof of concept algorithm. We present the proposed decision process and compare the components to R2P2 to promote confidence in human-machine teaming. To ground the research in existing theory, we

compare the proposed decision process to the online optimization algorithm proposed by Belmega et al. (2018). Next, we develop an algorithm as a proof of concept to show computational tractability. We implement the algorithm as a simulation model using Python and route a single transport against unknown demand and a single adversary, for simplicity.

The environment we develop is a 35×35 grid to simulate the operating area assigned to the medical transport. The process in the algorithm uses Boolean logic to control the decision flow. We use a single objective routing strategy, greedy on distance to the closest request, to route the transport through the network. We simulate the adversary using a random path to ensure coherent movement within the model. As the adversary traverses the model, a penalty is assigned to five levels of surrounding edges to represent a danger area, which affects the computation of the transport's shortest path. Three different variations of the environment are used to compare the adversary base model results with two modifications where the adversary is not present. This comparison shows the difference between perfect information and uncertainty in the model. The first modified environment routes the transport on the shortest path against the requests that are fulfilled in the base model. The results show the difference in distance traveled by the transport. The second modified environment routes the transport on the shortest path against the list of all the requests generated in the simulation run. The results show the difference in the number of requests fulfilled. We present the results of a single simulation run in the context of the operational problem to show an example of the interpretation and the importance of adversary modeling in warfare simulations.

Current medical modeling strategies, such as the Joint Medical Planning Tool, use either straight line Euclidean or Great-Circle distance which does not appropriately model the increased time required or distance traveled by transport in a contested environment due to maneuver. This presents a risk when planning against the results of the simulation, either in the underestimation of distance traveled or the overestimation of requests completed. To offset risk, we use the results of multiple simulation runs to propose a process for the development of a distance multiplier. This multiplier could be used in simulation modeling to estimate the presence of an adversary without a complete redesign. Finally, we discuss the use of unmanned and autonomous systems in the logistics network and the implications for medical evacuation routing.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 2: Background

In this chapter, we provide the context for the problem addressed in this thesis and background on the military decision process and computational approach used to develop the proposed decision process. This section is not designed to give a comprehensive review of the topics, but to provide a general definition of the approach and its application in the context of this thesis. The military decision process we use in this thesis is R2P2 to provide a common language to discuss the approach the proposed decision process takes to understanding and developing solutions to the problem. The primary computational approach we use is online optimization to ground our research in established theory.

2.1 Challenges in Wartime Medical Routing

The primary challenges in medical evacuation routing under wartime conditions can be summarized as: (i) unknown demand locations and quantities, (ii) time perishable cargo, and (iii) an unpredictable and dynamic threat environment.

2.1.1 Unknown Demand

Medical demand is, by nature, unpredictable. While observable events can signal the potential for demand, the uncertainty of warfare and dangers of operating in the at-sea industrial environment present special circumstances that increase the uncertainty of demand. Traditional vehicle routing problems rely on known demand locations and quantities to route against. Dynamic routing problems, such as ambulance routing, typically benefit from historical geographic trends and population densities. This information can enable the use of machine learning algorithms and advanced positioning of transports. At-sea medical networks, especially under wartime conditions, are continuously evolving. Therefore they are a memoryless process that does not benefit from historical information. Simply, routing on events that may signal demand will miss equally catastrophic medical circumstances that occur by chance and without an observable event from the transport's perspective.

2.1.2 Time-Perishable Cargo

Time perishable cargo refers to the medical requests in the network. The mechanism and severity of injury can make the medical request time perishable. There are five categories of evacuation precedence within the military medical system (Joint Chiefs of Staff 2018, pp. A10-A11). Of the five categories of evacuation precedence, this we focus on Priority I-Urgent and Priority IA-Urgent-Surgical evacuations. In both of these categories, the patient must be evacuated within a certain period to optimize their outcome. This period can be adjusted by positioning additional medical capability onboard the transport or at the location of injury. Special considerations, such as sea-state, last mission assignment, and available onboard facilities will impact the ability to position additional medical capability onboard the transport. Finite medical capabilities in theater will affect the presence and level of medical capability at the point of injury. The patient timeline may change in transit presenting a constraint that should be considered in the context of the operational situation. Medical operations are a risk mitigation measure for the commander. As such the considerations for time-based evacuation may be superseded by mission necessity, as dictated by the evolution of the operational picture.

2.1.3 Unpredictable Dynamic Threat Environment

The unpredictable threat environment refers to the inability to accurately predict adversary actions and the interventions of weather on military operations. The dynamic threat environment refers to the ever changing environment where the threat evolves including periods of both lessened and heightened threat conditions due to the mobility of maritime and littoral threats. Enemy actions can be anticipated to some degree through the use of intelligence operations. Adversary deception and counter-intelligence operations however can impact the ability to accurately predict enemy actions. Weather, while predictable to a certain degree, may still have unpredictable impacts on the execution of military operations and the functionality of certain equipment or capabilities. For example, the T-EPF has sea-state limitations that impact the ability to operate safely. In high sea-states, surgery at-sea is difficult or impossible, negating the benefits of forward surgical care. In this situation, the commander will need to balance the operational need for force preservation against the medical need for evacuation. On one end is maximum preservation of the transport, which results in no medical evacuations with the possibility of threat. On the other end, is high

risk to the transport as it braves adversary fire or poor sea-states to retrieve the casualty within the time window. For these and other reasons, traditional vehicle routing problems are insufficient to adequately model the considerations for routing against uncertainty in this context.

2.2 Operational Vignette

To conceptualize the analytical approach, it is important to understanding how the operational environment is structured. The operational vignette we developed is notional and provides an opportunity for us to present the context of the research using a common mental construct with the reader. Military officers well versed in the operational employment of naval forces may approach operating area assignment and force employment differently. Consider the following notional vignette:

The Joint Force Maritime Component Commander is tasked with conducting operations in a distributed maritime environment against an adversary in support of a larger campaign. Sea control has not yet been achieved, but both belligerents are successfully conducting sea denial in the operating area. Air operations in support of medical evacuation are not practical due to the threat environment, distance to the theater hospital, and the demand for aviation in support of combat operations. The maritime logistics task force has operational control over the theater-level surface medical transports. Tactical control has been delegated to requesting units during casualty transfer only. The capability assigned to the mission is a T-EPF Flight II Medical Variant equipped with limited capacity Role II surgical care (Joint Chiefs of Staff 2018). In this future scenario, the T-EPF is capable of manned or autonomous navigation; however, it is always manned with medical and support personnel during medical missions. Given the fuel limitations, the T-EPF has been assigned a 122,500 square nautical mile (NM) (350 x 350 NM) operating area with both adversary and friendly forces present. The location of friendly and adversary forces is unpredictable, requiring a high level of agility to respond to requests. The primary mission assigned to the T-EPF is to respond to requests for medical evacuation while avoiding interaction with enemy forces and staying over the fuel red-line. The designated refuel location is collocated with the hospital. Friendly forces in the area include surface ships and expeditionary advanced bases on small islands. Due to communication jamming and the need for low electronic emissions, the tactical picture is updated by onboard sensors and passively received burst transmissions

at certain waypoints. The transmissions update the location of the adversary and initiate new mission orders, such as evacuation requests. The collection of mission orders received is stored as a mission set. The transmission data is deconflicted with onboard sensors and prior information on the mission set and environment. Based on the data available and the requirements of the mission, a decision must be made on how the mission will be executed.

2.3 U.S. Marine Corps Rapid Response Planning Process

In data science, mathematical models are often viewed as a black box by those unfamiliar with the computational process (Petch et al. 2022; Lawless et al. 2019). This is a challenge for commanders that rely on these black box processes to make decisions. To enable transparency, the decision process proposed in this thesis was consciously constructed based on known decision algorithms to convey the structure in a manner familiar to military commanders. Given the mathematical nature of our approach to the operational problem, grounding in an established optimization process is also important to ensure the process is computationally tractable and can be implemented. Using this model of development, the developers and operators can understand how the proposed decision process approaches the operational problem presented.

The proposed decision process generally follows the standard military planning process; however, the implementation more closely mirrors the time-condensed planning process used by the USMC. The Rapid Response Planning Process, noted as R2P2 going forward, is a six step process used for time-constrained operations (U.S. Marine Corps 2022, App H, p. 108). The steps of the planning process are:

- Step 1. Problem Framing.
- Step 2. Course of Action (COA) Development.
- Step 3. COA Wargaming.
- Step 4. COA Comparison and Decision.
- Step 5. Orders Development.
- Step 6. Transition.

While the process mirrors the deliberate USMC Planning Process, in practice, the COA Wargaming and COA Comparison and Decision steps are condensed into a single step.

Figure 2.1 shows the comparison of the USMC Planning Process and R2P2. The steps of the planning process will be referenced in greater detail in the explanation of the proposed decision process in Chapter 3.

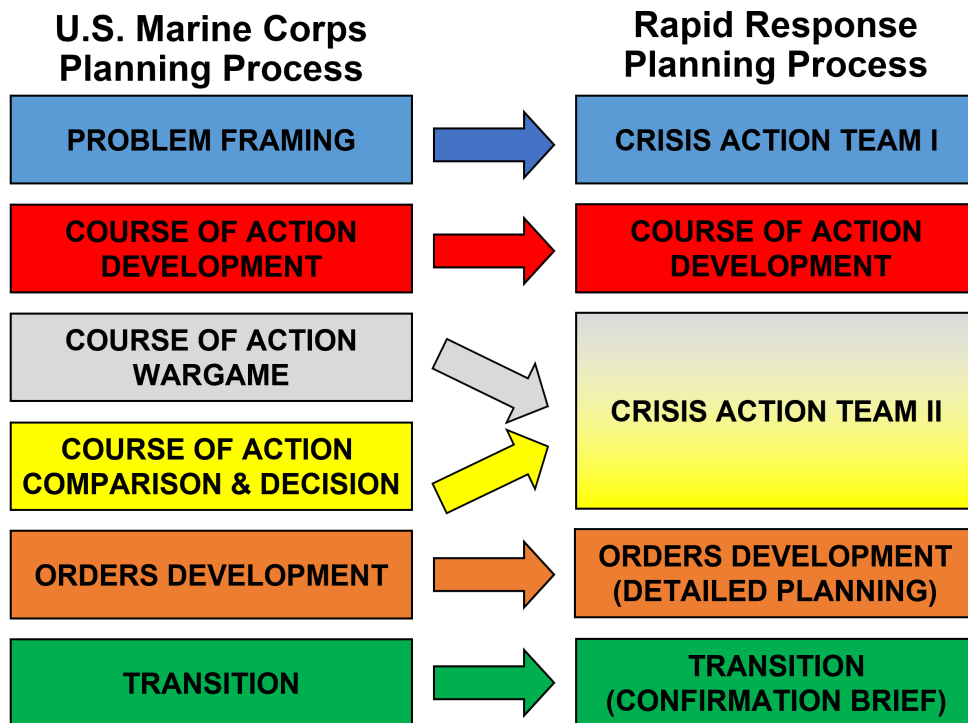


Figure 2.1. USMC deliberate and rapid response planning process comparison. Source: U.S. Marine Corps (2022).

The planning cell for R2P2 is known as the Crisis Action Team (CAT), an integrated multi-disciplined group consisting of both USMC and U.S. Navy (USN) operations and support personnel. The CAT has a core group of planners, with a listing of ancillary specialties that are included based on the mission and the commander’s discretion. The composition of the CAT is designed to give the commander the right mix of specialties, based on the understanding of the problem, to develop a COA that appropriately produces the desired effect to achieve the objective. To accelerate the planning process, the CAT operates based on pre-planned capabilities rather than novel development.

2.4 Network Optimization

Network optimization is the practice of applying network flows algorithms to a problem where the goal is to move some entity from one point to another in an underlying network as efficiently as possible (Ahuja et al. 1993, p. 1). This approach has many applications from modeling the flow of electricity to routing delivery vehicles. For this thesis, network optimization, specifically online network optimization, serves as the overarching analytical approach to move the medical transport through the underlying at-sea network.

2.4.1 Traditional Vehicle Routing

The vehicle routing problem is a combinatorial optimization problem which is a generalization of the traveling salesperson problem (Ahuja et al. 1993, pp. 623-625). The problem is concerned with routing vehicles through a network based on known costs and flow capacities with the objective of finding the optimal set of vehicle routes. The generalized approach, noted initially as the truck dispatching problem, was introduced by Dantzig and Ramser (1959) in the context of routing a fleet of fuel trucks from bulk locations to service stations.

On the surface, the routing of medical vehicles could be done within the context of the vehicle routing problem. There are many adjustments to the algorithm that have been developed to account for capacity, time windows, and penalties on missed pick-ups or deliveries (Google Developers 2022). The basic structure of these algorithms assumes the pick-up locations, drop-off locations, and quantities are known in advance to optimally route the vehicles through the network. In a maritime scenario under wartime conditions, medical demand is unknown in advance and subject to uncertainty. During exploratory analysis and initial modeling for this thesis, we attempted multiple variation of the vehicle routing problem, however, the traditionally structured vehicle routing problem lacked the ability to account for uncertainty in demand and adversary movements. Since the incomplete information hindered the development of an optimal route, another layer was considered necessary above the vehicle routing algorithm to account for the changes in information within the model.

2.4.2 Online Optimization

Online optimization is concerned with optimization problems that have incomplete knowledge of the future environment (Hazan 2021, p.169). As the events progress and information is obtained, the problem is re-evaluated in the context of the new environment. As a control layer over the vehicle routing problem, we found the general online optimization approach provided a better process than the vehicle routing problem alone. In a general sense, online optimization consists of the following four step process (Belmega et al. 2018, p. 13):

Step 1. Select an action.

Step 2. Incur loss.

Step 3. Update the action.

Step 4. Repeat until criteria met or end of loop.

The online decision process proposed by Belmega et al. (2018) more closely follows a trial and error online learning approach to online optimization. This is where the decision process we propose in this thesis and the online decision process proposed by Belmega et al. (2018) diverge. The approach to online optimization we use for the proposed decision process is different from online learning, where the entity observes and stores previous actions for later consideration. While there is the possibility to extend the research in this thesis to multiple objective optimization that may benefit from machine learning, the structure for this decision process synthesizes online optimization concepts and Boolean logic, not online learning. The discounting of previous actions is purposefully done since new information and an evolving environment negate the benefits of memory resulting in, simply, a new problem being solved every time.

2.4.3 Online Optimization Under Uncertainty

In a maritime environment under wartime conditions, the demand for medical evacuation and the movements of the adversary will not be precisely known in advance. Given this challenge, the vehicle routing problem alone is a poor fit for this context, however online optimization provides a method to update the vehicle routing problem as information becomes available. Since historical information provides little benefit in the dynamic environment of maritime conflict, online machine learning is a poor fit for single objective optimization in this context. The computational structure for this thesis therefore is an online optimized vehicle routing

problem using Boolean logic decision path controls as the overarching analytical approach to develop the path as the transport moves through the network.

CHAPTER 3: Proposed Decision Process

In this chapter we describe the proposed decision process for medical evacuation in the maritime given uncertainty in demand and adversary movement. We begin with a description of the process and compare the components to the steps in R2P2. This allows us to discuss the proposed process using an existing decision process familiar to operational commanders. Next, we present an algorithm developed using the decision process as a proof of concept to showcase the computational tractability. The algorithm, implemented in Python, is designed to simulate medical evacuation under wartime conditions in a distributed and contested maritime environment. We discuss the interpretation of the results of the proof of concept and the application of multiple simulation runs to the development of a distance multiplier for estimating the impacts of adversary presence in Chapter 4.

3.1 Process Overview

Broadly, the process we propose in this thesis to route under uncertainty consists of the following five steps:

- Step 1. Obtain the tactical picture and mission data.
- Step 2. Develop a routing strategy.
- Step 3. Evaluate the routing strategy against decision criteria.
- Step 4. Route to the next waypoint.
- Step 5. Repeat until success or failure criteria met.

We describe each of these steps in detail in the sub-sections that follow. As identified in Liu et al. (2022), machine learning was not well suited when routing against a random adversary. This incompatibility is compounded by uncertain demand, especially in cases where historical actions are not representative of the future tactical picture. For this reason, the decision algorithm we showcase as a proof of concept relies on Boolean logic to decide how to proceed through the network. Many military systems use similar logic over artificial intelligence or machine learning, both for interpretability and the incompatibility of these approaches with uncertainty. “Essentially all models are wrong, but some are

useful” (Wasserstein 2010, p. 1). The algorithm we propose for the proof of concept does not purport to make all the appropriate assumptions for a real world scenario. What we do show is how an algorithm might be constructed using the proposed decision process to model or, at a high conceptual level, program, a transport to move through a network given random demand and unpredictable adversary movement.

3.1.1 Obtain the Tactical Picture

The first step is to obtain the tactical picture and receive external data to understand the operating environment. Marler (2022) noted the first of the progressive key factors for medical planning in the maritime is understanding the location in depth prior to the implications of possible events. This can be translated here to first understanding in detail the operating area from the friendly, adversary, spatial, and environmental perspectives. Pre-existing data is compared with new data obtained from onboard sensors to update the tactical picture of the environment. Inertial or satellite navigation systems update the position of the transport in the operating area. Depending on the context, onboard logistics, medical, and support teams update the tracking systems with time sensitive data relevant to routing, which may include approximated patient survival times or request expiration times. External burst transmissions are received with new missions and over the horizon friendly and adversary data, which are then deconflicted with current onboard information. At the completion of this step, the transport has the tactical picture and parameters that will be used as inputs to develop the routing strategy. Until success or failure at Step 5, this step is repeated to update the tactical picture at each iteration of the loop.

The criteria for the update loop can be set in multiple ways by the decision maker. Some possible update criteria could be upon receiving a mission, proximity to an adversary, or at certain waypoints. This decision affects the computational complexity of the problem where continuous updates and the choice of waypoint resolution affect the speed of the algorithm’s execution. High resolution produces the ability to optimize at closer intervals, but may take longer to produce a result, where lower resolution produces faster results at longer intervals, but present higher risk of adversary interaction and shorter time windows to fulfill requests. Operationally, waypoint resolution and re-routing intervals can impact the exposure of the transport to risk of adversary interaction or the time to route to a request. Therefore, the decision criteria for this step relies on the availability of current information

on the operating environment and the time interval between process execution loops.

Translation to R2P2

This step is comparable to the CAT I step at the beginning of R2P2 depicted in Figure 2.1. During CAT I, problem framing is conducted based on the mission from higher headquarters. The necessary background information and data are gathered by the planning team to understand the context for mission execution. Each discipline on the planning team provides their initial staff estimates. Initiating guidance is provided by the commander which functions as the criteria to develop and test the COAs. At the conclusion of CAT I, the planning team has the necessary information to develop the COAs for comparison.

3.1.2 Develop A Routing Strategy

Once the tactical picture and mission data are updated, a routing strategy is developed to the next destination in the network based on the current mission set. The mission set includes the set of unexpired requests that the transport could fulfill during the time step. Under the research assumptions and using a single objective optimization strategy, we found a Boolean logic-based decision path for developing the routing strategy best accommodated the limitations of an uncertain environment. The criteria for routing are evaluated in sequence, then an existing algorithm is used to develop a path to the next destination in the network based on the selected destination identified by analysis of the decision criteria. One example of an algorithm for optimal path routing used in the proof-of-concept later in this chapter is Dijkstra's for single source shortest path (Dijkstra 1959; Ahuja et al. 1993). This is not the only possible algorithm, but is sufficient for the purpose of this thesis. Using multiple objective optimization, the algorithm for routing can be greedy on multiple aspects of the mission. Although not considered in depth here, a weighted multiple objective algorithm may be used to consider alternative routing options. As a baseline, the greedy algorithm for our single objective model is based on distance, however, this could be expanded to include other decision criteria, such as risk or survival times, in addition to distance.

The output of this step is an optimal route given the input criteria. In many cases, there may be additional decision criteria that must be considered. The purpose of the next step is to take the routing strategy developed in this step and compare it to the additional decision

criteria. If the routing strategy fails against the decision criteria, this step is repeated and a new strategy is developed and tested.

Translation to R2P2

This step is comparable to the COA Development step in R2P2 in Figure 2.1. The planning team takes the information gathered from the commander, higher headquarters, and parallel and subordinate units then develops COAs to meet the objective. This varies from the deliberate planning of the USMC Planning Process since the initial source of capabilities considered are pre-planned responses, such as standing mission packages. This is similar to the Boolean logic which we use as a control since the development of the routing strategy relies on a pre-defined decision path rather than a novel approach. At the conclusion of this step, the planning team has COAs developed for wargaming, comparison, and decision during CAT II.

3.1.3 Evaluate the Routing Strategy

After the routing strategy is developed, a comparison is made against the decision criteria. The decision criteria can be qualitative factors the decision maker considers important based on a pre-developed mental schema. They can also be quantitative factors such as weighted greedy aspects of the algorithm. Additional criteria may include patient survival times, total requests within a certain vicinity, or risk on the route. For single objective optimization, such as the one we use in the proof of concept model later in this chapter, the greedy aspect is the distance to the closest request. Given the next request is unknown and may appear along the path to the current request, this criterion allows the transport to divert and pick up the new request on the way. In multiple objective routing, the decision criteria are weighted based on their importance. The decision maker may gain additional control of the quantitative decision criteria by using a goal programming approach instead of a simple weighted approach (Royset 2021, pp. 5-6).

The output of this step is the optimal routing strategy given both the quantitative and, if considered, qualitative decision criteria. Should no strategy meet the decision criteria, the decision process loops back to previous steps to evaluate (a) the basis for the decision criteria in Step 1 or (b) the method used for developing the routing strategy in Step 2. Should

multiple routing strategies meet the decision criteria, they are compared and one is selected to move to Step 4.

Translation to R2P2

This step is comparable to the COA Wargaming, Comparison, and Decision actions that occur within the CAT II step of R2P2, depicted in Figure 2.1. The developed COAs are tested and evaluated against the commander's decision criteria then compared. If a COA meets the criteria, it is selected and refined. If no COAs meet the criteria, then the previous steps are repeated until a decision is made. At the conclusion of this step, the planning staff refine the selected COA and complete the final detailed planning required for execution.

3.1.4 Route to Next Waypoint

Once the routing strategy meets the decision criteria, action is taken toward the next waypoint. The waypoint represents the next opportunity the transport will take to update the route. If the next waypoint is a pick-up location, then the support staff will update the related routing data and the mission set once the destination is reached. If the next waypoint represents the last destination, then the mission is complete upon arrival and the parameters are reset for the transport based on the services available at the destination. This could be offloading patients, cargo, or refueling. During the transit, the transport may run out of fuel or interact with the adversary. Depending on the penalty of the action, this may be handled in different ways. In the case of adversary interaction, the result may be destruction of the transport. In the case of exhausted fuel, the transport may go below the red line requiring intervention from the mission control location.

Translation to R2P2

This step is comparable to the Orders Development and Transition steps of R2P2, depicted in Figure 2.1. Detailed planning has been completed, orders have been developed and issued to executing units, and the confirmation brief is given by the mission commander to the task force commander. At the conclusion of the Transition step, the executing unit under the mission commander is taking actions, in line with the COA selected, towards the objective.

3.1.5 Repeat Until Success or Failure

The final step is to repeat the process until an intervention occurs that reports the success or failure of the mission. In this decision process, repeat implies the algorithm loops back to Step 1, Obtain the tactical picture and mission data. The primary purpose of this step is to complete the circle, however, it acknowledges that there is an action that has previously occurred and the criteria for the next action must be updated. This is the point when updates internal to the transport are made. For example, fuel would have been expended during the previous time step that needs to be tracked and accounted for. At the completion of this step, the transport has one of three outcomes:

- (1) Continued routing of the transport, therefore, the algorithm repeats.
- (2) Destruction, either by exhausting fuel or interaction with the adversary, which results in a mission failure and the transport concludes operations.
- (3) Arrival at a final destination, which is considered mission success and the transport is made ready for future operations by resetting the parameters, such as refueling and offloading.

Translation to R2P2

This step is comparable to the operational assessment and feedback loop that occurs within R2P2. At some point during the execution of the mission, an unexpected event may occur that triggers another iteration of R2P2. The difference here is the waypoint represents the pre-planned execution of the process until a conclusion is reached, where R2P2 is repeated based on events and the tactical or operational picture. This could be replicated by adjusting the re-evaluation criteria of the algorithm within this step. Instead of re-evaluation at a specific waypoint, the adjusted criteria may include a specific distance from the adversary, change in expiration time of cargo or a request, or update only upon receiving a new mission.

3.2 Proof of Concept

To test the computational tractability of the decision process we propose in this thesis we designed an algorithm for a simulation model and implemented it using Python. To explore the structure of the algorithm further, we use the networkx package (Hagberg et al. 2008) in Python to visualize the results of the model runs. The model consists

of four parts: Experimental Setup, Class Structure, Functions, and Simulation Execution. The Experimental Setup describes the structure of the grid-world to visually represent the model. The Class Structure uses an object oriented programming approach to consolidate the actions of each object within the model. The objects in the Class Structure include (i) Transport, (ii) Adversary, (iii) Medical Requests, and (iv) Hospital. Functions are external to the Class Structures and are used to develop routes. Finally, Simulation Execution presents the algorithm used to test the flow of the computational implementation of the proposed decision process. Maritime medical evacuation under wartime conditions is the context used for the simulation model, however, other logistics applications could use a similar structure to evaluate the decision process' applicability to other problem sets. Holistically, the model is designed to simulate the movement of a surface medical transport through waypoints in a grid network to pick up casualties from demand locations while avoiding adversary units. The overarching objective is the transfer of maximum number of requests to a hospital location given the fuel constraint. The primary object of focus in the model is the Transport class which has the objective of routing to pick up requests within the network given a discrete number of movements. The other objects are designed to support the movement of the Transport through the network. While extra parameters could be included, such as Adversary class fuel, they are left as an opportunity for future research if the level of resolution is necessary for the problem addressed. Each object will be discussed in further detail in the sub-sections that follow.

3.2.1 Experimental Setup

The experimental environment for this simulation model uses a 35×35 node-based grid world with horizontal, vertical, and diagonal bi-directional arcs. Figure 3.1 provides a visualization of how the objects can move within the grid-world network. In the operational context, the distance traveled across an arc is 10 NMs. From node 1 (top left corner), to each of the other three corners of the 35×35 grid requires the transport to travel 350 NMs. This simulates the idea of an operational box assignment by the commander and allows for an estimation of fuel usage in the model. The use of nodes instead of boxes allows the model to simulate greater maneuverability and provides a point of control given the scenario. Nodes are referred to as waypoints in the simulation model and operational context since they represent the points in the network to optimally route through or to and the points in the model that simulate operational control of the transport.

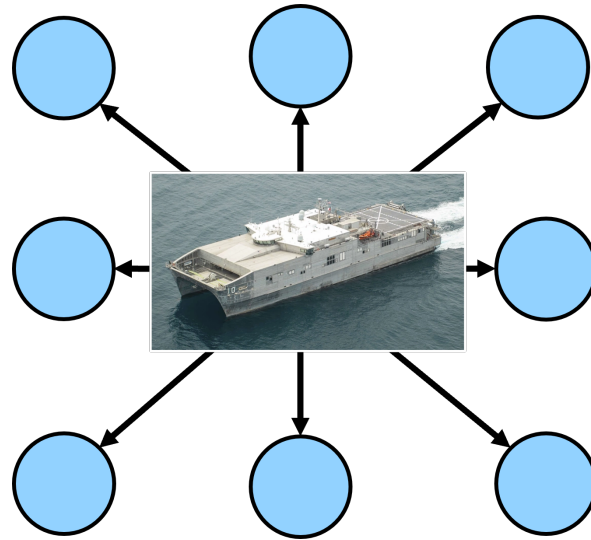


Figure 3.1. Visual depiction of node transitions in the simulation model. Source: Aistrup (2021); The appearance of DOD visual information does not imply or constitute DOD endorsement.

In a wartime scenario, and given the communications infrastructure in some austere locations, it is reasonable to assume that transport operations will occur in a degraded communication environments. For this reason, the waypoints also represent control points where the transport can receive burst transmissions to update the onboard tactical picture or receive new mission orders. The use of the waypoint structure in the maritime context is relevant since ships operate over greater distances and dispersion than ground units so modeling on a 2-D plane can be accomplished. In the ground context, modeling the 3-D characteristics of land is necessary for movement therefore this level of resolution is insufficient to appropriately model the operational area. Given this, the model, as structured, is relevant only to the maritime context. While more detailed modeling may be of greater concern for other maritime issues, to balance model resolution with computational complexity, the waypoint control method, as described, was deemed sufficient for this proof of concept.

Distance Disclaimer

While we discuss distances to provide operational context, we make no assumptions that the grid-world developed for our simulation model can simply be overlaid on a map and executed. Given the distances covered in this simulation model, at a minimum, the

difference between Euclidean and Great-Circle distance would have an impact. Future research would be required should this algorithm be applied to the routing of autonomous or semi-autonomous systems.

3.2.2 Class Structure

Four class structures are used to program the objects in the model which include: Transport, Adversary, Medical Evacuation, and Hospital. The Transport class represents the medical transport and is the primary focus for data collection and simulation design. The Adversary class represents the threat and adds penalties to the arcs to influence the Transport class routing decisions. The Medical Evacuation class represents friendly units requesting evacuation and are added to the mission set for the Transport class. Finally, the Hospital class serves as the destination for the requested evacuations and can be at a fixed node to represent a receiving port or transition between nodes to represent a receiving hospital ship.

Transport

The Transport class is the primary object of focus in this simulation model. The additional objects are designed to interact with the Transport class, thus, the actions of these objects that do not represent an effect on the Transport class are not simulated. An example of this is the fuel capacity or fighting capability of the adversary and friendly units. In the context of the operational problem, the Transport class is a sea-based medical evacuation capability equivalent to a T-EPF Flight II Medical Variant or other small to medium class ship no larger than an Freedom Class LCS. The operational assumptions for this class include:

- (1) The transport can connect to or retrieve patients from the maritime or maritime approachable location for patient transfer. An example of the former is a ship, submarine, or downed aviator, where the latter could be an expeditionary advanced base or littoral port.
- (2) The medical team can stabilize the patients enroute to the hospital. For the purposes of this proof of concept, patient quantities and survival times are not included in the parameters for this class.
- (3) The platform of the medical transport is control agnostic, simply, it can be viewed in either an autonomous or manned control capacity.

Initialization

To initialize this class in the model, a starting location, fuel capacity, adjacency list, penalty dictionary, and the first request location are provided to develop the initial route. The standard starting location is at node 1 (top left corner of the grid), but can be changed prior to starting the run. The fuel capacity is representative of the max distance that the Transport class can move. For example, this study assumes 10 NMs between waypoints which would, at 120 fuel units, equate to 1200 NMs of fuel capacity before the requirement to refuel.

Movement

An adjacency list is used to define where the medical transport can move. An associated cost dictionary defines the penalty of moving along the edges. The term edge is used to describe the bi-directional nature of the arcs in the model. The penalty cost dictionary used for the Transport class is developed by the Adversary class and represents danger areas for the transport, with the ring furthest away having the lowest penalty. The Transport class maintains a path list and a visited list to track where it has been, its current location, and where it will move next.

Update Path

To update the path for the Transport class, the class function requires the current mission set, the adjacency list, penalty dictionary developed by the Adversary class, and the location of the Hospital class. The algorithm for routing the Transport class is greedy on distance to nearest request rather than on the concentration of requests. This is due to the uncertain nature of medical demand, and the routing algorithm having incomplete information at the time of route computation. Details of the path update will be covered in Simulation Execution. The overarching approach to update the path is based on the calculated distance to the request, plus the distance from the request to the hospital. If the transport has sufficient fuel remaining, route to request; else go to the hospital.

Adversary

The Adversary class is primarily used to penalize the route of the Transport class to incite avoidance. In the context of the operational problem, the Adversary class is a Surface Action Group (SAG) consisting of Destroyer and Frigate type ships that have an operational mission to seek and destroy capital ships in the operating area. The operational assumptions for this class include:

- (1) The small, independently operating transport represents a target of opportunity if the transport enters the threat rings.
- (2) The adversary will not actively pursue the transport due to the presence of other, more lucrative targets.
- (3) Adversary submarines will not risk exposing their location to target the transport due to the risk of detection and the availability of other, higher priority targets.

Initialization

At initialization of the simulation, a random starting and ending point are generated for the Adversary class to follow using a random uniform $[0, 1]$ distribution where every node has an equal probability of selection. The shortest path between the two points is computed and becomes the starting path for the Adversary class. In the operational context, this path represents the current mission assigned to the adversary. An initial penalty dictionary is generated from the Adversary class position in the network using a five layer progressive penalty build with the furthest ring having the smallest penalty. This simulates a threat area for the Transport class out to five nodes or 50 NMs from the adversary's current location. The penalty is assigned to the edges connecting the nodes to the Adversary's position, not the nodes themselves.

Movement

The Adversary class follow the same movement pattern as the Transport class, however the movement uses a fixed, non-penalized cost dictionary. As the Adversary class moves through the network, the penalty cost dictionary used by the Transport class is updated to reflect the change in penalty on the edges connected to the Adversary class location and a reset of the edges that are now unaffected.

Update Path

At the end of the generated path, the Adversary class develops a new path using the end point of the old path as the starting point of the next. A new ending point is randomly generated using the same random uniform $[0, 1]$ construct in the initialization step. The shortest path between the two locations is computed and becomes the new path for the adversary. This cycle continues in perpetuity until the end of the simulation. In the operational context, this represents a change in orders for the adversary unit from their higher headquarters and is assumed to be unknown to the medical transport and friendly units in the operating area.

We use this structure to represent a semi-stochastic adversary that follows a deterministic path with random end points. This approach makes more sense than the random appearance of the adversary within the network since friendly units would have some intelligence on the position and movement of a surface based adversary unit. This is contrasted with the possible position rings that are used for submarine tracking, which would be better represented by a threat mapping approach to penalized routing.

Medical Requests

The Medical Requests class is used to generate stochastic demand for the Transport class. In the context of the operational problem, the Medical Requests class consists of the friendly surface ships, submarines, and small island locations in the operating area that have a medical event requiring evacuation to a higher level of care. Due to the nature of medical events at-sea and under wartime conditions, advanced locations, quantities, and concentrations of demand will be unknown. In some situations, known or anticipated events, such as the proximity of friendly and adversary forces can give indications that there may be demand in the future, however, even under these conditions, casualties are not a foregone conclusion. A parallel can be made to the ambulance routing problem, however, this problem has the benefit of historical geographic information and scheduled, predictable mass events that may present a higher risk. The maritime wartime environment does not benefit from this form of historical information, hence, demand remains largely unpredictable. Operating in a distributed, mobile manner further complicates maritime medical evacuation since, to appropriately position the transport, the distance to each requesting force will require a continuous distance optimization problem. This is not operationally feasible since, while not in service, the medical transport will likely be positioned in a location of safety for force preservation, not optimally positioned for response.

Initialization

At initialization of the simulation, the fixed probability of a new request during a time step is defined and an initial request is generated from the set of nodes using a random uniform $[0, 1]$ distribution where every node has an equal probability of selection. This initial request serves to activate the Transport class in the model to begin routing to the first request. This criteria can be refined by establishing a fixed or mobile set of nodes to represent the locations of ships or islands, but this structure was sufficient for this proof of

concept since the simulation model assumes small island less than 10 NMs and disbursed ships and patrol craft within the operating area. The random generation of a location doesn't need to be viewed as fixed. They could also be considered as a rendezvous point, instead waiting for pick-up, since this would be more operationally realistic.

New Request

To generate demand past the initialization step in the model, a check is performed at the beginning of each time step to see if the Medical Requests class initiates a new request. A random uniform $[0, 1]$ is generated then compared to the fixed probability of a new request occurring during a time step. For the proof of concept, the threshold 0.1 is used. If the random uniform $[0, 1]$ generated is less than 0.1, then a new request is generated and a node is selected at random with equal probability; else, the simulation continues without a new request.

Hospital

The Hospital class is the end point or sink in the simulation model. After routing to pick-up a request, the Transport class will make a decision based on remaining fuel and the existence of other requests in the mission set to either route to another request location or to the Hospital class to drop off patients and reset the parameters, effectively ending the simulation. In the context of the operational problem, the hospital location could be fixed, representing a port or land transfer location, or mobile representing a casualty receiving and treatment ship, such as a hospital ship or amphibious landing ship, with embarked medical teams. Per the assumptions in the Transport class, the Hospital class location is assumed to have the structure to enable the connection and transfer of patient from the transport to the hospital. For the purposes of this proof of concept, the hospital is in a fixed location, however, the model can be structured to simulate a mobile location, such as a hospital ship.

3.2.3 Functions

The primary functions outside of the class structure in the simulation model are shortest path and penalty generation.

Shortest Path

The shortest path problem can be easily formulated as:

$$\begin{aligned}
 \min \quad & \sum_{(i,j) \in A} c_{ij} x_{ij} \\
 s.t. \quad & \sum_{j:(i,j) \in A} x_{ij} - \sum_{j:(j,i) \in A} x_{ji} = \begin{cases} 1, & \text{if } i = s, \\ -1, & \text{if } i = t, \\ 0, & \text{otherwise.} \end{cases}, \quad \forall i \in N, \\
 & x_{ij} \geq 0, \quad \forall (i, j) \in A.
 \end{aligned}$$

Let N be the set of nodes and A the set of arcs. Also, denote $s \in N$ as the source/start node and $t \in N$ as the terminal/end node. Then, let the binary variables x_{ij} signal whether the arc is part of the shortest path connecting s to t . Each arc is associated with a cost of traversing it, denoted by c_{ij} . Hence, minimizing the total cost (i.e., identifying the shortest path), is done by minimizing $\sum_{(i,j) \in A} c_{ij} \cdot x_{ij}$. Moreover, the traditional flow preservation constraints must be included. The shortest path for this simulation was calculated using Dijkstra's algorithm (Dijkstra 1959), a poly-time algorithm that allows us to dispel with the formulation.

Penalty Generation

The penalty generation function is designed to change the cost dictionary used by Dijkstra's algorithm in the simulation model. Two cost dictionary data structures are used in the simulation. In the baseline cost dictionary the value 1 is set $\forall (i, j) \in A$. The penalty cost dictionary is adjusted based on the location of the adversary. Since the risk-based goal of the routing strategy is to avoid the adversary, penalty rings are developed to force the transport away. In the simulation, the penalty generation function takes the current position of the adversary and adds a progressive penalty to the surrounding edges connecting the defined nodes. For the proof of concept, a five ring cookie cutter sensor is used to simulate the sensor range of the adversary and vary the level of threat for the transport based on proximity. In the operational context, this represents a radius of 50 NMs from the adversary. This value could be debated by officers of the craft but is sufficient for this proof of concept.

3.2.4 Simulation Execution

In this section we walk through the simulation execution for the proof of concept simulation model. Figure 3.2 visually depicts the execution of the simulation flow algorithm.

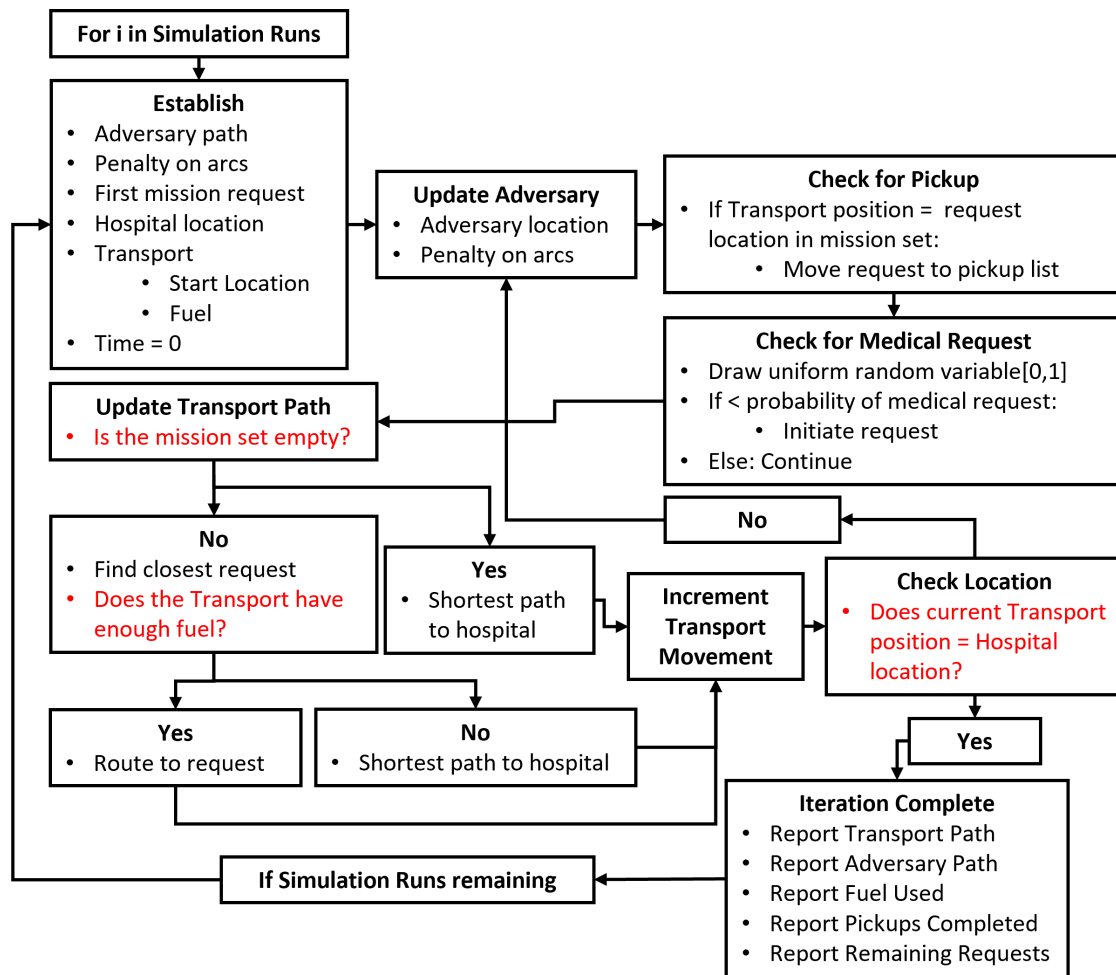


Figure 3.2. Process flow for proof of concept simulation execution. Note: Red text represents binary decisions made in the simulation model.

The sections that follow discuss the execution of the processes in the algorithm in further detail. This algorithm serves as a computational representation of the decision process we propose in this thesis. The execution of the algorithm in Python serves as the proof of concept that the process can be implemented in a programming language.

Simulation Setup

To setup the simulation, we first define the number of simulation runs. A simulation run begins with the execution of the run setup and concludes when the transport reaches the hospital. Next, we define the number of time steps and the probability of a medical request occurring during a time step. Defining the number of time steps is used as a control measure to protect against an infinite loop. Given this, and to ensure the simulation runs through, the number of time steps is greater than the transport fuel. Not all time steps may be used since the code breaks upon arrival at the hospital regardless of the number of remaining time steps. Finally, the number of threat rings for the adversary is defined, up to five rings or 50 NMs in the operational context.

Run Setup

The run setup begins with establishing the information necessary for the Transport class to begin routing. To develop the penalty dictionary, the initial path for the Adversary class is established using the initialization process discussed previously. Next, the first request is developed by the Medical Requests class and added as the first entry in the mission set. The Hospital location is assigned based on the structure. If fixed, then a node is assigned and does not change. If mobile, then a path is developed based on the criteria defined. Finally, the parameters for the Transport class are developed, including the starting location, the fuel capacity, and the path to the first request in the mission set.

Simulation Run

The simulation run algorithm proceeds as follows:

algorithm transport routing under uncertain demand and adversary movement

begin

establish run setup

set max time steps

set first medical request and probability of medical request in a time step

set hospital location or path

set transport start location and fuel capacity

set threat ring penalties and non-penalized edge cost

establish first adversary path

```

for  $i$  in time steps do
  begin
    increment the adversary 1 step along the path list
    generate penalty dictionary based on adversary position
    if transport location = mission set entry then
      append entry to picked-up list and remove entry from mission set list
    generate uniform random number between 0 and 1
    if number  $\leq$  probability of medical request during a time step then
      draw a node  $\in N$  with uniform probability of selection
      append the node to the mission set list
    update transport path
    begin
      if length of mission set = 0 then
        route to the hospital
      else length of mission set > 0 then
        compute the shortest path to all locations in the mission set
        find the entry in the mission set with the shortest distance
        compute the shortest path from the entry to the hospital
        if (distance from transport to entry) + (distance from entry to hospital) <
        remaining fuel then
          route to the entry
        else route to the hospital
      end
    increment the transport 1 step along the path list
    update transport current location
    append transport current location to visited list
    remove entry from path list
    if transport current location = hospital location
      break to end simulation run
    else increment  $i$  in time steps and continue
  end
  record data from simulation run
end

```

The simulation run begins with movement of the adversary along its defined path. If the path is empty, then a new path is developed at this time. The edges connected to the adversary's current location are penalized in progression out five nodes to update the penalty dictionary used by the transport. Next, a check is performed to see if the transport is located at the same node as an entry in the mission set. If it is, the request is removed from the mission set and transferred to the pickup set. Then a decision is made to generate a medical request during the current time step. A number is drawn from a uniform random $[0, 1]$ then compared to the probability of a medical evacuation request occurring within a time step. If the value is less than the probability, a request is generated. If it is greater, no new requests occur during the time step. The next steps correspond to the decision process we propose in this study. Now that the tactical picture for the time step is available to the transport, the following process is conducted:

- (1) If the mission set is empty, route to the hospital.
- (2) Else, if the mission set is not empty:
 - a. Compute the shortest path to all locations in the mission set and append the distances to a list.
 - b. Find the request corresponding to the minimum distance on the list, which is the closest request, and identify the entry in the mission set.
 - c. Find the distance from the entry to the hospital.
 - d. If the remaining fuel is greater than the combined distance from the current position to the request location and the request location to the hospital then route to the request location.
 - e. Else, if the remaining fuel is less than the combined distance, then route to the hospital.

Note the transport shortest path process uses the penalty cost dictionary while the adversary shortest path uses the baseline cost dictionary. Once the path is updated, the transport moves forward one step along the developed path. A check is then made to see if the current position matches the location of the hospital. If the current position matches, the simulation ends. If the position does not match, then the loop repeats until the criteria is met.

The *repeat* component of the simulation run represents the embedded use of online optimization in the simulation model as opposed to simple offline vehicle routing. At the end of the loop, the algorithm has completed the actions based on the data available at the

time of loop initiation. The transport then takes in new information about the environment and the stochastic demand. A new loop is then initiated with this new data to take another step towards the transport's objective. Under simple vehicle routing conditions with perfect information, the transport would continue without recomputing the route, hence not responding to new adversary positioning or stochastic demand. Under online optimization, the transport begins with imperfect information and optimizes the route at each defined time interval to adjust to the new data made available at the beginning of the time step. In the next chapter we will discuss the results of the proof of concept simulation and propose an application of the results from multiple simulation runs to the development of a distance multiplier for estimating the impact of adversary presence in simulation modeling.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4: Simulation Results and Discussion

In this chapter, we discuss the results of the single run of the proof of concept simulation described in Chapter 3. We present model modifications from baseline to compare the distance traveled and the number of requests fulfilled given the presence or non-presence of an adversary and the nature of information. We reemphasize the point estimate values from the single simulation runs are not important in this context, however the implications of the magnitude and direction of difference are. Each section therefore provides results in the operational context presented in Chapter 2 to provoke thought on the topic and incite future research into specific situations where this approach to simulation modeling may apply. Next, we use a data farming approach to multiple runs of the proof of concept simulation to propose a model for the construction of a distance multiplier for estimating the impacts of adversary presence in simulation modeling. Finally, we discuss the implications of this research and the medical considerations for the use of unmanned and autonomous systems.

4.1 Proof of Concept Simulation Results

The primary result of the simulation is the path taken by each object in the model. For this proof of concept, one transport and one adversary are simulated with the hospital in a fixed position. The parameters for the simulation are as follows:

Max Time Steps: 200 Iterations

Probability of Medical Request: 0.1

Hospital Location: Bottom Right Corner

Transport Start Location: Top Left Corner

Transport Fuel Capacity: 120 Movements

Threat Ring Arc Penalties:

(1) 100 (2) 50 (3) 25 (4) 10 (5) 5

Non-Penalized Arcs: 1.

Figures 4.1 to 4.3 are visualizations from a single simulation run developed using networkx (Hagberg et al. 2008). The blue area represents the nodes available to traverse in the network.

The grey nodes are the path taken by the transport during the simulation run, while the red nodes are the path taken by the adversary. The green circles denote the pick-ups that occurred, while the orange circles are the requests remaining at the end of the simulation run. In the sections that follow, we discuss the results of three different simulation structures, which use the same initiating data, to compare two different aspects of the operational problem, distance and mission accomplishment.




Figure 4.1. Visualization of the baseline simulation results for routing under uncertain demand and adversary threat. Legend: Grey Path-Transport, Red Path-Adversary, Green Circles-Completed Requests, Orange Circles-Missed Requests.

4.1.1 Baseline Simulation

The baseline simulation structure is transport routing under uncertain demand and adversary threat. The resulting visualization of this structure is presented in Figure 4.1. The next two structures were used to compare the baseline results to different scenarios which include perfect information and no adversary. Modification (1) routes the transport against the list of locations where pick-ups occurred in the baseline model. The intent of modification

(1) is to compare the distance traveled by the transport under different threat conditions. Modification (2) routes the transport against the list of all the locations where requests were made in the baseline model. The intent of modification (2) is to compare the number of request fulfilled by the transport under different threat conditions. Figures 4.2 and 4.3 are a visualization of the results of modifications (1) and (2) respectively.




Figure 4.2. Visualization of simulation modification (1) results with perfect information for routing against the pick-up list from the baseline simulation only without adversary threat. Legend: Grey Path-Transport, Red Path-Adversary, Green Circles-Completed Requests, Orange Circles-Missed Requests.

4.1.2 Simulation Modification (1)

First, we compare the results of the baseline simulation and modification (1). In the baseline simulation, the total distance traveled by the transport is 117 units under adversary threat and with unknown future information on demand. In modification (1), without the threat of an adversary and with perfect future information, the total distance traveled is 78 unit, which is a difference of 39. If the results from this single simulation run are put into an operational context, at a distance of 10 NMs between waypoints, the transport travels an

additional 390 NMs. For mission planning, this means the exclusion of the adversary in modeling could have impacts to the allocation of fuel and add transit time to the next role of medical care, potentially affecting patient outcomes.




Figure 4.3. Visualization of simulation modification (2) results with perfect information for routing against the full mission set from the baseline simulation without adversary threat. Legend: Grey Path-Transport, Red Path-Adversary, Green Circles-Completed Requests, Orange Circles-Missed Requests.

4.1.3 Simulation Modification (2)

Next, we compare the results of the baseline simulation and modification (2). In the baseline simulation under adversary threat and with unknown future information on demand, the total number of requests completed in the mission set is 14, with six remaining at the end of the simulation run. In modification (2), without the threat of an adversary and with perfect future information, the total number of requests completed in the mission set is 17, with three remaining. If the results from this single simulation run are put into an operational context, the transport is able to respond to three less requests given the fuel restrictions than if the adversary is not present, potentially overestimating patient survival in the operating

area. In a comparison of Figure 4.1 and Figure 4.3, it should be noted that the transport in Figure 4.3 picks up requests that are unreachable due to adversary interference in Figure 4.1. If the adversary had been present, the transport in Figure 4.3 would have had a higher chance of detection and interaction with the adversary.

4.1.4 Proof of Concept Model Applications

There are multiple possible approaches for constructing a proof of concept model to test the computational tractability of the decision process we propose in this thesis. Our approach uses a baseline model with modifications to explore two aspects of the operational problem, distance and mission accomplishment. We accomplish this through simulation modeling and visual exploration of individual simulation run results. Another approach is data farming which requires significant refinement for application. In the next section, we begin the exploration of a possible application of multiple runs of the proof of concept simulation model using data farming, however, due to the need for specific parameter information, the applicability is only conceptual in nature at this time. We therefore only propose an approach to this concept and leave refinement as an opportunity for future research.

4.2 Distance Multiplier

In this section, we propose an application of multiple runs of the proof of concept simulation model using a data farming approach. Data farming is the action of developing a simulation model to grow data, as opposed to analyzing preexisting data to gain insights. Many current military logistics models, such as the Joint Medical Planning Tool, use a straight line approach, such as Euclidean or Great-Circle distance, to route supplies and casualties to different demand locations without consideration for the maneuver associated with the presence of an adversary. Since reprogramming current simulation models is an expensive undertaking, we propose using a simple simulation model and a data farming approach to estimate the impacts of a contested environment. The farmed data can be used to develop a multiplying factor to apply to the Euclidean or Great-Circle distance to estimate adversary presence. This approach may enable a more accurate representation of time and distance, given an adversary threat in the interim between model development phases.

4.2.1 Approach

The fundamental approach to developing a multiplier is to first develop a representative simulation model with the presence of the adversary and without. This must be done using an estimate of the tactical actions the ship is allowed to take and the risk the commander is willing to accept. This may be best accomplished using a wargaming approach to obtain the qualitative nuances of the problem for conversion to quantitative use. For this research, we use the standing assumptions and the proof of concept simulation model as our representative simulation to discuss our approach to development without a complete algorithm re-design. Refer to Chapter 3 for a description of the proof of concept development. Note, the simulation model for this analysis must focus on the unit the multiplier is applied to. For example, our focus is on the medical transport. In this case, wargaming and the representative simulation must be designed with the medical transport at the center of development. Additionally, given the commander will likely have different impressions on the level of risk they are willing to take with different units, application of a pre-developed distance multiplier to another type of unit or different operational situation may not appropriately capture the risk-based approach to routing the commander intends to use in the actual situation.

4.2.2 Conceptual Development Process

The following conceptual process for the development of a distance multiplier may be helpful to the future researcher:

1. Research the modeling program the distance multiplier will be used within to identify how it will be implemented in the simulation model.
2. Identify the unit of interest within the modeling program the distance multiplier will be applied to.
3. Conduct a subject matter expert event, such as a wargame, to experiment with different risk parameters. During wargame design, ensure the scenarios presented to the participants have a pre-structured approach to convert the results of the wargame into implementable quantitative modeling parameters for use in the representative simulation.
4. Construct the representative simulation with and without the presence of the adversary in the model, similar to the approach we used for the proof of concept develop for this thesis. Our recommendation is using the baseline simulation and modification

- (1), since modification (2) addresses a different problem.
5. Farm multiple runs of the representative simulation for analysis.
 6. Conduct data analysis on the farmed data. Develop different levels of risk for variable parameters in the simulation. Point estimates are not advisable. Risk variation offer the opportunity for stochastic implementation of the distance multiplier in the pre-existing simulation model.
 7. Test the risk based multipliers in the pre-existing simulation model to ensure the results make sense given the scenario.

We recommend that, where possible, the developer makes an effort to visualize the results of the representative simulation. We found during the development phases for this thesis that minor adjustments to the algorithm have significant impacts on the output of the model, such as the order of a calculation. Without visualization, these issues may not have been discovered.

4.2.3 Data Farming and Analysis Example

To provide an abbreviated example of this process, we farm the data from the proof of concept simulation by conducting 1000 simulation runs from the baseline and modified scenarios. Depending on the development environment used for the simulation model, internal analysis or export may be required to put the data into a format suitable for analysis. A best practice is to first understand the data required to answer the research questions, then design the results to limit derived values from raw data that could otherwise have been captured with good simulation design. For our analysis, the data points most applicable to the development of a distance multiplier are:

- (1) Total Distance Traveled by the Transport (Integer);
- (2) Transport's Remaining Fuel (Integer);
- (3) Over Fuel Indicator (Binary: 0 if under, 1 if over);
- (4) Total Pickups from the Mission Set (Integer);
- (5) Total Requests Remaining in the Mission Set (Integer); and
- (6) Total Simulation Execution Time (Float).

Parameters (1) and (2) are used to compute the difference in distance traveled between the baseline simulation and modification (1). Parameter (3) is used as an indicator of

mission failure, where refueling was required. Parameters (4) and (5) are used to verify both variations of the simulation are using the same data, and parameter (6) is used to explore the run-time in the model. Parameter (6) is useful when planning the time associated with simulation execution development. If time is compressed, such as in crisis planning, fewer replications may be used, while in deliberate planning situations, the time horizon may allow for more replications and varied representative simulation parameters to compare results.

Simulation Run Time

For data analysis, the results from the 1000 simulation runs using Python were captured as an .xlsx file and imported into R for analysis. After removing the error records 946 full runs were available for analysis. The mean simulation run time across the 946 iterations was 2.30 minutes with a standard deviation of 0.64 minutes. The maximum simulation run time was 4.50 minutes and the minimum was 0.34 minutes.

Mission Accomplishment

After cleaning the data and conducting exploratory analysis, we turn our focus to the difference in distance traveled between the baseline simulation and modification (1). To begin shaping the distance multiplier and to understand the parameters in the simulation model, we explore mission accomplishment in the baseline simulation. The transport successfully completed the mission within the fuel constraints in 829 runs and fell below the fuel red line in 117 runs. In the operational context, the lack of adversary modeling leads to an underestimation in the fuel requirement in 117 runs, given the uncertain demand in the simulation. This shows that defining the fuel red line in a wartime scenario is not a trivial matter. The potential to underestimate fuel consumption under these conditions can result in mission delays which extend the transfer time for critical patients.

Distance Multiplier Construction

Once we have an understanding of the simulation model data, we can begin exploring the construction of the distance multiplier. Let the length of the path traveled in the baseline scenario represent the value of the path with the adversary present. Let the length of the path traveled in the simulation with modification (1) represent the path that would be taken

in the pre-existing simulation model. Convert the values into a percent difference between the two paths using the following equation:

(a) = length of the path in the baseline simulation

(b) = length of the path in simulation modification (1)

$$\frac{a - b}{b}$$

Add this value to 1 to represent the addition of the additional distance to the base value in the scenario.

Length of Transport Path in Baseline Simulation

The mean length of the path traveled (fuel used) in the baseline simulation was 111 with a standard deviation of 16.68. The minimum path traveled was 39 units and the maximum was 171. Note that the minimum represents the situation where only one request occurs in the simulation and the maximum show the highest mission failure value over the 120 fuel units that occurred across the 946 runs.

Length of Transport Path in Simulation Modification (1)

The mean length of the path traveled (fuel used) in simulation modification (1) was 82 with a standard deviation of 17.22. The minimum path traveled was 36 units and the maximum was 119. Note that the minimum represents the situation where only 1 request occurs in the simulation. If the maximum path exceeds the fuel limit in this modification, an evaluation of the simulation algorithm should be conducted to gain insight on why this occurred. Given this modification is only using the pick-ups that occurred in the baseline simulation, a result over the maximum fuel defined in the simulation model would be counter intuitive.

Proposed Distance Multiplier

The proposed distance multiplier applies the equation above to each of the 946 simulation runs then constructs an estimate at different levels of risk. The mean difference was 0.39 with a standard deviation of 0.31 for the 946 runs. The minimum difference was 0 and the maximum was 2.04. The values we propose for use correspond to the density of the difference in distance traveled.

1. 50 percent quantile = 0.32 (High Risk).
2. 75 percent quantile = 0.55 (High Risk).
3. 85 percent quantile = 0.71 (Medium Risk).
4. 95 percent quantile = 0.98 (Low Risk).

To interpret the above results, consider a building density (cumulative density) of the values in the simulation. Choosing the 50 percent quantile represents high risk where we are only accounting for 50 percent of the cases that can occur. In contrast, choosing the 95 percent quantile, we account for 95 percent of the cases that can occur and would be considered low risk. Considering the maximum value in the simulation of 2.04 in this context, theoretically, we would account for the worst case scenario given the assumptions we make in the representative simulation model. To apply this in the pre-existing simulation model, add 1 to the value above. For example, to use the 50 percent quantile calculate 1 plus 0.32, which is 1.32. If the value for the distance in the existing model is 100 NMs, then multiply 1.32 times 100 resulting in 132 NMs to simulate the presence of the adversary at this level of risk.

4.3 Discussion

The simulation model described in the previous section provides a proof of concept to showcase how the decision process we propose in this thesis can be applied to route a medical transport through a network with uncertainty in demand and under adversary threat. George Box notes “all models are wrong, but some are useful” (Wasserstein 2010). The simulation model, as applied, has many limitations that represent opportunities for future work, however, it shows the possibilities that this line of research has when applied to wartime logistics routing. If refined, this decision process could inform the development of decision algorithms implemented in autonomous systems that operate under uncertainty.

An outcome of the exploratory research conducted is the application of machine learning to this problem. In the single objective scenario, given unpredictable demand and adversary movement, machine learning is a poor fit for the problem. Given the chaos that occurs in conflict, machine learning is ill prepared at this stage of development to predict, based on the lack of applicability of historical information on future demand. While predictive algorithms could be used on adversary and friendly force movements to anticipate the

location of mass casualty events resulting from force-on-force exchanges, the unanticipated accident onboard a ship, such as a fall, would not be included in the analysis. Given this, the use of Boolean logic was found to be a better alternative.

With more depth, higher layer weighted decisions may benefit from machine learning, such as those in multiple objective optimization. For the individual objectives, online optimization using Boolean logic seems to perform better. When presenting computational models to military commanders, it is useful, where possible, to convey the process in a context that is familiar to the commander. In this thesis, we compared the decision process to the USMC Rapid Response Planning Process. The intent is the comparison frames the process in a way that appeals to the commander's previously established mental schema. Conveying how the algorithm works in this manner provides the commander with a better understanding of the decision process and the ability to shape it to match the desired mission effect. The result of this development-operations interaction can further accelerate the decision loop and provide commanders with confidence in the machine when participating in human-machine teaming.

4.4 Medical Considerations for Autonomous Systems

The notional scenario we use in this thesis uses medical evacuation operations as the context. Therefore, it is important to understand the implications of unmanned and autonomous systems within this scenario. At the operational level of war, there are two principal issues of concern with the development of autonomous systems. The first is the high level of maneuverability possible which may function outside of a human-safe envelope. The second is the size and manning available onboard these systems to allow them to fulfill multiple missions. For medical transport, the system design must be such that it is considered a semi-manned capability, simply, it functions without an engineering or piloting crew onboard, but must be able to support embarked personnel.

The Combat Logistics Fleet plays a part in the evacuation of patients from combatant ships, especially in wartime. With the shift to unmanned logistics, the medical evacuation network loses some capacity to evacuate patients due to the lack of support onboard for embarked personnel, by design. As the military explores solutions to fill this gap with other transportation capabilities, some simple considerations for the use of unmanned or

autonomous systems should be:

1. Can the transport navigate under conditions that are safe for the movement of critically injured patients?
2. Are there facilities onboard, such as living quarters and a galley, for the patients and support staff?
3. Can the transport connect to the combatant ships it will support in a manner safe for patient transfer?
4. If a patient is placed onboard the transport without a medical attendant, what are the ethical implications of this action?

While we do not address these issues directly, the research surrounding this routing problem must consider the operation of the capability under these circumstances since the research community influences the development of new unmanned and autonomous systems. One approach we offer is to explore multi-use capabilities as a semi-manned autonomous systems rather than the blanket term unmanned.

CHAPTER 5: Conclusion and Future Research

5.1 Conclusion

In this thesis, we propose a decision process to build upon when considering routing strategies against unknown demand and adversary threats using the principles of online optimization. We compare the decision process to USMC R2P2 to show the parallels between military planning and our research. Additionally, we use the military planning process as a medium to converse with commanders not versed in programming about the structure of the decision process we propose in this thesis. This is done to convey in a common language how the decision process works to promote confidence in the machine aspect of human-machine teaming. In the proof of concept we showcase the computational tractability of the decision process by developing an algorithm and implementing it as a simulation model. We show how the results can be visualized to explore different aspects of the problem.

We propose a data farming approach to apply the results of multiple runs using the algorithm to develop a simple distance multiplier to approximate the presence of an adversary in current simulation models such as the Joint Medical Planning Tool. Our intention is to apply additional realism to existing simulation models between development phases. Finally, we discuss the implications of the potential application of the decision process to autonomous and unmanned systems for multi-mission use. Our premise being that combat system engineers should consider the physical space and navigation implications of humans onboard these highly maneuverable systems in scenarios where multi-mission use, including medical applications, may be considered.

5.2 Future Research

The Hybrid Fleet, Naval Expeditionary Health Service Support, and Contested Distributed Logistics are emerging areas within the military service that are undergoing a sea change. As such, we see many opportunities based on our research to further explore employment

strategies for maritime logistics and medical forces in support of new force design initiatives and distributed operations.

5.2.1 The Future Hybrid Fleet

Technological innovation is pushing Navy operating concepts towards a Hybrid Fleet of manned and autonomous systems to remain competitive in current and future environments (Chief of Naval Operations 2022). This shift in U.S. Naval force structure will require a shift in the approach to force employment. By nature, military operations function under conditions of uncertainty. Commanders that employ the Hybrid Fleet will need to be confident in the ability of machine in the human-machine team to shorten the decision cycle and effectively employ the Fleet under combat conditions. In the manned force, standardized training, universal task lists, and doctrine provide the operational commander with an expected level of performance under certain conditions. Simply, when a specific command is given, the commander assumes a specific action associated with that command is taken within the time frame defined. This expectation must remain unchanged in the Hybrid Fleet.

While the decision process we propose in this thesis is based on the U.S. Marine Corps Rapid Response Planning Process, other constructs, such as the Observe, Orient, Decide, and Act loop can also be used to convey the structure of the decisions made by the machine to the commander (Osinga 2007). Our research does not purport to define a new planning or cognitive construct for decision making but connects the computational construct for decisions under uncertainty to a familiar decision process to allow for a conversation, in a common language, between the autonomous system developer and the operational commander. To effectively employ emerging technologies, autonomous system developers must strive to make the decision process as clear as possible, so commanders have confidence in the human-machine team under the most adverse conditions.

To improve upon this research and to further the objectives of the Hybrid Fleet force design, we believe that multiple objective optimization is the next logical step to further implement commander's intent into the autonomous decision process. While our proof of concept algorithm was designed to showcase computational tractability, we believe that refinement to include a multiple objective approach, such as goal programming, offers an opportunity

to incorporate risk based strategies given the adversary rather than simple avoidance (Royset 2021, pp. 5-6). Under the medical context, a transition to multiple objective optimization also offers an avenue to influence routing based on patient concentrations and severity to improve outcomes.

5.2.2 Naval Expeditionary Health Service Support

In the simulation model for this research, we use rudimentary medical parameters since the focus is on routing the transport through the network under threat. Further research could explore methods for embedding adversary modeling in existing medical network planning tools, such as the Joint Medical Planning Tool. While the simplistic distance multiplier could be used as a stop-gap measure to simulate the presence of the adversary along a route, further research and varied simulation model parameters would be necessary to make the multiplier usable in medical simulation modeling. The decision process could also be applied to a larger modeling context where there are multiple adversaries and transports at the operational level of war, which will require an additional decision layer to select the transport to route to specific requests.

Future research and simulation modeling could look at the application in littoral and island locations. Our simulation model assumes generally open water and small island (<10 NM) conditions. A higher resolution model with more waypoints or different decision points may be necessary to simulate routing under littoral and island conditions. The comparison of the resolution requirements under different geographic conditions could influence the modeling assumptions of different geographic regions and contribute to online routing research generally.

5.2.3 Distributed and Contested Logistics Support

In addition to health service support planning, we see parallels in our research with general logistics routing under uncertainty. Adversary avoidance, whether under wartime or dangerous weather conditions, is not limited to medical routing. Push-Logistics networks also operate under a similar demand structure, and in some cases, urgency. The coupling of weather avoidance and humanitarian relief operations, both foreign and domestic, also represent a similar problem as conditions can change quickly and in unpredictable ways. Given these scenarios, we see the principles within our decision process as generally applicable

to logistics routing under uncertainty. Additionally, with the advent of unmanned logistics, the algorithmic considerations in this thesis may contribute to the further development of employment applications for these systems in future concepts.

5.2.4 Communication Strategies for Autonomous Systems

To frame our problem within the future operating environment we explore communication and control strategies for unmanned surface vessel routing. As future warfighting concepts consistently include degraded communications, novel ways to approach communicating with manned, unmanned, and autonomous craft must be explored. Our application of online optimization to the problem of uncertain demand and adversary presence brought forward multiple issues related to the control of manned and unmanned systems. This includes the use of waypoints, line of site communications, and positive control of the vessel at certain ranges. Further exploration of the operational implications of the communication infrastructure under denied conditions is needed to develop an applicable control strategy. This research has relevant implication for emerging medical technologies, such as unmanned evacuation and telemedicine capabilities, including broadcast telemetry and robotic treatment, as continuous control via communications channels should not be expected in the future contested environment.

List of References

- Ahuja R, Magnanti T, Orlin J (1993) *Network Flows: Theory, Algorithms, and Applications* (Prentice-Hall, New Jersey).
- Aistrup J (2021) USNS Burlington (T-EPF-10) Operates Off The Coast Of Haiti [Online Image, Public Domain]. Defense Visual Information Distribution Service. <https://www.dvidshub.net/image/6805334/usns-burlington-t-epf-10-operates-off-coast-haiti>.
- Belmega V, Mertikopoulos P, Negrel R, Sanguinetti L (2018) Online convex optimization and no-regret learning: Algorithms, guarantees and applications ArXiv:1804.04529.
- Chief of Naval Operations (2022) Navigation Plan. Technical report, Department of the Navy, https://media.defense.gov/2022/Jul/26/2003042389/-1/-1/1/NAVIGATION%20PLAN%202022_SIGNED.PDF.
- Dantzig G, Ramser J (1959) The truck dispatching problem. *Management Science* 6(1), <https://doi.org/10.1287/mnsc.6.1.80>.
- Dijkstra E (1959) A note on two problems in connexion with graphs. *Numer. Math.* 1, <https://doi.org/10.1007/BF01386390>.
- Gillingham B, Kurtz C, Daniels D, Marler K, Robertson A (2022) A Commentary on Baker, M., Baker J., & Burkle, F. (2022) The Hospital Ship as a Strategic Asset in 21st Century Foreign Policy and Global Health Crises]. *Military Medicine* usac(062), <https://doi.org/10.1093/milmed/usac062>.
- Google Developers (2022) Google OR-Tools Vehicle Routing. Accessed June 15, 2022, <https://developers.google.com/optimization/routing>.
- Hagberg A, Schult D, Swart P (2008) Exploring network structure, dynamics, and function using networkx. *Proceedings of the 7th Python in Science Conference (SciPy 2008)* (Pasadena, CA), <https://www.osti.gov/biblio/960616-exploring-network-structure-dynamics-function-using-networkx>.
- Hazan E (2021) *Introduction to Online Convex Optimization, 2nd Ed.* <https://arxiv.org/abs/1909.05207>.
- Joint Chiefs of Staff (2018) *Joint Health Services*. JP 4-02, Department of Defense, https://www.jcs.mil/Portals/36/Documents/Doctrine/pubs/jp4_02ch1.pdf.

- Lawless W, Mittu R, Sofge D, Hiatt L (2019) Artificial Intelligence, Autonomy, and Human-Machine Teams — Interdependence, Context, and Explainable AI. *AI Magazine* 40(3), <https://doi.org/10.1609/aimag.v40i3.2866>.
- Liu Y, Vogiatzis C, Yoshida R, Morman E (2022) Solving Reward-Collecting Problems with UAVs: A Comparison of Online Optimization and Q-Learning. *Journal of Intelligent & Robotic Systems* 104(35), <https://doi.org/10.1007/s10846-021-01548-2>.
- Marler K (2022) *Medical Planning Factors for Offshore Personnel Rescue: An Integrative Study*. Dissertation, Central Michigan University, Mount Pleasant, MI, ProQuest Dissertations and Theses database (2649062738).
- Osinga F (2007) *Science, Strategy, and War: The Strategic Theory of John Boyd* (Routledge, London).
- Petch J, Di S, Nelson W (2022) Opening the Black Box: The Promise and Limitations of Explainable Machine Learning in Cardiology. *Canadian Journal of Cardiology* 38, <https://www.sciencedirect.com/science/article/pii/S0828282X21007030>.
- Royset J (2021) Good and Bad Optimization Models: Insights from Rockafellians ArXiv:2105.06073.
- U.S. Marine Corps (2022) *Marine Corps Planning Process*. MCWP 5-10, Department of the Navy, <https://www.marines.mil/Portals/1/Publications/MCWP%205-10.pdf?ver=CK4C14H7VdKkvnhSHhJUg%3d%3d>.
- Wasserstein R (2010) George Box: a model statistician. *significance* 7(3), <https://doi.org/10.1111/j.1740-9713.2010.00442.x>.
- White House (2017) National Security Strategy of the United States of America. Technical report, Washington, DC, <https://permanent.fdlp.gov/lps90878/2017/NSSFfinal121820170905.pdf>.
- White House (2021) Interim National Security Strategic Guidance. Technical report, Washington, DC, <https://www.whitehouse.gov/wp-content/uploads/2021/03/NSC-1v2.pdf>.

Initial Distribution List

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California