

3-23-2017

Using Markov Decision Processes with Heterogeneous Queueing Systems to Examine Military MEDEVAC Dispatching Policies

Phillip R. Jenkins

Follow this and additional works at: <https://scholar.afit.edu/etd>



Part of the [Operational Research Commons](#)

Recommended Citation

Jenkins, Phillip R., "Using Markov Decision Processes with Heterogeneous Queueing Systems to Examine Military MEDEVAC Dispatching Policies" (2017). *Theses and Dissertations*. 797.
<https://scholar.afit.edu/etd/797>

This Thesis is brought to you for free and open access by AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact richard.mansfield@afit.edu.



Using Markov Decision Processes with
Heterogeneous Queueing Systems to Examine
Military MEDEVAC Dispatching Policies

THESIS

Phillip R. Jenkins, Capt, USAF
AFIT-ENS-MS-17-M-137

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views expressed in this document are those of the author and do not reflect the official policy or position of the United States Air Force, the United States Department of Defense or the United States Government. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

AFIT-ENS-MS-17-M-137

USING MARKOV DECISION PROCESSES WITH HETEROGENEOUS
QUEUEING SYSTEMS TO EXAMINE MILITARY MEDEVAC DISPATCHING
POLICIES

THESIS

Presented to the Faculty
Department of Operational Sciences
Graduate School of Engineering and Management
Air Force Institute of Technology
Air University
Air Education and Training Command
in Partial Fulfillment of the Requirements for the
Degree of Master of Science in Operations Research

Phillip R. Jenkins, BS
Capt, USAF

March 2017

DISTRIBUTION STATEMENT A
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

AFIT-ENS-MS-17-M-137

USING MARKOV DECISION PROCESSES WITH HETEROGENEOUS
QUEUEING SYSTEMS TO EXAMINE MILITARY MEDEVAC DISPATCHING
POLICIES
THESIS

Phillip R. Jenkins, BS
Capt, USAF

Committee Membership:

Lt Col Matthew J. Robbins, PhD
Chair

Dr. Brian J. Lunday
Member

Abstract

A major focus of the Military Health System is to provide efficient and timely medical evacuation (MEDEVAC) to battlefield casualties. Medical planners are responsible for developing dispatching policies that dictate how aerial military MEDEVAC units are utilized during major combat operations. The objective of this research is to determine how to optimally dispatch MEDEVAC units in response to 9-line MEDEVAC requests to maximize MEDEVAC system performance. A discounted, infinite horizon Markov decision process (MDP) model is developed to examine the MEDEVAC dispatching problem. The MDP model allows the dispatching authority to accept, reject, or queue incoming requests based on the request's classification (i.e., zone and precedence level) and the state of the MEDEVAC system. Rejected requests are rerouted to be serviced by other, non-medical military organizations in theater. Performance is measured in terms of casualty survivability rather than a response time threshold since survival probability more accurately represents casualty outcomes. A representative planning scenario based on contingency operations in southern Afghanistan is utilized to investigate the differences between the optimal dispatching policy and three practitioner-friendly myopic baseline policies. Two computational experiments, a two-level, five-factor screening design and a subsequent three-level, three-factor full factorial design, are conducted to examine the impact of selected MEDEVAC problem features on the optimal policy and the system level performance measure. Results indicate that dispatching the closest available MEDEVAC unit is not always optimal and that dispatching MEDEVAC units considering the precedence level of requests and the locations of busy MEDEVAC units increases the performance of the MEDEVAC system. These results inform the development and implementation of

MEDEVAC tactics, techniques, and procedures by military medical planners. Moreover, an open question exists concerning the best exact solution approach for solving Markov decision problems due to recent advances in performance by commercial linear programming (LP) solvers. An analysis of solution approaches for the MEDEVAC dispatching problem reveals that the policy iteration algorithm substantially outperforms the LP algorithms executed by CPLEX 12.6 in regards to computational effort. This result supports the claim that policy iteration remains the superlative solution algorithm for exactly solving computationally tractable Markov decision problems.

Keywords: Markov decision processes, medical evacuation, admission control, queueing, priority dispatching, policy iteration, and linear programming comparison

I dedicate this thesis to the men and women who have fought and died in service to the United State of America. My hope is that this research is utilized and continued in an effort to provide the most efficient medical evacuation system possible for those who risk their lives to defend our country.

Acknowledgements

There are no words that convey my deep gratitude and respect for my thesis advisor and mentor, Dr. Matthew Robbins. His untiring support and guidance over this past year is what made this thesis possible. I am thankful for having the opportunity to work under his advisement.

I would also like to acknowledge my committee member, Dr. Brian Lunday, for being the second reader, and I am gratefully indebted to him for his valuable comments that improved the presentation and contents of this thesis.

I would also like to thank my family. Their continuous encouragement and support kept me motivated throughout the development of this thesis. My father taught me at a young age that if something is worth doing, it is worth doing right. Moreover, my mother instilled a competitive mindset in me that has influenced me to do anything and everything to the best of my ability. I would especially like to thank my beautiful wife and daughter for their unwavering love and patience during the countless hours I spent working on this thesis. It is easy to say that I would not have made it this far without them by my side.

Phillip R. Jenkins

Table of Contents

	Page
Abstract	iv
Dedication	vi
Acknowledgements	vii
List of Figures	ix
List of Tables	x
I. Introduction	1
II. Literature Review	6
III. Problem Description	19
IV. Methodology	29
V. Testing, Analysis & Results	43
5.1 Representative Scenario	43
5.2 Representative Scenario Results	51
5.3 Experimental Design	59
5.4 Experimental Design Results	60
5.5 Excursion 1 - Request Arrival Rate	67
5.6 Excursion 2 - MEDEVAC Flight Speed	68
5.7 Excursion 3 - Intra-Zone Policies	70
5.8 Excursion 4 - 3×3 case	76
5.9 Policy Iteration versus Linear Programming	84
VI. Conclusion	88
Appendix A. Storyboard	92
Bibliography	93
Vita	99

List of Figures

Figure	Page
1	MEDEVAC Mission Timeline 20
2	MEDEVAC Queueing System 23
3	Afghanistan combat fatalities by province, 2001-2016 45
4	MEDEVAC and MTF locations with Casualty Cluster Centers..... 47
5	Sampled Casualty Events in Helmand and Kandahar 48
6	Comparison of MEDEVAC 1 busy probabilities 56
7	Comparison of MEDEVAC 2 busy probabilities 57
8	Comparison of zone-precedence queue lengths 59
9	Comparison of MEDEVAC busy probabilities..... 74
10	Comparison of zone-precedence queue lengths 75
11	MEDEVAC and MTF locations with Casualty Cluster Centers..... 78
12	Sampled Casualty Events in Helmand, Kandahar, and Zabol 79
13	Policy Iteration vs Linear Programming 85
14	Policy Iteration vs Linear Programming 86

List of Tables

Table		Page
1	9-Line MEDEVAC Request Proportions by Zone-Precedence Level	46
2	Expected Response Times (minutes)	49
3	Expected Service Times (minutes)	50
4	Immediate Expected Rewards	51
5	2×2 Case Parameters	52
6	Comparison of Dispatching Policies for Scenario 1	53
7	Comparison of Dispatching Policies for Scenario 2	54
8	Comparison of Dispatching Policies for Scenario 3	54
9	Comparison of Combined Average MEDEVAC Utilization	58
10	Comparison of Total Expected Discounted Rewards & Optimality Gaps.....	59
11	2^5 Full Factorial Screening Experimental Factor Levels.....	60
12	2^5 Full Factorial Screening Experiment Results	61
13	Multiple Linear Regression Analysis	62
14	3^3 Full Factorial Experimental Factor Levels.....	63
15	3^3 Full Factorial Experimental Results	64
16	Multiple Linear Regression Analysis	64
17	Comparison of Dispatching Policies for Scenario 1	65
18	Comparison of Average λ Performance Metrics	66
19	Comparison of Average δ Performance Metrics	66
20	Comparison of Average p_{k_1} Performance Metrics	66

Table		Page
21	Comparison of MEDEVAC Dispatching Policies for Priority Requests	68
22	MEDEVAC Helicopter Flight Speed Analysis	69
23	Comparison of Intra-Zone Dispatching Policies for Scenario 1	71
24	Comparison of Intra-Zone Dispatching Policies for Scenario 2	71
25	Comparison of Intra-Zone Dispatching Policies for Scenario 3	72
26	Comparison of Intra-Zone Dispatching Policies for Scenario 4	73
27	Comparison of Total Expected Discounted Rewards & Optimality Gaps.....	76
28	3×3 Case Parameters	77
29	Proportions of Zone-Precedence Level 9-Line MEDEVAC Requests.....	78
30	Expected Response Times (minutes)	80
31	Expected Service Times (minutes)	80
32	Immediate Expected Reward	80
33	Comparison of Dispatching Policies for Scenario 5	81
34	Comparison of Dispatching Policies for Scenario 6	82
35	Comparison of Dispatching Policies for Scenario 7	83
36	Comparison of Dispatching Policies for Scenario 8	84
37	Comparison of Total Expected Discounted Rewards & Optimality Gaps.....	84
38	Policy Iteration versus Linear Programming Computational Efficiency (s)	85

USING MARKOV DECISION PROCESSES WITH HETEROGENEOUS QUEUEING SYSTEMS TO EXAMINE MILITARY MEDEVAC DISPATCHING POLICIES

I. Introduction

The primary objective of a deployed military emergency medical services (EMS) system is to successfully evacuate casualties from the battlefield in a timely manner. Casualty evacuation (CASEVAC) and medical evacuation (MEDEVAC) are the two main options available for transporting combat casualties to a medical treatment facility (MTF). CASEVAC refers to the transport of casualties to an MTF via non-medical vehicles or aircraft *without* en route medical care by onboard medical professionals. Casualties transported via CASEVAC may not receive the necessary medical care or be transported to the appropriate MTF. As such, MEDEVAC is the more preferred and primary method of transporting combat casualties. MEDEVAC refers to the transport of casualties to an appropriate MTF via standardized medical evacuation platforms with onboard medical professionals who are equipped to provide en route medical care and emergency medical intervention (Department of the Army, 2014).

While MEDEVAC operations utilize several different types of evacuation platforms, this thesis focuses on the aerial aspect of MEDEVAC operations (i.e., aeromedical helicopter operations). Helicopters have the capability and flexibility to fly directly to a predetermined casualty collection point (CCP), meeting battlefield casualties when they are at their most vulnerable and critical stages, landing in an area where no other platform (e.g., ground vehicle or fixed-wing aircraft) could, or utiliz-

ing a rescue hoist to lift casualties to the helicopter. After securing the casualties, helicopters can fly directly to dedicated trauma centers or hospitals unencumbered by roads with speeds often exceeding 150 miles per hour, all while providing definitive en route care from trained and highly skilled medics (O'Shea, 2011). These helicopter capabilities greatly contribute to recent increases in casualty survivability rates.

Helicopter ambulances were first introduced in the military during the Korean conflict and immediately became a high visibility asset of the MEDEVAC system. By the end of the Vietnam War, the capabilities (i.e., speed and versatility) of helicopters in austere conditions far exceeded the capabilities of ground platforms. The ability to travel across terrain in remote areas not accessible to ground vehicles makes helicopters well suited for MEDEVAC operations (De Lorenzo, 2003; Clarke & Davis, 2012). The United States Army operates HH-60M helicopters specifically designed for the MEDEVAC mission. HH-60M helicopters come equipped with the necessary resources (e.g., oxygen generator, integrated EKG machine, electronically controlled litters, built-in external hoist, and an infrared system that can locate patients by their body heat) to give medical personnel the ability to simultaneously treat and transport casualties from a CCP to an appropriate MTF. The urgency of the MEDEVAC mission is critical to the survivability of battlefield casualties and the HH-60M helicopter has proved to be advantageous to the Army with its ability to launch in less than seven minutes (O'Shea, 2011). Eastridge *et al.* (2012) report that the survivability of combat casualties has continued to increase over time since World War II (WWII). Approximately 80% of casualties occurring on the battlefield survived in WWII, while 84% survived during the Vietnam War. An increase to 90% casualty survivability was observed in the continuous decade of war between 2001-2011. The improved casualty rates are attributed to improvements in the versatility and speed of MEDEVAC helicopters and the resulting decrease in the time required for casualties

to receive proper medical care (De Lorenzo, 2003).

Military medical planners are responsible for designing deployed MEDEVAC systems. An effective and efficient MEDEVAC system boosts *esprit de corps* of deployed military personnel, who understand that rapid and quality care will be provided if they are injured in combat (Department of the Army, 2014). Important decisions include determining where to locate MEDEVAC units and MTFs, identifying a MEDEVAC dispatching policy, and recognizing when redeployment of aeromedical helicopters is necessary and possible. The location of MEDEVAC units is usually determined while considering two objectives: maximizing coverage and minimizing response time subject to logistical, resource, and force protection constraints. Deciding which MEDEVAC unit to dispatch to a given service request is a vital aspect of any EMS, including a MEDEVAC system, and is the primary focus of this thesis. The military often defaults to a myopic dispatching policy wherein the closest available MEDEVAC unit is dispatched to retrieve combat casualties from a CCP regardless of the request's evacuation precedence category (e.g., Priority I - *Urgent*, Priority II - *Priority*, and Priority III - *Routine*). Redeployment of MEDEVAC units prior to returning to their originating base is possible but poses challenges due to the numerous resource and availability requirements (e.g., refueling, resupply, and armed escort). These reasons also make temporary relocation of idle MEDEVAC units uncommon within a theater of operations (Rettke *et al.*, 2016).

This thesis examines the MEDEVAC dispatching problem wherein a dispatching authority must decide which MEDEVAC unit to dispatch to a particular 9-line MEDEVAC request. The location of MTFs and MEDEVAC assets are known and all MEDEVAC helicopters are assumed to have the capability to meet the mission requirements of any 9-line MEDEVAC request. Redeployment is not considered. The reported dispatch policy is based on the location and status of MEDEVAC units, the

location of the casualty event, and the evacuation precedence category of the casualty event.

An infinite horizon, discounted Markov decision process (MDP) model is formulated to determine how to optimally dispatch MEDEVAC helicopters to casualty events occurring in combat to maximize the expected total discounted reward attained by the system. A computational example is applied to a MEDEVAC system in Afghanistan in support of combat operations. Comparisons are made between the myopic policy that is typically utilized in practice and the optimal policy derived from the formulated MDP model.

An important difference between this thesis and other papers in this research area is the incorporation of admission control and queueing. Admission control allows the dispatching authority to observe the current state of the MEDEVAC system before making the decision to accept or reject an incoming request. This gives the dispatching authority the power to reject incoming requests, reserving MEDEVAC units for higher precedence requests instead of satisfying all requests for service. The rejected requests are not simply discarded; rather, they are redirected to another servicing agency to be serviced (i.e., CASEVAC). If the dispatch authority allows a request to enter the MEDEVAC system but all MEDEVAC units are currently servicing other requests, the entering request will be allocated to a queue based on its precedence level and zone. Once a request has entered the system, it will be serviced; however, the dispatching authority dictates which available MEDEVAC unit will service each request in the system, regardless of when the request entered the system. For example, an *urgent* request will be serviced before a *routine* request regardless of the order in which they entered the system. It is important to note that MEDEVAC units will not interrupt service to a request in the case of a higher precedence request arriving. Once a MEDEVAC unit is assigned a specific request,

it will be considered unavailable until it completes the service of that request.

The remainder of this thesis is organized as follows: Chapter II provides a review of research relating to MEDEVAC systems, Chapter III presents a description of the MEDEVAC dispatching problem, Chapter IV describes the MDP formulation developed to determine an optimal MEDEVAC dispatch policy, and Chapter V covers an application of the formulated MDP based on a representative scenario in southern Afghanistan. Chapter VI concludes the thesis and proposes several directions for future research.

II. Literature Review

For nearly half a century, research has been conducted on optimizing both civilian and military emergency medical services (EMS) response systems. The main features of this research include determining the location of servers; dictating the number of servers per location, the server dispatch policy, and the size and number of response zones (if a partitioning strategy for the service area is implemented); identifying which performance measure to focus on as the objective: response time thresholds (RTTs) or patient survivability rates; and recognizing if and when server relocation is necessary due to either a service completion or an incoming service request. Another complicating feature concerns the location of hospitals. In research examining civilian EMS systems the locations of hospitals are usually given as fixed; however, in some military planning contexts the medical treatment facility (MTF) locations are not given. Military medical planners must decide where to best place MTF locations when designing a military medical evacuation (MEDEVAC) system (Rettke *et al.*, 2016). Operations research (OR) methods have been a popular choice amongst researchers when examining EMS systems. Applied OR methods include stochastic modeling, queueing, discrete optimization, and simulation modeling (Green & Kolesar, 2004).

Research on EMS operations can be traced back to the late 1960s and early 1970s. The research conducted in this field primarily focuses on the civilian sector and examines characteristics such as the optimal location (Bianchi & Church, 1988; Daskin & Stern, 1981; Jarvis, 1975), allocation (Berlin & Liebman, 1974; Baker *et al.*, 1989; Hall, 1972), dispatch (Ignall *et al.*, 1982; Swersey, 1982; Green & Kolesar, 1984), and relocation of emergency vehicles (Berman, 1981; Kolesar & Walker, 1974; Chaiken & Larson, 1972) to enhance the performance of the EMS system. While the goal of most OR research is to aid decision makers, implementing published models does not occur as frequently as one might hope. However, this does not seem to be the case

with emergency response systems research. Green & Kolesar (2004) give an account of how emergency service management research has impacted emergency response systems. Despite the substantial amount of research conducted on improving the performance of civilian EMS systems, little research exists seeking to improve the performance of military EMS (i.e., MEDEVAC) systems.

The research presented in this thesis examines the optimal dispatch of military EMS vehicles (i.e., HH-60M MEDEVAC helicopters) to prioritized requests for service. Consideration of the precedence category (e.g., Priority I - *Urgent*, Priority II - *Priority*, and Priority III - *Routine*) is important. A substantial amount of research seeks to improve the overall performance of EMS system, but most research endeavors do not account for the precedence of the call (Bandara *et al.*, 2014). When the precedence of the call is not considered, the default dispatching rule sends the closest available emergency response vehicle to satisfy required service requests with no regard as to how that specific vehicle’s absence impacts the overall EMS system. Sending the closest available vehicle to a service request regardless of other factors (e.g., precedence, or severity) is commonly referred to as a myopic policy. Many researchers (Carter *et al.*, 1972; Nicholl *et al.*, 1999; Kuisma *et al.*, 2004) show that myopic policies tend to be suboptimal. Incorporating precedence categories into the construction of dispatching policies can ultimately lead to more lives being saved on the battlefield.

Unlike previous work in this area, admission control and prioritized queueing are explicitly accounted for when formulating the Markov decision process (MDP) model of the dispatching problem. Descriptive queueing systems model a wide range of phenomena and are quite effective in predicting and evaluating the performance of an existing system (Stidham & Weber, 1993). The formulation and analysis of queueing system models help improve the design of the system being studied. Controlled

queueing systems consist of three components: controllers, queues, and servers. The absence of a system controller (i.e., the MEDEVAC dispatching authority) can lead to erratic system behavior with periods of long queues followed by periods where servers remain idle (Puterman, 1994). Admission control allows the system controller to observe the current state of the system when a call (i.e., a 9-line MEDEVAC request) arrives and on this basis decide whether to admit the call to the eligible job queue. Admission control offers the possibility of significantly improving performance as compared to state-independent rules (Efrosinin, 2004). If a call is admitted, it will eventually receive service while those rejected never enter the system. Queueing models have been utilized in a variety of applications. See Stidham (2002) for a survey of work that has significantly contributed to the queueing theory field and see Stidham & Weber (1993) for a survey of numerous models for the optimal control of networks of queues with a focus on optimal control policies and Markov decision theory.

Typically, optimization problems for controlled queueing systems are easier to handle when they are modeled in discrete time rather than continuous time (Efrosinin, 2004). Uniformization can be applied to a continuous-time MDP (CTMDP) model to obtain a model with constant transition rates so that results and algorithms for discrete-time discounted models may be applied directly (Puterman, 1994). More details on how the MEDEVAC system can be converted from a continuous-time problem to a discrete-time problem will be discussed in the methodology section. The process of converting continuous-time problems to discrete-time problems has been well established and can be seen in the works of Rosberg *et al.* (1982), Lippman (1975), and Serfozo (1979). It is also important to recognize that the optimal policies resulting for continuous and discrete time problems are equivalent (Puterman, 1994). The controlled queueing system presented in this thesis is comprised of prioritized queues

with heterogeneous servers modeled as a MDP due to the appropriate choices for the control sets and state spaces. The MEDEVAC units are considered heterogeneous servers due to the different service/response times of incoming 9-line MEDEVAC requests. For example, a MEDEVAC unit will have a different service/response time for a zone one 9-line MEDEVAC request as compared to a zone four 9-line MEDEVAC request.

Another key feature in EMS system research is the optimality criterion. The optimality criterion for the dispatching problem is based on the selection of the performance measure. It is important to select an appropriate EMS performance measure because it dictates how the EMS system's resources are utilized and hence directly impacts the patient survivability rate. The vast majority of EMS systems measure performance according to an RTT (McLay & Mayorga, 2010). RTT is commonly referred to as the number (or fraction) of calls that can be serviced within a predetermined and fixed time frame. A call must be serviced within its stated RTT to be considered covered. RTTs are usually preferred over other types of measures related to the outcome of a patient because they are easier to evaluate and the data is readily available. There is not an officially adopted standard for RTT, but most urban areas in civilian EMS systems require calls to be serviced within eight minutes and fifty-nine seconds (8:59) with at least a 90 percent compliance rate (Fitch & Griffiths, 2005). That is, an EMS system must respond to at least 90 percent of service requests within the given RTT of 8:59. Williams (2005) showed that of the 200 most populated cities in America, over three quarters of civilian EMS system respondents follow a standard of 8:59 or less. Unfortunately, a military EMS system would not be able to respond to *urgent* requests within 8:59 due to the dispersed disposition of forces in combat, distances that must be traveled, and inherent combat environment.

In 2009, Secretary of Defense Robert Gates mandated that the United States

MEDEVAC system follow what is colloquially known as the Golden-hour Rule. The golden-hour rule requires delivery of battlefield casualties to an appropriate MTF within one hour of a 9-line MEDEVAC request (Olson *et al.*, 2013). These RTTs are often employed as system performance measures for life-threatening (i.e., *urgent*) calls in both the civilian and military EMS systems, respectively. While RTTs seem to have many benefits, one common criticism relates to how well patient survivability rates (the underlining measure to be maximized in EMS systems) are captured when utilizing RTTs. For example, according to the commonly used civilian EMS system RTT of 8:59, a call is considered to be covered if the response time is within 8:59, but any response time greater than 8:59 (e.g., nine minutes) would not be considered covered. Fitch (2005) suggests that there is not a statistically significant difference in casualty survivability rates between these cases.

Another performance measure that has been utilized for the optimality criterion is patient survivability rates. Recent research suggests that performance measures based on patient survivability provide better results when compared to RTTs (Pons & Markovchick, 2002; Knight *et al.*, 2012; Erkut *et al.*, 2008). However, estimating patient survivability tends to be a difficult task due to the lack of available data (McLay & Mayorga, 2010). Another challenge associated with patient survivability is defining when it actually occurs. For battlefield casualties, a casualty is usually considered “survived” once the individual is discharged from the military medical system. The problem with this definition is that a casualty may not be discharged for several months and can transfer to different medical facilities and locations while being treated, making the task of tracking casualty survivability tedious and difficult (Rettke *et al.*, 2016). Even with these challenges, many researches (McLay & Mayorga, 2010; Bandara *et al.*, 2012; Mayorga *et al.*, 2013; Bandara *et al.*, 2014) utilize patient survivability as the performance measure in EMS systems. Their results sug-

gest that utilizing patient survivability is better suited for determining the number of patients that survive and ultimately helps in increasing the survivability of patients.

As such, one of the objectives of this thesis is to implement optimal dispatching policies for MEDEVAC systems that maximize the probability of casualty survivability with the inclusion of the degree of severity (e.g., *urgent*, *priority*, and *routine*) of the request. Similar to Erkut *et al.* (2008), this thesis applies a survivability function that is monotonically decreasing in response time to model the outcome of casualties. As noted before, one of the primary challenges of using patient survivability as the performance measure is obtaining empirical data to support the functional form. Research conducted by Eastridge *et al.* (2012) gives an extensive account of statistics associated with combat related deaths, but unfortunately the response times related to the deaths are not documented. The lack of response time data in Eastridge *et al.* (2012) research makes it unlikely to develop a survivability function that has a high level of confidence. Although Feero *et al.* (1995) give an account of EMS response times in relation to trauma patients to study how they affect survivability, their research focuses on civilian EMS systems wherein response times are typically under eight minutes. The time it takes MEDEVAC units to transport battlefield casualties to an appropriate MTF is typically much longer than civilian EMS response times due to the fact of MEDEVACs having to travel significantly further than civilian EMS units (Rettke *et al.*, 2016). The current MEDEVAC response time goal, as mandated by Secretary of Defense Gates in 2009, is to successfully respond and transport an *urgent* 9-line MEDEVAC request to the necessary medical facility within 60 minutes of being notified of the 9-line MEDEVAC request (Garrett, 2013). EMS systems often do not consider more than three precedence categories due to the fact that these classifications need to be made in a matter of seconds (Bandara *et al.*, 2012). This thesis focuses on the three primary evacuation precedence categories as applied in the

United States Army: *urgent*, *priority*, and *routine*. The RTTs for these evacuation precedence categories (i.e., 60 minutes, 240 minutes, and 24 hours, respectively; Department of the Army (2014)) are utilized in the development of the casualty outcome functions.

Bandara *et al.* (2012) describe research where the probability of patients surviving in EMS systems is greatly enhanced if the precedence category is considered when deciding which emergency response vehicle to dispatch. A discounted, infinite horizon MDP model is formulated and analyzed by Bandara *et al.* (2012) in which two types of calls (i.e., life-threatening and non life-threatening) are prioritized according to the urgency of the call. The results indicate sending the closest unit available, regardless of call precedence, is not always optimal. The analysis recommends sending the closest available (i.e., idle) unit when life-threatening calls are submitted and the next closest unit when non life-threatening calls are submitted, regardless of the order the calls arrived. The optimal policy for life-threatening calls is intuitive because faster response times result in a higher probability of patient survivability. An ordered list of which units to dispatch is created for non life-threatening calls. This study highlights that an optimal dispatching policy may recommend sending a more distant vehicle to service a less urgent call if closer units are more likely to receive life-threatening calls. This policy essentially rations closer units in anticipation of a more urgent request. Increasing the number of zones and EMS units may make the results less intuitive, but EMS systems still can benefit from the implementation of an optimal policy versus a myopic approach. It is observed that many lives can be saved without increasing the cost by implementing the optimal policy. Bandara *et al.* (2014) also consider the severity level of incoming calls when implementing dispatch policies. The authors develop a simulation model to evaluate how each dispatch policy affects the overall performance of EMS systems. Their model also measures performance in terms of the

probability of patient survivability because it more accurately reflects the outcome of patients. Several examples with different response strategies are evaluated, and the recorded results are similar to those found in Bandara *et al.* (2012), which indicate that dispatching the closest vehicle is not always the optimal action. Bandara *et al.* (2014) find that dispatching vehicles based on the urgency of the calls ultimately leads to an increase in the average survival probability of patients. Utilizing these results, the authors develop an easy-to-implement heuristic algorithm that can be applied to large-scale EMS systems.

Mayorga *et al.* (2013) also examine dispatching policies for EMS systems wherein the performance is measured in terms of patient survival probability. Before comparing the performance of different dispatching policies via a simulation model, the authors determine the number, size, and location of response districts by utilizing a constructive heuristic that incorporates adjusted expected coverage. Their research is the first to address the joint problem of finding appropriate dispatching decisions and response districts for both intra-district and inter-district situations. An intra-district policy refers to how calls are managed when there is at least one available emergency unit within the district, whereas an inter-district policy refers to how calls should be answered in the event that no emergency units are available within the district at the time the call occurs. Two types of dispatching policies are considered for intra-district situations: a myopic policy (i.e., the closest available vehicle services the call) and a heuristic policy developed by Bandara *et al.* (2014). While myopic policies are generally practiced by many EMS systems, the heuristic policy Bandara *et al.* (2014) developed helps balance the workload of emergency units and incorporates the urgency of calls when making dispatching decisions, which has been proven to increase patient survivability rates. Two different policies are considered for the inter-district situations. The first policy assumes that EMS resources (e.g.,

fire engines, ambulances, and police) from other counties will assist in servicing calls when all available ambulance units within the district are unavailable. The second policy is to send ambulances from other districts to service calls when all ambulances within the district are busy. The second policy utilizes a preference list of ambulances to cross districts and is constructed by applying the heuristic proposed by Bandara *et al.* (2014). The results from this work indicate that integrated districting and dispatching policies are a vital aspect in increasing the probability of survivability for patients.

McLay & Mayorga (2013b) formulate an MDP model to determine how to dispatch EMS units to requests categorized by an evacuation precedence in an optimal manner given that dispatch authorities make errors in correctly categorizing the true urgency of each request. Unlike Mayorga *et al.* (2013) and Bandara *et al.* (2012), McLay & Mayorga (2013b) focus on the evacuation precedence of patients with an objective of maximizing the long-run average utility of the system while considering the possibility of patient classification errors. The authors utilize an RTT as the optimality criterion versus a performance measure based on patient survivability. They also consider over-responding and under-responding to perceived patient risk when classification errors exist. McLay & Mayorga (2013b) find that dispatching the closest ambulance to service incoming calls, regardless of the call precedence, is not always best. The authors also note that over-responding is preferred when there is a high rate of classification errors while under-responding is preferred when there is a low rate of classification errors. McLay & Mayorga (2013a) propose a constrained variant of the Markov decision problem introduced in McLay & Mayorga (2013b) and formulate an equity-constrained linear programming model to solve the constrained problem. The authors examine how dispatching strategies impact server-to-customer systems (i.e., an EMS system) given a set of equity constraints. Four separate equity measures

are considered, two of which consider equity from the server perspective and two of which consider equity from the customer perspective. Their objective is to determine an optimal dispatching policy for balancing equity and efficiency when dispatching distinguishable servers to prioritized customers in service-to-customer systems that maximizes the long-run average customer utility. Results indicate that when either the equity of servers or customers is considered then the equity for both customers and servers is simultaneously improved.

EMS research exists that focuses specifically on military MEDEVAC systems. Zeto *et al.* (2006) develop a goal programming model that seeks to maximize the aggregate expected demands covered and minimize the spare capacities of air ambulances. The authors leverage Alsalloum & Rand (2006), examining both the problems of resource allocation and coverage in a three-phased approach. In the first phase, they characterize the demand for MEDEVAC missions using a multivariate hierarchical cluster analysis. In the second phase, they then estimate the parameters of the model via a Monte Carlo simulation. In the third phase, they utilize a bi-criteria model to emplace the minimum number of required aircraft at each location to maximize the probability of meeting the MEDEVAC demand in the Afghanistan theater. Bastian *et al.* (2012) investigate the capabilities required for MEDEVAC aircraft platforms to successfully perform the necessary duties and provide coverage within a brigade operating space. The authors develop a decision support tool that military medical planners can utilize to analyze the risk associated with different MEDEVAC strategies. Bouma (2005) develops a MEDEVAC and treatment capability optimization model that assists in the redistribution, realignment, and restructuring of medical materials and resources to help meet requirements in the area of operations. Fulton *et al.* (2009) evaluate the planning factors and rules of allocation associated with Army air ambulance companies. Military medical planners typically use the rules of

allocation, which are based on strategic planning documents, to estimate the number of MEDEVAC units required for tactical and operational scenarios. The authors quantitatively analyze different rules through a Monte Carlo simulation and record the impact that they have on major combat operations. The results indicate that 0.4 aircraft per admission would be a reasonable planning factor. Finkbeiner (2013) proposes a hybrid discrete-event simulation and queueing approach to identify the minimum number of aircraft needed to reach a predetermined level of aeromedical evacuation. An integer programming model is subsequently utilized to determine where to locate helicopters within the area of coverage. Sundstrom *et al.* (1996) incorporate linear programming techniques to develop a model based on the probabilistic location set-covering problem that provides the required numbers of MEDEVAC assets needed as well as the optimal positioning of those assets to ensure orderly transport of battlefield casualties to an appropriate medical facility.

The allocation of MEDEVAC units during steady-state combat operations is studied by Fulton *et al.* (2010) and Bastian (2010). Fulton *et al.* (2010) formulate a stochastic optimization model that manages the locations of deployable military hospitals, hospital beds, and both aerial and ground MEDEVAC units prior to the reception of a 9-line MEDEVAC request. Their model uses an objective of minimizing the total travel time, which is weighted by the urgency level of the casualty, from the POI to an appropriate MTF. The weights associated with the urgency levels of casualties are derived from historical data of patient injury severity scores collected from Operation Iraqi Freedom (OIF) combat operations. Bastian (2010) formulates a stochastic optimization goal programming model to meet three separate objectives: maximize the coverage of theater-wide casualty demand in Afghanistan, minimize the spare capacity of MEDEVAC units, and minimize the maximal MTF evacuation site vulnerability to enemy attack.

Keneally *et al.* (2016) examine MEDEVAC dispatch policies in the Afghanistan theater via an MDP model. The authors assume that each service call arrives sequentially and the locations of each service center are predetermined. Their work classifies service calls into three evacuation precedence categories: *urgent*, *priority*, and *routine*. They consider the possibility that an armed escort may be required to accompany the MEDEVAC unit. The authors utilize a reward function based off of RTT and conduct computational experiments wherein MEDEVAC units operate in support of Operation Enduring Freedom (OEF). The results highlight that the myopic policy (i.e., the default policy in practice) does not always lead to the optimal dispatching strategy. Grannan *et al.* (2015) develop a binary linear programming (BLP) model to determine where to locate and how to dispatch multiple types of military MEDEVAC air assets. A spatial queuing approximation model provides inputs to the BLP model. The BLP model incorporates the precedence of each service call to maintain a high likelihood of survival for the most urgent casualties. The overall objective is to maximize the proportion of high-precedence calls responded to within a pre-determined RTT.

Rettke *et al.* (2016) formulate an MDP model to examine the MEDEVAC dispatching problem. The problem instance size in this study is too large for an exact dynamic programming solution model, so the authors employ approximate dynamic programming (ADP) techniques to determine an optimal dispatch policy. The computational experiments in this study indicate that their ADP generated policy is nearly 31% better than the myopic policy. Military medical planners can use these results to improve existing MEDEVAC tactics and techniques. Lejeune & Margot (2016) propose a MEDEVAC model that considers endogenous uncertainty in the delivery times of casualties. The objective of their model is to provide prompt medical treatment and evacuation to soldiers injured in combat. The model determines where

to locate MEDEVAC units and MTFs. Moreover, it helps the dispatch authority in determining which helicopters to dispatch and which MTF each call should report to. Results indicate a reduction in battlefield deaths due to an increase in timely treatment to combat casualties when compared to a myopic policy.

III. Problem Description

One of the primary missions of the Army Health System (AHS) is to provide medical evacuation (MEDEVAC) across a wide range of military operations. The dedicated Army helicopters (i.e., rotary-wing aircraft or air ambulances) utilized in MEDEVAC missions are under the command of the general support aviation battalion (GSAB). Any use of air ambulances must first be coordinated with the supporting GSAB to ensure synchronized evacuation procedures are executed. The GSAB manages all activities related to the execution of aerial operations and serves as the primary decision-making authority for the military MEDEVAC system (Department of the Army, 2014). An Army aeromedical evacuation officer (AEO) that works within the GSAB acts as the MEDEVAC dispatching authority in a deployed military emergency medical service (EMS) system (Fish, 2014). AEOs direct the use of medical aircraft, personnel, and equipment in support of operational and strategic medical evacuations within a theater of operations.

When a casualty event occurs and a 9-line MEDEVAC request is submitted, the AEO must make a decision quickly as to which MEDEVAC unit (if any) to dispatch. The casualty survivability rate will decrease if there are delays in decision making. To complicate matters further, there are many situations where MEDEVAC units require a team of armed helicopters to escort them to the casualty site due to high threat level conditions (e.g., enemy troops in the area). Armed escort requirements can potentially increase the overall response time, which ultimately decreases the chances of casualties surviving. Therefore, it is vital that the GSAB implements a dispatching policy that results in rapid and quality transport of life-threatening battlefield casualties from the point-of-injury (POI) to the nearest, most appropriate MTF. The procedures outlined in the Army's Medical Evacuation Field Manual (Department of the Army, 2014) and the graphical representations that Keneally *et al.* (2016) and Rettke *et al.*

(2016) offer in their problem descriptions are utilized as a basis for the MEDEVAC mission timeline depicted in Figure 1.




Figure 1. MEDEVAC Mission Timeline

A 9-line MEDEVAC request is transmitted in a standardized message format with a prescribed amount of information that helps expedite the process of transporting casualties. When a 9-line MEDEVAC request is determined to be necessary, it should be transmitted over a secure communication system via a dedicated frequency. However, a 9-line MEDEVAC request can still be transmitted if no secure communication systems are available. In wartime conditions, the information required in a 9-line MEDEVAC request is reported in the following order: the location of the pickup site (i.e., POI or casualty collection point (CCP)), radio frequency and call sign, number of casualties by precedence, special equipment required, number of casualties by type, security of pickup site, method of marking pickup site, casualty nationality and status, and chemical, biological, radiological, and nuclear contamination. Either the senior military member or the senior medical person (if available) at the scene identifies

the evacuation precedence category of each casualty and determines whether a 9-line MEDEVAC request is necessary. The tactical situation and the condition of each casualty are taken into consideration when making this decision. The overall precedence of a 9-line MEDEVAC request is based on the most time sensitive precedence of the casualties. Correct category placement is vital and should not be overemphasized because it may burden the evacuation system due to aerial ambulances being a low-asset, high-demand resource that must be managed accordingly. The United States Army utilizes the following evacuation precedence categories when prioritizing casualties that require medical evacuation (Department of the Army, 2014):

1. **Priority I, Urgent:** Assigned to emergency cases that should be evacuated as soon as possible and within a maximum of 1 hour in order to save life, limb, or eyesight, to prevent complications of serious illness, or to avoid permanent disability.
2. **Priority II, Priority:** Assigned to sick and wounded personnel requiring prompt medical care. This precedence is used when the individual should be evacuated within 4 hours or his medical condition could deteriorate to such a degree that he will become an URGENT precedence, or whose requirements for special treatment are not available locally, or who will suffer unnecessary pain or disability.
3. **Priority III, Routine:** Assigned to sick and wounded personnel requiring evacuation but whose condition is not expected to deteriorate significantly. The sick and wounded in this category should be evacuated within 24 hours.

In a combat situation, requests for MEDEVAC units are typically made at the POI once enemy fire has been suppressed. MEDEVAC requests are transmitted through several layers of command before reaching an AEO working within the GSAB at higher headquarters. The specific information flow depends on the communication infrastructure within the command, the communication equipment available to the

requesting unit, and the command and control organization of the MEDEVAC system (Rettke *et al.*, 2016). Once the request has been made, casualties are transported to a CCP, which is a predesignated point along the evacuation route for collecting the wounded (Department of the Army, 2000). The time at which the MEDEVAC request reaches the AEO is denoted by T_1 .

Once the GSAB receives the 9-line MEDEVAC request, the AEO must then decide whether to immediately assign a MEDEVAC unit to the request, depending on any pre-existing requests in the MEDEVAC system, the location of the pick-up site, the number and precedence of the casualties, and the status of the MEDEVAC units. If the MEDEVAC system is burdened with a high number of requests, the AEO may reject the incoming request from entering the system and redirect the request to be handled by casualty evacuation (CASEVAC). Assuming the request enters the system, the AEO will wait for a suitable MEDEVAC to become available. The AEO assigns the MEDEVAC unit to the request at time T_2 along with an armed escort, if required, to service the request.

The amount of time between an AEO receiving the 9-line MEDEVAC request, T_1 , and the assignment of the MEDEVAC unit, T_2 , is the total wait time for the request in the MEDEVAC system. The wait time comprises the time required to determine which MEDEVAC unit to dispatch; whether an armed escort is required; which armed escort team to assign, if required; and the time required to transmit the request information to the assigned MEDEVAC unit and armed escort team, if required.

As stated earlier, once a 9-line MEDEVAC request is received by the GSAB, the AEO must decide whether the request should enter the MEDEVAC system or if the request should be serviced by another organization (i.e., CASEVAC). If the AEO allows the request to enter the MEDEVAC system and at least one suitable

MEDEVAC unit is available to service the request, another decision must be made as to whether the request should be assigned immediately or if the request should be placed in a queue based on the evacuation precedence category and location (i.e., zone) of the request. If the AEO allows a request to enter the MEDEVAC system and no suitable MEDEVAC units are available to service the request, then the request is placed in its respective zone-precedence queue. Figure 2 depicts the multiple-server, multiple-buffer queueing model employed in this thesis. The MEDEVAC queueing system represented in Figure 2 visually depicts the wait time between points T_1 and T_2 in Figure 1.




Figure 2. MEDEVAC Queueing System

Decision epochs occur when a 9-line MEDEVAC request is received by the GSAB or when a MEDEVAC unit completes a service request and becomes available. When a 9-line request is submitted and received by the GSAB, the AEO's decision consists of sending the just-arrived 9-line MEDEVAC request to its respective zone-precedence queue (if the queue is not full), immediately assigning an available MEDEVAC unit to service the request, or rejecting the request from ever entering the system. Once a MEDEVAC unit reaches service completion and at least one of the zone-precedence queues is not empty, the AEO must make a decision. The AEO's decision consists of either assigning a queued 9-line MEDEVAC request to one of the idle MEDEVAC

units or waiting for another (possibly higher precedence) request to enter the system or another MEDEVAC unit to reach service completion.

The information from the 9-line MEDEVAC request is transmitted to the assigned MEDEVAC unit through the command's communication system. T_3 denotes the time at which the assigned MEDEVAC unit departs its station for the CCP. The amount of time between the MEDEVAC unit being assigned the 9-line MEDEVAC request, T_2 , and the MEDEVAC unit departure, T_3 , is the total mission preparation time, which includes preparing the medical equipment, medical personnel, and helicopters for the MEDEVAC mission. Typically, if an armed escort is required, it will take off with the MEDEVAC unit at the staging area, but there are situations where the MEDEVAC unit must meet an armed escort at a predetermined rally point en route to the CCP. The MEDEVAC unit cannot land at a high threat level CCP site without an armed escort, which could lead to an increased total response time.

T_4 denotes the time at which the MEDEVAC unit lands at the CCP site. Upon arrival to the CCP site, the MEDEVAC unit immediately begins initial treatment and loads casualties. T_5 denotes the time at which the MEDEVAC unit departs the CCP site and proceeds towards an MTF. The destination MTF is selected in a deterministic manner based on the location of the CCP site. The MTF that is located closest to the CCP site is the one that the MEDEVAC unit departs to at time T_5 .

The MEDEVAC unit arrives at the MTF site at time T_6 . After arriving, the MEDEVAC unit immediately begins to unload casualties and transfers the responsibility of subsequent care of the casualties to the medical staff at the MTF. After all casualties have been unloaded, the MEDEVAC unit departs the MTF and travels back to its own staging area. Once a MEDEVAC unit has finished unloading and transferring the subsequent care of casualties to the MTF medical staff, it must return to its own staging area before being tasked to service another 9-line MEDEVAC

request. This requirement comes from concerns about low fuel levels, crew bed down limitations, on-board equipment configurations, and other logistical issues (Rettke *et al.*, 2016). Typically, MEDEVAC units need to return to their home staging areas to refuel before being dispatched for another mission. T_7 denotes the time at which the MEDEVAC unit departs the MTF.

The MEDEVAC unit arrives back at its staging area, completes refueling, and is staged for future missions at time T_8 . Once the MEDEVAC unit arrives back at its staging area the mission is considered complete. The MEDEVAC unit then becomes available for dispatch to another 9-line MEDEVAC request.

It is important to note that battlefield conditions (e.g., enemy disposition, required equipment being transported, weather conditions, and the air density due to flight altitude) are expected to affect the travel times from the MEDEVAC staging area to the CCP site, from the CCP site to the selected MTF location, and from the MTF location back to the MEDEVAC staging area.

Military medical planners must consider the measurement of MEDEVAC system performance when considering dispatch policies. In civilian operations, the efficacy of EMS systems has been a difficult area to evaluate due to the multitude of variables present (MacFarlane & Benn, 2003). The search for a reliable measure of performance remains a topic of interest in the EMS field (McLay & Mayorga, 2010). Practitioners and researchers employ various means of assessment. The most common method for evaluating EMS systems utilizes ambulance response times. EMS systems commonly define the response time as the time required to reach the patient after receiving the emergency call. Since EMS systems are evaluated on response time, one of their primary focuses is the rapid response to cardiac arrest situations. This emphasis exists because the ability to provide effective treatment to patients undergoing cardiac arrest is time-sensitive.

Another reason behind this rationale is as follows. If the EMS system has the capability to respond quickly to cardiac arrest patients, then it is more likely to be able to service similar life-threatening situations. Therefore, defining the response time for a civilian EMS system to be the time between receiving the emergency call and the time the first emergency response vehicle arrives on scene is quite intuitive.

Nonetheless, MEDEVAC system performance cannot be measured using the same evaluation criteria as the civilian EMS system. Several additional factors complicate the medical evacuation of a casualty from a battlefield. The travel times, load times, and unload times can be much greater and vary significantly more in military EMS systems when compared to a civilian EMS system. Moreover, the primary cause of death for battlefield casualties is blood loss, not cardiac arrest. Garrett (2013) indicates that blood loss is the primary cause of death for nearly 85% of soldiers killed in action (KIA). Due to this issue, some MEDEVAC units have been recently equipped with in-flight blood transfusion capabilities, but the majority are not, and there is a lack of data to confirm whether this addition improves the ability to handle casualties with severe blood losses (Malsby III *et al.*, 2013). Without sufficient data to determine the effectiveness of in-flight transfusion, there has not been a change in the MEDEVAC system's evaluation measure. Therefore, unlike civilian EMS systems, it is vital to stabilize and transport battlefield casualties to an appropriate MTF (e.g., one that has the capability and resources to perform necessary care such as blood transfusions) and into surgery rather than simply providing medical aid at the CCP. So, while civilian EMS systems measure performance by response time (i.e., the time it takes to reach the patient after obtaining the emergency call), military EMS systems are evaluated in terms of how long it takes to transport the casualties from the CCP to an MTF. Therefore, it is appropriate to define the response time for a MEDEVAC unit as $T_7 - T_2$. Moreover, the service time for a MEDEVAC unit

is defined as $T_8 - T_2$, which is commonly associated as the time expended to service a request.

The primary objective of the MEDEVAC system presented in this thesis is to dispatch MEDEVAC units in a way that maximizes the expected total discounted reward attained by the system. The dispatch authority (i.e, AEO) must make sequential decisions under uncertainty as to which available MEDEVAC unit to dispatch to service a 9-line MEDEVAC request. It is impossible to know exactly when and where casualty events will occur, which prevents the dispatch authority from having *a priori* information on subsequent 9-line MEDEVAC requests. The knowledge and details of any 9-line MEDEVAC request only become known to the MEDEVAC system upon receipt of the request. Once the GSAB receives the request and the AEO selects a MEDEVAC unit to dispatch, the assigned MEDEVAC unit must initiate mission protocols immediately. The mission protocols of a MEDEVAC unit include preparing medical personnel and equipment prior to departure, traveling to the CCP to pick up casualties, providing appropriate en route medical care, and transporting casualties to the nearest MTF in a rapid and efficient manner. Delaying any mission tasks negatively impacts the total response time and ultimately decreases the survivability rates of casualties awaiting service.

Both a dynamic and stochastic approach are needed when analyzing the dispatch of civilian and military emergency response vehicles. The stochastic aspect of this problem derives from the uncertainty concerning the manifestation of casualty events. Moreover, the dispatch, travel, and service times vary for each request and cannot be predicted precisely. When examining civilian EMS systems, the data relating to dispatch, travel, and service times are easily accessible and can be leveraged to parameterize decision models. Unfortunately, as noted earlier, one of the underlining challenges for medical planners in the military is having to develop and identify a

dispatching policy prior to commencement of combat operations. No casualty event data exists for such a situation. Therefore, this thesis utilizes a rubric that emulates the judgment and expertise of military planners with regard to the future interactions of enemy and friendly forces to identify the locations and arrivals of casualty events.

IV. Methodology

This chapter presents the Markov decision process (MDP) model of the military’s medical evacuation (MEDEVAC) dispatching problem. One of the key benefits of formulating an MDP model is that it provides a framework in which dynamic programming algorithms can be utilized to compute exact optimal policies. In most cases, MDP formulations have clear definitions for the state space, action space, rewards, transition probabilities, and optimality equations.

The objective of the MDP model formulated in this thesis is to determine which available MEDEVAC unit to dispatch in response to a 9-line MEDEVAC request submission with the purpose of maximizing the expected total discounted reward over an infinite horizon.

The MDP model assumes that 9-line MEDEVAC requests arrive according to a Poisson process with parameter λ that is denoted by $PP(\lambda)$. Military medical planners must ensure the MEDEVAC system is tailored to effectively support friendly forces within an assigned area of operations (AO) (Department of the Army, 2014). In large-scale combat operations, military medical planners should examine the expected conditions of the operation and carefully select an appropriate λ -value based on these conditions to investigate the peak hours of operation. Each casualty event that leads to a 9-line MEDEVAC request submission is categorized by its precedence level, which is determined by the senior military member and/or medical personnel at the site of injury.

The Army utilizes three casualty event precedence categories (i.e., *urgent*, *priority*, and *routine*) when submitting a 9-line MEDEVAC request (Department of the Army, 2014). A routine evacuation precedence level is assigned to casualties that are triaged as minimally injured (i.e., non-life-threatening), and typically results in standard ground or waterborne assets responding within 24 hours of the initial event

(De Lorenzo, 2003). Since the focus of this thesis is on the aerial aspect of MEDEVAC operations and routine 9-line requests typically do not utilize dedicated air evacuation assets, this thesis only considers 9-line MEDEVAC requests that have a precedence level of either urgent or priority.

The arrival of urgent and priority 9-line MEDEVAC requests from different zones is modeled utilizing a *splitting* technique. Splitting is generating two or more counting processes out of a single Poisson process (Kulkarni, 2009). Let the original counting process $\{N(t') : t' \geq 0\}$ denote the $PP(\lambda)$ that counts the number of 9-line MEDEVAC request arrivals to the general support aviation battalion (GSAB) that have taken place during the time interval $(0, t']$. The original counting process can be split into counting processes that are categorized by the zone $z \in \mathcal{Z} = \{1, 2, \dots, |\mathcal{Z}|\}$ and the precedence level $k \in \mathcal{K} = \{1, 2, \dots, |\mathcal{K}|\}$ of the request. The sets $\mathcal{Z}$ and $\mathcal{K}$ represent the set of zones and the set of precedence levels in the system, respectively. Let $\mathcal{R} = \{(z, k) : (z, k) \in \mathcal{Z} \times \mathcal{K}\}$ be the set of request categories. There is a total of $|\mathcal{R}| = |\mathcal{Z}||\mathcal{K}|$ request categories. The original process $\{N(t') : t' \geq 0\}$ is split into $|\mathcal{R}|$ independent processes $\{N_{zk}(t') : t' \geq 0\}, \forall (z, k) \in \mathcal{R}$. It is clear that

$$N(t') = \sum_{(z,k) \in \mathcal{R}} N_{zk}(t') \quad (1)$$

since each request belongs to one and only one category. The nature of the split processes $\{N_{zk}(t') : t' \geq 0\}, \forall (z, k) \in \mathcal{R}$ depends on how the requests are categorized. The process of categorizing each request is called the splitting mechanism. The Bernoulli splitting mechanism generates the split processes $\{N_{zk}(t') : t' \geq 0\}, \forall (z, k) \in \mathcal{R}$ given parameters $p_{zk} > 0, \forall (z, k) \in \mathcal{R}$ such that $\sum_{(z,k) \in \mathcal{R}} p_{zk} = 1$. Each request is independently categorized by its zone z and precedence level k combination with probability p_{zk} independent of everything else. The splitting mechanism allows the characterization of each split process $\{N_{zk}(t') : t' \geq 0\}, (z, k) \in \mathcal{R}$ as a

Poisson process with parameter λp_{zk} , which is denoted by $PP(\lambda p_{zk})$.

There may be times when a 9-line MEDEVAC request is admitted into the system, but all MEDEVAC units are currently servicing other requests. When this occurs, the submitted 9-line MEDEVAC request is placed in its respective zone-precedence queue to be serviced at a later time. Moreover, there may be system states wherein an idle MEDEVAC is available for assignment, but placing the submitted request in its respective zone-precedence queue rather than assigning the idle MEDEVAC to the request could prove more advantageous in the long run. For example, the decision not to assign an available MEDEVAC unit immediately could prove beneficial if a lower precedence request enters the system while many MEDEVAC units are busy. In such a situation, waiting for another MEDEVAC unit to become available before servicing the lower precedence request allows the idle MEDEVAC unit to be available for a possibly higher precedence request, yet to arrive.

The service time for a MEDEVAC unit comprises the time between the initial assignment notification and returning to the staging area. This thesis assumes that the service times of the MEDEVAC units are exponentially distributed. While this simplifying assumption may not be realistic, it is often utilized in related literature. For example, Jarvis (1985) performs several computational experiments, and the results suggest that the shape of the service-time distribution has little impact on the overall behavior of the system. Similarly, research by Gross & Harris (1998) also indicate the insensitivity of service time distributions to system performance. Moreover, McLay & Mayorga (2013b) perform simulation analyses utilizing different types of service time distributions to study the impact of modeling the system with exponential service times versus more realistic service times. Results indicate that the assumption of exponential service times does not significantly impact the optimal policies. This suggests that the optimal policies determined utilizing the MDP model from this the-

sis give military medical planners relevant insight as to how to dispatch MEDEVAC units despite the simplifying assumption of exponentially distributed service times.

Having introduced the characteristics of the arrival process and the nature of the service times, formulation of the MDP model can now proceed. The development of the MDP model components leverage Maxwell *et al.* (2010), Keneally *et al.* (2016), and Rettke *et al.* (2016). The decision epochs, state space, action space, transition probabilities, rewards, objective, and optimality equation are described in detail below.

The decision epochs of the MEDEVAC system are the points in time that require a decision. The set of decision epochs is denoted as $\mathcal{T} = \{1, 2, \dots\}$. Two event types in the MEDEVAC system constitute all decision epochs. The first event type is the submission of a 9-line MEDEVAC request. The second event type is the change in the status of a MEDEVAC unit from busy to available upon completing a mission.

The MEDEVAC system MDP model follows the properties of semi-Markov decision processes (SMDPs). SMDPs generalize MDPs by requiring the decision-maker to select a feasible action whenever the system changes, allowing the time spent in a specific state to follow an arbitrary probability distribution, and modeling the system evolution in continuous time (Puterman, 1994). The MEDEVAC system MDP model is viewed as a continuous time MDP (CTMDP), which is a special case of an SMDP wherein the inter-transition times are exponentially distributed and decisions are made at every transition. There are several different ways that CTMDPs can be analyzed, but the primary method utilized in this thesis is *uniformization*. Uniformization is applied to the CTMDP model to obtain an equivalent discrete-time discounted model with constant transition rates (Puterman, 1994). The transformation allows the results and algorithms for discrete-time MDP models to be applied directly.

The state $S_t \in \mathcal{S}$ describes the status of the entire MEDEVAC system at decision epoch $t \in \mathcal{T}$. The MEDEVAC system state is represented by the tuple $S_t = (M_t, Q_t, \hat{R}_t)$ wherein M_t represents the MEDEVAC status tuple at epoch t , Q_t represents the queue status tuple at epoch t , and $\hat{R}_t$ represents the request arrival status tuple at epoch t .

The MEDEVAC status tuple M_t describes the status of every MEDEVAC unit in the system at epoch t . The tuple M_t can be written as

$$M_t = (M_{tm})_{m \in \mathcal{M}}, \quad (2)$$

where $\mathcal{M} = \{1, 2, \dots, |\mathcal{M}|\}$ represents the set of MEDEVAC units in the system. The state variable $M_{tm} \in \{0\} \cup \mathcal{Z}$ contains the information pertaining to MEDEVAC unit $m \in \mathcal{M}$ at epoch t . Each MEDEVAC unit can either be idle or servicing a request in one of the zones in the system. When $M_{tm} = 0$, MEDEVAC unit m is idle. When $M_{tm} = z$, MEDEVAC unit m is servicing a request from zone $z \in \mathcal{Z}$.

The queue status tuple Q_t describes the status of every zone-precedence queue in the system at epoch t . The tuple Q_t can be written as

$$Q_t = (Q_{tzk})_{z \in \mathcal{Z}, k \in \mathcal{K}}. \quad (3)$$

The state variable $Q_{tzk} \in \{0, 1, \dots, q^{max}\}$ contains the information pertaining to the $(z, k) \in \mathcal{R}$ zone-precedence queue at epoch t . Each zone-precedence queue can hold no more than q^{max} requests at any point in time.

The request arrival status tuple $\hat{R}_t$ indicates whether there is a request arrival awaiting an admission decision at epoch t ; it also provides the zone and precedence level of the request arrival, given one is present at epoch t . Let $\hat{R}_t = (0, 0)$ when there

is not a request arrival at the GSAB at epoch t . Otherwise, let

$$\hat{R}_t = \left(\hat{Z}_t, \hat{K}_t \right)_{\hat{Z}_t \in \mathcal{Z}, \hat{K}_t \in \mathcal{K}}. \quad (4)$$

The random variable $\hat{Z}_t$ represents the zone of the request arrival, and the random variable $\hat{K}_t$ represents the precedence level of the request arrival at epoch t . At epoch t , the information in $\hat{Z}_t$ and $\hat{K}_t$ has just been realized and is no longer uncertain. However, $\hat{Z}_t$ and $\hat{K}_t$ are random variables at epochs $1, 2, \dots, t-1$ because the information they contain is still uncertain.

The size of the state space $\mathcal{S}$ depends on $|\mathcal{M}|, |\mathcal{Z}|, |\mathcal{K}|$, and q^{max} . The following expression indicates the cardinality of the state space for the MEDEVAC system:

$$|\mathcal{S}| = (1 + |\mathcal{Z}|)^{|\mathcal{M}|} (1 + q^{max})^{|\mathcal{Z}||\mathcal{K}|} (1 + |\mathcal{Z}||\mathcal{K}|). \quad (5)$$

Unfortunately, the size of the state space grows exponentially with respect to the number of state variables. This is commonly referred to as the curse of dimensionality and renders dynamic programming intractable for analyzing practical (i.e., large-scale) scenarios. The purpose of constructing and analyzing small problem instances is to determine if any insight concerning practical scenarios can be obtained by solving the small problem instances exactly utilizing dynamic programming.

Events are triggered when a 9-line MEDEVAC request is submitted to the system or if a busy MEDEVAC unit completes a service request and becomes available. An admission control decision only occurs when a 9-line MEDEVAC request is submitted to the system. A dispatching decision may be necessary when either of these two event types occur.

The MEDEVAC system employs an inter-zone policy regarding airspace access that allows any MEDEVAC unit to service any 9-line MEDEVAC request, regardless

of the zone from which the request originated. Once a MEDEVAC unit is tasked, it will be considered unavailable until the task is completed and the MEDEVAC unit has returned to its own staging area. While rerouting a MEDEVAC unit during mid-flight can be accomplished, potential delays and communication difficulties can create issues in the MEDEVAC system that may ultimately cost casualties their lives. Furthermore, most military operations do not utilize a MEDEVAC unit rerouting strategy during combat operations (Rettke *et al.*, 2016). Due to these reasons, rerouting MEDEVAC units mid-flight is not incorporated in this MDP model.

When a 9-line MEDEVAC request is submitted, the AEO must take into account the current state of the system and make an admission control and possibly a dispatching decision. There are three possible alternatives: allowing the request to enter its respective zone-precedence queue; assigning an available MEDEVAC unit to service the request immediately; or rejecting the request from entering the system, which forces the request to be serviced by an outside agency (i.e., CASEVAC). If a request arrival is present at epoch t and its queue is not full, i.e., $\hat{R}_t = (\hat{Z}_t, \hat{K}_t)$ and $Q_{t\hat{Z}_t\hat{K}_t} < q^{max}$, $\hat{Z}_t \in \mathcal{Z}$, $\hat{K}_t \in \mathcal{K}$, then the AEO can either accept or reject the request from entering the system. If the request is accepted, it can either be placed in its respective zone-precedence queue or an available MEDEVAC unit can be tasked to service the request immediately. Moreover, if a request arrival is present at epoch t and its queue is full, i.e., $\hat{R}_t = (\hat{Z}_t, \hat{K}_t)$ and $Q_{t\hat{Z}_t\hat{K}_t} = q^{max}$, $\hat{Z}_t \in \mathcal{Z}$, $\hat{K}_t \in \mathcal{K}$, then the AEO must reject the request from entering the system. Practically speaking, q^{max} should be set high enough so that requests are not routinely rejected due to a full queue.

Let the decision variable $x_t^{reject} \in \{\Delta, 0, 1\}$ denote the admission control decision at epoch t . If an arrival request is not present at epoch t , i.e., $\hat{R}_t = (0, 0)$, the only available decision is $x_t^{reject} = \Delta$, which indicates the system will continue to transition

without any impact from x_t^{reject} . When $x_t^{reject} = 0$, the arrival request at epoch t is admitted to the MEDEVAC system, whereas when $x_t^{reject} = 1$, the arrival request at epoch t is rejected from entering the MEDEVAC system.

Dispatching decisions may be required when either a 9-line request is submitted or a busy MEDEVAC unit completes a service request and becomes available. Let $\mathcal{I}(S_t) = \{m : m \in \mathcal{M}, M_{tm} = 0\}$ denote the set of idle MEDEVAC units available for dispatching when the state of the system is S_t at epoch t . Let $\mathcal{W}(S_t) = \{(z, k) : (z, k) \in \mathcal{R}, Q_{tzk} > 0\}$ denote the set of zone-precedence queues that have at least one casualty event awaiting service when the state of the system is S_t at epoch t . The dispatching decision is represented by the tuple $x_t^d = (x_t^{ar}, x_t^{qr})$ wherein x_t^{ar} represents the arrival request dispatch decision tuple and x_t^{qr} represents the queued requests dispatch decision tuple at epoch t .

The arrival request dispatch decision tuple x_t^{ar} describes the AEO's dispatching decision with regard to arrival requests at epoch t . The tuple x_t^{ar} can be written as

$$x_t^{ar} = (x_{tm}^{ar})_{m \in \mathcal{I}(S_t)}. \quad (6)$$

The decision variable $x_{tm}^{ar} = 1$ if MEDEVAC unit $m \in \mathcal{I}(S_t)$ is dispatched to service the arrival request $\hat{R}_t = (\hat{Z}_t, \hat{K}_t)$, where $\hat{Z}_t \in \mathcal{Z}$ and $\hat{K}_t \in \mathcal{K}$, at epoch t , and 0 otherwise.

The queued requests dispatch decision tuple, x_t^{qr} , describes the AEO's dispatching decision with regard to queued requests at epoch t . The tuple x_t^{qr} can be written as

$$x_t^{qr} = (x_{tmzk}^{qr})_{m \in \mathcal{I}(S_t), (z, k) \in \mathcal{W}(S_t)}. \quad (7)$$

The decision variable $x_{tmzk}^{qr} = 1$ if MEDEVAC unit $m \in \mathcal{I}(S_t)$ is dispatched to service a queued request from the (z, k) zone-precedence queue, where $(z, k) \in \mathcal{W}(S_t)$, at

epoch t , and 0 otherwise.

Let $x_t = (x_t^{reject}, x_t^d)$ denote a compact representation of the decision variables at epoch t . Several constraints bound the decisions being made at epoch t . The first constraint,

$$I_{\{\hat{R}_t \neq (0,0)\}} \sum_{m \in \mathcal{I}(S_t)} x_{tm}^{ar} + \sum_{m \in \mathcal{I}(S_t)} \sum_{(z,k) \in \mathcal{W}(S_t)} x_{tmzk}^{qr} \leq 1, \quad (8)$$

requires that there is at most one MEDEVAC unit dispatched at epoch t . The next constraint,

$$x_t^{reject} \leq 1 - \sum_{m \in \mathcal{I}(S_t)} x_{tm}^{ar}, \quad (9)$$

indicates that if an arrival request is present at epoch t and a MEDEVAC unit is tasked to service the arrival request at epoch t , as indicated by $x_{tm}^{ar} = 1$ for some $m \in \mathcal{I}(S_t)$, then the arrival request must enter the system, as indicated by $x_t^{reject} = 0$. Otherwise, $x_{tm}^{ar} = 0$ for all $m \in \mathcal{I}(S_t)$, and the arrival request is either queued (i.e., $x_t^{reject} = 0$) or rejected (i.e., $x_t^{reject} = 1$) from the system at epoch t . The set of available actions when a decision is required is denoted as follows

$$\mathcal{X}(S_t) = \begin{cases} (\Delta, (\{0\}^{|\mathcal{I}(S_t)|}, \{0, 1\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (0, 0), \mathcal{I}(S_t) \neq \emptyset, \mathcal{W}(S_t) \neq \emptyset \\ (\Delta, (\{0\}^{|\mathcal{I}(S_t)|}, \{0\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (0, 0), (\mathcal{I}(S_t) = \emptyset \text{ or } \mathcal{W}(S_t) = \emptyset) \\ (1, (\{0\}^{|\mathcal{I}(S_t)|}, \{0, 1\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), Q_{t\hat{Z}_t\hat{K}_t} = q^{max}, \mathcal{I}(S_t) \neq \emptyset, \mathcal{W}(S_t) \neq \emptyset \\ (1, (\{0\}^{|\mathcal{I}(S_t)|}, \{0\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), Q_{t\hat{Z}_t\hat{K}_t} = q^{max}, (\mathcal{I}(S_t) = \emptyset \text{ or } \mathcal{W}(S_t) = \emptyset) \\ (\{0, 1\}^{|\mathcal{I}(S_t)|}, (\{0, 1\}^{|\mathcal{I}(S_t)|}, \{0, 1\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), Q_{t\hat{Z}_t\hat{K}_t} < q^{max}, \mathcal{I}(S_t) \neq \emptyset, \mathcal{W}(S_t) \neq \emptyset \\ (\{0, 1\}^{|\mathcal{I}(S_t)|}, (\{0\}^{|\mathcal{I}(S_t)|}, \{0\}^{|\mathcal{I}(S_t)| \times |\mathcal{W}(S_t)|})), & \text{if } \hat{R}_t = (\hat{Z}_t, \hat{K}_t), Q_{t\hat{Z}_t\hat{K}_t} < q^{max}, \mathcal{I}(S_t) = \emptyset \end{cases} \quad (10)$$

where Constraints (8) and (9) must be satisfied. The first two cases in Equation (10) represent all feasible actions when the decision epoch occurs due to a MEDEVAC unit completing a service request and becoming available immediately, whereas the last four cases represent all feasible actions when the decision epoch occurs due to a 9-line MEDEVAC request submission.

State transitions are Markovian with two possible events dictating the transition.

The first event type is the submission of a 9-line MEDEVAC request. Recall that 9-line MEDEVAC requests arrive according to a $PP(\lambda)$. The second event type is the change in the status of a MEDEVAC unit from busy to available upon completing a mission. Let μ_{mz} denote the service rate of MEDEVAC unit $m \in \mathcal{M}$ when servicing a 9-line MEDEVAC request in zone $z \in \mathcal{Z}$. Let $\mathcal{B}(S_t) = \{m : m \in \mathcal{M}, M_{tm} \neq 0\}$ denote the set of busy MEDEVAC units when the state of the system is S_t at epoch t . If the MEDEVAC system is in pre-decision state S_t and action x_t is taken, the system will immediately transition to a post decision state S_t^x . The sojourn time in S_t^x (i.e., the time the system remains in post decision state S_t^x before transitioning to the next pre-decision state S_{t+1}) follows an exponential distribution with parameter $\beta(S_t, x_t)$. Simple calculations reveal that

$$\beta(S_t, x_t) = \lambda + \sum_{m \in \mathcal{B}(S_t)} \mu_{m, M_{tm}} + \sum_{m \in \mathcal{I}(S_t)} \mu_{m, \hat{Z}_t} x_{tm}^{ar} + \sum_{m \in \mathcal{I}(S_t)} \sum_{(z, k) \in \mathcal{W}(S_t)} \mu_{mz} x_{tmzk}^{qr}. \quad (11)$$

If $\mathcal{B}(S_t) = \emptyset$, $x_{tm}^{ar} = 0 \forall m \in \mathcal{I}(S_t)$, and $x_{tmzk}^{qr} = 0 \forall m \in \mathcal{I}(S_t), (z, k) \in \mathcal{W}(S_t)$, then $\beta(S_t, x_t)$ represents the sojourn time for the state-action pairs wherein the next decision epoch occurs upon the arrival of a 9-line MEDEVAC request. Otherwise, $\beta(S_t, x_t)$ represents the sojourn time for the state-action pairs wherein the next decision epoch occurs after either a 9-line MEDEVAC request arrives to the GSAB or one of the busy MEDEVACs completes a service request and becomes available. Let T_a denote the time until the next 9-line MEDEVAC request arrival. Let T_s denote the time until the next service completion. The time until the next decision epoch T_e satisfies $T_e = \min(T_a, T_s)$. Since both T_a and T_s follow an exponential distribution, standard calculations show that T_e follows an exponential distribution with parameter $\beta(S_t, x_t)$.

The probabilistic behavior of the process is summarized in terms of its infinitesimal

generator. The infinitesimal generator is a $|\mathcal{S}| \times |\mathcal{S}|$ matrix G with components:

$$G(S_{t+1}|S_t, x_t) = \begin{cases} -[1 - p(S_t|S_t, x_t)]\beta(S_t, x_t), & \text{if } S_{t+1} = S_t \\ p(S_{t+1}|S_t, x_t)\beta(S_t, x_t), & \text{if } S_{t+1} \neq S_t \end{cases} \quad (12)$$

wherein $p(S_{t+1}|S_t, x_t)$ denotes the probability that the system transitions to state S_{t+1} given that it is currently in state S_t and decision x_t is made. Note that $p(S_t|S_t, x_t) = 0$, which means that the system will transition to a different state at the end of a sojourn in state S_t .

Puterman (1994) argues that converting CTMDPs to equivalent discrete-time MDPs via the uniformization approach makes subsequent analysis easier to perform. To uniformize the system, the maximum rate of transition must be determined and is calculated by

$$\nu = \lambda + \sum_{m \in \mathcal{M}} \tau_m, \quad (13)$$

where

$$\tau_m = \max_{z \in \mathcal{Z}} \mu_{mz}, \quad \forall m \in \mathcal{M}. \quad (14)$$

The restriction that there are no self transitions from a state to itself is removed when uniformization is applied to the process. Applying uniformization gives the following transition probabilities:

$$\tilde{p}(S_{t+1}|S_t, x_t) = \begin{cases} 1 - \frac{[1 - p(S_t|S_t, x_t)]\beta(S_t, x_t)}{\nu}, & \text{if } S_{t+1} = S_t \\ \frac{p(S_{t+1}|S_t, x_t)\beta(S_t, x_t)}{\nu}, & \text{if } S_{t+1} \neq S_t. \end{cases} \quad (15)$$

This transformation may be viewed as inducing extra (i.e., “notional”) transitions from a state to itself. This modified process has the same probabilistic structure as the CTMDP.

The decision epochs in CTMDPs follow each state transition, and the times between decision epochs are exponentially distributed. Several factors impact the amount of reward gained from making a decision to service a 9-line MEDEVAC request. These factors include the zone and precedence level of the 9-line MEDEVAC request as well as the staging area of the servicing MEDEVAC unit. Let $c(S_t, x_t) = \psi_{mzk}$ denote the immediate expected reward (i.e., contribution) if MEDEVAC unit $m \in \mathcal{M}$ is dispatched to service a zone $z \in \mathcal{Z}$, precedence level $k \in \mathcal{K}$ 9-line MEDEVAC request (i.e., $x_{tm}^{ar} = 1$ or $x_{tmzk}^{qr} = 1$). The immediate expected reward is computed as follows:

$$\psi_{mzk} = \begin{cases} \delta e^{\frac{\zeta_{mz}}{60}}, & \text{if } k = 1 \text{ (i.e., } urgent) \\ e^{\frac{\zeta_{mz}}{240}}, & \text{if } k = 2 \text{ (i.e., } priority) \\ 0, & \text{otherwise,} \end{cases} \quad (16)$$

wherein ζ_{mz} is the expected response time when MEDEVAC $m \in \mathcal{M}$ is dispatched to service a request in zone $z \in \mathcal{Z}$, and $\delta \geq 1$ is a tradeoff parameter utilized to vary the urgent to priority immediate expected reward ratio. If a MEDEVAC unit is not dispatched to service a 9-line MEDEVAC request at epoch t then $c(S_t, x_t) = 0$.

Multiple casualties with different levels of severity may comprise a single 9-line MEDEVAC request. In practice, each casualty is assigned an evacuation precedence category, but in this model, the overall 9-line MEDEVAC request classification (i.e., the evacuation precedence category) is based on the most time-sensitive casualty within the request. The 9-line MEDEVAC classification should not be overemphasized because it may place a burden on the MEDEVAC dispatching system that could result in the loss of lives.

Let $h(S_t, x_t)$ denote the continuous expected holding cost accumulated when decision x_t is selected in state S_t . The MEDEVAC system incurs a holding cost based on the time requirements outlined in the Army's Medical Evacuation Field Manual. The

MEDEVAC system seeks to service urgent and priority 9-line MEDEVAC requests within 60 and 240 minutes from notification, respectively. Let ϕ_k denote the holding cost rate for holding a single precedence- k request in its queue between decision epochs. The holding cost rate ϕ_k can be written as

$$\phi_k = \xi \frac{\sum_{m \in \mathcal{M}} \sum_{z \in \mathcal{Z}} \psi_{mzk}}{|\mathcal{M}| |\mathcal{Z}|}, \forall k \in \mathcal{K}, \quad (17)$$

where $\xi \in [0, 1]$ is a parameter that scales the holding cost rate for a precedence- k request based on the average immediate expected reward over all possible MEDEVAC-zone combinations. Summing the holding costs over all zone-precedence queues gives the following expression

$$h(S_t, x_t) = \sum_{z \in \mathcal{Z}} \sum_{k \in \mathcal{K}} \phi_k Q_{tzk}. \quad (18)$$

Simple calculations show that if $\mathcal{W}(S_t) = \emptyset$ then $h(S_t, x_t) = 0$. That is, if no requests are queued, then no holding cost is incurred. Since the system does not change in the time between decision epochs, the expected discounted reward is

$$r(S_t, x_t) = c(S_t, x_t) - \frac{h(S_t, x_t)}{\alpha + \beta(S_t, x_t)}, \quad (19)$$

where $\alpha > 0$ denotes the continuous time discounting rate. Applying uniformization gives

$$\tilde{r}(S_t, x_t) \equiv r(S_t, x_t) \frac{\alpha + \beta(S_t, x_t)}{\alpha + \nu}. \quad (20)$$

Note that the uniformized rewards agree with the rewards in the CTMDP.

Let $X^\pi(S_t)$ be a policy (i.e., decision function) that prescribes AEO dispatch decisions for each state $S_t \in \mathcal{S}$. That is, $x = X^\pi(S_t)$ is the dispatching decision returned when utilizing policy π . The optimal policy π^* is sought from the class of policies $(X^\pi(S_t)_{\pi \in \Pi})$ that maximizes the expected total discounted reward earned by

the MEDEVAC system. The objective is expressed as

$$\max_{\pi \in \Pi} \mathbb{E}^{\pi} \left\{ \sum_{t=1}^{\infty} \gamma^{t-1} \tilde{r}(S_t, X^{\pi}(S_t)) \right\}, \quad (21)$$

where $\gamma = \frac{\nu}{\nu + \alpha}$ is the uniformized discount factor. The optimal policy is found by solving the Bellman equation

$$J(S_t) = \max_{x_t \in \mathcal{X}(S_t)} \left\{ \tilde{r}(S_t, x_t) + \gamma \sum_{S_{t+1} \in \mathcal{S}} \tilde{p}(S_{t+1} | S_t, x_t) J(S_{t+1}) \right\}. \quad (22)$$

The policy iteration algorithm is implemented in MATLAB to solve Equation (22) exactly. Policy iteration starts with an initial policy and then iteratively performs two steps: *policy evaluation*, which computes the expected total discounted reward of each state given the current policy, and *policy improvement*, which updates the current policy if any improvements are available (Puterman, 1994). The policy iteration algorithm terminates after the policy converges.

For comparison purposes, a linear programming (LP) model of the Markov decision problem is also constructed. Puterman (1994) notes that LP has not proven to be an efficient method for solving large discounted Markov decision problems. However, recent advancements in LP algorithms have increased the computational efficiency of LP (e.g., as indicated by the performance testing of CPLEX and Gurobi in Bixby (2012)) and make LP a more viable solution method for solving MDPs. Use of an efficient LP algorithm benefits the analysis of MDPs because it eases the inclusion of constraints and provides a better mechanism with which to conduct sensitivity analyses.

V. Testing, Analysis & Results

This chapter presents a representative military medical evacuation (MEDEVAC) planning scenario utilized both to demonstrate the applicability of the Markov decision process (MDP) model and to examine the behavior of the optimal dispatching policy. A series of sensitivity analyses and excursions identify the model parameters that significantly impact the optimal dispatching policy. Military medical planners should focus on these parameters when developing MEDEVAC dispatching policies. Moreover, this chapter compares the computational efficiency of policy iteration via MATLAB versus linear programming via CPLEX. The thesis utilizes a dual Intel Xeon E5-2650v2 workstation having 128 GB of RAM and MATLAB's Parallel Computing Toolbox to conduct the computational experiments and analysis outlined in this chapter.

5.1 Representative Scenario

As of 2017, the United States (U.S.) continues to conduct military operations in Afghanistan. The initial launch of U.S. military operations in Afghanistan began with the initiation of Operation ENDURING FREEDOM (OEF) on October 7, 2001 in response to the terrorist attacks on New York's World Trade Center and the Pentagon on September 11, 2001. OEF lasted a little over 13 years and officially ended when U.S. combat operations in Afghanistan were terminated on December 31, 2014. However, as part of Operation FREEDOM'S SENTINEL (OFS), U.S. military forces still remain in Afghanistan to participate in a coalition mission to train and assist the Afghan military and to conduct counter-terrorism operations against Al Qaeda (Department of Defense, 2016). While official U.S. combat operations are currently not being conducted in Afghanistan, military medical planners still prepare and plan

for potential combat scenarios in the event that a sudden change requires U.S. combat operations.

The computational examples in Bandara *et al.* (2012), Keneally *et al.* (2016), and Rettke *et al.* (2016) are leveraged as a basis for the representative scenario examined herein. This thesis considers a notional planning scenario in which a coalition of allied countries executes combat operations in response to an increase in terrorist attacks by remnants of Al-Qaeda militants in southern Afghanistan. For simplicity, this notional scenario (hereafter referred to as the 2×2 case) assumes a MEDEVAC system with two demand zones (i.e., the zones at which 9-line MEDEVAC requests originate) and two MEDEVAC unit staging areas (i.e., the locations in which the MEDEVAC units are stationed) with one medical treatment facility (MTF) co-located at each staging area.

Afghanistan is comprised of 34 provinces. Figure 3 utilizes the data from White (2016) to illustrate the war-related fatalities of allied forces in Afghanistan by province since the beginning of OEF until December 2016. The 2×2 case assumes that the southern region of Afghanistan is the area of operations (AO) and is divided into two separate demand zones: Helmand province (Zone 1) and Kandahar province (Zone 2). Two MEDEVAC units are considered with one being staged in Zone 1 (i.e., MEDEVAC 1) and the other being staged in Zone 2 (i.e., MEDEVAC 2). The placement of the staging areas and co-located MTFs represents a general realism based on the historical trends in enemy activity in southern Afghanistan.

As depicted in Figure 3, Helmand and Kandahar are the two provinces that have produced the most war-related fatalities in Afghanistan since the start of OEF with 956 and 558 killed in action (KIA), respectively (White, 2016). While these numbers do not account for every type of casualty (e.g., military wounded in action (WIA) and civilian casualties), they do provide a representative sample that is utilized as an




Figure 3. Afghanistan combat fatalities by province, 2001-2016

approximation of the threat level present in each zone. Moreover, these numbers are utilized to determine the proportion of 9-line MEDEVAC requests from each zone. Simple calculations yield that the proportion of requests coming from Zone 1 is $p_{z_1} = 0.6314$ and the proportion of requests coming from Zone 2 is $p_{z_2} = 1 - p_{z_1} = 0.3686$.

Each 9-line MEDEVAC request is independently categorized by its zone z (e.g., Helmand and Kandahar) and precedence level k (e.g., urgent, priority, and routine) combination. Fulton *et al.* (2010) report that the probability of a casualty event being classified with a precedence level of urgent, priority, or routine is 11%, 12%, and 77%, respectively. Recall that routine requests are assumed to be serviced by non-MEDEVAC units (i.e., casualty evacuation (CASEVAC)). The 2×2 case assumes that the proportion of requests classified with an urgent precedence level is $p_{k_1} = 0.5$ and the proportion of requests classified with a priority precedence level is $p_{k_2} = 1 - p_{k_1} =$

0.5. The proportion of each request categorization p_{zk} is found by multiplying the zone proportion with the precedence level proportion (e.g., $p_{11} = p_{z_1}p_{k_1}$). Table 1 shows the 2×2 case's request categorization proportions.

Table 1. 9-Line MEDEVAC Request Proportions by Zone-Precedence Level

Zone, z	Urgent	Priority
Zone 1 (Helmand)	0.3157	0.3157
Zone 2 (Kandahar)	0.1843	0.1843

Military medical planners estimate the arrival rate of 9-line MEDEVAC requests by estimating when and where future tactical level engagements will take place, along with the likelihood and severity of casualty events. The reward obtained for servicing a 9-line MEDEVAC request depends on the locations of the request, the servicing MEDEVAC unit, and the closest MTF. The response and service times described in Chapter III are generated by leveraging the procedure set forth in Keneally *et al.* (2016).

The procedure utilized to model future 9-line MEDEVAC requests avoids using current data from southern Afghanistan to maintain operational security. Indeed, actual data for current MEDEVAC unit, casualty event, and MTF locations are restricted. Instead, the spatial distribution of future 9-line MEDEVAC requests are modeled with a Monte Carlo simulation via a Poisson cluster process. Casualty cluster centers are selected by leveraging data from ICOS (2008) pertaining to insurgent attacks in southern Afghanistan resulting in death in 2007. It is assumed that all casualty events generated from the casualty cluster centers result in 9-line MEDEVAC requests. Moreover, the distribution of 9-line MEDEVAC request locations from a given casualty cluster center is generated on a uniform distribution with respect to the distance of the request to the casualty cluster center. Military medical planners must keep in mind that data will certainly change with respect to each unique conflict. Furthermore, the dispatching policy generated depends on the input data

and, therefore, must be relevant to the scenario being modeled to obtain meaningful results.

Figure 4 depicts the two zones (i.e., Helmand and Kandahar) in southern Afghanistan utilized to generate the data, as well as the MEDEVAC and MTF locations. Recall that the MEDEVAC and MTF locations are collocated for the 2×2 case. The collocated MEDEVAC and MTF locations in each zone are represented by blue stars. The casualty cluster centers in each zone are represented by red diamonds.




Figure 4. MEDEVAC and MTF locations with Casualty Cluster Centers

Figure 5 illustrates several casualty events resulting in 9-line MEDEVAC requests throughout southern Afghanistan within a 48-hour time period. The collocated MEDEVAC and MTF locations are still represented by blue stars. The casualty events are represented by red circles.




Figure 5. Sampled Casualty Events in Helmand and Kandahar

The data generated for the variables that comprise the response time vary with each mission and, therefore, are represented as random variables. The response time variables representing mission preparation time, travel time to casualty collection point (CCP), service time at CCP, travel time to MTF, and service time at the MTF are defined in Chapter III and described in detail in the following four paragraphs.

The mission preparation time is exponentially distributed with a mean of 10 minutes. The 2008 MEDEVAC after action report (AAR) estimates mission prep time to be 20 minutes (Bastian, 2010). This AAR, along with personal experiences, influences Bastian (2010) to model mission prep time with a mean of 20 minutes and standard deviation of 5 minutes. However, a more recent interview with a MEDEVAC pilot in O'Shea (2011) reports that with proper pre-planning procedures the mission prep

time is often less than 10 minutes.

The armed escort delay is exponentially distributed with a mean of 10 minutes. Garrett (2013) reports that there is a 31% chance that a MEDEVAC mission requires an armed escort. Moreover, of the missions requiring an armed escort, approximately 4% are delayed due to issues caused primarily by the escort aircraft. These percentages are included in the computation of the expected response times and, therefore, the expected rewards. The delay caused by armed escorts is an important feature of the MEDEVAC problem. This thesis applies the same armed escort assumptions found in Keneally *et al.* (2016), to which we refer a more interested reader for a more in depth description on how armed escorts impact this MDP model.

The flight speed, which accounts for the travel time to the CCP and the travel time to the MTF, is uniformly distributed between 120 and 193 knots with a resulting mean of 156.5 knots. This flight speed is based on currently fielded MEDEVAC helicopters (i.e., HH-60Ms) and on subject matter expertise (Bastian, 2010).

The service time at the CCP and the service time at the MTF are exponentially distributed with a mean of 10 minutes and five minutes, respectively. These times are determined by leveraging the data provided by in-theater MEDEVAC pilots and other subject matter experts described in Bastian (2010) and Keneally *et al.* (2016).

The just-described response time random variables, casualty cluster centers, and MEDEVAC staging areas are utilized in a Monte Carlo simulation to obtain a synthetic, but realistic, spatial distribution of future 9-line MEDEVAC requests and response time data. The means of the response times are computed and presented in Table 2.

Table 2. Expected Response Times (minutes)

Zone, z	MEDEVAC 1	MEDEVAC 2
1 (Helmand)	34.25	48.18
2 (Kandahar)	52.98	36.89

After the expected response times are computed, the expected service times can be computed by simply adding the appropriate expected response time to the MEDEVAC unit’s travel time back to its staging area. This travel time is defined in Chapter III and is based on the flight speed of the MEDEVAC helicopter. The distribution for the flight speed for the travel time to the staging area is the same as the flight speed distributions for the travel times to the CCP and MTF. The expected service times for the 2×2 case are provided in Table 3.

Table 3. Expected Service Times (minutes)

Zone, z	MEDEVAC 1	MEDEVAC 2
1 (Helmand)	34.25	67.28
2 (Kandahar)	72.13	36.89

Recall from Chapter IV that the MEDEVAC system employs an inter-zone policy regarding airspace access. This means that any MEDEVAC unit can service any 9-line MEDEVAC request, regardless of the zone from which the request originated. For example, the MEDEVAC unit staged in Helmand for the 2×2 case can service requests from both Helmand and Kandahar.

The thesis applies a survivability function that is monotonically decreasing in response time to compute the reward obtained from servicing a 9-line MEDEVAC request. The immediate expected reward for servicing a 9-line MEDEVAC request is determined by the precedence level and the response time of the request as indicated in Equation (16). For the 2×2 case, the immediate expected reward function utilizes $\delta = 10$, which rewards the servicing of *urgent* (i.e., $k = 1$) 9-line MEDEVAC requests much more than *priority* (i.e., $k = 2$) 9-line MEDEVAC requests. Table 4 summarizes the computed immediate expected rewards, ψ_{mzk} .

The continuous expected holding cost is computed based on the number of urgent and priority 9-line MEDEVAC requests that are in the queue between decision epochs. The 2×2 case utilizes $\xi = 0.20$, which scales the holding cost rate for a precedence-

Table 4. Immediate Expected Rewards

Zone, z	Precedence, k	MEDEVAC, m	
		1	2
1 (Helmand)	1 (Urgent)	5.65	4.48
	2 (Priority)	0.87	0.82
2 (Kandahar)	1 (Urgent)	4.14	5.41
	2 (Priority)	0.80	0.86

k request to be 20% of the average immediate expected reward over all possible MEDEVAC-zone combinations.

The 2×2 case assumes a high operations tempo with a baseline request arrival rate of $\lambda = \frac{1}{60}$, representing an average 9-line MEDEVAC request rate of one request per . The military intelligence community, operational planners, and medical planners should work together to determine a reasonable request arrival rate prior to a planned combat operation based on the equipment, size, and disposition of friendly and adversary forces.

5.2 Representative Scenario Results

A list of parameters associated with the 2×2 case are displayed in Table 5. Utilizing the parameter settings in Table 5 and the expected response times, expected service times, and immediate expected rewards computed in the previous section, the optimal policy for the 2×2 case is determined via policy iteration. Applying Equation (5) reveals that size of the state space for the 2×2 case is 58,320. This result indicates that even for this relatively simple scenario, the size of the state space is quite large.

For comparison purposes, three baseline dispatching policies are considered. The three baseline policies are all based on a classic inter-zone myopic policy. Recall that an inter-zone myopic policy sends the closest available MEDEVAC unit to service an incoming 9-line MEDEVAC request, regardless of the request’s zone or precedence level. All three baseline policies adopt this strategy when at least one MEDEVAC

Table 5. 2×2 Case Parameters

Parameter	Description	Setting
λ	9-line MEDEVAC request arrival rate	$\frac{1}{60}$
$ \mathcal{M} $	Total # of MEDEVACs	2
$ \mathcal{Z} $	Total # of zones	2
$ \mathcal{K} $	Total # of precedence levels	2
q^{max}	Max (z, k) queue length	5
γ	Uniformized discount factor	0.99
δ	Weight for urgent request	10
ξ	Scale for holding cost rate	0.2
p_{z_1}	Zone 1 proportion of requests	0.6314
p_{z_2}	Zone 2 proportion of requests	0.3686
p_{k_1}	Urgent proportion of requests	0.5
p_{k_2}	Priority proportion of requests	0.5

unit is available. The differences between the three baseline policies are found when both MEDEVAC units are busy. The first baseline policy (i.e., Myopic 1) will queue 9-line MEDEVAC requests if there are no available MEDEVAC units to service the request, regardless of the request's zone or precedence level. The second baseline policy (i.e., Myopic 2) will queue only urgent 9-line MEDEVAC requests if there are no available MEDEVACs to service the request, regardless of the urgent request's zone. The third baseline policy (i.e., Myopic 3) will not queue any 9-line MEDEVAC requests. If there are queued requests, the Myopic 1 and Myopic 2 dispatching policies service requests with a prioritized first-come-first-serve basis. The optimal policy's dispatching strategy, queue lengths, and MEDEVAC utilization rates are compared against the three baseline policies to obtain a better understanding of where similarities and differences exist. Moreover, the optimality gap for each baseline policy is computed to demonstrate whether a myopic policy is appropriate for the given 2×2 case.

The dispatching decisions for the optimal policy and three baseline policies are compared in three separate scenarios. Each scenario (i.e., Scenarios 1-3) considers a set of MEDEVAC system states with empty zone-precedence queues. The first

scenario (i.e., Scenario 1) considers a system state wherein both MEDEVAC units are idle, which can be represented as $S_t \in ((0, 0), (0, 0, 0, 0), \hat{R}_t)$. The dispatching policies for Scenario 1 are displayed in Table 6. Regardless of the zone or precedence level of the incoming 9-line MEDEVAC request, $\hat{R}_t$, all four policies react in a myopic fashion when the system is in state $S_t \in ((0, 0), (0, 0, 0, 0), \hat{R}_t)$, sending the closest MEDEVAC unit to service the request.

Table 6. Comparison of Dispatching Policies for Scenario 1

Policy	$\hat{R}_t$	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 1
	(1,2)	Dispatch MEDEVAC 1
	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
Myopic 1	(1,1)	Dispatch MEDEVAC 1
Myopic 2	(1,2)	Dispatch MEDEVAC 1
Myopic 3	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2

The second scenario (i.e., Scenario 2) considers a set of MEDEVAC system states wherein MEDEVAC 1 is idle and MEDEVAC 2 is busy, which can be represented as $S_t \in ((0, z), (0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2\}$. Moreover, the third scenario (i.e., Scenario 3) considers a set of MEDEVAC system states wherein MEDEVAC 1 is busy and MEDEVAC is idle, which can be represented as $S_t \in ((z, 0), (0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2\}$. The dispatching policies for Scenarios 2 and 3 are displayed in Tables 7 and 8, respectively. Contrary to the findings of Keneally *et al.* (2016) in their computational example, the best MEDEVAC unit to dispatch to service a 9-line MEDEVAC request does depend on the zone in which the busy MEDEVAC is currently servicing. Note that this is an observed result based on the parameter settings for the 2×2 case and that location-independent policies are a possibility, as seen in Keneally *et al.* (2016). In Tables 7 and 8 an asterisk (*) is placed next to the incoming requests, $\hat{R}_t$, that do not follow a myopic policy. It is expected that

a myopic policy will apply to all urgent 9-line MEDEVAC requests due to the life threatening nature of these requests and the accompanying high rewards for servicing them

Table 7. Comparison of Dispatching Policies for Scenario 2

Policy	$\hat{R}_t$	MEDEVAC 2 Servicing Zone 1	MEDEVAC 2 Servicing Zone 2
		Queue\Dispatch\Reject	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
	(1,2)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
	(2,1)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
	(2,2)*	Reject	Queue
Myopic 1	(1,1)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
Myopic 2	(1,2)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
Myopic 3	(2,1)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1
	(2,2)	Dispatch MEDEVAC 1	Dispatch MEDEVAC 1

Table 8. Comparison of Dispatching Policies for Scenario 3

Policy	$\hat{R}_t$	MEDEVAC 1 Servicing Zone 1	MEDEVAC 1 Servicing Zone 2
		Queue\Dispatch\Reject	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(1,2)*	Queue	Reject
	(2,1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
Myopic 1	(1,1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
Myopic 2	(1,2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
Myopic 3	(2,1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2

Consider the Scenario 2 results displayed in Table 7. The MEDEVAC system is in a state $S_t \in ((0, z), (0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2\}$. The optimal dispatching policy for $\hat{R}_t = (2, 2)$ depends on z , the zone where MEDEVAC 2 is currently servicing a request. If MEDEVAC 2 is servicing Zone 1 (i.e., $z = 1$) and $\hat{R}_t = (2, 2)$, then the optimal decision is to reject the request from entering the system and send the request to be serviced by CASEVAC. If MEDEVAC 2 is servicing Zone 2 (i.e., $z = 2$) and $\hat{R}_t = (2, 2)$, then the optimal decision is to accept and queue the request. Both of these decisions differ from the myopic decision (i.e., dispatch MEDEVAC 1 to

service the request). This shows that, if the system is in a Scenario 2 state and $\hat{R}_t = (2, 2)$, then the optimal policy will reserve MEDEVAC 1 for either an urgent 9-line MEDEVAC request or a Zone 1 request. The difference between rejecting or queueing the request is driven by the difference in expected service times. Recall that there is large difference in expected service times for MEDEVAC 2 to Zone 1 and MEDEVAC 2 to Zone 2; 67.28 minutes and 36.28 minutes, respectively.

Consider the Scenario 3 results displayed in Table 8. The MEDEVAC system is in a state $S_t \in ((z, 0), (0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2\}$. The optimal dispatching policy for $\hat{R}_t = (1, 2)$ depends on z , the zone where MEDEVAC 1 is currently servicing a request. If MEDEVAC 1 is servicing Zone 1 (i.e., $z = 1$) and $\hat{R}_t = (1, 2)$, then the optimal decision is to accept and queue the request. If MEDEVAC 2 is servicing Zone 2 (i.e., $z = 2$) and $\hat{R}_t = (1, 2)$, then the optimal decision is to reject the request from entering the system and send the request to be serviced by CASEVAC. Both of these decisions differ from the myopic decision (i.e., dispatch MEDEVAC 2 to service the request). This shows that, if the system is in a Scenario 3 state and $\hat{R}_t = (1, 2)$, then the optimal policy will reserve MEDEVAC 2 for either an urgent 9-line MEDEVAC request or a Zone 2 request. The difference between rejecting or queueing the request is driven by the difference in expected service times. Recall that there is large difference in expected service times for MEDEVAC 1 to Zone 1 and MEDEVAC 2 to Zone 1; 34.25 minutes and 72.13 minutes, respectively.

The workload of each MEDEVAC unit is an interesting performance measure, and substantial differences between the optimal and baseline policies are revealed when the probabilities of each MEDEVAC being busy are examined. Figures 6 and 7 display the long-run busy probabilities for MEDEVAC 1 and MEDEVAC 2, respectively.

Examination of Figures 6 and 7 indicate that the optimal policy prefers to have MEDEVAC units servicing their own zone (i.e., the zone in which they are staged).




Figure 6. Comparison of MEDEVAC 1 busy probabilities

This result aligns with intuition due to the large differences in expected service times mentioned earlier. Moreover, both figures show a substantial difference in the MEDEVAC long-run busy probabilities between the optimal and baseline policies. Consider Figure 7. All three baseline policies result in MEDEVAC 2 servicing requests in Zone 1 more than Zone 2, whereas MEDEVAC 2 is primarily busy servicing Zone 2 requests when implementing the optimal policy. This interesting result is driven in part by the higher proportion of requests arriving from Zone 1 ($p_{z_1} = 0.6314$) and because baseline policies are forced to send MEDEVAC 2 to service Zone 1 requests (when MEDEVAC 1 is busy and MEDEVAC 2 is available) whereas the optimal policy is allowed to queue or reject the request.

Another interesting result found from the analysis of MEDEVAC busy probab-




Figure 7. Comparison of MEDEVAC 2 busy probabilities

ities is the combined average utilization of both MEDEVACs for each policy. The data displayed in Table 9 represents the system when at least one MEDEVAC unit is being utilized, regardless of the zone being serviced. The optimal policy has the lowest combined average MEDEVAC utilization of 0.3569, meaning that the optimal policy utilizes each MEDEVAC in the most efficient manner. Moreover, the Myopic 1 policy has the highest combined average MEDEVAC utilization of 0.4073. This aligns with expectations because the optimal policy has control over admission, queueing, and dispatching rules whereas the baseline policies do not.

Another performance measure of interest is the average lengths of each zone-precedence queue. Obviously, the Myopic 3 policy will not have any queueing data, but comparisons can still be made between the optimal policy, the Myopic 1 policy,

Table 9. Comparison of Combined Average MEDEVAC Utilization

Policy	Utilization
Optimal	0.3569
Myopic 1	0.4073
Myopic 2	0.3798
Myopic 3	0.3591

and the Myopic 2 policy. Figure 8 displays the long-run average queue lengths for each zone-precedence queue. The average zone-precedence queue lengths for the optimal policy are strictly less than the baseline averages for every zone-precedence queue except for Q_{t12} . The optimal policy’s average Q_{t12} length is greater than the Myopic 1 policy’s Q_{t12} length. In comparison to the Myopic 1 policy, the optimal policy queues more Zone 1, priority requests to reduce Zone 1, urgent request wait times. This result is explained by the proportion of requests arriving from Zone 1 ($p_{z_1} = 0.6314$) and the MEDEVAC unit service times for Zone 1.

Lastly, the optimality gaps between the baseline policies and the optimal policy are examined. The expected total discounted reward for the optimal policy and baseline policies when the system is in an *empty* state $S^0 = ((0, 0), (0, 0, 0, 0), (0, 0))$ (i.e., both MEDEVAC units are idle, every zone-precedence queue is empty, and there are no 9-line MEDEVAC requests in the system) are displayed in Table 10, along with the optimality gaps associated with each baseline policy. The results indicate that the best baseline policy is Myopic 2, which has the smallest optimality gap of 0.74%. Without having the ability to queue any requests, the Myopic 3 policy performs worse than every other policy and has the largest optimality gap of 5.73%. While these optimality gaps may not seem large, over a long enough time period the optimal policy will save more lives.




Figure 8. Comparison of zone-precedence queue lengths

Table 10. Comparison of Total Expected Discounted Rewards & Optimality Gaps

Policy, π	$J^\pi(S^0)$	Optimality Gap
Optimal	63.50	N/A
Myopic 1	62.09	2.21%
Myopic 2	63.02	0.74%
Myopic 3	59.86	5.73%

5.3 Experimental Design

Since there are many parameters associated with the MEDEVAC system, a screening experiment is developed to reveal the parameters that significantly impact the value of the optimal dispatching policy. Leveraging the results found from the 2×2 case, a 2^5 full factorial screening experiment is generated to determine the relative significance of factors $\lambda, \delta, \xi, p_{z_1}$, and p_{k_1} . All five of these factors are important

MEDEVAC parameters of interest and have an initial screening design with low and high factor levels. For example, the intensity in which 9-line MEDEVAC requests arrive to the system, λ , is designed with low and high factor levels (e.g., $\frac{1}{75}$ and $\frac{1}{45}$, respectively) to determine if decreased or increased intensity of λ has a significant impact on the value of the optimal dispatching policy.

The 2^5 full factorial screening experimental factors and the levels associated with each factor are displayed in Table 11. Once the results from the 2^5 full factorial screening experiment are examined, the factors that have a statistically significant impact on the value of the optimal dispatching policy are analyzed via a three-level experiment with low, intermediate, and high factor levels.

Table 11. 2^5 Full Factorial Screening Experimental Factor Levels

Factor	Low Level	High Level
λ	$\frac{1}{75}$	$\frac{1}{45}$
δ	5	15
ξ	0.1	0.3
p_{z_1}	0.25	0.75
p_{k_1}	0.25	0.75

5.4 Experimental Design Results

Table 12 reports the results from the 2^5 full factorial screening experiment. Starting from the left, the first column indicates the run number. The second through sixth columns indicate the factor levels. The rightmost column indicates the dependent variable $J^{\pi^*}(S^0)$, where $J^{\pi^*}(S^0)$ is computed utilizing Equation (22) under the optimal policy, π^* . Recall that S^0 is the empty state (i.e., both MEDEVAC units are idle, every zone-precedence queue is empty, and there are no 9-line MEDEVAC requests in the system).

Multiple linear regression analysis is conducted to examine the relationship between the independent factors $\lambda, \delta, \xi, p_{z_1}$, and p_{k_1} and the dependent variable $J^{\pi^*}(S^0)$.

Table 12. 2^5 Full Factorial Screening Experiment Results

Run #	$\frac{1}{\lambda}$	δ	ξ	p_{z_1}	p_{k_1}	$J^{\pi^*}(S^0)$
1	45	5	0.10	0.25	0.25	31.63
2	45	5	0.10	0.25	0.75	53.02
3	45	5	0.10	0.75	0.25	32.68
4	45	5	0.10	0.75	0.75	55.36
5	45	5	0.30	0.25	0.25	28.95
6	45	5	0.30	0.25	0.75	48.53
7	45	5	0.30	0.75	0.25	30.00
8	45	5	0.30	0.75	0.75	50.78
9	45	15	0.10	0.25	0.25	64.74
10	45	15	0.10	0.25	0.75	151.40
11	45	15	0.10	0.75	0.25	67.03
12	45	15	0.10	0.75	0.75	157.26
13	45	15	0.30	0.25	0.25	61.01
14	45	15	0.30	0.25	0.75	141.10
15	45	15	0.30	0.75	0.25	63.13
16	45	15	0.30	0.75	0.75	147.13
17	75	5	0.10	0.25	0.25	23.16
18	75	5	0.10	0.25	0.75	38.98
19	75	5	0.10	0.75	0.25	23.68
20	75	5	0.10	0.75	0.75	40.20
21	75	5	0.30	0.25	0.25	22.16
22	75	5	0.30	0.25	0.75	37.28
23	75	5	0.30	0.75	0.25	22.76
24	75	5	0.30	0.75	0.75	38.59
25	75	15	0.10	0.25	0.25	46.74
26	75	15	0.10	0.25	0.75	109.51
27	75	15	0.10	0.75	0.25	48.04
28	75	15	0.10	0.75	0.75	113.05
29	75	15	0.30	0.25	0.25	45.02
30	75	15	0.30	0.25	0.75	105.67
31	75	15	0.30	0.75	0.25	46.42
32	75	15	0.30	0.75	0.75	109.33

The results from the multiple linear regression analysis are displayed in Table 13. Starting from the left, the first column lists the dependent factors. The second, third, fourth, and fifth columns list the estimated coefficients (Coef), standard errors (SE), test statistics (T), and p -values (P) associated with the dependent factors, respectively.

Table 13. Multiple Linear Regression Analysis

Factors	Coef	SE	T	P
Intercept	-76.00	17.18	-4.42	< 0.00
λ	2201.80	673.03	3.27	< 0.00
δ	5.62	0.60	9.39	< 0.00
ξ	-18.31	29.91	-0.61	0.55
p_{z_1}	4.57	11.97	0.38	0.71
p_{k_1}	92.51	11.97	7.73	< 0.00

The results from the multiple linear regression analysis in Table 13 report that the p -values associated with factors λ , δ , and p_{k_1} are all less than 0.01, which indicates that these factors are statistically significant in predicting $J\pi^*(S^0)$. Intuitively, these results make sense. The rate with which 9-line MEDEVAC requests arrive directly impacts the number of requests that can be serviced, resulting in more/less opportunity to earn rewards. Increasing or decreasing the weight and proportion of urgent requests also directly impacts the amount of reward earned by the system. Moreover, Table 13 reports the p -values associated with ξ and p_{z_1} are both greater than 0.05 and, therefore, do not provide enough evidence to assume that the factors ξ and p_{z_1} are statistically significant in predicting $J\pi^*(S^0)$. The reason that these factors are not significant could be due to the selected experimental design factor levels. Selecting a wider range in factor levels for ξ and p_{z_1} could result in them becoming significant. This model results in an adjusted $R^2 = 0.83$, indicating an adequate fit but a reduced model excluding factors ξ and p_{z_1} is tested to see if a better model can be obtained.

Utilizing the results from Table 13, a 3^3 full factorial experiment is generated to examine the differences between the optimal and baseline dispatching policies at different levels for factors λ , δ , and p_{k_1} . The goal of the 3^3 full factorial experiment is to gain insight as to when medical planners should avoid implementing myopic dispatching policies (e.g., Myopic 1, Myopic 2, and Myopic 3) and to understand how the changes in the factor levels for λ , δ , and p_{k_1} impact the optimal dispatching policy.

The 3^3 full factorial experimental factors and the levels associated with each factor are displayed in Table 14.

Table 14. 3^3 Full Factorial Experimental Factor Levels

Factor	Low Level	Intermediate Level	High Level
λ	$\frac{1}{75}$	$\frac{1}{60}$	$\frac{1}{45}$
δ	5	10	15
p_{k_1}	0.25	0.50	0.75

Table 15 reports the results from the 3^3 full factorial experiment. Starting from the left, the first column indicates the run number. The next three columns indicate the factor levels. The fifth column indicates the dependent variable $J^{\pi^*}(S^0)$, where $J^{\pi^*}(S^0)$ is computed utilizing Equation (22) under the optimal policy, π^* . The next three columns indicate the optimality gaps for the Myopic 1, Myopic 2, and Myopic 3 policies, respectively. The following four columns indicate the MEDEVAC busy probabilities when the system is operating under the optimal dispatching policy. The four rightmost columns indicate the average zone-precedence queue lengths when the system is operating under the optimal dispatching policy.

The results from Table 15 indicate that the Myopic 2 policy strictly outperforms the Myopic 3 policy. Moreover, the Myopic 2 policy strictly outperforms the Myopic 1 policy when $\frac{1}{\lambda} \in \{45, 60\}$, but not when $\frac{1}{\lambda} = 75$. These results indicate that medical planners should never employ the Myopic 3 policy because there is always a better policy to choose for any given set of parameter settings in Table 15. Additionally, the Myopic 1 policy outperforms the Myopic 2 policy in several instances when $\frac{1}{\lambda} = 75$ because as the inter-arrival time of 9-line MEDEVACs increases it becomes more beneficial to queue all requests versus just queueing urgent requests.

The MEDEVAC unit busy probabilities associated with each set of parameter settings in Table 15 also provide interesting results. MEDEVAC 1 is busy servicing Zone 1 requests substantially more than servicing Zone 2 requests for all 27 runs.

Table 15. 3^3 Full Factorial Experimental Results

Run #	$\frac{1}{\lambda}$	δ	p_{k_1}	$J^{\pi^*}(S^0)$	Optimality Gaps			MEDEVAC 1 Busy		MEDEVAC 2 Busy		Average Queue Lengths			
					Myopic 1	Myopic 2	Myopic 3	Zone 1	Zone 2	Zone 1	Zone 2	Q_{i11}	Q_{i12}	Q_{i21}	Q_{i22}
1	45	5	0.25	30.63	6.59%	1.05%	5.27%	0.21	0.04	0.17	0.14	0.02	0.03	0.01	0.04
2	45	5	0.5	41.38	7.11%	1.28%	5.77%	0.20	0.07	0.17	0.14	0.04	0.01	0.02	0.02
3	45	5	0.75	51.74	7.01%	2.93%	5.34%	0.21	0.04	0.14	0.15	0.05	0.00	0.06	0.00
4	45	10	0.25	47.25	6.55%	1.08%	7.39%	0.20	0.04	0.15	0.15	0.02	0.01	0.01	0.04
5	45	10	0.5	74.69	7.87%	2.56%	8.32%	0.20	0.06	0.11	0.14	0.03	0.00	0.02	0.00
6	45	10	0.75	99.92	7.40%	3.74%	6.73%	0.20	0.09	0.15	0.13	0.05	0.00	0.03	0.00
7	45	15	0.25	64.20	7.01%	1.61%	8.89%	0.20	0.03	0.05	0.16	0.01	0.00	0.01	0.03
8	45	15	0.5	108.30	8.41%	3.33%	9.55%	0.19	0.06	0.11	0.14	0.03	0.00	0.02	0.00
9	45	15	0.75	149.11	8.17%	4.68%	7.84%	0.17	0.08	0.13	0.11	0.04	0.00	0.02	0.00
10	60	5	0.25	26.49	1.41%	1.27%	4.47%	0.16	0.07	0.13	0.10	0.01	0.03	0.01	0.01
11	60	5	0.5	35.73	1.64%	0.58%	4.66%	0.16	0.04	0.12	0.11	0.02	0.01	0.01	0.02
12	60	5	0.75	44.85	1.64%	0.68%	4.50%	0.16	0.06	0.12	0.11	0.03	0.01	0.02	0.01
13	60	10	0.25	40.36	1.63%	0.71%	5.34%	0.16	0.02	0.11	0.11	0.01	0.02	0.01	0.03
14	60	10	0.5	63.50	2.21%	0.74%	5.73%	0.17	0.04	0.07	0.12	0.01	0.04	0.01	0.02
15	60	10	0.75	86.10	2.21%	1.11%	5.33%	0.16	0.06	0.10	0.11	0.03	0.00	0.02	0.00
16	60	15	0.25	54.38	2.02%	0.71%	6.03%	0.16	0.02	0.11	0.12	0.01	0.02	0.00	0.03
17	60	15	0.5	91.63	2.83%	1.21%	6.53%	0.16	0.04	0.07	0.11	0.01	0.00	0.01	0.00
18	60	15	0.75	127.54	2.54%	1.40%	5.75%	0.16	0.06	0.10	0.11	0.02	0.00	0.02	0.00
19	75	5	0.25	23.05	0.41%	1.40%	3.83%	0.13	0.05	0.09	0.09	0.01	0.02	0.00	0.01
20	75	5	0.5	31.04	0.41%	0.49%	3.79%	0.13	0.05	0.09	0.09	0.01	0.01	0.01	0.00
21	75	5	0.75	39.03	0.47%	0.28%	3.74%	0.13	0.04	0.09	0.09	0.02	0.00	0.01	0.01
22	75	10	0.25	34.90	0.38%	0.60%	4.05%	0.13	0.03	0.09	0.10	0.01	0.01	0.00	0.02
23	75	10	0.5	54.87	0.74%	0.35%	4.30%	0.13	0.03	0.08	0.09	0.01	0.01	0.01	0.01
24	75	10	0.75	74.62	0.77%	0.39%	4.14%	0.14	0.04	0.07	0.09	0.02	0.01	0.01	0.01
25	75	15	0.25	46.88	0.64%	0.48%	4.43%	0.13	0.02	0.08	0.09	0.00	0.01	0.00	0.02
26	75	15	0.5	78.87	1.09%	0.52%	4.72%	0.14	0.02	0.05	0.10	0.01	0.03	0.00	0.01
27	75	15	0.75	110.37	1.02%	0.58%	4.42%	0.13	0.04	0.07	0.09	0.01	0.00	0.01	0.00

MEDEVAC 2 is busy servicing each zone approximately the same. This result aligns with intuition because the proportion of requests arriving from Zone 1 ($p_{z_1} = 0.6314$) is greater than the proportion of requests arriving from Zone 2 ($p_{z_2} = 0.3686$).

Multiple linear regression analysis is conducted to confirm the statistically significant relationship between the independent factors λ, δ , and p_{k_1} and the dependent variable $J^{\pi^*}(S^0)$. The results from the multiple linear regression analysis are displayed in Table 16. Starting from the left, the first column lists the dependent factors. The second, third, fourth, and fifth columns list the estimated coefficients (Coef), standard errors (SE), test statistics (T), and p -values (P) associated with the dependent factors, respectively.

Table 16. Multiple Linear Regression Analysis

Factors	Coef	SE	T	P
Intercept	-75.72	12.50	-6.06	< 0.00
λ	2145.60	571.74	3.75	< 0.00
δ	5.64	0.51	10.98	< 0.00
p_{k_1}	92.26	10.27	8.98	< 0.00

The results from the multiple linear regression analysis in Table 16 show that the p -values associated with factors λ, δ , and p_{k_1} are all less than 0.01, which indicates that these factors are statistically significant in predicting $J^{\pi^*}(S^0)$. Moreover, this model resulted in an adjusted $R^2 = 0.89$, which is greater than the adjusted $R^2 = 0.83$ computed from Table 13. This result indicates that the updated model with factors λ, δ , and p_{k_1} is better in predicting $J^{\pi^*}(S^0)$ than the previous model with factors $\lambda, \delta, \xi, p_{z_1}$, and p_{k_1} .

An interesting observation found from the 3^3 full factorial experiment is that the optimal dispatching policy aligns with the myopic policy when the MEDEVAC system is in a Scenario 1 state for 26 out of the 27 runs. Table 17 reports the optimal and baseline dispatching policies for the single run that the optimal dispatching policy does not act myopically. The optimal dispatching policy will reject precedence level two requests (i.e., priority requests) when the system is in a Scenario 1 state and $\lambda = \frac{1}{45}$, $\delta = 15$, and $p_{k_1} = 0.75$. This result is intuitive because the inter-arrival times of the requests have increased from one every 60 minutes to one every 45 minutes, the immediate expected reward for servicing urgent requests is substantially higher than servicing priority requests, and there is a much higher rate of urgent requests arriving to the system rather than priority requests.

Table 17. Comparison of Dispatching Policies for Scenario 1

Policy	$\hat{R}_t$	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 1
	(1,2)*	Reject
	(2,1)	Dispatch MEDEVAC 2
	(2,2)*	Reject
Myopic 1	(1,1)	Dispatch MEDEVAC 1
Myopic 2	(1,2)	Dispatch MEDEVAC 1
Myopic 3	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
Settings: $\lambda = \frac{1}{45}$, $\delta = 15$, and $p_{k_1} = 0.75$		

Tables 18, 19, and 20 display the average results for fixed λ -, δ -, and p_{k_1} - parameter values from Table 15. This aggregated information indicates how each factor impacts the system performance metrics.

Table 18. Comparison of Average λ Performance Metrics

Run #'s	$\frac{1}{\lambda}$	$J^{\pi^*}(S^0)$	Optimality Gaps			MEDEVAC 1 Busy		MEDEVAC 2 Busy		Average Queue Lengths			
			Myopic 1	Myopic 2	Myopic 3	Zone 1	Zone 2	Zone 1	Zone 2	Q_{t11}	Q_{t12}	Q_{t21}	Q_{t22}
1-9	45	74.14	7.35%	2.47%	7.23%	0.20	0.06	0.13	0.14	0.03	0.01	0.02	0.01
10-18	60	63.40	2.01%	0.94%	5.37%	0.16	0.05	0.10	0.11	0.02	0.01	0.01	0.01
19-27	75	54.85	0.66%	0.57%	4.16%	0.13	0.04	0.08	0.09	0.01	0.01	0.01	0.01

Table 19. Comparison of Average δ Performance Metrics

Run #'s	δ	$J^{\pi^*}(S^0)$	Optimality Gaps			MEDEVAC 1 Busy		MEDEVAC 2 Busy		Average Queue Lengths			
			Myopic 1	Myopic 2	Myopic 3	Zone 1	Zone 2	Zone 1	Zone 2	Q_{t11}	Q_{t12}	Q_{t21}	Q_{t22}
1-3,10-12,19-21	5	35.99	2.97%	1.11%	4.60%	0.17	0.05	0.13	0.11	0.02	0.01	0.02	0.01
4-6,13-15,22-24	10	64.02	3.31%	1.26%	5.70%	0.17	0.04	0.10	0.12	0.02	0.01	0.01	0.01
7-9,16-18,25-27	15	92.36	3.75%	1.61%	6.46%	0.16	0.04	0.09	0.11	0.02	0.01	0.01	0.01

Table 20. Comparison of Average p_{k_1} Performance Metrics

Run #'s	p_{k_1}	$J^{\pi^*}(S^0)$	Optimality Gaps			MEDEVAC 1 Busy		MEDEVAC 2 Busy		Average Queue Lengths			
			Myopic 1	Myopic 2	Myopic 3	Zone 1	Zone 2	Zone 1	Zone 2	Q_{t11}	Q_{t12}	Q_{t21}	Q_{t22}
1,4,7,10,13,16,19,22,25	0.25	40.90	2.96%	0.99%	5.52%	0.17	0.04	0.11	0.12	0.01	0.02	0.01	0.02
2,5,8,11,14,17,20,23,26	0.50	64.45	3.59%	1.23%	5.93%	0.17	0.04	0.10	0.11	0.02	0.01	0.01	0.01
3,6,9,12,15,18,21,24,27	0.75	87.03	3.47%	1.76%	5.31%	0.16	0.06	0.11	0.11	0.03	0.00	0.02	0.00

The results from Table 18 indicate that, as the inter-arrival time of 9-line MEDEVAC requests decreases, the expected total discounted reward (i.e., $J^{\pi^*}(S^0)$) and the optimality gaps between the optimal policy and the baseline policies increases. I.e., the myopic policies increasingly underperform the optimal policy as the frequency of 9-line MEDEVAC requests increases. The observed MEDEVAC busy probabilities have the same patterns as mentioned in the description of the results for Table 15. Moreover, the average urgent queue lengths are always greater than or equal to the priority queue lengths.

The results from Table 19 indicate that as the ratio of urgent to priority immediate expected reward decreases, the expected total discounted reward (i.e., $J^{\pi^*}(S^0)$) and the optimality gaps between the optimal policy and the baseline policies increases. The observed MEDEVAC busy probabilities have the same patterns as mentioned

in the description of the results for Table 15. Moreover, the average urgent queue lengths are always greater than or equal to the priority queue lengths.

The results from Table 20 indicate that as the proportion of urgent 9-line MEDEVAC requests increases, the expected total discounted reward (i.e., $J^{\pi^*}(S^0)$) increases. However, this same pattern is not observed when comparing the optimality gaps. The observed MEDEVAC busy probabilities have the same patterns as mentioned in the description of the results for Table 15.

The average optimality gaps in Tables 18-20 indicate that the Myopic 2 policy is the best myopic policy, on average. This results provides medical planners with an easy-to-implement policy that performs fairly close to the optimal policy. This is useful because the optimal policy may not be easy-to-implement or practical for certain scenarios.

5.5 Excursion 1 - Request Arrival Rate

The section considers the impact that the arrival rate λ has on the optimal policy when the MEDEVAC system is in a Scenario 1 state $S_t \in ((0, 0), (0, 0, 0, 0), \hat{R}_t)$. The same parameter settings from the 2×2 case are utilized for the request arrival rate excursion except for λ ; see Table 5 for a descriptive list of the parameters and the settings associated with each one. The computational results indicate that the optimal policy dispatches the closest MEDEVAC unit when the system is in a Scenario 1 state with an urgent 9-line MEDEVAC request arrival (i.e., $S_t \in ((0, 0), (0, 0, 0, 0), (z, 1))$ where $z \in \{1, 2\}$), regardless of the request arrival rate λ . However, this same result does not hold true for when the system is in a Scenario 1 state with a priority 9-line MEDEVAC request arrival (i.e., $S_t \in ((0, 0), (0, 0, 0, 0), (z, 2))$ where $z \in \{1, 2\}$). The dispatching policies for when the system is in a Scenario 1 state with a priority 9-line request are displayed in Table 21.

Table 21. Comparison of MEDEVAC Dispatching Policies for Priority Requests

$\frac{1}{\lambda}$	Optimal Policy		Myopic Policy	
	$\tilde{R}_t = (1, 2)$	$\tilde{R}_t = (2, 2)$	$\tilde{R}_t = (1, 2)$	$\tilde{R}_t = (2, 2)$
21	Reject	Reject	MEDEVAC 1	MEDEVAC 2
22	Reject	Reject	MEDEVAC 1	MEDEVAC 2
23	Reject	Reject	MEDEVAC 1	MEDEVAC 2
24	Reject	Reject	MEDEVAC 1	MEDEVAC 2
25	Reject	Reject	MEDEVAC 1	MEDEVAC 2
26	Reject	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2
27	Reject	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2
28	Reject	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2
29	MEDEVAC 1	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2
30	MEDEVAC 1	MEDEVAC 2	MEDEVAC 1	MEDEVAC 2

The results from Table 21 indicate that when $\frac{1}{\lambda} \leq 25$ the optimal policy is to reject priority 9-line MEDEVAC requests, regardless of the zone where the request originated from. For $\frac{1}{\lambda} \in \{26, 27, 28\}$ the optimal policy is to reject Zone 1 priority 9-line MEDEVAC requests and to dispatch MEDEVAC 2 to Zone 2 priority requests. Lastly, when $\frac{1}{\lambda} \geq 29$ the optimal policy dispatches MEDEVAC units in a myopic manner. These results indicate that the optimal policy reserves MEDEVAC units for urgent requests as the inter-arrival time of 9-line MEDEVAC requests decreases.

5.6 Excursion 2 - MEDEVAC Flight Speed

This section considers the impact of replacing the currently fielded HH-60M MEDEVAC helicopter with a more efficient (i.e., faster flight speed) aeromedical aircraft. The same parameter settings from the 2×2 case are utilized for the MEDEVAC flight speed excursion; see Table 5 for a descriptive list of the parameters and the settings associated with each one. The HH-60M MEDEVAC helicopter still utilizes a power plant that was designed prior to 1989 (Leoni, 2007). There are significantly faster experimental rotary wing aircraft that could potentially be put into service to replace the HH-60M (Rettke *et al.*, 2016). It is reasonable to assume new rotary wing

aircraft designs have 25%-50% increased average flight speeds when compared to the HH-60M MEDEVAC helicopter.

To examine the impact of employing new rotary wing aircraft, the mean of the flight speed random variable is adjusted while all of the other random variables modeling the MEDEVAC process remain the same. Incorporating this change leads to immediate changes to response and service times, along with the immediate expected reward. It is expected that as the mean flight speed increases, the optimal dispatching policy will deploy MEDEVAC units in a more myopic fashion resulting in decreased optimality gaps for the Myopic 1, Myopic 2, and Myopic 3 dispatching policies. Moreover, another interesting scenario examined is when the mean flight speed decreases, which can occur due to potential maintenance issues or environmental issues within the area of operations. With limited resources, it is reasonable to assume that slower HH-60M MEDEVAC helicopters would still be utilized in a high intensity conflict.

Table 22 reports the results obtained by increasing and decreasing the mean flight speed, where flight speed is indicated as a percentage increase over the flight speed of the currently employed HH-60M MEDEVAC helicopter.

Table 22. MEDEVAC Helicopter Flight Speed Analysis

Flight Speed	Optimality Gaps		
	Myopic 1	Myopic 2	Myopic 3
-50%	17.07%	7.95%	10.93%
-25%	5.02%	1.80%	6.56%
0%(i.e., current)	2.21%	0.74%	5.73%
25%	1.12%	0.37%	5.30%
50%	0.63%	0.24%	5.06%

As expected, the results from Table 22 indicate that as the mean flight speed of the MEDEVAC helicopter increases, the optimality gaps for the Myopic 1, Myopic 2, and Myopic 3 policies all decrease. This shows that if a new rotary wing aircraft is fielded for MEDEVAC purposes, then the optimal dispatching policy will deploy

MEDEVAC units in a more myopic fashion. Moreover, the results indicate that as the mean flight speed of the MEDEVAC helicopter decreases, the optimality gaps for the Myopic 1, Myopic 2, and Myopic 3 policies all increase. This is an important observation. Military medical planners must take flight speed issues into consideration when developing dispatching policies. These results should also persuade military medical planners to consider changing dispatching policies during steady state combat operations if the mean flight speed of the MEDEVAC helicopters being utilized decreases due to atmospheric, environmental, or mechanical issues.

5.7 Excursion 3 - Intra-Zone Policies

This section considers the impact of replacing the MEDEVAC system’s *inter-zone* policy with an *intra-zone* policy with regards to airspace access. The same parameter settings from the 2×2 case and the MEDEVAC flight speed excursion are utilized for the intra-zone policies excursion; see Table 5 for a descriptive list of the parameters and the settings associated with each one. An intra-zone policy prevents MEDEVAC units from operating in zones outside of the zone in which they are staged. Military situations may arise that force strict adherence to an intra-zone policy. For example, an execution of a specific, short-duration combat operation may enforce an intra-zone policy to reduce the risk of collisions and fratricide (Keneally *et al.*, 2016). Moreover, when separate branches of the U.S. military (i.e., Army and Marines) and/or allied countries are working together in a combat environment, perhaps for the first time, an intra-zone policy restricting MEDEVAC units to serve their own zone may be enforced due to chain of command, communication, and/or political realities (Keneally *et al.*, 2016).

To examine the impact of enforcing an intra-zone policy, each MEDEVAC unit is restricted to operate in their own zone while all other random variables modeling

the MEDEVAC process remain the same. The queueing strategies associated with each baseline policy remain the same. Recall that when both MEDEVAC units are busy the Myopic 1 policy queues all incoming requests, the Myopic 2 policy only queues incoming urgent requests, and the Myopic 3 policy does not queue any incoming requests. Tables 23-25 report the dispatching policies associated with being in Scenarios 1-3, respectively.

Table 23. Comparison of Intra-Zone Dispatching Policies for Scenario 1

Policy	$\hat{R}_t$	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 1
	(1,2)	Dispatch MEDEVAC 1
	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
Myopic 1	(1,1)	Dispatch MEDEVAC 1
Myopic 2	(1,2)	Dispatch MEDEVAC 1
Myopic 3	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2

Regardless of the zone or precedence level of the incoming 9-line MEDEVAC request, $\hat{R}_t$, all four policies react in a myopic fashion when the system is in Scenario 1, sending the closest MEDEVAC unit to service the request.

Table 24. Comparison of Intra-Zone Dispatching Policies for Scenario 2

Policy	$\hat{R}_t$	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 1
Myopic 2	(1,2)	Dispatch MEDEVAC 1
	(2,1)	Queue
	(2,2)	Reject
Myopic 1	(1,1)	Dispatch MEDEVAC 1
	(1,2)	Dispatch MEDEVAC 1
	(2,1)	Queue
	(2,2)	Queue
Myopic 3	(1,1)	Dispatch MEDEVAC 1
	(1,2)	Dispatch MEDEVAC 1
	(2,1)	Reject
	(2,2)	Reject

Table 25. Comparison of Intra-Zone Dispatching Policies for Scenario 3

Policy	$\hat{R}_t$	Queue\Dispatch\Reject
Optimal	(1,1)	Queue
Myopic 2	(1,2)	Reject
	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
Myopic 1	(1,1)	Queue
	(1,2)	Queue
	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
Myopic 3	(1,1)	Reject
	(1,2)	Reject
	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2

Tables 23-25 indicate that the intra-zone optimal dispatching policy and the intra-zone Myopic 2 dispatching policy (i.e., queue urgent) dispatch MEDEVAC units in the same manner for Scenarios 1-3. Moreover, it is observed that when a MEDEVAC is busy and a request from the MEDEVAC's zone arrives to the system, the intra-zone optimal dispatching policy always rejects priority requests from entering the system. The difference between the intra-zone optimal dispatching policy and the intra-zone Myopic 2 dispatching policy is observed when there is at least one urgent 9-line MEDEVAC request in the queue, the MEDEVAC unit able to service the urgent queued request is busy, and there is an incoming request associated with that zone. Many states satisfy this description. Let such states be denoted as *Scenario 4 states*. Table 26 reports the dispatching policies associated with being in Scenario 4 when either: MEDEVAC 1 is busy, there is an urgent Zone 1 MEDEVAC request in the queue (i.e., $Q_{t11} = 1$) and a Zone 1 MEDEVAC request is submitted; or MEDEVAC 2 is busy, there is an urgent Zone 2 MEDEVAC request in the queue (i.e., $Q_{t21} = 1$), and a Zone 2 MEDEVAC request is submitted.

Table 26 indicates that if the MEDEVAC system is in a Scenario 4 state, the optimal policy will reject all incoming requests from the zone with the busy MEDEVAC

Table 26. Comparison of Intra-Zone Dispatching Policies for Scenario 4

Policy	$\hat{R}_t$	MEDEVAC 1 Busy & $Q_{t11} = 1$	MEDEVAC 2 Busy & $Q_{t21} = 1$
		Queue\Dispatch\Reject	Queue\Dispatch\Reject
Optimal	(1,1)	Reject	N/A
	(1,2)	Reject	N/A
	(2,1)	N/A	Reject
	(2,2)	N/A	Reject
Myopic 1	(1,1)	Queue	N/A
	(1,2)	Reject	N/A
	(2,1)	N/A	Queue
	(2,2)	N/A	Reject
Myopic 2	(1,1)	Queue	N/A
	(1,2)	Reject	N/A
	(2,1)	N/A	Queue
	(2,2)	N/A	Reject
Myopic 3	(1,1)	N/A	N/A
	(1,2)	N/A	N/A
	(2,1)	N/A	N/A
	(2,2)	N/A	N/A

VAC and the queued urgent request. Conversely, the intra-zone Myopic 2 policy will queue all incoming urgent requests. While rejecting an urgent request may not align with expectations, holding more than one request in the queue is detrimental due to the MEDEVAC units being restricted to service only their own zones. If such a decision is not desired by command authorities, the holding cost rate for urgent requests should be updated to be less detrimental to system performance or the value of servicing urgent requests should be increased to discourage rejecting urgent requests from entering the system.

Figure 9 displays the long-run busy probabilities for each MEDEVAC unit. As expected, the results from Figure 9 indicate that MEDEVAC 1 is busier than MEDEVAC 2, regardless of which intra-zone policy is utilized. These results occur because the proportion of Zone 1 requests ($p_{z_1} = 0.6314$) is greater than the proportion of Zone 2 requests ($p_{z_2} = 0.3686$).

Figure 10 displays the long-run average queue lengths for each zone-precedence




Figure 9. Comparison of MEDEVAC busy probabilities

queue when the system is operating under intra-zone policies. As expected, the average zone-precedence queue lengths for the optimal policy are strictly less than the baseline averages for every zone-precedence queue. Another observation from Figure 10 is that the proportion of Zone 1 queued requests is greater than the proportion of Zone 2 queued requests. Again, this can be explained due to the proportion of Zone 1 requests ($p_{z_1} = 0.6314$) being greater than the proportion of Zone 2 requests ($p_{z_2} = 0.3686$).

Lastly, the optimality gap between the intra-zone optimal policy and the intra-zone baseline policies is examined. The expected total discounted reward for the intra-zone optimal policy and intra-zone baseline policies when the MEDEVAC system is in State S^0 is displayed in Table 27, along with the optimality gaps associated with each intra-




Figure 10. Comparison of zone-precedence queue lengths

zone baseline policy. The results indicate that the best intra-zone baseline policy is Myopic 2, which has the smallest optimality gap of 7.45%. The intra-zone Myopic 1 policy performs worse than every other policy and has the largest optimality gap of 23.64%. These results indicate that when intra-policy restrictions are enforced then the myopic dispatching policies substantially under-perform when compared to the optimal policy. The 2×2 case optimality gaps displayed in Table 10 are substantially less than the optimality gaps for the inter-zone policies. The Myopic 2 policy has the best optimality gaps for both the 2×2 case and the inter-zone policy excursion. However, the optimality gap for the Myopic 2 policy in the 2×2 case is 0.74% whereas the Myopic 2 optimality gap in the inter-zone policy excursion is 7.45%. Moreover, there is an even large difference between the Myopic 1 policies (2.21% versus 23.64%).

These results show that intra-zone policies perform substantially worse than inter-zone policies.

Table 27. Comparison of Total Expected Discounted Rewards & Optimality Gaps

Policy, π	$J^{\pi^*}(S^0)$	Optimality Gap
Optimal	60.23	N/A
Myopic 1	45.99	23.64%
Myopic 2	55.74	7.45%
Myopic 3	54.69	9.25%

5.8 Excursion 4 - 3×3 case

This section expands the 2×2 case by incorporating an additional zone and MEDEVAC unit. For simplicity, the expanded 2×2 case is referred to as the 3×3 case. The 3×3 case assumes that the southern region of Afghanistan is the AO and is divided into three separate demand zones: Helmand province (Zone 1), Kandahar province (Zone 2), and Zabol province (Zone 3). Three MEDEVAC units are considered with one being staged with a collocated MTF in Zone 1 (i.e., MEDEVAC 1), one being staged with a collocated MTF in Zone 2 (i.e., MEDEVAC 2), and one being stage *without* a collocated MTF (i.e., MEDEVAC 3). The placement of the MEDEVAC unit staging areas and MTFs in Zones 1 and 2 are the same as the 2×2 case and the placement of the MEDEVAC unit staging area for Zone 3 represents a general realism based on the historical trends in enemy activity in Zabol.

A list of parameters associated with the 3×3 case are displayed in Table 28. Utilizing the parameter settings in Table 28 and the expected response times, expected service times, and immediate expected rewards computed for the 3×3 case (described in detail in following paragraphs), the optimal policy is determined via policy iteration. Applying Equation (5) reveals shows that the size of the state space for the 3×3 case is 326,592. This is a substantial increase from the 2×2 case (58,320).

Table 28. 3×3 Case Parameters

Parameter	Description	Setting
λ	9-line MEDEVAC request arrival rate	$\frac{1}{60}$
$ \mathcal{M} $	Total # of MEDEVACs	3
$ \mathcal{Z} $	Total # of zones	3
$ \mathcal{K} $	Total # of precedence levels	2
q^{max}	Max (z, k) queue length	2
γ	Uniformized discount factor	0.99
δ	Weight for urgent request	10
ξ	Scale for holding cost rate	0.2
p_{z_1}	Zone 1 proportion of requests	0.5836
p_{z_2}	Zone 2 proportion of requests	0.3407
p_{z_3}	Zone 3 proportion of requests	0.0757
p_{k_1}	Urgent proportion of requests	0.5
p_{k_2}	Priority proportion of requests	0.5

Recall that Helmand and Kandahar are the two most war-related, fatality-producing provinces in Afghanistan since the start of OEF with 956 and 558 KIA, respectively. During this same period, 124 KIAs occurred in the Zabol province (White, 2016). These numbers are utilized to determine the proportion of 9-line MEDEVAC requests from each zone. Simple calculations yield that the proportion of requests coming from Zone 1 is $p_{z_1} = 0.5836$, the proportion of requests coming from Zone 2 is $p_{z_2} = 0.3407$, and the proportion of requests coming from Zone 3 is $p_{z_3} = 0.0757$.

The 3×3 case assumes that the proportion of requests classified with an urgent precedence level is $p_{k_1} = 0.5$ and the proportion of requests classified with a priority precedence level is $p_{k_2} = 1 - p_{k_1} = 0.5$. Recall that the proportion of each request categorization p_{zk} is found by multiplying the zone proportion with the precedence level proportion (e.g., $p_{11} = p_{z_1}p_{k_1}$). Table 29 shows the 3×3 case's request categorization proportions.

The 3×3 case utilizes the same procedures as the 2×2 case to model future 9-line MEDEVAC requests and to compute expected response times, expected service times, and immediate expected rewards.

Table 29. Proportions of Zone-Precedence Level 9-Line MEDEVAC Requests

Zone, z	Urgent	Priority
1 (Helmand)	0.2918	0.2918
2 (Kandahar)	0.17035	0.17035
3 (Zabol)	0.03785	0.03785

Figure 11 depicts the three zones (i.e., Helmand, Kandahar, and Zabol) in southern Afghanistan utilized to generate the data, as well as the MEDEVAC and MTF locations. Recall that the MEDEVAC and MTF locations for Zones 1 and 2 are collocated. These collocated MEDEVAC and MTF locations are represented by blue stars. The location of the MEDEVAC unit without a collocated MTF is represented by a blue square. The casualty cluster centers in each zone are represented by red diamonds.




Figure 11. MEDEVAC and MTF locations with Casualty Cluster Centers

Figure 12 illustrates several casualty events resulting in 9-line MEDEVAC requests throughout southern Afghanistan within a 48-hour time period. The collocated MEDEVAC and MTF locations are still represented by blue stars. Moreover, the location of the MEDEVAC unit without a collocated MTF is still represented by a blue square. The casualty events are represented by red circles.




Figure 12. Sampled Casualty Events in Helmand, Kandahar, and Zabol

Tables 30-32 respectively report the expected response times, expected service times, and immediate expected rewards for the 3×3 case, computed utilizing the model parameter values shown in Table 28.

The same three inter-zone baseline policies (i.e., Myopic 1, Myopic 2, Myopic 3) from the 2×2 case are utilized for comparison purposes. The dispatching policies for the optimal policy and three baseline policies are compared in four separate scenarios.

Table 30. Expected Response Times (minutes)

Zone, z	MEDEVAC 1	MEDEVAC 2	MEDEVAC 3
1 (Helmand)	34.25	48.18	74.75
2 (Kandahar)	52.98	36.89	56.13
3 (Zabol)	102.25	83.91	57.54

Table 31. Expected Service Times (minutes)

Zone, z	MEDEVAC 1	MEDEVAC 2	MEDEVAC 3
1 (Helmand)	34.25	67.28	120.31
2 (Kandahar)	72.13	36.89	83.26
3 (Zabol)	121.40	83.91	84.67

Table 32. Immediate Expected Reward

Zone, z	Precedence, k	MEDEVAC, m		
		1	2	3
Zone 1 (Helmand)	1(Urgent)	5.65	4.48	2.88
	2 (Priority)	0.87	0.82	0.73
Zone 2 (Kandahar)	1 (Urgent)	4.14	5.41	3.92
	2 (Priority)	0.80	0.86	0.79
Zone 3 (Zabol)	1 (Urgent)	1.82	2.47	3.83
	2 (Priority)	0.65	0.71	0.79

Each scenario (i.e., Scenarios 5-8) considers a set of MEDEVAC system states with empty zone-precedence queues.

Scenario 5 considers a system wherein every MEDEVAC unit is idle, which can be represented by $S_t \in ((0, 0, 0), (0, 0, 0, 0, 0, 0), \hat{R}_t)$. The dispatching policies for Scenario 5 are displayed in Table 33. An asterisk (*) is placed next to the incoming requests, $\hat{R}_t$, that do not follow a myopic policy. The results indicate that when the MEDEVAC system is in Scenario 5 the optimal policy reacts myopically for every type of incoming request except for when $\hat{R}_t = (1, 2)$. The optimal policy dispatches MEDEVAC 2 to service priority 9-line MEDEVAC requests originating from Zone 1 when all MEDEVAC units are idle and there are no queued requests. This indicates that the optimal policy reserves MEDEVAC 1 for urgent 9-line MEDEVAC requests. This aligns with expectations because Zone 1 has the highest proportion of 9-line

MEDEVAC requests ($p_{z_1} = 0.5836$).

Table 33. Comparison of Dispatching Policies for Scenario 5

Policy	$\hat{R}_t$	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 1
	(1,2)*	Dispatch MEDEVAC 2
	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
	(3,1)	Dispatch MEDEVAC 3
	(3,2)	Dispatch MEDEVAC 3
Myopic 1	(1,1)	Dispatch MEDEVAC 1
Myopic 2	(1,2)	Dispatch MEDEVAC 1
Myopic 3	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
	(3,1)	Dispatch MEDEVAC 3
	(3,2)	Dispatch MEDEVAC 3

Scenario 6 considers a set of MEDEVAC system states wherein MEDEVACs 1 and 2 are idle and MEDEVAC 3 is busy, which can be represented by $S_t \in ((0, 0, z), (0, 0, 0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2, 3\}$. The dispatching policies associated with being in a Scenario 6 state are displayed in Table 34. The results indicate that the best MEDEVAC unit to dispatch to service a 9-line MEDEVAC request does not depend on the zone in which the busy MEDEVAC (i.e., MEDEVAC 3) is currently servicing when the system is in a Scenario 6 state. Moreover, the MEDEVAC system reacts in a myopic fashion for every incoming 9-line MEDEVAC request except for when $\hat{R}_t = (3, 2)$. The optimal policy queues 9-line MEDEVAC requests when $S_t \in ((0, 0, z), (0, 0, 0, 0, 0, 0), (3, 2)), z = 1, 2, 3$ (i.e., MEDEVACs 1 and 2 are idle, MEDEVAC 3 is busy, there are no queued requests, and a request originates from Zone 3 with a priority precedence level). This shows the optimal policy reserves MEDEVAC 2 for 9-line requests originating from its own zone or urgent 9-line MEDEVAC requests from Zone 3.

Scenario 7 considers a set of MEDEVAC system states wherein MEDEVACs 1 and 3 are idle and MEDEVAC 2 is busy, which can be represented by $S_t \in$

Table 34. Comparison of Dispatching Policies for Scenario 6

Policy	$\hat{R}_t$	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 1
	(1,2)	Dispatch MEDEVAC 1
	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
	(3,1)	Dispatch MEDEVAC 2
	(3,2)*	Queue
Myopic 1	(1,1)	Dispatch MEDEVAC 1
Myopic 2	(1,2)	Dispatch MEDEVAC 1
Myopic 3	(2,1)	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2
	(3,1)	Dispatch MEDEVAC 2
	(3,2)	Dispatch MEDEVAC 2

$((0, z, 0), (0, 0, 0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2, 3\}$. The dispatching policies associated with being in a Scenario 7 state are displayed in Table 35. The results indicate that the best MEDEVAC unit to dispatch to service a 9-line MEDEVAC request does not depend on the zone in which the busy MEDEVAC (i.e., MEDEVAC 2) is currently servicing when the system is in State 3. Moreover, the MEDEVAC system reacts in a myopic fashion for every incoming 9-line MEDEVAC request except for when $\hat{R}_t = (2, 2)$. The optimal policy dispatches MEDEVAC 3 to service 9-line MEDEVAC requests when $S_t = ((0, z, 0), (0, 0, 0, 0, 0, 0), (2, 2)), z = 1, 2, 3$ (i.e., MEDEVACs 1 and 3 are idle, MEDEVAC 2 is busy, there are no queued requests, and a request originates from Zone 2 with a priority precedence level). This shows the optimal policy reserves MEDEVAC 2 for 9-line requests originating from its own zone or urgent 9-line MEDEVAC requests from Zone 2.

Scenario 8 considers a set of MEDEVAC system states wherein MEDEVACs 2 and 3 are idle and MEDEVAC 1 is busy, which can be represented by $S_t \in ((z, 0, 0), (0, 0, 0, 0, 0, 0), \hat{R}_t)$ where $z \in \{1, 2, 3\}$. The dispatching policies associated with being in a Scenario 8 state are displayed in Table 36. Unlike the policies from Scenarios 6 and 7, the results indicate that the best MEDEVAC unit to dispatch

Table 35. Comparison of Dispatching Policies for Scenario 7

Policy	$\hat{R}_t$	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 1
	(1,2)	Dispatch MEDEVAC 1
	(2,1)	Dispatch MEDEVAC 1
	(2,2)*	Dispatch MEDEVAC 3
	(3,1)	Dispatch MEDEVAC 3
	(3,2)	Dispatch MEDEVAC 3
Myopic 1	(1,1)	Dispatch MEDEVAC 1
Myopic 2	(1,2)	Dispatch MEDEVAC 1
Myopic 3	(2,1)	Dispatch MEDEVAC 1
	(2,2)	Dispatch MEDEVAC 1
	(3,1)	Dispatch MEDEVAC 3
	(3,2)	Dispatch MEDEVAC 3

to service a 9-line MEDEVAC request does depend on the zone in which the busy MEDEVAC (i.e. MEDEVAC 1) is currently servicing when the system is in a Scenario 8 state. The optimal dispatching policy for $\hat{R}_t = (1, 2)$ depends on z , the zone where MEDEVAC 1 is currently servicing a request. If MEDEVAC 1 is servicing Zone 1 (i.e., $z = 1$) and $\hat{R}_t = (1, 2)$, then the optimal decision is to dispatch MEDEVAC 2 to service the request. If MEDEVAC 1 is servicing either Zone 2 or Zone 3 (i.e., $z = 2$ or $z = 3$, respectively) and $\hat{R}_t = (1, 2)$, then the optimal decision is to dispatch MEDEVAC 3 to service the request. Recall that the service time of MEDEVAC 1 servicing Zone 1 is 34.25 minutes, which is substantially less than when MEDEVAC 1 is servicing Zone 2 (67.28 minutes) or Zone 3 (120.31 minutes). The optimal policy reserves MEDEVAC 2 when MEDEVAC 1 is servicing either Zone 2 or Zone 3 because of these long service times.

The optimality gap between the optimal policy and the baseline policies is examined. The expected total discounted reward for the optimal policy and baseline policies when the system is in an *empty* state $S^0 = ((0, 0, 0), (0, 0, 0, 0, 0, 0), (0, 0))$ (i.e., empty queues, idle MEDEVACs, no incoming 9-line MEDEVAC requests) are displayed in Table 37, along with the optimality gaps associated with each baseline

Table 36. Comparison of Dispatching Policies for Scenario 8

Policy	$\hat{R}_t$	MEDEVAC 1 Servicing Zone 1	MEDEVAC 1 Servicing Zone 2	MEDEVAC 1 Servicing Zone 3
		Queue\Dispatch\Reject	Queue\Dispatch\Reject	Queue\Dispatch\Reject
Optimal	(1,1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(1,2)*	Dispatch MEDEVAC 2	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3
	(2,1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(3,1)	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3
	(3,2)	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3
Myopic 1	(1,1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
Myopic 2	(1,2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
Myopic 3	(2,1)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(2,2)	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2	Dispatch MEDEVAC 2
	(3,1)	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3
	(3,2)	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3	Dispatch MEDEVAC 3

policy. The results indicate that the best baseline policy is Myopic 2, which has the smallest optimality gap of 3.48%. Without having the ability to queue any requests, the Myopic 3 policy performs worse than every other policy and has the largest optimality gap of 5.79%.

Table 37. Comparison of Total Expected Discounted Rewards & Optimality Gaps

Policy, π	$J^\pi(S^0)$	Optimality Gap
Optimal	56.41	N/A
Myopic 1	54.04	4.19%
Myopic 2	54.44	3.48%
Myopic 3	53.14	5.79%

5.9 Policy Iteration versus Linear Programming

This section compares the computational efficiency between policy iteration via MATLAB and linear programming (LP) via CPLEX for the MEDEVAC dispatching problem. Since each solution algorithm determines the optimal dispatching policy, the focus of the analysis is on how long it takes each algorithm to solve. Comparisons are made on the same computer and on the same problem instances after they have been loaded into memory. The problem instances are generated by adjusting the q^{max} parameter in the 2×2 case. Table 38 reports the total time in seconds required to find the optimal policy for each algorithm, and Figure 13 depicts the results from

Table 38 to visually show the differences in computational time.

Table 38. Policy Iteration versus Linear Programming Computational Efficiency (s)

$ \mathcal{S} $	Policy Iteration	CPLEX (Dual)	CPLEX (Primal)
720	0.03	0.07	1.17
3645	0.10	0.47	4.26
11520	0.35	3.13	38.15
28125	1.04	13.32	196.50
58320	2.23	55.74	656.16
108045	5.12	134.03	1981.55
184320	10.91	216.66	4782.90
295245	17.37	309.83	9754.49
450000	47.86	412.66	17037.00




Figure 13. Policy Iteration vs Linear Programming

The results from Table 38 and Figure 13 indicate that solving the MEDEVAC dispatching problem with CPLEX utilizing a primal simplex optimizer is substantially

worse than solving the problem with either policy iteration or CPLEX utilizing a dual simplex optimizer. Figure 14 excludes the results from CPLEX utilizing a primal simplex optimizer to provide a better visual comparison of policy iteration versus CPLEX utilizing a dual simplex optimizer.




Figure 14. Policy Iteration vs Linear Programming

The results from Figure 14 indicate that policy iteration substantially outperforms CPLEX even when the more theoretically appropriate dual simplex optimizer is utilized. Moreover, the results from Table 38, Figure 13, and Figure 14 show that the gaps between each algorithm increase as $|S|$ increases, indicating that larger, small-scale problems (i.e., ones that can still be solved to optimality) should be solved via policy iteration. These results comport with the findings of Puterman (1994) and Powell (2011).

LP problems can be stated in primal or dual form. Moreover, the optimal solution (if one exists) of the dual has a direct relationship to an optimal solution of the primal LP model. The dual simplex optimizer in CPLEX takes advantage of this relationship, but still reports the solution to a given problem in terms of the primal model. For the primal LP model of the MDP, the number of rows (i.e., inequality constraints) is equal to $|\mathcal{S}| \times \prod_{S_t \in \mathcal{S}} |\mathcal{X}(S_t)|$ (i.e., the number of state-action combinations). The number of columns (i.e., the number of variables) is equal to $|\mathcal{S}|$. Modern LP solvers can handle problems with tens of thousands of constraints without difficulty (Powell, 2011). Based on the sizes of the state and action space, it may be more efficient to solve the problems utilizing the dual formulation of the LP model resulting in $|\mathcal{S}|$ rows and $|\mathcal{S}| \times \prod_{S_t \in \mathcal{S}} |\mathcal{X}(S_t)|$ columns in the constraint matrix. Despite the greatly increased computational efficiency in LP algorithms reported in Bixby (2012), the results from this analysis indicate that policy iteration substantially outperforms LP via CPLEX (for both primal and dual simplex optimizers) for the MEDEVAC dispatching problem.

VI. Conclusion

This thesis examines the medical evacuation (MEDEVAC) dispatching problem. The objective of this research is to determine how to optimally dispatch MEDEVAC units to 9-line MEDEVAC requests to improve the performance of a deployed medical service system and ultimately maximize battlefield casualty survivability rates. A discounted, infinite horizon Markov decision process (MDP) is developed to enable examination of many different military medical planning scenarios. The MDP model incorporates admission control and queueing, which allows the dispatching authority to accept, reject, or queue incoming 9-line MEDEVAC requests based on the request's classification (i.e., zone and precedence level) and the state of the MEDEVAC system. Rejected requests are not simply discarded; rather, they are redirected to another servicing agency, such as casualty evacuation (CASEVAC), to be serviced. The MDP model also accounts for the severity of each call (i.e., urgent and priority) and applies a survivability function that is monotonically decreasing in response time to model the outcome of casualties. While response time thresholds (RTTs) are typically utilized to measure system performance for emergency medical systems, this thesis measures performance in terms of casualty survivability since survival probability more accurately mirrors casualty outcomes. To demonstrate the applicability of the MDP model and to examine the behavior of the optimal dispatching policy, a notional military planning scenario based on contingency operations in southern Afghanistan is developed. A series of sensitivity analyses and computational excursions identify the model parameters that significantly impact the optimal dispatching policy. Moreover, this thesis compares the computational efficiency of policy iteration via MATLAB versus linear programming via CPLEX, utilizing either of two embedded simplex implementation methodologies

The immediate expected reward obtained from servicing a specific 9-line MEDEVAC

VAC request depends on the locations of the request and the servicing MEDEVAC unit's staging area, along with the precedence level of the request. The total holding cost that the MEDEVAC system incurs during each state transition depends on the total number of queued requests and the precedence level of each queued request in the MEDEVAC system. Decisions are made when either a 9-line MEDEVAC request is submitted to the system or when a MEDEVAC unit finishes servicing a request. The dispatching authority examines the entire state of the MEDEVAC system when a decision is required.

Results indicate that dispatching the closest available MEDEVAC unit (i.e., a myopic policy) is not always optimal. Instead, dispatching MEDEVAC units considering the entire MEDEVAC system state (i.e., the MEDEVAC units' status, number and precedence level of queued requests, and location and precedence of the incoming request) increases the casualty survivability. The optimality gaps between the myopic policies examined and the optimal policy range between 0.74% and 5.73% when inter-zone policies are allowed and 7.45% and 23.64% when intra-zone policies are enforced. Over a protracted conflict, these policies will substantially decrease the survivability rates of battlefield casualties, and, therefore, implementation of optimal policies should be considered by medical planners. Myopic policies are often utilized in military practice because they are relatively easy to implement and they perform fairly well as long as the arrival of 9-line MEDEVAC requests occur less frequently. Of the myopic policies tested, the Myopic 2 policy performs the best in both the 2×2 case and 3×3 case with optimality gaps of 0.74% and 3.48%, respectively. Moreover, results confirm the criticality of the MEDEVAC helicopter's flight speed. Current flight speeds can decrease due to atmospheric, environmental, or mechanical issues. If these problems arise during combat operations and degrade the flight speed of the MEDEVAC helicopters, then myopic policies perform even worse when

compared to the optimal policy. For example, if the current flight speeds of MEDEVAC helicopters decrease by 50%, a myopic policy that queues all requests when no MEDEVAC units are available has a 17.07% optimality gap, which is substantially more than the baseline optimality gap of 2.21%. These results suggest that medical planners should consider changing dispatching policies during combat operations if one or more of these problems arise and negatively impact the flight speed of the MEDEVAC helicopters being utilized. Conversely, current flight speeds can increase if new rotary wing aircraft are employed in combat operations. If this were to occur, initial results indicate that as the flight speed increases, the performance gap between myopic policies and the optimal policy decreases. For example, if the current flight speeds of MEDEVAC helicopters increases by 50%, a myopic policy that queues only urgent requests when no MEDEVAC units are available only has a 0.24% optimality gap, which is substantially less than the baseline optimality gap of 0.74%. This comparison informs current MEDEVAC helicopter designs and development and provides promising results in terms of saving lives with a faster MEDEVAC helicopter.

The research presented in this thesis is of interest to both military and civilian medical planners and dispatch authorities. Medical planners can apply the MDP model developed to compare different dispatching policies for a variety of planning scenarios with fixed medical treatment facility (MTF) and MEDEVAC staging (i.e., hospital and ambulance for civilian sector) locations. Moreover, medical planners can evaluate different location schemes for the medical assets (e.g., MTFs, hospitals, MEDEVACs, and ambulances) to maximize the overall performance of the medical system.

One limiting assumption associated with the MDP model developed is that MEDEVAC units are required to return to their own staging areas to refuel and replenish medical supplies after unloading casualties at an MTF prior to servicing a queued

request. During combat operations, there are typically bases that have collocated MEDEVAC units and MTFs. It is reasonable to assume that MEDEVAC units staged in different zones can refuel and replenish medical supplies at these locations and immediately proceed to servicing a queued request instead of having to return back to their own staging areas first. The MDP model restricts MEDEVAC units from refueling at different locations as a simplifying assumption. Modifying the problem formulation and the corresponding MDP model to allow for refueling, replenishing of supplies, and the ability to immediately service queued requests after casualty delivery at an MTF with a collocated MEDEVAC unit would certainly reduce the response time for many 9-line MEDEVAC requests. This modification is a planned extension for future research.

Another insight drawn from this thesis is that the computational difficulty in solving the MEDEVAC dispatching problem increases substantially as the size of the state space grows. The computational efficiency of policy iteration via MATLAB is compared to linear programming (LP) via CPLEX. The results reveal that, although great strides have been accomplished in improving the performance of LP algorithms, policy iteration still outperforms LP algorithms by a substantial amount. Nevertheless, as the size of the state space grows exponentially, the use of exact dynamic programming techniques becomes intractable. This makes more realistic, large-scale problem instances impossible to analyze via exact algorithms. A planned extension to this work involves incorporating several approximate dynamic programming (ADP) algorithms to resolve the well known curse of dimensionality issue. While the representative scenario analyzed is not of a large-scale scenario, important insights are still drawn concerning the differences between the optimal policy and standard myopic policies utilized today. These insights should be taken into consideration by military medical planners and utilized when planning for major combat operations.

Using Markov Decision Processes with Heterogeneous Queueing Systems to Examine Military MEDEVAC Dispatching Policies

Capt Phillip R. Jenkins, Advisor: Lt Col Matthew J. Robbins, Ph.D., Reader: Dr. Brian J. Lunday



Appendix A. Storyboard

BACKGROUND

- The primary objective of a deployed military emergency medical services system is to evacuate battlefield casualties in an efficient and timely manner
- Military medical planners must consider how medical evacuation (MEDEVAC) units will be dispatched prior to and during military engagement in major combat operations
- MEDEVAC units rapidly transport battlefield casualties from a pre-determined casualty collection point (CCP) to the nearest medical treatment facility (MTF)
- The MEDEVAC dispatching problem is modeled as a discounted, infinite horizon Markov decision process (MDP)

PROBLEM DESCRIPTION



- When a casualty event occurs and a 9-line MEDEVAC request is submitted, a decision must be made quickly as to which MEDEVAC unit to dispatch
- Decision is complicated by
 - Threat level at CCP
 - Location of request
 - Location of MTFs
 - Location of MEDEVAC units
 - Severity of casualties
- Evacuation Categories
 - Priority 1 (Urgent): ASAP or within 60 minutes
 - Priority 2 (Priority): within 4 hours
 - Priority 3 (Routine): within 24 hours

CONTACT INFORMATION


Lt Col Matthew J. Robbins, Ph.D.
Department of Operational Sciences, AFIT

RESEARCH OBJECTIVES

- Determine the MEDEVAC dispatching policy that maximizes expected total discounted reward over an infinite horizon
- Compare optimal policy against practitioner-friendly baseline policies
- Examine sensitivity of model parameters
- Compare computational efficiency of policy iteration and linear programming (LP) algorithms





COMPUTATIONAL EXAMPLE



- A representative planning scenario based on notional contingency operations in southern Afghanistan due to an increase in terrorist attacks by remnants of Al Qaeda militants is considered
- Collocated MEDEVAC units and MTFs are staged in the two most war-related fatality producing provinces since 2001, Helmand and Kandahar. Moreover, a MEDEVAC unit is staged in Zabol without a collocated MTF

RESULTS AND CONCLUSIONS

- Dispatching the closest available MEDEVAC unit is not always optimal
- The optimal dispatching policy and the optimality gaps of the baseline policies are sensitive to
 - Airspace restrictions (i.e., inter vs intra)
 - Request arrival rate and flight speed
 - Proportion and weight of urgent requests
- Myopic 2 (Queue Urgent) is the best performing myopic policy, on average
- Policy iteration substantially outperforms LP algorithms in terms of computational efficiency

MDP MODEL

- Decision Epochs: $t \in \mathcal{T} = \{1, 2, \dots\}$
- States: $S_t = (M_t, Q_t, \hat{R}_t)$
 - MEDEVAC status: $M_t = (M_{im})_{m \in \mathcal{M}}$
 - Queue status: $Q_t = (Q_{i+k})_{i \in \mathcal{Z}, k \in \mathcal{K}}$
 - Arrival status: $\hat{R}_t = (\hat{R}_i, \hat{K}_i)_{\hat{R}_i \in \mathcal{Z}, \hat{K}_i \in \mathcal{K}}$
- Decisions: $x_t = (x_t^{evac}, x_t^d)$
 - Admission Control: x_t^{evac}
 - Dispatching: $x_t^d = (x_t^{d^*}, x_t^{d^*})$
- Transitions: Markovian and occur when
 - A 9-line MEDEVAC request is submitted
 - A busy MEDEVAC unit completes service
- Rewards: Immediate and continuous, $r(S_t, x_t)$
 - Immediate reward for servicing requests
 - Continuous cost for queued requests
- Objective: $\max_{x \in \mathcal{E}^*} \left\{ \sum_{t=1}^{\infty} \gamma^{t-1} r(S_t, X^*(S_t)) \right\}$

INTERESTED PARTIES

- This research interests
 - Military medical planning community
 - United States Army Medical Evacuation Potency Directorate
 - Military acquisition specialists
- Inform development of tactics, techniques, and procedures for MEDEVAC dispatching authorities
- Military acquisition specialists can utilize these results to examine the impact of improved MEDEVAC helicopter speed and capacity

FUTURE RESEARCH

- Increase problem size and apply approximate dynamic programming techniques
- Allow dispatching directly from MTF
- Allow dynamic in-flight rerouting
- Explicitly model MEDEVAC and MTF capacity
- Explicitly model batch arrival of casualties for each 9-line MEDEVAC request

Bibliography

- Alsalloum, Othman Ibraheem, & Rand, Graham K. 2006. Extensions to emergency vehicle location models. *Computers & Operations Research*, **33**(9), 2725–2743.
- Baker, Joanna R, Clayton, Edward R, & Taylor III, Bernard W. 1989. A non-linear multi-criteria programming approach for determining county emergency medical service ambulance allocations. *Journal of the Operational Research Society*, **40**(5), 423–432.
- Bandara, Damitha, Mayorga, Maria E, & McLay, Laura A. 2012. Optimal dispatching strategies for emergency vehicles to increase patient survivability. *International Journal of Operational Research*, **15**(2), 195–214.
- Bandara, Damitha, Mayorga, Maria E, & McLay, Laura A. 2014. Priority dispatching strategies for EMS systems. *Journal of the Operational Research Society*, **65**(4), 572–587.
- Bastian, Nathaniel D. 2010. A robust, multi-criteria modeling approach for optimizing aeromedical evacuation asset emplacement. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, **7**(1), 5–23.
- Bastian, Nathaniel D, Fulton, Lawrence V, Mitchell, Robert, Pollard, Wayne, Wierschem, David, & Wilson, Ronald. 2012. The future of vertical lift: initial insights for aircraft capability and medical planning. *Military Medicine*, **177**(7), 863–869.
- Berlin, Geoffrey N, & Liebman, Jon C. 1974. Mathematical analysis of emergency ambulance location. *Socio-Economic Planning Sciences*, **8**(6), 323–328.
- Berman, Oded. 1981. Repositioning of distinguishable urban service units on networks. *Computers & Operations Research*, **8**(2), 105–118.
- Bianchi, Geoffrey, & Church, Richard L. 1988. A hybrid FLEET model for emergency medical service system design. *Social Science & Medicine*, **26**(1), 163–171.
- Bixby, Robert E. 2012. A brief history of linear and mixed-integer programming computation. *Documenta Mathematica, Extra Volume: Optimization Stories*, 107–121.
- Bouma, Matthew F. 2005. *Medical evacuation and treatment capabilities optimization model*. Ph.D. thesis, Monterey, California. Naval Postgraduate School.
- Carter, Grace M, Chaiken, Jan M, & Ignall, Edward. 1972. Response areas for two emergency units. *Operations Research*, **20**(3), 571–594.
- Chaiken, Jan M, & Larson, Richard C. 1972. Methods for allocating urban emergency units: a survey. *Management Science*, **19**(4), 110–130.

- Clarke, Jonathan E, & Davis, Peter R. 2012. Medical evacuation and triage of combat casualties in Helmand Province, Afghanistan: October 2010-April 2011. *Military Medicine*, **177**(11), 1261–1266.
- Daskin, Mark S, & Stern, Edmund H. 1981. A hierarchical objective set covering model for emergency medical service vehicle deployment. *Transportation Science*, **15**(2), 137–152.
- De Lorenzo, Robert A. 2003. Military casualty evacuation: MEDEVAC. *Aeromedical evacuation: Management of acute and stabilized patients*, Springer Verlag, New York, 45–59.
- Department of Defense. 2016. *Defense Casualties Analysis System (DCAS) Operation FREEDOM'S SENTINEL (OFS)*. https://www.dmdc.osd.mil/dcas/pages/casualties_ofs.xhtml. Online; accessed 22 December 2016.
- Department of the Army, United States. 2000 (April). *Field Manual 8-10-6, Medical Evacuation in a Theater of Operations*.
- Department of the Army, United States. 2014 (August). *Field Manual 4-02.2, Medical Evacuation*.
- Eastridge, Brian J, Mabry, Robert L, Seguin, Peter, Cantrell, Joyce, Tops, Terrill, Uribe, Paul, Mallett, Olga, Zubko, Tamara, Oetjen-Gerdes, Lynne, Rasmussen, Todd E, *et al.* 2012. Death on the battlefield (2001–2011): implications for the future of combat casualty care. *Journal of Trauma and Acute Care Surgery*, **73**(6), S431–S437.
- Efrosinin, Dmitry. 2004. *Controlled queueing systems with heterogeneous servers*. Ph.D. thesis, Univ. Trier, FB 4, Informatik.
- Erkut, Erhan, Ingolfsson, Armann, & Erdoğan, Güneş. 2008. Ambulance location for maximum survival. *Naval Research Logistics (NRL)*, **55**(1), 42–58.
- Feero, Stan, Hedges, Jerris R, Simmons, Erik, & Irwin, Lisa. 1995. Does out-of-hospital EMS time affect trauma survival? *The American Journal of Emergency Medicine*, **13**(2), 133–135.
- Finkbeiner, Scott C. 2013. *Urgent Aeromedical Evacuation Network Capacity Planning*. Ph.D. thesis, AFIT MS Thesis.
- Fish, Peter N. 2014. *Army Medical Officer's Guide*. Stackpole Books.
- Fitch, Jay. 2005. Response times: myths, measurement & management. *JEMS: a Journal of Emergency Medical Services*, **30**(9), 47–56.
- Fitch, Jay, & Griffiths, Keith. 2005. EMS in critical condition: Meeting the challenge[Item No. E-43338]. *IQ Report*, **37**(5), 1–16.

- Fulton, Larry, McMurry, Pat, & Kerr, Bernie. 2009. A Monte Carlo simulation of air ambulance requirements during major combat operations. *Military Medicine*, **174**(6), 610–614.
- Fulton, Lawrence V, Lasdon, Leon S, McDaniel, Reuben R, & Coppola, M Nicholas. 2010. Two-stage stochastic optimization for the allocation of medical assets in steady-state combat operations. *The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology*, **7**(2), 89–102.
- Garrett, Michael X. 2013. *USCENTCOM Review of MEDEVAC Procedures in Afghanistan*. Tech. rept. United States Central Command.
- Grannan, Benjamin C, Bastian, Nathaniel D, & McLay, Laura A. 2015. A maximum expected covering problem for locating and dispatching two classes of military medical evacuation air assets. *Optimization Letters*, **9**(8), 1511–1531.
- Green, Linda, & Kolesar, Peter. 1984. A comparison of the multiple dispatch and M/M/c priority queueing models of police patrol. *Management Science*, **30**(6), 665–670.
- Green, Linda V, & Kolesar, Peter J. 2004. Anniversary article: Improving emergency responsiveness with management science. *Management Science*, **50**(8), 1001–1014.
- Gross, Donald, & Harris, Carl M. 1998. *Fundamentals of queueing theory*. Hoboken, New Jersey: John Wiley & Sons.
- Hall, William K. 1972. The application of multifunction stochastic service systems in allocating ambulances to an urban area. *Operations Research*, **20**(3), 558–570.
- ICOS. 2008. *Afghanistan-Pakistan Insurgent activities in Afghanistan and Pakistan (2007)*. http://www.icosgroup.net/wp-content/gallery/taliban-presence/016_map.png. Online; accessed 4 January 2017.
- Ignall, E, Carter, G, & Rider, K. 1982. An algorithm for the initial dispatch of fire companies. *Management Science*, **28**(4), 366–378.
- Jarvis, James P. 1985. Approximating the equilibrium behavior of multi-server loss systems. *Management Science*, **31**(2), 235–239.
- Jarvis, James Patrick. 1975. *Optimization in stochastic service systems with distinguishable servers*. Ph.D. thesis, Massachusetts Institute of Technology.
- Keneally, Sean K, Robbins, Matthew J, & Lunday, Brian J. 2016. A Markov decision process model for the optimal dispatch of military medical evacuation assets. *Health Care Management Science*, **19**(2), 111–129.

- Knight, Vincent Anthony, Harper, Paul Robert, & Smith, Leanne. 2012. Ambulance allocation for maximal survival with heterogeneous outcome measures. *Omega*, **40**(6), 918–926.
- Kolesar, Peter, & Walker, Warren E. 1974. An algorithm for the dynamic relocation of fire companies. *Operations Research*, **22**(2), 249–274.
- Kuisma, Markku, Holmström, Peter, Repo, Jukka, Määttä, Teuvo, Nousila-Wiik, Maria, & Boyd, James. 2004. Prehospital mortality in an EMS system using medical priority dispatching: a community based cohort study. *Resuscitation*, **61**(3), 297–302.
- Kulkarni, Vidyadhar G. 2009. *Modeling and analysis of stochastic systems*. Boca Raton, Florida: CRC Press.
- Lejeune, Miguel A, & Margot, F. 2016. *Aeromedical Battlefield Evacuation under Endogenous Uncertainty*. Tech. rept. Carnegie Mellon University, Pittsburg, PA, USA.
- Leoni, Ray D. 2007. *Black hawk: the story of a world class helicopter*. American Institute of Aeronautics.
- Lippman, Steven A. 1975. Applying a new device in the optimization of exponential queuing systems. *Operations Research*, **23**(4), 687–710.
- MacFarlane, C, & Benn, CA. 2003. Evaluation of emergency medical services systems: a classification to assist in determination of indicators. *Emergency Medicine Journal*, **20**(2), 188–191.
- Malsby III, Robert F, Quesada, Jose, Powell-Dunford, Nicole, Kinoshita, Ren, Kurtz, John, Gehlen, William, Adams, Colleen, Martin, Dustin, & Shackelford, Stacy. 2013. Prehospital blood product transfusion by US Army MEDEVAC during combat operations in Afghanistan: a process improvement initiative. *Military Medicine*, **178**(7), 785–791.
- Maxwell, Matthew S, Restrepo, Mateo, Henderson, Shane G, & Topaloglu, Huseyin. 2010. Approximate dynamic programming for ambulance redeployment. *INFORMS Journal on Computing*, **22**(2), 266–281.
- Mayorga, Maria E, Bandara, Damitha, & McLay, Laura A. 2013. Districting and dispatching policies for emergency medical service systems to improve patient survival. *IIE Transactions on Healthcare Systems Engineering*, **3**(1), 39–56.
- McLay, Laura A, & Mayorga, Maria E. 2010. Evaluating emergency medical service performance measures. *Health Care Management Science*, **13**(2), 124–136.

- McLay, Laura A, & Mayorga, Maria E. 2013a. A dispatching model for server-to-customer systems that balances efficiency and equity. *Manufacturing & Service Operations Management*, **15**(2), 205–220.
- McLay, Laura A, & Mayorga, Maria E. 2013b. A model for optimally dispatching ambulances to emergency calls with classification errors in patient priorities. *IIE Transactions*, **45**(1), 1–24.
- Nicholl, Jon, Coleman, Patricia, Parry, Gareth, Turner, Janette, & Dixon, Simon. 1999. Emergency priority dispatch systems a new era in the provision of ambulance services in the UK. *Pre-Hospital Immediate Care*, **3**(2), 71–75.
- Olson, Jr, Chris M, Bailey, Jeffrey, Mabry, Robert, Rush, Stephen, Morrison, Jonathan J, & Kuncir, Eric J. 2013. Forward aeromedical evacuation: A brief history, lessons learned from the Global War on Terror, and the way forward for US policy. *Journal of Trauma and Acute Care Surgery*, **75**(2), S130–S136.
- O’Shea, Brian. 2011. Saving Lives on the Battlefield. *Military Medical/CBRN Technology*, **15**(6), 8–13.
- Pons, Peter T, & Markovchick, Vincent J. 2002. Eight minutes or less: does the ambulance response time guideline impact trauma patient outcome? *The Journal of Emergency Medicine*, **23**(1), 43–48.
- Powell, Warren B. 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Princeton, New Jersey: John Wiley & Sons.
- Puterman, Martin L. 1994. *Markov decision processes: Discrete Stochastic Dynamic Programming*. Hoboken, New Jersey: John Wiley & Sons.
- Rettke, Aaron J, Robbins, Matthew J, & Lunday, Brian J. 2016. Approximate dynamic programming for the dispatch of military medical evacuation assets. *European Journal of Operational Research*, **254**(3), 824–839.
- Rosberg, Zvi, Varaiya, P, & Walrand, J. 1982. Optimal control of service in tandem queues. *IEEE Transactions on Automatic Control*, **27**(3), 600–610.
- Serfozo, Richard F. 1979. Technical Note An Equivalence Between Continuous and Discrete Time Markov Decision Processes. *Operations Research*, **27**(3), 616–620.
- Stidham, Jr, Shaler. 2002. Analysis, design, and control of queueing systems. *Operations Research*, **50**(1), 197–216.
- Stidham, Jr, Shaler, & Weber, Richard. 1993. A survey of Markov decision models for control of networks of queues. *Queueing Systems*, **13**(1-3), 291–314.

- Sundstrom, Scott C, Blood, Christopher G, & Matheny, Serge A. 1996. The optimal placement of casualty evacuation assets: a linear programming model. *Pages 907–911 of: Proceedings of the 28th conference on Winter simulation*. IEEE Computer Society.
- Swersey, Arthur J. 1982. A Markovian decision model for deciding how many fire companies to dispatch. *Management Science*, **28**(4), 352–365.
- White, Michael. 2016. *Operation ENDURING FREEDOM (OEF), Fatalities by Provinces*. <http://icasualties.org/OEF/ByProvince.aspx>. Online; accessed 22 December 2016.
- Williams, Dave. 2005. 2004 JEMS 200-city survey. *JEMS: a Journal of Emergency Medical Services*, **30**(2), 42–60.
- Zeto, John, Yamada, Wade, & Collins, Gale. 2006. Optimizing the Emplacement of Scarce Resources to Maximize the Expected Coverage of a Geographically Variant Demand Function. *Center for Army Analysis*.

Vita

Captain Phillip R. Jenkins attended Lowndes High School, Georgia and graduated in 2008. He accomplished his undergraduate studies at Ohio University with a Bachelor of Science degree in Mathematics in December 2012. Phillip commissioned into the United States Air Force as an Operations Research Analyst in December 2012.

Captain Jenkins' first assignment was to Air Force Global Strike Command A9 at Barksdale Air Force Base, Louisiana. While stationed at Barksdale, Phillip served as a Long Range Systems Strike Analyst and as the Chief of Force Structure Assessments.

In August 2015, Phillip entered the Air Force Institute of Technology's Graduate School of Engineering and Management at Wright-Patterson AFB, Ohio. Upon earning his Master's degree in Operations Research, he will continue his research at the Air Force Institute of Technology as a Doctor of Philosophy student with a specialty area in optimization and a minor in applied statistics.

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.						
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From — To)		
23-03-2017		Master's Thesis		Sept 2015 — Mar 2017		
4. TITLE AND SUBTITLE Using Markov Decision Processes with Heterogeneous Queueing Systems to Examine Military MEDEVAC Dispatching Policies				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Jenkins, Phillip R., Capt, USAF				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way WPAFB OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENS-MS-17-M-137		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) US Army Medical Evacuation Proponency Directorate Col Michael F. Breslin, Chief, MEPD 4103 Gladiator Street Fort Rucker, AL 36305 michael.f.breslin.mil@mail.mil				10. SPONSOR/MONITOR'S ACRONYM(S) MEPD		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A: APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT A major focus of the Military Health System is to provide efficient and timely medical evacuation (MEDEVAC) to battlefield casualties. The objective of this research is to determine how to optimally dispatch MEDEVAC units in response to requests to maximize system performance. A Markov decision process model is developed to examine the MEDEVAC dispatching problem. A representative planning scenario is utilized to investigate the differences between the optimal policy and three practitioner-friendly myopic baseline policies. Two computational experiments, a two-level, five-factor screening design and a subsequent three-level, three-factor full factorial design, are conducted to examine the impact of selected MEDEVAC problem features on the optimal policy and the system level performance measure. Results indicate that dispatching the closest available MEDEVAC unit is not always optimal. These results inform the development and implementation of MEDEVAC tactics, techniques, and procedures by military medical planners. Moreover, an open question exists concerning the best exact solution approach for solving Markov decision problems due to recent advances in performance by commercial linear programming (LP) solvers. An analysis of solution approaches for the MEDEVAC dispatching problem reveals that the policy iteration algorithm substantially outperforms the LP algorithms executed by CPLEX 12.6 in regards to computational effort.						
15. SUBJECT TERMS Markov decision processes, medical evacuation, admission control, queueing, priority dispatching, policy iteration, and linear programming comparison						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Dr. Matthew J. Robbins, AFIT/ENS	
U	U	U	UU	112	19b. TELEPHONE NUMBER (include area code) (937) 255-3636, x4539;matthew.robbs@afit.edu	